# HP Data Protector A.06.11

# Zero downtime backup integration guide

## for Oracle, SAP R/3, Microsoft SQL Server, Microsoft Exchange Server, and Microsoft Volume Shadow Copy Service

B 6 9 6 0 - 9 0 1 6 3

Part number: B6960-90163
First edition: September 2009

**Legal and notice information**

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

Printed in the US

# Contents

# Figures

# Tables

# Publication history

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

Table 1 Edition history

| Part number | Guide edition | Product |
|---|---|---|
| B6960-90114 | October 2004 | Data Protector Release A.05.50 |
| B6960-96013 | July 2006 | Data Protector Release A.06.00 |
| B6960-96047 | November 2008 | Data Protector Release A.06.10 |
| B6960-90163 | September 2009 | Data Protector Release A.06.11 |

# About this guide

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

## Intended audience

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

- Basic Data Protector functionality
- Database administration

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

It is also recommended to read the *HP Data Protector zero downtime backup concepts guide* for fundamentals of Data Protector integrations with disk arrays.

## Documentation set

Other documents and online Help provide related information.

### Guides

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

You can find these documents from the Manuals page of the HP Business Support Center website:

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

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

- *HP Data Protector concepts guide*

  This guide describes Data Protector concepts and provides background information on how Data Protector works. It is intended to be used with the task-oriented online Help.

- *HP Data Protector installation and licensing guide*

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

- *HP Data Protector troubleshooting guide*

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

- *HP Data Protector disaster recovery guide*

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

- *HP Data Protector integration guides*

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

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

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

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

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

  - *HP Data Protector integration guide for IBM applications: Informix, DB2, and Lotus Notes/Domino*

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

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

    This guide describes the integrations of Data Protector with VMware Virtual Infrastructure, Sybase, Network Node Manager, Network Data Management Protocol Server, and Citrix XenServer.

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

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

- *HP Data Protector integration guide for HP Reporter*

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

- *HP Data Protector integration guide for HP Operations Manager for UNIX*

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

- *HP Data Protector integration guide for HP Operations Manager for Windows*

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

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

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

- *HP Data Protector zero downtime backup concepts guide*

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

- *HP Data Protector zero downtime backup administrator's guide*

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

- *HP Data Protector zero downtime backup integration guide*

  This guide describes how to configure and use Data Protector to perform zero downtime backup, instant recovery, and standard restore of Oracle, SAP R/3, Microsoft Exchange Server, and Microsoft SQL Server databases. The guide also describes how to configure and use Data Protector to perform backup and restore using the Microsoft Volume Shadow Copy Service.

- *HP Data Protector MPE/iX system user guide*

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

- *HP Data Protector Media Operations user guide*

  This guide provides tracking and management of offline storage media. It describes the tasks of installing and configuring the application, performing daily media operations and producing reports.

- *HP Data Protector product announcements, software notes, and references*

  This guide gives a description of new features of HP Data Protector A.06.11. It also provides information on installation requirements, required patches, and limitations, as well as known issues and workarounds.

- *HP Data Protector product announcements, software notes, and references for integrations to HP Operations Manager, HP Reporter, HP Performance Manager, HP Performance Agent, and HP Service Information Portal*

  This guide fulfills a similar function for the listed integrations.

- *HP Data Protector Media Operations product announcements, software notes, and references*

  This guide fulfills a similar function for Media Operations.

- *HP Data Protector command line interface reference*

  This guide describes the Data Protector command-line interface, command options and their usage as well as provides some basic command-line examples.

## Online Help

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

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

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

# Documentation map

## Abbreviations

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

| Abbreviation | Guide |
|---|---|
| CLI | Command line interface reference |
| Concepts | Concepts guide |
| DR | Disaster recovery guide |
| GS | Getting started guide |
| Help | Online Help |
| IG-IBM | Integration guide for IBM applications: Informix, DB2, and Lotus Notes/Domino |
| IG-MS | Integration guide for Microsoft applications: SQL Server, SharePoint Portal Server, Exchange Server, and Volume Shadow Copy Service |
| IG-O/S | Integration guide for Oracle and SAP |
| IG-OMU | Integration guide for HP Operations Manager for UNIX |
| IG-OMW | Integration guide for HP Operations Manager for Windows |
| IG-PM/PA | Integration guide for HP Performance Manager and HP Performance Agent |
| IG-Report | Integration guide for HP Reporter |
| IG-SIP | Integration guide for HP Service Information Portal |
| IG-Var | Integration guide for VMware Virtual Infrastructure, Sybase, Network Node Manager, Network Data Management Protocol Server, and Citrix XenServer. |

| Abbreviation | Guide |
|---|---|
| Install | Installation and licensing guide |
| MO GS | Media Operations getting started guide |
| MO RN | Media Operations product announcements, software notes, and references |
| MO UG | Media Operations user guide |
| MPE/iX | MPE/iX system user guide |
| PA | Product announcements, software notes, and references |
| Trouble | Troubleshooting guide |
| ZDB Admin | ZDB administrator's guide |
| ZDB Concept | ZDB concepts guide |
| ZDB IG | ZDB integration guide |

## Map

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

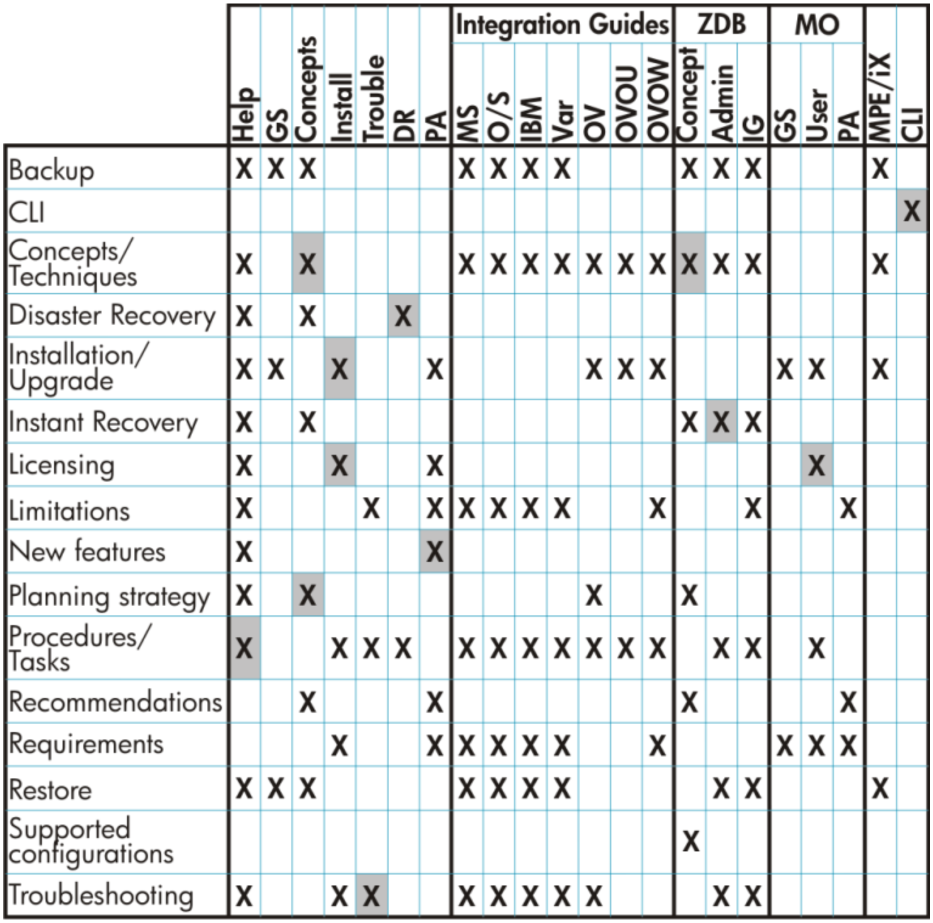

| | | | | | | | | Integration Guides | | | | | | | ZDB | | | MO | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | Help | GS | Concepts | Install | Trouble | DR | PA | MS | O/S | IBM | Var | OV | OVOU | OVOW | Concept | Admin | IG | GS | User |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

## Integrations

Look in these guides for details of the following integrations:

| Integration | Guide |
|---|---|
| HP Operations Manager for UNIX/for Windows | IG-OMU, IG-OMW |
| HP Performance Manager | IG-PM/PA |
| HP Performance Agent | IG-PM/PA |

| Integration | Guide |
|---|---|
| HP Reporter | IG-R |
| HP Service Information Portal | IG-SIP |
| HP StorageWorks Disk Array XP | all ZDB |
| HP StorageWorks Enterprise Virtual Array (EVA) | all ZDB |
| HP StorageWorks Virtual Array (VA) | all ZDB |
| IBM DB2 UDB | IG-IBM |
| Informix | IG-IBM |
| Lotus Notes/Domino | IG-IBM |
| Media Operations | MO User |
| MPE/iX system | MPE/iX |
| Microsoft Exchange Server | IG-MS, ZDB IG |
| Microsoft Exchange Single Mailbox | IG-MS |
| Microsoft SQL Server | IG-MS, ZDB IG |
| Microsoft Volume Shadow Copy Service (VSS) | IG-MS, ZDB IG |
| NDMP Server | IG-Var |
| Network Node Manager (NNM) | IG-Var |
| Oracle | IG-O/S |
| Oracle ZDB | ZDB IG |
| SAP DB | IG-O/S |
| SAP R/3 | IG-O/S, ZDB IG |

| Integration | Guide |
|---|---|
| Sybase | IG-Var |
| EMC Symmetrix | all ZDB |
| VMware | IG-Var |

# Document conventions and symbols

**Table 2 Document conventions**

| Convention | Element |
|---|---|
| Blue text: Table 2 on page 27 | Cross-reference links and e-mail addresses |
| Blue, underlined text: http://www.hp.com | Website addresses |
| *Italic* text | Text emphasis |
| `Monospace` text | - File and directory names<br>- System output<br>- Code<br>- Commands, their arguments, and argument values |
| *`Monospace, italic`* text | - Code variables<br>- Command variables |
| **`Monospace, bold`** text | Emphasized monospace text |

△ CAUTION:

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

IMPORTANT:

Provides clarifying information or specific instructions.

NOTE:

Provides additional information.

TIP:

Provides helpful hints and shortcuts.

## Data Protector graphical user interface

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

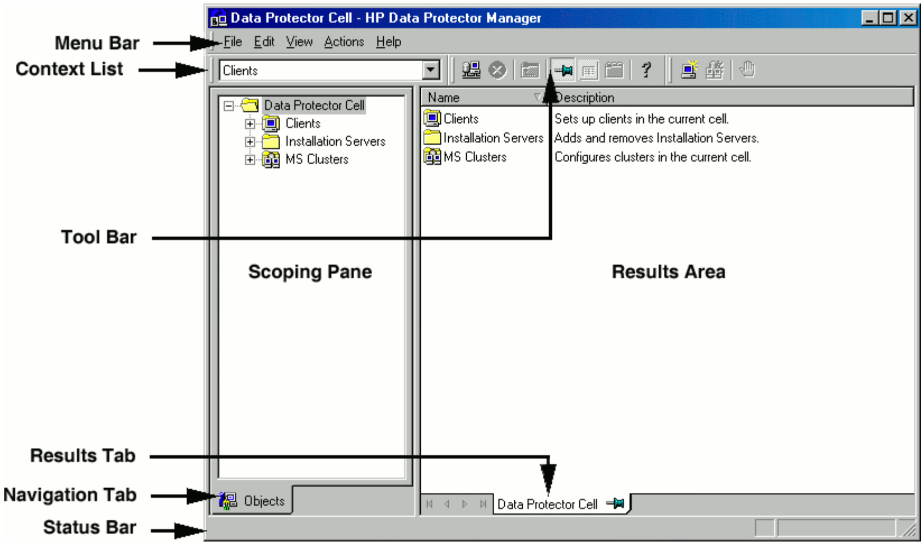

| Name | Description |
|---|---|
| Clients | Sets up clients in the current cell. |
| Installation Servers | Adds and removes Installation Servers. |
| MS Clusters | Configures clusters in the current cell. |

**Figure 1 Data Protector graphical user interface**

# General information

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

# HP technical support

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

http://www.hp.com/support

Before contacting HP, collect the following information:

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

# Subscription service

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

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

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

# HP websites

For additional information, see the following HP websites:

- http://www.hp.com
- http://www.hp.com/go/software
- http://www.hp.com/support/manuals
- http://h20230.www2.hp.com/selfsolve/manuals
- http://www.hp.com/support/downloads

# Documentation feedback

HP welcomes your feedback.

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

# 1 Data Protector Oracle ZDB integration

## Introduction

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

### Supported disk arrays

The following disk arrays can be used for ZDB of Oracle:

- EMC Symmetrix (EMC)
- HP StorageWorks Disk Array XP (XP)
- HP StorageWorks Virtual Array (VA)
- HP StorageWorks Enterprise Virtual Array (EVA)

NOTE:

With the Data Protector EMC integration, only ZDB to tape is supported. Consequently, instant recovery is not supported.

### Advantages

The advantages of using Data Protector Oracle ZDB integration are:

- ZDB reduces the performance degradation of the application system.

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

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

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

### ZDB methods and Oracle versions

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

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

## Backup and restore types

### Backup

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

- Online ZDB to disk, ZDB to tape, and ZDB to disk+tape.

  During the creation of a replica, the database on the application system is in hot backup mode. If a ZDB-to-tape or a ZDB-to-disk+tape session is being performed, the streaming of the data to tape media is subsequently performed on the backup system.
- Offline ZDB to disk, ZDB to tape, and ZDB to disk+tape.

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

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

IMPORTANT:

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

NOTE:

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

## Restore

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

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

Table 3 on page 34 provides an overview of recovery methods, depending on the type of backup that was performed and type of recovery required.

### Table 3 Oracle recovery methods

| Disk array | Backup type | Recover the whole database until | | Recover a part of database until now |
|---|---|---|---|---|
| | | Now | A point in time, logseq/thread, or SCN number | |
| **XP, VA, EVA, EMC** | **ZDB to tape - online** | Restore | Restore | Restore |
| | **ZDB to tape - offline** | Restore | Restore[1] | Restore |
| **XP, VA, EVA** | **ZDB to disk - online** | Instant recovery+ database recovery | Instant recovery+ database recovery | N/A |
| | **ZDB to disk - offline** | Instant recovery | Instant recovery+ database recovery[1] | N/A |
| | **ZDB to disk+tape - online** | • Restore<br>or<br>• Instant recovery+ database recovery | • Restore<br>or<br>• Instant recovery+ database recovery | Restore |
| | **ZDB to disk+tape - offline** | • Restore<br>or<br>• Instant recovery | • Restore<br>or<br>• Instant recovery+ database recovery[1] | Restore |

[1] The database must be put in archive mode

## Legend

*Restore*

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

*Instant recovery + database recovery*

The following three options are possible:

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

*Instant recovery*

Perform instant recovery without database recovery.

See the *HP Data Protector zero downtime backup concepts guide* for an overview of ZDB concepts and terminology.

# Integration concepts

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

## Components

The software components involved in backup and restore processes are:

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

## Integration functionality overview

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

- Database startup and shutdown
- Backups (backup and copy)

- Recovery (restore, recovery, and duplication)

## How does the integration work?

`Ob2rman.pl` executes RMAN, which directs the Oracle server processes on the target database to perform backup, restore and recovery. RMAN maintains the required information about the target databases in the recovery catalog, the Oracle central repository of information, and in the control file of a particular target database.

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

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

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

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

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

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

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

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

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

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

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

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

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

## Non-ZDB flow

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

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

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

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

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

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

## Restore flow

A restore session can be started using:

- Data Protector GUI
- RMAN CLI

You must specify which objects are to be restored.

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

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

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

The concept of Oracle integration, data and the control flow are shown in Figure 2 on page 38, and the related terms are explained in the following table.

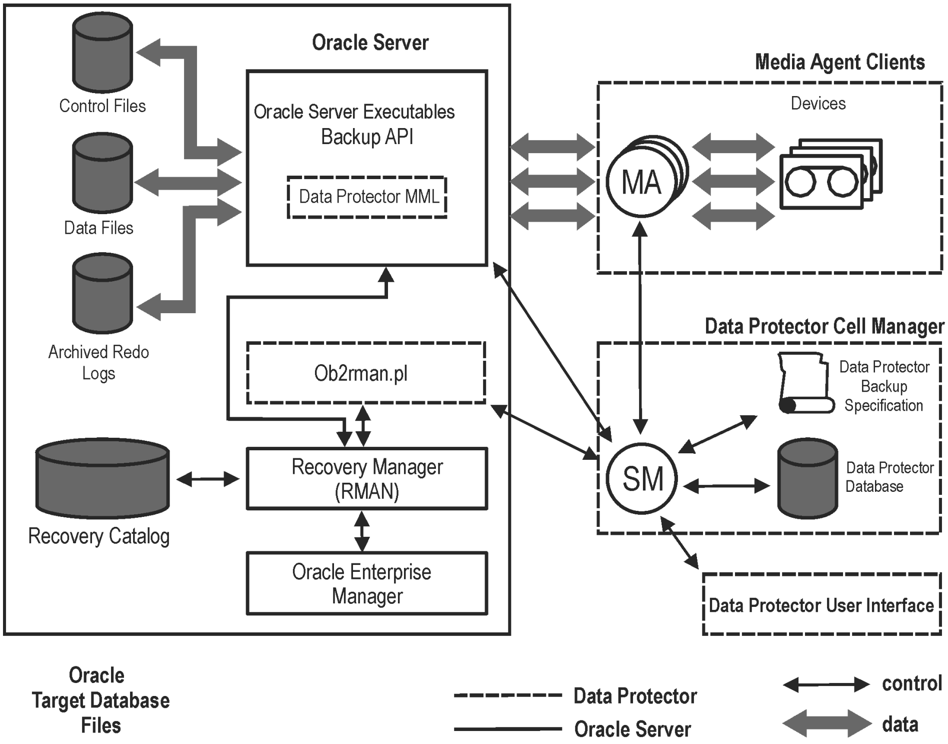

**Figure 2 Data Protector Oracle integration concept**

Oracle 10g/11g database files can also be part of ASM configuration.

### Legend

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

*RMAN*
The Oracle Recovery Manager.

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

*Backup API*
The Oracle-defined application programming interface.

*IDB*
The IDB where all the information about Data Protector sessions, including session messages, objects, data, used devices, and media is written.

*MA*
The Data Protector General Media Agent, which reads and writes data from and to media devices.

## Oracle backup set ZDB concepts

See the *HP Data Protector zero downtime backup concepts guide* for a general description of ZDB-to-disk, ZDB-to-tape, and ZDB-to-disk+tape and instant recovery concepts.

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

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

- Oracle control file, online redo log files, and SPFILE reside on a **different** volume group (if LVM is used) or source volume than Oracle datafiles.

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

By default, instant recovery for such a configuration is *not* enabled. You can enable instant recovery by setting the ZDB_ORA_INCLUDE_CF_OLF, ZDB_ORA_INCLUDE_SPF, and ZDB_ORA_NO_CHECKCONF_IR omnirc variables to 1. See "ZDB integrations omnirc variables" on page 425.

IMPORTANT:

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

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

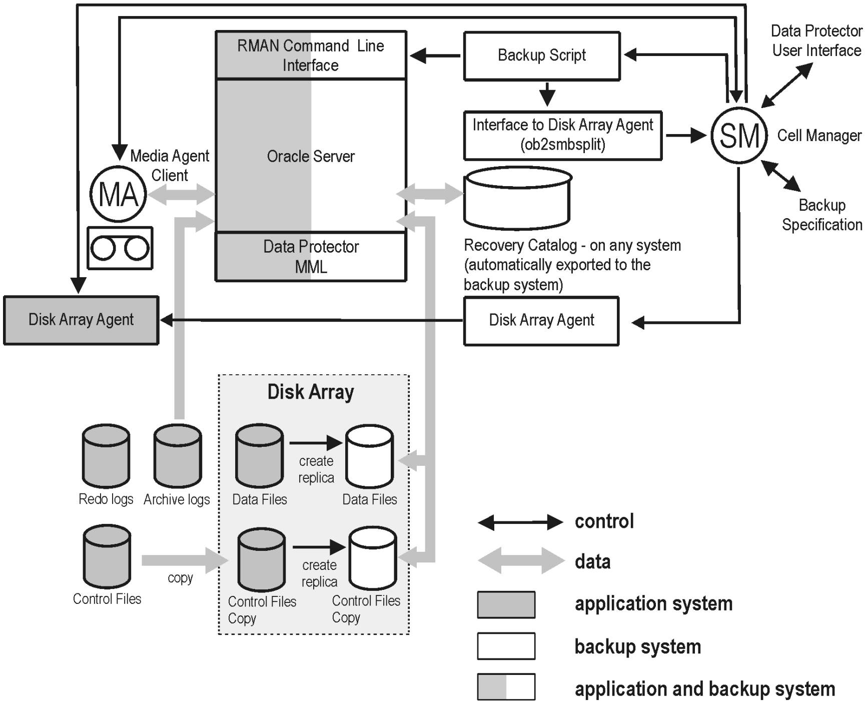

## Figure 3 Oracle backup set ZDB concept

Figure 3 on page 40 presents only the default integration behavior, where Oracle control file, online redo log files, and SPFILE reside on a different volume group (if LVM is used) or source volume than Oracle datafiles. Oracle 10g/11g database files can also be part of ASM configuration, but see the limitation for ASM configuration.

For more information on an alternative Oracle backup and restore concept, see “ZDB integrations omnirc variables” on page 425.

## Legend

*MA*

The General Media Agent writes data from a replica to backup media. The General Media Agent typically resides on the backup system.

*SM*

The session manager controls backup and restore sessions and writes session information to the IDB.

*Disk Array Agent*

The disk array Agents (ZDB Agents) are SYMA (on EMC), SSEA (on XP), SNAPA (on VA), and SMISA (on EVA).

*Data Protector MML*

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

## Backup process

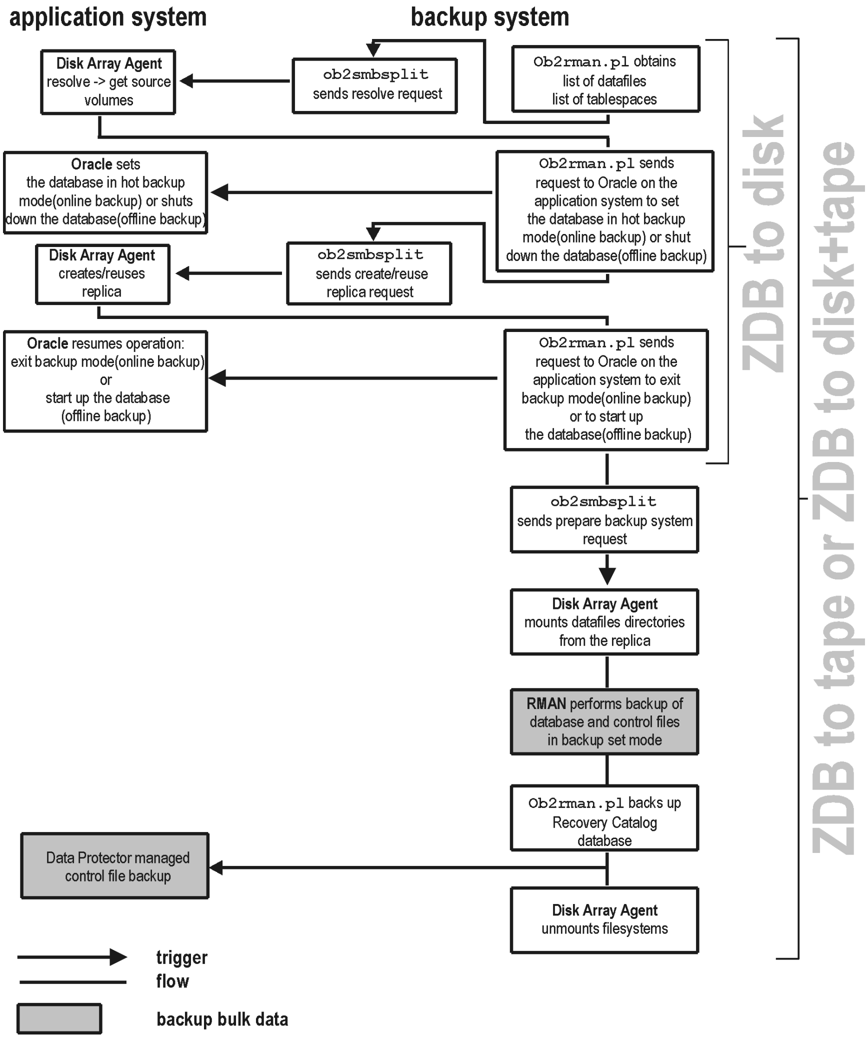
application system
backup system
Disk Array Agent resolve -> get source volumes
ob2smbsplit sends resolve request
Ob2rman.pl obtains list of datafiles list of tablespaces
Oracle sets the database in hot backup mode(online backup) or shuts down the database(offline backup)
Ob2rman.pl sends request to Oracle on the application system to set the database in hot backup mode(online backup) or shut down the database(offline backup)
Disk Array Agent creates/reuses replica
ob2smbsplit sends create/reuse replica request
Oracle resumes operation: exit backup mode(online backup) or start up the database (offline backup)
Ob2rman.pl sends request to Oracle on the application system to exit backup mode(online backup) or to start up the database(offline backup)
ob2smbsplit sends prepare backup system request
Disk Array Agent mounts datafiles directories from the replica
RMAN performs backup of database and control files in backup set mode
Ob2rman.pl backs up Recovery Catalog database
Data Protector managed control file backup
Disk Array Agent unmounts filesystems
ZDB to disk
ZDB to tape or ZDB to disk+tape
trigger
flow
backup bulk data

**Figure 4 Oracle backup set ZDB flow**

NOTE:

ZDB Agents are SYMA on EMC, SSEA on XP, SNAPA on VA, and SMISA on EVA.

See the *HP Data Protector zero downtime backup concepts guide* for a general description of ZDB and instant recovery concepts.

See the *HP Data Protector zero downtime backup administrator's guide* for a general description of the ZDB-to-disk, ZDB-to-tape, and ZDB-to-disk+tape session flows and for the explanation of actions triggered by ZDB options.

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

Operations on a replica (mounting, activating volume/disk groups,...) described below are dependent on or triggered by ZDB options. See the *HP Data Protector zero downtime backup administrator's guide* for more information on these options.

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

  When performing an *offline* ZDB session, `ob2rman.pl` shuts down the Oracle database, starts the procedure to create a replica of the source volumes on which the database is installed; and after the replica is created, starts the Oracle database.
- `Ob2rman.pl` then starts the procedure to prepare the replica on the backup system. In this step, volume/disk groups on the backup system are enabled and, unless the database is installed on raw partitions, the mount points with the Oracle database files are mounted.
- A ZDB Agent then mounts the database on the backup system to the mount points with the same names (created by Data Protector) as on the application system.

NOTE:

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

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

NOTE:

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

# Oracle proxy-copy ZDB concepts

See the *HP Data Protector zero downtime backup concepts guide* for a general description of ZDB-to-disk, ZDB-to-tape, ZDB-to-disk+tape, and instant recovery concepts.

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

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

- Oracle control file, online redo log files, and SPFILE reside on a **different** volume group (if LVM is used) or source volume than Oracle datafiles.

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

  By default, instant recovery is *not* enabled, even if this option is selected in the GUI. You can enable instant recovery by setting the `ZDB_ORA_INCLUDE_CF_OLF`, `ZDB_ORA_INCLUDE_SPF`, and `ZDB_ORA_NO_CHECKCONF_IR` `omnirc` variables to `1`. See "ZDB integrations omnirc variables" on page 425.

IMPORTANT:

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

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

Figure 5 on page 46 shows the architecture of the Data Protector Oracle ZDB integration. The figure illustrates the configuration, in which the backup isperformed on the backup system. It presents the default integration behavior, where Oracle control file, online redo log files, and SPFILE reside on different disk array source volumes than the Oracle data files. For more information on alternative Oracle backup and restore concepts, see "ZDB integrations omnirc variables" on page 425.

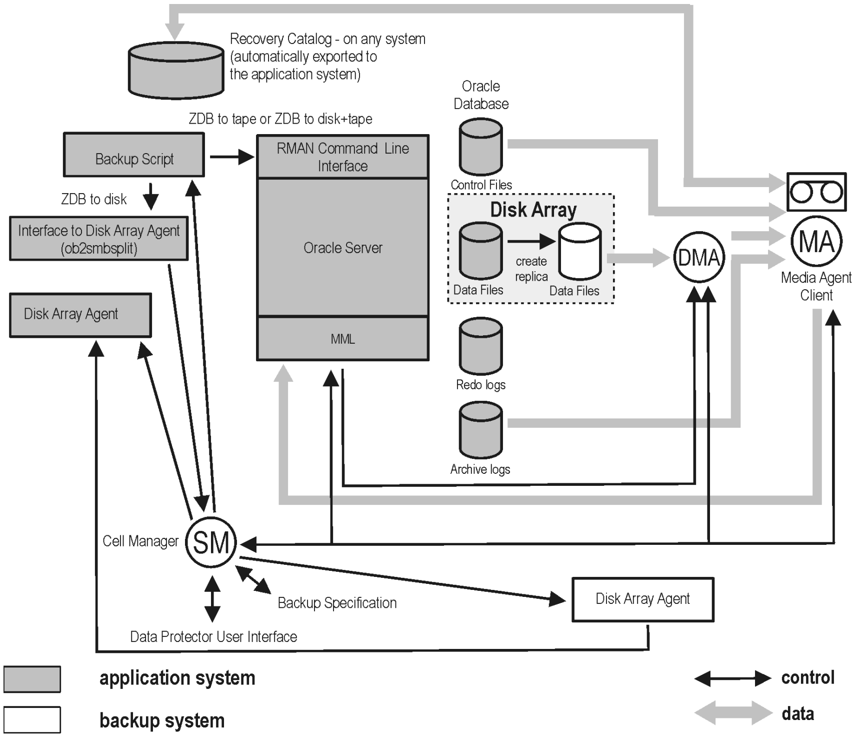

## Figure 5 Oracle proxy-copy ZDB concept

*MA*

The General Media Agent writes data from a replica to backup media. The General Media Agent typically resides on the backup system.

*SM*

The session manager controls backup and restore sessions and writes the session information to the IDB.

*Disk Array Agent*

The disk array Agents (ZDB Agents) are SYMA (on EMC), SSEA (on XP), SNAPA (on VA), and SMISA (on EVA).

*MML*

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

## Backup process

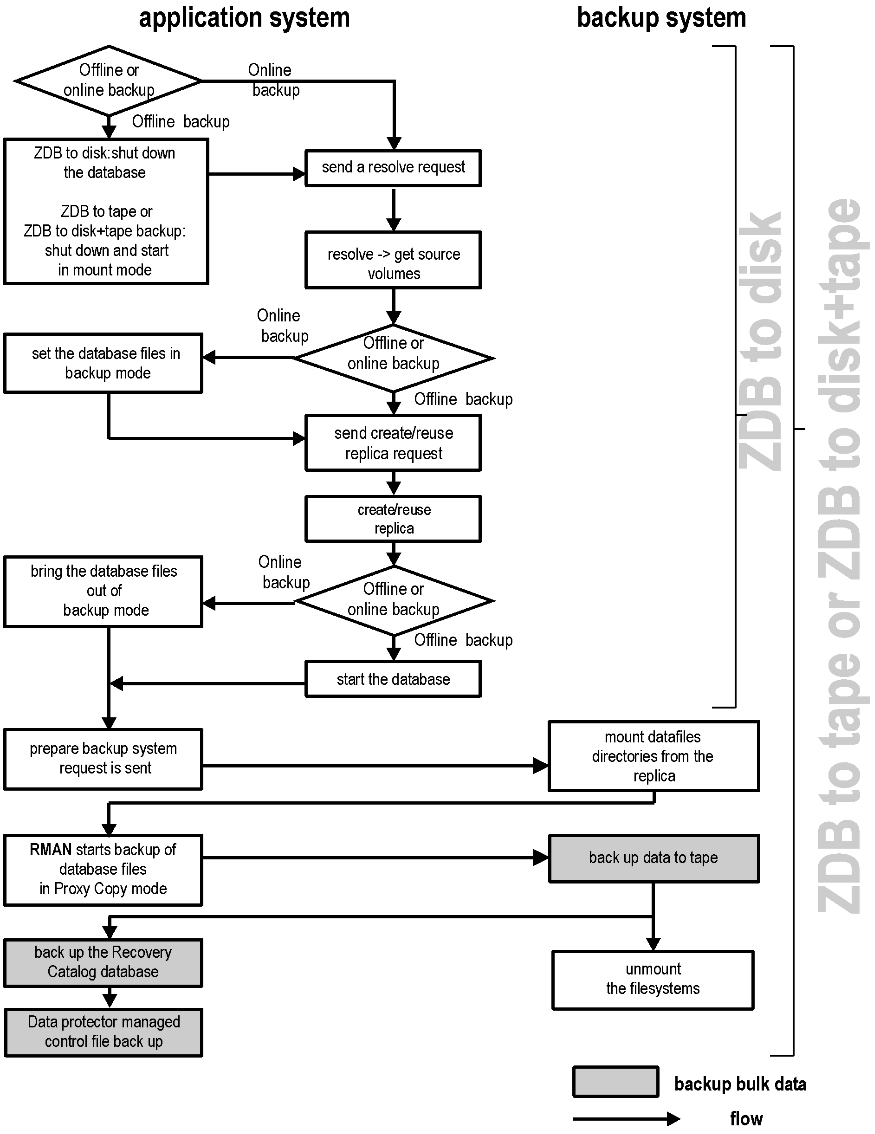
application system
backup system
Offline or online backup
Online backup
Offline backup
ZDB to disk:shut down the database
ZDB to tape or ZDB to disk+tape backup: shut down and start in mount mode
send a resolve request
resolve -> get source volumes
Online backup
Offline or online backup
Offline backup
set the database files in backup mode
send create/reuse replica request
create/reuse replica
Online backup
Offline or online backup
Offline backup
bring the database files out of backup mode
start the database
prepare backup system request is sent
mount datafiles directories from the replica
RMAN starts backup of database files in Proxy Copy mode
back up data to tape
back up the Recovery Catalog database
unmount the filesystems
Data protector managed control file back up
ZDB to disk
ZDB to tape or ZDB to disk+tape
backup bulk data
flow

**Figure 6 Oracle proxy-copy ZDB flow**

See the *HP Data Protector zero downtime backup administrator's guide* for a general description of the ZDB-to-disk, ZDB-to-tape, and ZDB-to-disk+tape sessions flows and for an explanation of actions triggered by ZDB options.

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

Operations on a replica (mounting, activating volume/disk groups...) described below are dependent on or triggered by ZDB options. See the *HP Data Protector zero downtime backup administrator's guide* for more information on these options.

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

  In the case of an *offline* ZDB-to-disk session, the database is shut down.
- Data Protector retrieves a list of files or raw disks to be included in the replica creation from the Oracle database and starts the resolving process. The list is used only to determine the source volumes to be replicated.

  In the case of a *ZDB-to-disk* session, if the location for control file copy is specified during configuration, Data Protector makes a copy of the control file to the specified directory on the application system. This directory has to reside on a disk array source volume.
- In the case of an *online* backup, the Oracle target database is switched into the backup mode.
- `ob2smbsplit` or MML starts the procedure to create a replica of the source volumes on which the database is installed.
- In the case of an *online* backup, the database files are taken out of the backup mode after the replica is created.

  In the case of an *offline* backup, the Oracle database is started by `ob2rman.pl` after the replica is created.
- Data Protector (for ZDB to disk) or MML (for ZDB to tape or ZDB to disk+tape) starts the procedure to prepare the replica on the backup system. In this step, volume/disk groups on the backup system are enabled (UNIX systems) and, unless the database is installed on raw disks, the mount points containing the Oracle database files are mounted.
- A ZDB Agent then mounts the database on the backup system to the mount points with the same names (created by Data Protector) as on the application system.
- If a ZDB-to-disk session is being performed, at this point the remaining ZDB options are processed and details of the session are written to the ZDB database. The session then finishes. The following steps in this description are not performed; therefore, RMAN is not given any information about the ZDB-to-disk session.
- MML on the application system sends a request to the Data Protector **data movement agent** (DMA) on the backup system to back up the datafiles to tape.

- The DMA reads data from the backup system and sends it to the General Media Agent to write the actual data to the backup device.

  DMA's role is also to disable the General Media Agent requests from accessing the application system. Thus, the database runs on the application system with greatly reduced performance degradation since the backup is performed on the backup system.
- At the end of data transfer, the backup system is disabled (filesystems are unmounted on all platforms and volume/disk groups deactivated on UNIX) and links are re-established.
- The recovery catalog and the control file are backed up automatically after the target database backup is finished on the backup system. However, you can disable this when creating a backup specification.

NOTE:

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

# Configuring the integration

## Prerequisites

- It is assumed that you are familiar with the Oracle database administration and the basic Data Protector functionality.
- You need a license to use the Data Protector ZDB integration with Oracle. Additional licenses are required for instant recovery and for the online extension. See the *HP Data Protector installation and licensing guide* for information about licensing.
- Before you begin, ensure that you have correctly installed and configured the Oracle Server and Data Protector systems. See the:
  - For an up-to-date list of supported versions, platforms, devices, and other information, see the latest support matrices at http://www.hp.com/support/manuals.
  - *HP Data Protector installation and licensing guide* for instructions on how to install Data Protector on various architectures and how to install a Data Protector disk array integration (EMC, XP, VA, or EVA) with Oracle.

    - *Oracle Recovery Manager User's Guide and References* for Oracle concepts and backup/recovery strategies.
    - *Oracle Backup and Recovery Guide* for the configuration and use of Recovery Manager, as well as for Oracle backup terminology and concepts.
    - *Oracle Enterprise Manager User's Guide* for information about backup and recovery with the Oracle Enterprise Manager, as well as information about SQL*Plus.

- A Data Protector disk array integration (EMC, XP, VA, or EVA) must be correctly installed and configured. For installation, see the *HP Data Protector installation and licensing guide*. For configuration, see the *HP Data Protector zero downtime backup administrator's guide*.
- The Oracle Server software must be installed on the application system and the Oracle target database must be open or mounted there.
- The Oracle recovery catalog database must be properly configured and open.
- Oracle net services must be properly configured and running (on the application system) for the Oracle target database and the recovery catalog. The net services are needed for the Data Protector Oracle agent to be connected to the Oracle database on the application system through Oracle.

  See the *Oracle Recovery Manager User's Guide and References* for more information about different connection options.

  See "Troubleshooting" on page 145 for details about how to check the prerequisites listed above.

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

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

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

- In case of Real Application Cluster (RAC), each node must have a dedicated disk for storing archive logs. Such disks must be NFS mounted on all other RAC nodes.

  However, if the archive logs are not on a NFS mounted disk, you must modify the archive log backup specification. See Problem on page 159.

## Limitations

- The MAXPIECESIZE RMAN parameter option is not supported because the restore of multiple backup pieces created during a backup is not possible using the Data Protector Oracle integration.
- The Oracle recovery catalog database must be used as RMAN repository for backup and restore operations. ZDB using the Oracle control file are not supported. This is set when configuring the database. See "Configuring Oracle databases" on page 60.
- The Oracle database identifier (DBID) must be a unique in a Data Protector cell. If you clone a database you must change the DBID.
- Preview is not possible for Oracle ZDB and restore sessions.
- Using the Oracle proxy-copy ZDB method, individual tablespaces or datafiles cannot be backed up during a ZDB-to-disk or ZDB-to-disk+tape session (instant recovery enabled), only the whole database can be backed up.
- The Oracle backup set ZDB method is not supported on Windows.
- The Oracle backup set ZDB method is supported on UNIX raw logical volumes only if these were created with LVM or VxVM.
- The single host configuration (BC1, TF/1) is not supported for Oracle backup set ZDB sessions.
- Object copying and object mirroring is not supported for ZDB to disk.
- Recovery files residing in the **flash recovery area** (Oracle 10g/11g only) cannot be backed up using ZDB.
- For backing up the ASM files, the following limitations apply:
  - Only the backup set ZDB method is supported.
  - On VA and EVA, the ASM files can be backed up using ZDB to tape only, because instant recovery for the ASM files is not supported on these arrays.
  - On XP, instant recovery is supported only if atomic split configuration is used. For more information, see Enabling instant recovery in an ASM configuration for the XP array.
- ***Oracle Data Guard:*** Standby database is not supported for ZDB.

## Before you begin

- Test whether the Oracle Server system and the Cell Manager communicate properly: Configure and run a Data Protector filesystem backup and restore on the Oracle Server system.

- Identify the Oracle database *user* that will be used by Data Protector for backup. This user must have the `SYSDBA` privilege granted. For example, it could be the Oracle user `sys`, which is created during database creation.

  See the Oracle documentation for more information on user privileges in Oracle.
- In case of the backup set method, if the Oracle database is installed on symbolic links, create these symbolic links on the backup system too.
- If an Oracle 10g/11g **automatic storage management** (**ASM**) instance manages files for more than one database, it is recommended to create a separate ASM disk group for each database.
- From the application system, using SQL*Plus, connect to the target database and recovery catalog by specifying the user, password, and net connect identifier. Connect to the target database as the database administrator and to the recovery catalog database as the recovery catalog owner.

### Example

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

```
sqlplus /nolog
```

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

- For *online backup* only, enable the Oracle automatic log archiving:
  1. Shut down the Oracle target database instance on the application system.
  2. Back up the entire database using a filesystem backup.

3. Select the location for archive logs:
    - If SPFILE is used:

      Run:

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

      Run:

      `log_archive_start=true`

      `log_archive_dest=`*`path_to_archive_logs`*

      The default path of the file is:

      ***Windows:*** *`ORACLE_HOME`*`\database\init`*`DB_NAME`*`.ora`

      ***UNIX:*** *`ORACLE_HOME`*`/dbs/init`*`DB_NAME`*`.ora`

      where *`DB_NAME`* is the name of the Oracle database instance.
4. Mount the target database and to enable the archive log mode, start SQL*Plus and type:

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

### Example

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

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

5. Back up the entire database.

## Backup set method

For backup set method:

- Ensure that the Oracle software on the backup system and application system have the same directory structure. That means that `ORACLE_HOME` for both Oracle installations has to be identical.
- Ensure that the following files are the same on the application system and the backup system. Check also that the permissions are identical as on the application system:
  - `names.ora`

    Default path: *`ORACLE_HOME`*`/network/admin/names.ora`
  - `init`*`DB_NAME`*`.ora`

    Default path: *`ORACLE_HOME`*`/dbs/init`*`DB_NAME`*`.ora`.
  - `orapw`*`DB_NAME`*

    Default path: *`ORACLE_HOME`*`/dbs/orapw`*`DB_NAME`*
  - `admin/`*`DB_NAME`*

    Default path: *`ORACLE_BASE`*`/admin/`*`DB_NAME`*

  Ensure that the Oracle net services on the application system and the backup system have the same directory structure. This can be accomplished by either NFS sharing of the files, manually copying the files from the application system to the backup system, or by using the UNIX `rdist` or `tar` commands to distribute the files from the application system.
- Test whether the Oracle user can log in to the Oracle target database as the Oracle database administrator and to the Oracle recovery catalog database as the Oracle recovery catalog owner from the backup system:
  1. Export `ORACLE_HOME`, `DB_NAME`, and on UNIX also `SHLIB_PATH` variables.
  2. Using SQL*Plus, connect to the Oracle recovery catalog database by specifying the user (recovery catalog owner), password, and net connect identifier.

3. Connect to the Oracle target database locally using the Oracle Net software as the Oracle database administrator with the SYSDBA role.

### Example

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

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

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

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

- Test whether the user root and the Oracle administrator (for example, the user oracle) can connect to the target database and the recovery catalog database using the RMAN command on the backup system:

1. Log in as the Oracle database administrator to the backup system (for example, the user oracle).

2. Execute the RMAN command and connect to the target database and the recovery catalog database.

### Example

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

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

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

```
RMAN> exit
Recovery Manager completed.
```

## Enabling instant recovery in an ASM configuration for the XP array

On UNIX systems, the ASM configuration on XP array can be supported for instant recovery, provided that support for the atomic split configuration on XP array is enabled. To support ASM for instant recovery, the following prerequisites must be met:

- The ASM files you will use for backup must be located on raw disks. They should not reside on raw logical volumes.
- The XP array which you will use for backing up the ASM files must support the atomic split configuration.
- The XP mirror groups must be created in the group mode with a specific CT number assigned. Note that the CT number 0 is not supported.
- The Data Protector `omnirc` variable `SSEA_ATOMIC_SPLIT` must be enabled (set to `1`). Note that by default the variable is disabled (set to `0`).

## Cluster-aware clients

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

***Windows:*** `set OB2BARHOSTNAME=`*`virtual_server_name`*

***UNIX:*** `export OB2BARHOSTNAME=`*`virtual_server_name`*

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

## Linking Oracle Server with the Data Protector MML

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

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

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

## Configuring Oracle user accounts

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

- Oracle operating system user account

  For details, see "Configuring Oracle operating system user accounts" on page 57.
- Oracle database user accounts

  For details, see "Configuring Oracle database users accounts" on page 59.

## 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 clients, run:

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

or

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

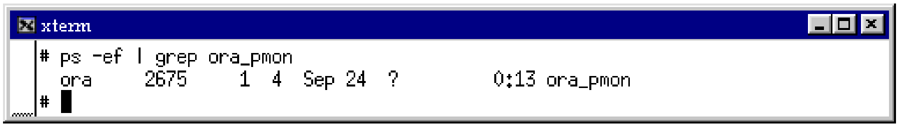

```
# ps -ef | grep ora_pmon
  ora      2675     1  4  Sep 24  ?         0:13 ora_pmon
# 
```

**Figure 7 Finding the Oracle user**

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

| Clients | Description |
|---|---|
| UNIX clients | Add the OSDBA user account and `root` user account from both the application system and backup system to the Data Protector `admin` or `operator` user group. The OSDBA user on the backup system must have the same numerical user ID and group ID as the OSDBA user on the application system (for example, `uid=101(ora) gid=101(dba)`).<br>TIP:<br>To find the user ID, connect to a client under this user account and run:<br>`#id` |
| Windows clients | On Windows clients, Data Protector connects to the Oracle database using the Data Protector `Inet` service on the related client. By default, the service runs under the `Local System account`, which is automatically added to the Data Protector `admin` user group. However, if you have restarted the Data Protector `Inet` service on the application system and backup system under OSDBA user accounts, you need to add the new users to the Data Protector `admin` or `operator` user group. |

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

NOTE:

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

## Clusters

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

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

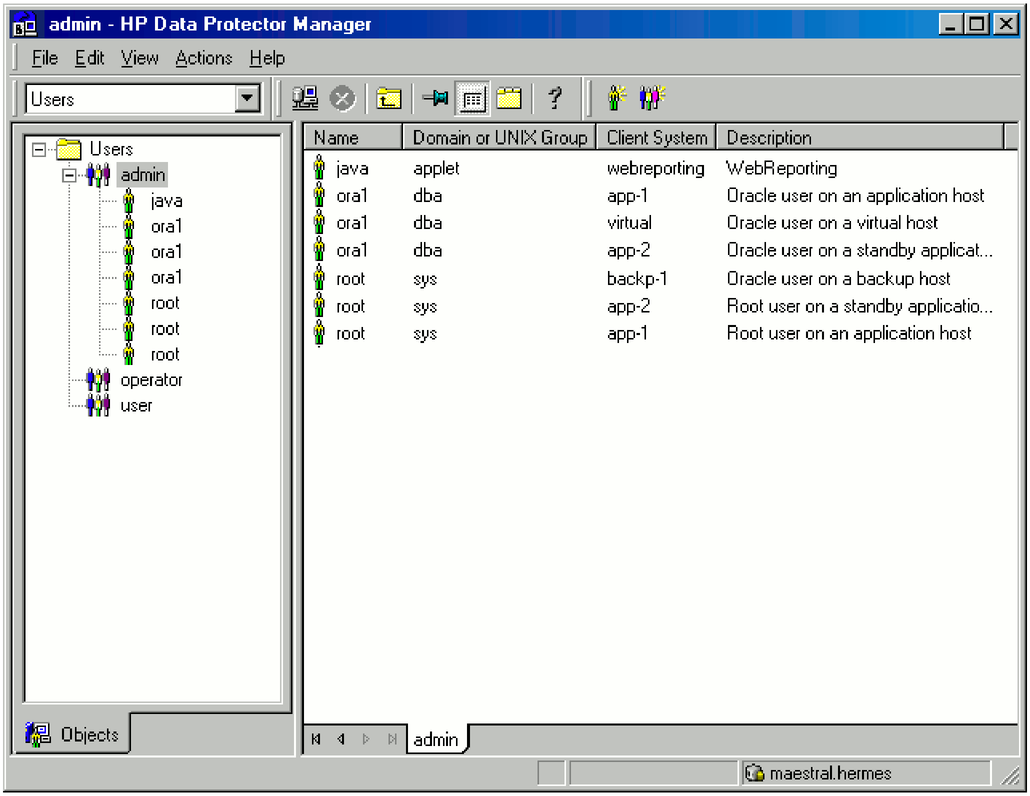

admin - HP Data Protector Manager
File Edit View Actions Help
Users
Objects
admin
maestral.hermes

- Users
  - admin
    - java
    - ora1
    - ora1
    - ora1
    - root
    - root
    - root
  - operator
  - user

| Name | Domain or UNIX Group | Client System | Description |
|---|---|---|---|
| java | applet | webreporting | WebReporting |
| ora1 | dba | app-1 | Oracle user on an application host |
| ora1 | dba | virtual | Oracle user on a virtual host |
| ora1 | dba | app-2 | Oracle user on a standby applicat... |
| root | sys | backp-1 | Oracle user on a backup host |
| root | sys | app-2 | Root user on a standby applicatio... |
| root | sys | app-1 | Root user on an application host |

**Figure 8 Example user configuration in a cluster environment**

## Configuring Oracle database users accounts

Identify or create the following Oracle database user accounts. You need to provide these user accounts when you configure the Oracle database as described in "Configuring Oracle databases" on page 60.

**Table 4 Oracle database user accounts**

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

| User | Description |
|---|---|
| Recovery catalog user | The owner of the recovery catalog (for example, `rman`). Required to log in to the catalog database. Needed if you use the recovery catalog. |
| Standby database user | Required to log in to the standby database. Applicable only in Oracle Data Guard environments. Needed to back up the standby database. |

## Configuring Oracle databases

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

- Oracle Server home directory
- Login information to the target database
- Optionally, login information to the recovery catalog database
- Optionally, login information to the standby database
- Backup method to be used and the related options

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

Ensure that the database is open during the configuration procedure.

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

### Using the Data Protector GUI

Configure an Oracle database when you create the first ZDB backup specification for the database. Start with the procedure described in “Creating backup specifications” on page 76 and at Step 6 on page 86 proceed as follows:

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

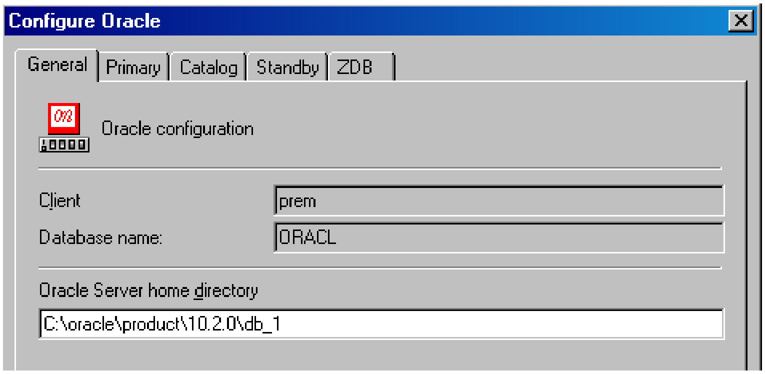

Client: prem

Database name: ORACL

Oracle Server home directory: C:\oracle\product\10.2.0\db_1

Figure 9 Configuring Oracle - General (Windows)

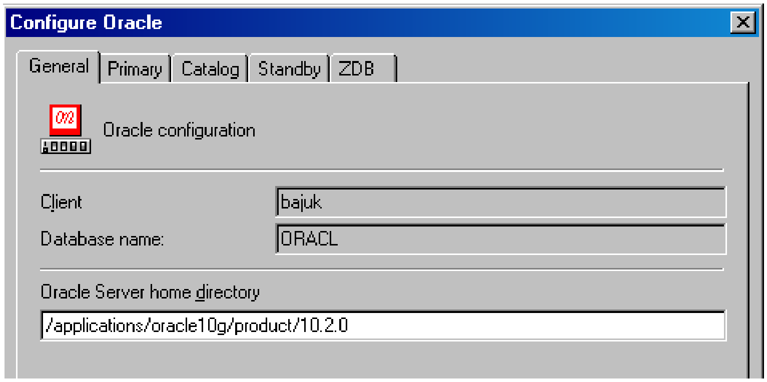

Client: bajuk

Database name: ORACL

Oracle Server home directory: /applications/oracle10g/product/10.2.0

Figure 10 Configuring Oracle - General (UNIX)

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

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

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

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

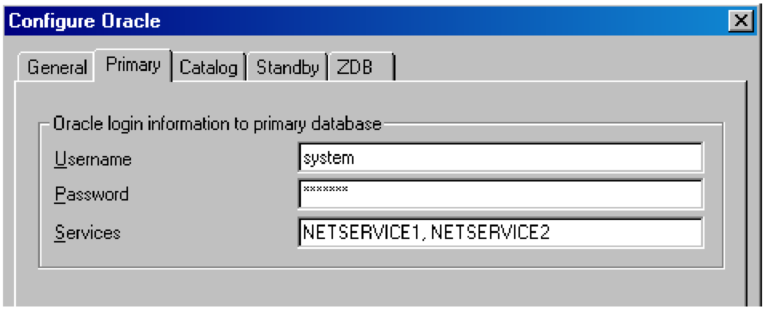

Oracle login information to primary database

Username: system

Password: *******

Services: NETSERVICE1, NETSERVICE2

**Figure 11 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.

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

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

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

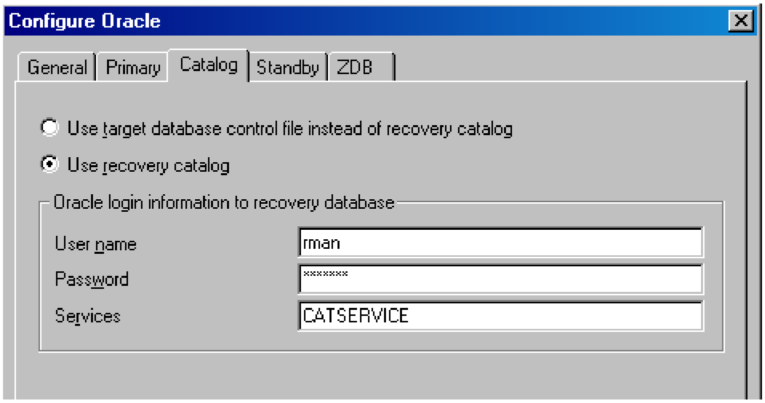

○ Use target database control file instead of recovery catalog

● Use recovery catalog

Oracle login information to recovery database

User name: rman

Password: *******

Services: CATSERVICE

Figure 12 Configuring Oracle - Catalog

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

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

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

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

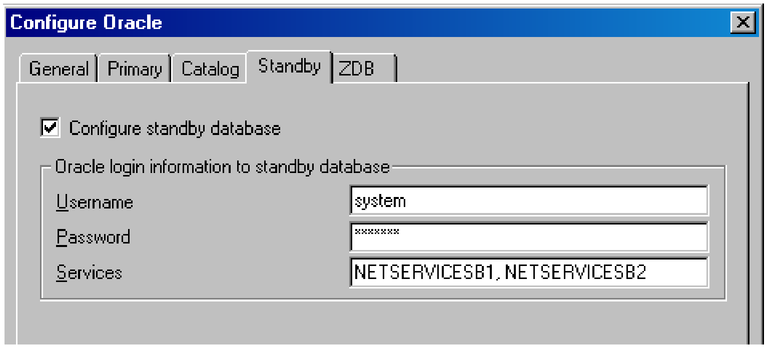

☑ Configure standby database

Oracle login information to standby database

Username: system

Password: *******

Services: NETSERVICESB1, NETSERVICESB2

Figure 13 Configuring Oracle - Standby

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

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

   If you do not specify the location, `o2rman.pl` will copy the copy of the control file from the application system to the backup system when it is needed. Thus, you do not need to create an additional disk for this location if you do not need the control file copy on a replica.

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

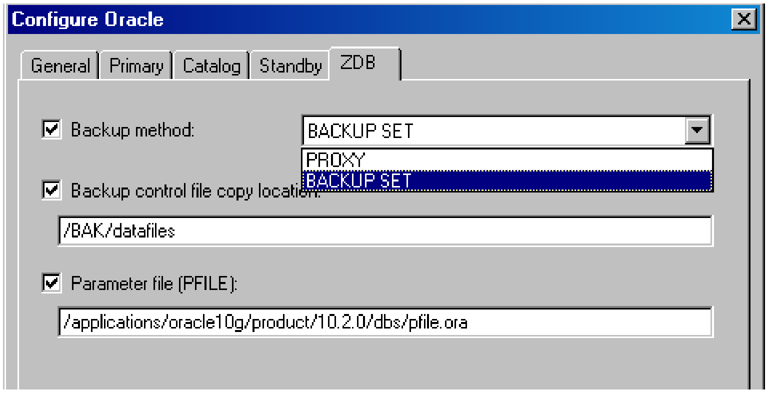

☑ Backup method: BACKUP SET

☑ Backup control file copy location:

/BAK/datafiles

☑ Parameter file (PFILE):

/applications/oracle10g/product/10.2.0/dbs/pfile.ora

**Figure 14 Configuring Oracle - ZDB**

   Click **OK**.

The Oracle database is configured. Exit the GUI or proceed with creating the backup specification at Step 7 on page 87.

## Using the Data Protector CLI

1. ***UNIX only:*** Log in to the Oracle Server system with an OSDBA user account.

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

   ***Windows:*** `Data_Protector_home\bin`

   ***HP-UX, Solaris:*** `/opt/omni/lbin`

   run:

   ***Windows:***

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

   ***UNIX:***

   ```
   util_oracle8.pl -config -dbname DB_NAME -orahome
   ORACLE_HOME PRIMARY_DB_LOGINCATALOG_DB_LOGIN
   ZDB_OPTIONS[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, ...]
```

*ZDB_OPTIONS* are:

`-zdb_method {PROXY | BACKUP_SE T}`

`[-ctlcp_location` *`BACKUP_CONTROL_FILE_COPY_LOCATION`*`]`

`[-pfile` *`PARAMETER_FILE`*`]`

`[-bkphost` *`BACKUP_SYSTEM`*`]`

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

If your ZDB method is *backup set*, you must provide the *`BACKUP_SYSTEM`* information.

## Parameter description

*`CLIENT_NAME`*

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

***RAC:*** The virtual server of the Oracle resource group.

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

*`DB_NAME`*

Name of the database to be configured.

*`ORACLE_HOME`*

Pathname of the Oracle Server home directory.

*`PRIMARY_USERNAME PRIMARY_PASSWORD`*

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

*`primary_net_service_name_1`*` [,`*`primary_net_service_name_2`*`, ...]`

Net services names for the primary database.

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

*`CATALOG_USERNAME CATALOG_PASSWORD`*

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

*`catalog_net_service_name`*

Net service name for the recovery catalog.

*`STANDBY_USERNAME STANDBY_PASSWORD`*

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

*`standby_net_service_name_1`*` [,`*`standby_net_service_name_2`*`, ...]`

Net services names for the standby database.

*BACKUP_CONTROL_FILE_COPY_LOCATION*

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

*PARAMETER_FILE*

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

*BACKUP_SYSTEM*

Name of the backup system.

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

## Example

The following example represents configuration on HP-UX or Solaris of an Oracle database and its recovery catalog with the backup set method used and the parameter file location specified.

The following names are used in the example:

- database name: `oracl`
- primary user name: `system`
- primary password: `manager`
- primary net service name 1: `netservice1`
- primary net service name 2: `netservice2`
- recovery catalog user name: `rman`
- recovery catalog password: `manager`
- recovery catalog net service name: `catservice`
- backup system name: `bcksys`

## Syntax

```
/opt/omni/lbin/util_oracle8.pl -config -dbname oracl -orahome
/app10g/oracle10g/product/10.1.0 -prmuser system -prmpasswd
manager -prmservice netservice1,netservice2 -rcuser rman
-rcpasswd manager -rcservice catservice -zdb_method BACKUP_SET
```

```
-pfile /app10g/oracle10g/product/10.1.0/dbs/pfile.ora -bkphost
bcksys
```

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

## Checking the configuration

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

### Using the Data Protector GUI

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

IMPORTANT:

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

### Using the Data Protector CLI

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

2. From the directory:

   ***Windows:*** `Data_Protector_home\bin`

   ***HP-UX, Solaris:*** `/opt/omni/lbin`

   run:

   ***Windows:***

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

   ***UNIX:***

   ```
   util_oracle8.pl -CHKCONF_SMB -dbname DB_NAME
   ```

### Handling errors

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

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

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

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

IMPORTANT:

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

## Checking configuration for instant recovery

Check if the Oracle configuration is suitable for instant recovery.

On the application system, from the directory:

***Windows:*** `Data_Protector_home\bin`

***HP-UX, Solaris:*** `/opt/omni/lbin`

run:

***Windows:***

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

***UNIX:***

```
util_oracle8.pl -CHKCONF_IR -dbname DB_NAME
```

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

- Reconfigure the Oracle database instance. See "Reconfiguring an Oracle instance for instant recovery" on page 421 on how to move the control files and redo logs to source volumes that are not replicated.

  Or:

- Set the `ZDB_ORA_INCLUDE_CF_OLF`, `ZDB_ORA_INCLUDE_SPF`, and `ZDB_ORA_NO_CHECKCONF_IR` `omnirc` variables and ignore the warning. However, note that the control file, SPFILE, and online redo logs are overwritten during instant recovery. See "ZDB integrations omnirc variables" on page 425 on how to set the `omnirc` variables.

## Setting environmental variables

Use environmental variables to modify backup environment to suit your needs. Environmental 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 on how environmental variables affect your environment, see Table 5.

NOTE:

Environmental variables are not supported on HP OpenVMS systems.

### Table 5 Environmental variables

| Environmental 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. |

To set environmental 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 Environmental Variables**.

2. In the Advanced dialog box, specify the variable name, its value, and click **Add**. See Figure 15.

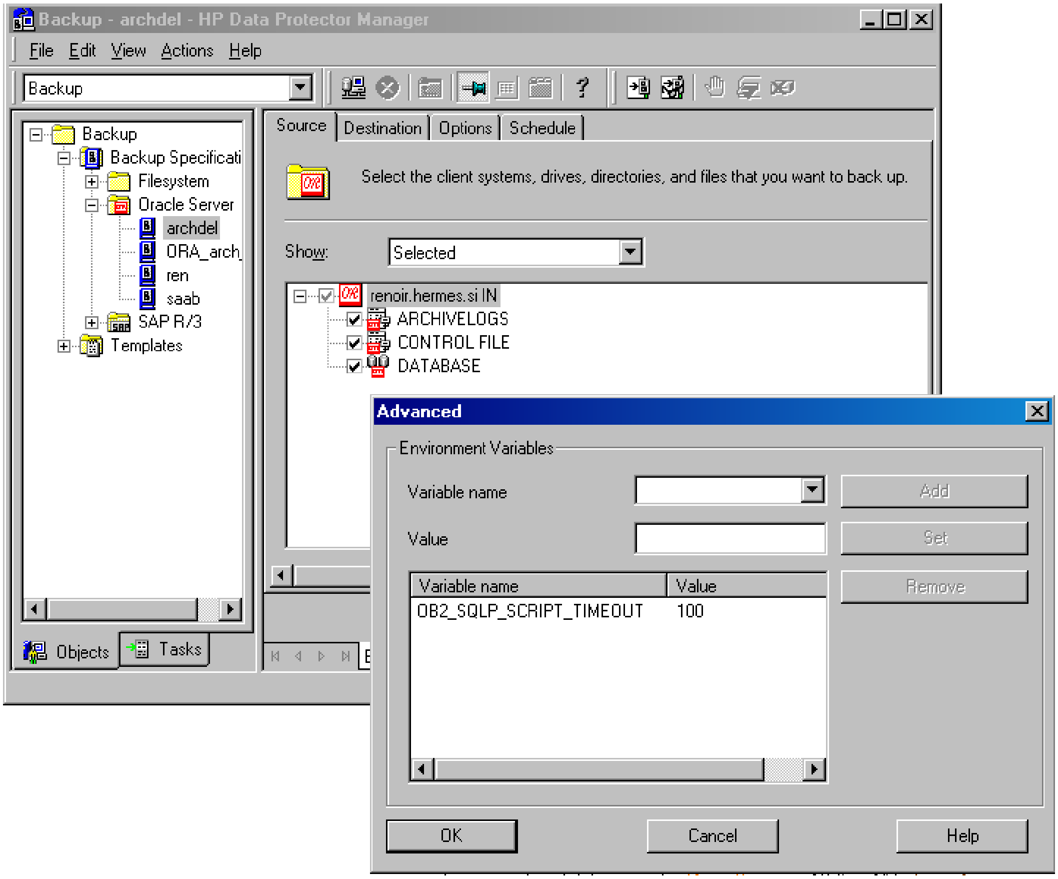

Select the client systems, drives, directories, and files that you want to back up.

Show: Selected

| Variable name | Value |
|---|---|
| OB2_SQLP_SCRIPT_TIMEOUT | 100 |

**Figure 15 Setting environmental variables**

Click **OK**.

## Using the Data Protector CLI

From the directory:

***Windows:*** `Data_Protector_home\bin`

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

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

run:

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

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

### Example

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

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

## Switching between Oracle backup methods

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

IMPORTANT:

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

To switch between the backup methods:

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

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

   ***UNIX:*** `/etc/opt/omni/server/barlists/oracle8`
3. Re-configure the database with the *new method* selected while creating a new Oracle ZDB specification.

4. Optionally, if you switch *from backup set to proxy-copy*, you may:

   a. On the Cell Manager, remove the file:

      ***Windows:*** `Data_Protector_home\Config\Server\Integ\Config\Oracle8\`*`client_name`*`%init`*`DB_NAME`*`_bckp.ora`

      ***UNIX:*** `/etc/opt/omni/server/integ/config/Oracle8/`*`client_name`*`%init`*`DB_NAME`*`_bckp.ora`

   b. Remove the Oracle software from the backup system.

5. Perform ZDB of the entire database.

IMPORTANT:

If you need to perform a restore from a time between the start and the end of the first backup of the entire database using the new backup method, RMAN may try to use backup files from old method through a channel allocated for the files from the old method and the restore will fail. See Problem on page 156 on how to restore such a backup.

# Backup

To configure an Oracle ZDB, perform the following steps:

1. Configure the devices you plan to use for a backup. See the online Help index: “configuring devices” for instructions.

   For a ZDB to disk, you also need to configure a backup device, as you will have to select it while configuring a backup specification. Otherwise, you cannot configure a backup specification for a ZDB to disk.

2. Configure media pools and media for a backup. See the online Help index: “creating media pools” for instructions.

3. Configure a non-ZDB backup specification and run the backup of Oracle data on the application system to verify that you have properly configured the Oracle environment. See the *HP Data Protector integration guide for Oracle and SAP* on how to create a non-ZDB backup specification.

4. Create a Data Protector Oracle ZDB backup specification. See “Creating backup specifications” on page 76.

# Creating backup specifications

## Online ZDB

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

## Offline ZDB

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

## Cluster-aware clients

Before you perform an *offline* ZDB in a cluster environment, take the Oracle Database resource offline and bring it back online after the replica is created. This can be done using the Oracle `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.

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

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

## Procedure

To create an Oracle ZDB backup specification:

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **Oracle Server**, and click **Add Backup**.

3. In the Create New Backup dialog box, select the following:

### Backup set method

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

### Proxy-copy method

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

### On EMC and XP

In the **Backup type** drop-down list, select the **Split mirror backup** option and in the **Sub type** drop-down list, select the split-mirror agent that is installed on the application and the backup systems (**EMC Symmetrix** or **HP StorageWorks XP**). See Figure 16 on page 77.

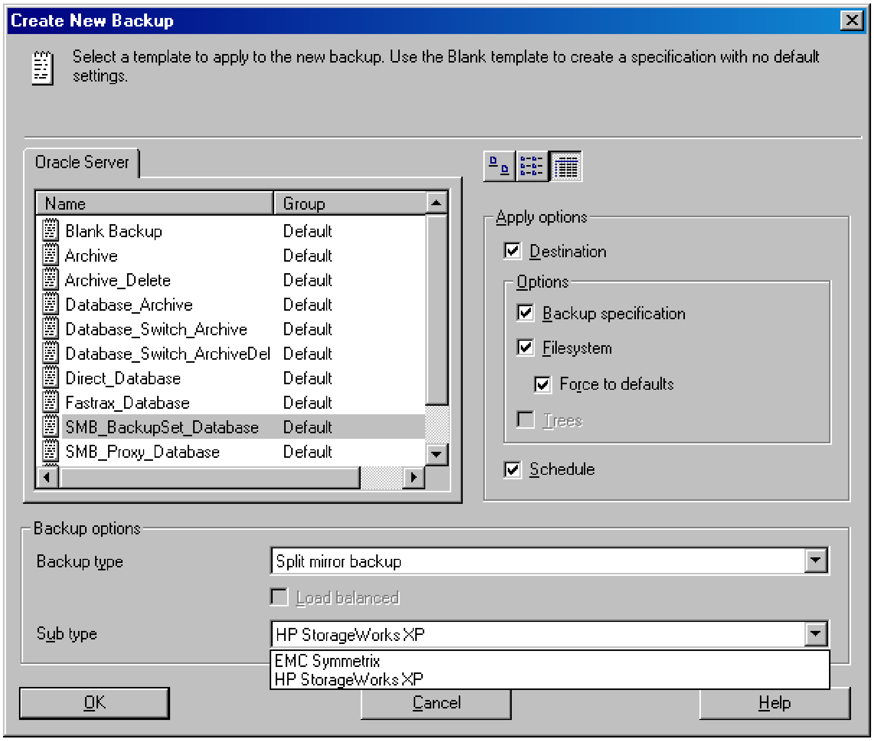

Create New Backup

Select a template to apply to the new backup. Use the Blank template to create a specification with no default settings.

Oracle Server

| Name | Group |
|---|---|
| Blank Backup | Default |
| Archive | Default |
| Archive_Delete | Default |
| Database_Archive | Default |
| Database_Switch_Archive | Default |
| Database_Switch_ArchiveDel | Default |
| Direct_Database | Default |
| Fastrax_Database | Default |
| SMB_BackupSet_Database | Default |
| SMB_Proxy_Database | Default |

Apply options

☑ Destination

Options

☑ Backup specification

☑ Filesystem

☑ Force to defaults

□ Trees

☑ Schedule

Backup options

Backup type: Split mirror backup

□ Load balanced

Sub type: HP StorageWorks XP

EMC Symmetrix
HP StorageWorks XP

**Figure 16 Selecting an Oracle ZDB template and Split mirror backup**

### On VA and EVA

In the **Backup type** drop-down list, select the **Snapshot backup** option and in the **Sub type** drop-down list, select the snapshot agent you have installed on the application and the backup system (**HP StorageWorks VA** or **HP StorageWorks EVA SMI-S**). See Figure 17 on page 78.

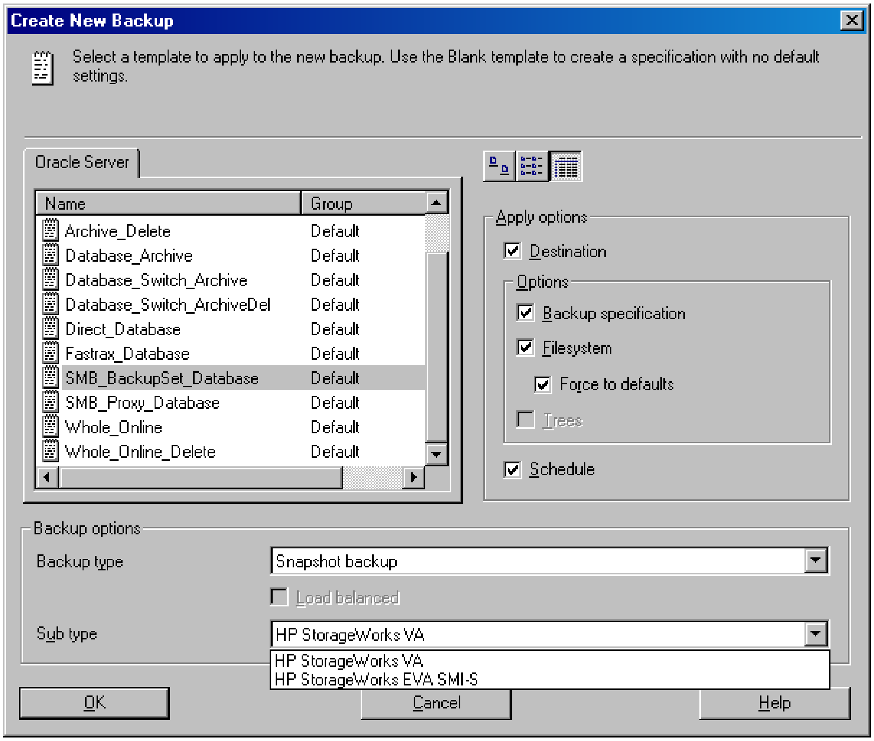

Create New Backup

Select a template to apply to the new backup. Use the Blank template to create a specification with no default settings.

Oracle Server

| Name | Group |
|---|---|
| Archive_Delete | Default |
| Database_Archive | Default |
| Database_Switch_Archive | Default |
| Database_Switch_ArchiveDel | Default |
| Direct_Database | Default |
| Fastrax_Database | Default |
| SMB_BackupSet_Database | Default |
| SMB_Proxy_Database | Default |
| Whole_Online | Default |
| Whole_Online_Delete | Default |

Apply options

☑ Destination

Options

☑ Backup specification

☑ Filesystem

☑ Force to defaults

☐ Trees

☑ Schedule

Backup options

Backup type: Snapshot backup

☐ Load balanced

Sub type: HP StorageWorks VA

**Figure 17 Selecting an Oracle ZDB template and Snapshot backup**

Click **OK**.

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

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

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

   Select other disk array specific backup options (see Figure 18 on page 79 for EMC, Figure 19 on page 81 for XP, Figure 20 on page 82 for VA, or Figure 21 on page 83 for EVA backup options). For detailed information on the options, press **F1**.

## On EMC

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

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

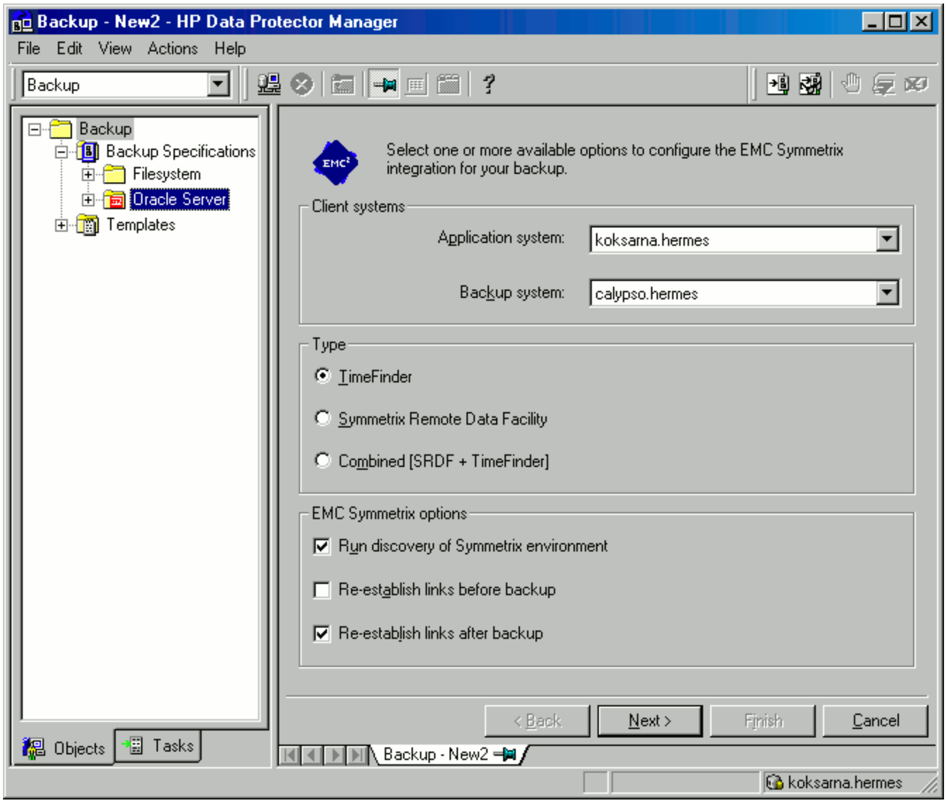

Backup - New2 - HP Data Protector Manager

Select one or more available options to configure the EMC Symmetrix integration for your backup.

Client systems

Application system: koksarna.hermes

Backup system: calypso.hermes

Type

● TimeFinder

○ Symmetrix Remote Data Facility

○ Combined [SRDF + TimeFinder]

EMC Symmetrix options

☑ Run discovery of Symmetrix environment

□ Re-establish links before backup

☑ Re-establish links after backup

Figure 18 EMC backup options

## On XP

To enable instant recovery, leave the **Track the replica for instant recovery** option selected. Note that restore using Data Protector will not be possible if this option is cleared.

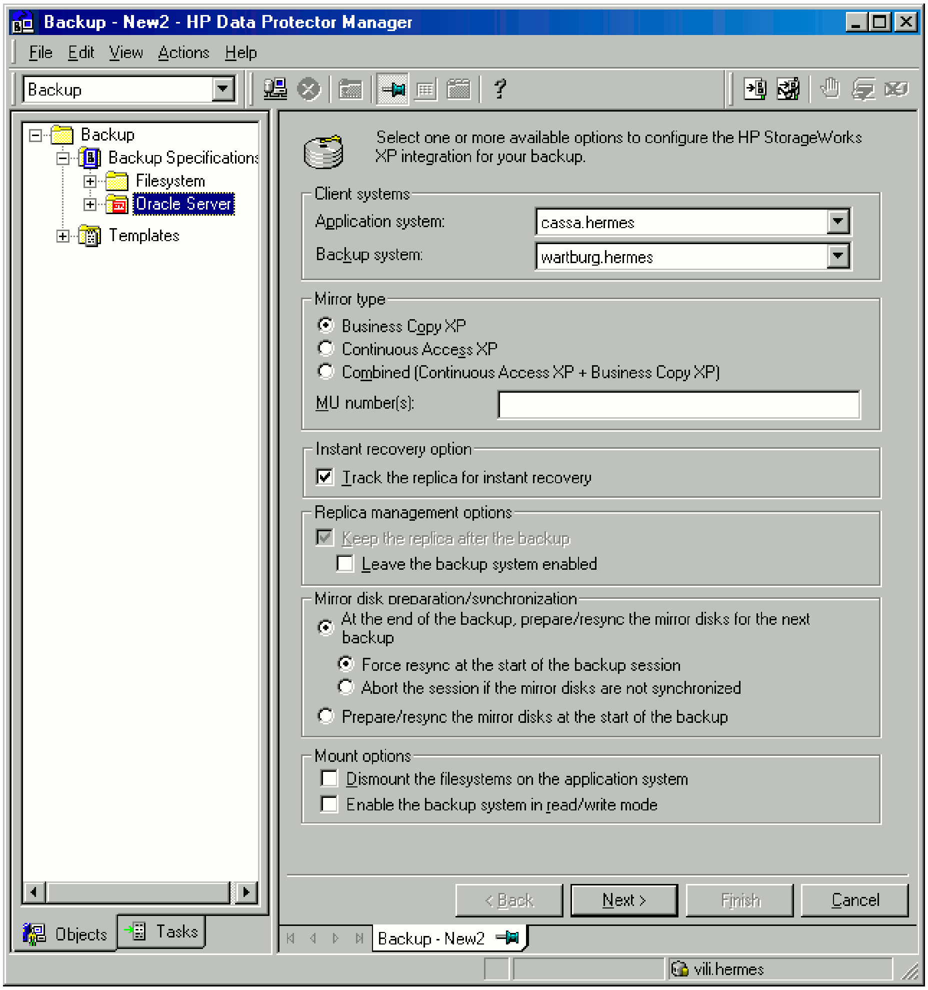

Select one or more available options to configure the HP StorageWorks XP integration for your backup.

Client systems

Application system: cassa.hermes

Backup system: wartburg.hermes

Mirror type

● Business Copy XP

○ Continuous Access XP

○ Combined (Continuous Access XP + Business Copy XP)

MU number(s):

Instant recovery option

☑ Track the replica for instant recovery

Replica management options

☑ Keep the replica after the backup

□ Leave the backup system enabled

Mirror disk preparation/synchronization

● At the end of the backup, prepare/resync the mirror disks for the next backup

● Force resync at the start of the backup session

○ Abort the session if the mirror disks are not synchronized

○ Prepare/resync the mirror disks at the start of the backup

Mount options

□ Dismount the filesystems on the application system

□ Enable the backup system in read/write mode

**Figure 19 XP backup options**

### On VA

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

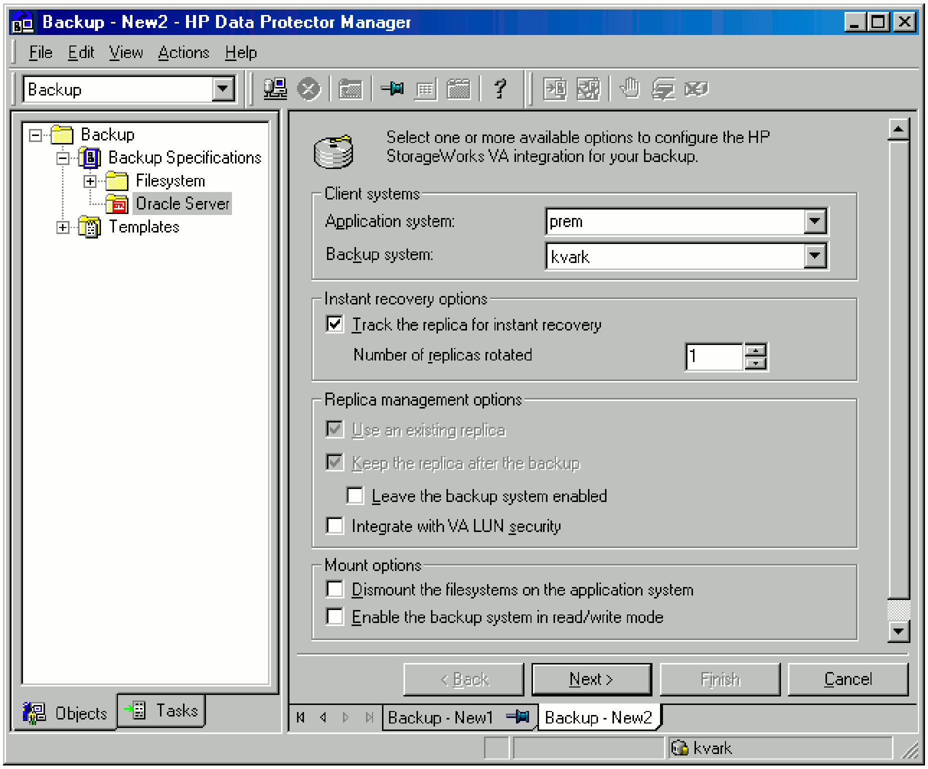

Select one or more available options to configure the HP StorageWorks VA integration for your backup.

Client systems

Application system: prem

Backup system: kvark

Instant recovery options

☑ Track the replica for instant recovery

Number of replicas rotated 1

Replica management options

☑ Use an existing replica

☑ Keep the replica after the backup

☐ Leave the backup system enabled

☐ Integrate with VA LUN security

Mount options

☐ Dismount the filesystems on the application system

☐ Enable the backup system in read/write mode

## Figure 20 VA backup options

### On EVA

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

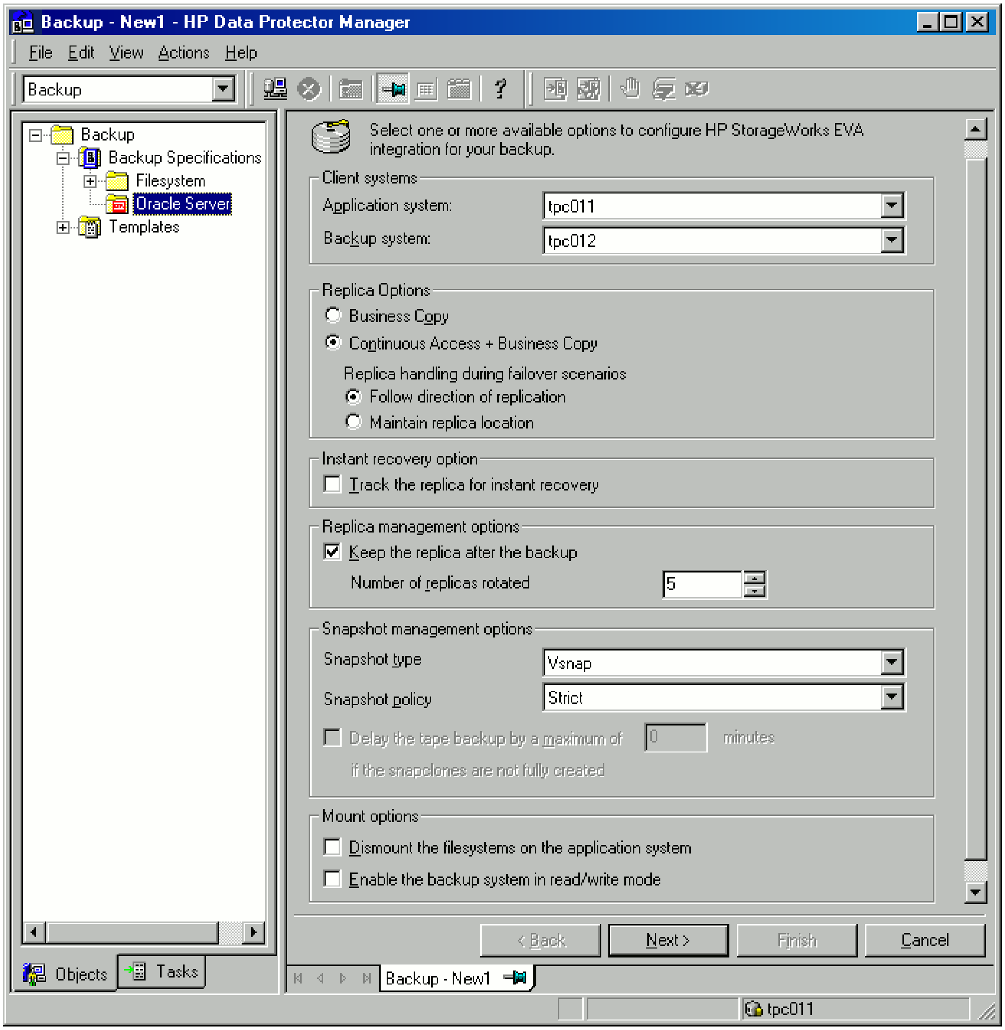

Select one or more available options to configure HP StorageWorks EVA integration for your backup.

Client systems

Application system: tpc011

Backup system: tpc012

Replica Options

○ Business Copy

● Continuous Access + Business Copy

Replica handling during failover scenarios

● Follow direction of replication

○ Maintain replica location

Instant recovery option

□ Track the replica for instant recovery

Replica management options

☑ Keep the replica after the backup

Number of replicas rotated 5

Snapshot management options

Snapshot type Vsnap

Snapshot policy Strict

□ Delay the tape backup by a maximum of 0 minutes if the snapclones are not fully created

Mount options

□ Dismount the filesystems on the application system

□ Enable the backup system in read/write mode

## Figure 21 EVA backup options

Click **Next**.

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

   The database name can be obtained using SQL*Plus:

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

NOTE:

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

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

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

- ***UNIX:*** 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 57.
- ***Windows Server 2008:*** It is not mandatory to specify these options and if they are not specified, the backup runs under the Local System Account.

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

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

NOTE:

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

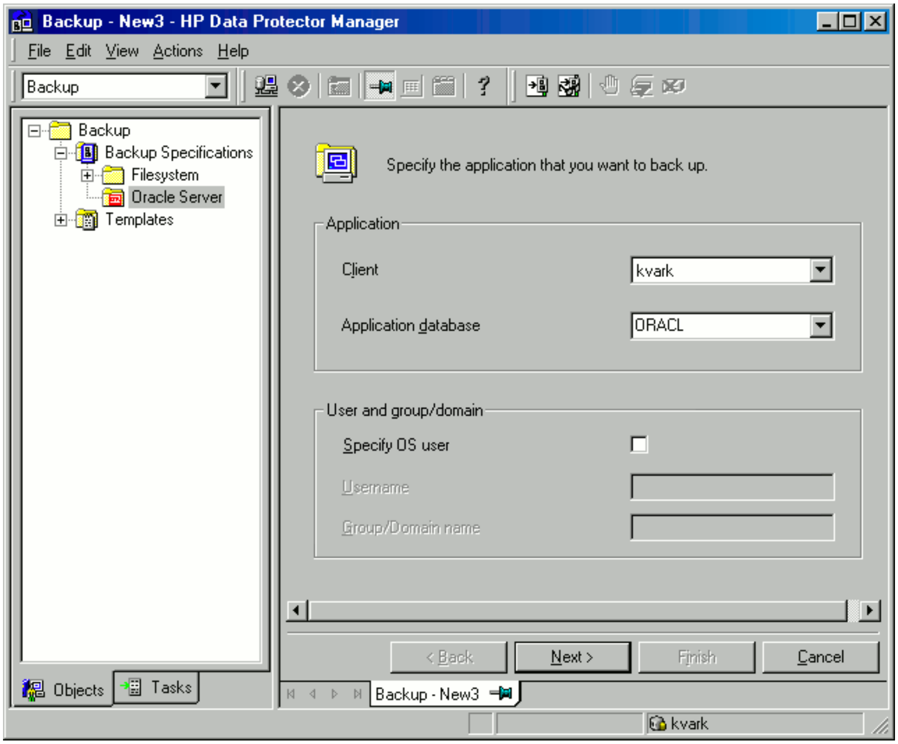

Specify the application that you want to back up.

Application

Client: kvark

Application database: ORACL

User and group/domain

Specify OS user □

Username

Group/Domain name

## Figure 22 Specifying an Oracle Server system (Windows)

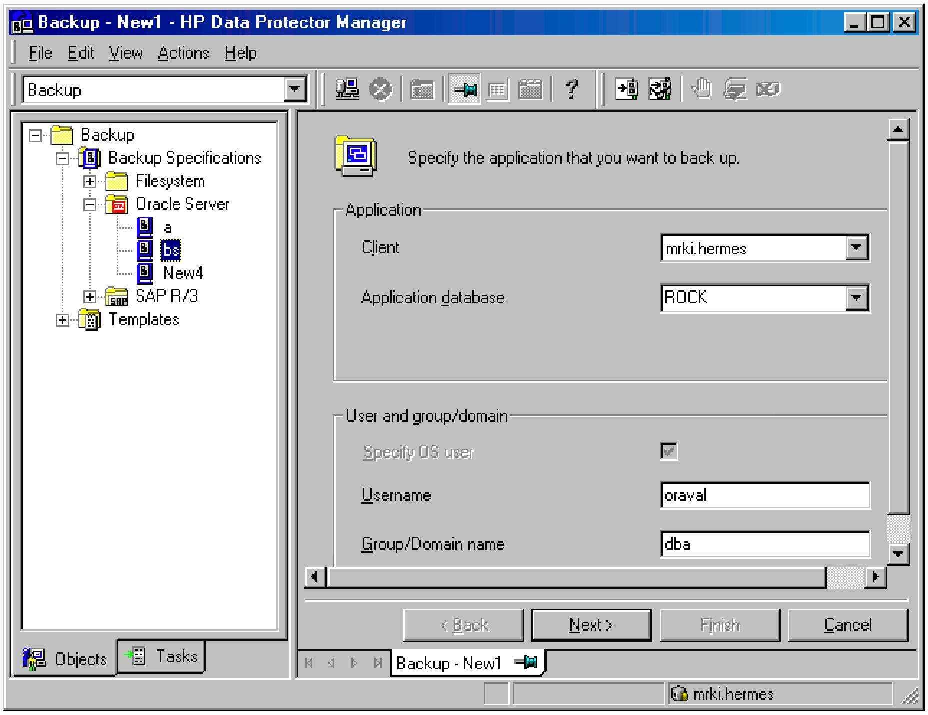

Specify the application that you want to back up.

Application

Client: mrki.hermes

Application database: ROCK

User and group/domain

Specify OS user ☑

Username: oraval

Group/Domain name: dba

Figure 23 Specifying an Oracle Server system (UNIX)

Click **Next**.

NOTE:

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

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

6. If the Oracle database is not configured yet for use with Data Protector, the Configure Oracle dialog box is displayed. Configure the Oracle database for use with Data Protector as described in "Configuring Oracle databases" on page 60.

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

NOTE:

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

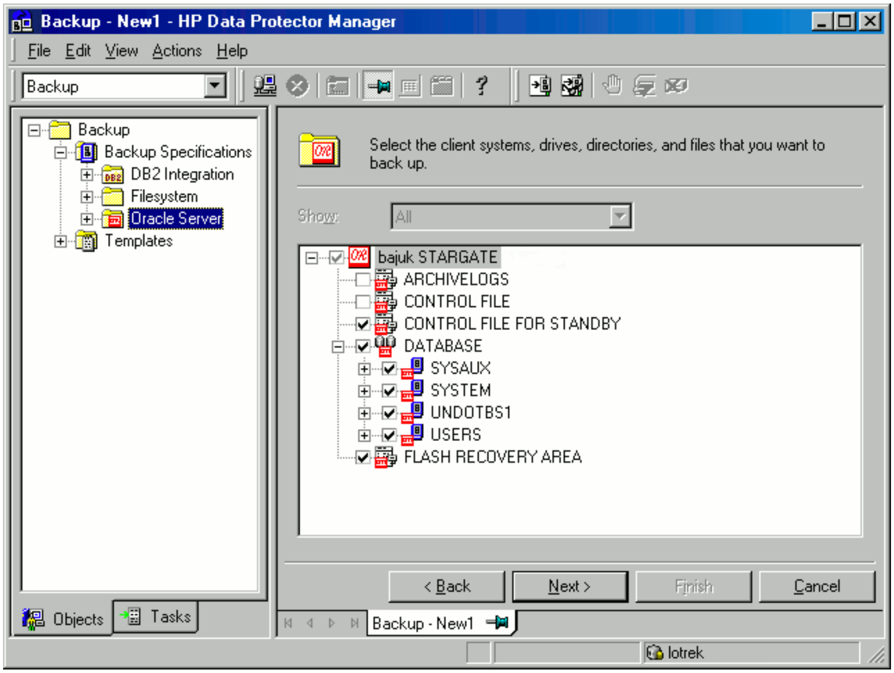

Select the client systems, drives, directories, and files that you want to back up.

Show: All

- ☑ bajuk STARGATE
  - ☐ ARCHIVELOGS
  - ☐ CONTROL FILE
  - ☑ CONTROL FILE FOR STANDBY
  - ☑ DATABASE
    - ☑ SYSAUX
    - ☑ SYSTEM
    - ☑ UNDOTBS1
    - ☑ USERS
  - ☑ FLASH RECOVERY AREA

**Figure 24 Selecting backup objects**

Click **Next**.

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

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

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

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

   NOTE:

   Object mirroring is not supported for ZDB to disk.

   Click **Next** to proceed.

9. Set the backup options.

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

## Offline ZDB

To perform an offline ZDB, select the **Backup offline** option in the Application Specific Options dialog box. This option stops the database before creating a replica, and restarts it after the replica is created. Note that if a ZDB-to-tape or ZDB-to-disk+tape session is being performed, the database is not offline during the actual backup to tape. See Figure 25 on page 89.

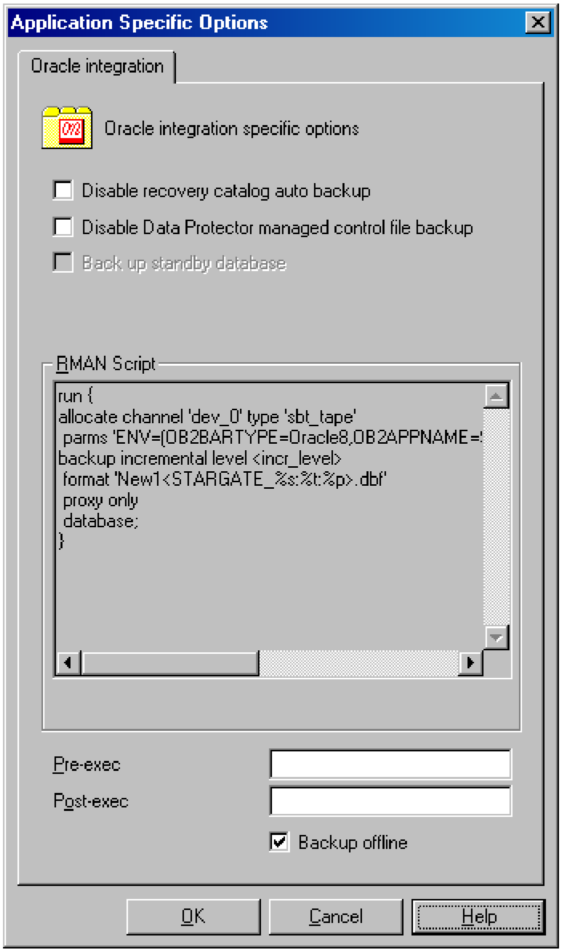
Application Specific Options
Oracle integration
Oracle integration specific options
OK
Cancel
Help

□ Disable recovery catalog auto backup

□ Disable Data Protector managed control file backup

□ Back up standby database

RMAN Script

```
run {
allocate channel 'dev_0' type 'sbt_tape'
 parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=
backup incremental level <incr_level>
 format 'New1<STARGATE_%s:%t:%p>.dbf'
 proxy only
 database;
}
```

Pre-exec

Post-exec

☑ Backup offline

**Figure 25 Backup offline option**

For information on other Application Specific Options, see Table 6 on page 91 or press **F1**.

Click **Next**.

10. Optionally, schedule the backup. For more details, see "Scheduling backup specifications" on page 98.

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

    Click **Next**.

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

IMPORTANT:

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

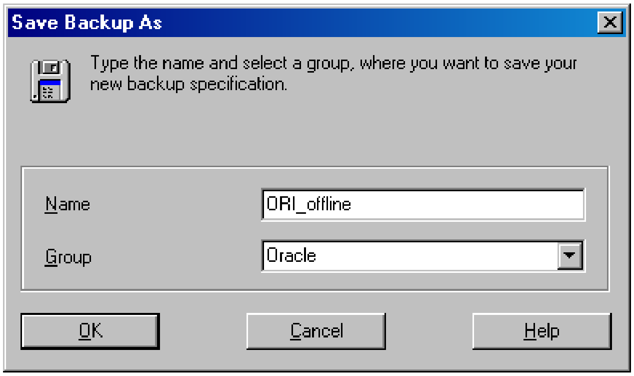

Save Backup As

Type the name and select a group, where you want to save your new backup specification.

Name: ORI_offline

Group: Oracle

**Figure 26 Saving the backup specification**

Click **OK**.

To start the backup, see "Starting backup sessions" on page 97.

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

TIP:

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

### Table 6 Oracle backup options

| | |
|---|---|
| **Disable recovery catalog auto backup** | By default, Data Protector backs up the recovery catalog after every ZDB to tape or ZDB to disk+tape. Select this option to disable backup of the recovery catalog. |
| **Disable Data Protector managed control file backup** | By default, Data Protector backs up the Data Protector managed control file after every ZDB to tape or ZDB to disk+tape. Select this option to disable backup of the Data Protector managed control file. |
| **Back up standby database** | This option is ignored for ZDB. |
| **RMAN Script** | You can edit the Oracle RMAN script section of the Data Protector Oracle backup specification. The script is created by Data Protector during the creation of a backup specification and reflects the backup specification’s selections and settings. You can edit the script only after the backup specification has been saved. For information on how to edit the RMAN script section, see “Editing the Oracle RMAN script” on page 93. |

| | |
|---|---|
| **Pre-exec, Post-exec** | Specify a command or RMAN script that will be started by `ob2rman.pl` on the Oracle Server system before the backup (`pre-exec`) or after it (`post-exec`). RMAN scripts must have the `.rman` extension. Do not use double quotes.<br>For example, you can provide scripts to shut down and start an Oracle instance. For examples of shut-downing and starting an Oracle instance on UNIX, see "Examples of pre-exec and post-exec scripts on UNIX" on page 92.<br>Provide the pathname of the command or RMAN script. |
| **Backup offline** | Select this option to perform an offline ZDB session. This option stops the database before creating a replica, and restarts it after the replica is created. See Figure 25 on page 89. |

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

### Pre-exec example

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

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

### Post-exec example

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

```
#!/bin/sh
export ORACLE_HOME=$2
export ORACLE_SQLNET_NAME=$1
if [ -f $ORACLE_HOME/bin/sqlplus ]; then
```

```
$ORACLE_HOME/bin/sqlplus << EOF
connect sys/manager@$ORACLE_SQLNET_NAME as sysdba
startup
EOF
echo "Oracle database \"$DB_NAME\" started."
exit 0
else
echo "Cannot find Oracle SQLPLUS ($ORACLE_HOME/bin/sqlplus)."
exit 1
fi
```

## Editing the Oracle RMAN script

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

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

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

### Limitations

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

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

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

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

To edit an Oracle RMAN script, click **Edit** in the **Application Specific Options** window (see Figure 31 on page 107), edit the script, and then click **Save** to save the changes to the script.

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

## Data Protector RMAN script structure

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

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

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

NOTE:

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

IMPORTANT:

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

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

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

### Proxy-copy

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

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

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

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

NOTE:

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

On HP StorageWorks XP, for an Oracle direct backup, add `OB2DMP=1`.

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

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

    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.

  - In case of an Oracle proxy-copy ZDB-to-disk+tape or ZDB-to-tape session, the `PROXY ONLY` option is required. Only one `BACKUP` command with the `proxy only` option is permitted and only one additional backup command for backing up the control file is permitted.
  - The RMAN `datafile` *`tablespace_name*datafile_name`* command.

- If the Archived Redo Logs were selected for a backup, **an RMAN backup statement for the backup of Oracle archive logs**.

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

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

  NOTE:

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

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

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

    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.

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

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

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

```
run {
allocate channel 'dev_0' type 'sbt_tape' parms
'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=DIPSI,OB2BARLIST=New1,
OB2SMB=1,OB2DMAP=1)';
```

```
backup incremental level <incr_level>
format 'New1<DIPSI_%s:%t:%p>.dbf'
proxy only
database
;
backup format 'New1<DIPSI_%s:%t:%p>.dbf' controlfile;
}
```

# Starting backup sessions

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

## Backup methods

- Schedule a backup of an existing Oracle ZDB backup specification using the Data Protector Scheduler. See "Scheduling backup specifications" on page 98.
- Start an interactive backup of an existing Oracle ZDB backup specification using the Data Protector GUI or the Data Protector CLI. See "Running an interactive backup" on page 100.

# Considerations

Before running an Oracle ZDB session, consider the following:

- It is not possible to start a ZDB, restore, or instant recovery sessions using the same source volume on the application system at the same time. A ZDB, restore, or instant recovery session must be started only after the preceding session that is using the same source volume on the application system has finished the ZDB session or restore; otherwise, the session will fail.
- For the backup set method, if the Oracle database is installed on symbolic links, then these symbolic links have to be also created on the backup system.
- On XP, if the LVM mirroring configuration is used, Data Protector displays a warning during a backup because the volume group source volumes on the application system do not have their BC pairs assigned. This message should be ignored.
- If the control file, SPILE, or online redo logs are on the same source volumes as the datafiles and the **Track the replica for instant recovery** option is selected, the backup session will be aborted. In this case, you need to either reconfigure the database or set the `ZDB_ORA_INCLUDE_CF_OLF`, `ZDB_ORA_INCLUDE_SPF`, and `ZDB_ORA_NO_CHECKCONF_IR` `omnirc` variables. See "Reconfiguring an Oracle instance for instant recovery" on page 421 or "ZDB integrations omnirc variables" on page 425.

- With the Oracle backup set ZDB method, you must manually re-synchronize the recovery catalog database with the current control file after you modified the physical schema of a database (for example, if you add or drop a tablespace, add a new datafile, or add or drop a rollback segment) and before you start the next backup.

## Scheduling backup specifications

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

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

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

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

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

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

   In the case of a ZDB-to-disk or a ZDB-to-disk+tape session, specify the **Split mirror/snapshot backup** option. See Figure 27 on page 99.

NOTE:

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

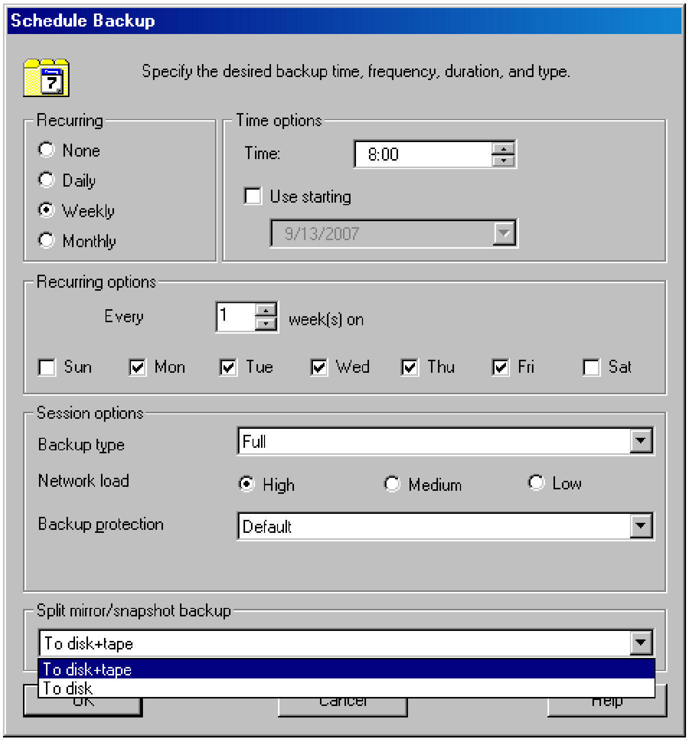

Schedule Backup

Specify the desired backup time, frequency, duration, and type.

Recurring

○ None

○ Daily

● Weekly

○ Monthly

Time options

Time: 8:00

□ Use starting

9/13/2007

Recurring options

Every 1 week(s) on

□ Sun ☑ Mon ☑ Tue ☑ Wed ☑ Thu ☑ Fri □ Sat

Session options

Backup type: Full

Network load: ● High ○ Medium ○ Low

Backup protection: Default

Split mirror/snapshot backup

To disk+tape

To disk+tape

To disk

**Figure 27 Selecting ZDB to disk or ZDB to disk+tape session using the Data Protector scheduler**

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

## Running an interactive backup

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

### Starting a backup using the GUI

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

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

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

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

   In the case of a ZDB-to-disk or a ZDB-to-disk+tape session, specify the **Split mirror/snapshot backup** option. See Figure 28 on page 101.

NOTE:

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

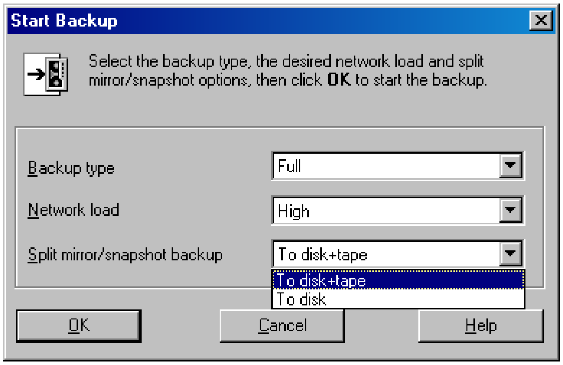

Start Backup

Select the backup type, the desired network load and split mirror/snapshot options, then click **OK** to start the backup.

Backup type: Full

Network load: High

Split mirror/snapshot backup: To disk+tape

**Figure 28 Selecting ZDB to disk or ZDB to disk+tape session when starting an interactive backup**

Click **OK**.

## Starting a backup using the CLI

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

```
omnib -oracle8_list Name
```

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

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

where *`Name`* is the name of the backup specification. For more information on the `omnib` command, see its man page or the *HP Data Protector command line interface reference*.

NOTE:

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

# Restore

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

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

Using the Data Protector GUI, you can also **duplicate** a production database. See "Duplicating an Oracle database" on page 115.

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

- Standard restore from backup media to the application system on LAN. See "Restoring from backup media to the application system on LAN" on page 104.
- Instant recovery. See "Instant recovery and database recovery" on page 138.

See also Table 3 on page 34 for an overview of recovery methods depending on the backup type and type of recovery.

## Microsoft Cluster Server clients

Before you start restoring a cluster-aware Oracle server, take the Oracle Database resource offline using, for example, the **Cluster Administrator** utility. See Figure 29 on page 103.

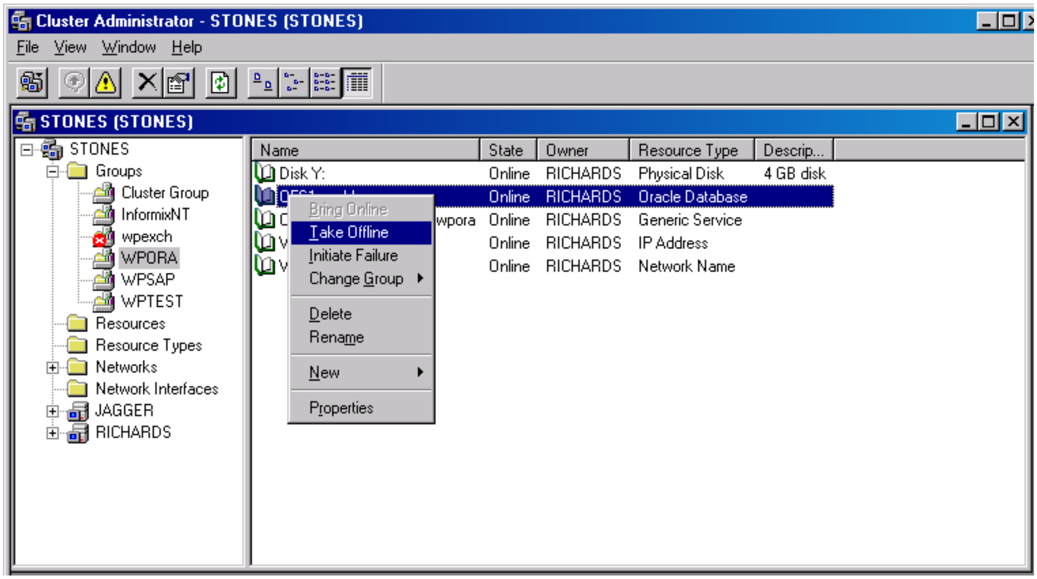

- STONES
  - Groups
    - Cluster Group
    - InformixNT
    - wpexch
    - WPORA
    - WPSAP
    - WPTEST
  - Resources
  - Resource Types
  - Networks
  - Network Interfaces
  - JAGGER
  - RICHARDS

| Name | State | Owner | Resource Type | Descrip... |
|---|---|---|---|---|
| Disk Y: | Online | RICHARDS | Physical Disk | 4 GB disk |
| O | Online | RICHARDS | Oracle Database | |
| C wpora | Online | RICHARDS | Generic Service | |
| V | Online | RICHARDS | IP Address | |
| V | Online | RICHARDS | Network Name | |

### Figure 29 Taking the Oracle resource group offline

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

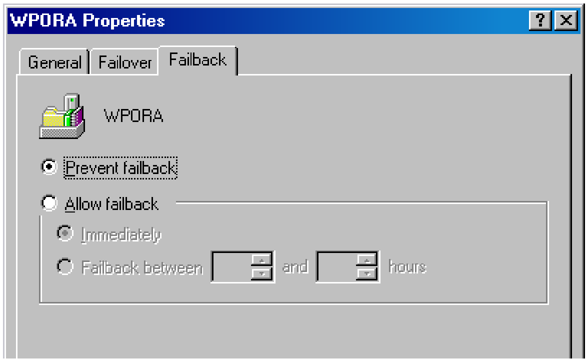

WPORA

● Prevent failback

○ Allow failback

● Immediately

○ Failback between and hours

### Figure 30 Checking properties

## MC/ServiceGuard clients

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

```
run {
allocate channel dev1 type 'sbt_tape' parms
'SBT_LIBRARY=Path_to_Data_Protector_MML,
ENV=(OB2BARHOSTNAME=virtual.domain.com)';
restore datafile '/opt/ora9i/oradata/MAKI/example02.dbf';
release channel dev1;
}
```

## Prerequisites

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

# Restoring from backup media to the application system on LAN

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

- Data Protector GUI. See "Restoring Oracle using the Data Protector GUI" on page 104.
- RMAN. See "Restoring Oracle using RMAN" on page 123.

# 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 161.

## Restoring database items in a disaster recovery

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

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

## Changing the database state

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

**Table 7 Required database states**

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

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

To put the database into the correct state, run:

```
sqlplus /nolog
```

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

```
SQL>shutdown immediate;
```

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

```
SQL>startup nomount;
```

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

```
SQL>startup mount;
```

## Restoring the recovery catalog database

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

To restore the recovery catalog database:

1. Ensure that the recovery catalog database is in the **Open** state.
2. Remove the recovery catalog from the database (if it exists), using the RMAN command `DROP CATALOG`.
3. In the Data Protector GUI, switch to the **Restore** context.

4. Under **Restore Objects**, expand **Oracle Server**, expand the client 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.

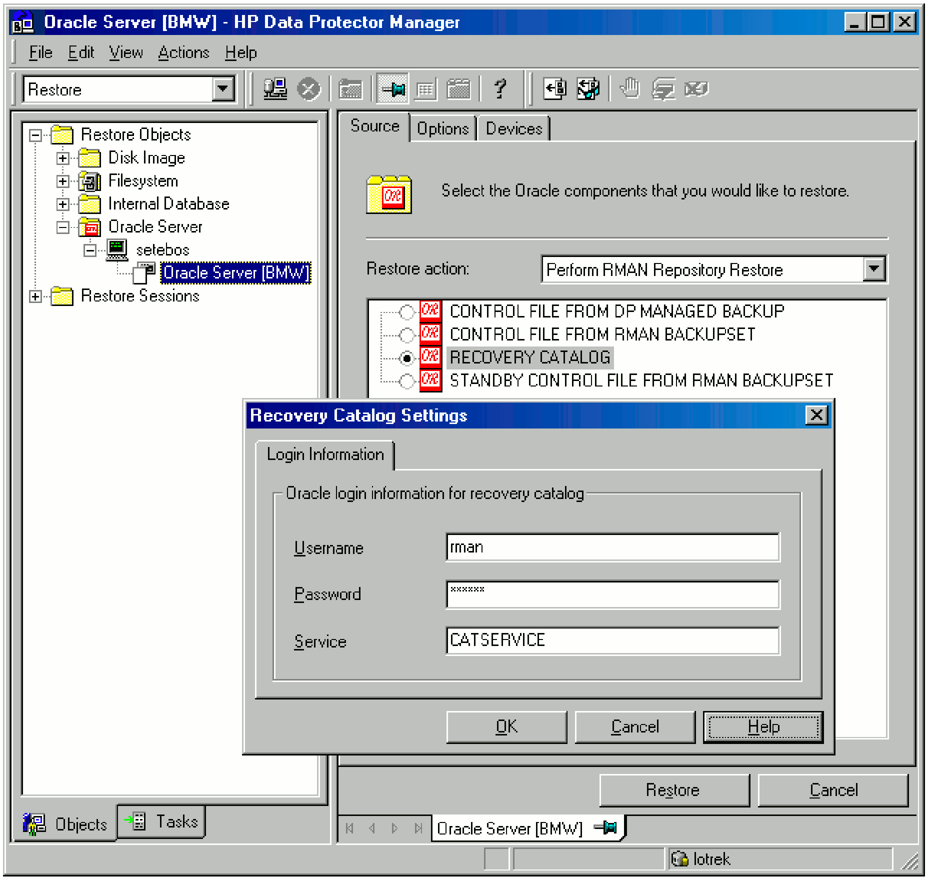

Select the Oracle components that you would like to restore.

Restore action: Perform RMAN Repository Restore

- ○ CONTROL FILE FROM DP MANAGED BACKUP
- ○ CONTROL FILE FROM RMAN BACKUPSET
- ● RECOVERY CATALOG
- ○ STANDBY CONTROL FILE FROM RMAN BACKUPSET

Recovery Catalog Settings

Login Information

Oracle login information for recovery catalog

Username: rman

Password: ******

Service: CATSERVICE

## Figure 31 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 119.

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 (`ctrl`*`DB_NAME`*`.dbf`) are restored to:

  ***Windows:*** *`Data_Protector_home`*`\tmp`

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

  After the restore, run the following script:

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

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

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

  The recovery catalog *is* required.

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

To restore the control file:

1. Open the `sqlplus` window and put the database in the nomount state. See "Changing the database state" on page 105.
2. In the Data Protector GUI, switch to the **Restore** context.
3. Under **Restore Objects**, expand **Oracle Server**, expand the client on which the database, for which you want to restore the control file, resides, and then click the database.
4. In the **Restore Action** drop-down list, select **Perform RMAN Repository Restore**.

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

   Set the other restore options. See "Restore, recovery, and duplicate options" on page 119 for information.
6. Click **Restore**.

Proceed with restoring the Oracle database objects.

## Restoring Oracle database objects

Before you restore Oracle database objects, ensure that you have an up-to-date version of the recovery catalog database and the control file. They contain the database structure information. If you do not have up-to-date versions of these files, restore them as described in "Restoring the recovery catalog database" on page 105 and "Restoring the control file" on page 108.

To restore Oracle database objects:

1. Put the database in the mount state. See "Changing the database state" on page 105.
2. In the Data Protector GUI, switch to the **Restore** context.
3. Under **Restore Objects**, expand **Oracle Server**, expand the client on which the database, for which you restore the database objects, resides, and then click the database.

4. In the **Restore action** drop-down list, select the type of restore you wish to perform. For information on the options, see “Restore, recovery, and duplicate options” on page 119.

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 123.

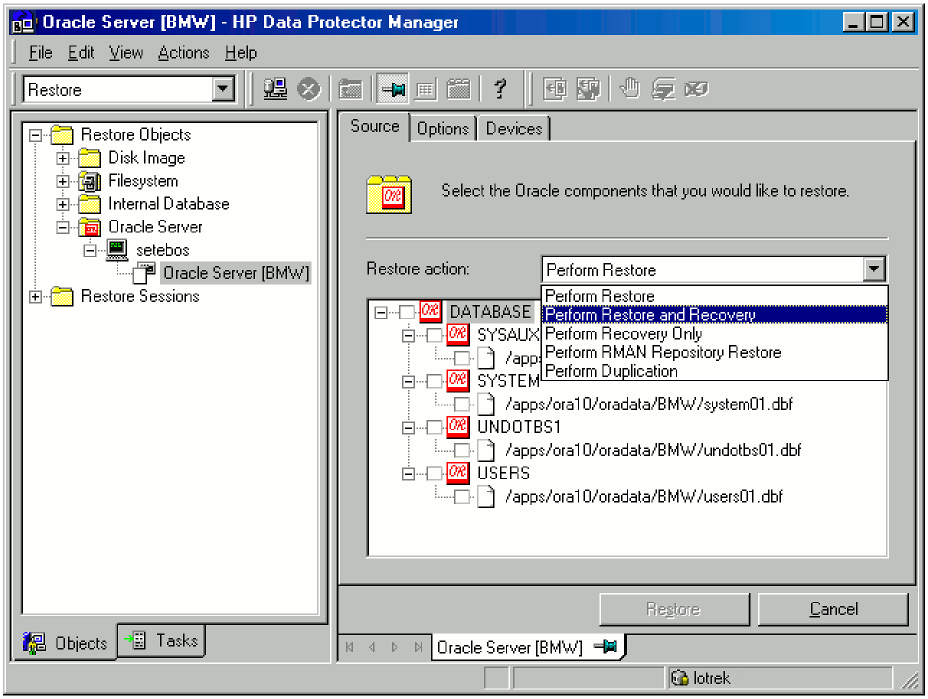

Select the Oracle components that you would like to restore.

Restore action: Perform Restore

/apps/ora10/oradata/BMW/system01.dbf

/apps/ora10/oradata/BMW/undotbs01.dbf

/apps/ora10/oradata/BMW/users01.dbf

Figure 32 Source page

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

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

   NOTE:

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

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

   Set the other restore options. See “Restore, recovery, and duplicate options” on page 119 for information.

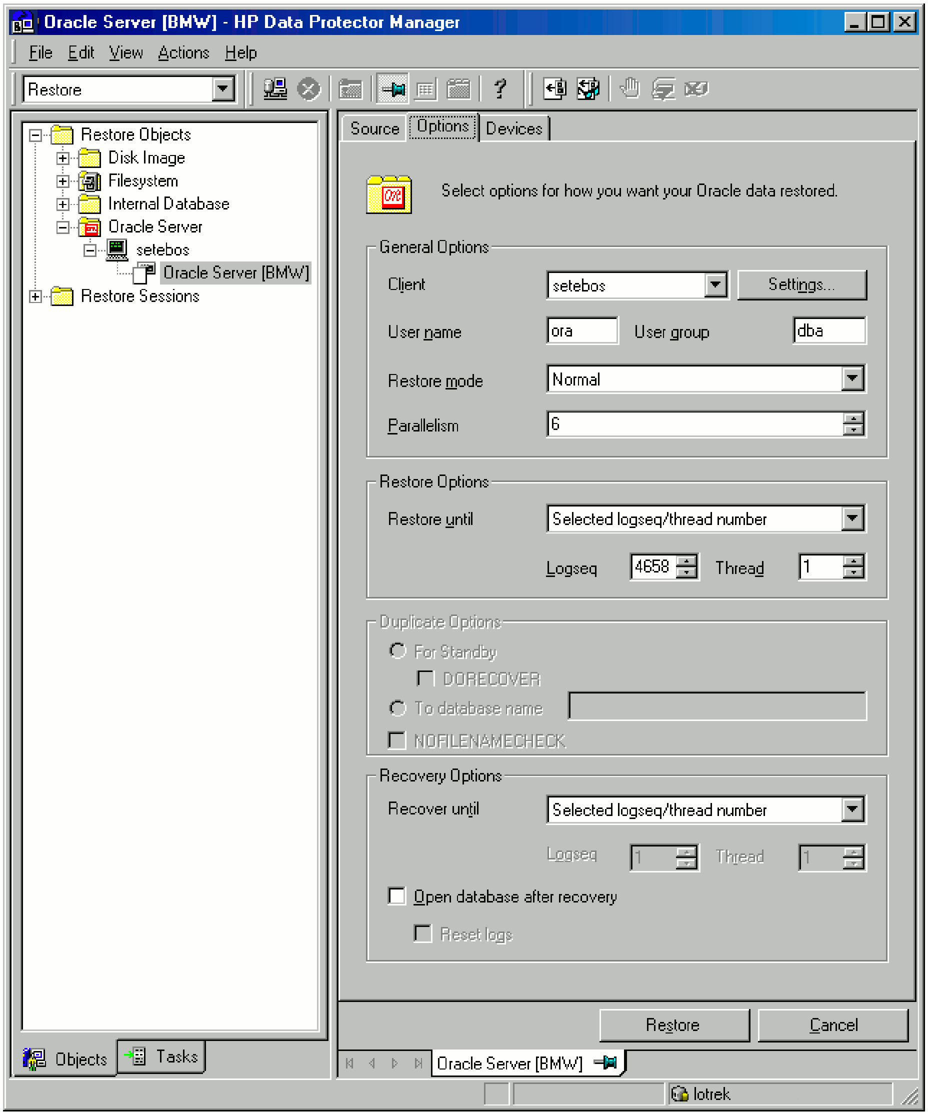

Select options for how you want your Oracle data restored.

General Options

Client: setebos

User name: ora

User group: dba

Restore mode: Normal

Parallelism: 6

Restore Options

Restore until: Selected logseq/thread number

Logseq: 4658 Thread: 1

Duplicate Options

○ For Standby

□ DORECOVER

○ To database name

□ NOFILENAMECHECK

Recovery Options

Recover until: Selected logseq/thread number

Logseq: 1 Thread: 1

□ Open database after recovery

□ Reset logs

## Figure 33 Options page

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

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

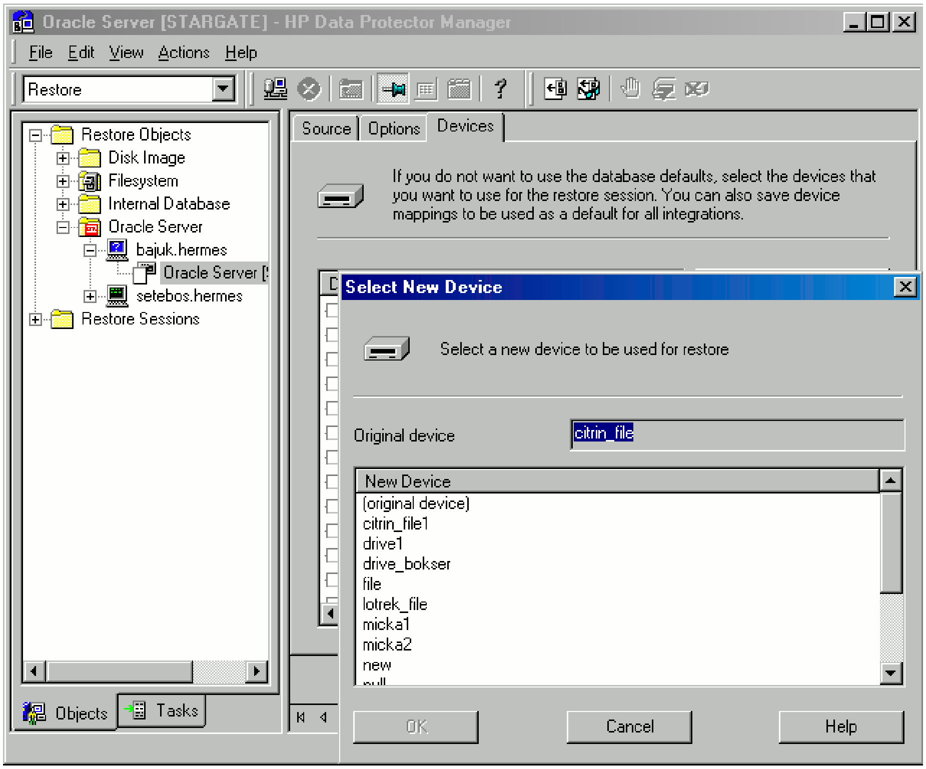

If you do not want to use the database defaults, select the devices that you want to use for the restore session. You can also save device mappings to be used as a default for all integrations.

Select a new device to be used for restore

Original device: citrin_file

New Device
(original device)
citrin_file1
drive1
drive_bokser
file
lotrek_file
micka1
micka2
new
null

**Figure 34 Devices page**

8. Click **Restore**.

After the restore:

1. Put the database in the correct state.

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

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

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

   ```
   rman target Target_Database_Login catalog
   Recovery_Catalog_Login
   ```

   ```
   RMAN> RESET DATABASE;
   ```

   ```
   RMAN> exit
   ```

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

   Open a command line window and enter the following commands:

   ```
   sqlplus /nolog
   ```

   ```
   SQL>recover database;
   ```

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

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

## Restoring tablespaces and datafiles

To restore tablespaces and datafiles:

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

   ```
   sqlplus /nolog
   ```

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

   ```
   SQL>alter database datafile 'datafile name' offline;
   ```

   If you are restoring a tablespace enter:

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

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

   Open a command line window and enter the following commands:

   ```
   sqlplus /nolog
   ```

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

   If you are restoring a datafile enter:

   ```
   SQL>alter database datafile 'datafile_name' online;
   ```

   If you are restoring a tablespace enter:

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

## Duplicating an Oracle database

Perform a production database duplication to create:

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

### Prerequisites

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

LOG_FILE_NAME_CONVERT initialization parameters appropriately. Thus, the target database files will not be overwritten by the duplicate database files.

### Limitations

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

To duplicate a production database:

1. On the client where the selected database will be duplicated, put the Oracle auxiliary database instance in the nomount state. See "Changing the database state" on page 105.
2. In the Context List of the Data Protector GUI, click **Restore**.
3. Under **Restore Objects**, expand **Oracle Server**, expand the client on which the production database resides, and then click the production database which you want to duplicate. If there are several such clients, select the client on which you want the Data Protector Oracle integration agent (ob2rman.pl) to be started.
4. In the **Restore Action** drop-down list, select **Perform Duplication**.

5. In the **Options** page, from the **Client** drop-down list, select the client on which the Data Protector Oracle integration agent (`ob2rman.pl`) will be started.

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

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

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

   Set duplicate options. For information, see "Duplicate options" on page 120 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.

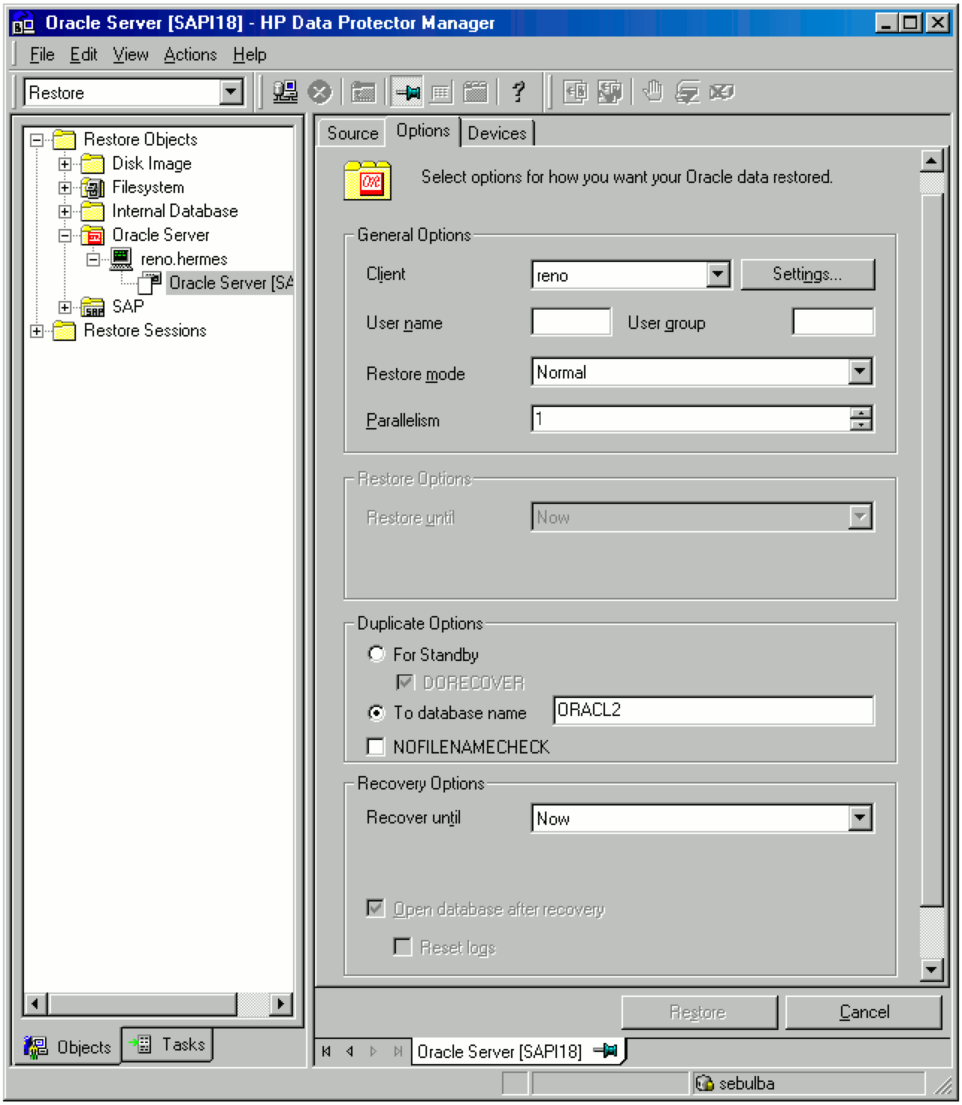

Select options for how you want your Oracle data restored.

General Options

Client: reno

User name:

User group:

Restore mode: Normal

Parallelism: 1

Restore Options

Restore until: Now

Duplicate Options

○ For Standby

☑ DORECOVER

● To database name: ORACL2

□ NOFILENAMECHECK

Recovery Options

Recover until: Now

☑ Open database after recovery

□ Reset logs

Figure 35 Oracle duplicate options

6. Click **Restore**.

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

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

# Restore, recovery, and duplicate options

## Restore action options

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

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

**Perform Restore**

Use this option to only restore (but not recover) the database objects using Data Protector. After restore, recover the database manually using RMAN. For information on recovering the database using RMAN, see "Restoring Oracle using RMAN" on page 123.

**Perform Restore and Recovery**

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

**Perform Recovery Only**

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

**Perform RMAN Repository Restore**

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

**Perform Duplication**

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

## General options

**Client**

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

**Settings**

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

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

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

**User name, User group** (UNIX systems only)

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

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

**Restore mode**

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

- Normal

  This option should be used when a conventional backup or ZDB using the backup set method was performed.

- Proxy copy

  This option should be used when the original Oracle backup was made using the Oracle RMAN proxy-copy method, such as ZDB of Oracle 9i.

This option is disabled when you perform recovery only.

**Parallelism**

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

In case of `Normal` restore mode, to optimize restore performance, specify the same number of data streams as were used during the backup. For example, if you set the backup concurrency to 3, set the number of parallel data streams to 3 as well. Note that if a very high number of parallel data streams is specified this may result in a resource problem because too much memory is being used.

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

## Duplicate options

Available if **Perform Duplication** was selected.

**For Standby**

Select this option to create a standby database.

Default: selected.

**DORECOVER**

Available if **For Standby** was selected.

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

**To database name**

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

**NOFILENAMECHECK**

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

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

Default: not selected.

## Restore and recovery options

**Restore until**

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

- **Now**

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

- **Selected time**

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

- **Selected logseq/thread number**

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

- **Selected SCN number**

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

**Recover until**

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

- **Now**

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

- **Selected time**

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

- **Selected logseq/thread numbe**r

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

- **Selected SCN number**

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

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

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

**Open database after recovery**

Opens the database after a recovery is performed.

**Reset logs**

Resets the archive logs after the database is opened.

*Always* reset the logs:

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

*Do not* reset the logs:

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

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

NOTE:

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

# Restoring Oracle using RMAN

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

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

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

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

The following examples of restore are given:

- "Example of full database restore and recovery" on page 126
- "Example of point-in-time restore" on page 127
- "Example of tablespace restore and recovery" on page 129
- "Example of datafile restore and recovery" on page 131
- "Example of archive log restore" on page 135

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

## Preparing the Oracle database for restore

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

## Prerequisites

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

- Make sure that the recovery catalog database is open. If the recovery catalog database cannot be brought online, you will probably need to restore the recovery catalog database. See "Restore" on page 102 for details on how to restore the recovery catalog database.
- Check which ZDB method (proxy-copy or backup set) was used for the backup session that you plan to restore.
- Control files must be available. If the control files are not available, you must restore them. See the *Oracle Recovery Manager User's Guide and References* for more details.

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

  When you are sure that the recovery catalog database files are in place, start the recovery catalog database.
- Make sure that the following environment variables are set:
  - `ORACLE_BASE`
  - `ORACLE_HOME`
  - `ORACLE_TERM`
  - `DB NAME`
  - `PATH`
  - `NLS_LANG`
  - `NLS_DATE_FORMAT`

## Windows example

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

### UNIX example

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

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

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

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

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

## Connection strings used in the examples

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

- Target connection string for target database:

  `sys/manager@PROD`

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

- Recovery catalog connection string for recovery catalog database:

  `rman/rman@CATAL`

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

## SBT_LIBRARY parameter

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

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

## Example of full database restore and recovery

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

1. Log in to the Oracle RMAN:
   - On Windows: `ORACLE_HOME\bin\rman target sys/manager@PROD catalog rman/rman@CATAL`
   - On UNIX: `ORACLE_HOME/bin/rman target sys/manager@PROD catalog rman/rman@CATAL`
2. Start the full database restore and recovery:

   For a non-ZDB or ZDB backup set session:

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

   For a ZDB proxy-copy session:

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

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

1. Create a file `restore_database` in the `/var/opt/omni/tmp` (UNIX systems) or `Data_Protector_home\tmp` directory.

2. Start the full database restore:

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

## 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:

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

2. Start the point-in-time restore:

   For a non-ZDB or ZDB backup set session:

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

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

   ```
   run{
   allocate channel 'dev1' type 'sbt_tape' parms
   'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
   ENV=(OB2BARTYPE=Oracle8,OB2PROXYCOPY=1, OB2APPNAME=DB_NAME)';
   allocate channel 'dev2' type 'sbt_tape' parms
   'SBT_LIBRARY=/opt/omni/lib/libob2oracle8.so,
   ENV=(OB2BARTYPE=Oracle8, OB2APPNAME=DB_NAME)';
   set until time 'Mar 14 2006 11:40:00';
   restore database;
   release channel 'dev1';
   recover database;
   sql 'alter database open';
   release channel 'dev2';
   }
   ```
3. After you have performed a point-in-time restore, reset the database in the Recovery Catalog.

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

1. Create a file `restore_PIT` in the `/var/opt/omni/tmp` or *`Data_Protector_home`*`\tmp` directory.

2. Start the point-in-time restore:
   - On Windows: `ORACLE_HOME\bin\rman target sys/manager@PROD catalog rman/rman@CATAL cmdfile=Data_Protector_home\tmp\restore_PIT`
   - On UNIX: `ORACLE_HOME/bin/rman target sys/manager@PROD catalog rman/rman@CATAL cmdfile=/var/opt/omni/tmp/restore_PIT`

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

2. Start the tablespace restore and recovery.

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

     For a non-ZDB or ZDB backup set session:

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

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

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

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

     For a non-ZDB or ZDB backup set session:

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

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

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

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

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

## 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.
   - On Windows: *`ORACLE_HOME`*`\bin\rman target sys/manager@PROD catalog rman/rman@CATAL`
   - On UNIX: *`ORACLE_HOME`*`/bin/rman target sys/manager@PROD catalog rman/rman@CATAL`

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

For a non-ZDB or ZDB backup set session:

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

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

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

Windows

For a non-ZDB or ZDB backup set session:

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

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

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

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

UNIX

For a non-ZDB or ZDB backup set session:

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

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

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

Windows

For a non-ZDB or ZDB backup set session:

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

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

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

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

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

## Example of archive log restore

To restore an archive log:

1. Log in to the Oracle RMAN:
   - On Windows: `ORACLE_HOME\bin\rman target sys/manager@PROD catalog rman/rman@CATAL`
   - On UNIX: `ORACLE_HOME/bin/rman target sys/manager@PROD catalog rman/rman@CATAL`
2. Start the archive log restore:

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

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

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

## 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:
   - On Windows: `ORACLE_HOME\bin\rman target sys/manager@TIN`
   - On UNIX: `ORACLE_HOME/bin/rman target sys/manager@TIN`

2. Run:

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

NOTE:

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

## Restoring using another device

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

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

`"DEV 1" "DEV 2"`

where

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

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

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

### Example

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

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

## Instant recovery and database recovery

For general information on instant recovery, see the *HP Data Protector zero downtime backup concepts guide* and the *HP Data Protector zero downtime backup administrator's guide*. For information on instant recovery in cluster environment (Cluster File System (CFS), MC/ServiceGuard, and Microsoft Cluster Server), see *HP Data Protector zero downtime backup administrator's guide*.

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

IMPORTANT:

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

### Prerequisites

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

### Limitations

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

### RAC preparation steps

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

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

The instant recovery procedure is the same as without RAC.

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

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

   ```
   export OB2BARHOSTNAME=virtual_hostname
   ```

## Instant recovery using the Data Protector GUI

To perform an instant recovery:

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

   For example:

   ```
   sqlplus
   ```

   ```
   sql> shutdown immediate
   ```

   ```
   sql> exit
   ```

2. In the Context List, click **Instant Recovery**.
3. Expand **Oracle Server** and select the ZDB-to-disk or ZDB-to-disk+tape session from which you want to perform the restore.

4. In the **Source** tab, select the objects to recover. Only whole databases can be selected. For StorageWorks Virtual Array and HP StorageWorks Disk Array XP, it is recommended to leave the **Keep the replica after the restore** option set to enable a restart of an instant recovery session. For HP StorageWorks Enterprise Virtual Array, replica is kept on the array only if the **Copy replica data to the source location** is selected.

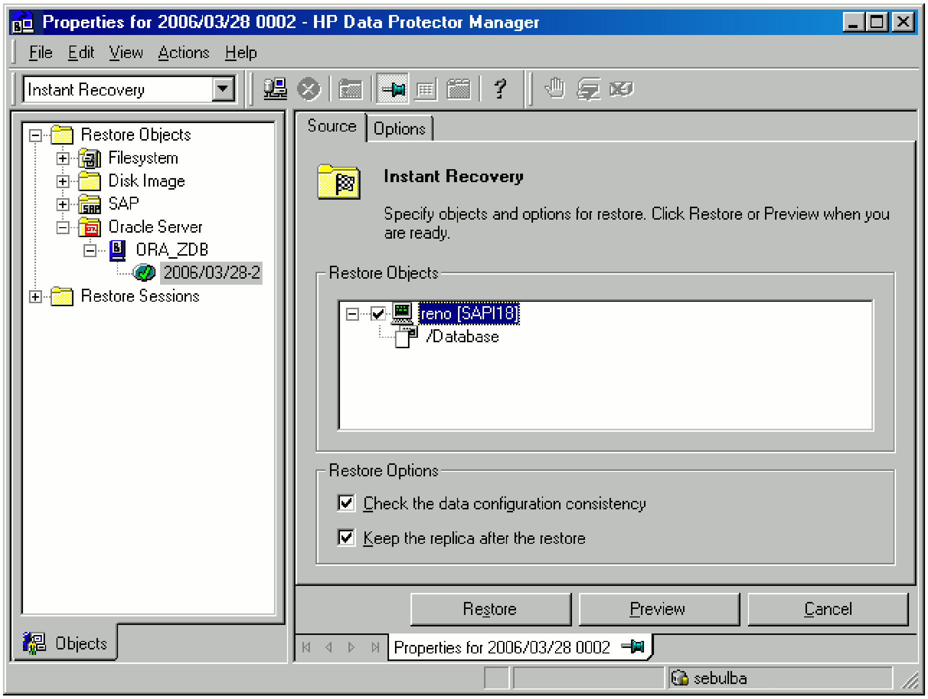

Properties for 2006/03/28 0002 - HP Data Protector Manager

**Instant Recovery**

Specify objects and options for restore. Click Restore or Preview when you are ready.

Restore Objects

- reno [SAPI18]
  - /Database

Restore Options

☑ Check the data configuration consistency

☑ Keep the replica after the restore

Figure 36 Selecting backup sessions (XP example)

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

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

> NOTE:
>
> You can perform a database recovery at a later time either from the Data Protector Manager Restore Context or manually using the RMAN CLI. See "Oracle database recovery after the instant recovery" on page 142.

   - To perform a database recovery immediately after an instant recovery, click on the **Options** tab, select **Recovery** and then select the database recovery options. For a recovery until a selected time, logseq/thread number, or SCN number, it is recommended to reset the log files. See Figure 37 on page 141 and "Restore, recovery, and duplicate options" on page 119 for details on available options.

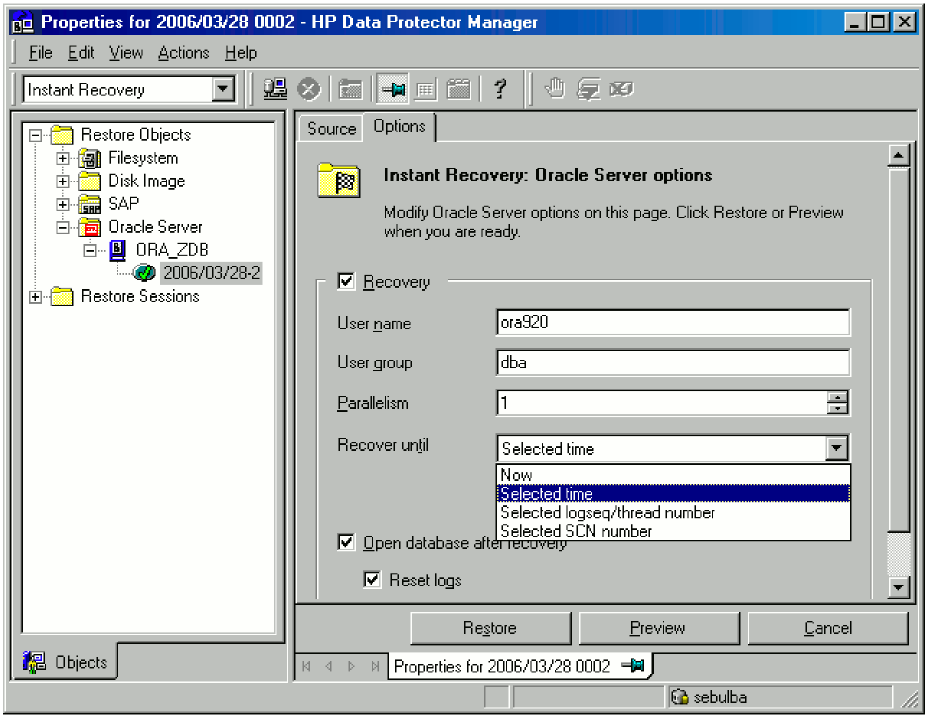

Properties for 2006/03/28 0002 - HP Data Protector Manager

Instant Recovery: Oracle Server options

Modify Oracle Server options on this page. Click Restore or Preview when you are ready.

☑ Recovery

User name: ora920

User group: dba

Parallelism: 1

Recover until: Selected time

☑ Open database after recovery

☑ Reset logs

Figure 37 Oracle recovery options

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

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

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

```
rman target Target_Database_Login catalog
Recovery_Catalog_Login
```

```
RMAN> RESET DATABASE;
```

```
RMAN> exit
```

## Oracle database recovery after the instant recovery

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

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

   ```
   startup mount
   ```

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

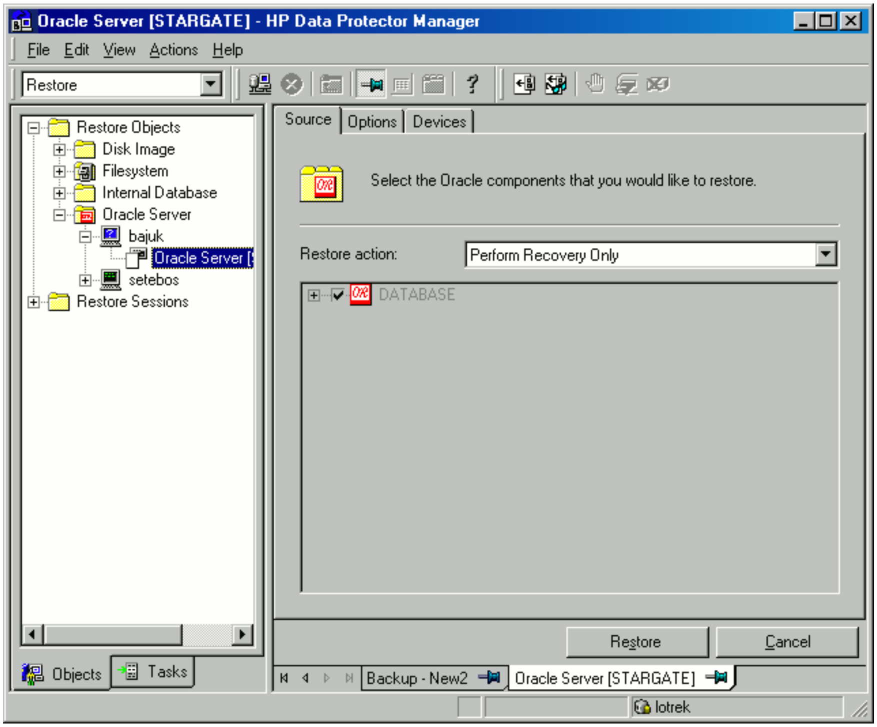

Select the Oracle components that you would like to restore.

Restore action: Perform Recovery Only

**Figure 38 Selecting the database for recovery**

     b. In the **Options** tab, select the recovery options. For details, see "Restore, recovery, and duplicate options" on page 119.
     c. Click **Restore**.
   - Perform a manual database recovery using RMAN.

     Run the following RMAN script to recover the database:

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

For additional examples on how to recover the database after an instant recovery, see "Restoring Oracle using RMAN" on page 123.

## Oracle in Veritas Cluster instant recovery

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

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

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

1. On the application system, enter the following commands to disable the two Veritas Cluster resources:

   a. `hares -offline application_resource_name -sys system`

      where *application_resource_name* is the name of the Veritas Cluster application resource for the Oracle application and *system* is the name of the active node.

      `hares -offline mountpoint_resource_name -sys system`

      where *mountpoint_resource_name* is the name of the Veritas Cluster mountpoint resource for the Oracle database files and *system* is the name of the active node.

   b. `hares -modify application_resource_name Enabled 0`

      where *application_resource_name* is the name of the Veritas Cluster application resource for the Oracle application.

      `hares -modify mountpoint_resource_name Enabled 0`

      where *mountpoint_resource_name* is the name of the Veritas Cluster mountpoint resource for the Oracle database files.

2. Perform an instant recovery.

3. If you performed only an instant recovery without the database recovery, use RMAN as described in "Oracle database recovery after the instant recovery" on page 142 to bring the Oracle database to a consistent state.

4. On the application system, enter the following commands to enable the two Veritas Cluster resources:

   a. `hares -modify` *`mountpoint_resource_name`* `Enabled 1`

      where *`mountpoint_resource_name`* is the name of the Veritas Cluster mountpoint resource for the Oracle database files.

      `hares -modify` *`application_resource_name`* `Enabled 1`

      where *`application_resource_name`* is the name of the Veritas Cluster application resource for the Oracle application.

   b. `hares -online` *`application_resource_name`* `-sys` *`system`*

      where *`application_resource_name`* is the name of the Veritas Cluster application resource for the Oracle application and *`system`* is the name of the active node.

      `hares -online` *`mountpoint_resource_name`* `-sys` *`system`*

      where *`mountpoint_resource_name`* is the name of the Veritas Cluster mountpoint resource for the Oracle database files and *`system`* is the name of the active node.

## Aborting sessions

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

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

For details on how to set environmental variables, see "Setting environmental variables" on page 71. For details on how to set the corresponding `omnirc` options, see the online Help index: "omnirc option". Note that environmental variables override omnirc options.

## Troubleshooting

This section contains a list of general checks and verifications and a list of problems you might encounter when using the Data Protector Oracle integration. You can start at “Problems” on page 154 and if you cannot find a solution there, perform general checks and verifications.

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

For general ZDB, restore, and instant recovery troubleshooting information, see the troubleshooting sections in the *HP Data Protector zero downtime backup administrator's guide*.

## Before you begin

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

## Checks and verifications

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

If your configuration, backup, or restore failed:

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

       ```
       ORACLE_HOME="ORACLE_HOME"
       export ORACLE_HOME
       DB_NAME="DB_NAME"
       export DB_NAME
       ```

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

       ```
       setenv ORACLE_HOME "ORACLE_HOME"
       setenv DB_NAME "DB_NAME"
       ```

     ***Windows:*** Set the *ORACLE_HOME* and *DB_NAME* variables.

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

   ```
   sqlplus /nolog
   ```

3. Start SQL*Plus and type:

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

   If this fails, open the Oracle target database.

- On the application system, verify that you can access the recovery catalog (if used) and that it is opened as follows:

  1. Export or set the *`ORACLE_HOME`* and *`DB_NAME`* variables as described in Step 1 on page 146.
  2. Start SQL*Plus from the `bin` directory in the *`ORACLE_HOME`* directory:

     ```
     sqlplus /nolog
     ```

  3. Start SQL*Plus and type:

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

  If this fails, open the recovery catalog.

- Verify that the listener is correctly configured for the Oracle target database and the recovery catalog database. This is required to properly establish network connections:

  1. Export or set the *`ORACLE_HOME`* variable as described in Step 1 on page 146.

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

```
lsnrctl status service
```

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

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

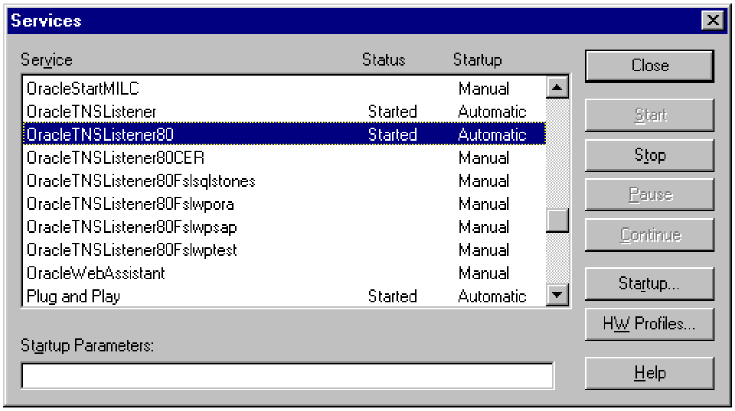

Services

| Service | Status | Startup |
|---|---|---|
| OracleStartMILC | | Manual |
| OracleTNSListener | Started | Automatic |
| OracleTNSListener80 | Started | Automatic |
| OracleTNSListener80CER | | Manual |
| OracleTNSListener80Fslsqlstones | | Manual |
| OracleTNSListener80Fslwpora | | Manual |
| OracleTNSListener80Fslwpsap | | Manual |
| OracleTNSListener80Fslwptest | | Manual |
| OracleWebAssistant | | Manual |
| Plug and Play | Started | Automatic |

Startup Parameters:

### Figure 39 Checking the status of the Oracle listener

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

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

```
sqlplus /nolog
```

4. Start SQL*Plus and type:

```
connect Target_Database_Login
```

```
exit
```

and then

```
connect Recovery_Catalog_Login
```

```
exit
```

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

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

    Export or set the *ORACLE_HOME* and *DB_NAME* variables as described in Step 1 on page 146.

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

    ```
    > exit
    ```

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

  - If you do not use the recovery catalog database:

    Export or set the *ORACLE_HOME* and *DB_NAME* variables as described in Step 1 on page 146.

    ```
    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.

    See the section "Recovery Manager Connection Options" in the *Oracle Backup and Recovery Guide* for information.

- If you use the recovery catalog database, verify that the target database is registered in the recovery catalog:
  1. Export or set the *ORACLE_HOME* variable as described in Step 1 on page 146.
  2. Start SQL*Plus from the `bin` directory in the *ORACLE_HOME;* directory:

     ```
     sqlplus /nolog
     ```

  3. Start SQL*Plus and type:

     ```
     connect Recovery_Catalog_Login;
     ```

     ```
     select * from rc_database;
     ```

     ```
     exit
     ```

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

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

  If you use the recovery catalog:

  1. Export or set the *ORACLE_HOME* variable as described in Step 1 on page 146.
  2. Start RMAN from the `bin` directory in the *ORACLE_HOME* directory:

     ```
     rman target Target_Database_Login catalog
     Recovery_Catalog_Login
     ```

  If you do not use the recovery catalog:

  1. Export or set the *ORACLE_HOME* variable as described in Step 1 on page 146.
  2. Start RMAN from the `bin` directory in the *ORACLE_HOME* directory:

     ```
     rman target Target_Database_Login nocatalog
     ```

  An example of the RMAN backup script is presented below:

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

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

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

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

Additionally, if your configuration or backup failed:

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

  See the *HP Data Protector installation and licensing guide* for details.
- Check if the `SYSDBA` privilege is granted to the Oracle administrator.
- If you have special Oracle environment settings, ensure that they are entered in the Data Protector Oracle configuration files on the Cell Manager. See the `util_cmd` man page or the *HP Data Protector command line interface reference*

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

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

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

  Ensure that the hostname defined in the backup specification as a system to be backed up is the name of the application system.

- On Windows, check the `Data Protector Inet` service startup parameters on the Oracle Server system:

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

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

- Examine the system errors reported in the following file on the application system (Oracle proxy-copy ZDB method) or backup system (Oracle backup set ZDB method):

  ***HP-UX , Solaris:*** `/var/opt/omni/log/debug.log`

  ***Windows:*** `Data_Protector_home\log\debug.log`

Additionally, if your backup or restore failed:

- Test the Data Protector internal data transfer using the `testbar2` utility:
  1. Verify that the Cell Manager name is correctly defined on the Oracle Server system. Check the following file, which contains the name of the Cell Manager system:

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

     ***Windows:*** `Data_Protector_home\Config\client\cell_server`

2. From the `bin` directory in the *`ORACLE_HOME`* directory, run:

   ***If backup failed:***

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

   ***If restore failed:***

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

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

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

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

- Check if the user under which the backup or restore session was started has appropriate Oracle permissions (for example, belongs to the `DBA` group). This user must also be in the Data Protector `operator` or `admin` user group.
- Check that the respective Data Protector user group has the `See private objects` user right enabled.
- ***If backup failed:***

  Create an Oracle backup specification to back up to a null device or file. If the backup succeeds, the problem may be related to the backup devices. See the *HP Data Protector troubleshooting guide* for instructions on troubleshooting devices.
- ***If restore failed:***

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

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

Additionally, if your restore failed:

- Verify that an object exists on the backup media.

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

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

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

- Ensure that the database is in the correct state.

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

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

    Open a command window and enter the following:

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

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

    Open a command window and enter the following:

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

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

  For information, see "Restoring Oracle using RMAN" on page 123.

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

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

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

  Open a SQLplus window and use the following command:

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

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

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

# Problems

## Problem

**SQL*Plus is unable to connect to destination**

## Action

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

## Problem

**The following error is displayed:** `ORA-12532: : invalid argument.`

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

## Action

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

## Problem

**Backup set ZDB is aborted after 10 minutes**

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

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

The ZDB session then aborts with the following message:

```
Bar backup se0ssion was started but no client connected in
600 seconds.
```

## Action

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

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

  `SmWaitForFirstClient=`*`minutes`*

- If you performed a clean installation:

  `SmWaitForFirstBackupClient=`*`minutes`*

See the *HP Data Protector troubleshooting guide* for more information on the global options file.

### Problem

**Backup Set ZDB fails after changing the physical schema of the Database**

Backup fails after you have modified the physical schema of a database, for example, if you added or dropped a tablespace, added a new datafile, or added or dropped a rollback segment. Depending on the performed modification, different error messages are displayed, for example:

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

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

### Action

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

### Problem

**Proxy copy restore fails**

Proxy copy restore fails with the following error:

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

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

### Action

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

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

At the `RMAN>` prompt, enter

```
list backup;
```

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

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

```
resync catalog;
```

After the synchronization is performed, restore should be possible.

## Problem

### Restore after a switch between ZDB methods fails

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

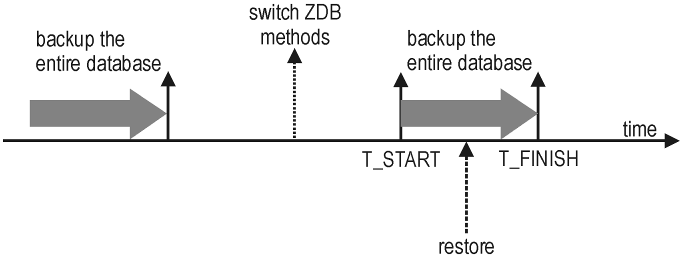

**Figure 40 Restore after a switch between ZDB methods fails**

## Action

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

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

```
run {
ALOCATE CHANNEL 'dev_0' TYPE 'sbt_tape'
PARMS 'SBT_LIBRARY=Path_to_Data_Protector_MML,
ENV(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME)';
ALOCATE CHANNEL 'dev_1' TYPE 'sbt_tape'
PARMS 'SBT_LIBRARY=Path_to_Data_Protector_MML,
ENV(OB2BARTYPE=Oracle8,OB2APPNAME=DB_NAME,
OB2PROXYCOPY=1)';
RESTORE DATAFILE list_of_backup_set_backups UNTIL
T_RESTORE CHANNEL 'dev_0';
RESTORE DATAFILE list_of_proxy-copy_backups UNTIL
T_RESTORE CHANNEL 'dev_1';
RELEASE 'dev_0';
RECOVER DATABASE UNTIL T_RECOVER ...
RELEASE 'dev_1';
}
```

Where:

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

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

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

## Problem

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

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

```
*RETVAL*8422
```

## Action

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

## 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 UNIX, Data Protector reports “Cannot allocate/attach shared memory”**

Backup fails and the following error message is displayed:

```
Cannot allocate/attach shared

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

## Action

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

## Problem

**Backup fails after a point in time restore and recovery**

The following error is displayed:

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

### Action

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

```
rman target Target_Database_Login catalog
Recovery_Catalog_Login
```

```
RMAN> RESET DATABASE;
```

```
RMAN> exit
```

### Problem

**Oracle online backup fails with the following error:**

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

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

### Action

To catalog the *controlfilecopy* to the RMAN catalog:

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

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

### Problem

**Backup of archive logs on RAC cannot be performed**

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

### Action

Edit the archive logs backup specification:

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

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

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

## Problem

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

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

## Action

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

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

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

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

Problem

**Binary util_orarest is missing**

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

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

Action

To resolve the problem:

1. Close Data Protector.
2. Set the environment variable on the system where the Cell Manager resides:

   `OB2_ORARESTHOSTNAME =` *`target_Oracle_host`*
3. Restart Data Protector and try to restore the database items again.
4. When the restore is complete, close Data Protector and re-set the following environment variable:

   `OB2_ORARESTHOSTNAME =` *`empty`*
5. Restart Data Protector.

Problem

**How to modify the RMAN restore script**

When you start a restore of an Oracle database using the Data Protector GUI or CLI, an RMAN restore script is created, which is instantly run, so you cannot edit it first.

Action

To edit the script before it is run, set the Data Protector `omnirc` variable `OB2RMANSAVE` to point to an existing directory. When the variable is set and you start a restore, the RMAN restore script, which is created at run time, is saved to the specified location under the name `RMAN_restore_`*`backup_specification_name`*`.rman`, and the actual restore is skipped. Then you can edit the script and run it manually afterwards. On how to set the omnirc variable, see the online Help index: “omnirc options”.

To start a restore using Data Protector again, clear the OB2RMANSAVE variable by deleting its content or commenting or removing the whole variable. If you comment or remove the variable on a Windows client, restart the Data Protector Inet service for the settings to take effect.

# 2 Data Protector SAP R/3 ZDB integration

## Introduction

This chapter explains how to configure and use the Data Protector SAP R/3 ZDB integration (**SAP R/3 ZDB integration**). It describes concepts and methods you need to understand to back up and restore the following files of the SAP R/3 database environment (**SAP R/3 objects**):

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

Data Protector supports offline and online backups. During an online backup, the SAP R/3 application is actively used.

Data Protector offers interactive and scheduled backups of the following types:

- ZDB to disk
- ZDB to tape
- ZDB to disk+tape

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

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

Table 8 on page 164 shows which restore methods are available, depending on the ZDB session you restore from.

### Table 8 Backup and restore sessions

| ZDB type | Restore methods |
|---|---|
| ZDB to tape | Standard restore |
| ZDB to disk | Instant recovery |
| ZDB to disk+tape | Standard restore, instant recovery |

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

### Table 9 Supported disk arrays and array configurations

| Supported array | Supported configurations |
|---|---|
| EMC Symmetrix | TimeFinder, SRDF, combined SRDF+TimeFinder |
| HP StorageWorks Disk Array XP (XP) | BC, CA, combined BC+CA |
| HP StorageWorks Enterprise Virtual Array (EVA) | BC, combined CA+BC |
| HP StorageWorks Virtual Array (VA) | BC |

This chapter provides information specific to the Data Protector SAP R/3 ZDB integration. For general Data Protector procedures and options, see the *online Help*. For details on ZDB terminology, ZDB types, advantages of offline and online backups, and instant recovery concepts, see the *HP Data Protector zero downtime backup concepts guide*.

# Integration concepts

Figure 41 on page 165 shows the architecture of the Data Protector SAP R/3 ZDB integration. The figure illustrates the preferred configuration, in which the Oracle control file, online redo log files, and Oracle SPFILE reside on a different volume group than the Oracle data files.

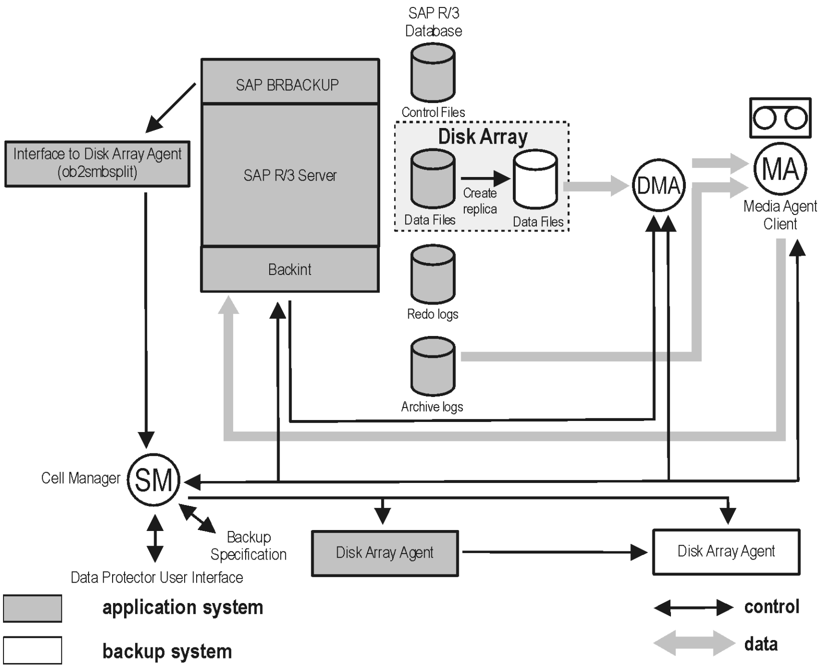

**Figure 41 SAP R/3 integration architecture**

For other supported configurations, see "ZDB integrations omnirc variables" on page 425.

## ZDB flow

For details on how ZDB options affect the ZDB flow, including mounting, activating volume/disk groups and so on, see the *HP Data Protector zero downtime backup administrator's guide*.

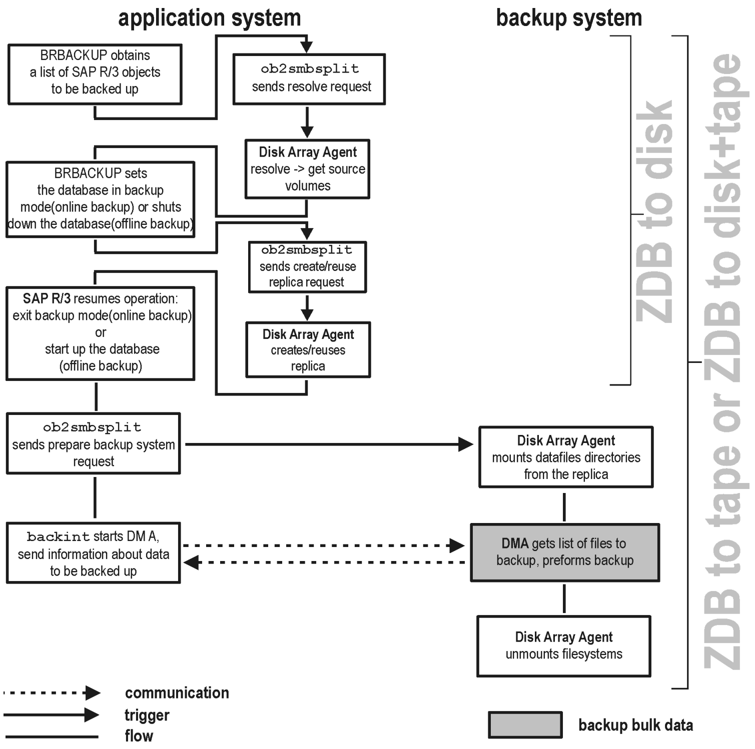
application system
backup system
BRBACKUP obtains a list of SAP R/3 objects to be backed up
ob2smbsplit sends resolve request
Disk Array Agent resolve -> get source volumes
BRBACKUP sets the database in backup mode(online backup) or shuts down the database(offline backup)
ob2smbsplit sends create/reuse replica request
Disk Array Agent creates/reuses replica
SAP R/3 resumes operation: exit backup mode(online backup) or start up the database (offline backup)
ob2smbsplit sends prepare backup system request
Disk Array Agent mounts datafiles directories from the replica
backint starts DMA, send information about data to be backed up
DMA gets list of files to backup, preforms backup
Disk Array Agent unmounts filesystems
ZDB to disk
ZDB to tape or ZDB to disk+tape
communication
trigger
flow
backup bulk data

## Figure 42 SAP R/3 ZDB session flow (BRBACKUP started on the application system)

1. The SAP R/3 backup specification is read and the Data Protector `omnisap.exe` program is started on the application system.

2. `Omnisap.exe` starts BRBACKUP, which switches the database to the backup mode (online backup) or shuts down the database (offline backup), and starts the split command on the application system, with the list of files to be included in the replica creation.

   The database is put out of the backup mode (online backup) or is restarted (offline backup) after the replica is created.

NOTE:

Data Protector can use the `splitint` interface (if BRTOOLS supports it) to reduce the time during which the database is in the backup mode.

3. The Data Protector `ob2smbsplit` command resolves the backup configuration, creates a replica, and preparing the replica for backup.

   The mountpoints for the backed up object are created on the backup system.

   The backup volume/disk groups are activated and the filesystems are mounted on the backup system.
4. ***ZDB to disk only:*** The remaining ZDB options are processed and the details on the session are written to the ZDB database. The session then finishes.
5. ***ZDB to tape, ZDB to disk+tape:*** BRBACKUP starts the Data Protector `backint` program, which starts establishing a connection between the Data Protector Data Movement Agents (DMA) on the backup system and the General Media Agents (MA). The process is coordinated by the Data Protector Backup Session Manager (BSM). When the connection is established, the data specified for backup is streamed to tape.
6. When the data transfer completes, the backup system is disabled (filesystems are unmounted on all platforms and volume/disk groups deactivated on UNIX).
7. ***EMC and XP only:*** Links are re-established, depending on how ZDB options are specified.

# Data Protector SAP R/3 configuration file

Data Protector stores the integration parameters for every configured SAP R/3 database in the following file on the Cell Manager:

- On UNIX:
  `/etc/opt/omni/server/integ/config/SAP/`*`client_name%ORACLE_SID`*
- On Windows:
  *`Data_Protector_home`*`\Config\Server\Integ\Config\Sap\`*`client_name%ORACLE_SID`*

The parameters stored are:

- Oracle home directory
- encoded connection string to the target database
- BRTOOLS home directory
- the variables which need to be exported prior to starting a backup
- SAPDATA home directory
- user name and user group
- temporary directory used for the copy of the control file or redo logs
- list of control files and redo logs that will be copied to a safe location

- character set (ORA_NLS_CHARACTERSET)
- speed parameters (time needed for a specific file to back up - in seconds)
- manual balancing parameters

The configuration parameters are written to the Data Protector SAP R/3 configuration file:

- during configuration of the integration
- during creation of a backup specification
- when the configuration parameters are changed

IMPORTANT:

To avoid problems with your backups, take extra care to ensure the syntax and punctuation of your configuration file match the examples.

NOTE:

You can set up the parameters in the `Environment` section (sublist) of the file by referring to other environment variables in the following way:

```
SAPDATA_HOME=${ORACLE_HOME}/data
```

## Syntax

The syntax of the Data Protector SAP R/3 configuration file is as follows:

```
ORACLE_HOME='ORACLE_HOME';
ConnStr='ENCODED_CONNECTION_STRING_TO_THE_TARGET_DATABASE';
BR_directory='BRTOOLS_HOME';
SAPDATA_HOME='SAPDATA_HOME';
ORA_NLS_CHARACTERSET='CHARACTER_SET';
OSUSER='USER_NAME';
OSGROUP='USER_GROUP';
Environment={
[ENV_var1='value1';]
[ENV_var2='value2';
...]
}
SAP_Parameters={backup_spec_name=('-concurrency #_of_concurrency
' | '-time_balance' | '-load_balance' | '-manual_balance');
}
speed={
```

```
AVERAGE=1;
'filename'=#_of_seconds_needed_to_back_up_this_file;
}
compression={'filename'=size_of_the_file_in_bytes_after_the
_compression;
}
manual_balance={backup_specification_name={
'filename'=device_number;
}
}
```

The ORA_NLS_CHARACTERSET parameter is set automatically by Data Protector during SAP R/3 database configuration. For details on how to configure SAP R/3 database for use with Data Protector, see "Configuring SAP R/3 databases" on page 182.

## Example

This is an example of the file:

```
ORACLE_HOME='/app/oracle805/product';
ConnStr='EIBBKIBBEIBBFIBBGHBBOHBB
QDBBOFBBCFBBPFBBCFBBIFBBGFBBDGBBBFBBCFBBDFBBCFBB';
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. The command resides in the *`Data_Protector_home`*`\bin` (Windows systems) or `/opt/omni/lbin` (HP-UX, Solaris systems) directory.

### Cluster-aware clients

In a cluster environment, the environment variable `OB2BARHOSTNAME` must be defined as the virtual hostname before running the `util_cmd` command from the command line (on the client). The `OB2BARHOSTNAME` variable is set as follows:

- On UNIX: `export OB2BARHOSTNAME=`*`virtual_hostname`*
- On Windows: `set OB2BARHOSTNAME=`*`virtual_hostname`*

### The `util_cmd` synopsis

The syntax of the `util_cmd` command is as follows:

`util_cmd -getconf[ig] SAP `*`oracle_instance`*` [-local `*`filename`*`]`

`util_cmd -getopt[ion] [SAP `*`oracle_instance`*`] `*`option_name`*
`[-sub[list] `*`sublist_name`*`] [-local `*`filename`*`]`

```
util_cmd -putopt[ion] [SAP oracle_instance] option_name
[option_value] [-sub[list] sublist_name] [-local filename]
```

where:

*option_name* is the name of the parameter

*option_value* is the value for the parameter

`[-sub[list] sublist_name]` specifies the sublist in the configuration file to which a parameter is written to or taken from.

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

- When it is used with the `-getconf[ig]` option, it specifies the filename for the output of the command to be written to. If the `-local` option is not specified, the output is written to the standard output.
- When it is used with the `-getopt[ion]`, it specifies the filename of the file from which the parameter and its value are to be taken and then written to the standard output. If the `-local` option is not specified, the parameter and its value are taken from the Data Protector SAP R/3 configuration file and then written to the standard output.
- When it is used with the `-putopt[ion]` option, it specifies the filename for the output of the command to be written to. If the `-local` option is not specified, the output is written to the Data Protector SAP R/3 configuration file.

NOTE:

If you are setting the *option_value* parameter as a number, the number must be put in single quotes, surrounded by double quotes.

## Return values

The `util_cmd` command displays a short status message after each operation (writes it to the standard error):

- `Configuration read/write operation successful.`

  This message is displayed when all the requested operations have been completed successfully.

- `Configuration option/file not found.`

  This message is displayed when either an option with the specified name does not exist in the configuration, or the file specified as the `-local` parameter does not exist.

- `Configuration read/write operation failed.`

  This message is displayed if any fatal errors occurred, for example: the Cell Manager is unavailable, the Data Protector SAP R/3 configuration file is missing on the Cell Manager, etc.

## Setting parameters

To set the Data Protector `OB2OPTS` and the Oracle `BR_TRACE` parameters for the Oracle instance `ICE` that is run by SAP R/3, use the following commands on the Data Protector SAP R/3 client:

### Windows

```
Data_Protector_home\bin\util_cmd -putopt SAP ICE OB2OPTS
'-debug 1-200 debug.txt' -sublist Environment
```

```
Data_Protector_home\bin\util_cmd -putopt SAP ICE BR_TRACE
"'10'" -sublist Environment
```

### HP-UX, Solaris

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

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

## Retrieving parameters

To retrieve the value of the `OB2OPTS` parameter for the Oracle instance `ICE`, use the following command on the Data Protector SAP R/3 client:

- On Windows: `Data_Protector_home\bin\util_cmd -getopt SAP ICE OB2OPTS -sublist Environment`
- On HP-UX, Solaris: `/opt/omni/lbin/util_cmd -getopt SAP ICE OB2OPTS \ -sublist Environment`

## Listing parameters

To list all the Data Protector SAP R/3 configuration file parameters for the Oracle instance `ICE`, use the following command on the Data Protector SAP R/3 client:

- On Windows: `Data_Protector_home\bin\util_cmd -getconf SAP ICE`
- On HP-UX, Solaris: `/opt/omni/lbin/util_cmd -getconf SAP ICE`

### Deleting parameters

To remove the value of the `OB2OPTS` parameter for the Oracle instance `ICE`, use the following command on the Data Protector SAP R/3 client:

- On Windows: `Data_Protector_home\bin\util_cmd -putopt SAP ICE OB2PTS "" -sublist Environment`
- On HP-UX, Solaris: `/opt/omni/lbin/util_cmd -putopt SAP ICE OB2OPTS "" -sublist Environment`

# Configuring the integration

To configure the integration:

1. Configure the required user accounts. See "Configuring user accounts" on page 175.
2. Configure SQL*Net V2 or Net8 TNS listener. See "Configuring SQL*Net V2 or Net8 TNS listener" on page 177.
3. Check the connection to the Oracle database from the application system. See "Checking the connection" on page 178.
4. Enable the use of the authentication password file. See "Authentication password file" on page 178.
5. Optionally, set the archived logging mode to enable online backups. See "Enabling archived logging" on page 179.
6. Share directories on the application system. See "Sharing directories on the application system" on page 180.
7. Configure every SAP R/3 database you intend to back up from or restore to. See "Configuring SAP R/3 databases" on page 182.
8. Configure the SAP R/3 parameter file. See "Configuring the SAP R/3 parameter file" on page 190.

## Prerequisites

- Ensure that you have correctly installed and configured the SAP R/3 application. The database used by the SAP R/3 application must be an Oracle database. If any other database is used, you can back it up using the corresponding Data Protector integration. It is assumed that you are familiar with the SAP R/3 application and Oracle database administration.

    - For supported versions, platforms, devices, and other information, see the latest support matrices at http://www.hp.com/support/manuals.
    - For information on installing, configuring, and using the SAP R/3 application and the SAP backup and restore tools (BRBACKUP, BRRESTORE, and BRARCHIVE), see the SAP R/3 application documentation.
- Ensure that you have a license to use the Data Protector SAP R/3 ZDB integration. For information, see the *HP Data Protector installation and licensing guide*.
- Ensure that you have correctly installed Data Protector.
    - For information on how to install the Data Protector disk array integration (EMC, XP, VA, or EVA) with SAP R/3 in various architectures, see the *HP Data Protector installation and licensing guide*.
    - On how to configure the Data Protector ZDB integration (EMC, XP, VA, or EVA), see the *HP Data Protector zero downtime backup administrator's guide*.
    - For information on the Data Protector Cell Manager package configuration in the MC/SG cluster, see the online Help index: “MC/ServiceGuard integration”.

NOTE:

You cannot run ZDB sessions in the RMAN mode.

## Before you begin

- Configure devices and media for use with Data Protector.
- To test whether the SAP R/3 system and the Cell Manager communicate properly, configure and run a Data Protector filesystem backup and restore.
- ***Windows systems except Windows Server 2008:*** Restart the Data Protector `Inet` service under the Oracle operating system user account described in “Configuring user accounts” on page 175. For information on changing the Data Protector Inet account, see the online Help index: “changing Data Protector Inet account”.

  If there are several SAP R/3 instances running on the same system with different SAP administrator accounts configured for each instance, create an additional, common SAP administrator account. Configure the Data Protector `Inet` service to use this account as the service startup account.

## Cluster-aware clients

- Configure SAP R/3 databases only on one cluster node, since the configuration files reside on the Cell Manager.

  ***UNIX:*** During the configuration, Data Protector creates a link to the Data Protector `backint` and `splitint` programs on the currently active node. On all the other nodes, do it manually. Run:

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

  If `splitint` is supported by BRTOOLS, also run:

  ```
  ln -s /opt/omni/lbin/ob2smbsplit \
  /usr/sap/ORACLE_SID/sys/exe/run/splitint
  ```

  ***Windows:*** During the configuration, Data Protector copies the Data Protector `backint` and `ob2smbsplit.exe` programs (the latter only if `splitint` is supported by BRTOOLS) from *`Data_Protector_home`*`\bin` to the directory that stores the SAP backup tools and renames `ob2smbsplit.exe` to `splitint.exe`. This is done only on the currently active node. On the other node, do it manually.

- If you intend to use the Data Protector CLI, set the Data Protector environment variable `OB2BARHOSTNAME` to the virtual server name as follows:

  ***Windows:*** `set OB2BARHOSTNAME=`*`virtual_server_name`*

  ***UNIX:*** `export OB2BARHOSTNAME=`*`virtual_server_name`*

- ***Tru64:*** Create the following links:

  ```
  ln -s /sapfiles/admin/dbs/initsap.dba initSAP.dba
  ln -s /sapfiles/admin/dbs/initsap.ora initSAP.ora
  ln -s /sapfiles/admin/dbs/initsap.sap initSAP.sap
  ```

NOTE:

SAP recommends to install SAP backup utilites 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:<br>• ***UNIX systems:*** dba and sapsys<br>• ***Windows systems:*** ORA_DBA and ORA_*SID*_DBA local groups<br>For example, user ora*SID*.<br>***UNIX systems only:*** Ensure that this user is the owner of the filesystem or of the raw logical volume on which the database is mounted. The minimum permissions should be 740. |
| 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:<br>• sysdba<br>• sysoper<br>For example, user system. |

Add the following user accounts to the Data Protector admin or operator user group:

- Oracle operating system user account
- ***UNIX systems only:*** User account root (for both the application system and backup system)

In cluster environments, add these user accounts to the Data Protector admin or operator user group for the following clients:

- virtual server
- every node in the cluster

For information on adding Data Protector users, see the online Help index: “adding users”.

## Configuring SQL*Net V2 or Net8 TNS listener

1. Ensure that the `listener.ora` and `tnsnames.ora` files on the application system are configured as shown in the following example. The files are located in:

   ***UNIX:*** *`ORACLE_HOME`*`/network/admin`

   ***Windows:*** *`ORACLE_HOME`*`\network\admin`

### Example

Oracle instance: `PRO`

Application system: `alpha.hp.com`

| | |
|---|---|
| `listener.ora` | `LISTENER =`<br>`(DESCRIPTION_LIST =`<br>`(DESCRIPTION =`<br>`(ADDRESS =`<br>`(PROTOCOL = TCP)  (HOST = alpha.hp.com)`<br>`(PORT = 1522)`<br>`)`<br>`)`<br>`)`<br><br>`SID_LIST_LISTENER =`<br>`(SID_LIST =`<br>`(SID_DESC =`<br>`(GLOBAL_DBNAME = PRO)`<br>`(SID_NAME = PRO)`<br>`(ORACLE_HOME = /app/oracle815/product)`<br>`)`<br>`)` |
| `tnsnames.ora` | `PRO =`<br>`(DESCRIPTION =`<br>`(ADDRESS_LIST =`<br>`(ADDRESS = (PROTOCOL = TCP)`<br>`(HOST = alpha.hp.com)  (PORT = 1522))`<br>`)`<br>`(CONNECT_DATA =  (SERVICE_NAME = PRO)`<br>`)`<br>`)` |

2. Start the SQL*Net V2 or Net8 TNS listener by running the following on the Oracle Server system:

   ***UNIX:*** `ORACLE_HOME/bin/lsnrctl start`

   ***Windows:*** `ORACLE_HOME\bin\lsnrctl start`

## Checking the connection

To check the connection to the Oracle instance from the application system:

1. Log in to the application system as the Oracle OS user.
2. Export/set the *`ORACLE_HOME`* and *`ORACLE_SID`* variables.
3. Start `sqlplus`.
4. Connect to the Oracle target database as the Oracle database user, first with the `sysdba` role and then with the `sysoper` role.

### Example

For the following configuration:

Oracle instance: `PRO ORACLE_HOME`: `/app/oracle816/product`

run:

```
id
uid=102(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 on the application system.
2. In the `init`*`ORACLE_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 `init`*`ORACLE_SID`*`.ora` file, set

   `log_archive_start = true`

   and specify the `log_archive_dest` option.

### Example

This is an example of the `init`*`ORACLE_SID`*`.ora` file for the Oracle instance `PRO`:

```
# @(#)initSID.ora    20.4.6.1    SAP    98/03/30
#####################################################
#   (c)Copyright SAP AG, Walldorf
#####################################################
. . . .
. . . . . . . . . .
. . . . . . . .
. . . . . . . .
### ORACLE Authentication Password File
remote_login_passwordfile = exclusive
### ORACLE archiving
log_archive_dest = /oracle/PRO/saparch/PROarch
log_archive_start  = true
. . . .
```

2. Mount the Oracle database and start the archived logging mode using the Oracle Server Manager. Run:

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

### Example

For the Oracle instance `PRO`, run:

***UNIX:*** `export ORACLE_SID=PRO`

***Windows:*** `set ORACLE_SID=PRO`

```
sqlplus /nolog
SQLPLUS> connect user/passwd@PRO;
Connected.
SQLPLUS> startup mount
ORACLE instance started.
Total System Global Area      6060224 bytes
Fixed Size                      47296 bytes
Variable Size                 4292608 bytes
Database Buffers              1638400 bytes
Redo Buffers                    81920 bytes
Database mounted.
SQLPLUS> alter database archivelog;
Statement processed.
SQLPLUS> archive log start;
Statement processed.
SQLPLUS> alter database open;
```

## Sharing directories on the application system

The following directories on the application system must be accessible:

- `sapbackup`
- `sapreorg`
- Oracle home directory
- BR*Tools home directory

NOTE:

The `sapreorg` and BR*Tools home directories must be accessible only if you want to run SAP compliant ZDB sessions (BRBACKUP is started on the backup system and not on the application system).

## UNIX application system

1. Share the directories on the application system through NFS with root permissions.

   For example, suppose that the `sapbackup` directory on the application system points to `/oracle/SID/sapbackup` and the backup system is `backup.company.com`. To share the `sapbackup` directory:

   ***HP-UX:*** In the file `/etc/exports` on the application system, add the line:

   ```
   /oracle/SID/sapbackup -root=backup.company.com
   ```

   ***Solaris:*** In the file `/etc/dfs/dfstab` on the application system, add the line:

   ```
   share -F nfs -o root=backup.company.com
   /oracle/SID/sapbackup
   ```

2. Mount the directories on the backup system. Ensure that you have the same directory structure on both the application and backup system.

   For example, suppose that you have an HP-UX application system `app.company.com`. To mount the `/oracle/SID/sapbackup` directory on the backup system, add the following line to the file `/etc/fstab` on the backup system:

   ```
   app.company.com:/oracle/SID/sapbackup
   /oracle/SID/sapbackup nfs     defaults 0 0
   ```

## Windows application system

- On the application system, find the location of the `sapbackup` and `sapreorg` directories. If they reside inside the SAP data home directory, share this directory under any name you want. Then, on the backup system, create the `HKEY_LOCAL_MACHINE\SOFTWARE\SAP\ORACLE_SID\Environment\SAPDATA_HOME` Windows registry key, specifying the SAP data home directory path as seen from the backup system.

  For example, suppose your application system is `mycomputer.company.com` and the SAP data home is `K:\oracle\my_instance`, which you share under the name my_SAPinstance. Then, the SAPDATA_HOME Windows registry key on the backup system must have the value \\mycomputer.company.com\my_SAPinstance.

If the sapbackup and sapreorg directories do not reside together, share each directory separately and also create a separate Windows registry key on the backup system (SAPBACKUP, SAPREORG).

- Share the Oracle home directory on the application system and specify its path in the ORACLE_HOME Windows registry key on the backup system, similarly as described above.
- Ensure that the BR*Tools home directory on the application system is accessible from the backup system as follows:

  \\*application_system*\sapmnt*SAP_SID*\SYS\exe\run

TIP:

Instead of creating registry keys, you can also set the Data Protector SAP R/3 integration environment variables (ORACLE_HOME, SAPDATA_HOME, SAPBACKUP, SAPREORG).

## Choosing authentication mode

Data Protector SAP R/3 ZDB integration supports two authentication modes for accessing Oracle databases that are used by SAP R/3:

- database authentication mode
- operating system authentication mode

With database authentication mode, you need to re-configure the SAP R/3 integration for an SAP R/3 database with the new Oracle login information each time the corresponding Oracle database user account changes. Such a reconfiguration is not needed if operating system authentication mode is used.

You select the preferred authentication mode when you configure a particular SAP R/3 database.

## Configuring SAP R/3 databases

You need to provide Data Protector with the following configuration parameters:

- Oracle Server home directory
- SAP R/3 data home directory

- Optionally, if you choose database authentication mode, Oracle database user account. The user account is used by BRBACKUP and BRARCHIVE during backup.
- Directory in which the SAP backup utilities are stored

Data Protector then creates the configuration file for the SAP R/3 database on the Cell Manager and verifies the connection to the database. On UNIX, Data Protector also creates a soft link for the `backint` program from the directory that stores the SAP backup utilities to `/opt/omni/lbin`.

On Windows, Data Protector copies the `backint` and programs (the latter only if `splitint` is supported by BR*Tools) from *`Data_Protector_home`*`\bin` to the directory that stores the SAP backup tools and renames `ob2smbsplit.exe` to `splitint.exe`.

To configure an SAP R/3 database, use the Data Protector GUI or CLI.

## Before you begin

- Ensure that the SAP R/3 database is open.

## Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **SAP R/3**, and click **Add Backup**.
3. In the **Create New Backup** dialog box, select the template.

   In the **Backup type** drop-down list, select the **Split mirror backup**.

   Click **OK**.
4. In **Application system**, select the SAP R/3 client to be backed up. In cluster environments, select the virtual server.

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

   Specify other disk array specific backup options (see Figure 48 on page 196 for EMC, Figure 49 on page 197 for XP, Figure 50 on page 198 for VA, or Figure 51 on page 199 for EVA). For details on the options, press **F1**.

NOTE:

***XP, VA, and EVA only:*** To enable instant recovery, select **Track the replica for instant recovery**.

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

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

   - ***UNIX:*** In **Username**, type the Oracle OS user described in "Configuring user accounts" on page 175. In **Group/Domain name**, type `dba`.
   - ***Windows Server 2008:***

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

   Ensure that this user has been added to the Data Protector `admin` or `operator` user group and has the SAP R/3 backup rights. This user becomes the backup owner.

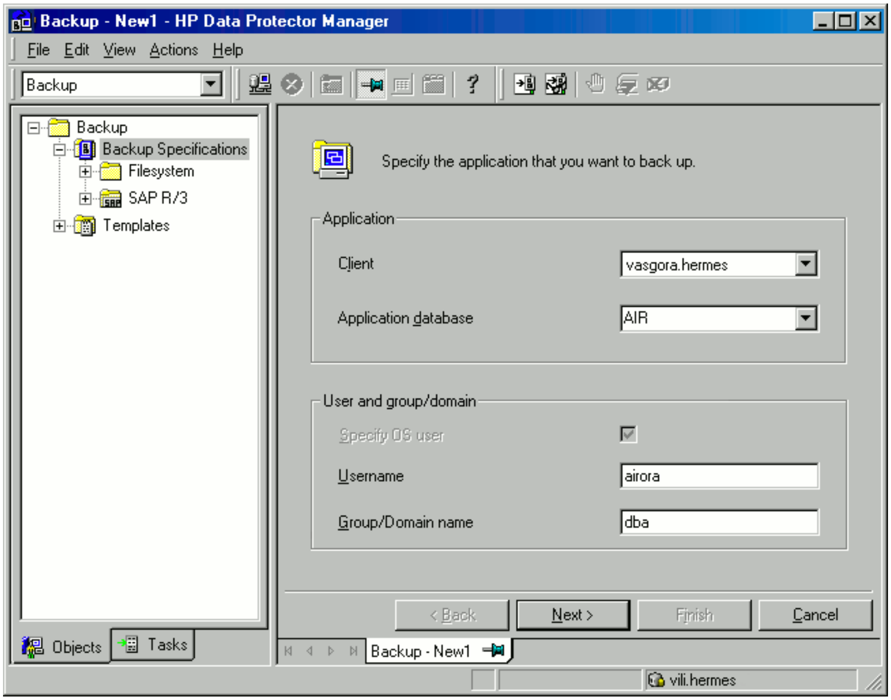

Specify the application that you want to back up.

Application

Client: vasgora.hermes

Application database: AIR

User and group/domain

Specify OS user ☑

Username: airora

Group/Domain name: dba

**Figure 43 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 175.
   - **Service**: Specify the Oracle service name.

   In **Backup and restore executables directory**, specify the pathname of the directory in which the SAP backup utilities reside. By default, the utilities reside in:

   ***UNIX:*** `/usr/sap/`*`ORACLE_SID`*`/SYS/exe/run`

   ***Windows:*** `\\`*`SAP_system`*`\sapmnt\`*`ORACLE_SID`*`\sys\exe\run`

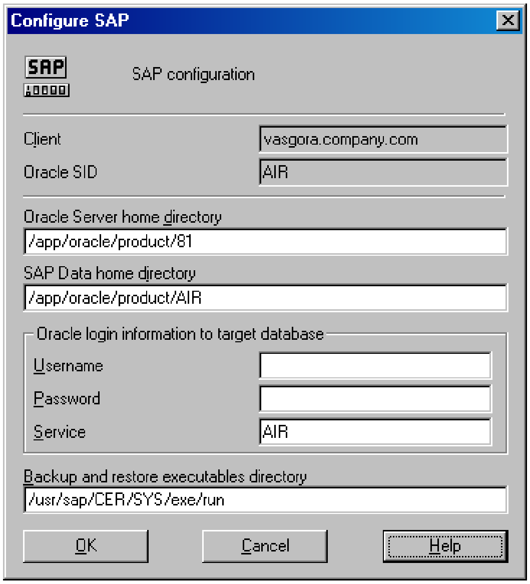

Configure SAP

SAP configuration

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:

Password:

Service: AIR

Backup and restore executables directory: /usr/sap/CER/SYS/exe/run

**Figure 44 Configuring an SAP R/3 database on a UNIX system (operating system authentication mode)**

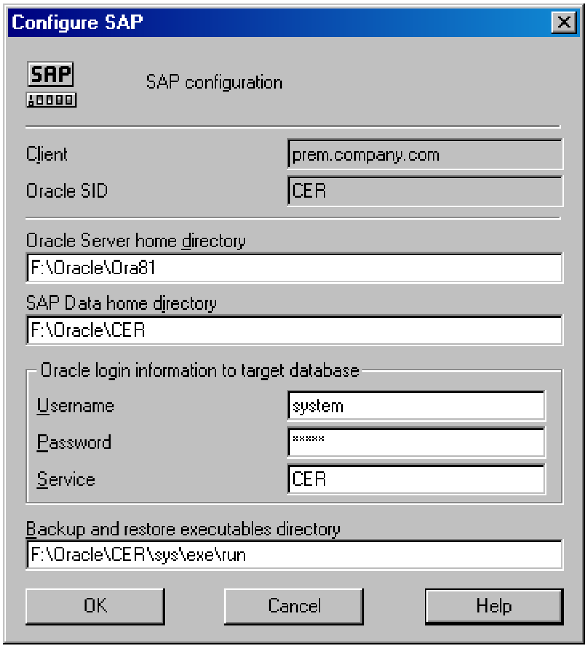

Client: prem.company.com

Oracle SID: CER

Oracle Server home directory: F:\Oracle\Ora81

SAP Data home directory: F:\Oracle\CER

Oracle login information to target database

Username: system

Password: *****

Service: CER

Backup and restore executables directory: F:\Oracle\CER\sys\exe\run

**Figure 45 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 Step 7 on page 201.

## Using the Data Protector CLI

1. Log in to the SAP R/3 system using the Oracle operating system user account.
2. At the command prompt, change current directory to the following directory:

   ***Windows systems:*** `Data_Protector_home\bin`

   ***HP-UX, Solaris systems:*** `/opt/omni/lbin`

3. Run:

```
util_sap.exe -CONFIG ORACLE_SID ORACLE_HOME
targetdb_connection_string SAPTOOLS_DIR
[SAPDATA_HOME][SQL_PATH]
```

### Parameter description

*ORACLE_SID*

Oracle instance name.

*ORACLE_HOME*

Pathname of the Oracle Server home directory.

*targetdb_connection_string*

This argument value determines the authentication mode used for accessing the Oracle database:

- To select the database authentication mode, specify the login information to the target database in the format *user_name/password@Oracle_service*.
- To select the local operating system authentication mode, specify only the character `/`.
- To select the remote operating system authentication mode, specify the login information to the target database in the format `/@`*Oracle_service*.

*SAPTOOLS_DIR*

Pathname of the directory that stores the SAP backup utilities.

*SAPDATA_HOME*

Pathname of the directory where the SAP R/3 data files are installed. By default, this parameter is set to *ORACLE_HOME*.

The message `*RETVAL*0` indicates successful configuration.

## Handling errors

If you receive the message `*RETVAL*`*error_number* where *error_number* is different than zero, an error occurred.

To get the error description:

***Windows:***

```
Data_Protector_home\bin\omnigetmsg 12 error_number
```

which is located on the Cell Manager.

***HP-UX, Solaris:*** Run:

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

TIP:

To get a list of Oracle instances that are used by the SAP R/3 application, run:

```
util_sap.exe -APP
```

To get a list of tablespaces of an Oracle instance, run:

```
util_sap.exe -OBJS0 ORACLE_SID
```

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

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

## Checking the configuration

You can check the configuration of an SAP R/3 database after you have created at least one backup specification for this database. Use the Data Protector GUI or CLI.

### Using the Data Protector GUI

1. In the Context List, select **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **SAP R/3**. Click the backup specification to display the Oracle instance to be checked.
3. Right-click the Oracle instance and click **Check configuration**.

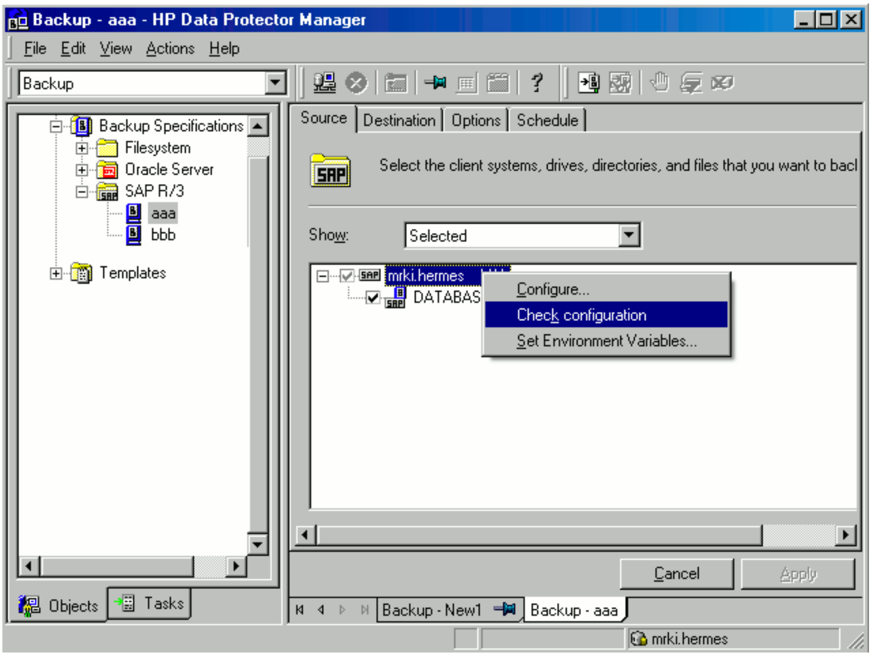

Select the client systems, drives, directories, and files that you want to bacl

**Figure 46 Checking the SAP R/3 configuration**

## Using the Data Protector CLI

Log in to the SAP R/3 system as the Oracle OS user. From the directory:

***Windows:*** `Data_Protector_home\bin`

***HP-UX, Solaris:*** `/opt/omni/lbin`

run:

```
util_sap.exe -CHKCONF ORACLE_SID
```

where *ORACLE_SID* is the name of the Oracle instance.

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

If you receive the message `*RETVAL*`*error_number* where *error_number* is different than zero, an error occurred. On how to get the error description, see “Handling errors” on page 187.

To check if the SAP R/3 configuration is suitable for instant recovery, from the directory,

***Windows:*** `Data_Protector_home\bin`

***HP-UX, Solaris:*** `/opt/omni/lbin`

run:

```
util_sap.exe -CHKCONF_IR ORACLE_SID [-verbose]
```

The `-verbose` option creates a file with a list of control files and redo log files that are on the same source volumes as the database files. If this list in not empty, a warning is displayed, stating that instant recovery is impossible. On how to solve this problem, see "Reconfiguring an Oracle instance for instant recovery" on page 421.

## Configuring the SAP R/3 parameter file

To configure the integration, you need to set some parameters in the SAP R/3 parameter file on both the application system and backup system. The file template is located in:

***UNIX:*** `ORACLE_HOME/dbs/initORACLE_SID.sap`

***Windows:*** `ORACLE_HOME\database\initORACLE_SID.sap`

### Table 10 SAP parameter file settings

| Parameter | Value/Description |
|---|---|
| `split_cmd` | On the application system:<br>***UNIX:*** `"/opt/omni/lbin/ob2smbsplit $"`<br>***Windows:*** `"Data_Protector_home\bin\ob2smbsplit $"`<br>On the backup system, you do not need to set the parameter.<br>BRBACKUP uses this parameter to trigger the replica creation. At run time, the optional sign "`$`" is replaced with the name of the text file containing the names of files to be backed up.<br>***Windows only:*** If the pathname contains spaces, use Windows short names instead. |
| `primary_db` | On the application system: `LOCAL`<br>On the backup system: name of the service used for connecting to the Oracle database.č<br>This parameter defines the service name of the Oracle database to link the backup system to the application system. |

# Backup

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

- ZDB to disk
- ZDB to tape
- ZDB to disk+tape

To configure a backup, create a ZDB backup specification.

Archived logs can only be backed up in a ZDB to disk+tape, ZDB to tape, or non-ZDB (standard backup) session. If you try to back up archived logs in a ZDB to disk session, the session fails.

TIP:

Create a separate (non-ZDB) backup specification for backing up archived logs.

What is backed up depends on your selection in the backup specification. For details, see Table 11 on page 191.

**Table 11 What is backed up**

| Selected items | Backed up files |
|---|---|
| **ARCHIVELOGS** | • offline (archived) redo logs<br>• control files |
| **DATABASE** or individual tablespaces | • data files<br>• control files<br>• SAP R/3 logs and parameter files<br>• online redo logs (only during offline backups) |

You can specify SAP R/3 backup options in two different ways:

- Using the BRBACKUP options.
- Using the SAP parameter file.

NOTE:

The BRBACKUP options override the settings in the SAP parameter file.

You can specify BRBACKUP options when you create a backup specification. If no options are specified, the SAP R/3 application refers to the current settings in the SAP parameter file. In such a case, before running a backup, ensure that the SAP parameter file is correctly configured.

TIP:

When you create a backup specification, select a backup template that already contains the desired BRBACKUP options.

## Considerations

- Before you start a backup, ensure that the SAP R/3 database is in the `open` or `shutdown` mode.
- ZDB, restore, and instant recovery sessions that use the same source volume on the application system cannot run simultaneously.
- You cannot start a ZDB to disk session if another session is backing up the archived logs, even if the Oracle data files and the archived logs reside on different source volumes.
- ***ZDB to disk only:*** BRBACKUP displays the following error, which should be ignored.

  ```
  Ignoring BRBACKUP return value 5 because it is disk-only
  backup
  ```
- ***XP only:*** If the LVM Mirroring configuration is used, Data Protector displays a warning during a backup because the volume group source volumes on the application system do not have their BC pairs assigned. The message should be ignored.
- ***ZDB to disk only:*** Archived logs cannot be backed up. To back up archived logs, create a non-ZDB (standard) SAP R/3 backup specification. For information, see the *HP Data Protector integration guide for Oracle and SAP*.
- Configurable backup modes are supported only through using the templates.

- By default, Data Protector supports all BRTOOL options except `-a` and `-b`. To enable support for `-a` and `-b`, set the `OB2BRTNOSECU` `omnirc` variable to `1`. On how to set the `omnirc` variable, see the online Help index: “omnirc options”.

## Creating backup specifications

Create a backup specification using the Data Protector Manager.

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **SAP R/3**, and click **Add Backup**.

3. In the **Create New Backup** dialog box, select a template (Figure 47 on page 194) and click **OK**.

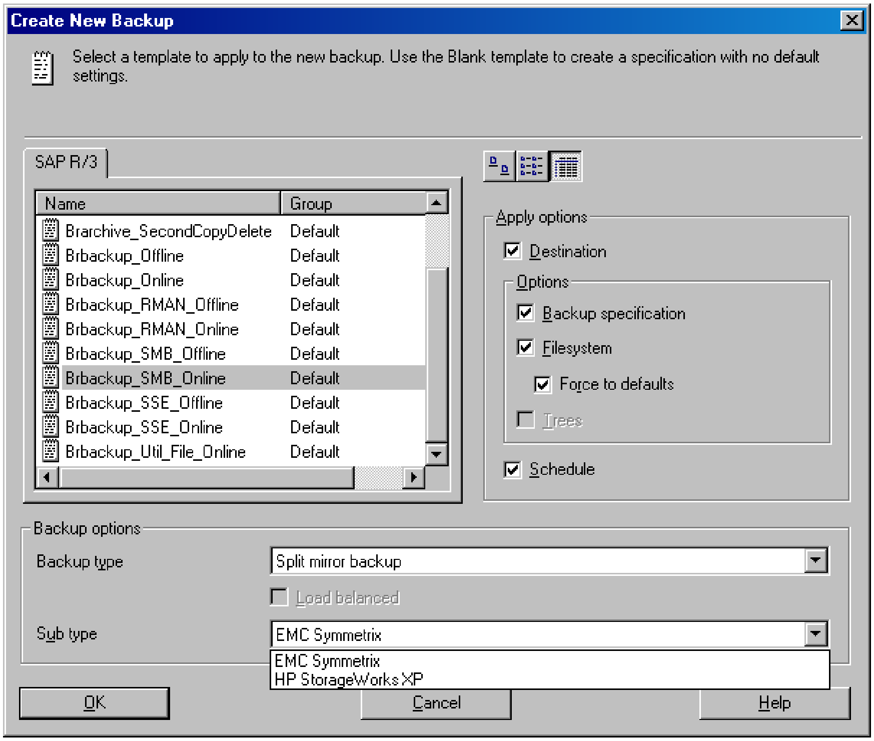

Create New Backup

Select a template to apply to the new backup. Use the Blank template to create a specification with no default settings.

SAP R/3

| Name | Group |
|---|---|
| Brarchive_SecondCopyDelete | Default |
| Brbackup_Offline | Default |
| Brbackup_Online | Default |
| Brbackup_RMAN_Offline | Default |
| Brbackup_RMAN_Online | Default |
| Brbackup_SMB_Offline | Default |
| Brbackup_SMB_Online | Default |
| Brbackup_SSE_Offline | Default |
| Brbackup_SSE_Online | Default |
| Brbackup_Util_File_Online | Default |

Apply options

☑ Destination

Options

☑ Backup specification

☑ Filesystem

☑ Force to defaults

☐ Trees

☑ Schedule

Backup options

Backup type: Split mirror backup

☐ Load balanced

Sub type: EMC Symmetrix

EMC Symmetrix
HP StorageWorks XP

OK | Cancel | Help

## Figure 47 Selecting a template

## Table 12 Backup templates

| | |
|---|---|
| **Blank SAP Backup** | No predefined options. |
| **Brbackup_SMB_Offline** | Used for an offline ZDB (split mirror or snapshot backup). The database is stopped during the creation of a replica. |
| **Brbackup_SMB_Online** | Used for an online ZDB (split mirror or snapshot backup). The database is active during the creation of a replica. |
| **Brbackup_SPLITINT_Offline** | Used for an offline ZDB (split mirror or snapshot backup) using `splitint`. The database is stopped during the creation of a replica. The database is |

| | |
|---|---|
| | offline for a shorter period of time than if SMB_offline was used, but `splitint` must be supported by BRTOOLS. |
| **Brbackup_SPLITINT_Online** | Used for an online ZDB (split mirror or snapshot backup) using `splitint`. The database is active during the creation of a replica. The database is in the backup mode for a shorter period of time than if SMB_online was used, but `splitint` must be supported by BRTOOLS. |

In **Backup type**, select:

***EMC and XP:*** **Split mirror backup**

***VA and EVA:*** **Snapshot backup**

In **Sub type**, select the appropriate disk array agent.

4. In **Application system**, select the SAP R/3 client to be backed up. In cluster environments, select the virtual server.

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

   Specify other disk array specific backup options (see Figure 48 on page 196 for EMC, Figure 49 on page 197 for XP, Figure 50 on page 198 for VA, or Figure 51 on page 199 for EVA for EVA). For details on the options, press **F1**.

NOTE:

***XP, VA, and EVA only:*** To enable instant recovery, leave **Track the replica for instant recovery** selected.

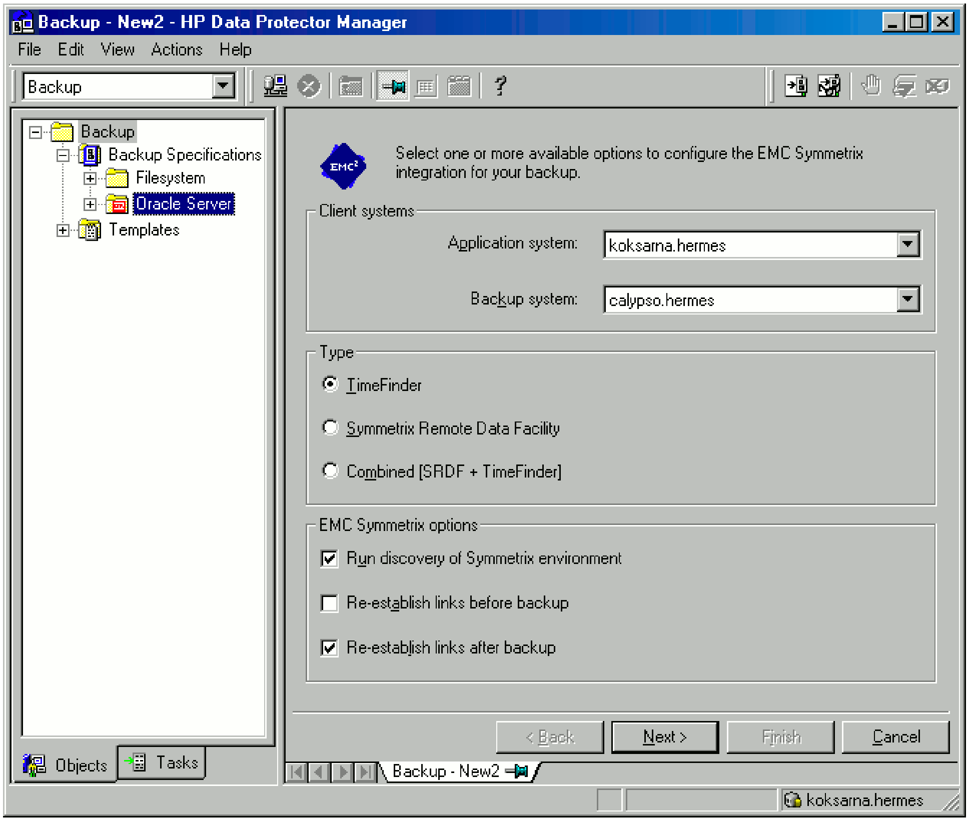

Select one or more available options to configure the EMC Symmetrix integration for your backup.

Client systems

Application system: koksarna.hermes

Backup system: calypso.hermes

Type

● TimeFinder

○ Symmetrix Remote Data Facility

○ Combined [SRDF + TimeFinder]

EMC Symmetrix options

☑ Run discovery of Symmetrix environment

□ Re-establish links before backup

☑ Re-establish links after backup

## Figure 48 EMC backup options

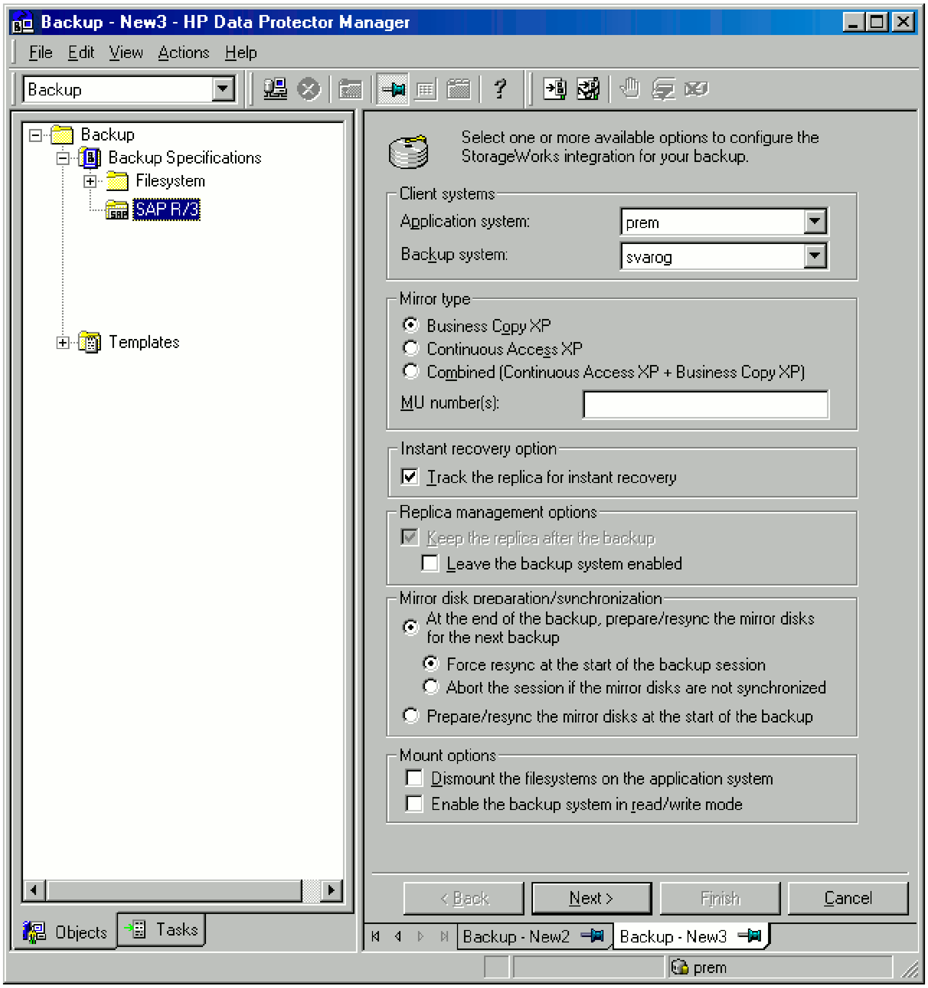

Select one or more available options to configure the StorageWorks integration for your backup.

Client systems

Application system: prem

Backup system: svarog

Mirror type

● Business Copy XP

○ Continuous Access XP

○ Combined (Continuous Access XP + Business Copy XP)

MU number(s):

Instant recovery option

☑ Track the replica for instant recovery

Replica management options

☑ Keep the replica after the backup

□ Leave the backup system enabled

Mirror disk preparation/synchronization

● At the end of the backup, prepare/resync the mirror disks for the next backup

● Force resync at the start of the backup session

○ Abort the session if the mirror disks are not synchronized

○ Prepare/resync the mirror disks at the start of the backup

Mount options

□ Dismount the filesystems on the application system

□ Enable the backup system in read/write mode

Figure 49 XP backup options

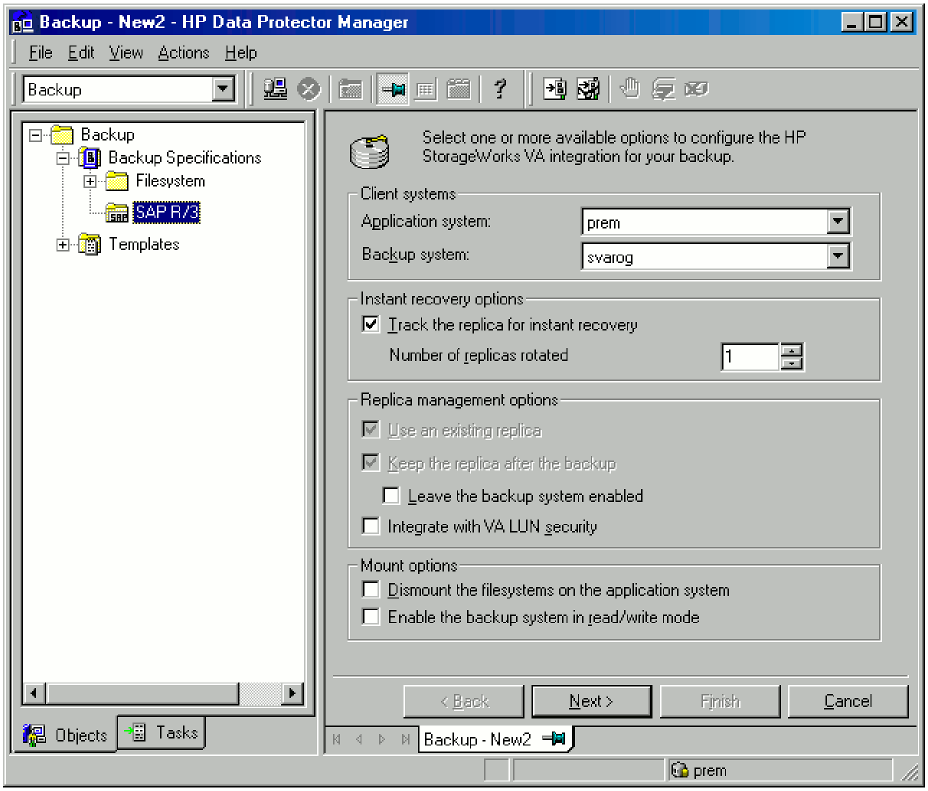

Select one or more available options to configure the HP StorageWorks VA integration for your backup.

Client systems

Application system: prem

Backup system: svarog

Instant recovery options

☑ Track the replica for instant recovery

Number of replicas rotated: 1

Replica management options

☑ Use an existing replica

☑ Keep the replica after the backup

☐ Leave the backup system enabled

☐ Integrate with VA LUN security

Mount options

☐ Dismount the filesystems on the application system

☐ Enable the backup system in read/write mode

Figure 50 VA backup options

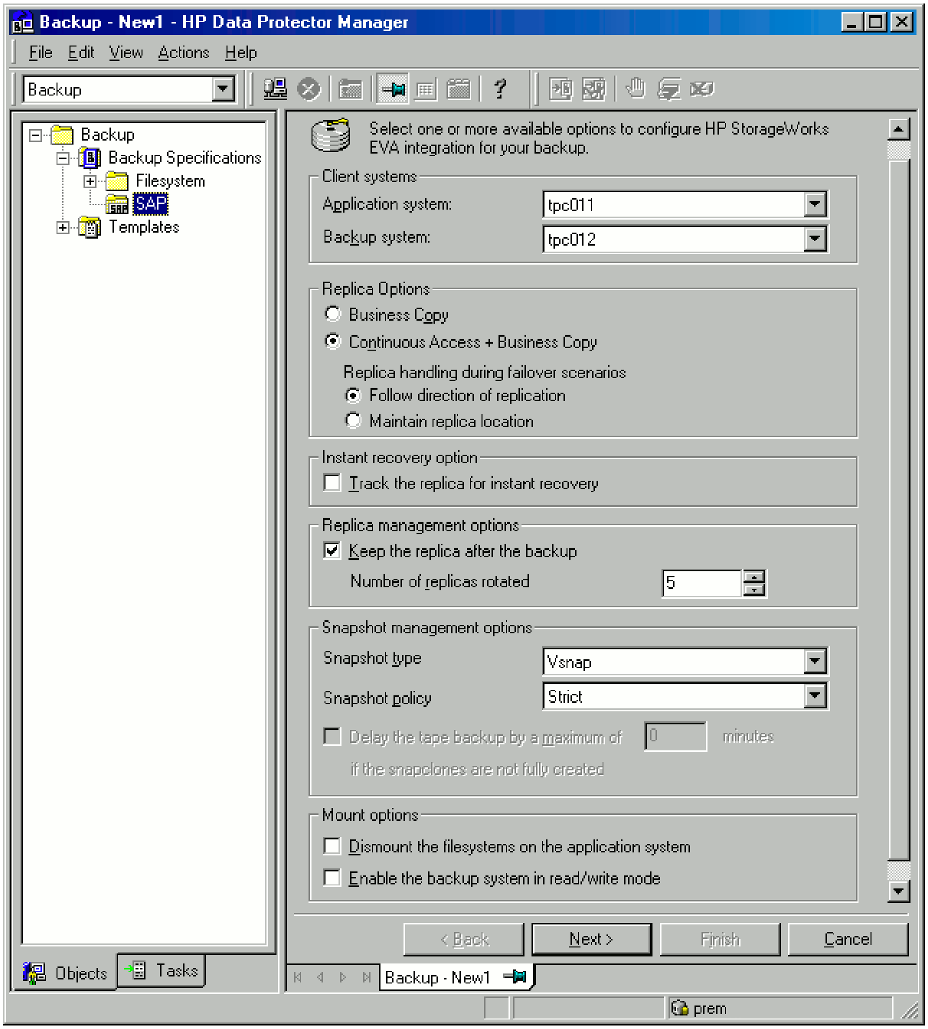

Select one or more available options to configure HP StorageWorks EVA integration for your backup.

Client systems

Application system: tpc011

Backup system: tpc012

Replica Options

○ Business Copy

● Continuous Access + Business Copy

Replica handling during failover scenarios

● Follow direction of replication

○ Maintain replica location

Instant recovery option

□ Track the replica for instant recovery

Replica management options

☑ Keep the replica after the backup

Number of replicas rotated 5

Snapshot management options

Snapshot type Vsnap

Snapshot policy Strict

□ Delay the tape backup by a maximum of 0 minutes if the snapclones are not fully created

Mount options

□ Dismount the filesystems on the application system

□ Enable the backup system in read/write mode

## Figure 51 EVA backup options

Click **Next**.

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

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

   - ***UNIX:*** In **Username**, type the Oracle OS user described in "Configuring user accounts" on page 175. In **Group/Domain name**, type `dba`.
   - ***Windows Server 2008:***

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

   Ensure that this user has been added to the Data Protector `admin` or `operator` user group and has the SAP R/3 backup rights. This user becomes the backup owner.

   Click **Next**.

6. If the SAP R/3 database is not configured yet for use with Data Protector, the **Configure SAP** dialog box is displayed. Configure it as described in "Configuring SAP R/3 databases" on page 182.

7. Select SAP R/3 objects to be backed up. You can select individual tablespaces, data files, or archived logs.

NOTE:

If you plan to do instant recovery, select the whole **DATABASE** item.

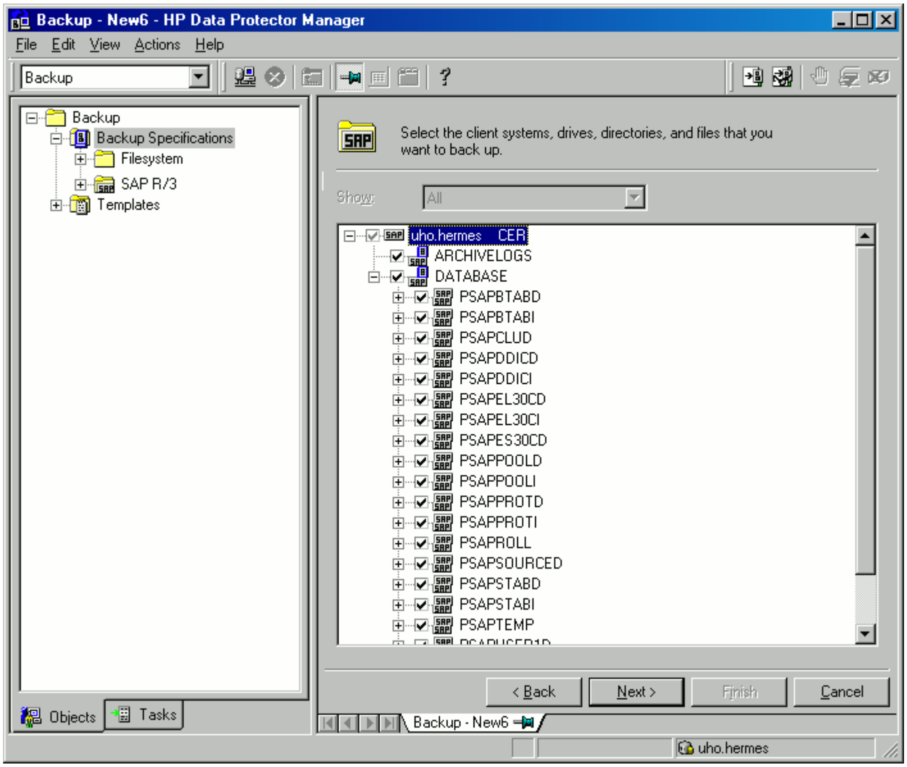

Select the client systems, drives, directories, and files that you want to back up.

Show: All

**Figure 52 Selecting backup objects**

Click **Next**.

8. Select devices to use for the backup.

   To specify device options, right-click the device and click **Properties**. Specify the number of parallel backup streams in the **Concurrency** tab and the media pool.

   NOTE:

   Parallelism (the number of streams your SAP R/3 database is backed up with) is set automatically. If load balancing is used, the number of streams equals the sum of concurrencies of the selected devices.

   Click **Next**.

9. Set backup options. For information on the application specific options, see Table 13 on page 204.

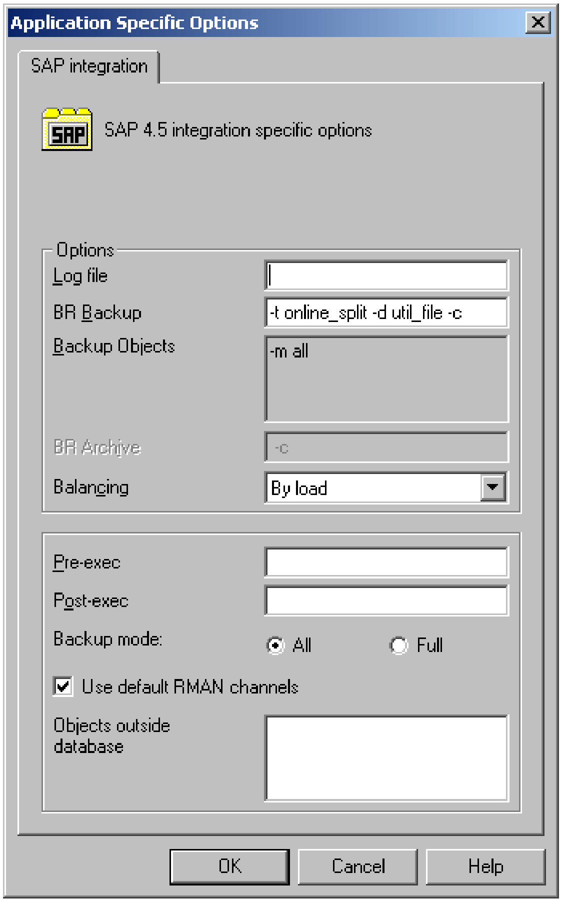
Application Specific Options
SAP integration
SAP 4.5 integration specific options
OK
Cancel
Help

Options

Log file:

BR Backup: -t online_split -d util_file -c

Backup Objects: -m all

BR Archive: -c

Balancing: By load

Pre-exec:

Post-exec:

Backup mode: ● All ○ Full

☑ Use default RMAN channels

Objects outside database:

Figure 53 Application specific options

Click **Next**.

10. Optionally, schedule the backup. See “Scheduling backup specifications” on page 206.

    Click **Next**.

11. Save the backup specification, specifying a name and a backup specification group.

TIP:

Preview your backup specification before using it for real. See “Previewing backup sessions” on page 207.

### 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.<br>For example, for online backup using the `splitint` interface, type `-t online_mirror` . If `splitint` is not supported by BRTOOLS, type `-t online_split`.<br>To run BRBACKUP under a different Oracle database user than the one specified during the configuration, type `-u` *`user_name`*. |
| **Backup Objects** | Lists BRBACKUP options passed by `omnisap.exe`. The list is displayed after you save the backup specification. |
| **BR Archive** | Specifies BRARCHIVE options.<br>Not applicable for ZDB to disk. |
| **Balancing: By Load** | Groups files into subsets of approximately equal sizes. The subsets are then backed up concurrently by Data Protector `sapback` programs.<br>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 `compression` section of the Data Protector SAP R/3 configuration file. See “Data Protector SAP R/3 configuration file” on page 167. |

| Option | Description |
|---|---|
| **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 `sapback` programs. Data Protector automatically stores backup speed information in the `speed` section of the Data Protector SAP R/3 configuration file. It uses this information to optimize the backup time.<br>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 214.<br>Not applicable for ZDB to disk. |
| **Balancing: None** | No balancing is used. The files are backed up in the same order as they are listed in the internal Oracle database structure. To check the order, use the Oracle Server Manager SQL command:<br>`select * from dba_data_files` |
| **Pre-exec, Post-exec** | The command specified here is started by `omnisap.exe` on the SAP R/3 system before the backup (`pre-exec`) or after it (`post-exec`). Do not use double quotes. Provide only the name. The command must reside in the directory:<br>***Windows:*** *`Data_Protector_home`*`\bin`<br>***HP-UX, Solaris:*** `/opt/omni/bin` |
| **Backup mode** | Not applicable for ZDB. |
| **Use default RMAN channels** | Not applicable for ZDB. |
| **Objects outside database** | Specifies non-database files of the Oracle SAP R/3 environment to be saved.<br>Save these files in a separate backup session. |

NOTE:

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

## Modifying backup specifications

To modify your backup specification, click its name in the Scoping Pane of the Backup context, then click the appropriate tab, and apply the changes.

## Scheduling backup specifications

You can run unattended backups at specific times or periodically. For details on scheduling, see the online Help index: “scheduled backups”.

### Scheduling example

To schedule ZDB to disk+tape at 8:00, 13:00, and 18:00 during week days:

1. In the **Schedule** property page, select the starting date in the calendar and click **Add** to open the **Schedule Backup** dialog box.
2. Under **Recurring**, select **Weekly**. Under **Time options**, select **8:00**. Under **Recurring Options**, select **Mon**, **Tue**, **Wed**, **Thu**, and **Fri**.

   In **Split mirror/snapshot backup** drop-down list, select **To disk+tape**. See Figure 54 on page 207.

   Click **OK**.
3. Repeat Step 1 on page 206 and Step 2 on page 206 to schedule backups at 13:00 and 18:00.
4. Click **Apply** to save the changes.

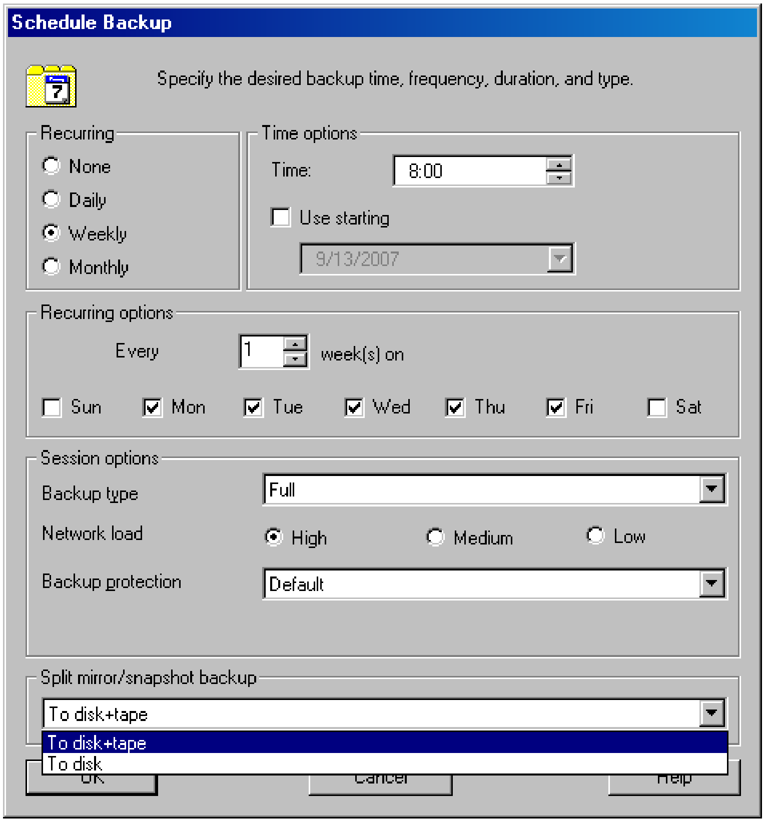

Schedule Backup

Specify the desired backup time, frequency, duration, and type.

Recurring

○ None
○ Daily
● Weekly
○ Monthly

Time options

Time: 8:00

□ Use starting

9/13/2007

Recurring options

Every 1 week(s) on

□ Sun ☑ Mon ☑ Tue ☑ Wed ☑ Thu ☑ Fri □ Sat

Session options

Backup type: Full

Network load: ● High ○ Medium ○ Low

Backup protection: Default

Split mirror/snapshot backup

To disk+tape

To disk+tape
To disk

**Figure 54 Scheduling backups**

## Previewing backup sessions

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

### Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **SAP R/3**. Right-click the backup specification you want to preview and click **Preview Backup**.
3. Specify **Backup type** and **Network load**. Click **OK**.

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

### Using the Data Protector CLI

From the directory:

***Windows:*** `Data_Protector_home\bin`

***HP-UX, Solaris:*** `/opt/omni/bin/`

run:

```
omnib -sap_list backup_specification_name -test_bar
```

### What happens during the preview?

The `omnisap.exe` command is started, which starts the Data Protector `testbar` command to test the following:

- The syntax of the backup specification
- If devices are correctly specified
- If the necessary media are in the devices

## Starting backup sessions

Interactive backups are run on demand. They are useful for urgent backups or restarting failed backups.

### Backup methods

Start a backup of SAP R/3 objects in any of the following ways:

- Using the Data Protector GUI.
- Using the Data Protector CLI.
- Using the SAP BR*Tools.

### Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, and then **SAP R/3**. Right-click the backup specification you want to start and click **Start Backup**.

3. Specify **Network load**. Click **OK**.

   NOTE:

   Only the `Full` backup type is supported.

   ***ZDB to disk, ZDB to disk+tape only:*** Specify the **Split mirror/snapshot backup** option.

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

## Using the Data Protector CLI

From the directory:

***Windows:*** `Data_Protector_home\bin`

***HP-UX, Solaris:*** `/opt/omni/bin/`

run:

***ZDB to tape, ZDB to disk+tape:***

```
omnib -sap_list backup_specification_name
```

***ZDB to disk:***

```
omnib -sap_list backup_specification_name -disk_only
```

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

## Using the SAP BRTOOLS

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

2. Export/set the following environmental variables:

   `ORACLE_SID=`*`SAP_instance_name`*

   `ORACLE_HOME=`*`Oracle_software_home_directory`*

   `[SAPBACKUP_TYPE=OFFLINE]`

      Default is `ONLINE`.

   `SAPDATA_HOME=`*`database_files_directory`*

   `SAPBACKUP=`*`BRTOOLS_logs_and_control_file_copy_directory`*

   `SAPREORG=`*`BRSPACE_logs_directory`*

   `OB2BARLIST=`*`backup_specification_name`*

      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.

   `[OB2BARHOSTNAME=`*`application_system_name`*`]`

      Required if you are logged in to the backup system. Optional if you want to specify a virtual server name in cluster environments.

   `[OB2BACKUPHOSTNAME=`*`backup_system_name`*`]`

      Required if you are logged in to the application system.

   `OB2SMB=1`

      Specifies a ZDB session.

   `[OB2SMBIR=1]`

      Specifies tracking replica for instant recovery.

   `[ZDB_ORA_INCLUDE_CF_OLF=1]`

      Required if you are logged in to the backup system. For details, see "ZDB integrations omnirc variables" on page 425.

   `[OB2DISKONLY=1]`

      Specifies a ZDB-to-disk session.

   If you are logged on to the backup system, ensure that the `NLS_LANG` environmental variable is set to the same value as the `NLS_LANG` environmental variable on the application system.

   Alternatively, these variables can be specified in the backint parameter file. If this is required, the location of the file must be specified in the SAP configuration file using the *`util_par_file`* parameter:

   ```
   util_par_file = path\filename
   ```

If you do not supply the path, the system searches for the parameter file in the directory:

***UNIX:*** *ORACLE_HOME*/dbs

***Windows:*** *SAPDATA_HOME*\database

3. Run the BRBACKUP command. The command syntax depends on whether you are logged in to the application system or backup system:

   Application system:

   ```
   brbackup -t {online_split | offline_split | online_mirror
   | offline_mirror} [-q split] -d util_file -m all -c -u
   user/password
   ```

   Backup system:

   ```
   brbackup -t {online_mirror | offline_mirror} [-q split]
   -d util_file -m all -c -u user/password
   ```

   The `-q split` option is required if `OB2DISKONLY` is set to `1`.

## Configuring SAP compliant ZDB sessions

SAP R/3 standards recommend that, in ZDB sessions that use the `splitint` backup interface, BRBACKUP is started on the backup system and not on the application system. You can configure Data Protector to comply with these standards by setting the Data Protector `OB2_MIRROR_COMP` environmental variable to `1`. The variable is saved in the Data Protector SAP R/3 instance configuration file. Consequently, in all `splitint` ZDB sessions for this SAP R/3 instance, BRBACKUP will be started on the backup system. By default, BRBACKUP is started on the application system.

Set the `OB2_MIRROR_COMP` environmental variable using the Data Protector GUI or CLI.

NOTE:

If no backup specification for the related SAP R/3 instance exists, you cannot use the Data Protector CLI to set the `OB2_MIRROR_COMP` variable.

### Using the Data Protector GUI

You can set the `OB2_MIRROR_COMP` variable when you create a backup specification or modify an existing one:

1. Proceed to the Source page of the backup specification.

   NOTE:

   In environments in which the control file and datafiles reside on the same source disk, Data Protector does not let you proceed to the Source page if the **Track the replica for instant recovery** option is selected. Specifically, the Data Protector instant recovery check fails. In such a case, clear the option first. You can select the option later if needed, when the OB2_MIRROR_COMP variable is already set and, consequently, the instant recovery check is no longer performed.

   Right click the SAP R/3 instance at the top and click **Set Environmental Variables**.

2. In the Advanced dialog box, set OB2_MIRROR_COMP to 1. See Figure 55 on page 213.

   Click **OK**.

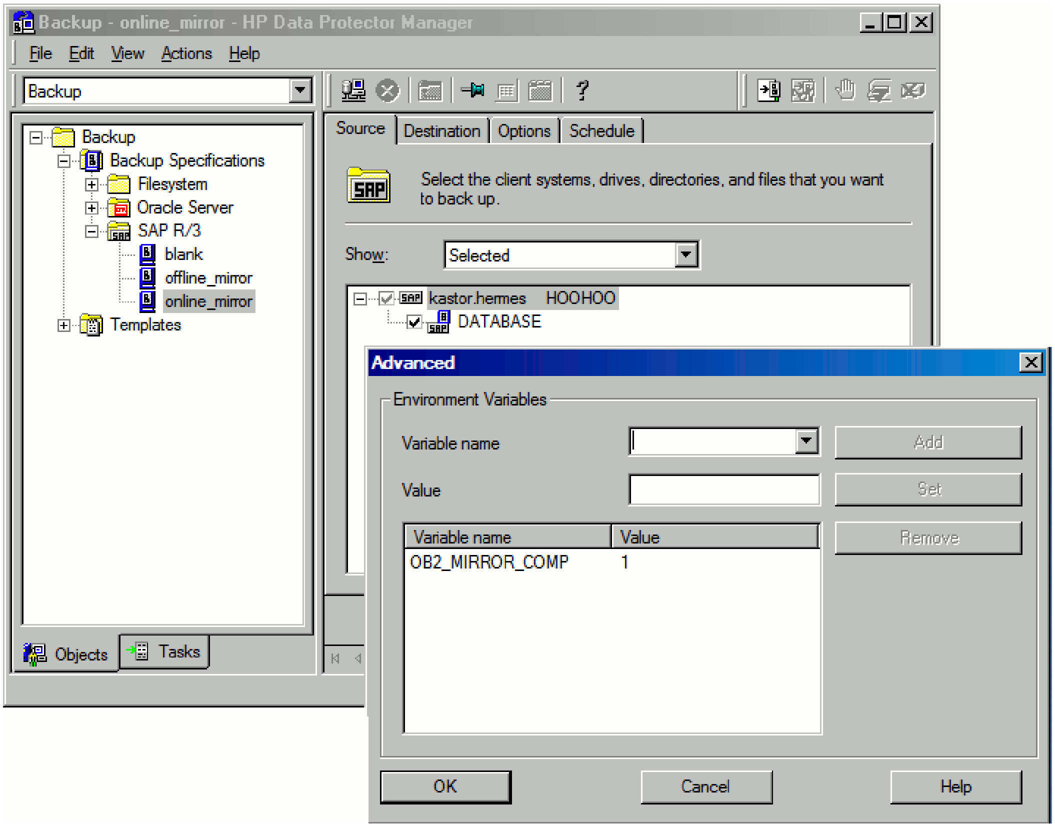

Select the client systems, drives, directories, and files that you want to back up.

Show: Selected

| Variable name | Value |
|---|---|
| OB2_MIRROR_COMP | 1 |

Figure 55 Setting environmental variables

## Using the Data Protector CLI

From the directory:

***Windows:*** `Data_Protector_home\bin`

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

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

run:

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

## Manual balancing

Manual balancing means that you manually group files into subsets, which are then backed up in parallel. To group files into subsets, add the `manual_balance` section to the Data Protector SAP R/3 configuration file as described in the following example.

### Example

Suppose that we have a backup specification named `SAP-R3` with the following files to be backed up: `fileA`, `fileB`, `fileC`, `fileD` . To group the files into three subsets (`0={fileA, fileC}`, `1={fileB}`, `2={fileD}`), add the following lines to the Data Protector SAP R/3 configuration file:

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

When you group files into subsets, consider the following:

- Use only one file from the same hard disk at a time.
- The number of files in a subset must be equal to or smaller than the sum of the concurrencies of all devices specified for backup.
- If the backup specification contains files that are not allocated to any subset, Data Protector automatically adds these files to the list of files to be backed up using the load balancing principle. Before the backup, this list is logged in:

  ***UNIX:*** `*ORACLE_HOME*/sapbackup/.*.lst`

  ***Windows:*** `*SAPDATA_HOME*\sapbackup\*.lst`

## Restore

You can restore SAP R/3 objects using any of the following methods:

- **Standard restore:** Data is restored from backup media created in ZDB to tape, ZDB to disk+tape, and non-ZDB (standard backup) sessions. See "Standard restore" on page 216.
- **Instant recovery:** Data is restored from a replica created in *online* ZDB–to–disk or ZDB–to–disk+tape sessions. See "Instant recovery" on page 221.

After the restore, you can recover the database to a specific point in time using the SAP BRTOOLS interface. Instant recovery method enables you to restore and recover the database within the same session. However, you can only restore (and recover) the whole database. To restore only a part of the database or the archived logs, use the standard restore method.

Table 14 on page 215 shows which restore methods are available, depending on the backup session you restore from.

### Table 14 SAP recovery methods

| Disk arrays | Backup type | Recovery of the whole database | | Recovery of a part of the database |
|---|---|---|---|---|
| | | Until now | To a point in time, logseq/thread or SCN number | |
| XP, VA, EVA, EMC | **ZDB to tape - online** | Restore | Restore | Restore |
| | **ZDB to tape - offline** | Restore | Restore | Restore |
| XP, VA, EVA | **ZDB to disk - online** | Instant recovery + database recovery | Instant recovery + database recovery | N/A |
| | **ZDB to disk - offline** | N/A | N/A | N/A |
| | **ZDB to disk+tape - online** | • Instant recovery + database recovery<br>or<br>• Restore | • Instant recovery + database recovery<br>or<br>• Restore | Restore |
| | **ZDB to disk+tape - offline** | Restore | Restore | Restore |

### Table 15 Legend

| | |
|---|---|
| Restore | You can do a standard restore from the Data Protector media using the Data Protector GUI or the SAP BRTOOLS. After the restore, you can recover the database using the SAP BRTOOLS. |

| | |
|---|---|
| Instant recovery + database recovery | You can do an instant recovery. You can include database recovery in the instant recovery session or do it afterwards, using the SAP BRTOOLS. |
| N/A | Not available. |

## Considerations

- SAP R/3 tablespaces located on raw partitions cannot be restored using the Data Protector GUI. Workaround: Use SAP restore commands (for example, `brrestore`).
- If your Oracle database is localized, you may need to set the appropriate Data Protector encoding before you start a restore. For details, see "Localized SAP R/3 objects" on page 228.
- Restore preview is not supported.

## Standard restore

Restore SAP R/3 objects using the Data Protector Manager.

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand **SAP R/3**, expand the client (backup system) from which the data was backed up, and then click the Oracle instance you want to restore.

3. In the **Source** page, select SAP R/3 files to be restored.

   To restore a file under a different name or to a different directory, right-click the file and click **Restore As/Into**.

   To restore a file from a specific backup session, right-click the file and click **Restore Version**.

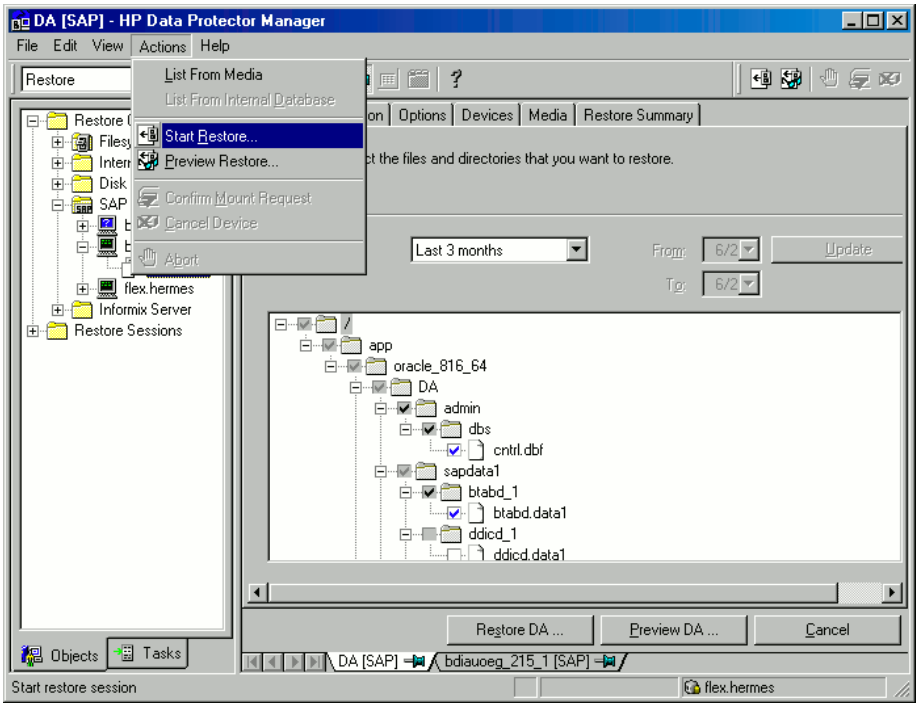

Figure 56 Selecting objects for restore

4. In the **Destination** tab, select the client to restore to (**Target client**). By default, this is the application system. See Figure 57 on page 218.

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

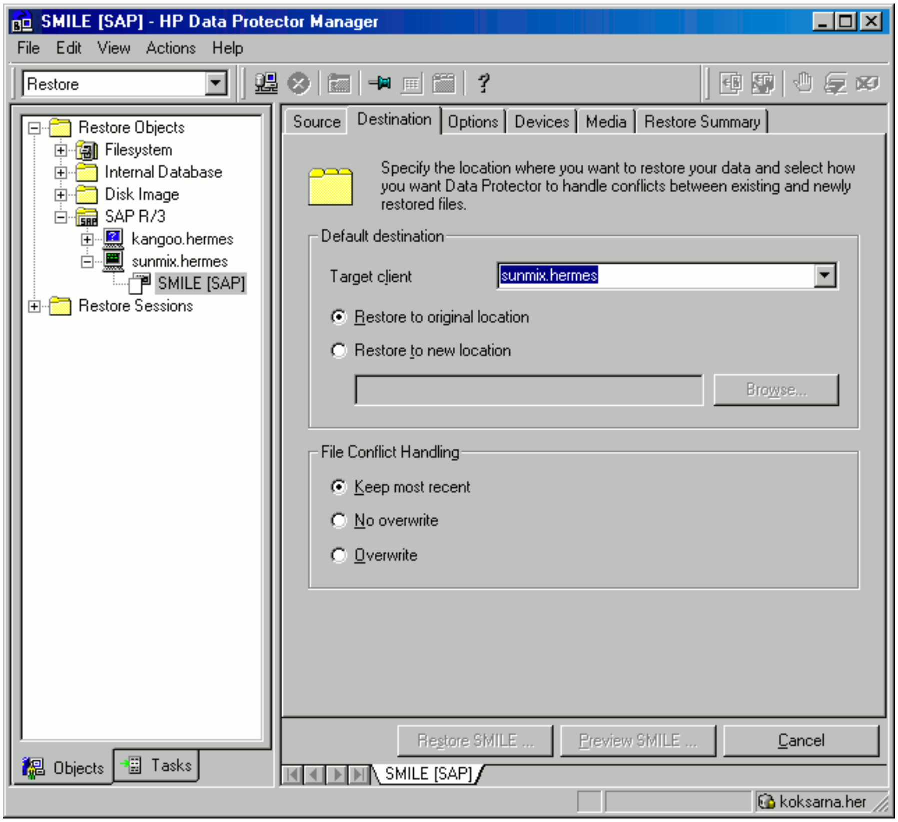

SMILE [SAP] - HP Data Protector Manager

Specify the location where you want to restore your data and select how you want Data Protector to handle conflicts between existing and newly restored files.

Default destination

Target client: sunmix.hermes

● Restore to original location

○ Restore to new location

File Conflict Handling

● Keep most recent

○ No overwrite

○ Overwrite

Figure 57 Selecting the target client

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

10. ***EMC and XP:*** This step is relevant only if you have both the `EMC Symmetrix Agent` and `HP StorageWorks XP Agent` components installed on the application system.

    ***EMC:*** Select `EMC Symmetrix restore`.

    ***XP:*** Select `StorageWorks XP restore`. See Figure 58 on page 219.

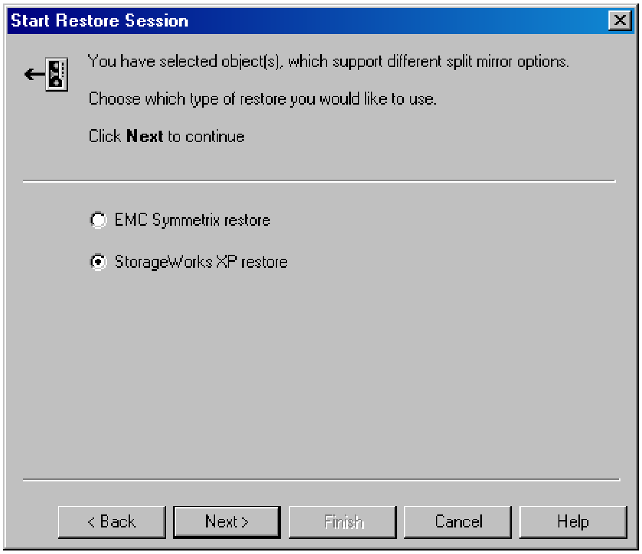

Start Restore Session

You have selected object(s), which support different split mirror options.

Choose which type of restore you would like to use.

Click **Next** to continue

○ EMC Symmetrix restore

● StorageWorks XP restore

**Figure 58 Selecting the XP restore**

Click **Next**.

11. ***EMC and XP:*** In the `EMC Symmetrix mode` or `Mirror mode` drop-down list, select **Disabled**. This sets the restore from backup media to the application system directly. See Figure 59 on page 220.

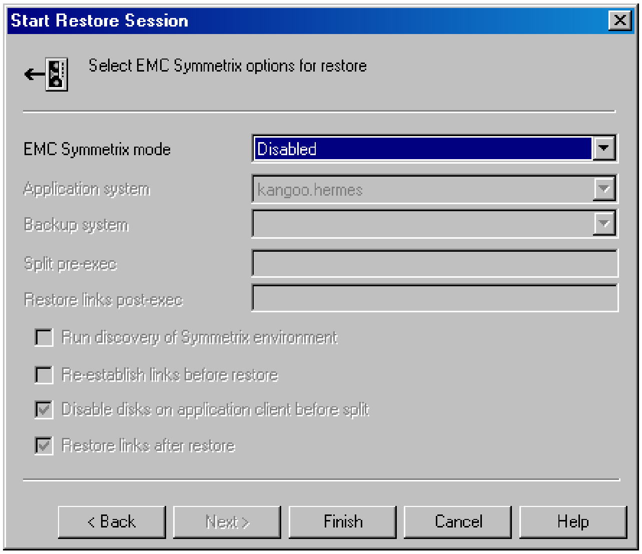

Start Restore Session

Select EMC Symmetrix options for restore

EMC Symmetrix mode: Disabled

Application system: kangoo.hermes

Backup system:

Split pre-exec:

Restore links post-exec:

☐ Run discovery of Symmetrix environment

☐ Re-establish links before restore

☑ Disable disks on application client before split

☑ Restore links after restore

**Figure 59 EMC - Selecting restore to the application system directly**

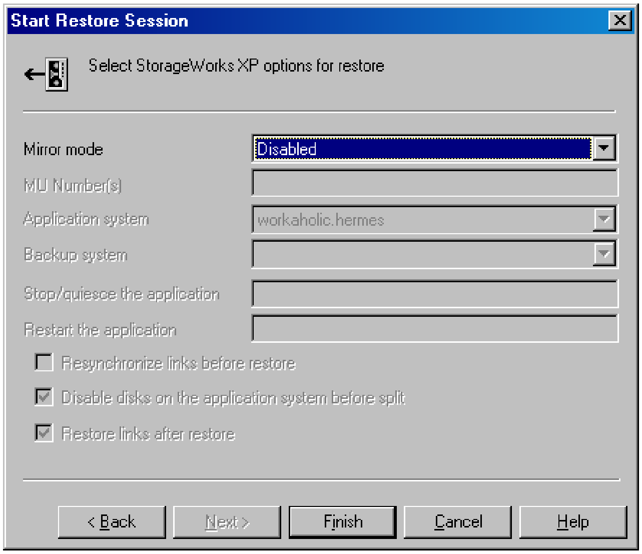

Start Restore Session

Select StorageWorks XP options for restore

Mirror mode: Disabled

MU Number(s):

Application system: workaholic.hermes

Backup system:

Stop/quiesce the application:

Restart the application:

☐ Resynchronize links before restore

☑ Disable disks on the application system before split

☑ Restore links after restore

**Figure 60 XP - Selecting restore to the application system directly**

12. Click **Finish** to start the restore.

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

## Instant recovery

For general information on instant recovery, see *HP Data Protector zero downtime backup concepts guide* and *HP Data Protector zero downtime backup administrator's guide*. For information on instant recovery in cluster environment (Cluster File System (CFS), MC/ServiceGuard, and Microsoft Cluster Server), see *HP Data Protector zero downtime backup administrator's guide*.

You can start an instant recovery using the Data Protector GUI or CLI.

### Considerations

- The database recovery part is performed after the instant recovery procedure. During database recovery, archive log backups performed after the ZDB are restored from tape by the SAP BR*Tools utilities. If selected, the logs are reset and the database is opened.
- If the replica to be used for instant recovery contains the control file, first see "Instant recovery from replicas containing the control file" on page 226.

### Instant recovery using the Data Protector GUI

To perform an instant recovery:

1. Shut down the Oracle database using `sqlplus`:

   For example:

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

2. In the Data Protector Manager Context List, select **Instant Recovery**.
3. Expand **SAP R/3** and select the ZDB-to-disk or ZDB-to-disk+tape session from which you want to perform the restore.

4. In the **Source** tab, select the objects to recover. Only whole databases can be selected.

   For HP StorageWorks Disk Array XP, it is recommended to leave the **Keep the replica after the restore** option selected to enable a restart of an instant recovery session. The option is selected by default, except for an offline backup where the database was in NOARCHIVELOG mode during the backup. For HP StorageWorks Enterprise Virtual Array, replica is kept on the array only if the **Copy replica data to the source location** is selected.

   Set the HP StorageWorks Enterprise Virtual Array or HP StorageWorks Disk Array XP options. For details, press **F1**.

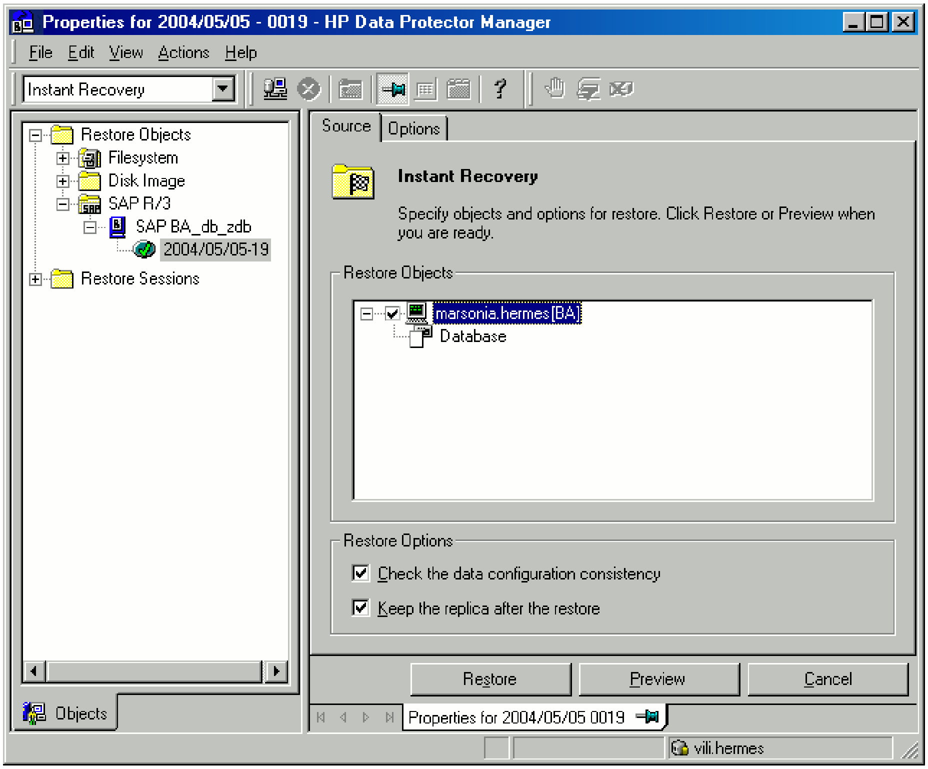

**Instant Recovery**

Specify objects and options for restore. Click Restore or Preview when you are ready.

Restore Objects

☑ marsonia.hermes [BA]
Database

Restore Options

☑ Check the data configuration consistency

☑ Keep the replica after the restore

**Figure 61 SAP R/3 source options**

5. At this point, you can decide whether to perform a database recovery immediately after an instant recovery or not:
   - To perform only an instant recovery, click **Restore**.
   - To automatically perform a database recovery after an instant recovery, select the recovery options. For details refer to "Database recovery options" on page 223.

     Click **Restore**.

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

## Database recovery options

**User name** (UNIX systems only)

Specifies the user name under which the instant recovery is performed. The user needs to be a member of the DBA group.

**User group** (UNIX systems only)

Specifies the user group the user in the **User name** field belongs to.

NOTE:

The **User name** and the **User group** must be the same as defined in the backup ownership.

**Recovery**

Enables database recovery after instant recovery.

**Recover until**

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

The following options are available:

**Now**: All existing archive logs are applied.

**Selected time**: Only archive logs until the specified time are applied.

**Selected logseq/ thread number**: Specifies an incomplete recovery. Only archive logs with a lower or equal number than the specified log sequence or thread number are applied.

**Selected SCN number**: Only archive logs until the specified SCN number are applied.

**Open database after recovery**

Opens the database after the recovery was performed.

**Reset logs**

Resets the archive logs after the database is opened. This option is by default not selected if the **Recover until** option is set to **Now**.

The following are Oracle recommendations on when to reset the logs:

*Always* reset the logs:

- After an incomplete recovery, that is, if not all archive redo logs are applied.
- If a backup of the control file is used for recovery.

*Do not* reset the logs:

- After a complete recovery, when the control file is not used.
- If the archive logs are used for a standby database. However, if you must reset the archive logs, recreate the standby database.

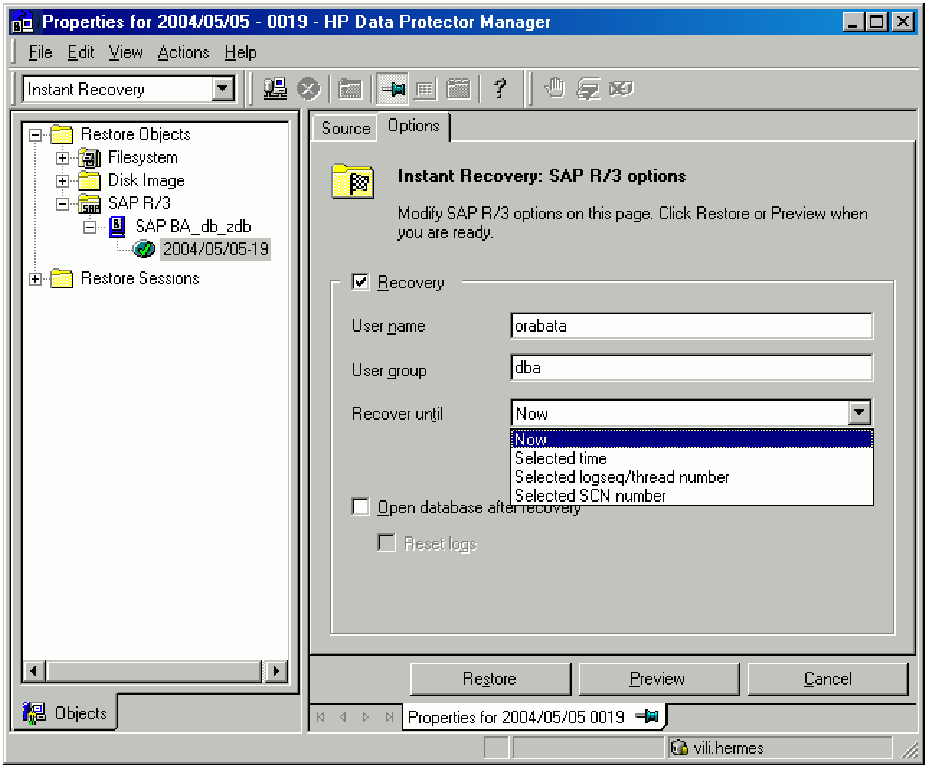

Instant Recovery: SAP R/3 options

Modify SAP R/3 options on this page. Click Restore or Preview when you are ready.

☑ Recovery

User name: orabata

User group: dba

Recover until: Now

- Now
- Selected time
- Selected logseq/thread number
- Selected SCN number

□ Open database after recovery

□ Reset logs

Figure 62 SAP R/3 recovery options

## Instant recovery using the Data Protector CLI

From the directory:

***Windows***: *Data_Protector_home*\bin

***HP-UX***: /opt/omni/bin/

run:

```
omnir -sap
```

-host *ClientName*

-session *SessionID*

{*DISK_ARRAY_XP_OPTIONS* | *ENTERPRISE_VIRTUAL_ARRAY_OPTIONS*}

[-appname *DB_Name*

```
-user UserName -group GroupName

[-recover {now | time MM/DD/YY hh:mm:ss | logseq LogSeqNumber
thread ThreadNumber | SCN Number} [-open [-resetlogs]]]

]

[-preview]
```

The order of options is important. On Windows clients, the user name and group name options are not required. For a detailed description of the options, see the *HP Data Protector command line interface reference*.

## Instant recovery from replicas containing the control file

During an instant recovery from a replica that contains the control file, the current control file and, possibly, online redo logs get overwritten. Therefore, before you start the session, copy the current control file and online redo logs to a safe location to be able to do a database recovery afterwards.

A replica contains the control file if it is created in any of the following sessions:

- *Online* ZDB session with the `omnirc` variable `ZDB_ORA_INCLUDE_CF_OLF` set to `1`
- *Online* SAP compliant ZDB session
- *Offline* ZDB session (any configuration)

NOTE:

An *offline* ZDB session also contains online redo logs. You can use such a session to restore the SAP R/3 database to a point in time at which the backup was performed. In this case, you do not need to follow the steps below.

To restore and recover the database:

1. Copy the current control files and online redo logs to a safe location.
2. Perform instant recovery (without database recovery). Use the Data Protector GUI or CLI.
3. Copy the current control files and online redo logs back to their original location.
4. Mount the target database.

5. Restore missing archived redo logs required for database recovery.

   Example:

   ```
   # sqlplus user/password@net_service_name
   SQL> select SEQUENCE#, NAME from V$ARCHIVED_LOG  where
   (NEXT_TIME>to_date('2007/10/03') and (FIRST_CHANGE#<='until_SCN');
   # brrestore -a log_no,... -d util_file -c force -u user/password
   ```

6. Recover the target database.

   Example:

   ```
   # rman target user/password@net_service_name
   RMAN> run{
   2> allocate channel dbrec type disk;
   3> recover database until scn until_SCN;
   4> release channel dbrec;
   5> }
   ```

## Restoring using another device

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

### Using the Data Protector GUI

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

### Using the Data Protector CLI or SAP commands

If you are restoring using the Data Protector CLI or SAP R/3 commands, specify the new device in the file:

***Windows:*** `Data_Protector_home\Config\Server\cell\restoredev`

***UNIX:*** `/etc/opt/omni/server/cell/restoredev`

Use the format:

`"DEV 1" "DEV 2"`

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

IMPORTANT:

Delete this file after use.

On Windows, use the Unicode format for the file.

## Localized SAP R/3 objects

Oracle Server uses its own encoding, which may differ from the encoding used by the filesystem. In the Backup context, Data Protector displays the logical structure of the Oracle database (with Oracle names) and in the Restore context, the filesystem structure of the Oracle database. Therefore, to display non-ASCII characters correctly, ensure that the Data Protector encoding matches with the Oracle Server encoding during backup and with the filesystem encoding during restore. However, the incorrect display does not impact the restore.

***UNIX:*** To be able to switch between the Data Protector encodings, start the GUI in UTF-8 locale.

***Windows:*** If the current values of DBCS and the default Windows character set for non-Unicode programs do not match, problems arise. See "ZDB, restore, or instant recovery sessions fail due to invalid characters in filenames" on page 240.

# Monitoring sessions

You can monitor currently running sessions in the Data Protector GUI. When you run an interactive backup or a restore session, a monitor window shows you the progress of the session. Closing the GUI does not affect the session.

You can also monitor sessions from any Data Protector client with the `User Interface` component installed, using the Monitor context.

On how to monitor a session, see the online Help index: "viewing currently running sessions".

System messages generated during backups are sent to both the SAP R/3 and the Data Protector monitor. However, mount requests are sent only to the Data Protector monitor.

# Troubleshooting

This section lists general checks and verifications plus problems you might encounter when using the Data Protector SAP R/3 integration.

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

For general ZDB, restore, and instant recovery related troubleshooting, see the troubleshooting sections in the *HP Data Protector zero downtime backup administrator's guide*.

See also the troubleshooting section in the SAP R/3 chapter of the *HP Data Protector integration guide*.

## Before you begin

- Ensure that the latest official Data Protector patches are installed. See the online Help index: "patches" on how to verify this.
- See the *HP Data Protector product announcements, software notes, and references* for general Data Protector limitations, as well as recognized issues and workarounds.
- For an up-to-date list of supported versions, platforms, and other information, see the latest support matrices at http://www.hp.com/support/manuals.

## General troubleshooting

### Problem

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

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

```
*RETVAL*8422
```

### Action

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

### 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 175.

## Verifying the prerequisites (Oracle side)

Perform the following verification steps, in numerical order, to verify that Oracle is installed properly:

1. On the application system, verify that the target database is online, as follows:

   ***UNIX***:

   ```
   export ORACLE_SID=Oracle_SID
   export ORACLE_HOME=Oracle_home_path
   $ORACLE_HOME/bin/sqlplus
   ```

   ***Windows***:

   ```
   set ORACLE_SID=Oracle_SID
   set ORACLE_HOME=Oracle_home_path
   %ORACLE_HOME%\bin\sqlplus
   ```

   At the SQLPlus prompt, type:

   ```
   connect user/passwd@service
   select * from dba_tablespaces
   exit;
   ```

   Try starting the target database.

2. In order to establish the TNS network connection, verify that Net8 software is configured correctly for the target database, as follows:

   - On the application system, perform the following:

     ***UNIX***:

     ```
     $ORACLE_HOME/bin/lsnrctl status service
     ```

     ***Windows***:

     ```
     %ORACLE_HOME%\bin\lsnrctl status service
     ```

     If it fails, either start the TNS listener process or refer to the Oracle documentation on how to create the TNS configuration file (LISTENER.ORA).

   - On the application system, perform the following. Use `sqlplus`:

     ***UNIX***:

     ```
     export ORACLE_SID=Oracle_SID
     export ORACLE_HOME=Oracle_home_path
     $ORACLE_HOME/bin/sqlplus
     ```

     ***Windows***:

     ```
     set ORACLE_SID=Oracle_SID
     set ORACLE_HOME=Oracle_home_path
     %ORACLE_HOME%\bin\sqlplus
     ```

     At the SQLPlus prompt, type:

     ```
     connect user/passwd@service;
     exit;
     ```

     If it fails, refer to the Oracle documentation on how to create the TNS configuration file (TSNAMES.ORA).

## Verifying the prerequisites (SAP side)

Before you begin the steps in this section, be sure you have completed all the steps in “Verifying the prerequisites (Oracle side)” on page 230.

Perform the following verification steps, in numerical order, to verify that SAP is installed properly:

1. On the application system, verify a backup directly to disk, as follows:

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

   If it fails, refer to the SAP Online help for instructions on how to execute a backup to disk using the SAP backup utility.

2. On the application system, verify a restore from the disk, as follows:

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

   If it fails, refer to the SAP Online help for instructions on how to execute a restore to disk using the SAP restore utility.

3. On the application system, verify that SAP is configured properly, as follows:

   Move the original `backint`. Create a test script with the name `backint` in the directory with the SAP backup utility, with the following entries:

   ```
   #!/usr/bin/sh
   echo "Test backint called as follows:"
   echo "$0 $*"
   echo "exiting 3 for a failure"
   exit 3
   ```

   Export all environment variables required by the SAP (SAPDATA_HOME, SAPBACKUP...) and then start the command with the backup owner user:

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

   or, if Data Protector uses `splitint`:

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

   If you receive arguments from `backint`, that means SAP is properly configured for backup using `backint`. Otherwise, you should reconfigure SAP.

## Verifying the configuration

Before you begin this section, be sure that you completed all the steps provided in the sections "Verifying the prerequisites (Oracle side)" on page 230 and "Verifying the prerequisites (SAP side)" on page 231.

Perform the following verification steps, in numerical order, to verify that Data Protector is configured properly:

1. On the application system, verify a Data Protector filesystem backup of the SAP Database Server:

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

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

   If it fails, refer to the *HP Data Protector troubleshooting guide* for help with troubleshooting a filesystem backup.

2. Verify the environment variable on the application system:

   If you have to export some variables before starting the SAP backup utilities, Oracle Server Manager, or the TNS listener, set these environment variables using the Data Protector GUI.

3. Verify the permissions of the SAP user on application system:

   SAP user permissions must be set to enable you to perform an SAP backup or restore with Data Protector. See “Configuring user accounts” on page 175. Use the `testbar2` to check the permissions:

   - Login in as an SAP user
   - Execute `/opt/omni/bin/testbar2 -perform:checkuser`

   If the user account has all the required permissions, you will see only the usual messages displayed on the screen.

4. Examine the system errors:

   System errors are reported in the following file on the Oracle Server:

   ```
   /var/opt/omni/log/debug.log
   ```

## Verifying the backup configuration

Before you begin this section, be sure that you completed all the steps provided in the sections “Verifying the prerequisites (Oracle side)” on page 230 and “Verifying the prerequisites (SAP side)” on page 231.

Perform the following verification steps, in numerical order, to verify that Data Protector is configured properly:

1. Verify the Data Protector SAP ZDB configuration on the application system:

   Execute the following command:

   ***HP-UX***:`/opt/omni/lbin/util_sap -CHKCONF` *`ORACLE_SID`*

   ***Windows***: *`Data_Protector_home`*`\bin\util_sap.exe -CHKCONF` *`ORACLE_SID`*

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

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

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

   which is located on the Cell Manager.

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

2. Verify the SAP user.

   Check that the respective user group has the `See Private Objects` user right selected. Also refer to "Configuring user accounts" on page 175.

3. On the application system, verify the backup using `testbar2`:

   Execute the following to ensure that communication within Data Protector is established:

   - Create a non-ZDB backup specification on the application system.
   - Run:

     ```
     /opt/omni/bin/testbar2 -type:SAP -appname:ORACLE_SID
     -perform:backup -file:file_name -bar barlist_name
     ```

   If it fails, check the errors and try to fix them or call a support representative for assistance.

4. On the application system, verify the backup using backint:

   Execute the following command to ensure that communication within Data Protector is established and that a backup of files can be performed:

   - Create a non-ZDB backup specification on the backup system.

   -
     ```
     export OB2BARLIST=barlist_name
     export OB2APPNAME=ORACLE_SID
     /opt/omni/lbin/backint -f backup -t file -u ORACLE_SID
     -i input_file
     ```

     where *input_file* is the file containing the full pathnames for backup.

   If it fails, check the errors and try to fix them or call a support representative for assistance.

## Verifying restore

Before you begin this section, be sure that you completed all the steps provided in the sections "Verifying the prerequisites (Oracle side)" on page 230 and "Verifying the prerequisites (SAP side)" on page 231.

Perform the following verification steps, in numerical order, to verify that Data Protector is configured properly:

1. Verify the user for the restore

   Verify that the user specified for the restore session is the user of the backup session and that they belong to the Data Protector operator or admin group. Check that the respective user group has the `See private objects` user right selected.

2. Verify that files are backed up and in the Data Protector database:

   - Using the `omnidb` command;

     See the appropriate man page on using the `omnidb` command.

   - Using `backint`;

     SAP tools also use this command to make a query.

     ```
     /opt/omni/lbin/backint -f inquiry -u ORACLE_SID -i
     input_file
     ```

     where *input_file* is what will be queried. `Backint` expects a list of files in the following format:

     ```
     backint_ID_1 pathName_1

     backint_ID_2 pathName_2

     backint_ID_3 pathName_3
     ```

     To retrieve the *backint_ID* numbers, enter the following command:

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

     or, alternatively, you can just specify `#NULL` as *backint_ID_1* in the *input_file*. In this case, the latest backup session for the file is used for the restore.

   If it fails, proceed as follows:

   - Check the backup session - was it successful?
   - Check the user rights. Was the query started under the correct SAP user account?
   - Call a support representative for assistance.

3. Verify the restore using Data Protector or CLI:

   If it fails, proceed as follows:

   - Check the backup session - was it successful?
   - Check that the files are in the Data Protector database.
   - Check the user rights. Was the restore started under the correct SAP user account?
   - Call a support representative for assistance.

4. Verify the restore using `testbar2`:

   Execute the following to ensure that restore is possible:

   ```
   /opt/omni/bin/testbar2 -type:SAP -appname:ORACLE_SID
   -perform:restore -file:file_name -bar barlist_name -object
   objectName
   ```

   If it fails, proceed as follows:

   - Check the backup session - was it successful?
   - Check that the files are in the Data Protector database.
   - Check the user rights. Was the a restore started under the correct SAP user account
   - Call a support representative for assistance.

5. Verify the restore using `backint`:

   `backint` is the same command used by the SAP backup utility.

   ```
   /opt/omni/lbin/backint -f restore -u ORACLE_SID -i
   input_file
   ```

   where *input_file* specifies what will be restored; `backint` expects the list of files in the following format:

   ```
   backint_ID_1 pathName_1 [targetDirectory_1]
   backint_ID_2 pathName_2 [targetDirectory_2]
   backint_ID_3 pathName_3 [targetDirectory_3]
   ```

   To retrieve the *backint_ID* numbers, enter the following command:

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

   or, alternatively, you can just specify `#NULL` as *backint_ID_1* in the *input_file*. In this case, the latest backup session for the file is used for the restore.

   If it fails, proceed as follows:

   - Check the backup session - was it successful?
   - Check that the files are in the Data Protector database.
   - Check the user rights. Was the restore started under the correct SAP user account?
   - Call a support representative for assistance.

## Configuration and backup problems

The following list gives a description of problems and actions to be taken to resolve them:

- **The Server Manager is unable to connect to the destination**

  Check whether the Oracle TNS listener process is up and running. Check whether there are any environment variables required for a successful remote connection to the target database; for example, *TNS_ADMIN* and *SHLIB_PATH*. Set these environment variables using the Data Protector GUI.

  Refer to Data Protector SAP R/3 configuration file for more information on the Data Protector SAP configuration file.

- **Configuration procedure fails**

  Check whether the Oracle Server is up and running.

  Check the login information for the target from the application system using Oracle Server Manager. If you cannot log in, then perform the following actions:

  Check whether `sysoper` and `sysdba` rights are set for the Oracle administrator user.

  Examine system errors reported in:

  - On UNIX:

    `/var/opt/omni/log/debug.log`

    `/var/opt/omni/log/sap.log`

    `/var/opt/omni/log/oracle8.log`

  - On Windows:

    `Data_Protector_home\log\debug.log`,

    `Data_Protector_home\log\sap.log`

    `Data_Protector_home\log\oracle8.log`

  If you have special Oracle environment settings, ensure that they are registered in the `Environment` sublist of the Data Protector SAP configuration file:

  `/etc/opt/omni/server/integ/config/SAP/`*`client_name%ORACLE_SID`* (UNIX Cell Manager), or

  *Data_Protector_home*`\Config\server\integ\config\sap\`*`client_name%ORACLE_SID`* (Windows Cell Manager).

  Refer to "Data Protector SAP R/3 configuration file" on page 167 for more information on the Data Protector SAP configuration file.

- **Starting the backup fails**

On UNIX systems, check the output of the following command on the application system:

```
/opt/omni/lbin/util_sap.exe -CHKCONF ORACLE_SID
```

In case of an error, the error number is displayed in the form:

```
*RETVAL*Error_number
```

To get the error description, start the following command on the application system:

```
/opt/omni/lbin/omnigetmsg 12 Error_number
```

On Windows systems, perform the following procedure using the Data Protector GUI:

1. In the Context List, select **Backup**.
2. In the Scoping Pane, expand **Backup**, **Backup Specifications**, and then **SAP R/3**. A list of SAP backup specifications is displayed.
3. In the Scoping Pane, select the failed backup specification and right-click on the SAP R/3 server item in the Results Pane to display a pop-up menu.
4. From the pop-up menu, select **Check Configuration**.

A short description of the problems and how to resolve them is displayed.

- **Backup does not work**
  - Check whether the Cell Manager is correctly set on the application system. The file /etc/opt/omni/client/cell_server (UNIX systems) or HKEY_LOCAL_MACHINE\SOFTWARE\Hewlett-Packard\OpenView\OmniBackII\Site\CellServer (Windows systems) must contain the name of the Cell Manager.
  - Check that the primary_db parameter in the init*ORACLE_SID*.sap file on the application system is set to LOCAL.
  - On UNIX systems, check whether the users are properly configured in user groups. Both the UNIX Oracle administrator (ora*ORACLE_SID*) and UNIX SAP administrator (*ORACLE_SID*adm) have to be in the Data Protector operator class.
  - On UNIX systems, check whether the permissions of the *SAPDATA_HOME*/sapbackup/ directory are set to 755.
  - On Windows systems, check that the user account that started the Data Protector Inet service is added in the Data Protector operator class.
- **Backup fails with "Connect to database instance failed"**

  If you start a backup while the database instance is in the unmount or mount mode, the session fails with a message similar to the following:

```
BR0301E SQL error -1033 at location BrDbConnect-2
ORA-01033: ORACLE initialization or shutdown in progress
BR0310E Connect to database instance HOOHOO failed
```

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

### Problem

**Configuration fails due to a script failure**

During configuration of an SAP R/3 database, Data Protector reports the following error:

```
Integration cannot be configured.

Script failed. Cannot get information from remote host.
```

### Action

If the SAP R/3 database is located on a Windows system, check the environment settings and ensure Data Protector Inet is running under a user account which has the required privileges. For details, see "Before you begin" on page 174.

If the SAP R/3 database is located on a UNIX system, resolve the problem by reviewing the user account configuration. For details, see "Configuring user accounts" on page 175.

## Restore problems

### Problem

**ZDB, restore, or instant recovery sessions fail due to invalid characters in filenames**

On Windows systems, where the Oracle Database Character Set (DBCS) is not set to the same value as the default Windows character set for non-Unicode programs, and where SAP tools are used to create Oracle datafiles, ZDB, restore, and Instant Recovery fail if the datafiles contain non-ASCII or non-Latin 1 characters.

### Actions

Use any of the following solutions:

- For new Oracle installations, set the DBCS to UTF-8.
- If you do not use other non-Unicode programs, set the language for non-Unicode programs to the same value as DBCS.

- Do not use non-ASCII or non-Latin 1 characters for filenames.

## Problem

### Restore of SAP R/3 tablespaces located on raw partitions fails

When restoring SAP tablespaces that are located on raw partitions using the Data Protector GUI, the restore fails with a message similar to the following:

```
[Major] From: VRDA@joca.company.com "SAP" Time: 5/9/06 3:33:51
PM /dev/sapdata/rsapdata Cannot restore -> rawdisk section
![Warning] From: VRDA@joca.company.com "SAP" Time: 5/9/06
3:42:45 PM Nothing restored.
```

## Action

Use SAP commands (for example, `brrestore`) to restore these tablespaces.

# 3 Data Protector Microsoft SQL Server ZDB integration

## Introduction

This chapter explains how to configure and use the Data Protector Microsoft SQL Server ZDB integration. It describes the concepts and methods you need to understand to back up and restore the Microsoft SQL Server (**SQL Server**) database objects.

During the backup, an SQL Server snapshot is made (the database files are frozen and any transactions to them are cached), so the database is highly available (*online* backup). The I/O to it is suspended during the time it takes to create a **replica** (split the mirror disks or create snapshots).

NOTE:

SQL Server snapshot is an SQL Server related term and does not mean the same as a disk array snapshot.

The following disk arrays and array configurations are supported:

| Supported array | Supported configurations |
|---|---|
| EMC Symmetrix (EMC) | Dual host TimeFinder |
| HP StorageWorks Disk Array XP (XP) | BC, CA, combined BC+CA |
| HP StorageWorks Enterprise Virtual Array (EVA) | BC, combined CA+BC |
| HP StorageWorks Virtual Array (VA) | BC |

All ZDB types (ZDB to tape, ZDB to disk, and ZDB to disk+tape) are supported by this integration. For ZDB types description, see the *HP Data Protector zero downtime backup concepts guide* and the online Help.

Using Data Protector, you can restore your SQL Server data:

- From backup media to the application system on LAN (standard restore).
- Using the instant recovery functionality.

The following table gives an overview of SQL Server recovery methods:

| ZDB type | Recovery method |
|---|---|
| ZDB to tape | Standard restore |
| ZDB to disk | Instant recovery |
| ZDB to disk+tape | Standard restore, instant recovery |

For a description of ZDB and instant recovery concepts, see the *HP Data Protector zero downtime backup concepts guide*.

# Integration concepts

Data Protector integrates with SQL Server through the Data Protector `sql_bar.exe` executable, installed on SQL Server. During backup, `sql_bar.exe`, started on the application system, connects to SQL Server to find the locations of the database files. The integration then backs up SQL Server database(s), which are replicated within a disk array.

During restore, `sql_bar.exe` connects to SQL Server to receive the restore data, which is then written to disks.

ZDB process depends on whether you replicate your data in a split mirror or snapshot replication and on the selected ZDB type. Restore process depends on the restore type - standard restore or instant recovery. See the *HP Data Protector zero downtime backup concepts guide* for a detailed description of replication techniques and ZDB and restore processes.

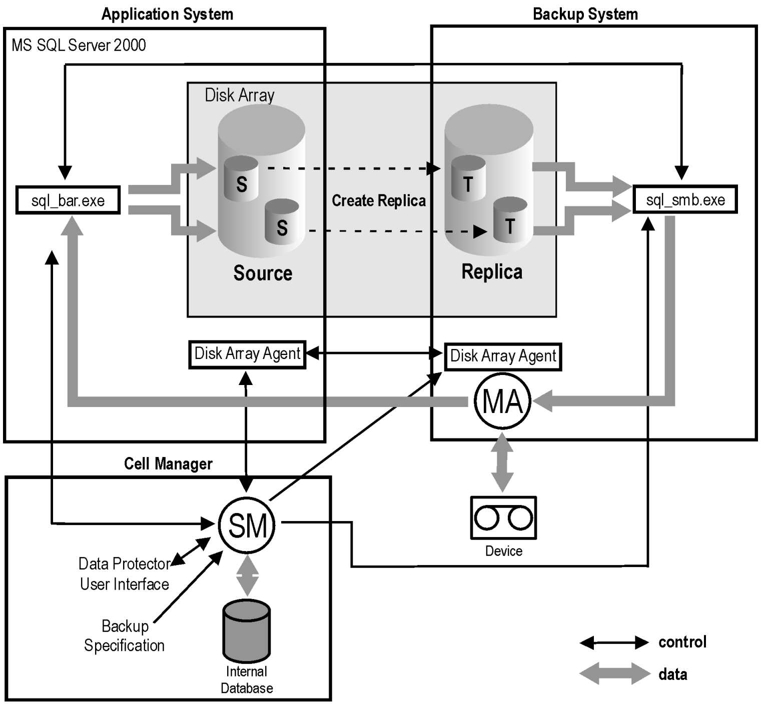

**Figure 63 Backup and restore concepts**

# Configuring the integration

## Prerequisites

- You need a license to use the SQL Server ZDB integration. See the *HP Data Protector installation and licensing guide* for information.
- Ensure that you correctly installed and configured SQL Server.
  - For supported versions, platforms, devices, and other information, see the latest support matrices at http://www.hp.com/support/manuals.
  - For information on installing, configuring, and using SQL Server, see the SQL Server documentation.
- Ensure that you correctly installed Data Protector. For information on installing Data Protector in various architectures and installing a Data Protector disk array integration (EMC, XP, VA, or EVA) with SQL Server, see the *HP Data Protector installation and licensing guide*.

Every SQL Server to be used with Data Protector must have the `MS SQL Integration` component installed.

- Install SQL Server on the application system. Install user databases on the disk array source volumes (system databases can be installed anywhere). If the system databases are also installed on a disk array, they *must* reside on *different* source volumes than user databases.

  If SQL Server is installed on the backup system as well, its databases *must* reside on source volumes *different* from the source volumes used for this integration. Drive letters/mount points assigned to those volumes must also be different from the drive letters/mount points assigned to the volumes used for this integration.

## Before you begin

- Configure devices and media for use with Data Protector. See the online Help index: “configuring devices” and “creating media pools” for instructions.

  For ZDB to disk, configure a backup device (for example, a standalone file device), as you will have to select it when configuring a backup specification. For information on configuring standalone devices, see the online Help index: “standalone devices”.
- If you plan to use **Integrated authentication** to connect to an SQL Server instance, you need to restart the Data Protector `Inet` service under a Windows domain user account that has the appropriate SQL Server permissions for running backups and restores. On Windows Server 2008, this is not required.

  For information on changing the user account under which the Data Protector `Inet` service is running, see the online Help index: “Inet, changing account”.
- To test whether SQL Server and Cell Manager communicate properly, configure and run a Data Protector filesystem ZDB and restore. See the online Help for instructions.

## Data Protector SQL Server configuration file

Data Protector stores integration parameters for every configured SQL Server on the Cell Manager in:

***HP-UX, Solaris:***

`/etc/opt/omni/server/integ/config/MSSQL/`*`client_name%instance_name`*

***Windows:***

*`Data_Protector_home`*`\Config\Server\Integ\Config\MSSQL\`*`client_name%instance_name`*

Configuration parameters are the username and password of the SQL Server user, who must have permissions to run backups and restores within SQL Server (assuming the standard security is used). They are written to the Data Protector SQL Server configuration file during configuration of the integration.

The content of the configuration file is:

```
Login='user';
Password='encoded_password';
Domain='domain';
```

IMPORTANT:

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

### Examples

- ***SQL Server authentication:***

  ```
  Login='sa';
  Domain='';
  Password='jsk74yh80fh43kdf';
  ```

- ***Windows authentication:***

  ```
  Login='Administrator';
  Domain='IPR';
  Password='dsjf08m80fh43kdf';
  ```

- ***Integrated authentication:***

  ```
  Login='';
  Domain='';
  Password='kf8u3hdgtfh43kdf';
  ```

## Configuring users

If you have restarted the Data Protector `Inet` service on the SQL Server system under a different user account, add this user to the Data Protector `admin` or `operator` Data Protector user group.

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

# Configuring SQL Server instances

An SQL Server instance is configured during the creation of the first backup specification. The configuration consists of setting the user account that Data Protector should use to connect to the SQL Server instance. The specified login information is saved to the Data Protector SQL Server instance configuration file on the Cell Manager.

NOTE:

Ensure that the user account to be used has appropriate SQL Server permissions for running backups and restores. Check the permissions using SQL Server Enterprise Manager.

You can change configuration by following instructions described in "Changing and checking configuration" on page 252.

## Prerequisites

- SQL Server must be online during configuration.
- Configuration must be performed for every SQL Server instance separately.

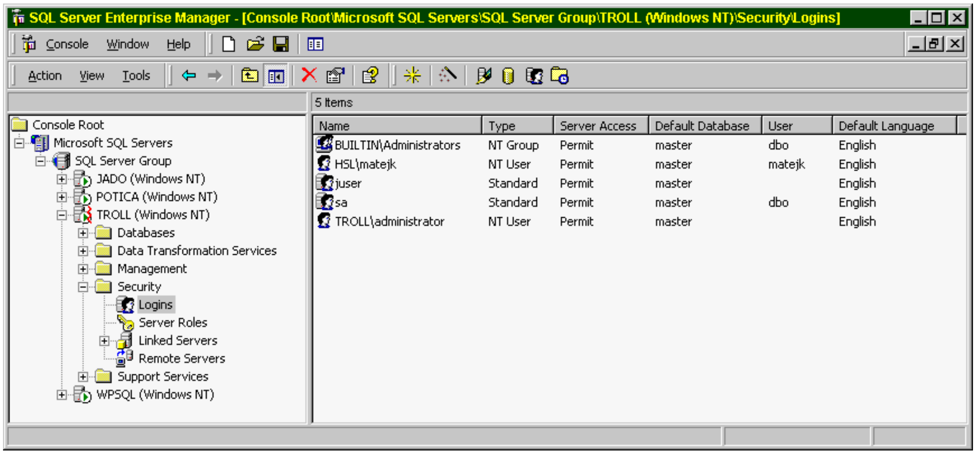

| Name | Type | Server Access | Default Database | User | Default Language |
|---|---|---|---|---|---|
| BUILTIN\Administrators | NT Group | Permit | master | dbo | English |
| HSL\matejk | NT User | Permit | master | matejk | English |
| juser | Standard | Permit | master | | English |
| sa | Standard | Permit | master | dbo | English |
| TROLL\administrator | NT User | Permit | master | | English |

### Figure 64 SQL Server users

Configure SQL Server instances using the Data Protector GUI or CLI.

## Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **MS SQL Server**, and click **Add Backup**.
3. In the **Create New Backup dialog box**, select the **Blank Microsoft SQL Server Backup** template and specify backup type. For details, see Step 3 on page 255.

   Click **OK**.
4. Specify the ZDB specific options. For details, see Step 4 on page 257.
5. In **Application database**, select or specify the name of the SQL Server instance.

   ***Windows Server 2008 only:*** If you intend to use the **Integrated authentication** option, specify the **User and group/domain** options. For information on the options, press **F1**.

   Click **Next**.

6. In the Configure MS SQL Server dialog box, specify the user account that Data Protector should use to connect to the SQL Server instance.

   - **SQL Server authentication**: SQL Server user account. Specify a username and password.
   - **Windows authentication**: Windows domain user account (preferred option). Specify a username, password, and the domain.
   - **Integrated authentication**: Select this option to enable Data Protector to connect to the SQL Server instance with the following Windows domain user account:
     - ***Windows Server 2008:*** The account specified in the **User and group/domain** options in the previous step or in the Client selection page.
     - ***Other Windows systems:*** The account under which the Data Protector `Inet` service on the SQL Server system is running.

   Ensure that the user account you specify has the appropriate permissions for backing up and restoring the SQL Server databases.

   See Figure 65 on page 250.

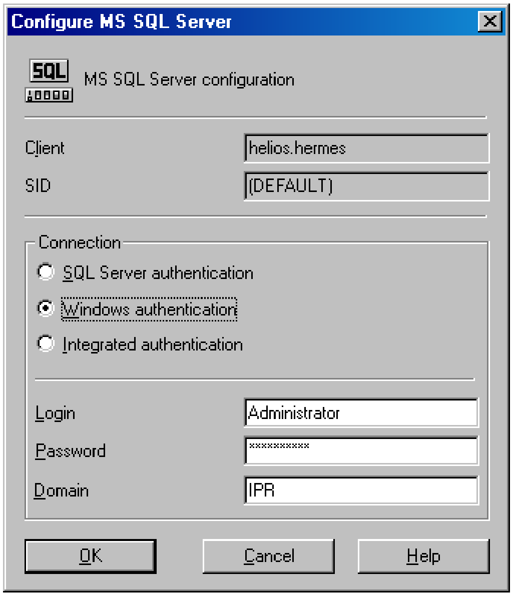

Configure MS SQL Server

MS SQL Server configuration

Client: helios.hermes

SID: (DEFAULT)

Connection

○ SQL Server authentication

● Windows authentication

○ Integrated authentication

Login: Administrator

Password: ***********

Domain: IPR

**Figure 65 Configuring SQL Server**

NOTE:

It is recommended that the SQL Server system administrator configures the integration.

For details about security, see the SQL Server documentation.

Click **OK** to confirm the configuration.

7. The SQL Server instance is configured. Exit the GUI or proceed with creating the backup specification at Step 7 on page 261.

## Using the Data Protector CLI

From the *Data_Protector_home*\bin directory, run:

```
sql_bar config [-appsrv:SQL_Server_client]
[-instance:instance_name] [-dbuser:SQL_Server_user
-password:password | -dbuser:Windows_user -password:password
-domain:domain]
```

### Parameter description

`-appsrv:SQL_Server_client`

Client system on which the SQL Server instance is running. This option is not required if you run the command locally.

`-instance:instance_name`

SQL Server instance name. If you omit this option, the default SQL Server instance is configured.

`-dbuser:SQL_Server_user -password:password`

SQL Server user account (**SQL Server authentication**)

`-dbuser:Windows_user -password:password -domain:domain`

Windows domain user account (**Windows authentication**)

NOTE:

If no user account is specified, Data Protector uses **Integrated authentication**.

The message `*RETVAL*0` indicates successful configuration.

## Changing and checking configuration

You can check and change configuration using the Data Protector GUI or CLI.

### Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **MS SQL Server**. Click a backup specification for which you want to change the configuration.
3. In the **Source** property page, right-click the SQL Server name and select **Configure**.
4. Configure SQL Server as described in "Configuring SQL Server instances" on page 248.

5. Right-click SQL Server and select **Check Configuration**. See Figure 66 on page 253.

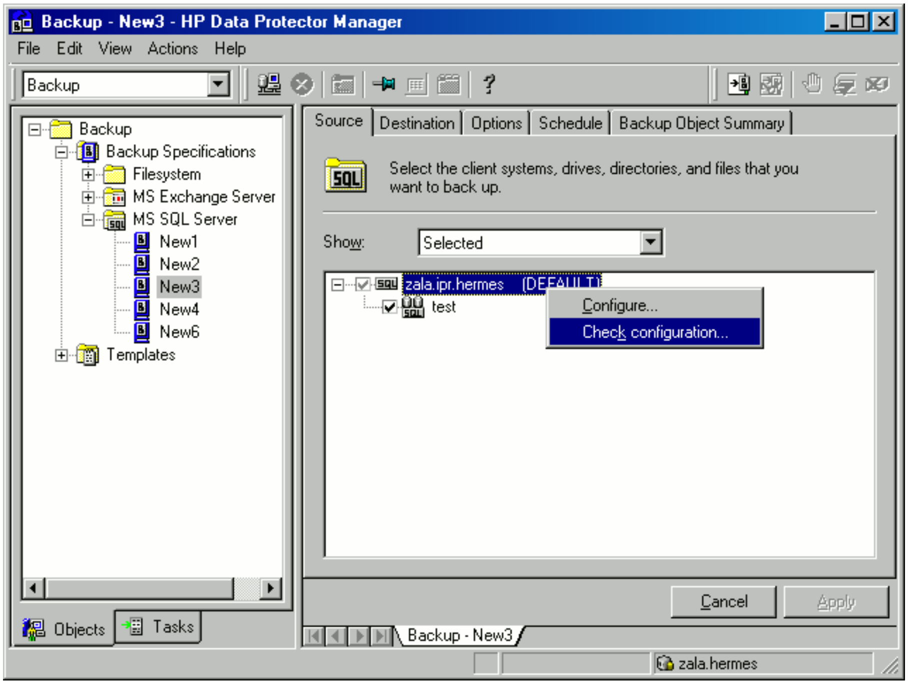

**Figure 66 Checking configuration**

## Using the Data Protector CLI

To change the configuration, run the command for configuring SQL Server instances again, entering different data.

To check configuration, run:

```
sql_bar chkconf [-instance:instance_name]
```

If the optional parameter `-instance:instance_name` is not specified, the default instance is checked.

If the integration is not properly configured, the command returns:

```
*RETVAL*8523
```

To get the information about the existing configuration, run:

```
sql_bar getconf [-instance:instance_name]
```

If `-instance:`*`instance_name`* is not specified, Data Protector returns configuration for the default instance.

# Backup

To run ZDB of an existing SQL Server ZDB specification:

- Schedule a backup using the Data Protector Scheduler.
- Start an interactive backup using the Data Protector GUI or CLI.

### Prerequisites

- The same drive letters, on which the source volumes to be replicated reside on the application system, must exist on the backup system to enable successful mounting of the target volumes. If the same drive letters do not exist on the backup system, the backup fails. The drive letters on the application and backup systems must be the same also in case of nested mount points.

### Considerations

Your session will fail if:

- You start ZDB, restore, or instant recovery using the same source volume on the application system at the same time. A session must be started only after the preceding session using the same source volume on the application system finishes.
- SQL Server services are not running when the backup starts.

To configure a ZDB, create a Data Protector SQL Server ZDB specification.

## Creating ZDB specifications

Create a ZDB specification, using the Data Protector Manager.

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, right-click **MS SQL Server**, and click **Add Backup**.

3. In the **Create New Backup** dialog box, select the **Blank Microsoft SQL Server Backup** template.

   ***XP, EMC:***

   In the **Backup type** drop-down list, select **Split mirror backup**. In the **Sub type** drop-down list, select either **EMC Symmetrix** or **HP StorageWorks XP**. See Figure 67 on page 255.

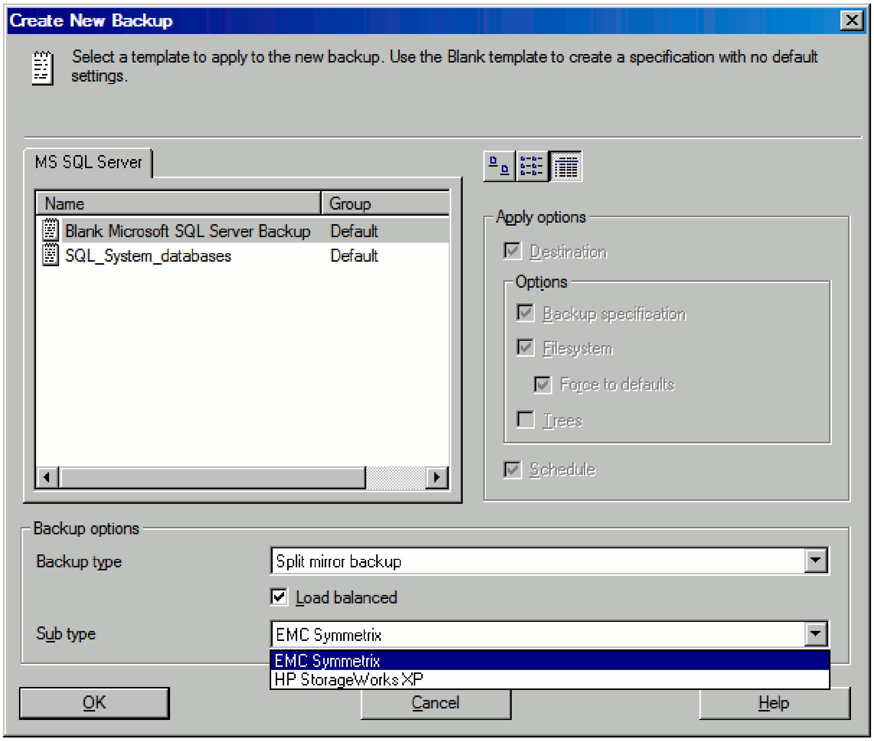

Create New Backup

Select a template to apply to the new backup. Use the Blank template to create a specification with no default settings.

MS SQL Server

| Name | Group |
|---|---|
| Blank Microsoft SQL Server Backup | Default |
| SQL_System_databases | Default |

Apply options

☑ Destination

Options

☑ Backup specification

☑ Filesystem

☑ Force to defaults

☐ Trees

☑ Schedule

Backup options

Backup type: Split mirror backup

☑ Load balanced

Sub type: EMC Symmetrix

## Figure 67 Selecting a split mirror backup

***VA, EVA:***

In the **Backup type** drop-down list, select **Snapshot backup**. In the **Sub type** drop-down list, select either **HP StorageWorks VA** or **HP StorageWorks EVA SMI-S**. See Figure 68 on page 256.

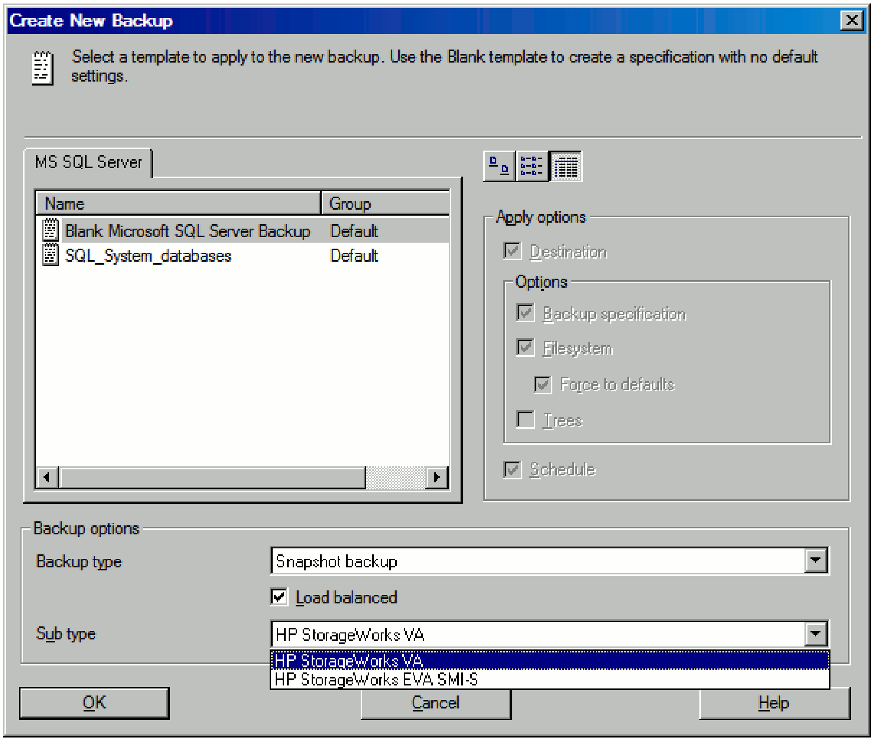

Create New Backup

Select a template to apply to the new backup. Use the Blank template to create a specification with no default settings.

MS SQL Server

| Name | Group |
|---|---|
| Blank Microsoft SQL Server Backup | Default |
| SQL_System_databases | Default |

Apply options

☑ Destination

Options

☑ Backup specification

☑ Filesystem

☑ Force to defaults

□ Trees

☑ Schedule

Backup options

Backup type: Snapshot backup

☑ Load balanced

Sub type: HP StorageWorks VA

HP StorageWorks VA

HP StorageWorks EVA SMI-S

Figure 68 Selecting a snapshot backup

Click **OK**.

4. Under **Client systems**, select the SQL Server system. In cluster environments, select the virtual server of the SQL Server resource group.

   In the **Backup system** drop-down list, select the backup system.

   Select other disk array specific backup options (see Figure 70 on page 258 for EMC, Figure 71 on page 259 for XP, Figure 72 on page 260 for VA, or Figure 72 on page 260 for EVA). For detailed information on options, press **F1**.

   ***EMC:***

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

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

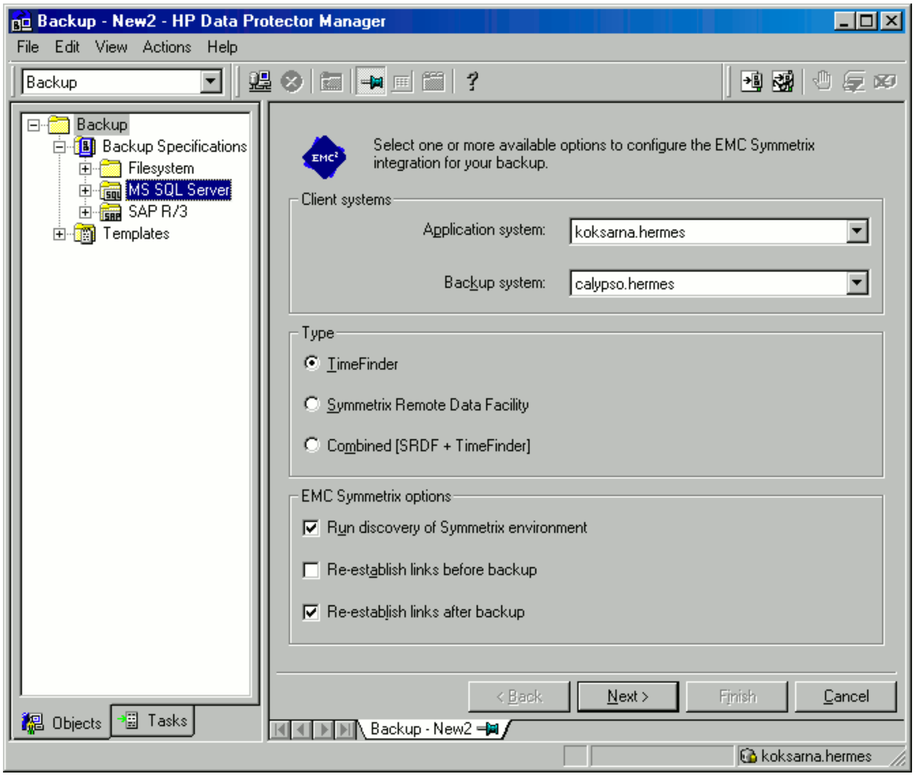

Select one or more available options to configure the EMC Symmetrix integration for your backup.

Client systems

Application system: koksarna.hermes

Backup system: calypso.hermes

Type

● TimeFinder

○ Symmetrix Remote Data Facility

○ Combined (SRDF + TimeFinder)

EMC Symmetrix options

☑ Run discovery of Symmetrix environment

☐ Re-establish links before backup

☑ Re-establish links after backup

**Figure 69 EMC backup options**

   ***XP:***

   To enable instant recovery, leave **Track the replica for instant recovery** selected.

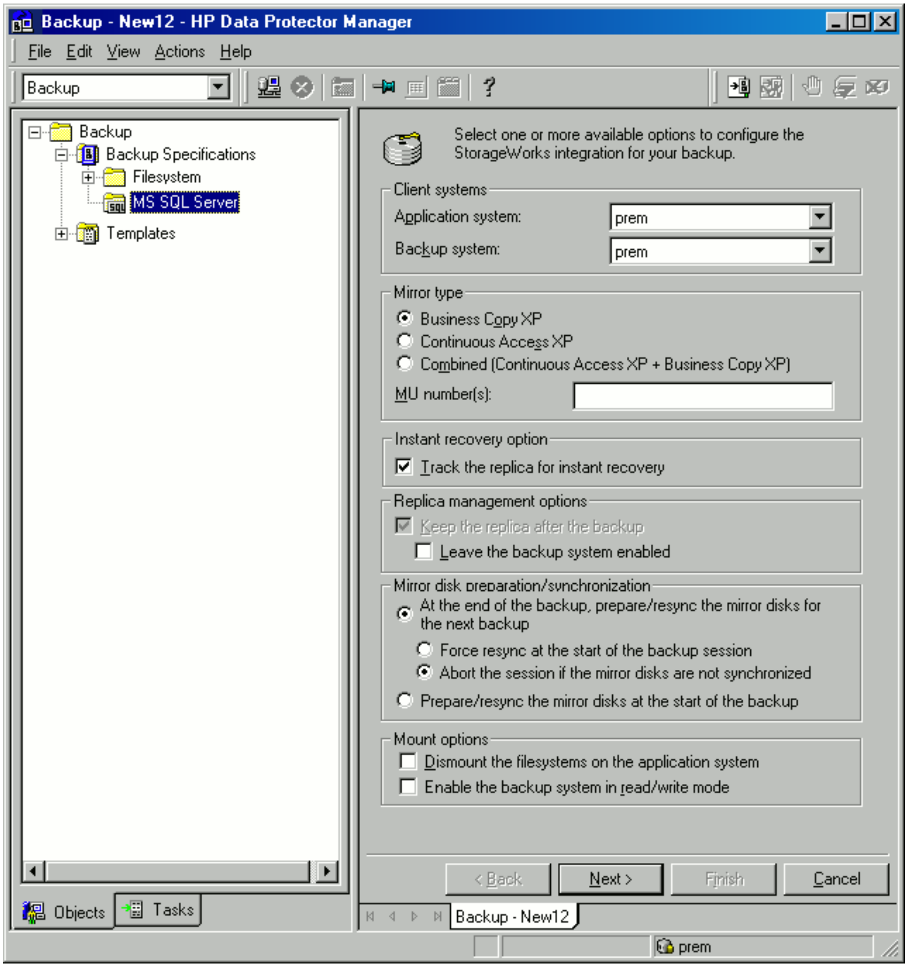

Select one or more available options to configure the StorageWorks integration for your backup.

Client systems

Application system: prem

Backup system: prem

Mirror type

● Business Copy XP

○ Continuous Access XP

○ Combined (Continuous Access XP + Business Copy XP)

MU number(s):

Instant recovery option

☑ Track the replica for instant recovery

Replica management options

☑ Keep the replica after the backup

□ Leave the backup system enabled

Mirror disk preparation/synchronization

● At the end of the backup, prepare/resync the mirror disks for the next backup

○ Force resync at the start of the backup session

● Abort the session if the mirror disks are not synchronized

○ Prepare/resync the mirror disks at the start of the backup

Mount options

□ Dismount the filesystems on the application system

□ Enable the backup system in read/write mode

## Figure 70 XP backup options

***VA:***

To enable instant recovery, leave **Track the replica for instant recovery** selected.

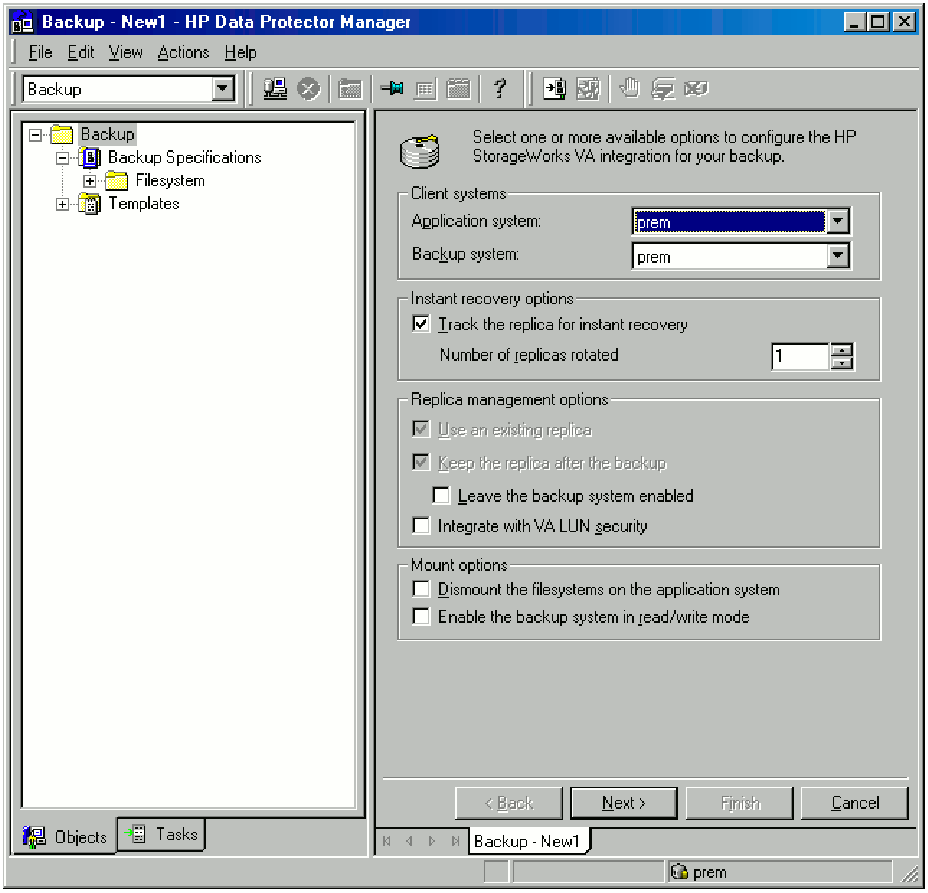

Select one or more available options to configure the HP StorageWorks VA integration for your backup.

Client systems

Application system: prem

Backup system: prem

Instant recovery options

☑ Track the replica for instant recovery

Number of replicas rotated: 1

Replica management options

☑ Use an existing replica

☑ Keep the replica after the backup

□ Leave the backup system enabled

□ Integrate with VA LUN security

Mount options

□ Dismount the filesystems on the application system

□ Enable the backup system in read/write mode

## Figure 71 VA backup options

***EVA:***

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

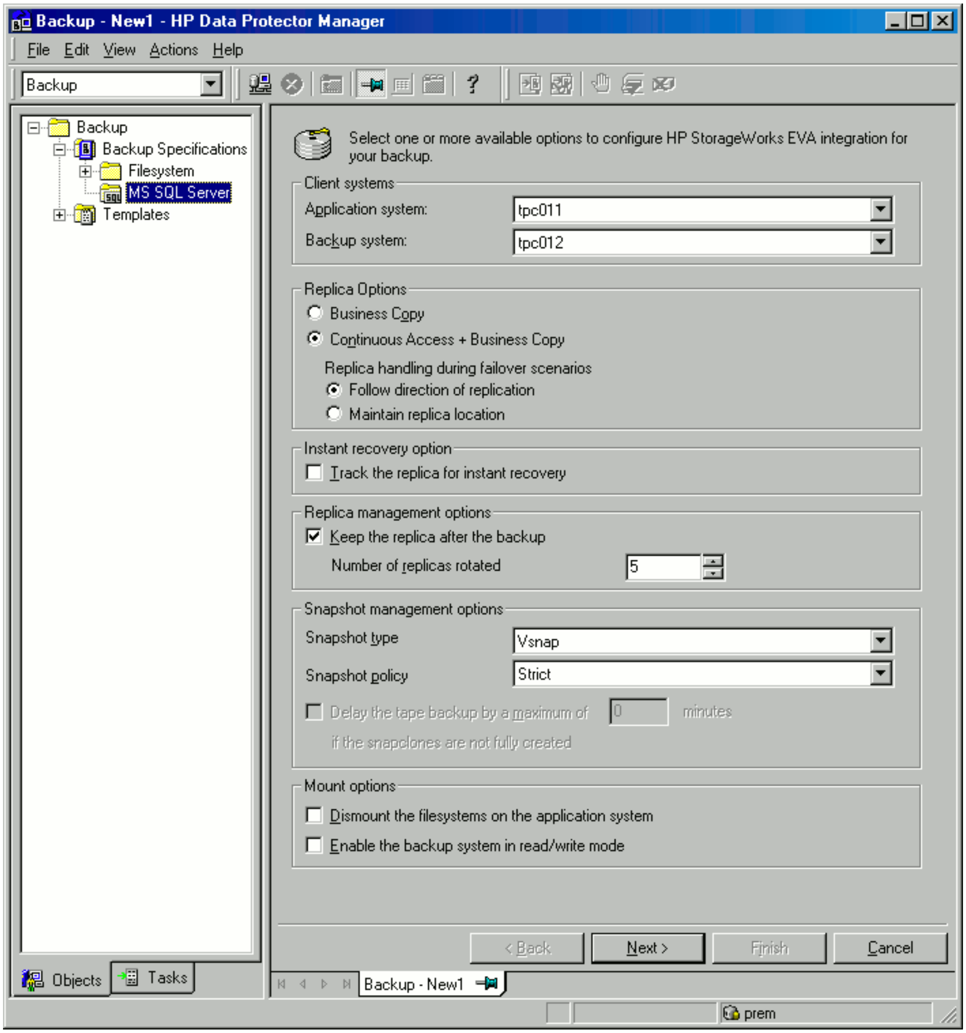

Select one or more available options to configure HP StorageWorks EVA integration for your backup.

Client systems

Application system: tpc011

Backup system: tpc012

Replica Options

○ Business Copy

● Continuous Access + Business Copy

Replica handling during failover scenarios

● Follow direction of replication

○ Maintain replica location

Instant recovery option

□ Track the replica for instant recovery

Replica management options

☑ Keep the replica after the backup

Number of replicas rotated 5

Snapshot management options

Snapshot type Vsnap

Snapshot policy Strict

□ Delay the tape backup by a maximum of 0 minutes if the snapclones are not fully created

Mount options

□ Dismount the filesystems on the application system

□ Enable the backup system in read/write mode

< Back | Next > | Finish | Cancel

**Figure 72 EVA backup options**

Click **Next**.

5. In **Application database**, specify the name of the SQL Server instance.

   ***Windows Server 2008 only:*** If you intend to use the **Integrated authentication** option, specify the **User and group/domain** options. For information on the options, press **F1**.

   Click **Next**.

6. If the client is not configured, the **Configure MS SQL Server** dialog box appears. Configure it as described in "Configuring SQL Server instances" on page 248.
7. Select the databases to be backed up.

IMPORTANT:

To enable instant recovery, create different backup specifications for user and system databases.

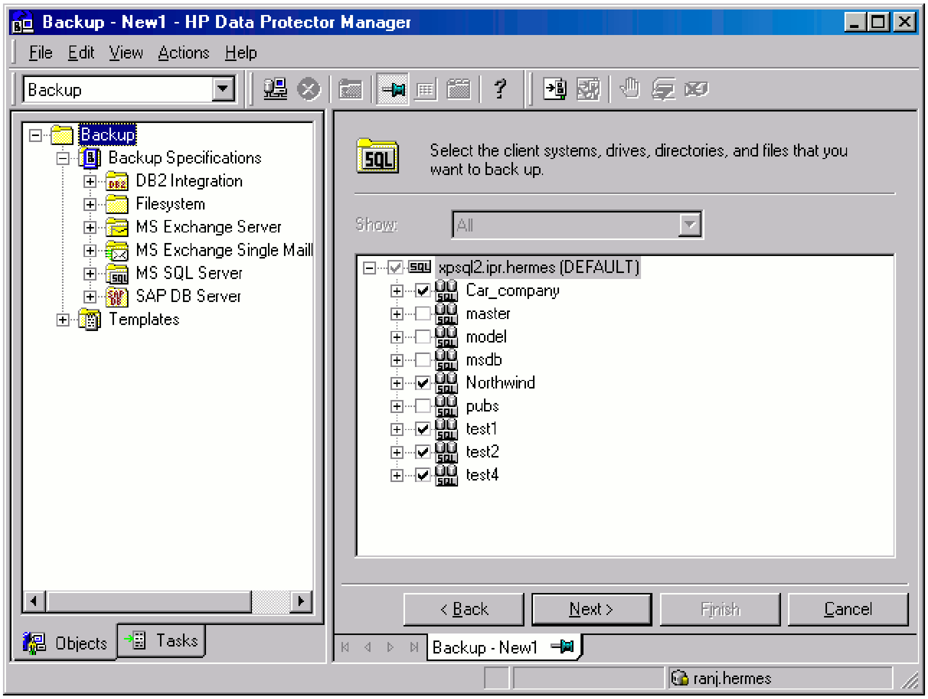

Select the client systems, drives, directories, and files that you want to back up.

Show: All

- ☑ xpsql2.ipr.hermes (DEFAULT)
  - ☑ Car_company
  - □ master
  - □ model
  - □ msdb
  - ☑ Northwind
  - □ pubs
  - ☑ test1
  - ☑ test2
  - ☑ test4

Figure 73 Selecting user databases

Click **Next**.

8. Select the devices. Click **Properties** to set the media pool and preallocation policy. The device concurrency is set to 1 and cannot be changed. For more information on options, press **F1**.

   To create additional backup copies (mirrors), specify the desired number by clicking **Add mirror**/**Remove mirror**. Select separate devices for each mirror. The minimum number of devices for mirroring equals the number of devices used for backup.

   For more information on object mirroring, see the online Help.

   NOTE:

   Object mirroring is not supported for ZDB to disk.

   Click **Next**.

9. Select backup options.

   For information on **Backup Specification Options** and **Common Application Options**, see the online Help.

   For information on **Application Specific Option**, see "SQL Server Specific Backup Options" on page 263.

   Click **Next**.

10. Optionally, schedule the backup. For information on scheduler, press **F1**.

    Note that only **Full** backup is performed.

11. Save the backup specification, specifying a name and backup specification group. You start the backup specification by clicking **Start Backup**.

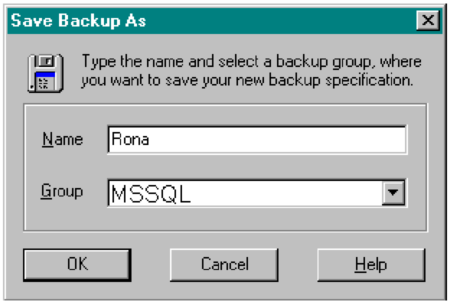

Save Backup As

Type the name and select a backup group, where you want to save your new backup specification.

Name: Rona

Group: MSSQL

Figure 74 Saving a backup specification

## SQL Server Specific Backup Options

SQL Server specific backup options are specified by clicking the **Advanced** tab in the **Application Specific Options** group box.

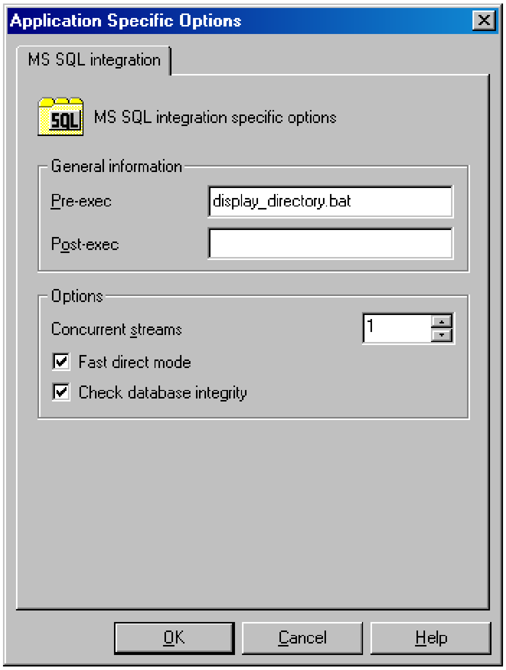

Application Specific Options

MS SQL integration

MS SQL integration specific options

General information

Pre-exec: display_directory.bat

Post-exec:

Options

Concurrent streams: 1

☑ Fast direct mode

☑ Check database integrity

### Figure 75 Application specific options

### Table 16 SQL Server backup options

| | |
|---|---|
| **Concurrent streams** | Sets the number of concurrent streams used to back up SQL Server databases from the replica to tape. Applicable for ZDB–to–tape and ZDB–to–disk+tape sessions. |
| **Fast direct mode** | Ignored for ZDB sessions. |
| **Check database integrity** | Performs data integrity validation before backup. If the check fails, the session completes with warnings. |
| **Pre-exec** | Specifies a command with arguments or a script started by `sql_bar.exe` on SQL Server before backup. Resides in the *`Data_Protector_home`*`\bin` directory. Only the filename must be provided in the backup specification. |
| **Post-exec** | Specifies a command with arguments or a script started by `sql_bar.exe` on SQL Server after backup. Resides in the *`Data_Protector_home`*`\bin` directory. Only the filename must be provided in the backup specification. |

NOTE:

Do not use double quotes (" ") in object-specific pre-exec and post-exec commands.

## Scheduling backups

You can run unattended ZDB at specific times or periodically. For details on scheduling, see the online Help index: “scheduled backups”.

NOTE:

You cannot run ZDB to disk or ZDB to disk+tape if **Track the replica for instant recovery** is not selected in the backup specification.

### Scheduling example

To schedule a database ZDB at 8:00, 13:00, and 18:00 during week days:

1. In the **Schedule** property page, select the starting date in the calendar and click **Add** to open the **Schedule Backup** dialog box.
2. Under **Recurring**, select **Weekly**. Under **Time options**, select **8:00**. Under **Recurring Options**, select **Mon**, **Tue**, **Wed**, **Thu**, and **Fri**.

   Click **OK**.
3. Repeat Step 1 on page 265 and Step 2 on page 265 to schedule backups at 13:00 and 18:00.
4. Click **Apply** to save the changes.

NOTE:

For ZDB sessions, the backup type is set to **Full**.

## Starting backup sessions

Interactive backups are run on demand. They are useful for urgent backups or restarting failed backups.

## Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, and then **MS SQL Server**. Right-click the backup specification you want to start and select **Start Backup**.
3. Select **Network load**. For information on network load, click **Help**. Click **OK**.

   For ZDB sessions, the backup type is set to **Full**.

   For ZDB to disk or ZDB to disk+tape, specify the **Split mirror/snapshot backup** option.

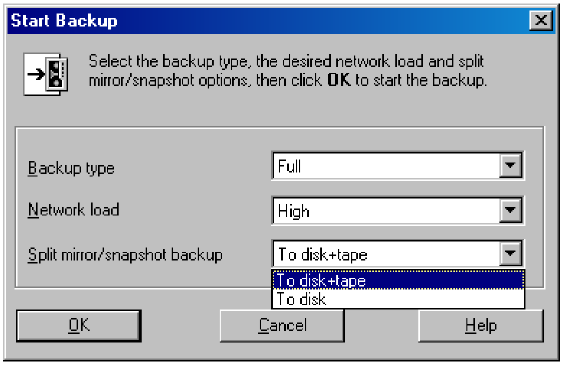

Start Backup

Select the backup type, the desired network load and split mirror/snapshot options, then click **OK** to start the backup.

Backup type: Full

Network load: High

Split mirror/snapshot backup: To disk+tape

To disk+tape
To disk

**Figure 76 Starting interactive backups**

## Using the Data Protector CLI

To start ZDB to tape or ZDB to disk+tape, run:

```
omnib -mssql_list ListName
```

To start ZDB to disk, run:

```
omnib -mssql_list ListName -disk_only
```

where `ListName` is the name of the backup specification. For more information on `omnib`, see its man page.

# Restore

Data Protector offers restore from backup media to the application system on LAN (standard restore), where you can select various restore options depending on your restore scenario, and instant recovery. For more information, see the following sections.

## Before you begin

- Before starting restore, verify that the database is not being in use.

## Standard restore

NOTE:

There is no need to create an empty database before restore, because the database and its files are generated automatically by SQL Server.

Proceed as follows using the Data Protector Manager:

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand **Restore Objects**, **MS SQL Server**, and then select the client (backup system) from which you want to restore. A list of backed up objects is displayed in the Results Area.

3. Select the SQL Server objects that you want to restore. See Figure 77 on page 268.

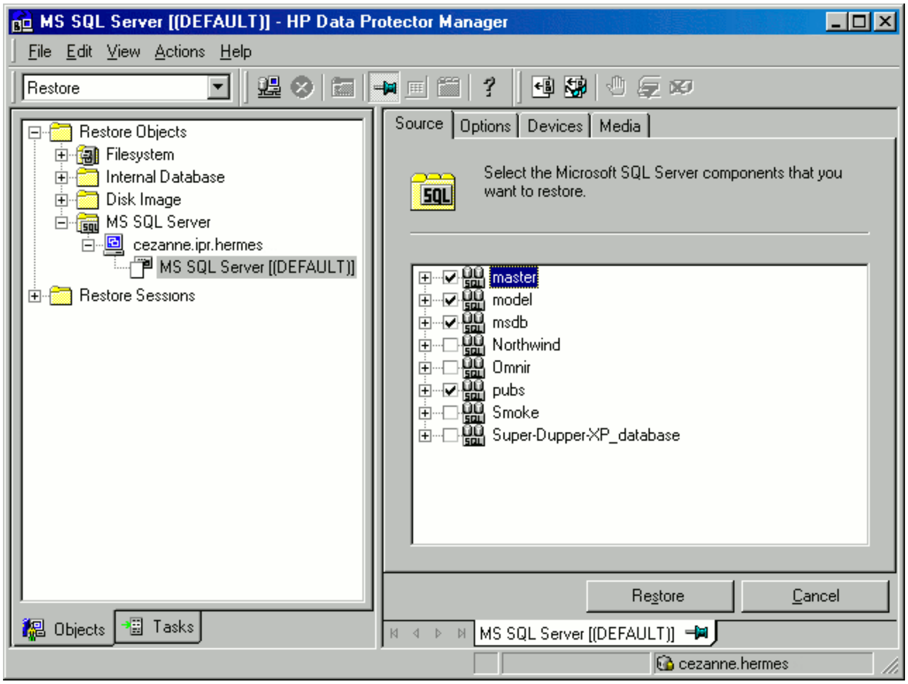

Select the Microsoft SQL Server components that you want to restore.

- ☑ master
- ☑ model
- ☑ msdb
- □ Northwind
- □ Omnir
- ☑ pubs
- □ Smoke
- □ Super-Dupper-XP_database

## Figure 77 Restore objects

To select backup object specific options, right-click the object and select **Properties**.

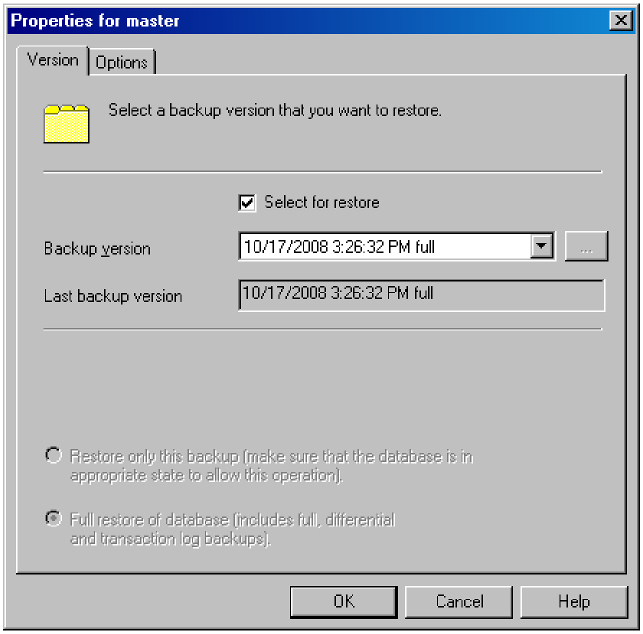

Select a backup version that you want to restore.

☑ Select for restore

Backup version: 10/17/2008 3:26:32 PM full

Last backup version: 10/17/2008 3:26:32 PM full

○ Restore only this backup (make sure that the database is in appropriate state to allow this operation).

● Full restore of database (includes full, differential and transaction log backups).

Figure 78 Selecting object specific options

You can select the version (backup date) from which you want to restore and choose SQL Server specific restore options. Note that some options are not available for restore of data files. See "Restore options" on page 271 for details.

Click **OK**.

4. If you want to restore your data to another client or instance, specify new locations for the databases in the **Options** property page. See "Restore options" on page 271.

NOTE:

When you click **Options**, the cell is browsed for running SQL Server instances that can become target instances for restore. If no instances are found, **Restore to another instance** is disabled and the message **There are no instances on this client system** is displayed.

Select one of the following **Restore actions**:

- **Restore data** (default). Select to restore the whole database.
- **Restore and display file list only**. Select if you do not know the original filenames. In this case, the files backed up in a particular session are displayed.
- **Restore and display headers only**. Select if you need specific details about backup. SQL Server header information is displayed.

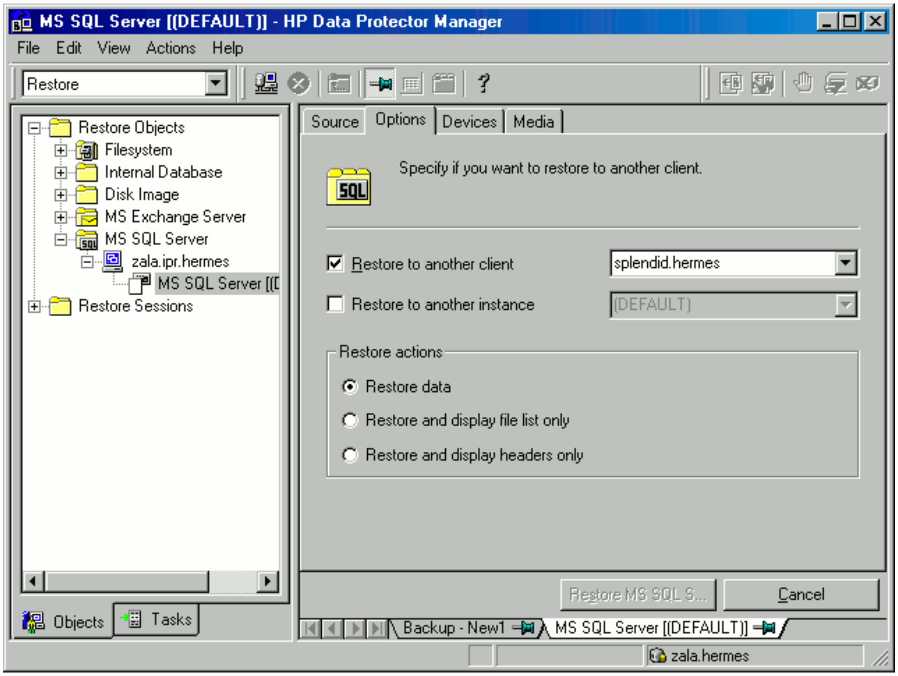

Specify if you want to restore to another client.

☑ Restore to another client: splendid.hermes

□ Restore to another instance: [DEFAULT]

Restore actions

● Restore data

○ Restore and display file list only

○ Restore and display headers only

**Figure 79 Restore options**

5. In the **Devices** page, select devices to use for the restore.

   The **Automatic device selection** option is selected by default, but it is recommended to select the **Original device selection** option.

IMPORTANT:

If you decide to select the **Automatic device selection** option, ensure that the number of available devices is equal to or greater than the number of devices that were used for backup.

NOTE:

For restore, you can use a different device than you used for backup. For information on restoring using another device, see the online Help index: "selecting, devices for restore".

6. Click **Restore MS SQL Server** and then **Next** to select **Report level** and **Network load**.

   Click **Finish** to start restore.

## Restore options

### Table 17 SQL Server restore options

| Option | Description |
|---|---|
| **Backup version** | Specifies the backup session from which the selected objects will be restored. |
| **Point-in-time restore** | This option is only available for database objects.<br>Specifies a point in time to which the database state will be restored (you also need to select **Backup version** and set **Stop at**). After recovery, the database is in the state it was at the specified date and time.<br>Only transaction logs written before the specified date and time are applied to the database. |

| Option | Description |
|---|---|
| **Stop at** | This option is only available for database objects.<br>Specifies the exact time when the rollforward of transactions will be stopped. Therefore, to enable database recovery to a particular point in time, backup you restore from must be a transaction log backup.<br>You cannot use this option with **NORECOVERY** or **STANDBY**. If you specify **Stop at** time that is after the end of **RESTORE LOG** operation, the database is left in a non-recovered state (as if **RESTORE LOG** is run with **NORECOVERY**). |
| **Restore only this backup** | If you restored a database version and left it in a non-operational or standby state, you can subsequently restore differential or transaction log backups one by one, leaving each version non-operational to restore additional backups. |
| **Full restore of the database** | All necessary versions are restored, including the latest full backup, the latest differential backup (if one exists), and all transaction log backups from the last differential up to the selected version. |
| **Force restore over the existing database** | Overwrites the existing database residing on the target SQL Server instance.<br>If a database with the same name as the one you are restoring already exists and has a different internal structure, SQL Server does not let you rewrite the database unless you select this option.<br>If you are restoring a data file from the PRIMARY group to an existing database, you *must* set this option for that data file. |
| **Recovery completion state** | Enables selecting the database state after recovery. You may select from:<br>• Leaving the database operational. Once the last transaction log is restored and the recovery completed, the database becomes operational.<br>• Leaving the database non-operational after the last transaction log is restored. You may restore additional transaction logs one by one.<br>• Leaving the database in read-only mode. You may restore additional transaction logs before the database is set to read-write mode.<br>This selection is only available for database objects. |

| Option | Description |
|---|---|
| **Restore database with a new name** | This option is only available for database objects.<br>Restores the database under a different name. Specify the database logical filename and the destination filename (suboptions of **Restore files to new locations**). |
| **Restore files to new locations** | Restores files to a new location. Specify the database logical filename and a destination target filename for the specified logical filename. Use this option if you restore data to another server, instance, or make a database copy on the same server. |

TIP:

To allow different restore scenarios, you can combine general restore options, such as **Restore database to another Microsoft SQL Server** and **Restore using a different device**, with object-specific restore options, such as **Point-in-time restore**, **Recovery completion state**, **Force restore over existing database**.

## Restoring to another SQL Serverinstance or/and another SQL Server

### Prerequisites

- Both SQL Servers must have the same local settings (code page and sort order). This information is displayed in the session monitor for each backup.
- The target SQL Server must be configured and reside in the same Data Protector cell as the original SQL Server. For configuration procedure, see "Creating ZDB specifications" on page 254.

1. Select the databases you want to restore and their versions.
2. Select the following:
   - To restore to another SQL Server client, select **Restore to another client** and the target client from the drop-down list.
   - To restore to another SQL Server instance, select **Restore to another instance**. If there are no instances in the drop-down list, enter the instance name yourself.
   - Ensure that the specified SQL Server instance exists on the target client. Otherwise, restore fails.
3. Specify new database locations.

4. Start restore. See “Restore” on page 266.

## Instant recovery

See the *HP Data Protector zero downtime backup concepts guide* and *HP Data Protector zero downtime backup administrator's guide* for general information on instant recovery.

### Prerequisites

- If you restore user databases, put the databases offline:
  1. Start SQL Server Enterprise Manager.
  2. Selecting the database and click **Action**.
  3. Select all tasks and take them offline.
- If you restore system databases, put SQL Server offline by starting SQL Server Enterprise Manager, right-clicking SQL Server, and clicking **Stop**.

Perform instant recovery using the Data Protector Manager:

1. In the Context List, click **Instant Recovery**.
2. Expand **MS SQL Server** and select the backup session (replica) from which you want to restore. By default, the database will be recovered until the last backed up transaction.

3. To recover user databases to a specific point in time:

   a. In the **Source** property page, under **Restore Objects**, right-click a database and click **Properties**.

   b. In the **Backup version** drop-down list, select the required replica. The latest version is selected by default.

      Select **Point in time restore**. From the **Stop at** drop-down list, select the point in time to which the transactions should be applied, and click **OK**. If no transaction logs are available, this option is disabled.

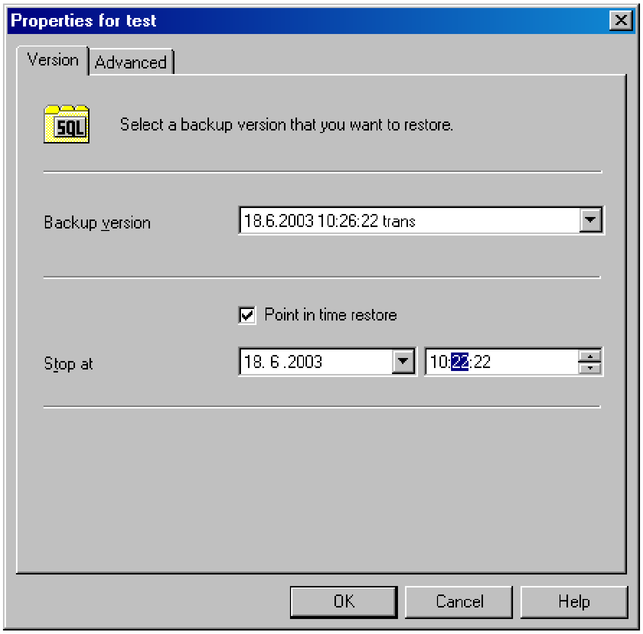

Select a backup version that you want to restore.

Backup version: 18.6.2003 10:26:22 trans

☑ Point in time restore

Stop at: 18. 6 .2003 10:22:22

**Figure 80 Point-in-Time restore**

To restore the database under a different name, click **Advanced** and select **Restore database with new name**. See Figure 81 on page 276.

IMPORTANT:

If the logical filename and physical filename are not listed, add them to the list. Specify the same names as used for ZDB; otherwise, instant recovery fails.

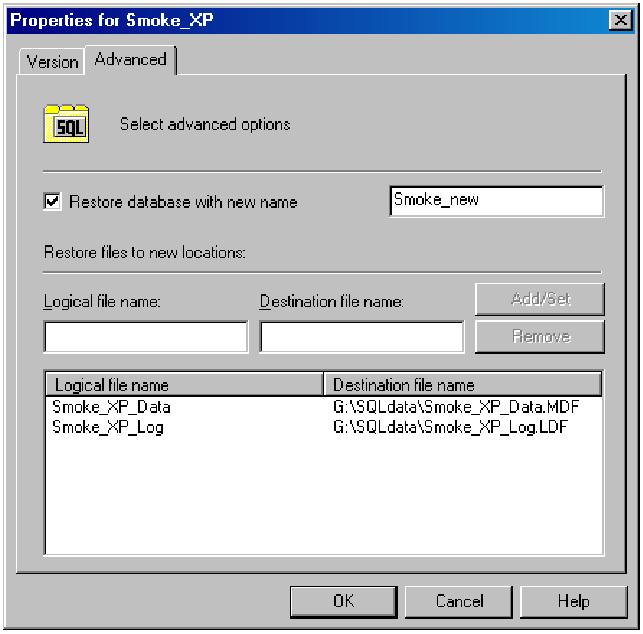

Select advanced options

☑ Restore database with new name: Smoke_new

Restore files to new locations:

Logical file name:

Destination file name:

| Logical file name | Destination file name |
|---|---|
| Smoke_XP_Data | G:\SQLdata\Smoke_XP_Data.MDF |
| Smoke_XP_Log | G:\SQLdata\Smoke_XP_Log.LDF |

**Figure 81 Restoring database with a new name**

4. Click **Restore MS SQL Server**.

   If you restore system databases, SQL Server displays errors because its services are offline. Therefore, when restore completes, start SQL Server manually using SQL Server Enterprise Manager.

# Monitoring sessions

You can monitor currently running or view previous sessions in the Data Protector GUI. When you run an interactive session, the monitor window shows you the session progress. Closing the GUI does not affect the session.

You can also monitor sessions using the **Monitor** context from any Data Protector client with the **User Interface** component installed.

For information on monitoring sessions, see the online Help index: “viewing currently running sessions” and “viewing finished sessions”.

# Troubleshooting

This section lists general checks and verifications, plus problems you might encounter when using the Data Protector SQL Server integration. Start at "Problems" on page 278. If you cannot find a solution there, perform general checks and verifications.

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

For general ZDB, restore, and instant recovery related troubleshooting, see the *HP Data Protector zero downtime backup administrator's guide*.

## Before you begin

- Ensure that the latest official Data Protector patches are installed. See the online Help index: "patches" on how to verify this.
- See the *HP Data Protector product announcements, software notes, and references* for general Data Protector limitations, as well as recognized issues and workarounds.
- See http://www.hp.com/support/manuals for an up-to-date list of supported versions, platforms, and other information.

## Checks and verifications

If your configuration, backup, or restore failed:

- Check that SQL Server services are running.
- Examine system errors reported in `*Data_Protector_home*\log\debug.log` on the SQL Server client.

  Additionally, check `errorlog` and `VDI.log` files in the `*MSSQL*\log` directory.
- Make a test filesystem backup and restore of the problematic client. For information, see the online Help.
- Check that every SQL Server used with Data Protector has the `MS SQL Integration` component installed.
- Connect to SQL Server via SQL Server Enterprise Manager using the same login ID as you specified in the Data Protector **Configuration** dialog box.
- Perform a database backup using SQL Server Enterprise Manager. If the backup fails, fix any SQL Server problems, and then perform a backup using Data Protector.

Additionally, if your backup failed:

- Verify the configuration file to check if the Cell Manager is correctly set on SQL Server.
- If you do not see the SQL Server instance as the application database when creating a backup specification, enter the instance name yourself. When "not-named instance" is not displayed, insert the `DEFAULT` string.
- If Data Protector reports that the integration is properly configured, verify that the SQL Server user has appropriate rights to access the required databases.

During restore, if the following error occurred when executing an SQL statement:

```
Error message: "Microsoft SQL-DMO (ODBC SQLState: 01000)?15[152:5]
1646 [Microsoft][ODBC SQL Server Driver][SQL Server]The master database
has been successfully restored. Shutting down SQL Server.[Microsoft]
[ODBC SQL Server Driver] [SQL Server]SQL Server is terminating this
process."
```

note that this is the expected behavior when the master database is restored in a single user mode.

## Problems

### Problem

**The integration is properly configured but the database backup fails after a timeout**

- With an error similar to:

```
[Warning] From: OB2BAR@computer1.com "MSSQL70"
Time: 3/14/2000 8:19:22 PM
Error has occurred while executing SQL statement.
Error message: 'Microsoft SQL-DMO (ODBC SQLState: 42000) Error
number: bc5
[Microsoft][ODBC SQL Server Driver][SQL Server]Backup or restore
operation terminating abnormally.'
[Critical] From: OB2BAR@computer1.com "MSSQL70"
Time: 3/14/00 8:19:24 PM
Received ABORT request from SM => aborting
```

- SQL Server error log contains an entry similar to:

```
2000-03-14 20:19:21.62 kernel
BackupVirtualDeviceSet::Initialize: Open failure on backup
device 'Data_Protector_master'.
Operating system error -2147024891(Access is denied.).
```

- SQL Server **VDI.LOG** file contains an entry similar to:

  ```
  2000/03/15 13:19:31 pid(2112)
  Error at BuildSecurityAttributes: SetSecurityDescriptorDacl
  Status Code: 1338, x53A Explanation: The security descriptor

  structure is invalid.
  ```

SQL Server service and `Data Protector Inet` are running under different accounts. The integration cannot access SQL Server due to security problems.

### Action

Restart the `Data Protector Inet` service under the same account as the SQL Server service is running.

### Problem

**Backup fails if the appropriate drive letter on the backup system does not exist**

Backup fails with an error, similar to:

```
[Major] From: SSEA@computer1.com ""  Time: 02-Feb-08 14:07:54
Filesystem \\.\Volume{ef58fe0e-b2b8-11db-aa08-000802804af6} could not be mounted
to Q:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\DATA\.
([87] The parameter is incorrect. ).
```

Backup fails because SSE agent tries to mount the filesystem to the drive letter which does not exist on the backup system. The drive letter must be the same as on the application system. SSE or SMIS agents always mount the filesystems to the same drive letters on the backup systems as if the `ZDB_PRESERVE_MOUNTPOINTS` `omnirc` variable is set to `1`.

### Action

For source volumes to be successfully replicated, create the same drive letters on the backup system as they are used on the application system for mounting the source volumes.

### Problem

**Database is left in unrecovered state after “Invalid value specified for STOPAT parameter” is reported**

The database remains in an unrecovered state as if the `RESTORE LOG` operation was run with `Leave the database non-operational`.

### Action

Recover the database to the latest point in time using SQL Query Analyzer:

```
RESTORE DATABASE database_name WITH RECOVERY
```

After the recovery, additional transaction logs cannot be applied.

### Problem

**Database restore fails with “Error has occurred while executing SQL statement”**

If the database to be restored is still running, an error similar to the following may be reported:

```
[Warning] From: OB2BAR@taz "(DEFAULT)"  Time: 3/12/2001 7:20:28 PM

Error has occurred while executing SQL statement

Error message: 'Microsoft SQL-DMO (ODBC SQLState: 42000)
Error number: c1d

[Microsoft][ODBC SQL Server Driver][SQL Server]Exclusive access
could not be obtained because the database is in use.

[Microsoft][ODBC SQL Server Driver][SQL Server]RESTORE DATABASE
is terminating abnormally.'
```

### Action

To put the database offline, start SQL Server Enterprise Manager, go to the database and click **Action**. Select all tasks and put them offline.

### Problem

**Transaction logs cannot be restored from tape**

The recovery completed successfully and the database was put in `norecovery` state, but transaction logs cannot be restored from tape.

### Action

Recover the database to the state of ZDB to disk using the SQL Query Analyzer:

```
RESTORE DATABASE database name WITH RECOVERY
```

After the recovery, additional transaction logs cannot be applied.

### Problem

**Instant recovery of SQL Server databases fails**

### Action

Restart SQL Server instance services after instant recovery completes. If restarting services does not automatically start the recovery of all system databases, start the SQL Server instance in a single-user mode and manually start the recovery of the master database. Do the same with other system databases. At the end, restart SQL Server instance services.

### Problem

**Instant recovery completes with errors**

The following errors appear:

```
[Critical] From: OB2BAR_Main@qawin46.dpqaexch.com "(DEFAULT)"
Time: 4/8/2006 7:01:42 PM

Microsoft SQL Server reported the following error during
login:

The object was not open.

[Warning] From: OB2BAR_Main@qawin46.dpqaexch.com "(DEFAULT)"
Time: 4/8/2006 7:01:42 PM

[152:9208] Data Protector is probably not configured for use
with SQL Server on this host.
```

These errors are displayed because SQL Server is offline.

### Action

After the recovery, start SQL Server manually using SQL Server Enterprise Manager.

### Problem

**Restore to another client in the Data Protector cell not configured for use with SQL Server fails**

### Action

Configure the SQL integration on this client (see "Configuring the integration" on page 245).

## Problem

**Database is left in unrecovered state after restore completed successfully**

If you set the time for `Stop at` beyond the end of the `RESTORE LOG` operation, the database remains in the unrecovered state as if the `RESTORE LOG` operation was run with `Leave the database non-operational`.

## Action

Recover the database to the latest point in time by using the SQL Query Analyzer:

```
RESTORE DATABASE database_name WITH RECOVERY
```

After the recovery, additional transaction logs cannot be applied.

## Problem

**Instant recovery of a Microsoft SQL Server database configured on a Microsoft Cluster Server system fails with "The physical filename may be incorrect"**

Instant recovery of a Microsoft SQL Server split-mirror backup in a Business Copy XP configuration on a Microsoft Cluster Server system fails with the following error:

```
[Microsoft][ODBC SQL Server Driver][SQL Server]Device
activation error.
```

```
The physical file name '<Data/Log filename>' may be incorrect.
```

## Action

Perform the following steps:

1. Using Microsoft SQL Server Enterprise Manager, detach the Microsoft SQL Server database that you want to recover.
2. Using Cluster Administrator, take the Microsoft SQL Server Disk resource offline.
3. On the application system, configure the `ZDB_TAKE_CLUSRES_ONLINE` `omnirc` variable.

   For details, see "ZDB integrations omnirc variables" on page 425.
4. Start instant recovery.
5. When the message `Please, take MS SQL cluster resources online` appears in the Data Protector GUI, bring the Microsoft SQL Server cluster resources online using Cluster Administrator.

## Problem

**Restoring a Microsoft SQL Server 2005 instance to an alternate location when full-text indexing is enabled fails**

When the Use full-text indexing option is enabled for a particular database in a Microsoft SQL Server 2005 instance, the restore session does not complete successfully, since restore of the full-text catalog of the SQL database fails. The session report contains warning messages about the full-text catalog file being used by the affected database.

## Action

To solve the problem:

1. In the HP Data Protector Manager, switch to the **Restore** context.
2. In the Scoping Pane, expand **Restore Objects** and then **MS SQL Server**. Select name of the Microsoft SQL Server for which you want to perform restore.
3. In the Results Area, double-click the bar name corresponding to the particular Microsoft SQL Server instance. A list of backed up objects gets displayed.
4. Select the desired Microsoft SQL Server database, right-click it, and click **Properties**.
5. In the Properties window, click the **Advanced** tab.
6. Select the **Restore database with new name** option, and enter the new database name in the text box.
7. For all logical file names that are already present on the list, update contents of the Destination file name column accordingly.
8. Add the full-text catalog to the list.

   In the Logical file name text box, enter the string `sysft_`*`Full-Text_Catalog_Name`*. In the Destination file name text box, enter the corresponding physical location.

   NOTE:

   The full-text catalog is always restored to its original location, regardless of the specified physical location.

9. Click **Add/Set**.

10. In the Version and Options property pages, specify the appropriate options. For details, see “ Standard restore” on page 267.
11. Click **OK** to close the Properties window.
12. In the Options, Devices, and Media property pages, specify the appropriate options. For details, see “ Standard restore” on page 267.
13. Click **Restore** and then **Next** to select the Report level and Network load.
14. Click **Finish** to start the restore session.

# 4 Data Protector Microsoft Exchange Server ZDB integration

## Introduction

This chapter explains how to configure and use the Data Protector Microsoft Exchange Server ZDB integration. It describes the concepts and methods you need to understand to back up and restore the Microsoft Exchange Server (**Exchange Server**) database objects.

During backup, the database is only stopped for the time it takes to create a **replica** (split mirror disks or create snapshots). If ZDB to tape is performed, backup is subsequently performed offline on the backup system.

The following disk arrays and array configurations are supported:

| Supported array | Supported configurations |
|---|---|
| HP StorageWorks Disk Array XP (XP) | BC, CA, combined BC+CA |
| HP StorageWorks Enterprise Virtual Array (EVA) | BC, combined CA+BC |
| HP StorageWorks Virtual Array (VA) | BC |

All ZDB types (ZDB to tape, ZDB to disk, and ZDB to disk+tape) are supported by this integration. For ZDB types description, see the *HP Data Protector zero downtime backup concepts guide* and the online Help.

Using Data Protector, you can restore Exchange Server data:

- From backup media to the application system on LAN (standard restore).

- Using instant recovery.

The following table gives an overview of Exchange Server recovery methods:

| | **Rollforward recovery** | **Point-in-time recovery** |
|---|---|---|
| **ZDB to tape** | Standard restore (database) + standard restore (transaction logs) | Standard restore (database) |
| **ZDB to disk** | Instant recovery (database) + standard restore (transaction logs) | Instant recovery (database) |
| **ZDB to disk+tape** | • Instant recovery (database) + standard restore (transaction logs)<br>• Standard restore (database) + standard restore (transaction logs) | • Instant recovery (database)<br>• Standard restore (database) |

For a description of ZDB and instant recovery concepts, see the *HP Data Protector zero downtime backup concepts guide*.

# Integration concepts

The Exchange ZDB integration backs up Microsoft Information Store (MIS), Key Management Service (KMS), and Site Replication Service (SRS), which are replicated within a disk array.

Data Protector integrates with Exchange Server through:

- The Data Protector `ese_bar.exe` executable, installed on Exchange Server and used for retrieving locations of files and folders pertaining to a storage group/store.
- The `omniEx2000.exe` command, responsible for mounting and dismounting databases and purging transaction logs after a transaction logs backup.

An Exchange ZDB specification is created using the `omnicreatedl` command. Based on this specification, you can perform a filesystem ZDB using the Data Protector GUI. During backup, `ese_bar.exe` resolves database objects and retrieves a list of files and folders to be backed up. The following happens:

- `Stop/quiesce application` is executed, which dismounts the databases and verifies their consistency using `esefile.exe` (Exchange 2000) or `eseutil.exe` (Exchange 2003). See Checking Exchange files for consistency.
- When the replica is created, `Restart the application` is executed, which mounts dismounted databases.
- If ZDB to tape is performed, Data Protector backs up Exchange databases on the backup system. Transaction logs are not deleted and can be backed up in a separate session using non-ZDB procedure.

ZDB process depends on whether you replicate your data in a split mirror or snapshot replication and on the selected ZDB type. Restore process depends on the restore type - standard restore or instant recovery. See the *HP Data Protector zero downtime backup concepts guide* for a detailed description of replication techniques, as well as ZDB and restore processes.

# Configuring the integration

## Prerequisites

- You need a license to use the Exchange Server ZDB integration. See the *HP Data Protector installation and licensing guide* for information.
- Ensure that you correctly installed and configured Exchange Server.
    - For supported versions, platforms, devices, and other information, see the latest support matrices at http://www.hp.com/support/manuals.
    - For information on installing, configuring, and using Exchange Server, see the Exchange Server documentation.
- Ensure that you correctly installed Data Protector. For information on installing Data Protector in various architectures and installing a Data Protector disk array integration (XP, VA, or EVA) with Exchange Server, see the *HP Data Protector installation and licensing guide*.

  Every Exchange Server system to be used with Data Protector must have the `MS Exchange Integration` component installed.
- Install Exchange Server on the application system. All parts of the Exchange database (Information Store (MIS), Key Management Service (KMS), and Site Replication Service (SRS)) must be installed on the disk array source volumes.
- Split mirror backup of Exchange Server 2003 databases requires Exchange Server 2003 SP2 or higher.

## Before you begin

- Configure devices and media for use with Data Protector. See the online Help index: "configuring devices" and "creating media pools" for instructions.

  For ZDB to disk, configure a backup device (for example, a standalone file device), as you will have to select it when configuring a backup specification. For information, see the online Help index: "standalone devices".

- To test whether Exchange Server and Cell Manager communicate properly, configure and run a Data Protector filesystem ZDB and restore. See the online Help for instructions.
- Before performing transaction logs backups, disable **circular logging** for all storage groups.

  If the application is cluster-aware, disable circular logging on all cluster nodes.

- Add the `*Exchange_home*\bin` directory to the Windows **Path** environment variable:

  1. In the Windows Explorer, right-click **My Computer** and click **Properties**.
  2. In the **Properties** dialog box, click **Advanced** and then **Environment Variables**.
  3. Select **Path** in the **System Variables** list and click **Edit**.
  4. Add `*Exchange_home*\bin` in the **Variable Value** text box and click **OK**.

  See Figure 82.

  If the integration is cluster-aware, perform this procedure on all cluster nodes.

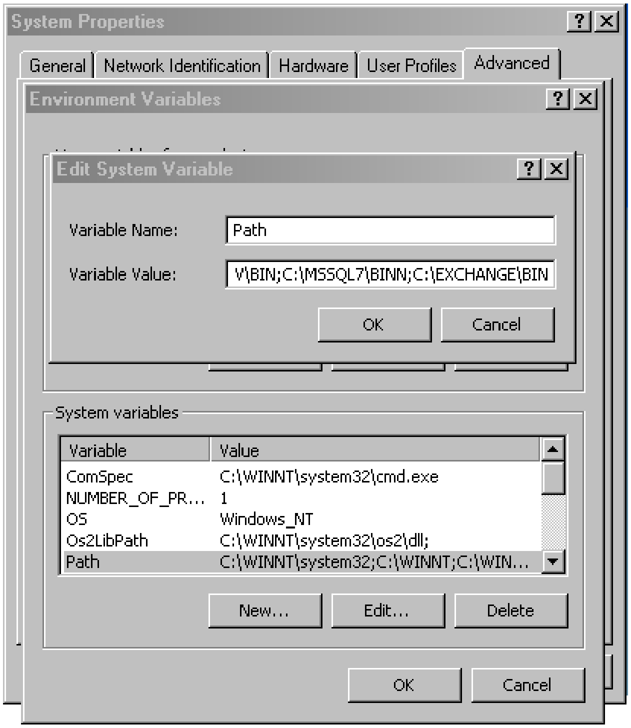

Variable Name: Path

Variable Value: V\BIN;C:\MSSQL7\BINN;C:\EXCHANGE\BIN

| Variable | Value |
|---|---|
| ComSpec | C:\WINNT\system32\cmd.exe |
| NUMBER_OF_PR... | 1 |
| OS | Windows_NT |
| Os2LibPath | C:\WINNT\system32\os2\dll; |
| Path | C:\WINNT\system32;C:\WINNT;C:\WIN... |

**Figure 82 Path system variable**

- In cluster environments, prior to configuring a ZDB specification, create a new physical resource (LUN for VA, virtual disk for EVA, and LDEV for XP) for the cluster and move the Exchange Server database and transaction log files to it.

  For information on moving the Exchange Server database and transaction log files, see the following documents available at http://support.microsoft.com:

  - *How to Move Exchange Databases and Logs in Exchange Server 2003 (821915)*
  - *XADM: How to Move Exchange Databases and Logs in Exchange 2000 Server (257184)*
  - *HOW TO: Add New Mailbox Stores in Exchange Server 2003 (821748)*
  - *HOW TO: Add New Mailbox Stores in Exchange 2000 (319218)*

# Backup

To run ZDB of an existing Exchange Server ZDB specification:

- Schedule a backup using the Data Protector Scheduler.
- Start an interactive backup using the Data Protector GUI or CLI.

### Limitations

- Backup preview is not supported.

### Considerations

- You can perform transaction logs backups only if circular logging is disabled for the involved Exchange Server.

  Circular logging is a Microsoft Exchange mode, where transaction logs are automatically overwritten when the data they contain is committed to the database.

  If enabled, this option reduces disk storage space requirements, but does not allow you to perform transaction logs backups.
- Exchange Server on the application system must be running. Otherwise, data consistency is not guaranteed.
- You cannot start ZDB, restore, or instant recovery using the same source volume on the application system at the same time. A session must be started only after the preceding session using the same source volume on the application system finishes; otherwise, the session fails.
- Transaction logs must be backed up using the common filesystem backup functionality.
- When backing up log files, select either **Log all** or **Log files** logging level in the **Advanced Filesystem Options** dialog box. This enables purging backed up log files from disk. For more information, see the online Help index: “filesystem options”.

## Configuring Exchange Server ZDB

To configure a ZDB:

1. Configure devices and media for backup.
2. Create a Data Protector Exchange Server ZDB specification.

### Creating ZDB specifications

Create a ZDB specification using the `omnicreatedl` command. After that, perform a filesystem ZDB using the Data Protector GUI.

The `omnicreatedl` command creates a ZDB specification with included `Stop/quiesce the application` and `Restart the application` scripts (`omniEx2000.exe`) for mounting/dismounting backed up databases and checking their consistency.

Additionally, `omnicreatedl` creates a transaction logs backup specification for each storage group specified in a ZDB specification (circular logging disabled). See Figure 98. The transaction logs backup specification includes the post-exec script for purging backed up log files.

TIP:

A transaction logs backup specification can either be triggered by post-exec (configured on the backup specification level) in the backup specification for database files backup (recommended), or started manually after the backup specification for database files backup is started. For more information, see the online Help index: “pre- and post-exec commands“.

The following is the synopsis of the `omnicreatedl` command:

```
omnicreatedl -ex2000 -datalistName [-deviceName]
{DISK_ARRAY_XP_OPTIONS | VIRTUAL_ARRAY_OPTIONS |
ENTERPRISE_VIRTUAL_ARRAY_OPTIONS} EXCHANGE_OPTIONS [-force]
[-virtualSrvName]
```

For a detailed synopsis of *DISK_ARRAY_XP_OPTIONS*, *VIRTUAL_ARRAY_OPTIONS*, *ENTERPRISE_VIRTUAL_ARRAY_OPTIONS*, and *EXCHANGE_OPTIONS*, see the `omnicreatedl` man page.

Descriptions of the parameters are presented in the tables below:

IMPORTANT:

If parameters contain spaces, use double quotes when specifying them in the `omnicreatedl` command. For example, `-storage_group "First Storage Group"`.

### Table 18 General options

| Parameter | Description |
|---|---|
| `-ex2000` | Instructs `omnicreatedl` to create a ZDB specification and transaction logs backup specification for all specified storage groups (circular logging disabled). |
| `-datalist` *Name* | Specifies the name of the ZDB specification created on the Cell Manager in *Data_Protector_home*`\config\server\datalists` (Windows) or `/etc/opt/omni/server/datalists` (UNIX). |
| `-device` *Name* | Specifies the backup device. If this option is not used, specify the backup device using the Data Protector GUI. |
| `-force` | Forces overwriting of an existing ZDB specification with the same name. |
| `-virtualSrv` *Name* | In cluster configurations, specifies the name of the Exchange Server virtual server. |

### Table 19 XP options

| Parameter | Description |
|---|---|
| `-split_mirror -sse` | Instructs `omnicreatedl` to create a split mirror ZDB specification. |

| Parameter | Description |
|---|---|
| `( -local` *`app_sys bck_sys`* `\| -remote` *`app_sys bck_sys`* `\| -combined` *`app_sys bck_sys`* `)` | Specifies XP configurations:<br>`-local` selects Business Copy (BC) XP.<br>`-remote` selects Continuous Access (CA) XP.<br>`-combined` selects Business Copy + Continuous Access (BC+CA) XP.<br>For all configurations, *`app_sys`* sets the application system and *`bck_sys`* sets the backup system.<br>In cluster environments, specify the virtual server hostname (rather than the physical node hostname). |
| `-mirrors` *`MU_Numbers`* | Sets the number of replicas in the replica set. Enter an integer number from 0 to 2 or any range/combination of integer numbers from 0 to 2 separated by a comma, for example:<br>`1`<br>`1-2`<br>`2,0,1`<br>A sequence does not set the order in which the replicas are used. The algorithm of using replicas is described in the *HP Data Protector zero downtime backup concepts guide*.<br>A range must be specified in ascending order.<br>If this option is not specified, MU# 0 is set. |
| `[ -split \| -establish ]` | This parameter is optional.<br>If `-split` is selected, mirrored disks in the replica are resynchronized with P-VOLs at the backup start.<br>If `-establish` is selected and the replica for the next backup is not synchronized, a resync is initiated before next backup.<br>Default: `-establish`. |
| `-leave_enabled_bs` | Available if `-keep_version` is specified.<br>By default, Data Protector dismounts filesystems on the backup system after each backup. If this option is specified, filesystems remain mounted after backup.<br>Thus, you can use the backup system for data warehouse activity, but not for instant recovery. |

| Parameter | Description |
|---|---|
| -instant_restore | This parameter is optional.<br>Select to perform ZDB to disk or ZDB to disk+tape and leave the replica on a disk array for instant recovery. If this option is not set, you cannot perform instant recovery from the replica created or reused in this session.<br>When -instant_restore is selected, omnicreatedl automatically sets the -keep_version option. |
| -keep_version | This parameter is optional.<br>If selected, participating pairs remain split after backup, enabling you to restore directly from the replica.<br>If not selected, participating disks are resynchronized after backup if one or no replica is set by -mirrors *MU_Numbers*. If more than one replica is set, the disks remain split.<br>If this option is not specified, you cannot specify -leave_enabled_bs. |

**Table 20 VA options**

| Parameter | Description |
|---|---|
| -snapshot | Instructs omnicreatedl to create a snapshot ZDB specification. |
| -va *app_sys bck_sys* | Specifies the application system *app_sys* and the backup system *bck_sys*. |
| -use_existing_snapshot | By default, this option is automatically selected if -instant_recoveryis set.<br>If configuring a ZDB to tape, select this option to reuse an existing replica.<br>The replica can be reused only if it already exists on a disk array. Only replicas not marked for instant recovery or replicas with snapshots not listed in the VA LUN exclude file can be reused. |

| Parameter | Description |
|---|---|
| -instant_recovery | This parameter is optional.<br>Select to perform a ZDB to disk or ZDB to disk+tape and leave the replica on a disk array for instant recovery. Also, specify -snapshots *number*. The options -use_existing_snapshotand -leave_versionare automatically selected.<br>If this option is not set, you cannot perform instant recovery from the replica created or reused in this session. |
| -snapshots *number* | This parameter is optional.<br>During ZDB, Data Protector creates a new replica and leaves it on the array until the specified -snapshots *number* is reached (specify if you selected -instant_recovery). After that, the oldest replica is reused.<br>Default: 1. Maximum: 1024. |
| -leave_version | This parameter is optional.<br>By default, this option is automatically selected if -instant_recovery is set.<br>For ZDB to tape, select it to keep the replica on a disk array after backup. This replica is not available for instant recovery, but can be reused for future backups using the same backup specification (-use_existing_snapshotselected).<br>If this option is not selected, the replica is deleted after backup. |
| -leave_enabled_bs | Available if -leave_versionis selected.<br>By default, Data Protector dismounts filesystems on the backup system after each backup. If this option is specified, filesystems remain mounted.<br>Thus, you can use the backup system for data warehouse activities, but not for instant recovery. |

| Parameter | Description |
|---|---|
| `-lun_security` | Specify this option to apply LUN security to child LUNs (target volumes or snapshots) created by the integration.<br>*If Secure Manager is activated, specify this option and configure passwords correctly; otherwise, the backup fails.* |

### Table 21 EVA options

| Parameter | Description |
|---|---|
| `-snapshot` | Instructs `omnicreatedl` to create a snapshot ZDB specification. |
| `-smis` *`app_sys bck_sys`* | Specifies the application system *`app_sys`* and the backup system *`bck_sys`*. |
| `-instant_recovery` | This parameter is optional.<br>Select to perform a ZDB to disk or ZDB to disk+tape and leave the replica on a disk array for instant recovery. Also, specify `-snapshots` *`number`*. The options `-snapshot_type clone` and `-snapshot_policy strict` are automatically selected.<br>If this option is not set, you cannot perform instant recovery from the replica created or reused in this session. |
| `-snapshots` *`number`* | This parameter is optional.<br>During ZDB, Data Protector creates a new replica and leaves it on the array until the specified `-snapshots` *`number`* is reached. After that, the oldest replica is reused.<br>Default: 1.<br>The maximum number for vsnaps and standard snapshots is 7. Data Protector does not limit the number of replicas rotated, but the session fails if the limit is exceeded. |
| `-snapshot_type {standard \| vsnap \| clone}` | Instructs Data Protector to create a particular type of EVA snapshots:<br>-`standard` creates snapshots with the pre-allocation of disk space.<br>-`vsnap` creates snapshots without the pre-allocation of disk space.<br>-`clone` creates a clone of a source volume (original virtual disk). |

| Parameter | Description |
|---|---|
| `-snapshot_policy {strict \| loose}` | Specifies a snapshot policy depending on the type of already existing snapshots for the same source volume.<br>When `-strict` is selected, Data Protector attempts to create snapshots as specified by `-snapshot_type`. If the source volumes used in the session have existing snapshots of a different type, the selected type is not created (the session is aborted).<br>When `-loose` is selected, Data Protector creates snapshots of a different type than specified by `-snapshot_type` (if this helps to complete a session successfully).<br>For example, if you select standard snapshots, but Data Protector detects that standard snapshots cannot be created because vsnaps/snapclones of the source volumes already exist in a replica set, it creates either vsnaps or snapclones instead of standard snapshots.<br>Note that Data Protector can use only one type of snapshots in a backup session. For example, if the source volumes used in a session have existing standard snapshots/vsnaps, the session is aborted. |
| `-wait_clonecopy` *`number`* | Available if `-snapshot_type clone` is selected.<br>Specify to prevent degradation of the application data access times by delaying moving data to tape until the cloning process completes (ZDB to tape, ZDB to disk+tape). Set the maximum waiting time. When the specified time is reached, backup to tape starts (even if cloning is not finished). |
| `-replica_conf {local \| combined}` | Specifies EVA configurations:<br>`local` selects Business Copy (BC) EVA.<br>`combined` selects Continuous Access + Business Copy (CA+BC) EVA. |

| Parameter | Description |
|---|---|
| -ca_failover_option {follow_replica_direction \| maintain_replica_location} | Available if combined configuration is selected.<br>Specify to control the replication direction after a failover.<br>Select follow_replica_direction to follow the replication direction and create replicas on the array remote to current source. A failover reverses the replication direction and the replicas are created on the array that was originally a source EVA.<br>Select maintain_replica_location to maintain the replica location and create replicas on the array remote to home. After a failover, replicas continue to be created on the destination array that has also become a source EVA.<br>When -ca_failover_option is selected, follow_replica_direction is set as default. |

### Table 22 Exchange Server options

| Parameter | Description |
|---|---|
| -annotation { MIS \| SRS \| KMS } | Specifies Exchange Server annotations: Microsoft Information Store (MIS), Site Replication Service (SRS), and Key Management Service (KMS).<br>Default: MIS. |
| -all_storage_groups | Creates a backup specification for all databases relating to MIS (specified by -annotation MIS). |
| -storage_group *Storage_Group_Name* | Creates a backup specification for all stores relating to the specified. Multiple declarations of -storage_groupcreate a backup specification for the selected storage groups. |
| -store *Store* | Creates a backup specification only for specified store(s) inside the storage group. To create a backup specification for many stores, specify a list of stores after the -store parameter.<br>You can obtain store names using Exchange System Administrator. |

For more information, see the omnicreatedl man page.

TIP:

To create an Exchange ZDB specification and a transaction logs backup specification, you can also use the `omniex2000SM.bat` file, located in the *Data_Protector_home*`\bin` directory on the application system. The script provides templates and examples of `omnicreatedl` usage. You can modify `omniex2000SM.bat` using any text editor. Uncomment (delete `@REM` before the command) the line with the appropriate command and edit parameters. See Figure 83.

To create an Exchange ZDB specification and a transaction logs backup specification using `omniex2000SM.bat`, make the necessary modifications and run: *Data_Protector_home*`\bin\omniex2000SM.bat`.

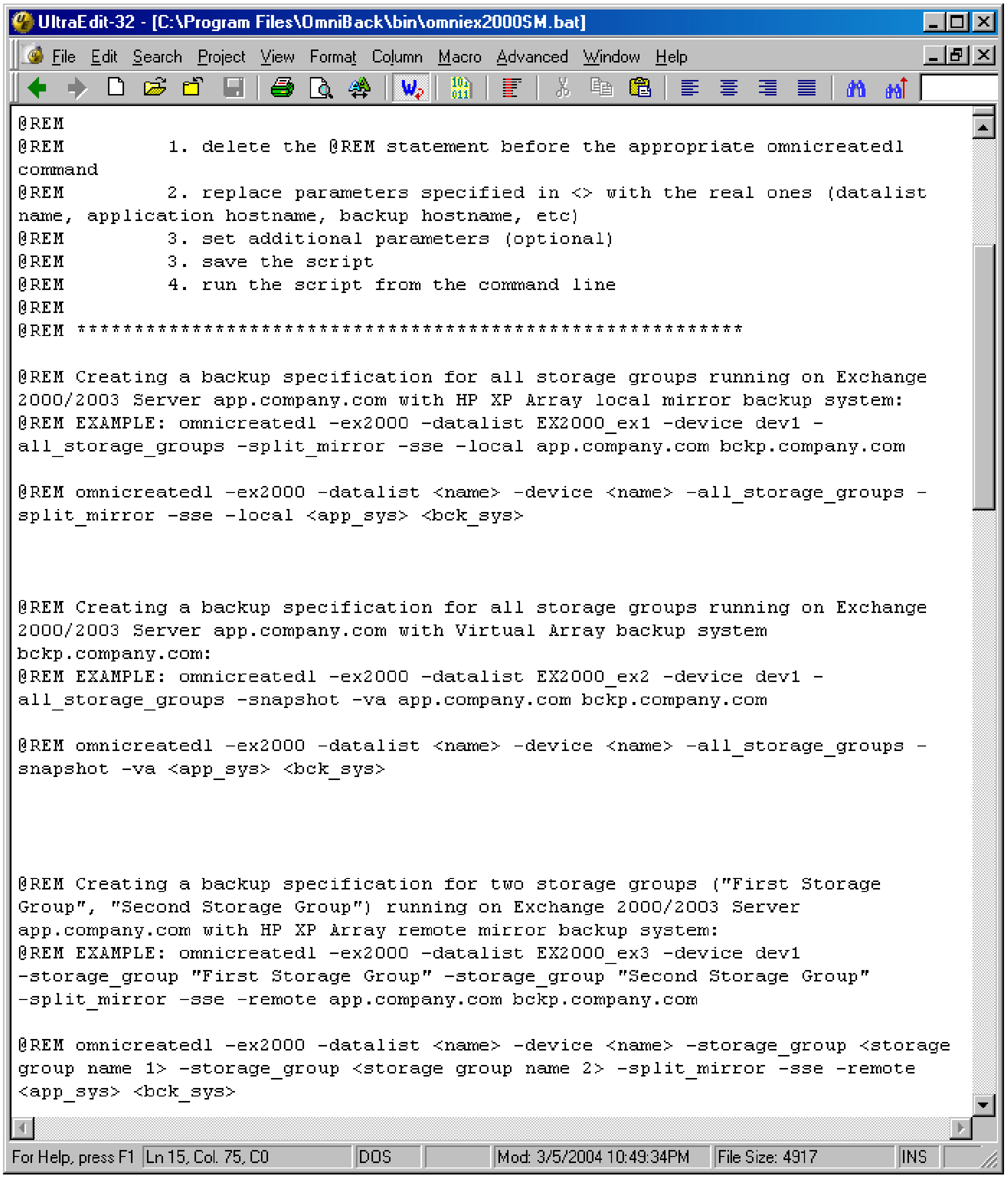

UltraEdit-32 - [C:\Program Files\OmniBack\bin\omniex2000SM.bat]
File Edit Search Project View Format Column Macro Advanced Window Help
For Help, press F1
Ln 15, Col. 75, C0
DOS
Mod: 3/5/2004 10:49:34PM
File Size: 4917
INS

```
@REM
@REM          1. delete the @REM statement before the appropriate omnicreatedl
command
@REM          2. replace parameters specified in <> with the real ones (datalist
name, application hostname, backup hostname, etc)
@REM          3. set additional parameters (optional)
@REM          3. save the script
@REM          4. run the script from the command line
@REM
@REM **********************************************************

@REM Creating a backup specification for all storage groups running on Exchange
2000/2003 Server app.company.com with HP XP Array local mirror backup system:
@REM EXAMPLE: omnicreatedl -ex2000 -datalist EX2000_ex1 -device dev1 -
all_storage_groups -split_mirror -sse -local app.company.com bckp.company.com

@REM omnicreatedl -ex2000 -datalist <name> -device <name> -all_storage_groups -
split_mirror -sse -local <app_sys> <bck_sys>



@REM Creating a backup specification for all storage groups running on Exchange
2000/2003 Server app.company.com with Virtual Array backup system
bckp.company.com:
@REM EXAMPLE: omnicreatedl -ex2000 -datalist EX2000_ex2 -device dev1 -
all_storage_groups -snapshot -va app.company.com bckp.company.com

@REM omnicreatedl -ex2000 -datalist <name> -device <name> -all_storage_groups -
snapshot -va <app_sys> <bck_sys>




@REM Creating a backup specification for two storage groups ("First Storage
Group", "Second Storage Group") running on Exchange 2000/2003 Server
app.company.com with HP XP Array remote mirror backup system:
@REM EXAMPLE: omnicreatedl -ex2000 -datalist EX2000_ex3 -device dev1
-storage_group "First Storage Group" -storage_group "Second Storage Group"
-split_mirror -sse -remote app.company.com bckp.company.com

@REM omnicreatedl -ex2000 -datalist <name> -device <name> -storage_group <storage
group name 1> -storage_group <storage group name 2> -split_mirror -sse -remote
<app_sys> <bck_sys>
```

Figure 83 omniex2000SM.bat file

## Examples of using `omnicreatedl`

***Example 1 - XP:***

To create a ZDB specification BS1, using device dev1, for all storage groups running on Exchange Server on the application system computer_app.company.com and backup system computer_bck.company.com using the BC configuration, run:

```
omnicreatedl -ex2000 -datalist BS1 -device dev1
-all_storage_groups -split_mirror -sse -local
computer_app.company.com computer_bck.company.com
```

***Example 2 - XP:***

To create a ZDB specification BS2, using device dev1, for two storage groups (First Storage Group and Second Storage Group) running on Exchange Server on the application system computer_app.company.com and backup system computer_bck.company.com using the CA configuration, run:

```
omnicreatedl -ex2000 -datalist BS2 -device dev1 -storage_group
"First Storage Group" -storage_group "Second Storage Group"
-split_mirror -sse -remote computer_app.company.com
computer_bck.company.com
```

***Example 3 - XP:***

To create a ZDB specification BS3, using device dev1, for two stores (Public and Mailbox, part of First Storage Group) running on Exchange Server on the application system computer_app.company.com and backup system computer_bck.company.com using the BC+CA configuration, run:

```
omnicreatedl -ex2000 -datalist BS3 -device dev1 -storage_group
"First Storage Group" -store "Public" "Mailbox" -split_mirror
-sse -combined computer_app.company.com
computer_bck.company.com
```

***Example 4 - VA:***

To create a ZDB-to-tape specification BS1, using the backup device dev1, for all storage groups running on Exchange Server on the application system computer1.company.com and backup system computer2.company.com, run:

```
omnicreatedl -ex2000 -datalist BS1 -device dev1 -snapshot -va
computer1.company.com computer2.company.com -leave_version
-storage_group "First Storage Group" -force
```

If the backup specification already exists, it will be overwritten. If it does not exist, omnicreatedl creates a transaction logs backup specification file First Storage Group (LOGS) computer1.company.com for First Storage Group log files backup. The replica is kept for future (non-instant recovery) use.

***Example 5 - VA:***

To create a ZDB-to-disk+tape or ZDB-to-disk specification `Exchange_example` to back up Site Replication Service to the backup device `dev1`, using the replica set with 5 replicas, run:

```
omnicreatedl -ex2000 -datalist Exchange_example -device dev1
-snapshot -va computer1.company.com computer2.company.com
-instant_recovery -snapshots 5 -annotation SRS
```

The `omnicreatedl` creates a transaction logs backup specification file `SRS (LOGS) computer1.company.com` for Site Replication Service log files backup (circular logging disabled). When `omnib` or Data Protector GUI is used to start backup, select ZDB to disk or ZDB to disk+tape.

***Example 6 - EVA:***

To create a ZDB-to-tape specification `BS1`, using the backup device `dev1`, for all storage groups running on Exchange Server on the application system `computer1.company.com` and backup system `computer2.company.com`, run:

```
omnicreatedl -ex2000 -datalist BS1 -device dev1 -snapshot
-smis computer_app.company.com computer_bck.company.com
-snapshot_type vsnap -snapshot_policy strict -storage_group
"First Storage Group"
```

The `omnicreatedl` creates a transaction logs backup specification file `First Storage Group (LOGS) computer1.company.com` for First Storage Group log files backup (if it does not already exist). Data Protector attempts to create vsnaps, and if they cannot be created, the session aborts.

***Example 7 - EVA:***

To create a ZDB-to-disk specification `Exchange_example` to back up Site Replication Service using the backup device `dev1`, using the replica set with 5 replicas, run:

```
omnicreatedl -ex2000 -datalist Exchange_example -device dev1
-snapshot -smis computer1.company.com computer2.company.com
-instant_recovery -snapshots 5 -annotation SRS
```

The `omnicreatedl` creates a transaction logs backup specification file `SRS (LOGS) computer1.company.com` for Site Replication Service log files backup (circular logging disabled). When `omnib` or Data Protector GUI is used to start backup, select ZDB to disk.

***Example 8 - EVA:***

To create a ZDB-to-disk+tape specification `Exchange_example` to back up Site Replication Service to the backup device `dev1`, using the replica set with 3 replicas and delay backup to tape for 50 minutes, run:

```
omnicreatedl -ex2000 -datalist Exchange_example -device dev1
-snapshot -smis computer1.company.com computer2.company.com
-instant_recovery -snapshots 3 -wait_clonecopy 50 -annotation
SRS
```

The `omnicreatedl` creates a transaction logs backup specification file `SRS (LOGS) computer1.company.com` for Site Replication Service log files backup (circular logging disabled). When `omnib` or Data Protector GUI is used to start backup, select ZDB to disk+tape.

## Modifying ZDB specifications

After you created an Exchange ZDB specification and a transaction logs backup specification using `omnicreatedl`, you can modify them using the Data Protector GUI.

NOTE:

To add/remove storage groups from a ZDB specification, create a new backup specification rather than modify the existing one. Changing the saved ZDB specification manually can impact `Stop/quiesce the application` and `Restart the application` parameters, which are automatically defined when a new backup specification is created using `omnicreatedl`.

Proceed as follows using the Data Protector Manager:

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

2. In the Scoping Pane, expand **Backup Specifications**, then **Filesystem**, and click a ZDB specification. See Figure 84.

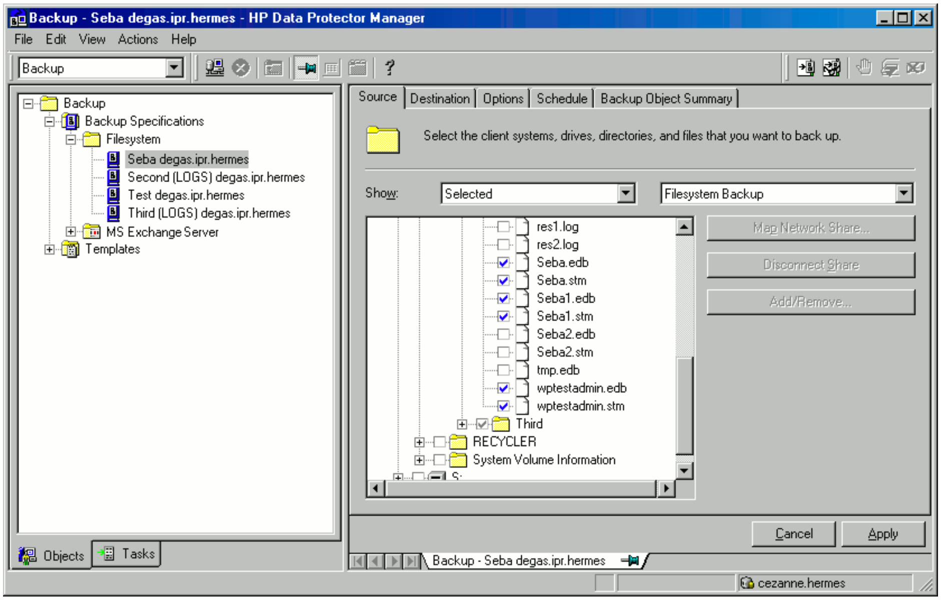

Select the client systems, drives, directories, and files that you want to back up.

Show: Selected

Filesystem Backup

**Figure 84 Exchange ZDB specification**

3. Modify backup devices, set ZDB options, schedule, and start backup. For a particular database, back up `.edb` and `.stm` files.

   To create additional backup copies (mirrors), specify the desired number by clicking **Add mirror**/**Remove mirror** under the **Destination** tag. Select separate devices for backup and each mirror.

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

NOTE:

Object mirroring is not supported for ZDB to disk.

## Checking Exchange files for consistency

Exchange Server allows you to perform a page level integrity check of the database files. It is run on the Exchange offline database during ZDB using `esefile` (Exchange 2000) and `eseutil` (Exchange 2003) utilities. For more information, see *Esefile Support Utility for Exchange Server 5.5 and Exchange 2000 Server (248406)* and

*How To: Use the Eseutil Utility to Detect File Header Damage in Exchange 2003 (825088)* available at http://support.microsoft.com.

IMPORTANT:

Page level integrity check is not supported on EVA.

Proceed as follows:

1. Copy the utility to the backup system.

   The `esefile` resides on the Exchange 2000 Installation CD at: `*CD*:\ENGLISH\EXCH2000\ENT\SUPPORT\UTILS\I386\ESEFILE`

   The `eseutil` resides in `*Exchange_Server_home*\bin` once the Exchange Server 2003 is installed.

   NOTE:

   On XP, do not mirror the location where the utility resides, to prevent the file from being overwritten when the mirror is synchronized.

2. Write a script that runs `esefile.exe` or `eseutil.exe` and save it, for example, as `integritycheck.bat`. Basically, the command looks like: *`Path_to_esefile`*`\esefile.exe /s` *`Exchange_Server_database_file`* (Exchange 2000) or *`Path_to_eseutil`*`\eseutil.exe /k` *`Exchange_Server_database_file`* (Exchange 2003) The `esefile.exe` utility with the `/s` option or `eseutil.exe` with the `/k` option tests the checksums on the Exchange Server database.

### Example

The `esefile.exe` resides in `C:\` directory on the backup system. Data Protector mounts the backup data to `C:\Program Files\Omniback\tmp` with the mountpoints

`zdb.hp.com\H\First\MailStore.edb`,

`zdb.hp.com\H\First\MailStore2.edb`, and

`zdb.hp.com\H\First\Public.edb`.

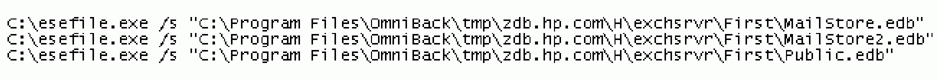

```
C:\esefile.exe /s "C:\Program Files\OmniBack\tmp\zdb.hp.com\H\exchsrvr\First\MailStore.edb"
C:\esefile.exe /s "C:\Program Files\OmniBack\tmp\zdb.hp.com\H\exchsrvr\First\MailStore2.edb"
C:\esefile.exe /s "C:\Program Files\OmniBack\tmp\zdb.hp.com\H\exchsrvr\First\Public.edb"
```

**Figure 85 Example script for MS Exchange Server 2000**

See the *HP Data Protector zero downtime backup administrator's guide* for information on the mountpoint creation. The commands in the script ensure that `*.edb` database files specified in the backup specification are checked for consistency.

3. Modify the ZDB specification to include `integritycheck.bat` as post-exec. Execute the script on the backup system. See the online Help index: “pre- and post-exec commands” for more information.
4. In the backup specification, select `Leave the backup system enabled`. This ensures that the database is available on the backup system when the integrity check is started.

After backup, `esefile` or `eseutil` perform page level integrity check against database files specified in `integritycheck.bat`.

## Scheduling backups

You can run unattended ZDB at specific times or periodically. For details on scheduling, see the online Help index: “scheduled backups”.

NOTE:

You cannot run ZDB to disk or ZDB to disk+tape if `-instant_restore` (XP) or `-instant_recovery` (VA, EVA) is not selected in the backup specification.

### Scheduling example

To schedule a database ZDB at 8:00, 13:00, and 18:00 during week days:

1. In the **Schedule** property page, select the starting date in the calendar and click **Add** to open the **Schedule Backup** dialog box.
2. Under **Recurring**, select **Weekly**. Under **Time options**, select **8:00**. Under **Recurring Options**, select **Mon**, **Tue**, **Wed**, **Thu**, and **Fri**.

   Click **OK**.
3. Repeat Step 1 and Step 2 to schedule backups at 13:00 and 18:00.
4. Click **Apply** to save the changes.

NOTE:

For ZDB sessions, the backup type is set to **Full**.

For ZDB to disk and ZDB to disk+tape, specify **Split mirror/snapshot backup**.

## Starting backup sessions

Interactive backups are run on demand. They are useful for urgent backups or restarting failed backups.

NOTE:

You cannot run ZDB to disk or ZDB to disk+tape if `-instant_restore` (XP) or `-instant_recovery` (VA, EVA) is not selected in the backup specification.

## Using the Data Protector GUI

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **Filesystem**. Right-click the backup specification you want to start and select **Start Backup**.
3. Select **Network load**. For information on network load, click **Help**. Click **OK**.

   For ZDB sessions, the backup type is set to **Full**.

   For ZDB to disk or ZDB to disk+tape, specify the **Split mirror/snapshot backup** option.

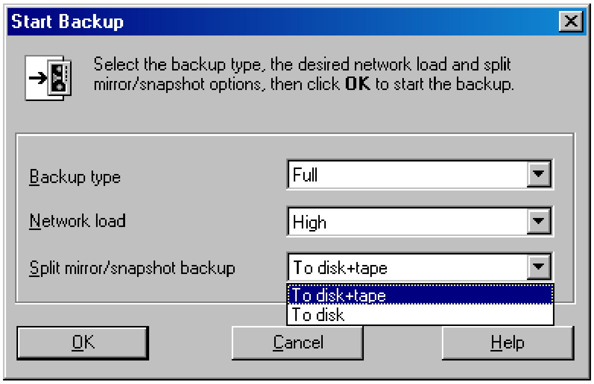

Start Backup

Select the backup type, the desired network load and split mirror/snapshot options, then click **OK** to start the backup.

Backup type: Full

Network load: High

Split mirror/snapshot backup: To disk+tape

**Figure 86 Starting interactive backups**

Click **OK**.

## Using the Data Protector CLI

To start ZDB to tape or ZDB to disk+tape, run:

```
omnib -datalist Name
```

To start ZDB to disk, run:

```
omnib -datalist Name -disk_only
```

where `Name` is the backup specification name. For more information on `omnib`, see its man page.

# Restore

Data Protector offers restore from backup media to the application system on LAN (standard restore), where you can select among various restore options depending on your restore scenario, and instant recovery. See the below sections for more information.

### Considerations

- Instant recovery restores data from a replica on the backup system to the source volumes on the application system. Therefore, you cannot *selectively* restore objects (storage groups, stores, or Exchange Server) that do not reside on separate source volumes.
- Transaction logs must be backed up to perform rollforward recovery.

NOTE:

In case of a disaster, Exchange Server must be installed and configured with the same database names and locations as before the disaster.

## Standard restore

You can restore Exchange objects to any Exchange Server with the same configuration as the original system within the Data Protector cell.

You can recover Exchange databases using:

- Point-in-time recovery

  Only data files (`.edb` and `.stm`) are restored. Databases are restored to their state at the backup time, and all data modified after that is lost.

- Rollforward recovery

  Restores Exchange database files and transaction logs, and then replays the transaction logs. Databases are recovered to their last consistent state.

TIP:

You can improve data transfer rate by configuring a backup device on the application system. See the online Help index: “configuring, backup devices” for details. You can also perform a restore using another device. See the online Help index “select, devices for restore”.

You can perform a filesystem restore of Exchange databases (`.edb` and `.stm` files) and transaction logs (`.log` files); however, to recover the database, you need to perform some additional steps. For details, see *Offline Backup and Restoration Procedures for Exchange (296788)* at http://support.microsoft.com.

For a detailed procedure on performing filesystem restore, see the online Help index: “standard restore procedure”. The procedures below provide only a general description of the restore process.

## Point-in-time recovery

IMPORTANT:

The procedure described below gives instructions on how to copy `.edb` and `.stm` files to the appropriate database location and is only part of the procedure described in *Offline Backup and Restoration Procedures for Exchange (296788)*.

Proceed as follows using the Data Protector Manager:

1. In the Context List, select **Restore**.

2. In the Scoping Pane, expand **Restore Objects**, **Filesystem**, and then select the client (backup system). Mark the drive letter on which the database resides.

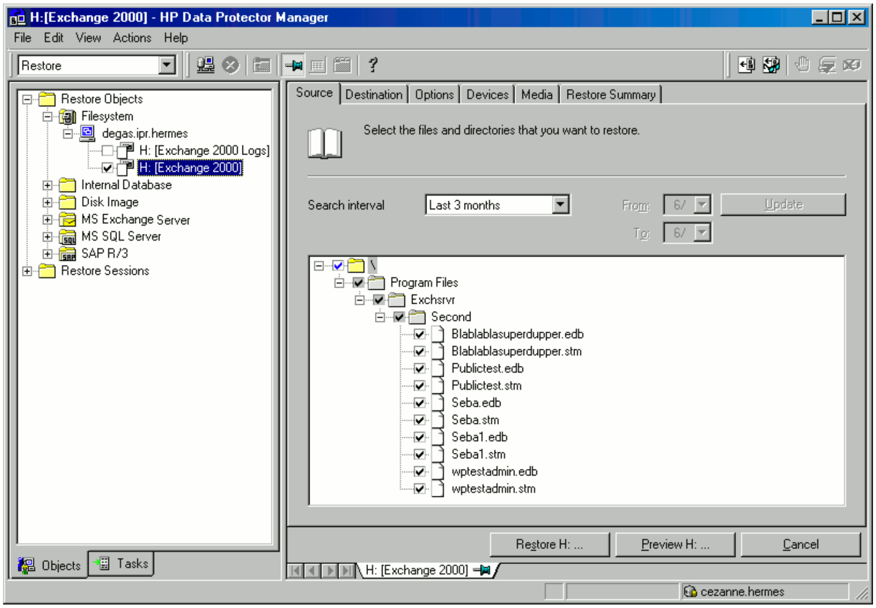

H:[Exchange 2000] - HP Data Protector Manager

- Restore Objects
  - Filesystem
    - degas.ipr.hermes
      - □ H: [Exchange 2000 Logs]
      - ☑ H: [Exchange 2000]
  - Internal Database
  - Disk Image
  - MS Exchange Server
  - MS SQL Server
  - SAP R/3
- Restore Sessions

Select the files and directories that you want to restore.

Search interval: Last 3 months

- ☑ \
  - ☑ Program Files
    - ☑ Exchsrvr
      - ☑ Second
        - ☑ Blablablasuperdupper.edb
        - ☑ Blablablasuperdupper.stm
        - ☑ Publictest.edb
        - ☑ Publictest.stm
        - ☑ Seba.edb
        - ☑ Seba.stm
        - ☑ Seba1.edb
        - ☑ Seba1.stm
        - ☑ wptestadmin.edb
        - ☑ wptestadmin.stm

## Figure 87 Restoring Exchange database

Exchange database consists of two files: *name*.edb and *name*.stm. For a particular database, select both files. You can restore a storage group by selecting the storage group folder when the entire storage group was backed up and all databases reside in the same directory.

3. Under the **Destination** tab, select the application system as the **Target client**.

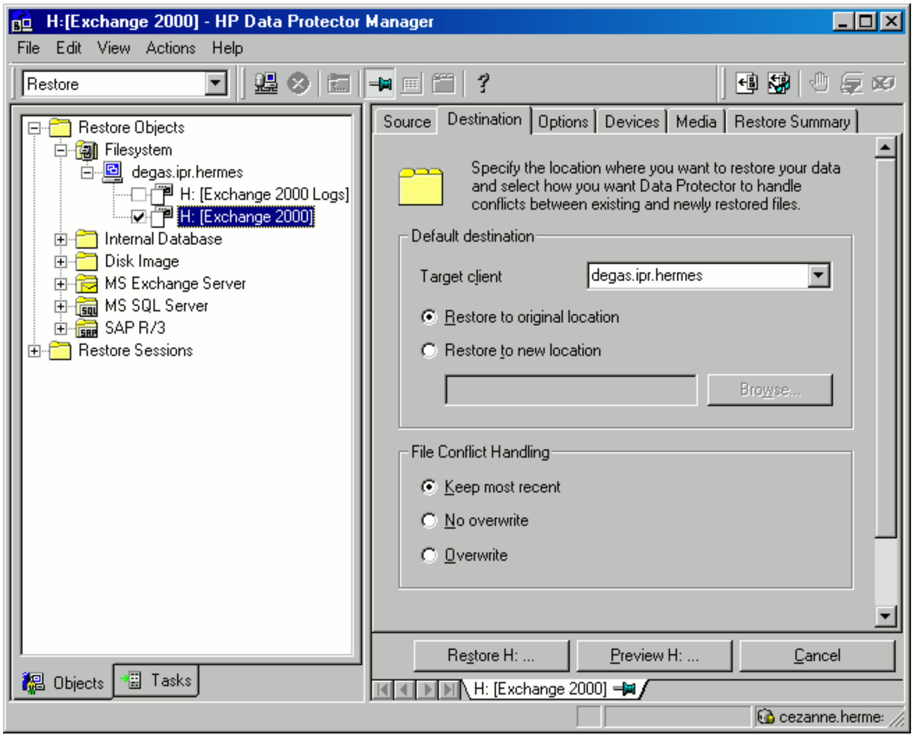

Specify the location where you want to restore your data and select how you want Data Protector to handle conflicts between existing and newly restored files.

Default destination

Target client: degas.ipr.hermes

● Restore to original location

○ Restore to new location

File Conflict Handling

● Keep most recent

○ No overwrite

○ Overwrite

## Figure 88 Selecting the application system

4. Set restore options. For information, see the online Help index: “restore, options“.

5. In the **Devices** page, select devices to use for the restore. Note that you can use a different device for restore than the one used for backup. For information, see the online Help index: “selecting, devices for restore”.

6. Click **Restore**. The **Start Restore Session** dialog box is displayed.

7. Specify **Report level** and **Network load**.

8. *XP, EMC:*

Select **StorageWorks XP restore**. Click **Next**.

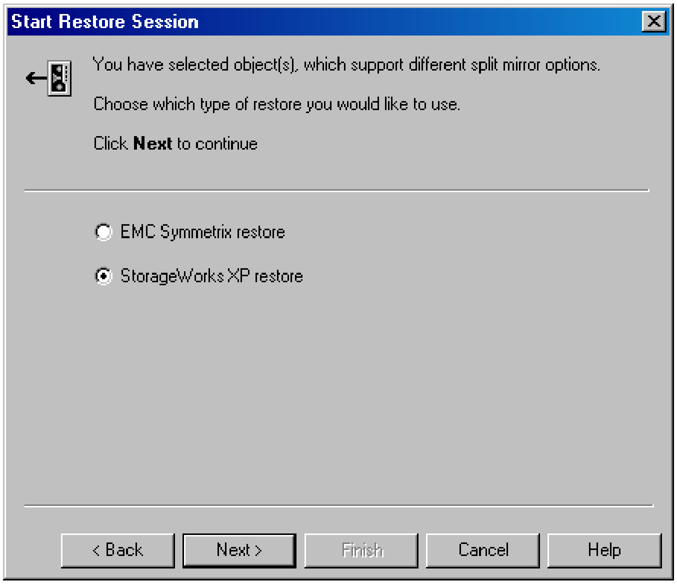

You have selected object(s), which support different split mirror options.

Choose which type of restore you would like to use.

Click **Next** to continue

○ EMC Symmetrix restore

● StorageWorks XP restore

Figure 89 Selecting XP restore

9. *XP:*

In the **Start Restore Session** dialog box, select **Disabled** in the **Mirror mode** drop-down list. This sets a direct restore from the backup media to the application system on LAN.

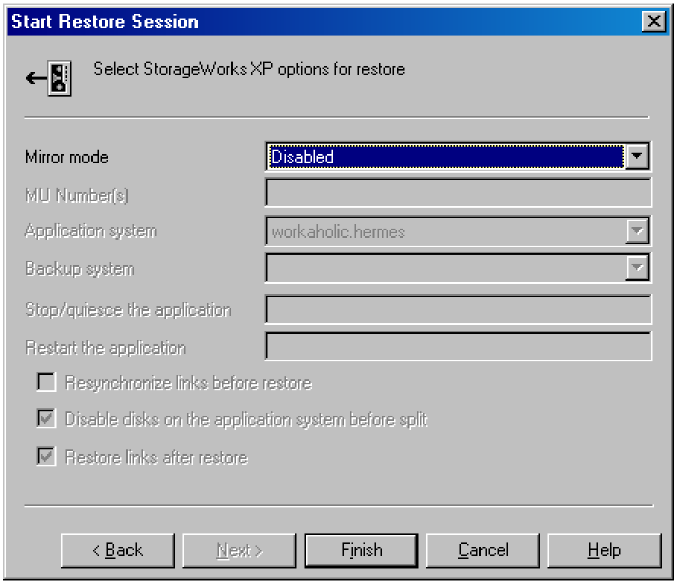

Start Restore Session

Select StorageWorks XP options for restore

Mirror mode: Disabled

MU Number(s):

Application system: workaholic.hermes

Backup system:

Stop/quiesce the application:

Restart the application:

□ Resynchronize links before restore

☑ Disable disks on the application system before split

☑ Restore links after restore

**Figure 90 XP Restore option**

10. Click **Finish** to start restore.

## Rollforward recovery

IMPORTANT:

The procedure described below gives instructions on how to copy `.edb`, `.stm`, and `.log` files to the appropriate database location and only part of the procedure described in the *Offline Backup and Restoration Procedures for Exchange (296788)*.

When following the Microsoft procedure, use the procedure below to copy `.edb`, `.stm` and `.log` files to the appropriate database location.

The procedure below needs to be utilized twice: for the database and for transaction logs restore.

Proceed as follows using the Data Protector Manager:

1. In the Context List, select **Restore**.

2. Proceed as follows:
    - To restore the database, expand **Restore Objects**, **Filesystem**, and the name of the backed up server. Then select the **[Exchange 2000]** object as shown in Figure 91.
    - To restore transaction logs, expand **Restore Objects**, **Filesystem**, and the name of the backed up server. Then select the **[Exchange 2000 Logs]** object as shown in Figure 92.

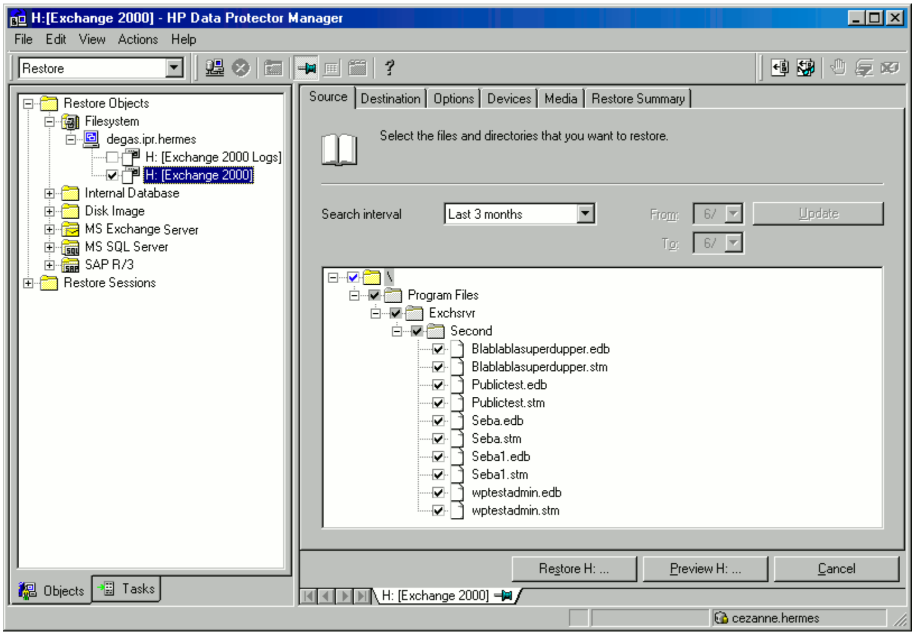
H:[Exchange 2000] - HP Data Protector Manager

Restore

- Restore Objects
  - Filesystem
    - degas.ipr.hermes
      - □ H: [Exchange 2000 Logs]
      - ☑ H: [Exchange 2000]
  - Internal Database
  - Disk Image
  - MS Exchange Server
  - MS SQL Server
  - SAP R/3
- Restore Sessions

Select the files and directories that you want to restore.

Search interval: Last 3 months

- ☑ \
  - ☑ Program Files
    - ☑ Exchsrvr
      - ☑ Second
        - ☑ Blablablasuperdupper.edb
        - ☑ Blablablasuperdupper.stm
        - ☑ Publictest.edb
        - ☑ Publictest.stm
        - ☑ Seba.edb
        - ☑ Seba.stm
        - ☑ Seba1.edb
        - ☑ Seba1.stm
        - ☑ wptestadmin.edb
        - ☑ wptestadmin.stm

H: [Exchange 2000]

cezanne.hermes

## Figure 91 Restoring Exchange database

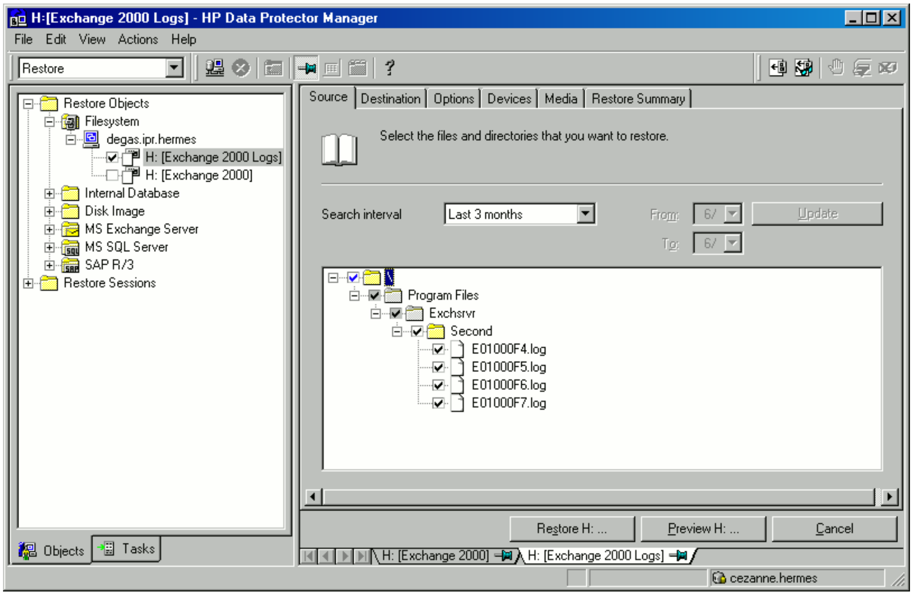

H:[Exchange 2000 Logs] - HP Data Protector Manager

Select the files and directories that you want to restore.

Search interval: Last 3 months

Figure 92 Restoring log files

3. Make the following selections in the Results Area:

   - To restore the Exchange database, select `.edb` and `.stm` files. For a particular database, select both files. You can restore a storage group by selecting the storage group folder when the entire storage group was backed up and all databases reside in the same directory.
   - To restore transaction logs, select `.log` files.

4. Select the application system as the **Target client** under the **Destination** tab.

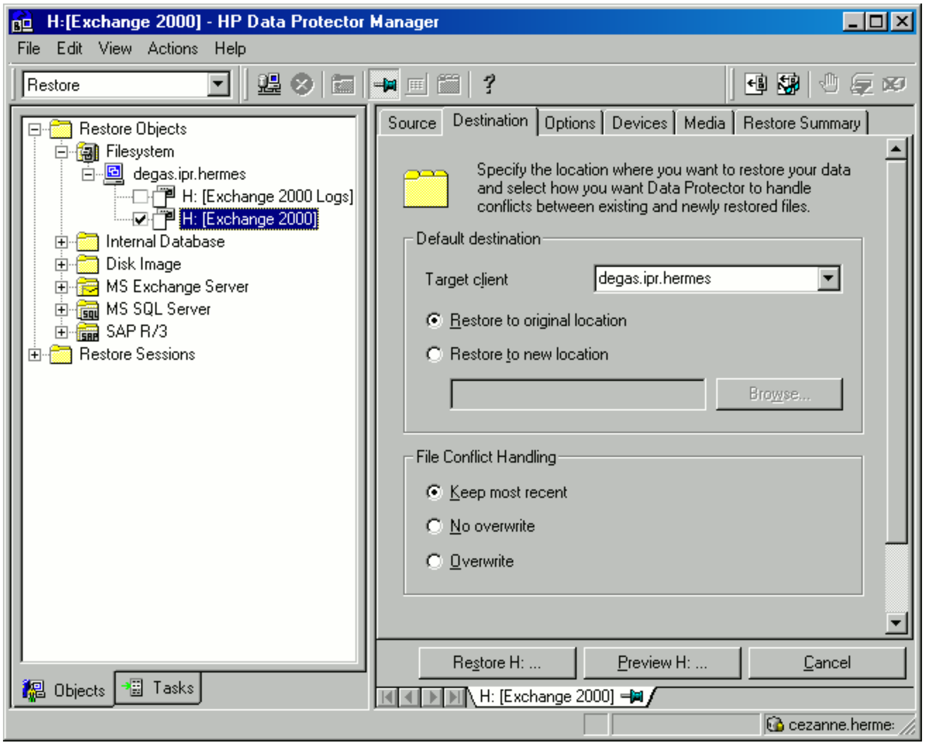

H:[Exchange 2000] - HP Data Protector Manager

- Restore Objects
  - Filesystem
    - degas.ipr.hermes
      - □ H: [Exchange 2000 Logs]
      - ☑ H: [Exchange 2000]
  - Internal Database
  - Disk Image
  - MS Exchange Server
  - MS SQL Server
  - SAP R/3
- Restore Sessions

Specify the location where you want to restore your data and select how you want Data Protector to handle conflicts between existing and newly restored files.

Default destination

Target client: degas.ipr.hermes

● Restore to original location

○ Restore to new location

File Conflict Handling

● Keep most recent

○ No overwrite

○ Overwrite

## Figure 93 Selecting the application system

5. Set restore options. For information, see the online Help index: “restore, options“.

6. Click **Restore**. The **Start Restore Session** dialog box is displayed.

7. Specify **Report level** and **Network load**. Click **Next**.

8. ***XP, EMC:***

Select **StorageWorks XP restore**. Click **Next**.

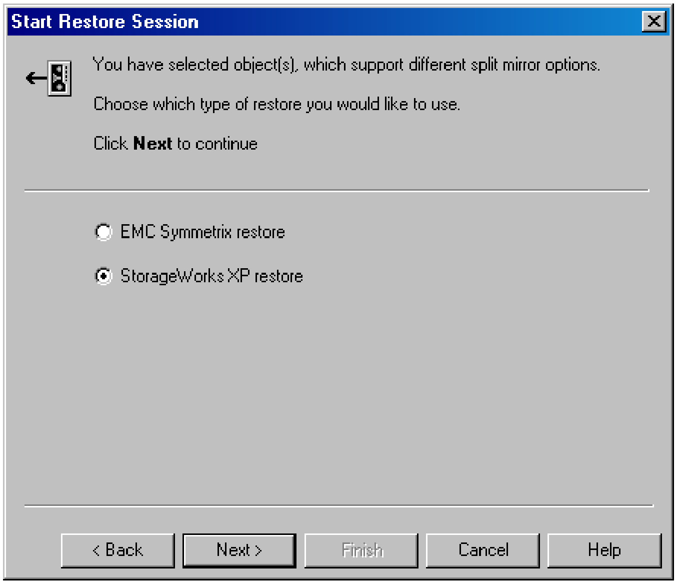

Start Restore Session

You have selected object(s), which support different split mirror options.

Choose which type of restore you would like to use.

Click **Next** to continue

○ EMC Symmetrix restore

● StorageWorks XP restore

Figure 94 Selecting XP restore

9. ***XP:***

   In the **Start Restore Session** dialog box, select **Disabled** in the **Mirror mode** drop-down list. This sets a direct restore from the backup media to the application system on LAN.

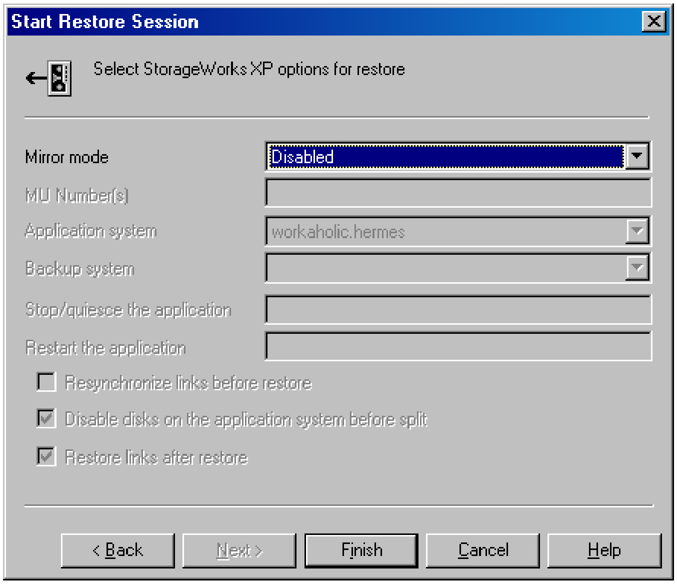

Start Restore Session

Select StorageWorks XP options for restore

Mirror mode: Disabled

MU Number(s):

Application system: workaholic.hermes

Backup system:

Stop/quiesce the application:

Restart the application:

□ Resynchronize links before restore

☑ Disable disks on the application system before split

☑ Restore links after restore

**Figure 95 XP Restore option**

10. Click **Finish** to start restore.

## Instant recovery

See the *HP Data Protector zero downtime backup concepts guide* and *HP Data Protector zero downtime backup administrator's guide* for general information on instant recovery.

You can recover Exchange databases using:

- Point-in-time recovery

  Databases are restored to their state at the backup time, and all data modified after that is lost.

- Rollforward recovery

  Restores Exchange database files and transaction logs, and then replays the transaction logs. Databases are recovered to their last consistent state.

## Point-in-time recovery

For detailed procedure on how to recover the Exchange Server, see *Offline Backup and Restoration Procedures for Exchange (296788)* at http://support.microsoft.com.

When following the Microsoft procedure, see the *HP Data Protector zero downtime backup administrator's guide* for information on copying backed up `.edb` and `.stm` files to the appropriate database using instant recovery.

## Rollforward recovery

For detailed procedure on how to recover the Exchange Server, see *Offline Backup and Restoration Procedures for Exchange (296788)* at http://support.microsoft.com.

When following the Microsoft procedure, see the *HP Data Protector zero downtime backup administrator's guide* for information on copying backed up `.edb` and `.stm` files to the appropriate database using instant recovery.

The procedure below gives instructions on how to restore Exchange transaction logs using the Data Protector GUI when following the Microsoft procedure:

1. In the Context List, select **Restore**.
2. In the Scoping Pane, expand **Restore Objects**, **Filesystem**, and then the name of the backed up server. Select the **[Exchange 2000 Logs]** object as shown in Figure 92.
3. Under the **Destination** tab, select the application system as the **Target client as shown in** Figure 93.
4. Set restore options. For information, see the online Help index: "restore, options".
5. Click **Restore**. The **Start Restore Session** dialog box is displayed.
6. Specify **Report level** and **Network load**. Click **Next**.

7. ***XP, EMC:***

Select **StorageWorks XP restore**. Click **Next**.

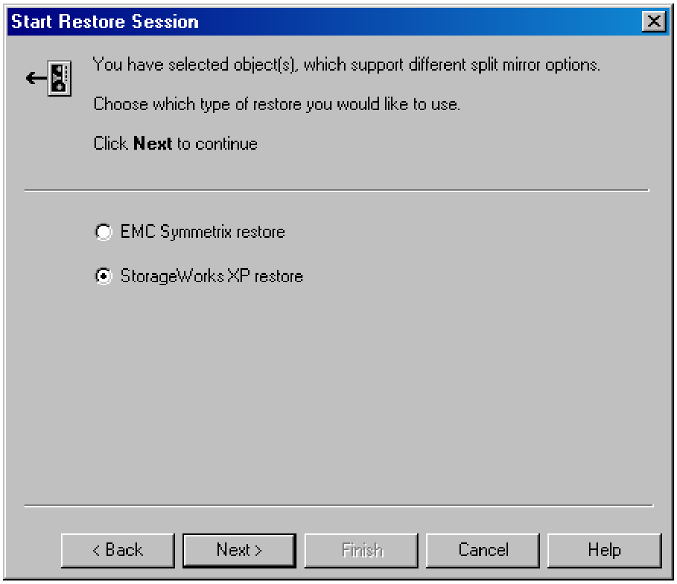

Start Restore Session

You have selected object(s), which support different split mirror options.

Choose which type of restore you would like to use.

Click **Next** to continue

○ EMC Symmetrix restore

● StorageWorks XP restore

**Figure 96 Selecting XP restore**

8. *XP:*

In the **Start Restore Session** dialog box, select **Disabled** in the **Mirror mode** drop-down list. This sets a direct restore from the backup media to the application system on LAN.

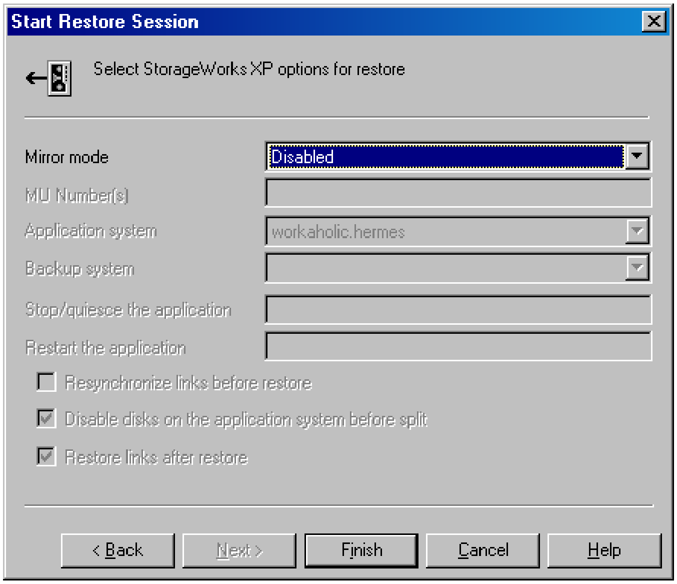

Start Restore Session

Select StorageWorks XP options for restore

Mirror mode: Disabled

MU Number(s):

Application system: workaholic.hermes

Backup system:

Stop/quiesce the application:

Restart the application:

☐ Resynchronize links before restore

☑ Disable disks on the application system before split

☑ Restore links after restore

Figure 97 XP Restore option

9. Click **Finish** to start restore.

# Troubleshooting

This section lists general checks and verifications, plus problems you might encounter when using the Data Protector Exchange Server integration. Start at Problems. If you cannot find a solution there, perform general checks and verifications.

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

For general ZDB, restore, and instant recovery related troubleshooting, see the *HP Data Protector zero downtime backup administrator's guide*.

## Before you begin

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

## Checks and verifications

If your configuration, backup, or restore failed:

- Check that Exchange Server services (Microsoft Exchange System Attendant and Microsoft Exchange Information Store) are running.
- Using Exchange System Manager, check that all stores to be backed up are mounted and all stores to be restored are dismounted.
- Perform a backup of the Exchange Information Store using Windows Backup. If the backup fails, fix Exchange Server problems first, and then perform a backup using Data Protector.
- Ensure that the Cell Manager is correctly set on Exchange Server by checking the following registry entry:

  ```
  HKEY_LOCAL_MACHINE\SOFTWARE\Hewlett-Packard\OpenView\OmniBack
  II\Site
  ```

  Its name and value must be `CellServer` and "*Cell Manager hostname*", respectively.
- Examine system errors reported in *Data_Protector_home*`\log\debug.log` on Exchange Server functioning as a Data Protector client.

  Additionally, examine the errors reported in the Windows Event log.
- Check if the following directories exist on the Data Protector Cell Manager:

  *Data_Protector_home*`\config\server\barlists\msese`

  *Data_Protector_home*`\config\server\barschedules\msese`
- Make a test filesystem backup and restore of the problematic client. For information, see the online Help.
- Create a backup specification to back up to a null or file device and run the backup. If the backup succeeds, the problem may be related to backup devices.

See the *HP Data Protector troubleshooting guide* for instructions on troubleshooting devices.

- Try to restart the Microsoft Exchange Server and start the backup again.
- Check that the *Exchange_home*\bin directory is added to the Windows Path environment variable. For details, see "Configuring the integration" on page 287.
- If you cannot mount the storage after a successful restore, check that LOGS storage on the same storage group is also restored.
- Define a directory for temporary log files in the **Restore** context. Check if the specified directory exists. If it does not, create it or specify another existing directory.
- To restore to another system, make sure Exchange Server is installed on that system and has the same organization and site names as the restored server.

## Problems

### Problem

**Exchange integration is unable to initiate clsEx2000 class**

The following error is reported in the debug.log file:

```
Unable to initiate clsEx2000 class. Error: -2147221164
```

### Action

Register omniex2000.dll by running the following command from *Data_Protector_home*\bin: regsvr32 omniex2000.dll

### Problem

**Omnicreatedl cannot create a backup specification for the log file backup**

The following error is reported:

```
Cannot create log datalist for name.
```

### Action

Check if circular logging is disabled on the Exchange Server for the storage group. If not, disable it.

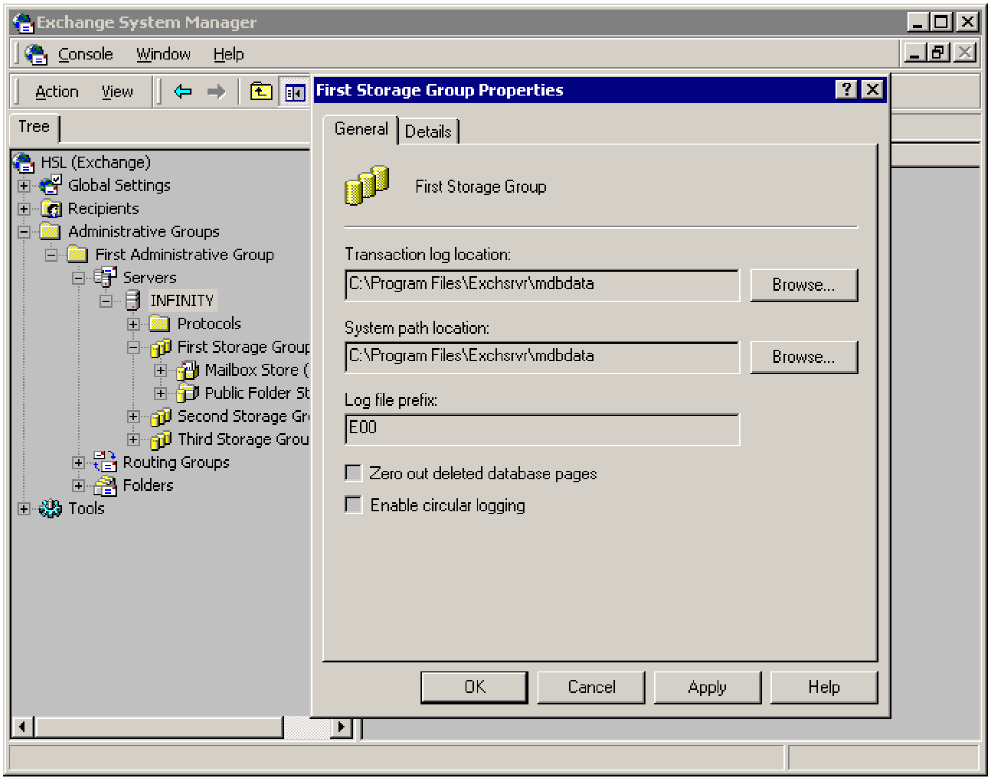

Transaction log location: C:\Program Files\Exchsrvr\mdbdata

System path location: C:\Program Files\Exchsrvr\mdbdata

Log file prefix: E00

☐ Zero out deleted database pages

☐ Enable circular logging

Figure 98 Disabling circular logging

## Problem

**Omnicreatedl cannot create a backup specification and the session is aborted**

The following error is reported:

```
[ERROR] Could not obtain any data for backup from host app_sys.
Make sure the Exchange server is running and Data Protector Exchange
Integration installed.
```

## Actions

- If the application is cluster-aware, specify the `-virtualSrv` parameter.
- Check the user rights on the application system client. You must have the `Save backup specification` user right to create and save a backup specification.

- Check that *Exchange_home*\bin is listed in the Path system variable. For instructions, see Creating ZDB specifications.

## Problem

### Error starting service

If the KMS service password is not kept on disk, the following error is reported:

```
[Major] From: OB2BAR_main@machine.com ""  Time: 10/20/2003 5:17:24
PMError starting service.
```

Reconfigure KMS to store the KMS service password on disk. See *XADM: How to Change the KMS Service Password Startup Location (196129)* for information. Then restart the backup.

# 5 Integrating the Data Protector ZDB integrations and Microsoft Volume Shadow Copy Service

## Introduction

A traditional backup process is based on the direct communication between the backup application and the application whose data is backed up. This backup method requires from the backup application an individual interface for each application it backs up.

The number of applications on the market is constantly increasing. The necessity of handling application specific features can cause difficulties in backup, restore, and storage activities. An effective solution to this problem is introducing a coordinator among the actors of the backup and restore process.

## Volume Shadow Copy Service

**Volume Shadow Copy Service (VSS)** is a software service introduced by Microsoft on Windows operating systems.

During backup, the service collaborates with a backup application (for example, Data Protector), applications to be backed up (usually database applications), shadow copy providers, and the operating system kernel to create a consistent, point-in-time shadow copy set, which can be backed up to backup media or can be kept on a disk array for instant recovery or data mining.

During restore, the service collaborates with the backup application and database application to prepare for the restore operation and to perform the recovery of the restored data.

Data Protector supports the integration with VSS.

The Data Protector Volume Shadow Copy integration provides a unified communication interface that can coordinate backup and restore of an application regardless of their specific features. With this approach, a backup application does not need to handle each application to be backed up specifically. However, the production application as well as the backup application must conform to the VSS specification.

Figure 99 on page 331 and Figure 100 on page 331 show the differences between the traditional backup model and the model with the Data Protector Microsoft Volume Shadow Copy integration.

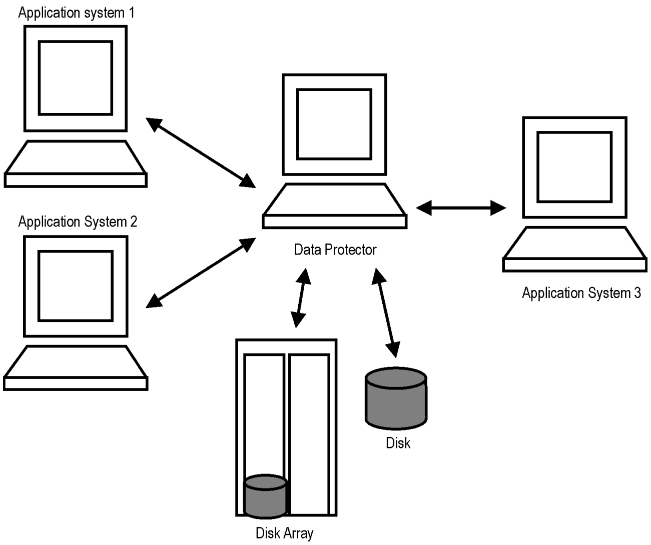

## Figure 99 Actors of the traditional backup model

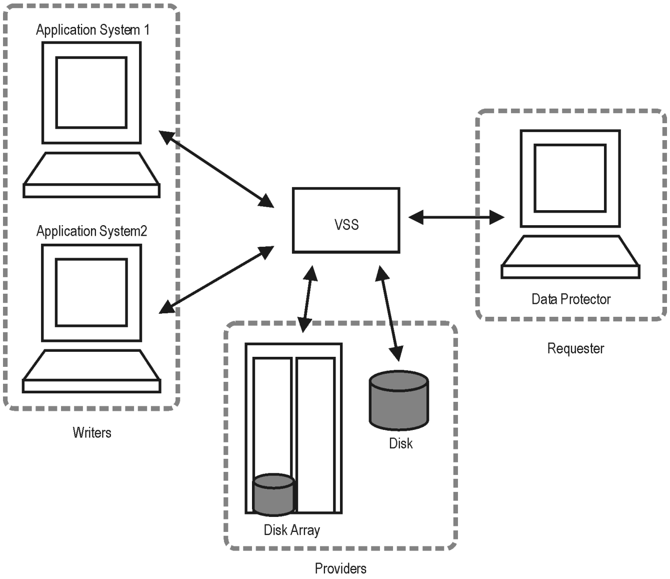

## Figure 100 Actors of the Data Protector VSS integration backup model

Without using the Volume Shadow Copy Service, Data Protector has to communicate with each application to be backed up individually. The Data Protector VSS integration introduces a unified backup and restore interface and provides the coordination among the participants of the backup and restore process.

## VSS backup types

The following backup types are available with the Data Protector VSS integration:

- Local backup

  The volume shadow copies are presented to the same application system from which they were created. The shadow copies are backed up to tape or to disk from the application system. If they are kept on the disk array, they can be mounted on the application system.

- Transportable backup

  Instead of presenting the volume shadow copies to the application system, they are presented to an alternate system, the backup system, from which they are backed up to backup media at leisurely pace. In this way, the impact of backup on performance on the application system is reduced to the creation of volume shadow copies. If the shadow copies are kept on the disk array, they can be mounted on the backup system. A VSS hardware provider that supports transportable snapshots is required.

## VSS restore

Using Data Protector, you can perform the restore:

- From backup media to the application system on LAN (standard restore).
- Using the instant recovery functionality.

  During instant recovery, the data is restored from backup replica to the source volumes on the application system. Only the needed differential and transaction log backups are restored from the backup medium.

### Instant recovery considerations

- The instant recovery functionality restores data from a replica on the backup system to the source volumes on the application system. Therefore, using instant recovery, it is *only* possible to *selectively* restore separate writer’s components if they reside on separate source volumes.

# Integration concepts

The Data Protector integration with the Microsoft Volume Shadow Copy Service provides full support for certified VSS writers.

Refer to the *HP Data Protector zero downtime backup concepts guide* for a general description of ZDB (split mirror or snapshot backup) and instant recovery concepts.

For a complete list of supported VSS writers and provider, see the latest support matrices at http://www.hp.com/support/manuals.

### Benefits of using the integration

Advantages of using the Data Protector VSS integration are the following:

- Unified backup interface is provided for all applications that provide a writer.
- Data integrity is provided on application level, because it is provided by the writers. No interference is needed from the backup application.

### VSSBAR agent

The central part of the integration is the **VSSBAR agent**, which links Data Protector with the Microsoft Volume Shadow Copy Service. Data Protector Microsoft Volume Shadow Copy integration uses the VSSBAR agent for automatic browsing of VSS-aware writers, coordinating backup and restore. VSSBAR agent is responsible for the following actions:

- detecting VSS writers
- examining and analyzing Writer Metadata Document (WMD)

  **Writer Metadata Document (WMD)** is metadata provided by each writer. Writers identify themselves by the metadata and instruct the backup application what to back up and how to restore the data. Thus, Data Protector follows the requirements provided by the writer when selecting the volumes to be backed up and the restore method.

- requesting shadow copy creation
- backing up writers’ data to media
- coordinating restore session start
- restoring the Writer Metadata Document
- restoring writer’s data from media

## Backup

During the Data Protector VSS integration backup, Data Protector does not interact directly with each writer, but through the VSS interface. It uses the VSSBAR agent to coordinate the backup process. The consistency of data is a responsibility of the VSS writer and not dependent on Data Protector functionality. The backup process of the VSS writers consists of the following phases:

1. When you select writers and components you want to back up and start a VSS integration backup, Data Protector communicates with the Volume Shadow Copy Service (backup coordinator) to notify that the backup is about to start.
2. The coordinator identifies all writers that support the VSS feature and passes the list of available writers and their characteristics (Writer Metadata Document) back to Data Protector.
3. Data Protector examines Writer Metadata and identifies the volumes that contain the data to be backed up. Then the VSS informs available writers about selected components.
4. Data Protector prepares a list of volumes (shadow copy set) that must be put into consistent state, and passes the list back to the coordinator for preparing a shadow copy.
5. The VSSBAR agent notifies the writers about the shadow copy creation. The VSS mechanism ensures that there are no writes on the volume while the shadow copy is being created.

NOTE:

When the VSSBAR agent creates a shadow copy of the volume, this volume is marked in order to avoid attempts to simultaneously create another shadow copy of the same volume. In order to prevent any deadlocks arising from volume locking, only a single VSSBAR agent at a time is allowed to define a shadow copy set.

6. When the writers are fully prepared for the consistent shadow copy backup, the VSSBAR agent passes shadow copy creation requests to VSS.
7. After a shadow copy is created, the VSS service returns the related information to Data Protector.

8. Depending on the backup options selected in the backup specification, the ZDB type selected in the backup specification or at the start of the backup, and the backup type (local or transportable), the following scenarios, which are not mutually exclusive, are possible.

   The replica can be:

   - Kept on a disk array for rotation and instant recovery purposes.
   - Kept on a disk array only for data mining.
   - Moved to a backup medium from the application system (local backup) or backup system (transportable backup).
   - Presented and mounted to the backup system specified in the backup specification. It can be mounted in read–only or read/write mode.
   - If not kept on the array, destroyed after the backup session completes (after data in the replica is moved to a backup medium).
   - Manipulated using the Data Protector `omnidbvss` command.

   Figure 101 on page 336 shows the relationships between the components of a local VSS backup. Figure 102 on page 336 shows the relationships between the components of a transportable VSS backup.

9. The related information about the completed backup session is recorded in the VSSDB.

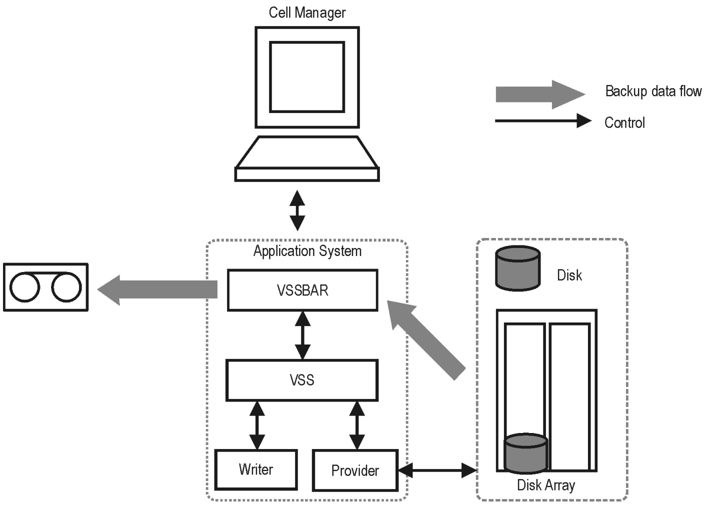

## Figure 101 Local VSS backup

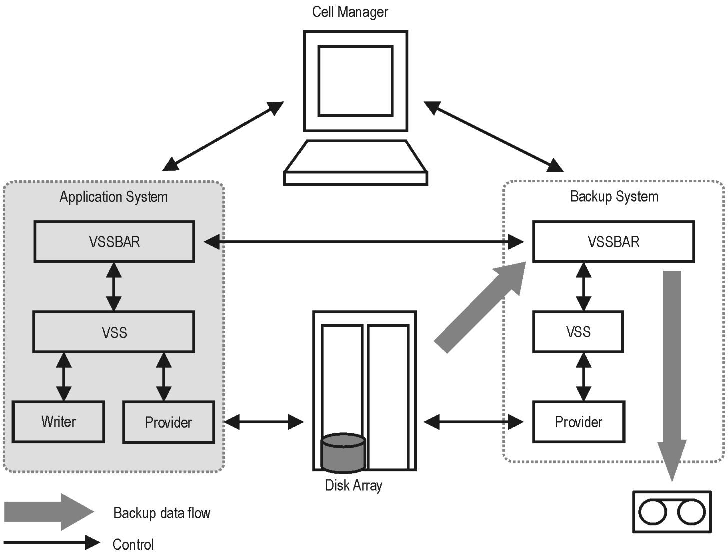

## Figure 102 Transportable VSS backup

# Data consistency

The **filesystem backup** does not guarantee application data consistency, only filesystem consistency. Application data consistency can be achieved only by using supported application writers, such as Microsoft Exchange Server writer.

## Backup with the Enterprise Virtual Array VSS and VDS hardware provider

There are three replica types available when backing up using the Enterprise Virtual Array (EVA) hardware provider:

- Standard snapshot
- Vsnap
- Snapclone

However, in the Data Protector GUI, you can select only between standard snapshot (**Snapshot (Differential)**) and snapclone (**Mirror/Clone (Plex)**). To create vsnaps, you must first configure hardware provider to create vsnaps.

For instant recovery, only snapclone (**Mirror/Clone (Plex)** replica type) is supported.

For more information, see "Snapshot replica" on page 361.

## Backup with the Disk Array XP VSS and VDS hardware provider

There are two configuration modes available for the backup with the Disk Array XP (XP) hardware provider:

- VSS compliant mode

  When the XP provider is in VSS compliant mode, the source volume (P-VOL) and its replica (S-VOL) are in simplex, unpaired state after backup. Therefore the number of replicas (S-VOLs per P-VOL) rotated is not limited. Restore from a backup in such a configuration is possible only by switching disks.

- Resync mode

  When the XP provider is in the resync mode, the source volume (P-VOL) and its replica (S-VOL) are in a suspended mirror relationship after backup. The maximum number of replicas (S-VOLs per P-VOL) rotated is three, provided that the 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.

# Instant recovery

After ZDB to disk+tape or ZDB to disk, restore can be performed using instant recovery. After ZDB to disk, instant recovery is the only possible restore method. Backup session information is saved in the IDB, and the array-specific information required for instant recovery is saved in the VSS database (VSSDB).

Instant recovery consists of two phases:

- An instant recovery session is started and the relevant data files are restored back to the system.
- A writer performs the instant recovery, which is based on the restored data. In case of MSDE, SQL Server 2005, and Microsoft Exchange Server 2003 and 2007 writers, database recovery is performed automatically by the database engine.

IMPORTANT:

During instant recovery, the whole replica is restored. This means not only are the originally selected backup objects restored, but the complete content of all the volume groups that contained them. After the instant recovery, the content is returned to the states when the replica was created.

For detailed information on the VSS integration instant recovery concepts, see the "Instant recovery" on page 391.

For detailed information on the general instant recovery concepts, see the *HP Data Protector zero downtime backup concepts guide* and *HP Data Protector zero downtime backup administrator's guide*.

## Standard restore from tape

Data Protector offers two restore modes:

- **component restore** using the VSS service
- **file restore** using the DMA instead of VSS.

By default, Data Protector restores writer components using the VSS service.

Instant recovery is also available within the Data Protector VSS integration. This functionality requires VDS hardware providers.

### Restoring components

During the restore procedure, the Data Protector VSS integration coordinates communication between Data Protector and the writers. In general, the restore flow consists of the following phases: preparing for restore, restoring components, and notifying the application writers that a restore has been completed. The restore procedure of the VSS-aware writers consists of the following phases.

1. Data Protector first restores the metadata, which was collected during the backup. Then it examines the metadata to identify the backup components and determine the restore method. It also checks if restore to specific volumes is possible.
2. Data Protector connects to the coordinator (VSS service) to notify that the restore is about to start, which in turn communicates with the writer. Data Protector restores the data from the backup media to the locations specified in the backup metadata. During the restore, Data Protector follows the writers' instructions regarding any additional checking or processing specified in the WMD.
3. After the data are successfully restored from the backup media, Data Protector informs the coordinator that the restore is completed and the writers can now access the newly-restored data and start the internal processing, for example recovery.

## Restoring files

For a successful restore of a VSS component, all files comprising this component must be restored. If a restore of a single file fails, the restore of the a whole component fails. Data Protector offers an additional restore mode for restoring single files that does not use the Microsoft Volume Shadow Copy Service, thus solving this problem. This mode can also be used for restoring to systems that do not support VSS or do not have a VSS writer installed.

When restoring files or a group of files, DMA is started and the files are restored using the standard Data Protector filesystem restore procedure.

IMPORTANT:

As the file restore mode does not utilize VSS services, additional tasks that are performed after a component restore – such as database recovery – are not performed and your application data may be left in an inconsistent state, requiring additional manual procedures before the application is recovered.

# VSS database

The VSS database (VSSDB) is an extension to the Data Protector internal database (IDB) on the Cell Manager. It holds the VSS integration-specific information about the VSS backup sessions and their replicas. The stored information can be used for the following purposes:

- Instant recovery
- Data mining (you can save backup components and writer metadata documents)

- Listing the backup sessions and viewing details about them
- Removing the backup sessions
- Enabling (presenting and mounting) and disabling (dismounting and unpresenting) the ZDB-to-disk and ZDB-to-disk+tape sessions' replicas

The VSSDB holds the VSS sessions' metadata. This metadata is stored in two parts of the VSSDB, persistent and the non-persistent. The **persistent** part holds information about all backup sessions, while the **non-persistent** part holds information only on instant recovery–enabled sessions (ZDB-to-disk and ZDB-to-disk+tape sessions). Once a replica from the instant recovery-enabled session is rotated out of the replica set, the information about the session is deleted from the non-persistent part of the VSSDB and the session's target volumes are removed from the disk array. The information about the session is still available in the persistent part of the VSSDB and can be deleted only manually using the `omnidbvss` command. Note that ZDB-to-tape sessions and sessions created using the software provider are always recorded only to the persistent part of the VSSDB.

You can query and manage the items of the VSSDB using the `omnidbvss` command. For information on the command syntax, description, and examples, see the *HP Data Protector command line interface reference*.

The following tasks can be performed using the `omnidbvss` command:

- Querying the VSSDB
- Deleting backup sessions
- Enabling and disabling replicas

## Querying the VSSDB

Using the `omnidbvss` command, you can list:

- All backup sessions, as well as their details.
- All backup sessions based on a specific backup specification, as well as their details.
- All backup sessions recorded in the persistent part of the VSSDB, which were created before a specified date, as well as their details.
- Details on a specific backup session, identified by the session ID.

Using the `omnidbvss` command, you can save the documents about the backup components and writer metadata to a specified directory.

See the `omnidbvss` man page for the command syntax and examples.

## Deleting backup sessions

Using the omnidbvss command, you can delete:

- A specific instant recovery-enabled backup session (a replica version), identified by the session ID, from the non-persistent part of the VSSDB and disk array, or only from the VSSDB.
- All instant recovery-enabled backup sessions based on a specific backup specification (a replica set) from the non-persistent part of the VSSDB and disk array, or only from the VSSDB.
- A specific backup session, identified by the session ID, from the persistent part of the VSSDB.
- All backup sessions based on a specific backup specification or that were created before a specified date from the persistent part of the VSSDB.

See the omnidbvss man page for the command syntax and examples.

## Enabling and disabling replicas

Using the omnidbvss command, you can present and mount (enable) on a backup system and unmount and unpresent (disable) from a backup system the replicas from:

- All instant recovery-enabled sessions.
- A specific instant recovery-enabled session, identified by the session ID.
- All instant recovery-enabled sessions based on a specific backup specification.

## Changes in the environment

During a VSS instant recovery sessions or the VSSDB maintenance sessions using the omnidbvss command (remove, disable, enable), the state of target volumes in a replica is checked prior to restore, remove, disable, and enable operations. Before restore, the state of the source volumes on the application system is also checked. If the check finds any changes that have not been tracked in the VSSDB, the session is aborted, notifying you about the specific changes. This check can be disabled by setting the OB2VSS_IGNORE_BACKUP_DISK_CHANGES and OB2VSS_IGNORE_SOURCE_DISK_CHANGES omnirc variables to 1.

During VSS backup session, a replica is created and left on a disk array until the specified number of replicas rotated is reached. After that, the next replica to be created replaces the oldest replica in the set. If the oldest replica in the set cannot be replaced because you have manually (not using Data Protector) unpresented the replica from the backup system and presented it on some other system, the backup session continues with creating a new replica and leaving the old one on the array. Note that in this case, if there is no space on the disk array for creating a new replica, the backup session also fails.

If you set the `OB2VSS_IGNORE_BACKUP_DISK_CHANGES` `omnirc` variable to `1`, the backup session continues by removing the old replica and creating a new one in its place.

## Microsoft Exchange Server 2007 writer concepts

This section gives details of additional features that Microsoft Exchange Server 2007 supports.

### Backup

Microsoft Exchange Server 2007 offers two models of replication for data protection that are supported by Data Protector.

- local continuous replication (LCR)

  With LCR, you can create and maintain an exact copy (LCR copy) of databases in a storage group. LCR copies are used in the event of data corruption, since you can switch the Exchange server to use LCR copies in only a few seconds. If an LCR copy is used for backup and is located on a different disk from the original data, the load on a production database is minimal.

- cluster continuous replication (CCR)

  CCR has the same characteristics as LCR. The only difference is that in the CCR environment, databases and transaction logs are replicated to separate servers. Therefore, CCR copies can be used for disaster recovery. You can perform VSS backups on the passive Exchange Server node where the CCR copy is located and thus reduce the load on the active node.

The replicated storage groups are represented as a new Exchange Server writer instance, Exchange Replication Service. They are backed up in the same way as original or production storage groups.

If using LCR or CCR with Standby Continuous Replication (SCR) configured, backups can only be performed on the source side of the SCR. Microsoft does not support backups on the SCR target side. For further information on SCR and supported SCR configurations, refer to the Microsoft website.

### Restore

With Microsoft Exchange Server 2007 writer, you can restore your data not only to the original location (from which the backup was performed) but also to a different location. You can restore:

- A whole storage group
- A single store

In both cases the respective LCR or CCR copies can also be restored.

You can restore data to:

- The original storage group
- A different storage group
- A non-Exchange location – With this restore method, after the restore is completed, the Recovery Storage Group (RSG) can be created automatically.
- A recovery server – This restore method restores data to different client and different storage group.

If you restore to a different storage group, you can access single mailboxes or individual e-mail messages at a different location without changing the original storage group content. Furthermore, if the whole server is destroyed, restoring to a different Exchange Server system (recovery server) will minimize the time window during which your mailboxes will be unavailable.

## Microsoft Hyper-V VSS writer concepts

Microsoft Hyper-V writer is the successor to the Microsoft Virtual Server 2005 writer and is supported on Microsoft Windows Server 2008. It is a VSS writer with similar set of functionality as Virtual Server. With both, it is possible to perform backups and restores of virtual machines. With Hyper-V, it is possible to perform online backups using hardware providers.

For backup and restore specifics of Hyper-V writer, see “Microsoft Hyper-V VSS writer backup specifics” on page 375 and “Microsoft Hyper-V VSS writer restore specifics” on page 391.

### Prerequisites

- Windows Server 2008 Service Pack 2 must be installed. For more information, see http://support.microsoft.com/kb/948465.

# Prerequisites and limitations

This is a list of prerequisites and limitations for the Data Protector Microsoft Volume Shadow Copy Service integration. Additional limitations and recommendations that are not directly connected to the integration (such as operating system and GUI limitations) and disk array limitations are listed in the *HP Data Protector product announcements, software notes, and references*.

## Prerequisites

- Before you begin, ensure that you have correctly installed and configured Data Protector, writers and providers. Refer to the:
  - http://www.hp.com/support/manuals for an up-to-date list of supported versions, platforms, devices, disk arrays, limitations, and other information.
  - *HP Data Protector installation and licensing guide* on how to install Data Protector on various architectures and how to install the Data Protector Microsoft Volume Shadow Copy Service integration.
  - Writers and shadow copy providers documentation for instructions on how to install and configure writers and providers on your system.
  - To be able to perform instant recovery–enabled backups and instant recovery, ensure that the source volumes do not reside on GPT (GUID Partition Table) disks. If the system where you want to install an application database has GPT disks, change them to MBR (Master Boot Record) disks using the Windows Management Console. For instructions on how to perform the change, see the Microsoft Windows Help.
- To enable successful shadow copy creation install the following Windows Server 2003 hotfixes: KB950903 and KB949002.

### Transportable backup prerequisites

- The backup system must be configured to accept connections from the application system.
- The VSS and VDS hardware providers are required for this type of backup and they must be installed and configured on the application and backup system. Refer to the provider documentation for details.

## Limitations

For a list of general Data Protector limitations, see the *HP Data Protector product announcements, software notes, and references*.

- Maximum 512 shadow copies per volume are allowed. The number of shadow copies per volume is limited by system resources.
- To run a VSS integration backup, the writer’s data must be on an NTFS filesystem. For hardware providers, this is not required.
- The VSS integration backup of writers which store their data on network shared volumes is not supported.
- The Data Protector Microsoft VSS integration does not provide any restore method for writers requesting a custom restore method. These writers are by default not presented by Data Protector.

  If a writer specifies a custom restore method, it is only possible to restore the writer's data as plain files using the Data Protector functionality. You can perform the custom restore manually. Refer to the writers documentation for additional information on the restore methods.
- Preview is not possible for VSS backup and restore sessions.
- VSS backups with hardware provider are only supported for volumes on disk arrays. This is also valid for system disks. Instant recovery of system disks backups is not possible.
- If the backup system is in a cluster and if you select the `Mount the replica on backup system` option, note that the replica disks are mounted to the active cluster node. In the case of a backup system failover, the mounted disks are not unmounted and moved to another node.
- During a local backup, with the option `Track the replica for instant recovery` selected, and with the option `Mount the replica` not selected, note the following:
  - On Windows Server 2003, the created snaphots are unpresented on the application system after the backup.
  - On Windows Server 2008, the created snaphosts are only put offline after the backup, while they stay presented on the application system.

### Enterprise Virtual Array hardware provider limitations

- With the EVA provider, if there are too many, for example more than four, source volumes selected for backup, the creation of snapshots may take too long and VSS may fail the snapshots creation process. It is strongly recommended not to specify too many objects in a backup specification, since they might reside on too many source volumes. For example, a Microsoft Exchange Server storage group should typically not reside on more than four source volumes, and only one storage group should be selected for backup if the storage group occupies four source volumes. The same limitation also applies to Microsoft SQL Server databases. If your backup session still fails, reduce the number of volumes further, for example to two.

## Disk Array XP hardware provider limitations

- If the Disk Array XP VSS hardware provider is in the VSS compliant mode, and there is a high read-write load the disk array, the replica creation may take a long time. Consequently, the backup session may fail.

  By default, Data Protector waits for a maximum of 45 minutes (2700 seconds) for mirrors to become synchronized. If synchronization is not completed when the waiting period expires, Data Protector aborts the backup session.

  You can increase the waiting period duration by setting the `OB2VSS_WAIT_TIMEOUT` `omnirc` variable. The variable value determines the period (in seconds) for which the Data Protector agent waits for mirrors that are created in one session to become fully synchronized. If the `OB2VSS_WAIT_TIMEOUT` variable is set to more than 2 hours (7200 seconds), you should additionally increase the value of the global variable `SmIdleTimeout` past the value of `OB2VSS_WAIT_TIMEOUT`. The value of `SmIdleTimeout` should be specified in minutes, rather than seconds.

  If backup sessions are lengthy and there is a possibility that the VSS backup window will not be met, it is highly recommended that you first switch the Disk Array XP VSS hardware provider to resync mode. Afterwards, consecutively run backup sessions until the backup session count reaches the current value of the **Number of replicas rotated** option. This will start creating disk mirrors on Disk Array XP that will be used in scheduled backups.

## ZDB limitations

- Dynamic disks (Logical Disk Manager partitions) can only be backed up with software providers.
- GPT disks are not supported for backup and instant recovery purposes.

## Instant recovery limitations

- Only writers that have data on volumes managed by VDS hardware providers can be restored using the instant recovery functionality.

- Not all VSS writers supported by Data Protector can be used with the instant recovery functionality. For a list of the VSS writers that support the instant recovery functionality, see the latest support matrices at http://www.hp.com/support/manuals.

# Configuring the integration

The Data Protector Microsoft Volume Shadow Copy integration by itself does not require any configuration steps neither on the Data Protector nor on the application side. However, some applications such as Microsoft Exchange Server may require additional configuration steps. For details, see the appropriate sections. You must also configure the VSS and VDS hardware providers.

VSS writers are either a part of Windows operating system or delivered with applications. Data Protector automatically detects writers when the VSS backup specification is created and registers them.

You may check which writers and providers are installed and registered on your system using the following Windows operating system command:

- For a list of writers: `VSSadmin list writers`
- For a list of VSS providers: `VSSadmin list providers`
- For instant recovery enabled backups, VDS hardware providers should be present in the list of installed software. Check **Control Panel** > **Add/Remove Programs**.

For instant recovery enabled backups, only non-system writers are supported, because instant recovery cannot be performed on a system disk.

### Prerequisites

- Disable the operating system option **Automatic mounting of new volumes**. At the command prompt, run:

  ```
  mountvol /N
  ```

- Do not manually mount target volumes that have been created by Data Protector.
- For cluster-aware clients, install the Data Protector VSS integration software component on all the cluster nodes.

## Configuring the application system for instant recovery-enabled backup sessions

To be able to perform instant recovery-enabled backups and then instant recovery, you must resolve the application system using the `omnidbvss` command. During the resolve operation, Data Protector contacts VDS to get information about the storage IDs of the source volumes presented on the application system. The information about the storageIDs is required during instant recovery, in case such information cannot be gathered at restore time.

If resolving is not performed before a backup session, one of the following happens:

- The application system is resolved automatically during the backup session, if the `OB2VSS_DISABLE_AUTO_RESOLVE` variable in the `omnirc` file is set to `0` (default). In this case, the backup time for creating a replica is prolonged.
- A ZDB-to-disk session fails notifying you to resolve the application system before the backup is run, if the `OB2VSS_DISABLE_AUTO_RESOLVE` variable is set to `1`.
- A ZDB-to-disk+tape session completes with the warning that only backup to tape was performed not leaving the replica on the disk array and thus disabling instant recovery, if the `OB2VSS_DISABLE_AUTO_RESOLVE` variable is set to `1`.

Always perform the resolve operation after:

- Installing or upgrading Data Protector.
- Your source volumes' configuration on the application system has changed (for example, you have modify the existing source volumes or you have presented new source volumes).
- You have added a new storage object (for example, a Microsoft Exchange Server storage group).

To resolve the application system, run the following on any VSS client in the Data Protector cell:

```
omnidbvss -resolve {-apphost ApplicationSystem | -all}
```

where *`ApplicationSystem`* is the name of the application system you want to resolve. In a cluster environment (for example, Microsoft Exchange Server CCR), provide the virtual server name for the *`ApplicationSystem`* parameter to enable the resolve operation on both cluster nodes.

To resolve all application systems in a cell simultaneously, use the option `-all`.

Alternatively, you can set the `omnirc` variable `OB2VSS_ALWAYS_RESOLVE_SOURCES` to `1`. In this case, the source volumes on the application system are resolved automatically during every instant recovery-enabled backup session, but the replica creation time is considerably prolonged. However, do not set this variable to `1` in a CCR Microsoft Exchange environment if you plan to back up only the database copy disks and leave the possibility for eventual instant recovery from these disk to the production database disks. In this case, the production database disks are never resolved during backup and this can cause such an instant recovery to fail.

## Configuring HP StorageWorks Enterprise Virtual Array

This section describes how to configure the Data Protector HP StorageWorks Enterprise Virtual Array (EVA) integration for use with the VSS integration.

### Prerequisites

- Install ***HP storage components and licenses:***
    - HP StorageWorks Virtual Controller Software (VCS or XCS) and Command View (CV) EVA. See the VCS and CV EVA documentation for installation instructions. For information on supported product versions, see the latest support matrices at http://www.hp.com/support/manuals.
    - HP Operations Manager (HP StorageWorks SMI-S EVA provider part).
    - HP StorageWorks Business Copy (BC) EVA microcode and license.
    - If the alternate paths solution is used, install HP MPIO DSM for EVA.
    - An EVA license for, at least, basic and snapshot operation.
- Install ***Data Protector components and licenses:***
    - An instant recovery license.
    - HP StorageWorks EVA SMI-S Agent. If this agent is not installed, only “Restore using Microsoft Virtual Disk Service” (switch of disks) is possible.

      For installation and licensing information, see the *HP Data Protector installation and licensing guide*.
- Make sure the same operating system (and its version) is installed on the application and backup systems.
- Connect EVA to the application and backup systems through the SAN. Backup system must be connected to the same SAN as the EVA.

For supported backup and connectivity topologies, see *HP Data Protector zero downtime backup administrator's guide*.

For general Data Protector and integration-specific limitations, see *HP Data Protector product announcements, software notes, and references*.

## Configuring the Integration

Before you start configuration, make sure you met the prerequisites described in "Prerequisites" on page 351. Then set the login information for SMIS-S EVA provider running on a management system.

To set, delete, list, or check the login information, use the `omnidbsmis —ompasswd` CLI command. For command syntax and examples, see the *HP Data Protector command line interface reference*.

The information you provide is kept in the ZDB database for EVA integration (SMISDB).

For each management systems you configure, the following information is stored:

- Hostname as recognized in the IP network.
- Port number through which SMI-S Agent communicates (default - 5988). If SMI-S provider accepts the SSL-based connection, the default port number is 5989.
- User name and encoded password for SMI-S EVA provider login.

SMISDB resides on the Cell Manager in: *Data_Protector_home*`\db40\smisdb`.

### Considerations

If a failover from the active to the standby management system happens, proceed as follows:

- If standby and failed management systems have the same hostname, no action is needed.
- If standby and failed management systems have different hostnames, remove the failed system from the Data Protector configuration, and then add the new management system.

IMPORTANT:

It is recommended to use the default port number (5988 for non-SSL and 5989 for SSL-based SMI-S provider connection settings).

It is also recommended to run `omnidbsmis -ompasswd -check [-host` *HostName*`]` before backup or instant recovery to verify the configuration of SMI-S EVA Provider.

## Configuring HP StorageWorks Disk Array XP

This section describes how to configure the Data Protector HP StorageWorks Disk Array XP integration for use with the Data Protector VSS integration.

### Prerequisites

- Install ***Disk Array XP components and licenses:***
    - RAID Manager Library on the application and backup systems. See the RAID Manager Library documentation for installation instructions.

      RAID Manager Library is firmware-dependent. Consult the HP sales representative for information on which version of RAID Manager Library to use.
    - Business Copy (BC) XP microcode and a license for at least basic and Business Copy operation.
    - XP VSS and VDS provider.
    - Device specific module (MPIO DSM).
- Install ***Data Protector components and licenses:***
    - An instant recovery license.
    - The HP StorageWorks XP Agent. If this agent is not installed, it is not possible to restore data from the disk array XP.

      For installation and licensing information, see the *HP Data Protector installation and licensing guide*.
- Make sure the same operating system and version is installed on both application and backup systems.
- Connect the application and backup systems to the same Disk Array XP.
- Assign LUNs to the respective ports.
- Pre-configure the LDEVs for the VSS snapshot creation and put them in the S-VOL host group. If you perform restore using VDS, you need to put new LDEVs in the S-VOL host group after the restore, since VDS restore switches the disks and the replica becomes the source volume. Add as many new LDEVS as were used for restore.

### Limitations

- You cannot switch between the VSS compliant mode and resync mode during execution of a backup session.

- If the XP provider is in resync mode, a maximum of three replicas (S-VOLs) can be created for a source volume (P-VOL).
- Second-level mirrors are not supported.

For supported backup and connectivity topologies, see the *HP Data Protector zero downtime backup administrator's guide*.

For general Data Protector and integration-specific limitations, see the *HP Data Protector product announcements, software notes, and references*.

# VSS and VDS provider configuration for ZDB and instant recovery

To be able to perform ZDB-to-disk, ZDB-to-disk+tape, and instant recovery, you must configure the VSS and VDS provider for EVA or XP. Perform the configuration task before you run the backup with the specified instant recovery options for the first time.

For details on how to configure the provider, see the hardware provider documentation.

## Configuring the VSS hardware provider for Enterprise Virtual Arrray

Open the HP StorageWorks *VSS Configuration Utility* for EVA and follow the steps below:

1. Ensure that the correct management appliance IP address is specified.
2. Click **Login** and in the Login dialog box, enter the username and password.
3. Select **SnapClone** as the Snapshot Type.
4. Click **Select Disk Group** and select the disk group with the disks that you intend to back up for instant recovery.

Open the HP StorageWorks *VDS Configuration Utility* for EVA and follow the steps below:

1. Ensure that the correct Appliance IP address is specified.
2. Click **Login** and in the Login dialog box, enter the username and password.
3. Click **Select Disk Group** and select the following disk groups:
   - The disk group with disks created by the VSS provider.
   - The disk group with the source disks.

After you once complete this procedure, the VSS and VDS providers are ready to be used for zero downtime backup and instant recovery.

## Configuring the VSS hardware provider for Disk Array XP

Before running ZDB sessions, open the HP StorageWorks XP VSS Hardware Provider Configuration Utility and choose the configuration mode for your backups: `VSS Compliant Mode` or `Resync Mode`.

NOTE:

Changing the modes between backup sessions may have an impact on restore. For details, see "HP StorageWorks Disk Array XP considerations" on page 396.

# Configuration check

With instant recovery, it is possible to *selectively* restore separate writer's components if they reside on separate source volumes. Instant recovery of separate components also requires that the volumes with components data should not contain any other data. The configuration check detects, whether there is more than one component on the volume and whether there is any other data besides the component's data.

At backup time, Data Protector can check whether the individual components can be selectively restored using the instant recovery functionality. At instant recovery time, Data Protector can check whether the data required for the components restore is available.

In case of the Microsoft Exchange Server Writer, it is checked, whether the whole storage group should be restored, or whether the separate database stores can be restored individually. In case of the MSDE Writer, it is checked, whether the user, system, and log files are on separate volumes.

If the check fails for a component during the backup, you will not be able to select this component for the instant recovery, you will need to recover all writer components. If the check fails during the instant recovery, the instant recovery session will fail.

## Configuration check modes

You can select between three configuration check modes:

- **Strict**

  If any file or folder on the volume does not belong to the component, the ZDB to disk or instant recovery fails.

- **Non-strict**

  If any folder on the volume does not belong to the component, the ZDB to disk or instant recovery fails.

- **Disabled**

  A check detects whether there is more than one component on the volume and whether there is any other data besides the component's data on the volume, but the session will not fail.

IMPORTANT:

If the configuration check is disabled for an instant recovery, you will lose the data that does not belong to a component, but resides on the same volume as the component. Disable configuration check only if instant recovery cannot be performed with an enabled configuration check and only if you are sure that this will not result in a loss of data during the instant recovery.

The configuration check should not be disabled except for the cases when instant recovery cannot be performed with an enabled configuration check. Due to the specific behavior, some writers (not applicable to the MSDE Writer and Microsoft Exchange Server writers) may create temporary files on components volumes during backup and instant recovery causing the check failure. In such cases, instant recovery is not possible without disabling the check.

## Microsoft Exchange Server writer specific configuration

A backup of the Microsoft Exchange database is considered successful only when consistency check of the replicated data files succeeds.

For instant recovery-enabled backups, it is recommended that the stores and transaction logs reside on different source volumes. Such an Exchange Server configuration enables you to restore individual stores and perform roll-forward recovery. Otherwise, if all data resides on the same source volume, all the data residing on this volume must be restored, thus only point-in-time recovery is possible.

### Prerequisites

- When configuring an Exchange Server 2007 LCR environment, ensure that the original database and the database copy of the Exchange Server have the same directory and file structure. Failing to do so will result in unsuccessful attempts to mount the database after restore of an LCR copy to the location where original database resides.

## Microsoft Exchange Server 2007 configuration in CCR environments

In CCR environments where a production database resides on a different EVA disk array from the database copy, and the backup system and the application system are not the same physical system, ensure that:

- The backup system is connected to the same disk array where the replicas (volume shadow copies) of backup objects (either production database or database copy) will be created. If you back up the production database, the backup system must be connected to the disk array on which the production database resides. If you back up the database copy, the backup system must be connected to the disk array on which the database copy resides.
- The VSS/VDS hardware provider is configured for the same Command View on both EVA disk arrays.

If the above prerequisites are not met, failed presentation of replicas to the backup system causes entire backup session to fail.

In cases where the application and backup systems are the same, ensure that the VSS/VDS hardware provider is configured for the same Command View on both EVAs. Otherwise, rotation and deletion of old replicas from other EVA disk array is not possible.

If a failover occurs, the backup using the same backup specification as before failover will fail, since the backup system is no longer connected to the disk array where the new replicas reside. In this case, you can:

- Perform manual failover to restore the same disk array connectivity configuration as before the first failover.
- Create a new backup specification with the new backup objects specified.
- Establish a connection between the backup system and the disk array where the new replicas reside.

For restore limitations in this scenario, see “Limitations” on page 398.

# Writers specifics

This section describes specific information about VSS writers, that you need to take into account before backing up or restoring the writers.

VSS writers either come with the Windows operating system or with applications. For a complete list of supported VSS writers and providers, see the latest support matrices at http://www.hp.com/support/manuals.

The Data Protector Microsoft VSS integration does not provide any restore method for writers requesting a custom restore. If a writer specifies a custom restore method, it is only possible to restore the writer's data as plain files using the Data Protector functionality. You can perform the custom restore manually. Refer to the writers documentation for additional information on the restore methods.

NOTE:

Writers requiring custom restore methods are by default not shown by Data Protector. The `omnirc` variable `OB2VSS_SHOWALLWRITERS` must be set to `1` for all writers to be displayed.

Table 23 on page 356 provides a description of VSS writers.

### Table 23 Writer description

| Writer name | Description | Restore method |
|---|---|---|
| Certificate Authority Writer | This is a system writer, used to back up and restore Certificate Authority (CA) Service database. This service issues, revokes, and manages certificates employed in public key-based cryptography technologies. | Files are restored after a reboot. |
| Cluster Service Writer | This VSS writer using a custom API, is used to back up and restore Cluster Service on Microsoft Cluster Server (MSCS). The Cluster Service is a component on Windows servers used to control server cluster activities on cluster nodes. It is fundamental to the operation of the cluster. | Custom restore method |
| COM+ REGDB Writer | This VSS writer using a custom API, is used to back up and restore COM+ Database Service. This service provides automatic distribution of events to subscribing COM+ components. | Custom restore method |

| Writer name | Description | Restore method |
|---|---|---|
| DHCP Jet Writer | This is a system writer, used to back up and restore DHCP Service database. DHCP Service provides dynamic IP address assignment and network configuration for Dynamic Host Configuration Protocol (DHCP) clients. | Files are restored after a reboot. |
| Event Log Writer | This is a system writer, used to back up and restore Event Logs. Event Logs are files where the Windows operating system saves information about events, such as starting and stopping services or the logging on and logging off of a user. | Files are restored after a reboot. |
| FRS Writer | This VSS writer using a custom API, is used to back up and restore File Replication Service data. File Replication Service is a multithreaded replication engine that replicates system policies and logon scripts stored in System Volume (`SYSVOL`). FRS can also replicate data for Distributed File System (Dfs), copy and maintain shared files and folders on multiple servers simultaneously. | Custom restore method |
| IIS Metabase Writer | This is a system writer, used to back up and restore Microsoft Internet Information Server (IIS). IIS 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). | Files are restored after a reboot. |
| MSDE Writer | This is a writer used to back up and restore Microsoft SQL Server 2000/2005. SQL Server is a database management system that can respond to queries from client machines formatted in the SQL language. | Refer to "MSDE writer restore specifics" on page 384. |
| Microsoft SQL 2005 Writer | This is a writer used to back up and restore Microsoft SQL Server 2005. SQL Server is a database management system that can respond to queries from client machines formatted in the SQL language. | Standard VSS restore and instant recovery. |

| Writer name | Description | Restore method |
|---|---|---|
| Microsoft Exchange Server Writer | This is a writer used to back up and restore Microsoft Exchange Server. Microsoft Exchange Server is a mail and groupware server. | Standard VSS restore and instant recovery. |
| Microsoft Virtual Server 2005 Writer | This is a writer used to back up and restore Microsoft Virtual Server 2005. Microsoft Virtual Server 2005 is a virtualization platform for Microsoft Windows servers. Data Protector supports live backup of individual virtual machines and the Virtual Server configuration, ensuring data consistency of the backup and restore. Hardware providers are not supported if a Virtual Server Machine is in online mode; use a software provider or put the Virtual Server Machine in offline mode. For details about Virtual Server online and offline mode, see the Microsoft Virtual Server documentation.<br>Cluster configurations are not supported, only individual nodes can be backed up. | Standard VSS restore and instant recovery. |
| Microsoft Hyper-V Writer | This is a writer used to back up and restore Microsoft Virtual Server 2008 Hyper-V configuration and individual or all virtual machines running on the server. Software and hardware providers are supported during online and offline backup.<br>Cluster-aware backups are not supported. | Standard VSS restore and instant recovery. |
| NTDS Writer | This is a system writer used to back up and restore Microsoft Active Directory on Windows servers. Active Directory Service is a Windows server directory service that enables you to manage data structures distributed over a network. For example, Active Directory Service stores information about user accounts, passwords, phone numbers, profiles, and installed services. It provides methods for storing directory data and making this data available to network users and administrators. | To restore Active Directory, boot into Directory restore mode. Files will be restored if they can be overwritten. |

| Writer name | Description | Restore method |
|---|---|---|
| Registry Writer | This VSS writer using a custom API, is used to back up and restore Windows Registry. Windows Registry is a database repository of information containing the Windows system configuration. | Custom restore method |
| Remote Storage Writer | This is a system writer used to back up and restore Remote Storage Service (RSS). RSS is used to automatically move infrequently accessed files from local to remote storage. Remote files are recalled automatically when the file is opened. | Files are restored after a reboot. |
| Removable Storage Manager Writer | This is a system writer used to back up and restore Removable Storage Manager Service. This service manages removable media, drives, and libraries. | Files are restored after a reboot. |
| System Writer | This is a system writer that backs up a specific set of Windows dynamic link libraries (DLL). | Files are restored after a reboot. |
| TermServLicencing Writer | This is a system writer that backs up Windows Terminal Services. These services provide a multi-session environment that allows client systems to access a virtual Windows desktop session and Windows-based programs running on the server. | Files are restored after a reboot. |
| WINS Jet Writer | This is a system writer, used to back up and restore Windows Internet Name Service (WINS). WINS is a dynamic replicated database service that can register and resolve NetBIOS names to IP addresses used on a TCP/IP network. | Files are restored after a reboot. |
| WMI Writer | This is a system writer, used to back up and restore Windows Management Instrumentation (WMI). WMI is a unified management infrastructure in Windows for monitoring system resources. | Files are restored after a reboot. |

# Backing up writers data

## Overview

To run backups and restores of the VSS writers and filesystems, you need to configure the Data Protector Microsoft Volume Shadow Copy Service integration backup specifications.

To configure the backup using the VSS integration, perform the following steps:

### Configuration steps

1. Configure disk arrays, devices, media, and media pools needed for the backup. See the online Help for instructions.
2. Create a Data Protector VSS backup specification specifying the VSS components to back up, the media and devices to which you want your data to be backed up, as well as the Data Protector backup options that define the behavior of your backup or restore session.

If you are using Data Protector VSS to perform a zero downtime backup (ZDB) of the MSDE Writer or Microsoft Exchange Server Writer, the Microsoft SQL Server or Microsoft Exchange Server has to be running on the application system for a backup to start.

## Replica types

The two basic families of replicas available with Data Protector are split mirror and snapshot.

### Split mirror replica

This type of replica is produced using mirroring technology, provided by HP StorageWorks Disk Array XP. Mirroring technology allows a duplicate (a mirror) of the filesystem or application data to be created during normal application use. When a replica is required, a disk from the S-VOL host group is first synchronized with the P-VOL, and after that the pair is split leaving a fixed copy, or an independent split mirror replica of the source volume. Depending on the configuration mode set in the VSS hardware provider, two types of mirror replicas can exist after the pair is split:

- If the provider is in *VSS compliant mode*, the replica is an independent copy or clone of the source volume with no pair relationship between the replica and its source volume.
- If the provider is in *resync mode*, the replica and its source volume are in a suspended mirror relationship, so you can only restore such a replica by re-synchronizing the replica with its source volume.

### Snapshot replica

This type of replica is produced using snapshot technology, provided by HP StorageWorks Enterprise Virtual Array.

With a snapshot replica, no synchronization is required before creating the fixed point-in-time copies, because a snapshot is a logical copy of pointers to the data rather than the data itself. Snapshot replicas are considered to be created almost instantaneously and are immediately available for use. However, with the snapclone type of replica, background processes may still run for some time afterwards.

A range of snapshot techniques is available, each with its own particular merits from the storage and data processing points of view. Data Protector supports the following snapshot replica types using the EVA hardware provider:

- Snapshot *with* pre-allocation of disk space (**standard snapshot**).
- Snapshot *without* pre-allocation of disk space (**vsnap** or virtually capacity-free snapshot).
- A full copy of the source volume (original virtual disk), independent of the original virtual disk (**snapclone**).

You select the replica type in the Data Protector GUI when creating a backup specification. However, you can choose between snapshot, clone, or default:

If you select **Snapshot (Differential)**, standard snapshots are created by default. To create vsnaps, configure the hardware provider to create vsnaps before creating a backup specification.

If you select **Mirror/Clone (Plex)**, a snapclone is created.

If you select **Default**, the replica type created will be the one which is configured in the hardware provider.

Note that for instant recovery, only snapclones are supported.

## HP StorageWorks EVA specifics

- When the cloning process for a source volume is in progress, another replica of the same source volume cannot be created.

- Only one type of target volumes per source volume can exist on an EVA disk array at the same time. For example, creation of a snapclone may fail if a snapshot or vsnap for the same source volume exists on the array. You should first delete such replicas using the Data Protector command `omnidbvss`, or using EVA Command View if they were not created by Data Protector.

### HP StorageWorks Disk Array XP specifics

- For backups created in the *resync mode*, the value of the **Number of replicas rotated** option is limited to a maximum of three replicas (S-VOLs) per source volume (P-VOL), provided that the `MU# range` option in the XP VSS hardware provider configuration is set to `0-2` or `0, 1, 2`. When the same P-VOL is used in more than one backup specification, the sum of values specified in the **Number of replicas rotated** option in all backup specifications should not exceed 3.

  This limitation also applies for the following backup scenario in a Microsoft Exchange Server CCR environment with the same backup system selected: Suppose you create three backups of a source volume on the production database system with one backup specification, and then create three backups of the source volume on the database copy system with a different specification. Together, this would create six replicas. However, Data Protector limits the number and supports only three replicas together in such a case. Therefore, you can choose to back up source volumes only on the production database system or only on the database copy system. Otherwise, after three replicas are created with one backup specification, the next backup with the second backup specification will fail.
- You can change the XP VSS hardware provider mode between different backup sessions using the same backup specification, but this is not recommended when you use replica rotation, because restore of such backup data may fail.

  For details, see "HP StorageWorks Disk Array XP considerations" on page 396.

## Creating backup specifications using GUI

The procedure below shows how to back up Microsoft VSS objects using the Data Protector GUI. Some writers have specific limitations. For writers specific limitations, see:

- For Microsoft Exchange Server specifics, see "Microsoft Exchange Server writer backup specifics" on page 370.

To create a new backup specification for the VSS integration, proceed as follows:

1. In the **HP Data Protector Manager**, switch to the **Backup** context.
2. In the Scoping Pane, expand **Backup Specifications**.

3. Right-click **MS Volume Shadow Copy Writers** and click **Add Backup**.

4. In the **Create New Backup** dialog box, select the backup type. You can choose between the following types:

   - **Local or network backup**

     This type is used for single host VSS backup. To perform a backup for instant recovery purposes, you need a hardware provider. Otherwise, no hardware provider is required for this type of backup.

   - **VSS transportable backup**

     Use this option to create shadow copies on the application system and present them to the backup system, which can perform the backup to tape.

     A hardware provider is required for this type of backup.

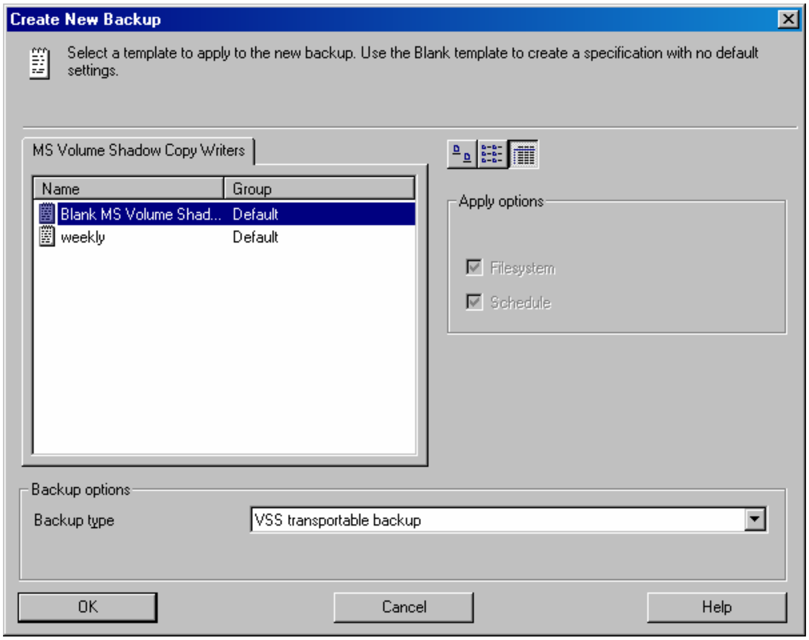

Select a template to apply to the new backup. Use the Blank template to create a specification with no default settings.

| Name | Group |
|---|---|
| Blank MS Volume Shad... | Default |
| weekly | Default |

Apply options

☑ Filesystem

☑ Schedule

Backup options

Backup type: VSS transportable backup

Figure 103 Selecting VSS transportable backup

5. Specify the following options:

   - The name of the application system.

     When backing up cluster-aware writers (such as SQL Server via the MSDE Writer, or Exchange Server in the LCR or CCR environment), specify the virtual server name given in the particular writer resource group as the application system.

   - For a transportable backup, the name of the backup system from where the shadow copy can be backed up to tape or where you want your shadow copies to be presented and mounted after the backup.
   - For local or network backup, the type of provider: select **Use hardware provider** if you want to enable instant recovery and specify other backup options.
   - To enable instant recovery, select **Track the replica for instant recovery**.
   - If you have selected **Track the replica for instant recovery**, specify the configuration check mode. Configuration check applies to the disk backups that are to be used for instant recovery and does not apply to the tape backup. For information on the options, see "Configuration check" on page 353.
   - If you do not select **Track the replica for instant recovery**, you can specify the replica type:

     **Default**

     The replica type created depends on the VSS hardware provider configuration.

     **Mirror/Clone (Plex)**

     A shadow copy, independent from its source volume is created:

     - With EVA hardware provider, snapclone is created.
     - With XP hardware provider, the replica created depends on the provider mode (clone in the case of VSS compliant mode, or mirror in the case of resync mode).

     **Snapshot (Differential)**

     A shadow copy, dependent on its source volume is created: by default, standard snapshot is created (with pre-allocated disk space).

     This replica type is not supported by XP hardware provider.

     A provider may support one or both types. If you select an unsupported replica type, the backup will fail.

     For more information on EVA provider replica types, see "Snapshot replica" on page 361.

   - Optionally, specify other replica management and mount options.

For descriptions of these options, press **F1**.

Click **Next**.

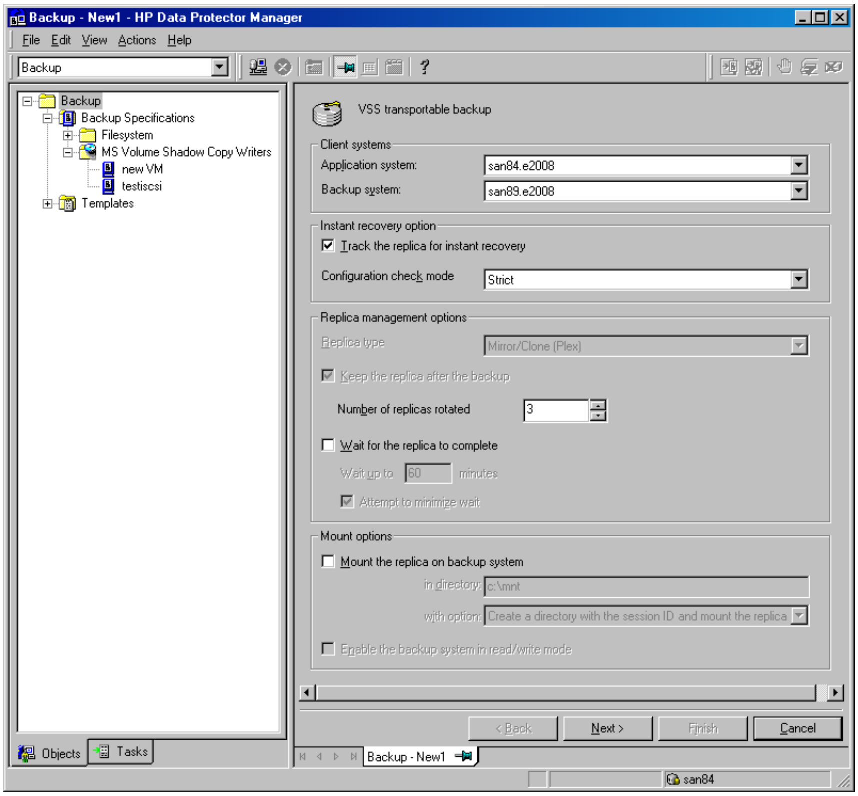

VSS transportable backup

Client systems

Application system: san84.e2008

Backup system: san89.e2008

Instant recovery option

☑ Track the replica for instant recovery

Configuration check mode: Strict

Replica management options

Replica type: Mirror/Clone (Plex)

☑ Keep the replica after the backup

Number of replicas rotated: 3

☐ Wait for the replica to complete

Wait up to 60 minutes

☑ Attempt to minimize wait

Mount options

☐ Mount the replica on backup system

in directory: c:\mnt

with option: Create a directory with the session ID and mount the replica

☐ Enable the backup system in read/write mode

### Figure 104 VSS transportable backup options

6. In this page, the selection of your VSS client is displayed.

   On Windows Server 2008, you can specify the **User and group/domain** options. For information on these options, press **F1**.

   Click **Next**.

7. Select the backup objects you want to back up. Make sure that in case of instant-recovery enabled sessions, you select all objects (for example all storage groups, all virtual machines ...) residing on a specific source volume that will be backed up.

   only a whole source volume to be backed up, regardless of the number of virtual machines residing on the volume.

   You can specify a **full client backup** by selecting the top-level item (the name of the client), a single writer or a writer's component backup by selecting a lower-level item.

   If full client is selected, Data Protector checks which writers exist on the client and backs up all of them at backup time.

NOTE:

On EVA disk array, to perform an instant recovery, do not select too many objects because they might reside on more than four source volumes. If the objects reside on more than four source volumes, create several backup specifications where objects do not reside on more than four source volumes.

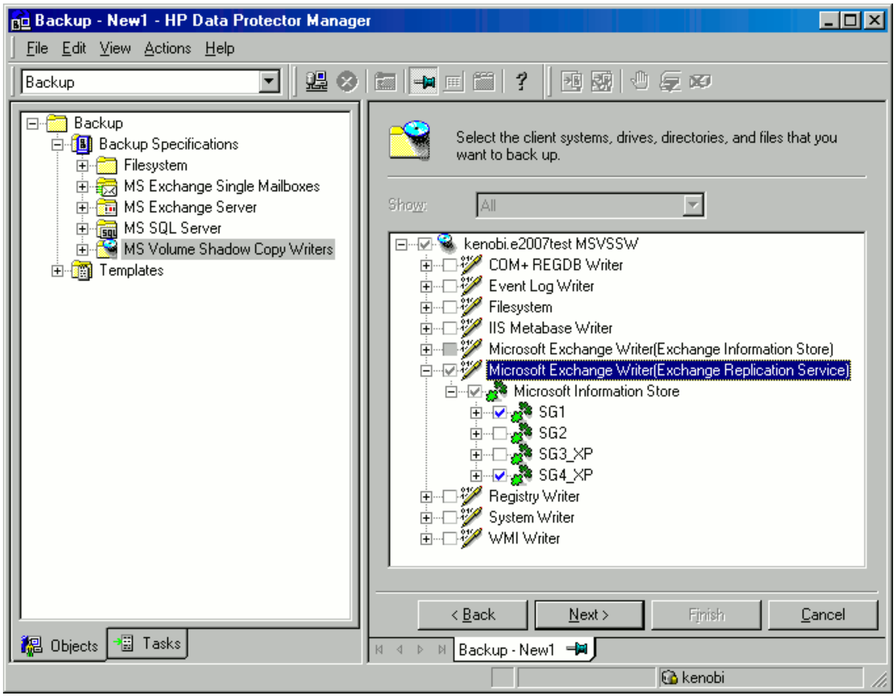
Select the client systems, drives, directories, and files that you want to back up.

Show: All

## Figure 105 Selecting Microsoft Exchange Server 2007 CCR copy backup objects

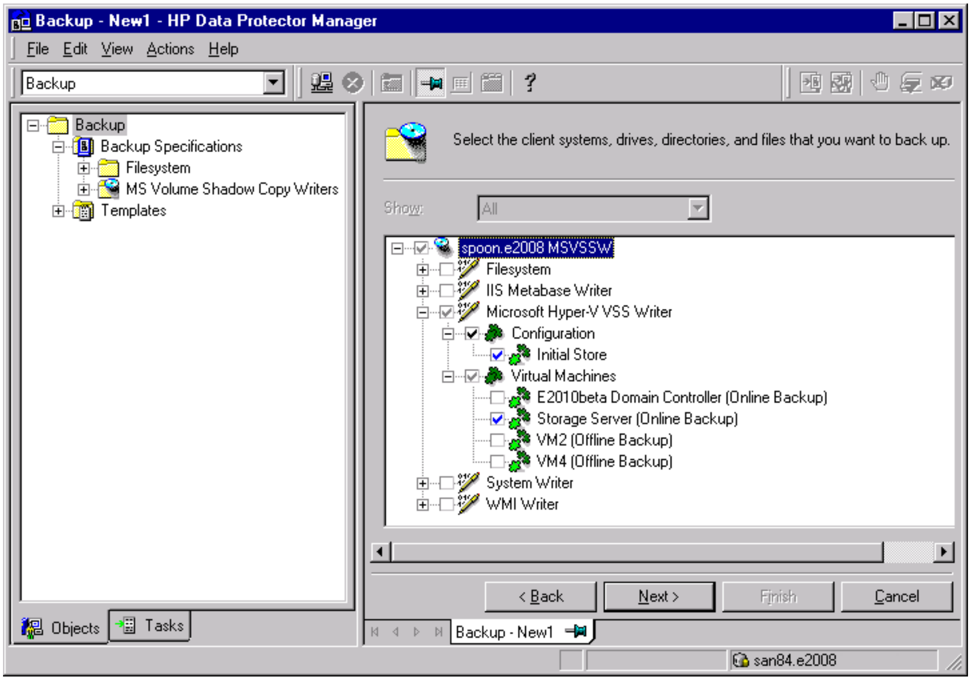

Select the client systems, drives, directories, and files that you want to back up.

- ☑ spoon.e2008 MSVSSW
  - □ Filesystem
  - □ IIS Metabase Writer
  - ☑ Microsoft Hyper-V VSS Writer
    - ☑ Configuration
      - ☑ Initial Store
    - ☑ Virtual Machines
      - □ E2010beta Domain Controller (Online Backup)
      - ☑ Storage Server (Online Backup)
      - □ VM2 (Offline Backup)
      - □ VM4 (Offline Backup)
  - □ System Writer
  - □ WMI Writer

## Figure 106 Selecting Microsoft Hyper–V VSS writer backup objects

If a writer requires all of its components to be backed up, lower-level items are automatically selected. If you select such a writer for backup, all its components will be backed up.

If a writer has no components to be backed up, it is not displayed in the list of writers, and is not backed up when the full client is selected.

The **Filesystem** item displays all mounted disks. If another disk is mounted to a directory on a disk, the parent disk name is displayed twice. The first name represents the parent disk name (for example `c:`), while the second name represents the container for the mountpoint (for example `c:\mnt\1`). To select the mounted disk, select the container for the mountpoint.

**Microsoft Exchange Writer:** Optionally, to specify options for consistency check of Microsoft Exchange Server Writer, right-click **Microsoft Exchange Writer** and click **Additional options**.

**Microsoft Exchange Server 2007 Writer in a CCR environment:** Optionally, to specify clustering options, right-click **Microsoft Exchange Writer** and click **Additional options**.

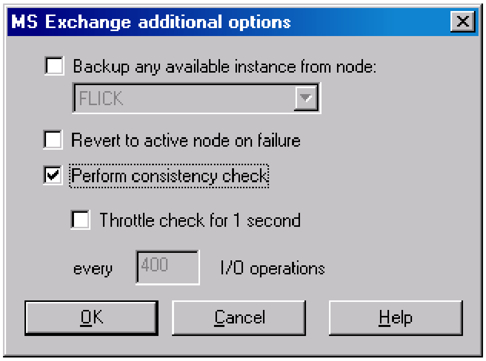

MS Exchange additional options

□ Backup any available instance from node:

FLICK

□ Revert to active node on failure

☑ Perform consistency check

□ Throttle check for 1 second

every 400 I/O operations

Figure 107 Additional options for Microsoft Exchange Server 2007 in a CCR environment

8. For backup to tape, select the devices you want to use for the backup to tape. Click **Properties** to set the device concurrency, media pool, and preallocation policy. For more information on these options, click **Help**. If you do not select a device, only backup to disk will be available.

IMPORTANT:

If you do not configure devices and do not select the option **Track the replica for instant recovery**, you will not be able to restore the data with Data Protector.

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

For detailed information on the object mirror functionality, see the online Help.

9. Following the wizard, select the backup options and schedule your backup.

   TIP:

   If you are not sure about selecting the backup options, keep the default values.

   For details about the options common to all Data Protector backup specifications, see the online Help.

10. Once you have defined all backup options and the schedule, you need to name and save the newly-created backup specification. You have now completed the creation of a Microsoft Volume Shadow Copy Writers backup specification.

11. You can review the newly-created and saved backup specification in the **Backup** context, under the specified group of backup specifications.

12. You can run backup using one of the following methods:
    - Schedule the backup of an existing Microsoft Volume Shadow Copy Writers backup specification using the Data Protector Scheduler.
    - Start an interactive backup of an existing Microsoft Volume Shadow Copy Writers backup specification.

## VSS specific backup options

### Table 24 VSS specific backup options

| Option | Description |
|---|---|
| **Pre-exec** | Specify a command that will be started by `vssbar.exe` on the application system before the backup. Do not use double quotes. Type only the name of the command, not the pathname. The command must reside in *`Data_Protector_home`*`\bin`. |
| **Post-exec** | Specify a command that will be started by `vssbar.exe` on the application system after the backup. Do not use double quotes. Type only the name of the command, not the pathname. The command must reside in *`Data_Protector_home`*`\bin`. |

## Microsoft Exchange Server writer backup specifics

The Microsoft Exchange Server Writer supports the following Microsoft Exchange backup types:

## Backup types

The Microsoft Exchange Server Writer supports the following Microsoft Exchange backup types:

- *Full* - backs up databases, transaction logs, and checkpoint files. The transaction logs are truncated.
- *Incremental* - backs up the transaction logs to record changes since the last full or incremental backup. The transaction logs are truncated.
- *Differential* - similar as incremental backup, but the transaction logs are not truncated.
- *Copy* - a Full backup, but the logs are not truncated. This type of backup is not intended for use in recovering failed systems.

## Limitations

- A combination of VSS snapshot backups and non-VSS backups (for example, incremental stream backups) is not supported.
- You can back up only the whole server or full storage groups. Single stores cannot be backed up.
- Circular logging must be disabled.
- With Exchange Server 2003, only one VSS backup session can be running on the same application system at once. With Exchange Server 2007, only one VSS backup session backing up a specific storage group can be running on the same application system at once. Consequently, if you start such a backup session while another one is in progress, the latter waits until the first one finishes.

## Recommendations

- With the HP StorageWorks EVA provider, an Exchange storage group should not reside on more than four source volumes. However, the recommended configuration is that transaction logs reside on one source volume and the stores on another source volume. This configuration enables you to perform a rollforward recovery if only the stores are lost.

  Create *one backup specification for each storage group* and schedule them in a row.

## Rollforward recovery

Rollforward of a storage group can be performed with incremental or differential ZDB to tape of the same objects as were backed up in instant recovery enabled sessions (ZDB to disk or ZDB to disk + tape). If the logs are still available on disk, the restore of last full backup will roll forward the store to the latest state (recorded in the logs).

## Consistency check

A backup of the Microsoft Exchange Server database is considered as successful only if the consistency check of the replicated datafiles succeeds. The consistency check is enabled by default. To disable the consistency check, click on a created backup specification, right-click **Microsoft Exchange Writer** in the Source tab, and then click **Additional options**. In this page, you can also specify to throttle the consistency check for a second after the specified number of input/output operations.

The consistency check can also be run before instant recovery.

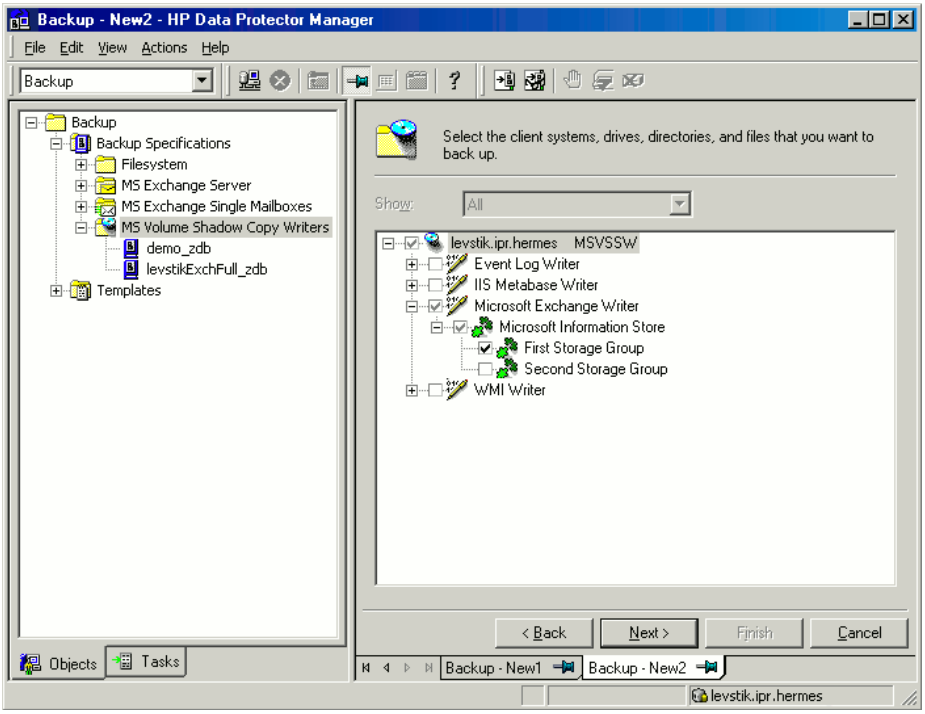

Select the client systems, drives, directories, and files that you want to back up.

- ☑ levstik.ipr.hermes MSVSSW
  - □ Event Log Writer
  - □ IIS Metabase Writer
  - ☑ Microsoft Exchange Writer
    - ☑ Microsoft Information Store
      - ☑ First Storage Group
      - □ Second Storage Group
  - □ WMI Writer

**Figure 108 Selecting Microsoft Exchange Server 2003 storage groups**

## Microsoft Exchange Server 2007 writer backup specifics

In LCR and CCR environments, the replicated storage groups are represented as a new instance of Exchange Server writer, **Exchange Replication Service**. The replicated storage groups are backed up in the same way as original (production) storage groups.

You can select any combination of storage groups for backup. However, you cannot select original and replicated storage group in the same backup specification. See Figure 109.

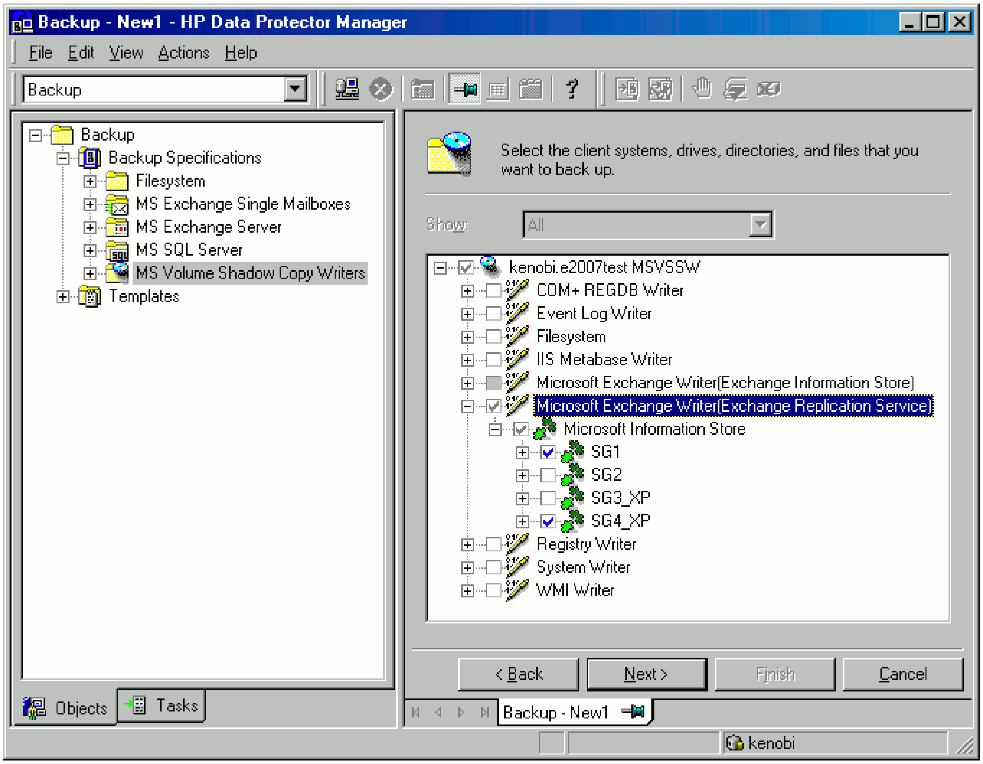

Select the client systems, drives, directories, and files that you want to back up.

- ☑ kenobi.e2007test MSVSSW
  - □ COM+ REGDB Writer
  - □ Event Log Writer
  - □ Filesystem
  - □ IIS Metabase Writer
  - ▨ Microsoft Exchange Writer(Exchange Information Store)
  - ☑ Microsoft Exchange Writer(Exchange Replication Service)
    - ☑ Microsoft Information Store
      - ☑ SG1
      - □ SG2
      - □ SG3_XP
      - ☑ SG4_XP
  - □ Registry Writer
  - □ System Writer
  - □ WMI Writer

## Figure 109 Selecting a replicated Microsoft Exchange Server 2007 storage group

In a CCR environment, a cluster node from which you want a backup to be performed can be selected, regardless of which instance (Information Store or Replication Service) resides on this node. If you select the cluster node, Data Protector backs up any available instance on this node ignoring the selection of the instance in the GUI even if, for example, you select the replicated storage group (Exchange Replication Service) as backup object.

To specify the cluster node from which you want to perform a backup of any instance residing on this node, right-click an Exchange writer and select the node in the MS Exchange additional options dialog box under **Back up any available instance from node**. See Figure 107.

When backing up a Replication Service instance, backup may fail due to any of the following reasons:

- The selected node is not available.
- The status of the storage group to be backed up is not "Healthy".
- Data Protector is not running on the selected node.
- `Vssbar.exe` cannot not be started on the selected node.

To avoid the session failing, select the option **Revert to active node on failure** in the same dialog box. The backup will be restarted on the original server (active cluster node) and the original storage group will be backed up. This option is ignored during backup of an Information Store instance.

## Prerequisites

- Before creating a backup specification and before running a backup, ensure that the Exchange Server 2007 copy status is “Healthy”. Otherwise, you cannot browse the objects to be backed up during creating backup specification, or the backup will fail. Note that the status “Initializing” is not acceptable.

## Limitations

- In CCR environments, if a replicated storage group (Replication Service instance) is selected for backup, no other VSS writer (for example, filesystem writer or SQL Server writer) can be selected to be backed up in the same backup session. Because in this case, two writers would be located on different systems, but Data Protector limits VSS backup sessions to involve only one system.
- In LCR and CCR environments, each storage group can have only one store.
- CCR clusters can have only two nodes.

## Considerations

- Transaction log files are truncated after each full or incremental Microsoft Exchange Server backup. The LCR and CCR clustering technologies, however, guarantee that logs that have not been replicated are not deleted. Thus, running backups in a mode that truncates logs may not actually free space. This may happen if replication of logs has not completed yet.
- Before using XP VSS hardware provider in the resync mode, consider all applicable limitations of this mode. The limitations are listed in “HP StorageWorks Disk Array XP specifics” on page 362 and “Limitations” on page 398. Due to these limitations, it is recommended to rather use XP provider in VSS compliant mode, or EVA VSS provider, or VSS software provider.

## Microsoft Hyper-V VSS writer backup specifics

With the Hyper-V VSS writer, it is possible to back up:

- The Hyper-V configuration
- Virtual machines

The following two types of backups are supported by the Data Protector VSS integration:

- *Online backup*

  Online backup of Hyper-V writer data is possible using a software provider or hardware provider.

  If using a hardware provider for an online backup process, Hyper-V writer creates a shadow copy, which is replicated together with the *vhd* file via the VSS hardware provider. Afterwards, the shadow copy is presented to the Hypervisor System. In the case of transportable backup, the shadow copy is afterwards unpresented from the Hypervisor System and presented to the backup system, specified in the backup specification.

  After restore of an online backup session, the virtual machine is always turned off, regardless of the state in which it was before the restore.

- *Offline backup*

  Offline backup is performed in the following cases:

  - The guest operating system – any other non-Microsoft guest operating system that is supported by Hyper-V – is not VSS enabled.
  - The guest operating system does not have Hyper-V VSS integration services installed.
  - The virtual machine to be backed up is turned off.

  Before offline backup, the virtual machine is automatically suspended (if not already) and resumed after the backup.

  After restore of an offline backup session, the virtual machine is suspended, regardless of the state it was in before the restore.

  The advantage of offline backup is that virtual machines are restored to the state they were in at backup time, including the state of applications running at the time of backup. This is possible because a virtual machine is in a suspended state during the backup.

### Prerequisites

- For online backup:

    - The guest operating system must have the Hyper-V VSS integration services installed and should not use dynamics disks.
    - The snapshot file (`avhd` file) needs to be configured on the same volume as the virtual disk file (`vhd` file).
    - The virtual machine to be backed up must be online.
- For offline backup, if the prerequisites for online backup are met, a virtual machine to be backed up must be put in the offline or suspended state manually.

### Limitations

- Cluster-aware backups are not supported.
- ZDB to disk, ZDB to disk+tape, and instant recovery of a cluster node is not supported.
- Only whole virtual machines can be backed up or restored. However, with ZDB to disk and ZDB to disk + tape only a whole source volume can be backed up, regardless of the number of virtual machines residing on the volume.
- Backup of virtual machines configured to use physical disks is not supported.
- Hyper-V writer data backup is not supported using the XP VSS hardware provider.

### Backup from a physical cluster node

When backing up from a cluster node consider the following:

- Offline backups trigger a failover.

  During offline backups the virtual machine to be backed up is suspended for a moment. The cluster server recognizes this as a failure and initiates the cluster failover. To avoid such a failover, perform one of the following:

    - Run only online backups.
    - Before running an offline backup, manually put the virtual machine into the "Saved" state using Failover Cluster Administrator.
- Whenever a failover happens, you need to use another backup specification.

  After a failover, the hostname where the virtual machine is running changes. Since this change is not reflected in the original backup specification, you need to create a new backup specification to specify the new hostname as the application system name.

## Scheduling the backup

For more detailed information on scheduling, refer to the online Help index: "scheduled backups".

To schedule a Microsoft Volume Shadow Copy Writers backup specification, perform the following steps in the Data Protector GUI:

1. In the **HP Data Protector Manager**, switch to the **Backup** context.
2. In the Scoping Pane, expand **Backup**, then **Backup Specifications**. Click **MS Volume Shadow Copy Writers**.

   A list of available backup specifications is displayed in the Results Area.
3. Double-click the backup specification you want to schedule and click the **Schedule** tab to open the **Schedule** property page.
4. In the **Schedule** property page, select a date in the calendar 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 for ZDB sessions that are to be used for instant recovery is set to **Full**.

   In the case of a ZDB-to-disk or a ZDB-to-disk+tape session, specify the **Split mirror/snapshot backup** option.

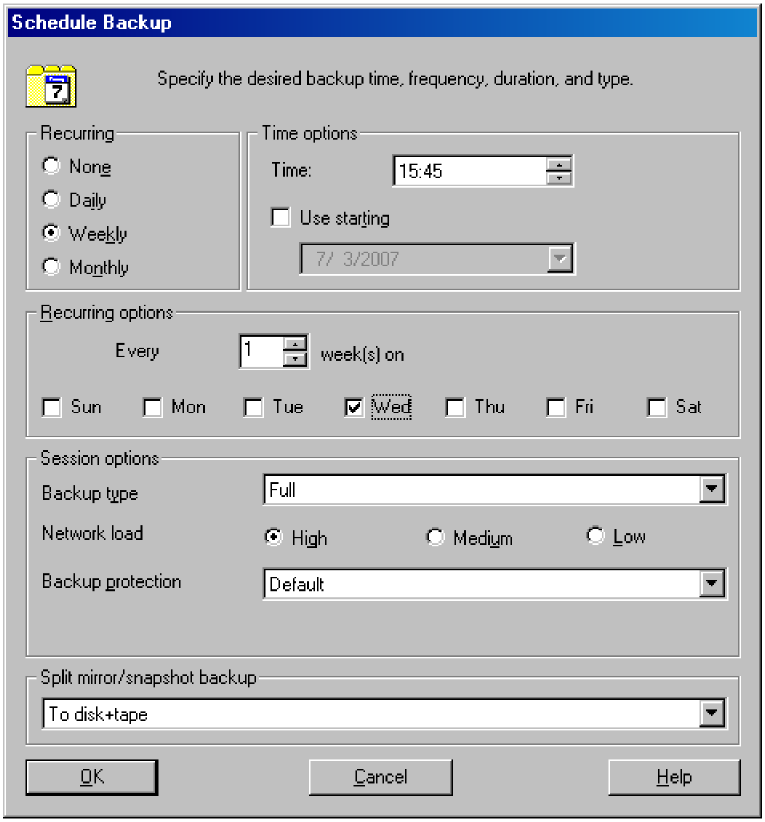

Schedule Backup

Specify the desired backup time, frequency, duration, and type.

Recurring
○ None
○ Daily
● Weekly
○ Monthly

Time options
Time: 15:45
□ 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: ● High ○ Medium ○ Low
Backup protection: Default

Split mirror/snapshot backup
To disk+tape

**Figure 110 Scheduling a backup**

6. Click **OK** to return to the **Schedule** property page.
7. Click **Apply** to save the changes.

## Running an interactive backup

An interactive backup can be started using the Data Protector GUI by following these steps:

1. In the **HP Data Protector Manager**, switch to the **Backup** context.
2. In the Scoping Pane, expand **Backup**; then expand the **Backup Specifications** and the **MS Volume Shadow Copy Writers** items.
3. Right-click the backup specification you want to run, and then select **Start Backup** from the pop-up menu.

   The **Start Backup** dialog box appears.

   Select the backup type and the network load (**High**, **Medium**, or **Low**).

   Note that the backup type for ZDB sessions that are to be used for instant recovery is set to **Full**.

   In the case of a ZDB-to-disk or a ZDB-to-disk+tape session, specify the **Split mirror/snapshot backup** option.

   Refer to online Help for a description of network load.
4. Click **OK**. Upon successful completion of the backup session, a `Session Completed Successfully` message appears.

# Restoring writers data

You can restore the Data Protector Microsoft Volume Shadow Copy Service integration backup objects using the Data Protector GUI.

Data Protector offers two methods for restoring the writers:

- From backup media to the application system on LAN (standard restore). Refer to "Restore procedure" on page 380.
- Using the instant recovery functionality with the Microsoft Exchange Server and Microsoft SQL Server ZDB integrations. Refer to ???.

NOTE:

Data Protector first restores the Writer Metadata collected during the backup time. This metadata contains the information about the backup components and the restore method. Data Protector performs restore according to the restore method specified by the writers.

## Limitations for custom restore

- Data Protector Microsoft VSS integration does not automatically provide any restore method for writers requesting custom restore. If a writer specifies custom restore method, it is only possible to restore the writer's data as plain files using the Data Protector file restore functionality. You can use the `Restore Into` option to specify an alternate restore path for these plain files. You can then perform the custom restore from these plain files manually. For information on writer’s custom restore, refer to the writers documentation.

## Restore procedure

The procedure below shows how to restore Microsoft VSS components using the Data Protector GUI. Some writers require custom restore procedures and/or have specific limitations. See also the appropriate sections:

- For Microsoft Exchange Server Writer specifics, see “Microsoft Exchange Server writer restore specifics” on page 386.
- For MSDE Writer specifics, see “MSDE writer restore specifics” on page 384.

To restore Microsoft VSS objects using the Data Protector GUI, proceed as follows:

1. In the **HP Data Protector**, switch to the **Restore** context.
2. Expand **Restore Objects**, expand **MS Volume Shadow Copy Writers**, expand the client from which you want to restore the data, and then click **MS Volume Shadow Copy Writers**. In the Results Area, a list of writers, which were backed up on this client, is displayed.
3. Select the Restore mode:
   - **Restore components**

     With this option selected, whole components are restored using the Volume Shadow Copy Service. Individual files cannot be selected for restore.
   - **Restore files to temporary location**

     With this option, you can select individual files or a group of files that were backed up using the selected writer. The files are restored using the Data Mover Agent and not the Volume Shadow Copy service.

4. In the Results Area, select the writers or writers' components (for component restore) or files or a group of files (for file restore mode).

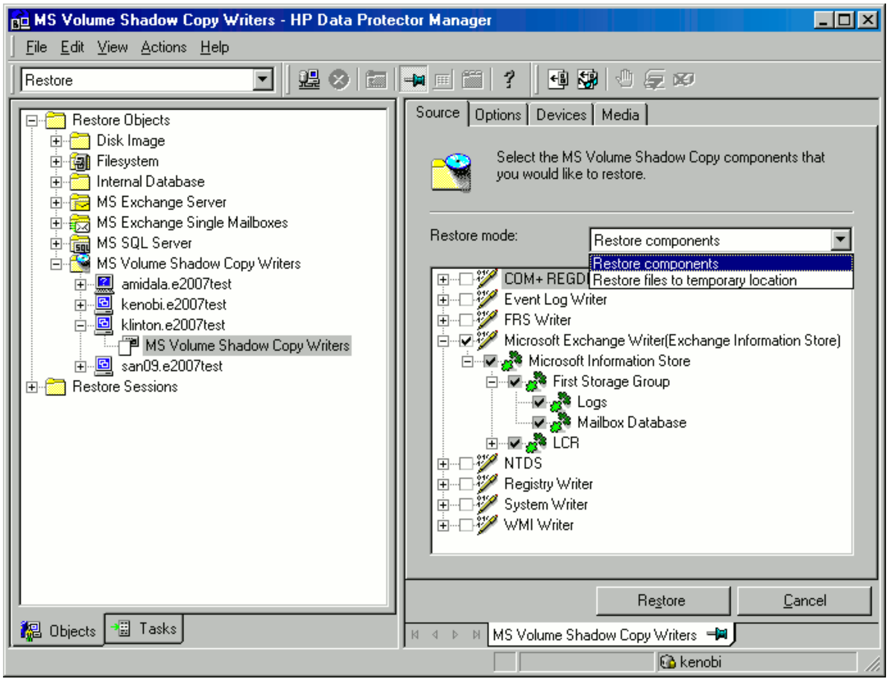

Select the MS Volume Shadow Copy components that you would like to restore.

Restore mode: Restore components

Figure 111 Restore objects

You can select the top-level item (full writer restore) or only specific components. If you select a full writer restore, but some components of this writer were not backed up in the same session, the unavailable components are shaded and you cannot select them. To select the version (the date of a backup), right-click the object name and click **Properties**. The last available backup version is selected by default, however, you can select a different version from the drop-down list.

**Exchange Server Writer:** Optionally, to specify options for the consistency check of a Microsoft Exchange writer, right-click the writer and click **Additional options**.

You can perform *point-in-time* or *rollforward* restore of Microsoft Exchange Server writer. For more information, see Microsoft Exchange Server writer restore specifics.

**Exchange Server 2007 Writer:**

If you restore an LCR or CCR copy to the original location, the restore will be performed to the original database (Exchange Information Store) and not to the database copy (Exchange Replication Service).

**a.** Right-click a storage group, a store, or logs and click **Restore as**.

**b.** In the MS Exchange additional options dialog box, select the target location for the components you want to restore: target server, target storage group, and target stores. The following options are available:

- **Restore to a different store**

  This option is selected by default.

  Select this option to select the target stores for each store to be restored (original stores). First, select the target system from the Target server name drop-down list, and then choose the desired pairs of stores by selecting entries in the Original and Target drop-down lists. Note that only stores cannot be restored, so logs are automatically selected in the restore session to different location.

- **Restore to a non-Exchange location**

  Select this option to restore your data to a non-Exchange location. In this case, the restored data will not be managed by Exchange Server and a Recovery Storage Group (RSG) will not be created. You can manually create an RSG after the restore session completes. First, select the target system from the Target server name drop-down list, and then choose the desired stores using the Original drop-down list.

- **Restore to a non-Exchange location and create RSG**

  Select this option to restore your data to a non-Exchange location. After restore, Data Protector will create a Recovery Storage Group called `DP RSG` on the target server. The selected stores and logs will be restored to this recovery group. First, select the target system from the Target server name drop-down list, and then choose the desired stores using the Original drop-down list.

By default, the storage group is restored in the `C:\Omni` directory. To select another location use the **Restore into location** option. Click **Browse** and select the desired location.

IMPORTANT:

The selected directory must be either empty or you can specify to create a new directory. If the directory is not empty, the restore session fails.

Note that you can only specify the restore location for the storage group and not for a specific store.

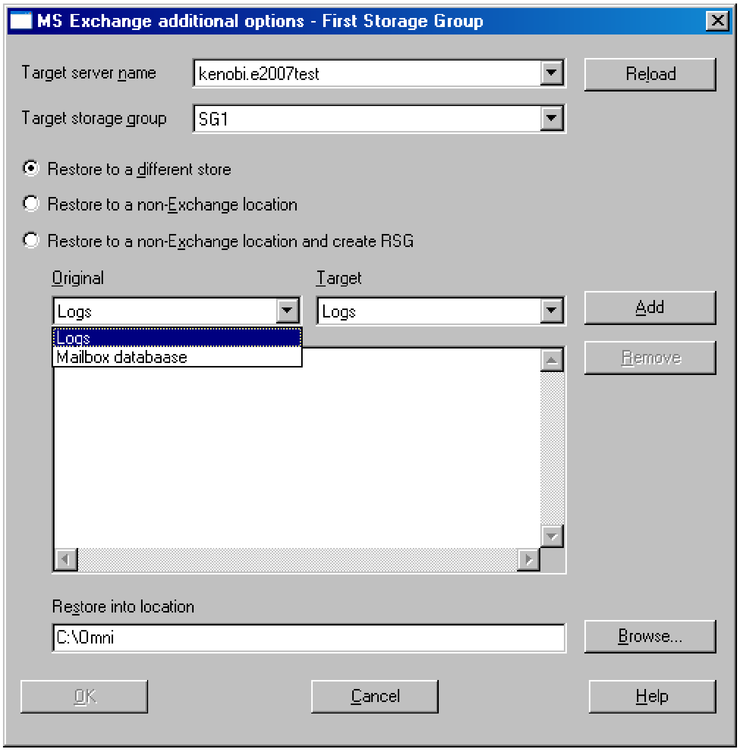

MS Exchange additional options - First Storage Group

Target server name: kenobi.e2007test

Target storage group: SG1

● Restore to a different store

○ Restore to a non-Exchange location

○ Restore to a non-Exchange location and create RSG

Original: Logs

Target: Logs

Logs
Mailbox databaase

Restore into location: C:\Omni

**Figure 112 Restore to different location options (Exchange Server 2007 Writer)**

5. In the **Options** property page, select the MS Volume Shadow Copy specific restore options. Refer to “Restore options” on page 384.

6. In the **Devices** and **Media** property pages, the devices and media for restore are automatically selected.

   Note that you can change the device used for the restore. Therefore, you have the possibility of using a different device for a restore than the one that was used for the backup. See the online Help index: “selecting devices for restore”.

7. Click **Restore**. Review your selection, and then click **Finish** to start a restore session.

   The restore session messages are displayed in the Results Area.

8. If you are restoring a VSS writer that requires a custom restore, continue manually, using the writers specific methods, if it is provided by a writer. Refer to the writers’ documentation.

## Restore options

The following restore options are specific to the Data Protector MS Volume Shadow Copy integration.

NOTE:

Do not use these options for Microsoft Exchange Server 2007 writer since there are other Exchange Server 2007 specific options provided for restore to another location.

**Restore to another client**

By default, the components or files are restored to the client from which the application data was backed up. However, you may restore the data to another VSS client if you specify the **Restore to another client** option. The new target Microsoft VSS client must be a part of the Data Protector cell. For component restore, it must also run on the same platform and have the MS Volume Shadow Copy Integration software component installed. For file restore, the MS Volume Shadow Copy Integration software component is not required.

**Restore into the following directory**

By default, you restore the data to the same directory from which it was backed up (it can be on the original client or on some other client which you selected).

However, if you specify the **Restore into the following directory** option, your data will be restored to another directory. When defining the restore location, you can specify the path to the directory where you want to restore your data.

# MSDE writer restore specifics

MSDE writer is used to back up and restore Microsoft SQL database.

IMPORTANT:

Before restoring the SQL system databases (`master`, `model`, `msdb` and `pub`), you have to stop the SQL service.

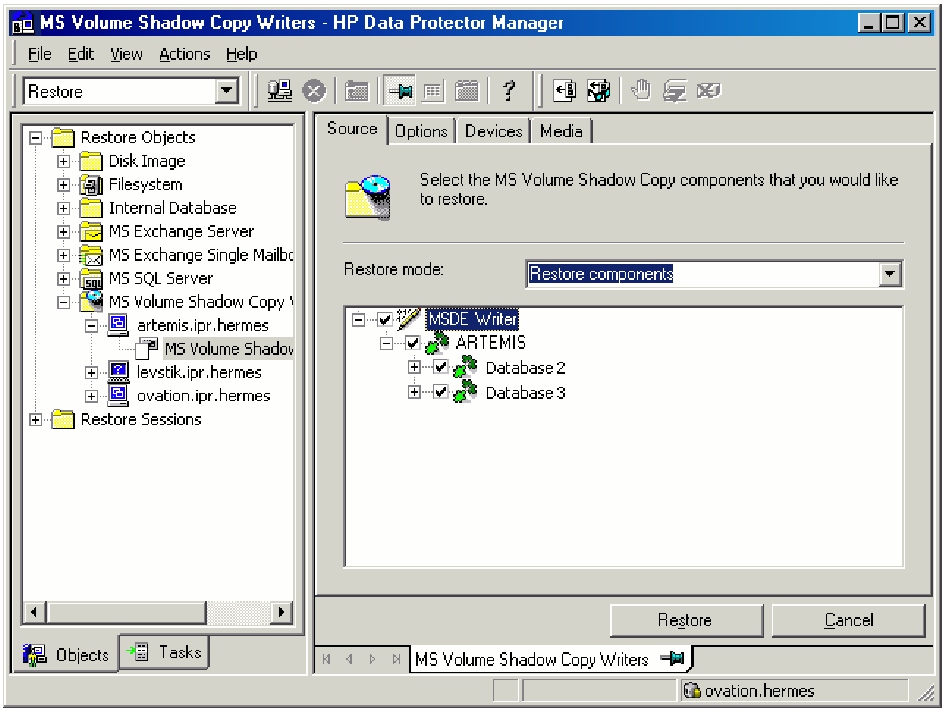

Select the MS Volume Shadow Copy components that you would like to restore.

Restore mode: Restore components

### Figure 113 MSDE writer

When you expand the MSDE Writer item in the Results Area, all Microsoft SQL Server instances are displayed. Each instance contains all databases it includes. System databases (`master`, `model`, `msdb` and `pub`) are always listed there.

IMPORTANT:

If system databases are restored, the whole internal database structure will be changed.

NOTE:

Only point-in-time restore is possible. Rollforward restore is not supported.

User databases will be restored only if it is possible to overwrite the files. MSDE writer will take the user databases offline before the restore, while SQL service will have to be stopped manually in order to restore the system databases.

## Microsoft Exchange Server writer restore specifics

Microsoft Exchange Server Writer is used to restore Microsoft Exchange Server database files.

When restoring from a Microsoft Exchange backup, the following two scenarios are possible:

- One or more databases are corrupted, but the log files are not damaged. In this case the database is restored and transaction logs are applied. See "Rollforward recovery from the loss of one or more databases" on page 386.
- The log files are corrupted or missing. In this case all databases and log files need to be restored. A rollforward recovery of the database is not possible. See "Point-in-time restore after loss of a log file" on page 387.

### Microsoft Exchange Server 2003 limitations

- Shadow copies cannot be restored to alternate locations on the backup system.
- You cannot restore the shadow copy to the Recovery Storage Group.
- Rollforward recovery cannot be performed after a point-in-time restore.
- Restore to Microsoft Exchange Server 2007 from backup made with Microsoft Exchange Server 2003 is not supported.

## Rollforward recovery from the loss of one or more databases

For a rollforward recovery:

1. In case of Microsoft Exchange Server 2003 writer, dismount all stores from the storage group in which the target store resides using Microsoft Exchange System Manager.

2. In the Data Protector GUI, switch to the **Restore** context. Expand **Restore Objects** and **MS Volume Shadow Copy Writers** and select the client from which you want to restore the data.

   In the Results Area, expand Microsoft Exchange Writer and select the stores you want to recover. The **Logs** component is shaded and cannot be selected. You cannot select versions of individual stores, because a rollforward recovery is performed only to the current storage group state.

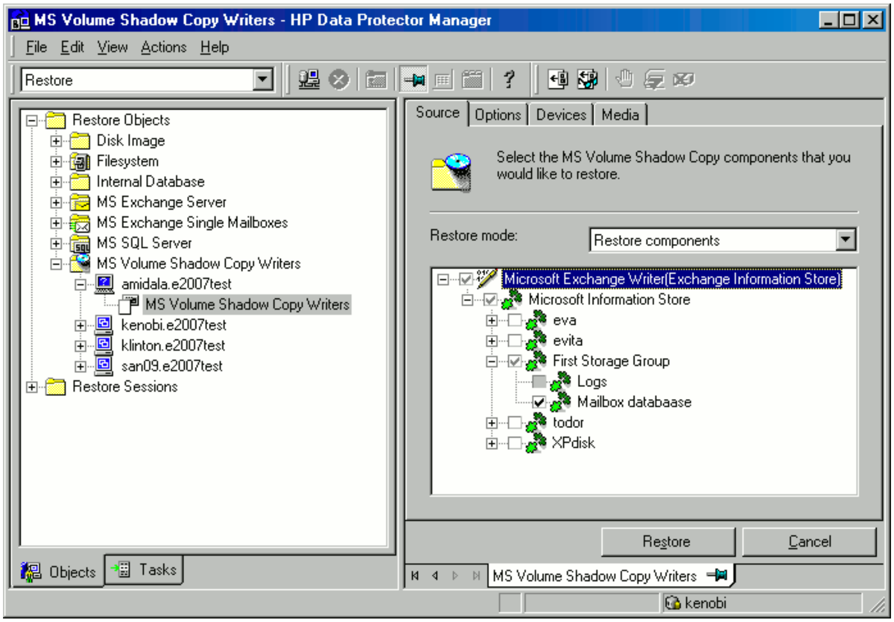

Select the MS Volume Shadow Copy components that you would like to restore.

Restore mode: Restore components

### Figure 114 Selecting Microsoft Exchange Server 2007 stores for rollforward recovery

3. Proceed as with general VSS writer restore. See "Restore procedure" on page 380 for the general VSS writer restore procedure.

4. Mount all stores from the storage group in which they reside using Exchange System Manager. Selected stores are recovered.

## Point-in-time restore after loss of a log file

To perform a point-in-time restore:

1. In case of Microsoft Exchange Server 2003 writer, start Exchange System Manager and check if the storage group is already dismounted. If not, dismount the whole group.

2. Switch to the **Restore** context. Expand **Restore Objects** and **Microsoft Volume Shadow Copy Writers** and select the client from which you want to restore the data.

   In the Results Area, expand Microsoft Exchange Writer and select the whole storage group. Do not select individual stores.

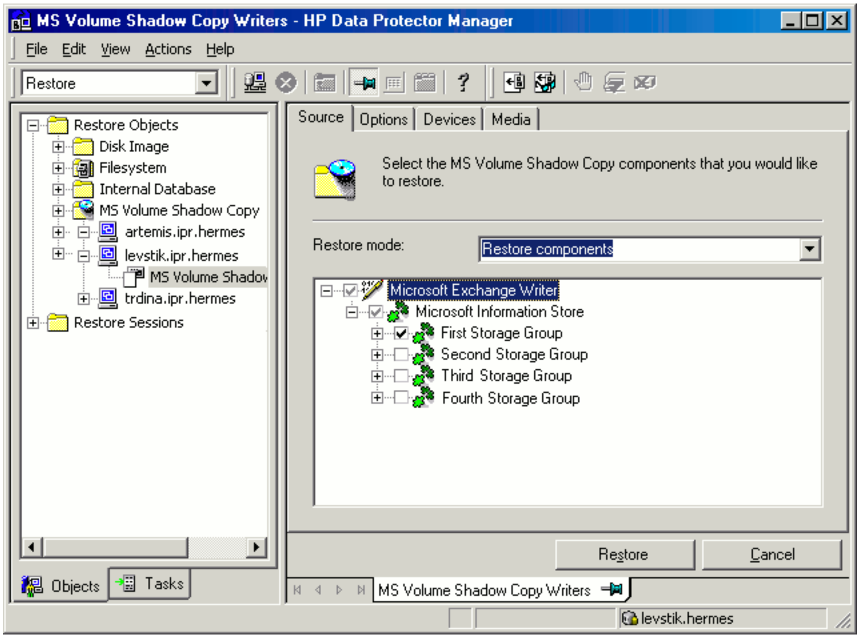

MS Volume Shadow Copy Writers - HP Data Protector Manager

Select the MS Volume Shadow Copy components that you would like to restore.

Restore mode: Restore components

- ☑ Microsoft Exchange Writer
  - ☑ Microsoft Information Store
    - ☑ First Storage Group
    - □ Second Storage Group
    - □ Third Storage Group
    - □ Fourth Storage Group

**Figure 115 Selecting Microsoft Exchange Server 2003 stores for point-in-time restore**

3. Proceed as with general VSS writer restore. See "Restore procedure" on page 380 for the general VSS writer restore procedure.

4. Mount the stores from the storage group in which the target stores reside using Exchange System Manager. All stores are mounted and put in the state as they were at the last selected full, incremental, or differential backup.

## Microsoft Exchange Server 2007 writer restore specifics

Microsoft Exchange Server 2007 writer offers an additional restore option:

- Restore to a different location than the location from where the backup was made (original location). After that, a recovery storage group can be created by Data Protector.

The following restore scenarios are supported:

- Restore of a storage group or individual stores to the original location.
- Restore of LCR and CCR copies (storage groups or individual stores) to the location where the original database resides.
- Restore of a storage group or individual stores and LCR and CCR copies (storage groups or individual stores) to a different location:
  - *Restore to a different storage group*

    Data can be restored to a **different storage group**. Exchange Server can then replay the restored logs from the original storage group on the target storage group to bring the restored database up to date. If you perform restore to a different storage group, you can access single mailboxes or individual e-mail messages on the alternate location without compromising availability of the original storage group with a potentially unsuccessful restore operation.
  - *Restore to a non-Exchange location*

    Data can be restored to a **non-Exchange location**. On Data Protector session level, this option is already available if you use **Restore into the following directory** option. On a storage group level, however, a storage group can be restored to a target directory, and another one to a different storage group in the same session if the target client system is the same. Individual stores can also be restored to a non-Exchange location. During the restore, data is restored into target component destination path.

    In this scenario, the restored data is not managed by Exchange Server.
  - *Restore to a non-Exchange location with creation of Recovery Storage Group (RSG)*

    Data can be restored to a non-Exchange location and a **Recovery Storage Group** can be created for the databases afterwards. Exchange Server recognizes the Recovery Storage Group `DP RSG` created by Data Protector.
  - *Restore to Recovery Server*

    If the Exchange Server system is unusable, an entire storage group can be restored to another Exchange Server system (**Recovery Server**) and some other storage group. If the target system is identical to original (the same host name, storage group names, store names, and the same directory structure), restore can be started as normal restore to original location. In case of differences, restore to a different location must be used.

  To restore to a different location, right-click the storage group, or store, or logs to be restored and click **Restore as**. Select target locations for all the components you want to restore.

You can restore different storage groups to original and different location in the same restore session if the target restore client system is the same for all storage groups involved.

During restore to original location, all stores in the storage group are dismounted, even if only one store from the storage group will be restored. After restore, the stores are left dismounted.

## Prerequisites

- For restore to a different location, `Inet` must run under a domain account which needs to be member of the following groups on the local system:
  - Administrators
  - Exchange Server Administrators
- While configuring restore to a different storage group in the Data Protector GUI, only mounted stores are displayed in the **Target** drop-down list. To enable all stores for target selection, mount the dismounted stores and click **Reload** in the **MS Exchange additional options**.

IMPORTANT:

During a restore to a non-Exchange location and creation of a Recovery Storage Group (RSG), Data Protector deletes an RSG that may already exists at the target location.

## Limitations

Due to Microsoft Exchange Server limitations:

- Restore to an LCR or CCR storage group copy (Exchange Replication Service instance) is not supported. Restore of an LCR or CCR copy to original location will be performed to the location where original database resides (Exchange Information Store).
- All stores in an RSG must originate from the same storage group.
- You cannot mount public folders to an RSG.
- Because only one storage group in an RSG is supported, only one storage group can be restored with the option **Restore to a non-Exchange location and create RSG**.
- When restoring to a different location, you cannot restore a specific database store without the transaction logs. When a store is selected for restore, logs are automatically included in the restore.

- If an RSG for the storage group you are restoring already exists on some system other than the target system, the RSG is not automatically deleted and so the creation of an RSG on the target system fails. In this case, you should manually delete the RSG on the other system.

## Microsoft Hyper-V VSS writer restore specifics

Restore of Hyper-V writer data is possible to the original or to a different location. During a restore to a different location, the Hyper-V writer checks whether a virtual machine with the same identity already exists on the system. If it does, the Hyper-V writer removes the virtual machine from the system before restore and imports the restored virtual machine. If such a virtual machine does not exist on the system, it means that a restore session to this system is in progress or that the virtual machine was already removed. It is possible to perform restore or instant recovery of Hyper-V virtual machines to any Hyper-V system, which has a Hyper-V writer.

### Restore from a physical cluster node

ZDB to tape sessions from a physical cluster node can be restored to any cluster node using a standard restore procedure. To successfully restore such a session to a cluster node, perform the following steps:

1. Delete the cluster group of the virtual machine to be restored using Failover Cluster Management.
2. Perform standard restore of the virtual machine on the cluster node where the virtual machine disk is active. For the standard restore procedure, see "Restore procedure" on page 380.
3. Re-create the virtual machine cluster group using Failover Cluster Management.

## Instant recovery

For general information and concepts on instant recovery, see the *HP Data Protector zero downtime backup concepts guide* and *HP Data Protector zero downtime backup administrator's guide*.

The Data Protector VSS integration provides instant recovery of MSDE Writer, SqlServer Writer, Microsoft Exchange Server 2003 Writer, and Microsoft Exchange Server 2007 Writer.

You can *selectively* restore separate writer's components only if they reside on separate target volumes. Instant recovery of separate components also requires that the target volumes with components data should not contain any other data. If you do not select all writer's components residing on the target volumes to be restored, the restore fails.

Before starting an instant recovery, Data Protector checks:

- Whether the data that is required for the components restore is available. If the check fails, the instant recovery session also fails, ensuring that no data that does not belong to the components is lost.
- If any changes that are not controlled by Data Protector happened in a disk array environment (for example, the target volumes to be restored are presented to some other system as it is recorded in the VSSDB, or the source volumes to be replaced are not presented on the application system). If Data Protector detects such changes (by comparing the state of the environment with data in the VSSDB), the instant recovery session aborts. However, you can force the session to continue by setting the `OB2VSS_IGNORE_BACKUP_DISK_CHANGES` and `OB2VSS_IGNORE_SOURCE_DISK_CHANGES` `omnirc` variables to `1`.

For Exchange Server 2007 writer specifics (prerequisites and limitations), see "Microsoft Exchange Server 2007 instant recovery prerequisites and limitations" on page 398.

## Overview

Instant recovery restores data directly from a replica to source volumes. All data in the replica is restored (regardless of selections during backup). For instant recovery concepts, see the *HP Data Protector zero downtime backup concepts guide*.

Instant recovery can be performed:

- **Using Microsoft Virtual Disk Service (VDS)**

  With this method, switch of a replica from the specified backup session with the source volumes is performed. Information about the replica is deleted from the VSSDB, therefore it is not possible to perform another instant recovery using that replica. For more information about the switch of disks method, see "Switch of disks" on page 394.

  In case of the Data Protector XP integration, this type of instant recovery can be performed if backup was run using the XP provider in the VSS compliant mode.

- **Using a Data Protector disk array integration** (HP StorageWorks EVA SMIS-S Agent or HP StorageWorks XP Agent)

## Limitations

- Instant recovery cannot be performed after a point-in-time recovery. To be able to perform instant recovery, you first need to run backup with the instant recovery options set.
- The number of replicas available for instant recovery is limited by **Number of replicas rotated**, which sets the size of the replica set. You can view these replicas in the GUI, in the **Instant Recovery** context by expanding **Restore Sessions**. Replicas are identified by the backup specification name and the session ID. Other information, such as time when a particular replica was created, is also provided. Alternately, you can use the `omnidbvss` command to list sessions. For information about `omnidbvss`, see the *HP Data Protector command line interface reference*.
- Concurrent instant recovery sessions utilizing the same backup system or the same application system are not supported. Restores utilizing the same backup system or the same application system must be performed serially.
- For instant recoveryof Exchange Server 2007 data to a different location, the Microsoft Virtual Disk Service, HP StorageWorks EVA, and HP StorageWorks Disk Array XP restore options are not available.

For additional Exchange Server 2007 writer limitations, see "Microsoft Exchange Server 2007 instant recovery prerequisites and limitations" on page 398.

## Instant recovery using Microsoft Virtual Disk Service

The following instant recovery method is possible using VDS:

### Switch of disks

With this method, the source volumes are unpresented and the target volumes (replica) are presented in the place of the source volume. You can select to retain the old source volumes. However, you cannot retain the replica and perform another instant recovery from this replica.

If the restore session fails or is aborted, the VSS integration reverts the application and backup environment back to the state in which they were before the instant recovery.

### Advantages

- Instant recovery is very fast.
- The original source volume can be retained after instant recovery.

- By selecting the **Restore using Microsoft Virtual Disk Service** option, you can restore your data without using HP StorageWorks EVA SMIS-S Agent or HP StorageWorks XP Agent.

### Disadvantages

- It is not possible to perform another instant recovery using the same replica.
- This method of instant recovery changes the physical location of the application data, since, after instant recovery, a replica becomes the source volume. The application starts running on the physical disks that were used for backup. If the selected replica and its source volume belong to separate disk groups, the two disk groups are also switched during instant recovery.

To use this method of instant recovery, in the Data Protector GUI, select the **Restore using Microsoft Virtual Disk Service** option.

## Instant recovery using HP StorageWorks EVA SMIS-S Agent

Depending on the instant recovery options you select in the Data Protector GUI, instant recovery can be performed in three different ways.

### Switch of disks

For description on this method, see "Switch of disks" on page 394.

To use this method of instant recovery, in the Data Protector GUI, select the **Restore using HP StorageWorks EVA** and **Switch to the replica** options.

### Copy of replica data with the source volume retained

With this method, a new copy from a replica is created in the same disk group in which the source volumes reside. When the copy is finished, a switch of disks between the source and the new copy is performed. The old source volumes are retained.

If the restore session fails or is aborted, the VSS integration reverts the application and backup environment back to the state in which they were before the instant recovery.

You can also select to retain the replica and thus enable to perform another restore from this backup.

### Advantages

- The original source volumes are retained after restore.

- Disk group of the source volume is not changed after instant recovery.
- Another restore from the same backup is possible.

### Disadvantages

- Restore is not as fast as with the "switch of disks" method.
- Physical location of the application production data changes.

To perform this type of restore, select the **Copy replica data to the source location** option and leave selected the **Retain source for forensics** option in the Data Protector GUI.

### Copy of replica data with the source volume not retained

With this method, the source volumes are directly overwritten by the replica by copying data from the target volumes back to the source volumes. The old source volumes are not retained and if the restore session fails after the copy process has already started, the original application data residing on the source volumes is lost. If you try to abort the session at this point, the abort operation is rejected and the session continues.

You can select to retain the replica and thus enable to perform another restore from this backup.

### Advantages

- Disk group of the source volume is not changed after instant recovery.
- Physical location of the application production data remains the same.
- Another restore from the same backup is possible.

### Disadvantages

- Restore is not as fast as with the "switch of disks" method.
- The original source volumes are lost during restore.

To perform this type of restore, select the **Copy replica data to the source location** option and clear the **Retain source for forensics** option in the Data Protector GUI.

### HP StorageWorks EVA considerations

- If you restore using HP StorageWorks EVA, before instant recovery, it is recommended to check if the SMI-S CIMOMs were configured properly in the Data Protector cell by running the `omnidbsmis -ompasswd -check [`*`HostName`*`]` command.

This command performs a health check of you environment, which may help identify such potential problems as wrong user name or password provided, a broken network connection, a DNS resolution problem, and so on.

- After instant recovery, restored filesystems are mounted to the same mount points or drive letters as they were at the backup time. If these mount points or drive letters have other filesystems mounted, these filesystems are automatically dismounted before instant recovery, and the restored filesystems are mounted afterwards.

## Instant recovery using HP StorageWorks Disk Array XP Agent

This type of instant recovery can be performed if backup was performed with the XP VSS hardware provider in the resync mode.

With this type of instant recovery, Disk Array XP Agent (SSEA) synchronizes an S-VOL with its P-VOL and then splits the pair during the instant recovery session.

If the option **Wait for the replica to complete** is selected, SSEA waits for re-synchronization or copy process to complete (normal restore), before making the source volume available. If this option is not selected, the source volume is immediately available while the re-synchronization or copy process is running in the background (quick restore).

If you try to abort the session after the copy process has started, the abort operation is rejected and the session continues.

IMPORTANT:

The original source volumes are not retained after instant recovery. If the session fails after the copy process has already started, the application data residing on the original source volumes is lost.

### HP StorageWorks Disk Array XP considerations

- Instant recovery depends on the XP VSS hardware provider mode which was used during backup. If the VSS compliant mode was used, you can restore your data only using VDS. If the resync mode was used, you can restore your data only using SSEA.
- Changing the selected XP VSS hardware provider mode between different backup sessions using the same backup specification is not recommended. If the mode is changed while the replica rotation count is set to more than `1`, instant recovery will fail in the following situations:

- If you perform a backup in one mode and then the same backup in the other mode, restore of the backup performed in the VSS compliant mode (the "switch of disks" restore) will fail because the relationship between the S-VOL and its P-VOL will be detected from the other backup performed in the resync mode. This pair relationship should not be removed during restore and thus a switch of the source volume with the replica cannot be performed.

  To restore such a backup, perform one of the following prior to restore:

  - Using the `omnidbvss` command, manually remove S-VOLs that were created during backups in resync mode.
  - Set the variable `OB2VSS_FORCE_INSTANT_RECOVERY` in the `omnirc` file and enable the restore option **Retain source for forensics** in the GUI. When this option is selected, Data Protector preserves the pair relationship between the S-VOL and its P-VOL created in the backup with resync mode.

- If you perform restore using re-synchronization (the "copy of replica data" restore) and the production source volume (P-VOL) is not presented on the system, the restore session is aborted. This may happen after at least two backup sessions are run using the same backup specification in different modes and you run a "switch of disks" restore before a "copy of replica data" restore.

  Under such circumstances, perform one of the following prior to restore:

  - On the application system, manually present the P-VOL.
  - Set the variable `OB2VSS_FORCE_INSTANT_RECOVERY` in the `omnirc` file.

The following restore considerations apply to restore using Disk Array XP Agent only:

- You cannot start another instant recovery session using the same disk on the application system at the same time. A session can be started only after the preceding session completes synchronization.
- After instant recovery, the restored filesystems are mounted to the drive letters as were used at the time when backup was run. If these drive letters have other filesystems mounted, these filesystems are automatically dismounted before instant recovery, and the restored filesystems are mounted afterwards.

## Microsoft Exchange Server writer instant recovery specifics

This section provides specific information on the point-in-time and roll-forward recovery of Microsoft Exchange Server data.

With Microsoft Exchange Server 2003, you can restore a whole storage group or individual stores only to the original location. With Microsoft Exchange Server 2007, you can restore a whole storage group or individual stores to the original location or to a different location. For details, see "Microsoft Exchange Server 2007 writer restore specifics" on page 388.

Note that restore to the original location will fail if not all objects residing on the target volume are selected for restore. Fore example, if you have four stores and transaction logs on the target volume, and you select only one store for restore, restore session will fail. Thus, the following configuration scenarios are possible when restoring Microsoft Exchange Serve data:

- Transaction logs and database stores are on the same target volume.

  You cannot select only database stores for instant recovery in the GUI or CLI. If transaction logs and/or database stores are lost, whole storage group (or all objects residing on the target volume) need to be recovered.

  In this case, you can perform only *point-in-time* recovery. Transaction logs will be replaced with the backed up transaction logs.

- Transaction logs and database stores are on different target volumes.

  You can select only the database stores for instant recovery in the GUI or CLI. If a database store is lost, it can be recovered separately provided that it resides alone on the target volume to be recovered. Otherwise, all stores residing on the target volume must be selected for restore. If transaction logs are lost, whole storage group should be recovered.

  In this case, you can perform either *point-in-time* or *rollforward* recovery.

  To perform a *point-in-time* recovery of Microsoft Exchange Server writer data, select the whole storage group. Transaction logs will be replaced with the backed up transaction logs.

  To perform a *rollforward* recovery, select only the database stores and original location. The existing transaction logs will be applied to restored databases. However, rollforward recovery will not be possible if *point-in-time* recovery of the same backup session was performed before.

## Microsoft Exchange Server 2007 instant recovery prerequisites and limitations

### Prerequisites

For prerequisites which apply also to instant recovery, see "Prerequisites" on page 390.

### Limitations

- In CCR environments where original database resides on a different EVA disk array than the database copy, restore of the database copy to the location where original database resides will fail, since restore to a different disk array is not supported. In this case, perform manual failover of Exchange Server before the restore.

- Restore to a different location can be only performed using Microsoft Virtual Disk Service.

  Therefore, with XP VSS hardware provider, restore to a different location is possible only after the backup session was run in the VSS compliant mode.

- With XP VSS hardware provider, if a backup of an LCR or CCR storage group copy is run in the resync mode, restore of data to the original location (original database or Exchange Information Store) fails.

  The reason for such a behavior is that source volumes, on which the database copy or Exchange Replication Service resides, are in the suspended mirror relationship with the ZDB replica, which should be restored to the original database.

  Under such circumstances, the restore process first unpresents the source volumes from the Exchange Server database copy, and presents them to the Exchange Server original database afterwards. These actions result in shortage of source disks presented to the database copy. After restore, you need to set up your LCR or CCR environment again.

  To enable such restore process, set the `omnirc` variable `OB2VSS_FORCE_INSTANT_RECOVERY` to `1`, perform restore, and after it, manually set up the LCR or CCR environment.

- With Microsoft Exchange Server 2007 running in a CCR environment, whose database copy resides on the backup system, there are particular situations where you need to consider additional steps for performing instant recovery:

  - When XP VSS hardware provider is used, and the application system, where the production database resides, and the backup system, where the database copy resides, are connected to different disks arrays that are not configured in the same SAN. In such a configuration, the application system does not see LUNs on the backup system and vice versa.
  - When EVA VSS hardware provider is used and the disk arrays of the application system and the backup system are not controlled by the same Command View.
  - When the XP VSS hardware provider is used in the resync mode, regardless of the number of disk arrays are used.

  In the above situations, there is only one instant recovery scenario. For a successful instant recovery, follow the steps:

  1. Fail over Exchange Server. This will cause your production database to reside on the former backup system, where the database copy resided before failover.

2. Perform instant recovery.
3. After the instant recovery session completes, fail back Exchange Server.

   This action makes all backup LUNs available on the same backup system where the database copy resides.

NOTE:

It is recommended that both production and replication server systems are using disks of the same disk array.

For additional restore limitations, which apply also to instant recovery, see "Limitations" on page 390.

### Database recovery

You can run a database recovery from the **Instant Recovery** context of the Data Protector GUI. This option is available, if you have created a separate backup specification for **Incremental**/**Differential** backup with the same object and description as you have in the backup specification for instant recovery. Such an **Incremental/Differential** backup is based on the **Full** backup with the selected instant recovery option. You can select an **Incremental/Differential** backup in the **Instant Recovery** context and start restore. Instant recovery will be performed and transaction logs will be automatically applied to the recovered storage group.

## Instant recovery procedure

### Prerequisites

- For Microsoft Exchange Server 2003, manually dismount the databases to be recovered.

### Procedure

1. In the **HP Data Protector Manager**, switch to the **Instant Recovery** context.

2. In the Scoping Pane, expand **MS Volume Shadow Copy Writers** under **Restore Objects**, expand the backup specification from which you want to restore, and click the backup session from which you want to restore.

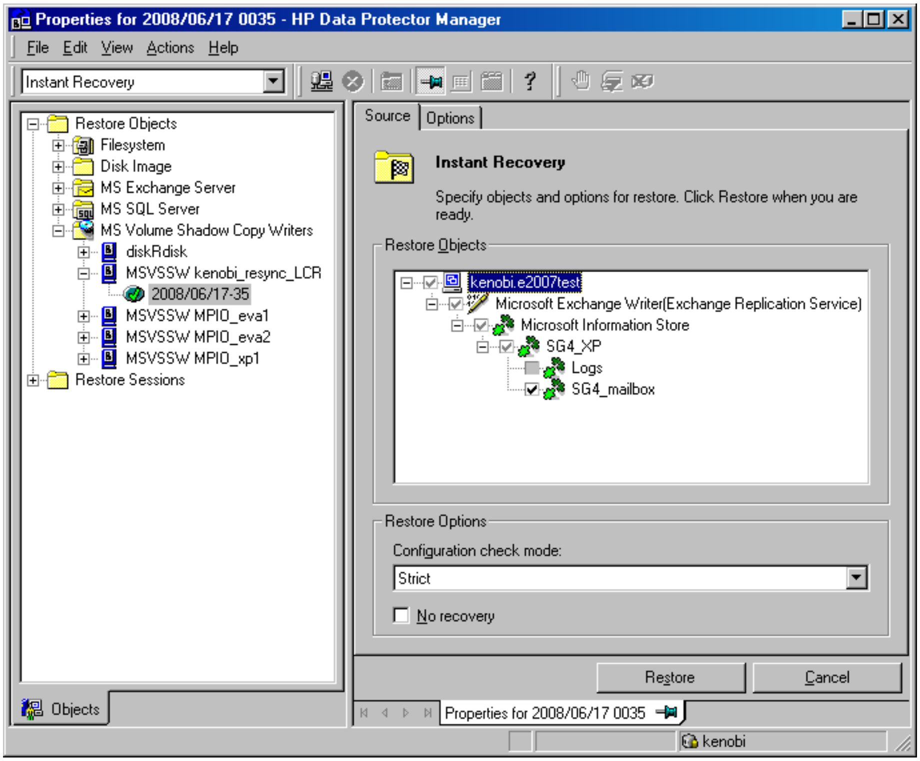

Properties for 2008/06/17 0035 - HP Data Protector Manager

**Instant Recovery**

Specify objects and options for restore. Click Restore when you are ready.

Restore Objects

- ☑ kenobi.e2007test
  - ☑ Microsoft Exchange Writer(Exchange Replication Service)
    - ☑ Microsoft Information Store
      - ☑ SG4_XP
        - □ Logs
        - ☑ SG4_mailbox

Restore Options

Configuration check mode: Strict

□ No recovery

Figure 116 Selecting writers components for instant recovery

3. In the Source property page, specify writers and/or components for recovery.

   Select the configuration check mode. For more information, see "Configuration check" on page 353.

   To perform additional steps after the instant recovery before recovering the database, select the **No recovery** option. Data Protector will not perform a database recovery and you can finish the recovery steps (such as applying transaction logs) later on manually.

   If the option is not selected, all recovery steps are finished. In such a case, additional tasks cannot be performed.

   This option is available only for the *SQLServer writer* and *Microsoft Exchange Server 2007 writer*. It is not supported for Microsoft Exchange writer, where the transaction logs are always applied.

   **Exchange Server Writer:** Optionally, to specify options for the consistency check of a Microsoft Exchange writer, right-click the writer and click **Additional options**.

   You can perform *point-in-time* or *rollforward* restore of Microsoft Exchange Server writer. For more information, see "Microsoft Exchange Server writer instant recovery specifics" on page 397.

   **Exchange Server 2007 Writer:**

   If you restore an LCR or CCR copy to the original location, note that restore will be performed to the original database (Exchange Information Store) and not to the database copy (Exchange Replication Service).

   a. Right-click a storage group, a store, or logs and click **Restore as**.

   b. In the MS Exchange Additional options dialog box, select the target location for the components you want to restore: target server, target storage group, and target stores. The following options are possible:

      - **Restore to a different store**

        This option is selected by default. It enables you to select the target stores for each store to be restored (original stores). Select the target server first and then from the Original and Target drop-down list, select the stores pairs. For example, to restore only one store, select only one target store for the original store you want to restore. Note that only stores cannot be restored, so logs are automatically selected in the restore session to different location.

      - **Restore to a non-Exchange location**

        Select this option to restore your data to a non-Exchange location. This means that the restored data will not be managed by Exchange Server and Recovery Storage Group (RSG) will not be created. You can create

RSG when data is restored. Select the target server first and then from the **Original** drop-down list, select the stores to be restored.

- **Restore to a non-Exchange location and create RSG**

  Select this option to restore your data to a non-Exchange location. After restore, Data Protector will create a Recovery Storage Group called **DP RSG** on the target server. The stores that you have selected and logs will be restored to this recovery group. Select the target server first and then from the **Original** drop-down list, select the stores to be restored.

Optionally, specify the mount point for the storage group to be restored in **Mount backup volume to the path**. Click **Browse** to specify the desired mount point. By default, the storage group is mounted to the `C:\Omni_Mnt` directory.

IMPORTANT:

The selected directory must be either empty or you can specify to create a new directory. If the directory is not empty, the instant recovery session fails.

Note that you can only specify the mount point for the storage group and not for a specific store.

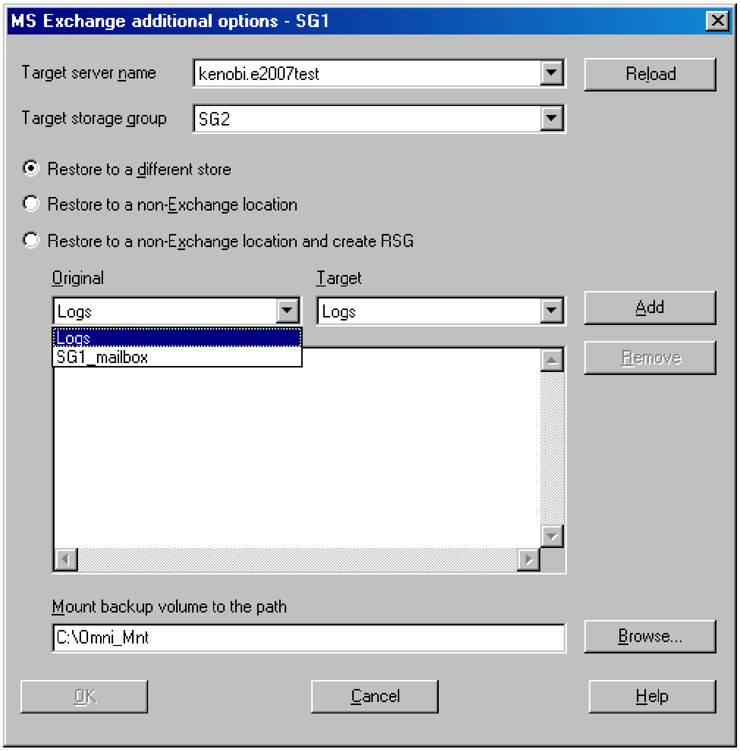

MS Exchange additional options - SG1

Target server name: kenobi.e2007test

Target storage group: SG2

● Restore to a different store

○ Restore to a non-Exchange location

○ Restore to a non-Exchange location and create RSG

Original: Logs

Target: Logs

Logs
SG1_mailbox

Mount backup volume to the path: C:\Omni_Mnt

Figure 117 Restore to different location options (Exchange Server 2007 Writer)

4. Optionally, under the **Options** tab, select VDS options or disk array options.

   For both instant recovery types, to retain the source volume, select **Retain source for forensics**.

   When **Restore using HP StorageWorks EVA** is selected, aditional options are available. In case of Microsoft Exchange Server 2007 writer restore to a different location, these options are not available.

   - The restore method: **Copy replica data to the source volume** or **Switch the replica**.

     If you select **Copy replica data to the source volume**, a full copy of the replica is created in the source storage (the copy is then switched with the source volume). A virtual copy is created and immediately available for use. However, in the background, the copy process runs for some time to copy all data from the replica to the source storage. To avoid the slowdown of the copy process caused by application use or the integrity check process (Exchange Server), use this option. The restore resumes after the time specified in **Wait up to *n* minutes** (default: 60) is reached (even if the copy process is not finished) or after the copy process has finished if it is finished before the specified time.

   - To delete the replica after the restore, select **Delete replica after successful restore.**

   For details, press **F1**.

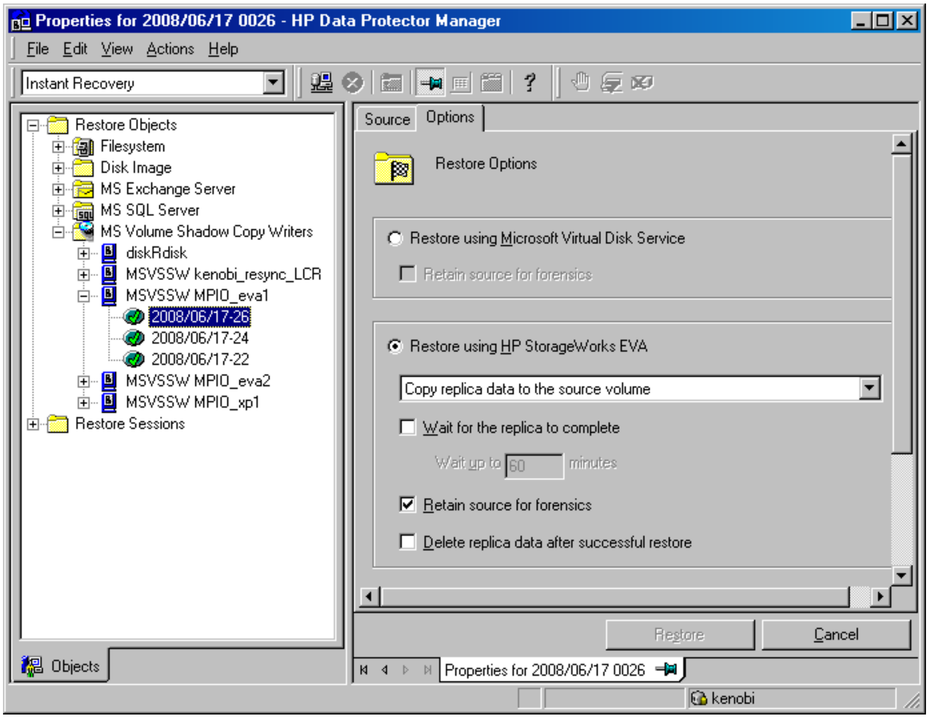
Properties for 2008/06/17 0026 - HP Data Protector Manager

Instant Recovery

- Restore Objects
  - Filesystem
  - Disk Image
  - MS Exchange Server
  - MS SQL Server
  - MS Volume Shadow Copy Writers
    - diskRdisk
    - MSVSSW kenobi_resync_LCR
    - MSVSSW MPIO_eva1
      - 2008/06/17-26
      - 2008/06/17-24
      - 2008/06/17-22
    - MSVSSW MPIO_eva2
    - MSVSSW MPIO_xp1
- Restore Sessions

Restore Options

○ Restore using Microsoft Virtual Disk Service

□ Retain source for forensics

● Restore using HP StorageWorks EVA

Copy replica data to the source volume

□ Wait for the replica to complete

Wait up to 60 minutes

☑ Retain source for forensics

□ Delete replica data after successful restore

Figure 118 Selecting instant recovery options (EVA integration)

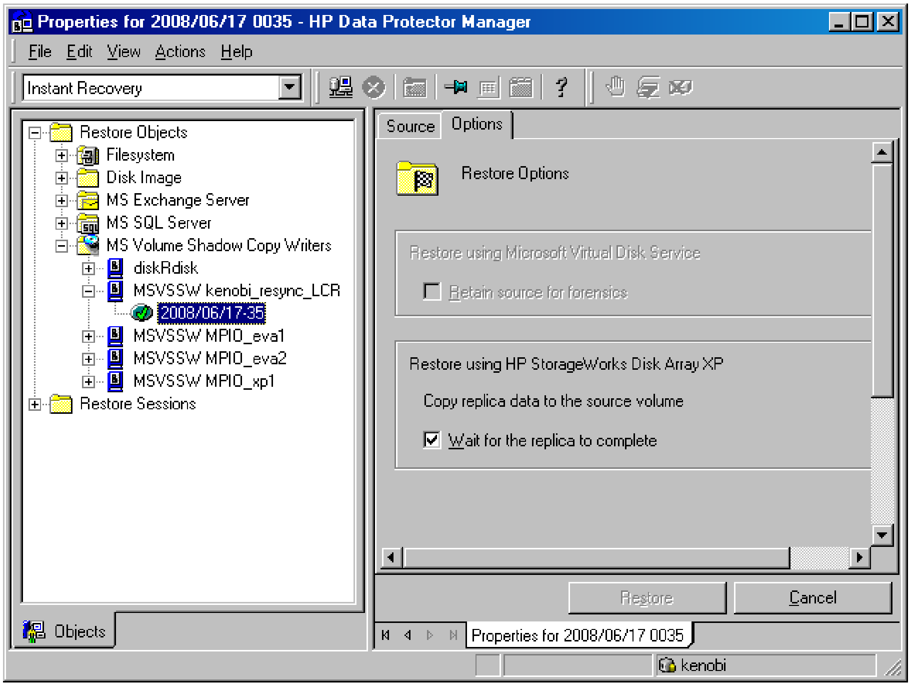

Restore using Microsoft Virtual Disk Service

☐ Retain source for forensics

Restore using HP StorageWorks Disk Array XP

Copy replica data to the source volume

☑ Wait for the replica to complete

Figure 119 Selecting instant recovery options (XP integration)

5. Click **Restore**.

6. If you restored Microsoft Exchange Server data, manually re-mount the database stores. In case of restore to Recovery Storage Group, re-mount the database stores in this RSG.

7. If you restored an LCR or CCR copy to the original database, perform additional steps:

   - In case of an LCR restore, it is recommended to *seed* the restored database to synchronize the original database with its copy.

     For more information, see the web page http://technet.microsoft.com/en-us/library/aa995973.aspx.

   - In case of a CCR restore, you may need to perform additional steps. For details, see "Additional steps after restore of a Microsoft Exchange Server CCR database copy to the original database" on page 408.

### Clusters

To perform an instant recovery in the cluster environment, use the above instant recovery procedure. Use virtual hostnames. For details on performing an instant recovery in cluster configurations, see the *HP Data Protector zero downtime backup administrator's guide*.

### Additional steps after restore of a Microsoft Exchange Server CCR database copy to the original database

In Microsoft Exchange Server CCR environments, after an instant recovery of a CCR database copy to the original database residing on the application system, additional steps are required.

△ CAUTION:

Do not move the clustered mailbox server to the other node (using the `Move-ClusteredMailboxServer` cmdlet) without first completing the procedure below. If you move the server at this point, data loss may occur.

Perform the following steps to enable normal operation of the original and copy databases:

1. On the application system, mount the restored store.
2. On the passive node, where the copy of the database exists, delete transaction logs.
3. On the passive node, use the `Update-StorageGroupCopy` cmdlet to seed the storage group copy or to re-synchronize the original storage group and its copy.
4. On the passive node, use the `Resume-StorageGroupCopy` cmdlet to resume the storage group copy.

# Troubleshooting

This section lists problems you might encounter when using the Data Protector Microsoft Volume Shadow Copy integration.

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

## Before you begin

- Ensure that the latest official Data Protector patches are installed. See the online Help index: "patches" on how to verify this.
- See the *HP Data Protector product announcements, software notes, and references* for general Data Protector limitations, as well as recognized issues and workarounds.
- See http://www.hp.com/support/manuals for an up-to-date list of supported versions, platforms, and other information.

## Checks and verifications

- On the application and backup systems, examine system errors reported in:

  *Data_Protector_home*\log\debug.log
- With the XP integration, ensure that RAID Manager Library is correctly installed on both the application and backup systems. In addition, check if the `libsvrrm.dll` file exists in the *RMLIB_home* directory.

## Common problems

### Problem

**Backup or instant recovery aborts due to VDS problems**

A backup or instant recovery session aborts with the following error:

```
Failed to load VDS service.
```

This error can appear as a result of abnormal termination of a VDS service.

### Actions

1. Stop the Virtual Disk Service (VDS):
   - Check that Virtual Disk Service is started, stop it using the command `net stop vds` or via the Control Panel.
   - If the above step does not help, stop the Virtual Disk Service by terminating the process `vds.exe` using the Task Manager. VDS will be started automatically as needed.
   - Alternatively, you could also log in to the system again and check if the system requests your confirmation to stop an abnormally terminated VDS. If it does, reconfigure the debugger to launch automatically to avoid similar problems in the future. This can be done by setting the `HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\AeDebug` value `Auto` to `1`.
2. Check your configuration.

   Successful Data Protector backup or restore completion is influenced by a number of factors imposed by non-Data Protector components, including Microsoft Exchange Server and VSS. One of the factors is timing, which in turn is impacted by environmental conditions such as I/O loads, storage device activity and concurrent Data Protector actions.

   To inhibit potential issues, it is recommended that you balance Data Protector activity across the available activity or maintenance window. If you encounter issues, *review and possibly reduce concurrent Data Protector activities* such as concurrent backups or restores.

## Backup problems

### Problem

**Microsoft Exchange Server 2003 writer backup failed**

When you start a backup of a Microsoft Exchange Server 2003 Writer, the following error is displayed:

```
[MAJOR]: Writer 'Microsoft Exchange Writer' failed to prepare files
for backup.
```

The reason may be the failure of the previous backup and, due to a Microsoft Exchange Server 2003 issue, the Exchange Server Writer could not perform a proper cleanup.

## Action

Restart the Information Store.

For more information, see the web page http://support.microsoft.com/kb/945424/en-us.

## Problem

### Microsoft Exchange Server 2003 aborts the backup if the shadow copy creation takes over 20 seconds

If an Exchange Server 2003 Writer is being backed up, the session can fail with VSSBAR reporting:

```
Snapshot could not be created.
```

In the application event log on the application system, the following event is recorded:

```
Event Type: Error
Event Source: ESE Event
Category: (16)
Event ID: 2004
Information Store (4916) Shadow copy 3 time-out (20000 ms).
```

## Action

The following can help to solve the problem:

- Limit the number of users that are accessing the management system.
- Reduce the number of volumes in a snapshot set. Create a backup specification dedicated to each storage group instead of one specification for the whole server. See “Microsoft Exchange Server writer backup specifics” on page 370.

## Problem

### Snapshots (volume shadow copies) creation fails with EVA provider

During a VSS backup, the following error is displayed:

```
[Critical] It was not possible to create Volume Shadow Copies.
```

If there are too many source volumes selected for backup (more than four) it may happen that the creation of snapshots takes too long and VSS fails the snapshots creation process.

### Action

Do not select too many objects in a backup specification because they might reside on more than four source volumes. This means that for example a Microsoft Exchange Server storage group should not reside on more than four source volumes and only one storage group should be specified for backup in this case. The same applies to a Microsoft SQL Server database.

### Problem

**No HP StorageWorks SMI-S EVA provider login entries are configured within SMISDB**

### Action

Add the login information for HP StorageWorks SMI-S EVA provider by running:

```
omnidbsmis -ompasswd -add HostName [-ssl] [-port PortNumber]
[-namespace Namespace] [-user Username] [-passwd Password]
```

### Problem

**Configuration of an HP StorageWorks SMI-S EVA CIMOM failed**

### Action

Run the following command to perform a health check of you environment, which may help identify any potential problems that occurred during the configuration:

```
omnidbsmis -ompasswd -check [HostName]
```

### Problem

**Volume shadow copies cannot be imported due to low space in registry and consequently backup fails**

If one of the following errors is reported during backup:

- By Data Protector:

  ```
  [Critical] It was not possible to import Volume Shadow Copies from
  application host.
  ```

- By Windows operating system:

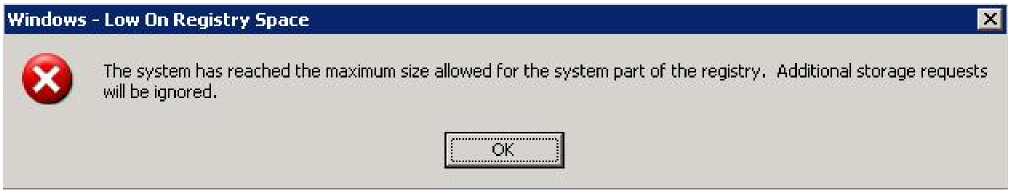

Windows - Low On Registry Space

The system has reached the maximum size allowed for the system part of the registry. Additional storage requests will be ignored.

**Figure 120 Error reported by registry**

it means that the Windows registry is out of space.

If there are too many entries in the Windows system registry, the registry runs out of space and consequently backup fails. Although Data Protector does not write anything to registry directly, the registry contains entries about all volumes and disks ever presented on the system recorded by the volume storage driver.

## Action

The registry should be cleaned up. For information on how to clean the registry, see the web page http://support.microsoft.com/kb/277222/en-us.

In case of problems with the installation of the Scrubber utility provided by Microsoft, contact Microsoft Product Support Services.

It is recommended to perform registry management tasks periodically to prevent the registry from filling.

## Problem

**Data Protector reports that a volume was not removed**

During a backup session, the following error is reported:

```
Volume 'StorageID', which is part of backup 'backupID', was not removed.
```

This error may appear if a replica on the disk array is removed, but its entry in the VSS database (VSSDB) remains intact.

This may happen when a backup session fails without creating a replica, but the reference that was created for it could not be removed due to network problems.

## Action

To remove the false entry from the VSSDB:

1. In the error messages, find the ID of the session which could not remove the target volume.

2. Run the following command, where *SessionID* is the session whose target volume(s) could not be removed:

   ```
   omnidbvss -remove session SessionID -reference
   ```

   Upon successful removal, the command should display the confirmation message `Removing references of session SessionID from VSSDB.`

## Problem

### Backup of a Microsoft Exchange Server 2007 CCR database copy fails

During a backup session of a database copy in a Microsoft Exchange Server CCR environment, Data Protector reports a major error notifying that the backup session has failed.

This problem may occur in CCR environments, due to a Microsoft Exchange Server 2007 issue with incorrect display of database copy states in Exchange Management Console. In such a case, a database copy in a "Failed" state may be displayed as "Healthy".

## Action

To make the Exchange Management Console show the real status of a database copy, perform either of the following actions:

- Update Microsoft Exchange Server 2007 with Service Pack 1.
- Perform re-seeding procedure as follows:
  1. On the passive node, suspend the replication by using `Suspend-StorageGroupCopy` cmdlet.
  2. Delete all log files from the `Logs` directory of the database copy.
  3. Seed the database copy or re-synchronize the original database and its copy by using the `Update-StorageGroupCopy` cmdlet.
  4. Resume the database copy by using the `Resume-StorageGroupCopy` cmdlet.
  5. On the passive node, check the status of the Exchange Replication Service by using `vssadmin list writers` command. If the status is not `stable`, restart Microsoft Exchange Replication Service.

## Problem

### Exchange Replication Service writer instance in LCR environment is not displayed in the Data Protector GUI

In the Data Protector GUI, while creating a backup specification, the `Microsoft Exchange Writer(Exchange Replication Service)` object is not displayed on the pane for selection of backup objects.

This problem may occur in LCR environments, due to a Microsoft Exchange Server 2007 issue with incorrect display of database copy states in Exchange Management Console. In such a case, a database copy in a "Failed" state may be displayed as "Healthy".

### Action

To make the Exchange Management Console show the real status of a database copy and to enable selection of the backup object, restart Microsoft Exchange Replication Service.

## Restore problems

### Problem

**After the restore of system writers was aborted, the Windows operating system is corrupted when you restart it**

If the restore of some system writers (for example, System Writer) is aborted for any reason (hardware or software failure, manually aborted, etc.), the Windows operating system may be corrupted after the restart (for example, the GUI or some system services cannot be started, etc.).

### Action

Depending on the nature of the corruption, repair or re-install the operating system from the Windows installation CD-ROM.

### Problem

**Instant recovery of the Microsoft Exchange Server Writer fails**

This problem may occur if the Microsoft Exchange Writer is not in the stable state. Check this by running `VSSadmin list writers` from the command prompt.

### Action

Bring the Exchange Server Writer to a stable state by restarting the Microsoft Exchange Information Store.

### Problem

**After restore, the system restart error is displayed**

In some cases, if the HBAs are changed on the application system, the software driver from previous HBAs can cause this error after restore.

### Action

Remove older drivers of uninstalled HBAs from the application system.

### Problem

**After an SQLServer writer instant recovery is restarted, the database cannot be brought online**

If an instant recovery of the SQLServer writer is aborted or fails, you can restart the session. However, the SQLServer writer may report an error, stating that the files cannot be prepared during the instant recovery and the session completes with errors.

The issue can appear if the disk was dismounted before the session was aborted and the files are not visible to the writer. As a result, an error is reported and the database can not be brought online after restore.

### Action

To recover the database after such session:

1. Detach from the database.
2. Attach to the database.

### Problem

**Creation of an RSG fails when an RSG for the same storage group already exists**

When you start a standard restore or instant recovery of the Microsoft Exchange Server 2007 writer with the option `Restore to a non-Exchange location and create RSG`, the following error is displayed:

```
[Major]
Application specific function 'PostRestoreEndExt' failed with error:
'The mailbox database that you specified is already associated with
a recovery mailbox database.
```

This error appears if an RSG for a storage group or mailbox database you restore is already created on some other system.

### Action

Since only one RSG can exist for the same storage group and since Data Protector can only delete an RSG that exists on the target restore location, you need to manually delete the RSG on the other system and restart the restore.

# User scenario for Microsoft Exchange Server 2003 backup and restore

This section provides examples of backup and restore policies for Microsoft Exchange Server 2003. Three examples are provided, one for ZDB (transportable VSS snapshots) and two for different types of restore.

## Example - VSS transportable backup

This example describes a backup scenario for Microsoft Exchange Server 2003 using VSS transportable backup together with HP StorageWorks EVA. The data is to be backed up on tape once a day for the storage group containing critical mailboxes and every second day for all other storage groups.

The storage groups should be backed up separately.

The following example setup is possible:

- A Microsoft Exchange Server 2003 is running on the application system and is connected to a HP StorageWorks Enterprise Virtual Array. A tape drive is connected to the application system for restore purposes.
- A separate backup system is connected to a tape library.
- Two storage groups are configured on the Microsoft Exchange Server 2003, each containing two stores.

  The first storage group, named Lists_Group, contains company-wide mailing lists. The storage group contains two stores, Staff_Store for general mailing lists and Team_Store for team mailing lists.

  The second storage group, named Mailbox_Group, contains a larger number of individual mailboxes for employees, which are given higher priority. This way the mailboxes are still available even if mailing lists are down. The mailboxes are set up in two separate stores, one store for each department, and named Sales_Store and Support_Store.
- Every storage group has a separate backup specification to reduce the time needed for the shadow copy creation. Refer to “Backup problems” on page 410.

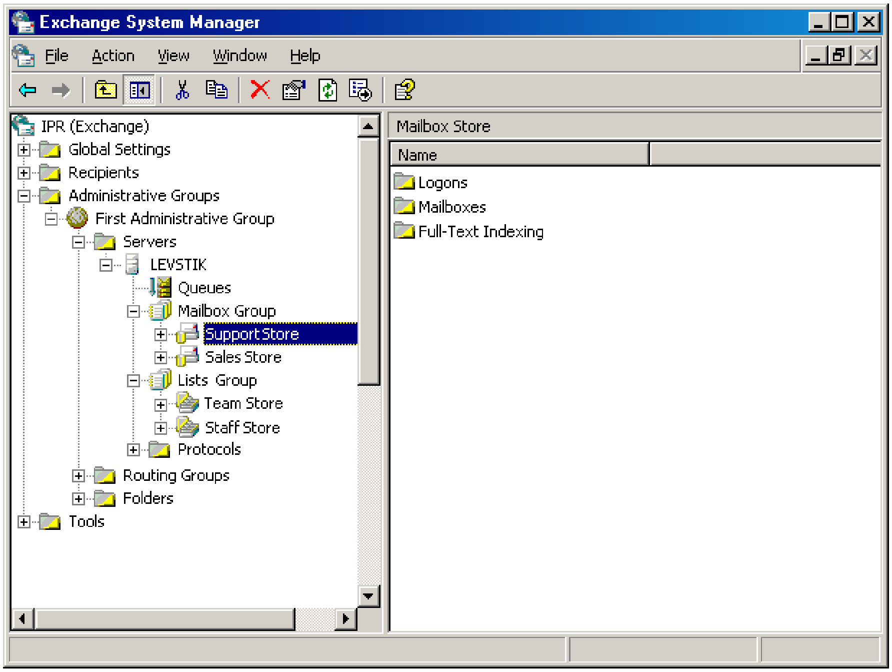

Figure 121 Microsoft Exchange Server 2003 storage groups

For the first storage group (Lists_Group), create the following backup specification:

1. Select **VSS transportable backup** as the backup type. Select the application and backup system and specify **Snapshot (Differential)** as the replica type.
2. Expand the Microsoft Exchange Server 2003 writer and select the first storage group (Lists_Group) for backup.
3. Using the Data Protector scheduler, schedule the backup specification to start a VSS transportable backup at night. In the Recurring box, specify **Full** backup **Weekly** on Wednesdays and Fridays and **Incremental** backup **Weekly** on other days. Set the backup protection to 2 days.

For the second storage group (Mailbox_Group), create the backup specification that can be used for instant recovery:

1. Select **VSS transportable backup** as the backup type. Select the application and backup system and set the **Track the replica for instant recovery** option.
2. Expand the Microsoft Exchange Server 2003 writer and select the second storage group (Mailbox_Group) for backup.

3. Using the Data Protector scheduler you need to schedule the backup specification to start a VSS transportable backup everyday. In the Recurring box, select **Daily**, set **Recurring** options to 1 day and set the backup protection to 4 days.

## Example restore scenario for Microsoft Exchange Server 2003

In this example the Microsoft Exchange Server 2003 is configured as in "Example - VSS transportable backup" on page 417 and the backup policy given in the example was implemented.

### Example 1. Example - restoring a single store

The second store (Support_Store) in the second storage group (Mailbox_Group), which contains user mailboxes, is damaged, but the transaction logs and other stores are not damaged. Therefore, a rollforward recovery will be performed only for this store.

To perform a rollforward recovery, proceed as follows:

1. Start the Microsoft Exchange Manager and unmount all stores in Mailbox_Group, but do not unmount the stores from the first storage group.
2. Start Data Protector, go to the **Instant Recovery** context and expand the application system. Select the last backup session of the backup specification with enabled instant recovery options and expand Microsoft Exchange Server 2003 Writer. Select Support_Store under Mailbox_Group for instant recovery.

   Click **Restore**.
3. After the session is restored, start the Microsoft Exchange Manager. Mount all stores in the second storage group.

### Example 2. Example - restoring a complete storage group after loss of transaction logs

In this example, the first store (Staff_Store) and transaction logs in the first storage group (Lists_Group) are damaged but the second store (Mailbox_Group) is intact. Since the transaction logs are damaged, a point-in-time restore of the whole storage group will be performed.

To restore the complete Lists_Group storage group, perform the following steps:

1. Start Exchange System Manager and check if the first storage group (Lists_Group) is already unmounted. If not, unmount the whole group.
2. Start Data Protector. In the Restore context expand the application system. Select the last backup session and then Microsoft Exchange Server 2003 Writer. Select the first storage group (Lists_Group).

   Start the restore.
3. Start Microsoft Exchange Manager and mount all stores in Lists_Group.

# A Appendix

## In this appendix

This appendix gives information on the following topics:

- "Reconfiguring an Oracle instance for instant recovery" on page 421
- "ZDB integrations omnirc variables" on page 425

## Reconfiguring an Oracle instance for instant recovery

If the control files or redo logs are located on the same volume group (if LVM is used) or source volume as the database files, the control files and online redo logs are overwritten during instant recovery. In such case, you may want to reconfigure the Oracle instance. Refer to "Oracle backup set ZDB concepts" on page 39 and to "Oracle proxy-copy ZDB concepts" on page 44 for details on the required configuration. For additional examples on how to move the redo logs and control files, refer to "Examples for moving the control files and redo logs to different locations " on page 423.

### Moving Online Redo Logs

To move the *online redo log files* from the source volumes to be replicated, to other locations:

1. List the online redo log files:

   ```
   $ sqlplus
   ```

   ```
   SQL> select member from v$logfile;
   ```

2. Shut down the database:

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

   ```
   SQL> shutdown
   ```

   ```
   SQL> exit
   ```

3. Move the log files to a different location using operating system tools.
4. Start the database in mount mode:

   ```
   $ sqlplus
   SQL> connect user/password@service as sysdba;
   SQL> startup mount;
   ```

5. Register the new locations for each moved file:

   ```
   SQL> alter database rename file 'OldPathName' to
   'NewPathName';
   ```

   where `OldPathName` and `NewPathName` are full paths to the log file.

6. Open the database in normal mode:

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

## Moving Control Files

To move the *control files* from the source volumes to be replicated, to other locations:

1. Determine if the database uses the SPFILE parameter:

   ```
   SQL> show parameter SPFILE
   ```

2. If the database does not use SPFILE:

   a. Shut down the database.

   ```
   SQL> shutdown
   ```

   b. Move the control files to a different location using operating system tools.

   c. Edit the `CONTROL_FILES` parameter in the database's initialization parameter file (usually located in the `$ORACLE_HOME/dbs/initSID.ora` directory) to change the existing control file names:

   ```
   control_files = ("NewPathName", ...)
   ```

   d. Restart the database:

   ```
   SQL> startup
   ```

   If the database uses SPFILE:

   a. Specify the new location for control files by running the following command:

   ```
   SQL> alter system set control_files='NewPathName1',
   'NewPathName2',..., scope=spfile
   ```

   b. Shut down the database.

```
SQL> shutdown
```

c. Move the control files to a different location.

d. Restart the database:

```
SQL> startup
```

# Examples for moving the control files and redo logs to different locations

## Example - Moving Online Redo Logs

In the following example for Oracle9i on HP-UX, the data files are on the same source volume as the control files and redo logs, which is `/opt/oracle/product/9.2.0`.

To move the *online redo log files* from `/opt/oracle/product/9.2.0` to `/oracle/logs` (which is not replicated):

1. List the online redo log files:

   ```
   $ sqlplus
   ```

   ```
   SQL> select member from v$logfile;
   ```

   ```
   /opt/oracle/product/9.2.0/oradata/redo01.log
   /opt/oracle/product/9.2.0/oradata/redo02.log
   /opt/oracle/product/9.2.0/oradata/redo03.log
   ```

   List the filenames and tablespaces to check whether they are on the same source volumes as the control files:

   ```
   SQL> select FILE_NAME,TABLESPACE_NAME,BYTES from
   dba_data_files;
   ```

   ```
   FILE_NAME
   ----------------------------------------------------------
   TABLESPACE_NAME                     BYTES
   ------------------------------ ----------
   /opt/oracle/product/9.2.0/oradata/system01.dbf
   SYSTEM                         419430400
   /opt/oracle/product/9.2.0/oradata/undotbs01.dbf
   UNDOTBS1                       377487360
   /opt/oracle/product/9.2.0/oradata/cwmlite01.dbf
   CWMLITE                         20971520
   ```

2. Shut down the database:

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

SQL> shutdown

SQL> exit
```

3. Move the log files to a different location.

```
$ mv /opt/oracle/product/9.2.0/oradata/redo* /oracle/logs
```

4. Start the database in mount mode:

```
$ sqlplus

SQL> connect user/ password@service as sysdba;

SQL> startup mount;
```

5. Rename the new locations for each moved file:

```
alter database rename file
'/opt/oracle/product/9.2.0/oradata/redo01.log' to
'/oracle/logs/redo01.log';

Database altered.

alter database rename file
'/opt/oracle/product/9.2.0/oradata/redo02.log' to
'/oracle/logs/redo01.log';

Database altered.

alter database rename file
'/opt/oracle/product/9.2.0/oradata/redo03.log' to
'/oracle/logs/redo01.log';

Database altered.
```

6. Open the database in normal mode:

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

## Example - Moving Control Files for Oracle9i

In the following example, the Oracle9i database uses SPFILE. To move the control files from `/opt/oracle/product/9.2.0/` to `/oracle/oractl`:

1. Determine if the database uses the SPFILE parameter:

   ```
   SQL> show parameter spfile;

   NAME                            TYPE            VALUE
   ------------------------- ----------- --------------------
   spfile                          string       ?/dbs/spfile@.ora
   ```

2. Specify the new location for the control files by running the following command (in a single line and without the "\" characters):

   ```
   SQL> alter system setcontrol_files='/oracle/logs/RCVCAT\
   /control01.ctl','/oracle/logs/RCVCAT/control02','/oracle\
   /logs/RCVCAT/control03.ctl' scope=spfile;
   ```

3. Shut down the database:

   ```
   SQL> shutdown
   ```

4. Move the control files to the new location:

   ```
   mv /opt/oracle/product/9.2.0/oradata/control*
   /oracle/oractl
   ```

5. Restart the database:

   ```
   SQL> startup
   ```

# ZDB integrations omnirc variables

The Data Protector ZDB integrations use environment variables, which can be set in the `/opt/omni/.omnirc` (on UNIX systems) or *`Data_Protector_home`*`\omnirc` file (on Windows systems), on both the application and backup systems. These variables are used for Data Protector ZDB integrations customizing. See the online Help index: “omnirc options” for information on how to use the `omnirc` file.

For information on Data Protector ZDB agents `omnirc` file variables, see the *HP Data Protector zero downtime backup administrator's guide*.

This section explains the `omnirc` file variables that can be set for the Data Protector ZDB integrations.

**ZDB_ORA_INCLUDE_CF_OLF**: An Oracle integration or SAP R/3 integration related variable.

NOTE:

This variable is not supported on EMC.

The default value is 0. Possible values are 0 and 1.

If an offline backup is performed using the Data Protector SAP R/3 integration, the variable is ignored and the integration behaves as if the variable was set to 1.

## Instant Recovery

The instant recovery process depends on whether the control file and redo logs reside on the same disk array source volume as datafiles or not:

- If this variable is set to 0 (default), during a ZDB session, Data Protector creates target volumes only for the source volumes containing Oracle datafiles. Target volumes for source volumes containing Oracle control file and Oracle online redo logs are not created.

  For Oracle proxy-copy or backup set ZDB and restore concepts when this variable is set to 0, see "Oracle backup set ZDB concepts" on page 39 and "Oracle proxy-copy ZDB concepts" on page 44. For SAP R/3 backup and restore concept when this variable is set to 0, see "Integration concepts" on page 164.
- If this variable is set to 1, Data Protector creates target volumes for all source volumes containing Oracle datafiles, Oracle control file, and if Oracle integration is used, Oracle online redo logs.

IMPORTANT:

If the ZDB_ORA_INCLUDE_CF_OLF variable is set to 1 the control files and redo logs are overwritten during instant recovery.

## Opening the Database on the Backup System

To successfully open the database on the backup system for *other* purposes than Data Protector, note:

- With Oracle proxy-copy ZDB method, set this variable to 1.
- With Oracle backup set ZDB method, used with the Oracle integration, you can always open the database on the backup system.

### Prerequisites

The prerequisites for this variable to be set to `1` are:

- Data Protector Oracle Server integration: Oracle datafiles, Oracle control file, and Oracle online redo logs must be installed on a disk array.
- Data Protector SAP R/3 integration: Oracle datafiles and Oracle control file must be installed on a disk array.

See Figure 122 on page 428 and Figure 123 on page 429 for Oracle backup and restore concepts when this variable is set to `1`. See Figure 124 on page 430 for SAP R/3 backup and restore concepts when this variable is set to `1`.

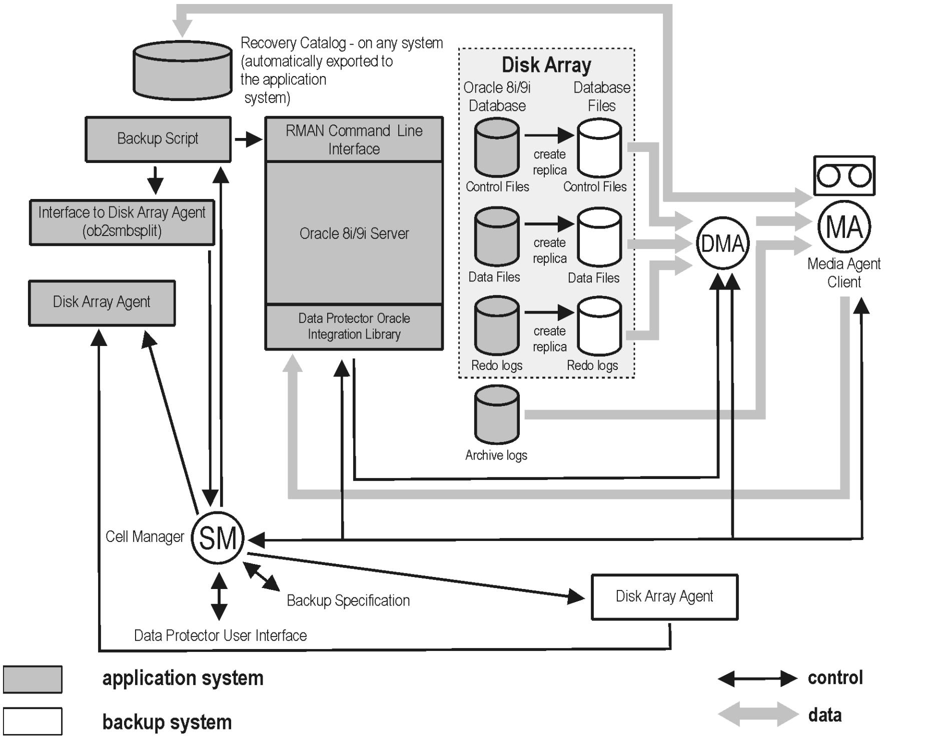

**Figure 122 Oracle proxy-copy ZDB and restore concepts when the ZDB_ORA_INCLUDE_CF_OLF variable is set to 1**

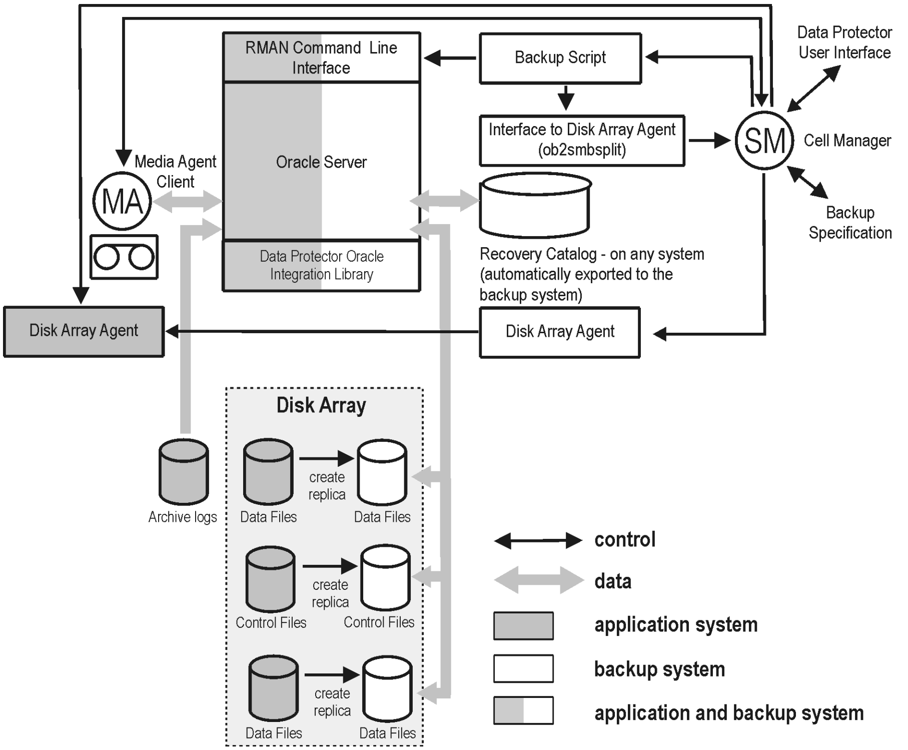

**Figure 123 Oracle backup set ZDB and restore concept when the ZDB_ORA_INCLUDE_CF_OLF variable is set to 1**

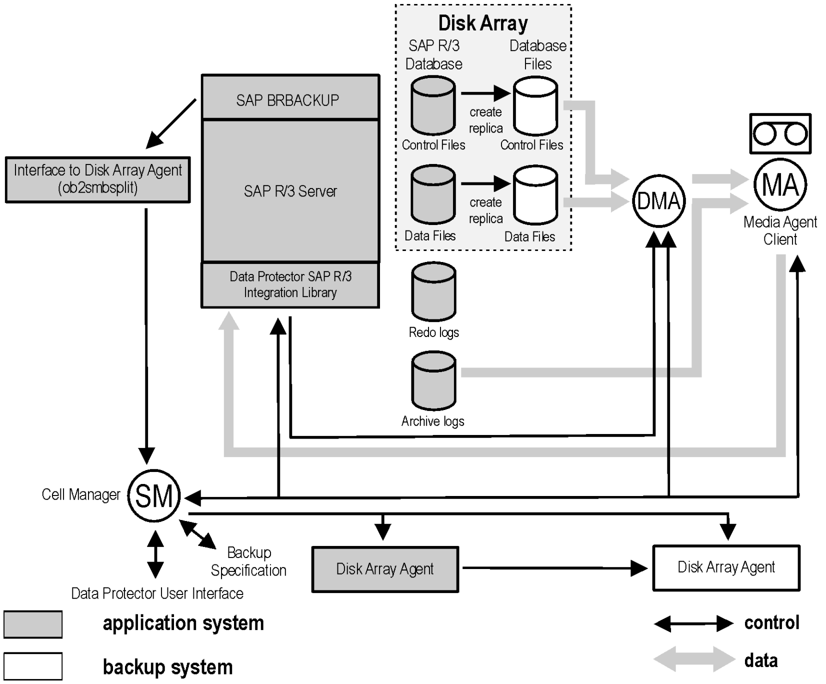

## Figure 124 SAP R/3 backup and restore concept when the ZDB_ORA_INCLUDE_CF_OLF variable is set to 1 with online backup, or in case of offline backup

**ZDB_ORA_INCLUDE_SPF**: Data Protector Oracle Server integration related variable.

The default value is `0`. Possible values are `0` and `1`.

The variable is ignored and the integration behaves as if the variable was set to `1` if offline backup is performed using the Data Protector SAP R/3 integration.

By default (if this variable is set to `0`), during a ZDB session, Data Protector checks if Oracle datafiles and the SPFILE are on the same source volumes. If the SPFILE and datafiles are on the same volumes and instant recovery is enabled, the ZDB session is aborted.

If the `ZDB_ORA_INCLUDE_SPF` variable set to `1`, Data Protector skips the check.

IMPORTANT:

If the `ZDB_ORA_INCLUDE_SPF` variable is set to `1` and the SPFILE is on the same source volumes as datafiles, it is overwritten during instant recovery.

**ZDB_ORA_NO_CHECKCONF_IR:** Data Protector Oracle Server integration related variable.

The default value is `0`. Possible values are `0` and `1`.

By default, the Oracle configuration is checked whether or not the control file, SPFILE, and online redo logs are on different volume groups than datafiles. To check the configuration, the CLI binary `omniresolve` is used internally. This binary needs to have the setuid bit set on UNIX. Setting this variable to `1`, the check will be skipped, and the `omniresolve` binary will not be used to check the Oracle configuration.

Note that checking Oracle for instant recovery suitability is important to be sure that the instant recovery will not overwrite the control file, online redo logs, and the SPFILE.

**OB2MARAWREAD_KB:** This variable sets the read block size for Oracle and SAP R/3 ZDB integrations on UNIX systems with Oracle tablespaces or datafiles installed on disk images and when using the proxy-copy method (when using DMA).

The default value is 64KB. The specified value must be in the range between 1KB an 1MB.

The specified size is automatically adjusted to a size which is a multiple of the block size. The values above 256KB could cause the DMA to fail.

**ZDB_TAKE_CLUSRES_ONLINE:** This variable specifies how many times Data Protector tries to connect to Microsoft SQL Server in case the first connection fails. A reconnection is triggered every 30 seconds. This means that Data Protector waits up to *`ZDB_TAKE_CLUSRES_ONLINE`* `x 30` seconds for the Microsoft SQL Server resources to start up.

# Glossary

**access rights**
*See* user rights.

**ACSLS**
*(StorageTek specific term)* The Automated Cartridge System Library Server (ACSLS) software that manages the Automated Cartridge System (ACS).

**Active Directory**
*(Windows specific term)* The directory service in a Windows network. It contains information about resources on the network and makes them accessible to users and applications. The directory services provide a consistent way to name, describe, locate, access and manage resources regardless of the physical system they reside on.

**AES 256–bit encryption**
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**
*(EMASS/GRAU specific term)* Automated Mixed-Media library.

**application agent**
A component needed on a client to back up or restore online database integrations.
*See also* Disk Agent.

**application system**
*(ZDB specific term)* A system the application or database runs on. The application or database data is located on source volumes.
*See also* backup system and source volume.

**archived redo log**
*(Oracle specific term)* Also called offline redo log. If the Oracle database operates in the ARCHIVELOG mode, as each online redo log is filled, it is copied to an archived log destination. This copy is the archived redo log. The presence or absence of

an archived redo log is determined by the mode the database is using:

- ARCHIVELOG - The filled online redo log files are archived before they are reused. The database can be recovered if an instance or a disk fails. A “hot” backup can be performed only when the database is running in this mode.
- NOARCHIVELOG - The filled online redo log files are not archived.

*See also* online redo log.

**archive logging**

*(Lotus Domino Server specific term)* Lotus Domino Server database mode where transaction log files are overwritten only after they have been backed up.

**ASR Set**

A collection of files stored on several diskettes, required for proper reconfiguration of the replacement disk (disk partitioning and logical volume configuration) and automatic recovery of the original system configuration and user data that was backed up during the full client backup. These files are stored as an ASR archive file on the Cell Manager in the directory `Data_Protector_program_data\Config\Server\dr\asr` (Windows Server 2008), `Data_Protector_home\Config\Server\dr\asr` (other Windows systesm), 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**

*(Oracle specific term)* Automatic Storage Management is an Oracle 10g/11g integrated filesystem and volume manager that manages Oracle database files. It eliminates complexity

associated with managing data and disk and provides striping and mirroring capabilities to optimize performance.

**automigration** *(VLS specific term)* The functionality that allows data backups to be first made to the VLS' virtual tapes and then migrated to physical tapes (one virtual tape emulating one physical tape) without using an intermediate backup application.
*See also* Virtual Library System (VLS) and virtual tape.

**BACKINT** *(SAP R/3 specific term)* SAP R/3 backup programs can call the Data Protector backint interface program via an open interface, which enables them to communicate with Data Protector software. For backup and restore, SAP R/3 programs issue orders for the Data Protector backint interface.

**backup API** The Oracle interface between the Oracle backup/restore utility and the backup/restore media management layer. The interface defines a set of routines to allow the reading and writing of data to the backup media, the creation, searching and removing the backup files.

**backup chain** *See* restore chain.

**backup device** A device configured for use with Data Protector, which can write data to and read data from storage media. This can be, for example, a standalone DDS/DAT drive or a library.

**backup generation** One backup generation includes one full backup and all incremental backups until the next full backup.

**backup ID** An identifier of an integration object that equals the session ID of the backup of this object. The backup ID is preserved when an object is copied, exported, or imported.

**backup object** A backup unit that contains all items backed up from one disk volume (logical disk or mount point). The backed up items can be any number of files, directories, or the entire disk or mount point. Additionally, a backup object can be a database/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 and mount point on UNIX). For integration objects — backup stream identification, indicating the backed up database/application items.

- Description: For filesystem objects — uniquely defines objects with identical client name and mount point. For integration objects — displays the integration type (for example, SAP or Lotus).
- Type: Backup object type. For filesystem objects — filesystem type (for example, WinFS). For integration objects — "Bar".

**backup owner**

Each backup object in the IDB has an owner. The default owner of a backup is the user who starts the backup session.

**backup session**

A process that creates a copy of data on storage media. The activities are specified in a backup specification or an interactive session. All clients configured in one backup specification are backed up together in one backup session using the same backup type (full or incremental). The result of a backup session is a set of media, which was written to, also called the backup or media set.
*See also* backup specification, incremental backup, and full backup.

**backup set**

A complete set of integration objects associated with a backup.

**backup set**

*(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.

**backup system**

*(ZDB specific term)* A system connected to target volumes of one or multiple application systems. The backup system is typically connected to a backup device to perform the backup of the data in a replica.

*See also* application system, target volume, and replica.

**backup types** *See* incremental backup, differential backup, transaction backup, full backup, and delta backup.

**backup view** Data Protector provides different views for backup specifications:
By Type - according to the type of data available for backups/templates. Default view.
By Group - according to the group to which backup specifications/templates belong.
By Name - according to the name of backup specifications/templates.
By Manager - if you are running MoM, you can also set the Backup view according to the Cell Manager to which backup specifications/templates belong.

**BC** *(EMC Symmetrix specific term)* Business Continuance are processes that allow customers to access and manage instant copies of EMC Symmetrix standard devices.
*See also* BCV.

**BC** *(HP StorageWorks Disk Array XP specific term)* The Business Copy XP allows to maintain internal copies of HP StorageWorks Disk Array XP LDEVs for purposes such as data backup or data duplication. The copies (secondary volumes or S-VOLs) can be separated from the primary volumes (P-VOLs) and connected to a different system for various purposes, such as backup and development. For backup purposes, P-VOLs should be connected to the application system, and one of the S-VOL mirror sets should be connected to the backup system.
*See also* HP StorageWorks Disk Array XP LDEV, CA, Main Control Unit, application system, and backup system.

**BC EVA** *(HP StorageWorks EVA specific term)* Business Copy EVA is a local replication software solution enabling you to create point-in-time copies (replicas) of the source volumes using the snapshot and clone capabilities of the EVA firmware.
*See also* replica, source volume, snapshot, and CA+BC EVA.

**BC Process** *(EMC Symmetrix specific term)* A protected storage environment solution that has defined specially configured EMC Symmetrix devices as mirrors or Business Continuance Volumes to protect data on EMC Symmetrix standard devices.
*See also* BCV.

**BC VA**

*(HP StorageWorks Virtual Array specific term)* Business Copy VA allows you to maintain internal copies of HP StorageWorks Virtual Array LUNs for data backup or data duplication within the same virtual array. The copies (child or Business Copy LUNs) can be used for various purposes, such as backup, data analysis or development. When used for backup purposes, the original (parent) LUNs are connected to the application system and the Business Copy (child) LUNs are connected to the backup system.
*See also* HP StorageWorks Virtual Array LUN, application system, and backup system.

**BCV**

*(EMC Symmetrix specific term)* Business Continuance Volumes, or BCV devices, are dedicated SLDs that are pre-configured in the ICDA on which the business continuation operation runs. BCV devices are assigned separate SCSI addresses, differing from the addresses used by the SLDs they mirror. The BCV devices are used as splittable mirrors of the primary EMC Symmetrix SLDs that need to be protected.
*See also* BC and BC Process.

**Boolean operators**

The Boolean operators for the full text search functionality of the online Help system are AND, OR, NOT, and NEAR. Used when searching, they enable you to define your query precisely by creating a relationship between search terms. If no operator is specified in a multi-word search, AND is used by default. For example, the query manual disaster recovery is equivalent to manual AND disaster AND recovery.

**boot volume/disk/ partition**

A volume/disk/partition with files required for the initial step of the boot process. Microsoft terminology defines the boot volume/disk/partition as a volume/disk/partition containing the operating system files.

**BRARCHIVE**

*(SAP R/3 specific term)* An SAP R/3 backup tool that allows you to archive redo log files. BRARCHIVE also saves all the logs and profiles of the archiving process.
*See also* BRBACKUP, and BRRESTORE.

**BRBACKUP**

*(SAP R/3 specific term)* An SAP R/3 backup tool that allows an online or offline backup of the control file, of individual data files, or of all tablespaces and, if necessary, of the online redo log files.
*See also* BRARCHIVE, and BRRESTORE.

**BRRESTORE**

*(SAP R/3 specific term)* An SAP R/3 tool that can be used to restore files of the following type:

- Database data files, control files, and online redo log files saved with BRBACKUP
- Redo log files archived with BRARCHIVE
- Non-database files saved with BRBACKUP

You can specify files, tablespaces, complete backups, log sequence numbers of redo log files, or the session ID of the backup.
*See also* BRBACKUP, and BRARCHIVE.

**BSM**

The Data Protector Backup Session Manager controls the backup session. This process always runs on the Cell Manager system.

**CA**

*(HP StorageWorks Disk Array XP specific term)* Continuous Access XP allows you to create and maintain remote copies of HP StorageWorks Disk Array XP LDEVs for purposes such as data duplication, backup, and disaster recovery. CA operations involve the main (primary) disk arrays and the remote (secondary) disk arrays. The main disk arrays contain the CA primary volumes (P-VOLs), which contain the original data and are connected to the application system. The remote disk arrays contain the CA secondary volumes (S-VOLs) connected to the backup system.
*See also* BC *(HP StorageWorks Disk Array XP specific term)*, Main Control Unit and HP StorageWorks Disk Array XP LDEV.

**CA+BC EVA**

*(HP StorageWorks EVA specific term)* The combination of Continuous Access (CA) EVA and Business Copy (BC) EVA enables you to create and maintain copies (replicas) of the source volumes on a remote EVA, and then use these copies as the source for local replication on this remote array.
*See also* BC EVA, replica, and source volume.

**CAP**

*(StorageTek specific term)* Cartridge Access Port is a port built into the door panel of a library. The purpose is to enter or eject media.

**catalog protection**

Defines how long information about backed up data (such as file names and file versions) is kept in the IDB.
*See also* data protection.

**CDB**

The Catalog Database is a part of the IDB that contains information about backup, restore, object copy, object consolidation, object verification, and media management sessions. Depending on the selected logging level, it also contains file names and file versions. This part of the database is always local to the cell.
*See also* MMDB.

**CDF file**

*(UNIX specific term)* A Context Dependent File is a file consisting of several files grouped under the same pathname. The system ordinarily selects one of the files using the context of the process. This mechanism allows machine dependent executables, system data, and device files to work correctly from all hosts in a cluster while using the same pathname.

**cell**

A set of systems that are under the control of a Cell Manager. The cell typically represents the systems on a site or an organizational entity, which are connected to the same LAN. Central control is available to administer the backup and restore policies and tasks.

**Cell Manager**

The main system in the cell where the essential Data Protector software is installed and from which all backup and restore activities are managed. The GUI used for management tasks can be located on a different system. Each cell has one Cell Manager system.

**centralized licensing**

Data Protector allows you to configure centralized licensing for the whole enterprise environment consisting of several cells. All Data Protector licenses are installed and kept on the Enterprise Cell Manager system. You can then allocate licenses to specific cells to suit your needs.
*See also* MoM.

**Centralized Media Management Database (CMMDB)**

*See* CMMDB.

**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**
*(Windows specific term)* 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'
- 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-aware application**

It is an application that supports the cluster Application Programming Interface. Each cluster-aware application declares its own critical resources (disk volumes (on Microsoft Cluster Server), volume groups (on MC/ServiceGuard), application services, IP names and addresses, and so on).

**cluster continuous replication**

*(Microsoft Exchange Server specific term)* Cluster continuous replication (CCR) is a high availability solution that uses cluster management and failover options to create and maintain an exact copy (CCR copy) of a storage group. A storage group is replicated to a separate server. CCR removes any single point of failure in your Exchange back-end servers. You can perform backups using VSS on your passive Exchange Server node where a CCR copy is located and thus reducing the load on the active node.

A CCR copy is used for disaster recovery since you can switch to the CCR copy in a few seconds. A replicated storage group is represented as a new instance of Exchange writer called Exchange Replication Service and can be backed up (using VSS) like an ordinary storage group.

*See also* Exchange Replication Service and local continuous replication.

**CMD Script for Informix Server**

*(Informix Server specific term)* A Windows CMD script that is created in INFORMIXDIR when an Informix Server database is configured. The CMD script is a set of system commands that export environment variables for Informix Server.

**CMMDB**

The Data Protector Centralized Media Management Database (CMMDB) is the result of merging MMDBs from several cells in the MoM environment. It allows you to share high-end devices and media across multiple cells in a MoM environment. One cell can control the robotics, serving the devices that are connected to systems in other cells. The CMMDB must reside on the Manager-of-Managers. A reliable network connection between the MoM cell and the other Data Protector cells is highly recommended

*See also* MoM.

**COM+ Class Registration Database**

*(Windows specific term)* The COM+ Class Registration Database and the Windows Registry store application attributes, class attributes, and computer-level attributes. This guarantees consistency among these attributes and provides common operation on top of these attributes.

**command-line interface (CLI)**

A set of DOS and UNIX like commands that you can use in shell scripts to perform Data Protector configuration, backup, restore, and management tasks.

**Command View (CV) EVA**

*(HP StorageWorks EVA specific term)* The user interface that enables you to configure, manage, and monitor your HP StorageWorks EVA storage system. It is used to perform various storage management tasks, for example, creating virtual disk families, managing storage system hardware, and creating snapclones and snapshots of virtual disks. The Command View EVA software runs on the HP Storage Management Appliance, and is accessed by a Web browser.
*See also* HP StorageWorks EVA SMI-S Agent and HP StorageWorks SMI-S EVA provider.

**Command View VLS**

*(VLS specific term)* A web browser-based GUI that is used to configure, manage, and monitor the VLS through a LAN.
*See also* Virtual Library System (VLS).

**concurrency**

*See* Disk Agent concurrency.

**control file**

*(Oracle and SAP R/3 specific term)* An Oracle data file that contains entries specifying the physical structure of the database. It provides database consistency information used for recovery.

**copy set**

*(HP StorageWorks EVA specific term)* A pair that consists of the source volumes on a local EVA and their replica on a remote EVA.
*See also* source volume, replica, and CA+BC EVA

**CRS**

The Cell Request Server process (service), which runs on the Data Protector Cell Manager, and starts and controls the backup and restore sessions. The service is started as soon as Data Protector is installed on the Cell Manager. On Windows systems, the CRS runs under the account of the user specified at installation time. On UNIX systems, it runs under the account `root`.

**CSM**

The Data Protector Copy and Consolidation Session Manager process controls the object copy and object consolidation sessions and runs on the Cell Manager system.

**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 stream**

Sequence of data transferred over the communication channel.

**Data_Protector_home**

On Windows Vista and Windows Server 2008, the directory containing Data Protector program files. On other Windows operating systems, the directory containing Data Protector program files and data files. 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**

On Windows Vista and Windows Server 2008, the directory containing Data Protector data files. 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.

**Data Replication (DR) group**

*(HP StorageWorks EVA specific term)* A logical grouping of EVA virtual disks. It can contain up to eight copy sets provided they have common characteristics and share a common CA EVA log.
*See also* copy set.

**database server**

A computer with a large database stored on it, such as the SAP R/3 or Microsoft SQL database. A server has a database that can be accessed by clients.

**Dbobject**

*(Informix Server specific term)* An Informix Server physical database object. It can be a blobspace, dbspace, or logical log file.

**DC directory**

The Detail Catalog (DC) directory contains DC binary files, which store information about file versions. It represents the DCBF part of the IDB, which occupies approximately 80% of the IDB. The default DC directory is called the `dcbf` directory and is located on the Cell Manager in the directory *Data_Protector_program_data*\db40 (Windows Server 2008), *Data_Protector_home*\db40 (other Windows systems), or `/var/opt/omni/server/db40` (UNIX systems). You can create more DC directories and use a custom location. Up to 50 DC directories are supported per cell. The default maximum size of a DC directory is 16 GB.

**DCBF**

The Detail Catalog Binary Files (DCBF) part of the IDB stores information about file versions and attributes. It occupies approximately 80% of the IDB. One DC binary file is created for each Data Protector medium used for backup. Its maximum size is limited by the file system settings.

**delta backup**

A delta backup is a backup containing all the changes made to the database from the last backup of any type.
*See also* backup types.

**device**

A physical unit which contains either just a drive or a more complex unit such as a library.

**device chain**

A device chain consists of several standalone devices configured for sequential use. When a medium in one device gets full, the backup automatically continues on a medium in the next device in the device chain.

**device group**

*(EMC Symmetrix specific term)* A logical unit representing several EMC Symmetrix devices. A device cannot belong to more than a single device group. All devices in a device group must be on the same EMC Symmetrix unit. You can use a device group to identify and work with a subset of the available EMC Symmetrix devices.

**device streaming**

A device is streaming if it can feed enough data to the medium to keep it moving forward continuously. Otherwise, the tape has to be stopped, the device waits for more data, reverses the tape a little and resumes to write to the tape, and so on. In other words, if the data rate written to the tape is less or equal the data rate which can be delivered to the device by the computer system, then the device is streaming. Streaming significantly improves the performance of the device and use of space.

**DHCP server**
A system running the Dynamic Host Configuration Protocol (DHCP) providing dynamic IP address assignment and network configuration for DHCP clients.

**differential backup**
An incremental backup that backs up changes made since the last full backup. To perform this type of backup, specify the Incr1 backup type.
*See 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.

**direct backup**
A SAN-based backup solution in which data movement directly from disk to tape (or to other secondary storage) is facilitated by the SCSI Extended Copy (Xcopy) command. Direct backup lessens the backup I/O load on systems in a SAN environment. The data movement is facilitated directly from disk to tape (or to other secondary storage) by the SCSI Extended Copy (XCopy) command. The command is provided by any element of the infrastructure including bridges, switches, tape libraries, and disk subsystems.
*See also* XCopy engine.

**directory junction**
*(Windows specific term)* Directory junctions use the reparse point concept of Windows. An NTFS 5 directory junction allows you to redirect a directory/file request to another location.

**disaster recovery**
A process to restore a client’s main system disk to a state close to the time when a (full) backup was performed.

**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.

**domain controller**

A server in a network that is responsible for user security and verifying passwords within a group of other servers.

**DR image**

Data required for temporary disaster recovery operating system (DR OS) installation and configuration.

**DR OS**

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 (typically a tape drive). It can also read the data from the medium and send it to the computer system.

**drive-based encryption**

Data Protector drive-based encryption uses the encryption functionality of the drive. While performing the backup, the drive encrypts both the data and the meta-data that is written to the medium.

**drive index**

A number that identifies the mechanical position of a drive inside a library device. This number is used by the robotic control to access a drive.

**dynamic client**

*See* client backup with disk discovery.

**EMC Symmetrix Agent (SYMA)** ***(EMC Symmetrix specific term)***

*See* Symmetrix Agent (SYMA).

**emergency boot file**

*(Informix Server specific term)* The Informix Server configuration file `ixbar.`*`server_id`* that resides in the directory *`INFORMIXDIR`*`/etc` (on Windows) or *`INFORMIXDIR`*`\etc` (on UNIX). *`INFORMIXDIR`* is the Informix Server home directory and *`server_id`* is the value of the `SERVERNUM` configuration parameter. Each line of the emergency boot file corresponds to one backup object.

**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 key 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 several `StoreID`s 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**

Also referred to as SCSI Exchanger.
*See also* library.

**exporting media**

A process that removes all data about backup sessions, such as systems, objects, and file names, which reside on the media from the IDB. Information about the media and their relation to a pool is also removed from the IDB. The data on the media remains unchanged.
*See also* importing media.

**Extensible Storage Engine (ESE)**

*(Microsoft Exchange Server specific term)* A database technology used as a storage system for information exchange in Microsoft Exchange Server.

**failover**

Transferring of the most important cluster data, called group (on Windows) or package (on UNIX) from one cluster node to another. A failover can occur mostly because of software or hardware failures or maintenance on the primary node.

**failover**

*(HP StorageWorks EVA specific term)* An operation that reverses the roles of source and destination in CA+BC EVA configurations.
*See also* CA+BC EVA.

**FC bridge**

*See* Fibre Channel bridge.

**Fibre Channel** An ANSI standard for high-speed computer interconnection. Using either optical or copper cables, it allows the high speed 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** *(Windows specific term)* 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 StorageWorks Disk Array XP specific term)* HP StorageWorks Disk Array XP allows up to three mirror copies of a primary volume and each of these copies can have additional two copies. The three mirror copies are called first-level mirrors. *See also* primary volume and MU number.

**flash recovery area**

*(Oracle specific term)* Flash recovery area is an Oracle 10g/11g managed directory, filesystem, or Automatic Storage Management disk group that serves as a centralized storage area for files related to backup and recovery (recovery files). *See also* recovery files.

**fnames.dat**

The `fnames.dat` files of the IDB contain information on the names of the backed up files. Typically, these files occupy about 20% of the IDB, if filenames are stored.

**formatting**

A process that erases any data contained on a medium and prepares it for use with Data Protector. Information about media (medium ID, description, and location) is saved in the IDB as well as on the respective media (media header). Data Protector media with protected data are not formatted until the protection expires or the media are unprotected/recycled.

**free pool**

An auxiliary source of media for use by media pools when they run out of media. The media pools must be configured to use free pools.

**full backup**

A backup in which all selected objects are backed up, whether or not they have been recently modified.
*See also* backup types.

**full database backup**

A backup of all data in a database, not only the data that has been changed after the last (full or incremental) backup of the database. A full database backup does not depend on any other backup.

**full mailbox backup**

A full mailbox backup is a backup of the entire mailbox content.

**full ZDB**

A ZDB to tape or ZDB to disk+tape session in which all selected objects are streamed to tape, even if there are no changes from the previous backup.
*See also* incremental ZDB.

**global options file**

A file that allows you to customize Data Protector. It explains the global options, which cover various aspects of Data

Protector, typically time-outs and limits, and affect the entire Data Protector cell. The file is located on the Cell Manager in the directory *Data_Protector_program_data*\Config\Server\Options (Windows Server 2008), *Data_Protector_home*\Config\Server\Options (other Windows systems), or /etc/opt/omni/server/options (HP-UX or Solaris systems).

**group** *(Microsoft Cluster Server specific term)* A collection of resources (for example disk volumes, application services, IP names, and addresses) that are needed to run a specific cluster-aware applications.

**GUI** A graphical user interface provided by Data Protector for easy access to all configuration, administration, and operation tasks. Besides the original Data Protector GUI that runs on Windows, Data Protector also provides a Java-based graphical user interface with the same look and feel, which runs on numerous platforms.

**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 on the Cell Manager in the directory *Data_Protector_program_data*\Config\Server\holidays (Windows Server 2008), *Data_Protector_home*\Config\Server\holidays (other Windows systems), or /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 Operations Manager**

HP Operations Manager provides powerful capabilities for operations management of a large number of systems and applications in a network. Data Protector provides an integration into this management product. This integration is implemented as a SMART Plug-In for HP Operations Manager management servers on Windows, HP-UX, Solaris, and Linux. Earlier versions of HP Operations Manager were called IT/Operation, Operations Center, Vantage Point Operations, and OpenView Operations.

**HP Operations Manager SMART Plug-In (SPI)**

A fully integrated, out-of-the-box solution which "plugs into" HP Operations Manager, extending the managed domain. Through the Data Protector integration, which is implemented as an HP Operations Manager SMART Plug-In, a user can have an arbitrary number of Data Protector Cell Managers monitored as an extension to HP Operations Manager.

**HP StorageWorks Disk Array XP LDEV**

A logical partition of a physical disk within an HP StorageWorks Disk Array XP. LDEVs are entities that can be replicated in the Continuous Access XP (CA) and Business Copy XP (BC) configurations, or can be used as standalone entities.
*See also* BC, CA *(HP StorageWorks Disk Array XP specific term)*, and replica.

**HP StorageWorks EVA SMI-S Agent**

A Data Protector software module that executes all tasks required for the HP StorageWorks Enterprise Virtual Array integration. With the EVA SMI-S Agent, the control over the array is established through HP StorageWorks SMI-S EVA provider, which directs communication between incoming requests and CV EVA.
*See also* Command View (CV) EVA and HP StorageWorks SMI-S EVA provider.

**HP StorageWorks SMI-S EVA provider**

An interface used for controlling HP StorageWorks Enterprise Virtual Array. SMI-S EVA provider runs as a separate service on the HP Storage Management Appliance system and acts as a gateway between incoming requests and Command View EVA. With the Data Protector HP StorageWorks EVA integration, SMI-S EVA provider accepts standardized requests from the EVA SMI-S Agent, communicates with Command View EVA for

information or method invocation, and returns standardized responses.
*See also* HP StorageWorks EVA SMI-S Agent and Command View (CV) EVA.

**HP StorageWorks Virtual Array LUN**
A logical partition of a physical disk within an HP StorageWorks Virtual Array. LUNs are entities that can be replicated in the HP StorageWorks Business Copy VA configuration, or can be used as standalone entities.
*See also* BC VA and replica.

**ICDA**
*(EMC Symmetrix specific term)* EMC's Symmetrix Integrated Cached Disk Arrays (ICDA) is a disk array device that combines a set of physical disks, a number of FWD SCSI channels, an internal cache memory, and control and diagnostic software commonly referred to as the microcode.

**IDB**
The Data Protector Internal Database is an embedded database located on the Cell Manager that keeps information regarding which data is backed up, on which media it is backed up, how backup and restore sessions are run, and which devices and libraries are configured.

**IDB recovery file**
An IDB file (obrindex.dat) with information about IDB backups, media, and devices used for the backup. This information can significantly simplify IDB recovery. It is recommended to relocate the file, together with IDB transaction logs, to a separate physical disk from other IDB directories, and, additionally, to make an additional copy of the file.

**importing media**
A process that re-reads all data about backup sessions which are on the medium back into the IDB. This then allows for fast and convenient access to the data on the media.
*See also* exporting media.

**incremental backup**
A backup that selects only files that have changed since a previous backup. Several levels of incremental backup are available, which enables detailed control of restore chain length.
*See also* backup types.

**incremental backup**
*(Microsoft Exchange Server specific term)* A backup of the Microsoft Exchange Server data that has changed since the last full or incremental backup. With the incremental backup, only the transaction log files are backed up.

*See also* backup types.

**incremental mailbox backup** An incremental mailbox backup backs up all the changes made to the mailbox after the last backup of any type.

**incremental1 mailbox backup** An incremental1 mailbox backup backs up all the changes made to the mailbox after the last full backup.

**incremental (re)-establish** *(EMC Symmetrix specific term)* A BCV or SRDF control operation. In BCV control operations, an incremental establish causes the BCV device to be synchronized incrementally and to function as an EMC Symmetrix mirrored medium. The EMC Symmetrix devices must have been previously paired. In SRDF control operations, an incremental establish causes the target (R2) device to be synchronized incrementally and to function as an EMC Symmetrix mirrored medium. The EMC Symmetrix devices must have been previously paired.

**incremental restore** *(EMC Symmetrix specific term)* A BCV or SRDF control operation. In BCV control operations, an incremental restore reassigns a BCV device as the next available mirror of the standard device in the pair. However, the standard devices are updated with only the data that was written to the BCV device during the time of the original pair split, and the data that was written to the standard device during the split is overwritten with data from the BCV mirror. In SRDF control operations, an incremental restore reassigns a target (R2) device as the next available mirror of the source (R1) device in the pair. However, the source (R1) devices are updated with only the data that was written to the target (R2) device during the time of the original pair split, and the data that was written to the source (R1) device during the split is overwritten with data from the target (R2) mirror.

**incremental ZDB** A filesystem ZDB to tape or ZDB to disk+tape session in which only changes from the last protected full or incremental backup are streamed to tape.
*See also* full ZDB.

**Inet** A process that runs on each UNIX system or service that runs on each Windows system in the Data Protector cell. It is responsible for communication between systems in the cell and for starting other processes needed for backup and restore. The Inet service is started as soon as Data Protector is installed on a system. The Inet process is started by the inetd daemon.

**Information Store**

*(Microsoft Exchange Server specific term)* The Microsoft Exchange Server service that is responsible for storage management. Information Store in Microsoft Exchange Server manages two kinds of stores: mailbox stores and public folder stores. A mailbox store consists of mailboxes that belong to individual users. A public folder store contains public folders and messages that are shared among several users.
*See also* Key Management Service and Site Replication Service.

**Informix Server**

*(Informix Server specific term)* Refers to Informix Dynamic Server.

**initializing**

*See* formatting.

**Installation Server**

A computer system that holds a repository of the Data Protector software packages for a specific architecture. The Installation Server is used for remote installation of Data Protector clients. In mixed environments at least two Installation Servers are needed: one for UNIX systems and one for Windows systems.

**instant recovery**

*(ZDB specific term)* A process in which a replica, produced by a ZDB-to-disk or a ZDB-to-disk+tape session, is used to restore the contents of the source volumes to their states at the time at which the replica was created, avoiding the need to perform a restore from tape. Depending on the application or database concerned, this may be all that is required, or other steps, such as the application of transaction log files, may be required for full recovery.
*See also* replica, zero downtime backup (ZDB), ZDB to disk, and ZDB to disk+tape.

**integration object**

A backup object of a Data Protector integration, such as Oracle or SAP DB.

**Internet Information Services (IIS)**

*(Windows specific term)* Microsoft Internet Information Services is a network file and application server that supports multiple protocols. Primarily, IIS transmits information in Hypertext Markup Language (HTML) pages by using the Hypertext Transport Protocol (HTTP).

**IP address**

An Internet Protocol address is a numeric address of a system used to uniquely identify the system on the network. The IP address consists of four groups of numbers separated by periods (full stops).

**ISQL**

*(Sybase specific term)* A Sybase utility used to perform system administration tasks on Sybase SQL Server.

**Java GUI Client**

The Java GUI Client is a component of the Java GUI that contains only user interface related functionalities and requires connection to the Java GUI Server to function.

**Java GUI Server**

The Java GUI Server is a component of the Java GUI that is installed on the Data Protector Cell Manager system. The Java GUI Server receives requests from the Java GUI Client, processes them and then sends the responses back to the Java GUI Client. The communication is done through Hypertext Transfer Protocol (HTTP) on port 5556.

**jukebox**

*See* library.

**jukebox device**

A device consisting of multiple slots used to store either optical or file media. When being used to store file media, the jukebox device is known as the “file jukebox device”.

**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.

**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.

**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.

**keystore**

All encryption keys are centrally stored in the keystore on the Cell Manager and administered by the Key Management Server (KMS).

**LBO**

*(EMC Symmetrix specific term)* A Logical Backup Object (LBO) is an object of data storage/retrieval in the EMC Symmetrix environment. It is stored/retrieved by EMC Symmetrix as one entity and can only be restored as a whole.

**library**

Also called autochanger, jukebox, autoloader, or exchanger. A library contains media in repository slots. Each slot holds one medium (for example, DDS/DAT). Media are moved between slots and drives by a robotic mechanism, allowing random access to media. The library can contain multiple drives.

**lights-out operation or unattended operation**

A backup or restore operation that takes place outside of normal business hours without an operator. This implies that no operator personnel is present to work with the backup application or service mount requests, for example.

**LISTENER.ORA**

*(Oracle specific term)* An Oracle configuration file that describes one or more Transparent Network Substrate (TNS) listeners on a server.

**load balancing**

By default, Data Protector automatically balances the usage of devices selected for backup, so that they are used evenly. Load balancing optimizes the device usage by balancing the number of objects written to each device. Since load balancing is done automatically during backup time, you do not need to manage how the data is actually backed up. You just specify the devices to be used. If you do not want to use load balancing, you can select which device will be used for each object in the backup specification. Data Protector will access the devices in the specified order.

**local and remote recovery**

Remote recovery is performed if all Media Agent hosts specified in the SRD file are accessible. If any of them fails, the disaster recovery process fails over to the local mode. This means that the target system is searched for locally attached devices. If only one device is found, it is automatically used. Otherwise, Data Protector prompts you to select the device, which will be used for restore.

**local continuous replication**

*(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 specific term)* A script provided by ON-Bar that you can use to start backing up logical log files when Informix Server issues a logfull event alarm. The Informix Server `ALARMPROGRAM` configuration parameter defaults to the *`INFORMIXDIR`*`/etc/log_full.sh`, where *`INFORMIXDIR`* is the Informix Server home directory. If you do not want logical logs to be backed up continuously, set the `ALARMPROGRAM` configuration parameter to *`INFORMIXDIR`*`/etc/no_log.sh`.

**logging level**

The logging level determines the amount of details on files and directories written to the IDB during backup, object copying, or object consolidation. You can always restore your data, regardless of the logging level used during backup. Data Protector provides four logging levels: Log All, Log Directories, Log Files, and No Log. The different logging level settings influence the IDB growth, backup speed, and the convenience of browsing data for restore.

**logical-log files**

This applies to online database backup. Logical-log files are files in which modified data is first stored before being flushed to disk. In the event of a failure, these logical-log files are used to roll forward all transactions that have been committed as well as roll back any transactions that have not been committed.

**login ID**

*(Microsoft SQL Server specific term)* The name a user uses to log on to Microsoft SQL Server. A login ID is valid if Microsoft SQL Server has an entry for that user in the system table syslogin.

**login information to the Oracle Target Database**

*(Oracle and SAP R/3 specific term)* The format of the login information is `user_name/password@service`, where:

- `user_name` is the name by which a user is known to Oracle Server and to other users. Every user name is associated with a password and both have to be entered to connect to an Oracle Target Database. This user must have Oracle SYSDBA or SYSOPER rights.
- `password` must be the same as the password specified in the Oracle password file (`orapwd`), which is used for authentication of users performing database administration.
- `service` is the name used to identify an SQL*Net server process for the target database.

**login information to the Recovery Catalog Database**

*(Oracle specific term)* The format of the login information to the Recovery (Oracle) Catalog Database is `user_name/password@service`, where the description of the user name, password, and service name is the same as in the Oracle SQL*Net V2 login information to the Oracle target database. In this case, `service` is the name of the service to the Recovery Catalog Database, not the Oracle target database. Note that the Oracle user specified here must be the owner of the Oracle Recovery Catalog.

**Lotus C API**

*(Lotus Domino Server specific term)* An interface for the exchange of backup and recovery information between Lotus Domino Server and a backup solution, like Data Protector.

**LVM**

A Logical Volume Manager is a subsystem for structuring and mapping physical disk space to logical volumes on UNIX systems. An LVM system consists of several volume groups, where each volume group has several volumes.

**Magic Packet**

*See* Wake ONLAN.

**mailbox**

*(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 StorageWorks Disk Array XP specific term)* An HP StorageWorks XP disk array that contains the primary volumes for the CA and BC configurations and acts as a master device. *See also* BC (HP StorageWorks Disk Array XP specific term), CA *(HP StorageWorks Disk Array XP specific term)*, and HP StorageWorks Disk Array XP LDEV.

**Manager-of-Managers (MoM)**

*See* MoM.

**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.

**MAPI**

*(Microsoft Exchange Server 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 sssion, the Disk Agent then writes the data to the disk. A Media Agent also manages the robotics control of a library.

**media allocation policy** Determines in which sequence media are used for backup. The Strict allocation policy directs Data Protector to prompt for a specific medium. The Loose policy directs Data Protector to prompt for any suitable medium. The Formatted First policy directs Data Protector to give preference to unknown media, even if unprotected media are available in the library.

**media condition** The quality of a medium as derived from the media condition factors. Heavy usage and age result in an increased number of read and write errors with tape media. Media need to be replaced when they are marked as POOR.

**media condition factors** The user-assigned age threshold and overwrite threshold used to determine the state of a medium.

**medium ID** A unique identifier assigned to a medium by Data Protector.

**media label** A user-defined identifier used to describe a medium.

**media location** A user-defined physical location of a medium, such as "building 4" or "off-site storage".

**media management session** A session performing some action on a medium, such as initializing, scanning the content, verifying data on a medium, or copying a medium.

**media pool** A set of media of the same type (such as DDS) used and tracked as a group. Media are formatted and assigned to a media pool.

**media set** The result of a backup session is data backed up on a group of media called media set. Depending on the media usage policy, several sessions can share the same media.

**media type** The physical type of media, such as DDS or DLT.

**media usage policy** The media usage policy controls how new backups are added to the already used media. It can be `Appendable`, `Non-Appendable`, or `Appendable for incrementals only`.

**merging**

This defines one mode to resolve file conflicts during restore. If the file to be restored already exists at the destination, the one with the more recent modification date is kept. Files not present on the disk are always restored.
*See also* overwrite.

**Microsoft Exchange Server**

A "client-server" messaging and a workgroup system that offers a transparent connection to many different communication systems. It provides users with an electronic mail system, individual and group scheduling, online forms, and workflow automation tools. It provides a developer with a platform on which to build custom information-sharing and messaging-service applications.

**Microsoft Management Console (MMC)**

*(Windows specific term)* An administration model for Windows-based environments. It provides a simple, consistent, and integrated administration user interface allowing management of many applications through the same GUI, provided that the applications adhere to the MMC model.

**Microsoft SQL Server**

A database management system designed to meet the requirements of distributed "client-server" computing.

**Microsoft Volume Shadow Copy Service (VSS)**

A software service that provides a unified communication interface to coordinate backup and restore of a VSS-aware application regardless of its specific features. This service collaborates with the backup application, writers, shadow copy providers, and the operating system kernel to implement the management of volume shadow copies and shadow copy sets.
*See also* shadow copy, shadow copy provider, replica, and writer.

**mirror *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)***

*See* target volume.

**mirror rotation *(HP StorageWorks Disk Array XP specific term)***

*See* replica set rotation.

**MMD**
The Media Management Daemon process (service) runs on the Data Protector Cell Manager and controls media management and device operations. The process is started when Data Protector is installed on the Cell Manager.

**MMDB**
The Media Management Database (MMDB) is a part of the IDB that contains information about media, media pools, devices, libraries, library drives, and slots configured in the cell, as well as the Data Protector media used for backup. In an enterprise backup environment, this part of the database can be common to all cells.
*See also* CMMDB, CDB.

**MoM**
Several cells can be grouped together and managed from a central cell. The management system of the central cell is the Manager-of-Managers (MoM). The cells are called MoM clients. The MoM enables you to configure and manage multiple cells from a central point.

**mount request**
A screen prompt that tells you to insert a specific medium into a device. Once you respond to the mount request by providing the required medium and confirm the mount request, the session continues.

**mount point**
The access point in a directory structure for a disk or logical volume, for example`/opt` or `d:`. On UNIX, the mount points are displayed using the `bdf` or `df` command.

**MSM**
The Data Protector Media Session Manager, which runs on the Cell Manager and controls media sessions, such as copying media.

**MU number**
*(HP StorageWorks Disk Array XP specific term)* Mirror Unit number. An integer number (0, 1 or 2), used to indicate a first-level mirror.
*See also* first-level mirror.

**multi-drive server**
A license that allows you to run an unlimited number of Media Agents on a single system. This license, which is bound to the IP address of the Cell Manager, is no longer available.

**obdrindex.dat**
*See* IDB recovery file.

**OBDR capable device** A device that can emulate a CD-ROM drive loaded with a bootable disk and can thus be used as a backup or boot device for disaster recovery purposes.

**object** *See* backup object.

**object consolidation** The process of merging a restore chain of a backup object, consisting of a full backup and at least one incremental backup, into a new, consolidated version of this object. The process is a part of the synthetic backup procedure. The result is a synthetic full backup of the specified backup object.

**object consolidation session** A process that merges a restore chain of a backup object, consisting of a full backup and at least one incremental backup, into a new, consolidated version of this object.

**object copy** A copy of a specific object version that is created during an object copy session or a backup session with object mirroring.

**object copy session** A process that creates an additional copy of the backed up data on a different media set. During an object copy session, the selected backed up objects are copied from the source to the target media.

**object copying** The process of copying selected object versions to a specific media set. You can select object versions from one or several backup sessions to be copied.

**object ID** *(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 (several minutes or hours). For instance, for backup to tape, until streaming of data to the tape is finished.
- For ZDB methods, the database is also put into the quiescent state, but for the period of the data replication process only (several seconds). Normal database operation can then be resumed for the rest of the backup process.

*See also* zero downtime backup (ZDB) and online backup.

**offline recovery**

Offline recovery is performed if the Cell Manager is not accessible, for example, due to network problems. Only standalone and SCSI library devices can be used for offline recovery. Recovery of the Cell Manager is always offline.

**offline redo log**

*See* archived redo log.

**ON-Bar**

*(Informix Server specific term)* A backup and restore system for Informix Server. ON-Bar enables you to create a copy of your Informix Server data and later restore the data. The ON-Bar backup and restore system involves the following components:

- the `onbar` command
- Data Protector as the backup solution
- the XBSA interface
- ON-Bar catalog tables, which are used to back up dbobjects and track instances of dbobjects through multiple backups.

**ONCONFIG**

*(Informix Server specific term)* An environment variable that specifies the name of the active ONCONFIG configuration file. If the ONCONFIG environment variable is not present, Informix Server uses the configuration values from the `onconfig` file in the directory `*INFORMIXDIR*\etc` (on Windows) or `*INFORMIXDIR*/etc/` (on UNIX).

**online backup** A backup performed while a database application remains available for use. The database is placed into a special backup mode of operation for the time period that the backup application requires access to the original data objects. During this period, the database is fully operational, but there may be a small performance impact and log files may grow very quickly.

- For simple backup methods (non ZDB), backup mode is required for the whole backup period (several minutes or hours). For instance, for backup to tape, until streaming of data to tape is finished.
- For ZDB methods, backup mode is required for the short period of the data replication process only (several seconds). Normal database operation can then be resumed for the rest of the backup process.

In some cases, transaction logs may also have to be backed up to allow a consistent database to be restored.
*See also* zero downtime backup (ZDB), and offline backup.

**online redo log** *(Oracle specific term)* Redo logs that have not been archived, but are either available to the instance for recording database activity or are filled and waiting to be archived or reused.
*See also* archived redo log.

**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**

Backup ownership affects the ability of users to see and restore data. Each backup session and all the data backed up within it is assigned an owner. The owner can be the user that starts an interactive backup, the account under which the CRS process is running, or the user specified as the owner in the backup specification options.
If a user starts an existing backup specification without modifying it, the backup session is not considered as interactive.
If a modified backup specification is started by a user, the user is the owner unless the following is true:

- The user has the Switch Session Ownership user right.
- The backup session owner is explicitly defined in the backup specification, where the username, group or domain name, and the system name are specified.

If a backup is scheduled on a UNIX Cell Manager, the session owner is root:sys unless the above conditions are true.
If a backup is scheduled on a Windows Cell Manager, the session owner is the user specified during the installation, unless the above conditions are true.

**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 Server 2008), *Data_Protector_home*\Config\Server\dr\p1s (other Windows systems), or /etc/opt/omni/server/dr/p1s (UNIX systems) with the filename recovery.p1s.

**package** *(MC/ServiceGuard and Veritas Cluster specific term)* A collection of resources (for example volume groups, application services, IP names and addresses) that are needed to run a specific cluster-aware application.

**pair status** *(HP StorageWorks Disk Array XP specific term)* A mirrored pair of disks can have various status values depending on the action performed on it. The three most important status values are:

- COPY - The mirrored pair is currently re-synchronizing. Data is transferred from one disk to the other. The disks do not contain the same data.
- PAIR - The mirrored pair is completely synchronized and both disks (the primary volume and the mirrored volume) contain identical data.
- SUSPENDED - The link between the mirrored disks is suspended. That means that both disks are accessed and updated independently. However, the mirror relationship is still maintained and the pair can be re-synchronized without transferring the complete disk.

**parallel restore** Restoring backed up data to multiple disks at the same time (that is, in parallel) by running multiple Disk Agents, that receive data from one Media Agent. For the parallel restore to work, select data that is located on different disks or logical volumes and during backup, the data from the different objects must have been sent to the same device using a concurrency of 2 or more. During a parallel restore, the data for multiple objects selected for restore is read from media at the same time, thereby improving performance.

**parallelism** The concept of reading multiple data streams from an online database.

**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 and as shell scripts on UNIX.
*See also* pre-exec.

**pre- and post-exec commands** Pre- and post-exec commands are used to perform additional action before and after a backup or restore session. They are not supplied by Data Protector. You need to create your own commands. They can be written as executables or batch files on Windows and as shell scripts on UNIX.

**prealloc list** A subset of media in a media pool that specifies the order in which media are used for backup.

**pre-exec** A backup option that executes a command or script before the backup of an object or before the entire session is started. Pre-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows and as shell scripts on UNIX.
*See also* post-exec.

**primary volume (P-VOL)** *(HP StorageWorks Disk Array XP specific term)* Standard HP StorageWorks Disk Array XP LDEVs that act as a primary volume for the CA and BC configurations. The P-VOL is located in the MCU.
*See also* secondary volume (S-VOL) and Main Control Unit (MCU).

**protection** *See* data protection and also catalog protection.

**public folder store**

*(Microsoft Exchange Server specific term)* The part of the Information Store that maintains information in public folders. A public folder store consists of a binary rich-text .edb file and a streaming native internet content .stm file.

**public/private backed up data**

When configuring a backup, you can select whether the backed up data will be:

- public, that is visible (and accessible for restore) to all Data Protector users
- private, that is, visible (and accessible for restore) only to the owner of the backup and administrators

**RAID**

Redundant Array of Inexpensive Disks.

**RAID Manager Library**

*(HP StorageWorks Disk Array XP specific term)* The RAID Manager Library is used internally by Data Protector on Solaris systems to allow access to HP StorageWorks Disk Array XP configuration, status, and performance data and to key HP StorageWorks Disk Array XP features through the use of function calls translated into a sequence of low level SCSI commands.

**RAID Manager XP**

*(HP StorageWorks Disk Array XP specific term)* The RAID Manager XP application provides an extensive list of commands to report and control the status of the CA and BC applications. The commands communicate through a RAID Manager instance with the HP StorageWorks Disk Array XP Disk Control Unit. This instance translates the commands into a sequence of low level SCSI commands.

**rawdisk backup**

*See* disk image backup.

**RCU**

*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.

**RDS**

The Raima Database Server process (service) runs on the Data Protector Cell Manager and manages the IDB. The process is started when Data Protector is installed on the Cell Manager.

**Recovery Catalog**

*(Oracle specific term)* A set of Oracle tables and views that are used by Recovery Manager to store information about Oracle databases. This information is used by Recovery Manager to manage the backup, restore, and recovery of Oracle databases. The recovery catalog contains information about:

- The physical schema of the Oracle target database
- Data file and archived log backup sets
- Data file copies
- Archived Redo Logs
- Stored scripts

**Recovery Catalog Database**

*(Oracle specific term)* An Oracle database that contains a recovery catalog schema. You should not store the recovery catalog in your target database.

**recovery files**

*(Oracle specific term)* Recovery files are Oracle 10g/11g specific files that reside in the flash recovery area: the current control file, online redo logs, archived redo logs, flashback logs, control file autobackups, datafile copies, and backup pieces. *See also* flash recovery area.

**RecoveryInfo**

When backing up Windows configuration files, Data Protector collects the information about the current system configuration (information on disk layout, volume, and network configuration). This information is needed for disaster recovery.

**Recovery Manager (RMAN)**

*(Oracle specific term)* An Oracle command-line interface that directs an Oracle Server process to back up, restore, or recover the database it is connected to. RMAN uses either the recovery catalog or the control file to store information about backups. This information can be used later in restore sessions.

**recycle**

A process that removes the data protection from all backed up data on a medium, allowing Data Protector to overwrite it during one of the next backups. Data that belongs to the same session(s) but resides on other media is also unprotected. Recycling does not actually alter the data on the medium.

**redo log**

*(Oracle specific term)* Every Oracle database has a set of two or more redo log files. The set of redo log files for the database is known as the database's redo log. Oracle uses the redo log to record all changes made to data.

**Remote Control Unit (RCU)**

*(HP StorageWorks Disk Array XP specific term)* The Remote Control Unit (RCU) acts as a slave of an MCU in a CA configuration. In bidirectional configurations, the RCU can act as an MCU.

**Removable Storage Management Database**

*(Windows specific term)* A Windows service used for managing removable media (such as tapes and disks) and storage devices (libraries). Removable Storage allows applications to access and share the same media resources.

**reparse point**

*(Windows specific term)* A system-controlled attribute that can be associated with any directory or file. The value of a reparse attribute can have user-controlled data. The format of the data is understood by the application that stored the data and a filesystem filter that was installed to interpret the data and process such files. Whenever the filesystem encounters a file with a reparse point, it attempts to find the filesystem filter associated with the data format.

**replica**

*(ZDB specific term)* An image, at a particular point in time, of the data in source volumes that contain user-specified backup objects. Depending on the hardware or software with which it is created, the image may be an independent exact duplicate (clone) of the storage blocks at the physical disk level (for example, a split mirror or snapclone), or a virtual copy (for example, a snapshot). From perspective of a basic operating system, the complete physical disk containing backup objects is replicated. However, if a volume manager is used on UNIX, the whole volume or disk group containing a backup object (logical volume) is replicated. If partitions are used on Windows, the whole physical volume containing the selected partition is replicated.
*See also* snapshot, snapshot creation, split mirror, and split mirror creation.

**replica set**

*(ZDB specific term)* A group of replicas, all created using the same backup specification.
*See also* replica and replica set rotation.

**replica set rotation**

*(ZDB specific term)* The use of a replica set for regular backup production: Each time the same backup specification requiring the use of a replica set is run, a new replica is created and added to the set, until the maximum number for the set is

reached. After that, the oldest replica in the set is replaced, maintaining the maximum number of replicas in the set.
*See also* replica and replica set.

**restore chain**
All backups that are necessary for a restore of a backup object to a certain point in time. A restore chain consists of a full backup of the object and any number of related incremental backups.

**restore session**
A process that copies data from backup media to a client.

**resync mode**
*(HP StorageWorks Disk Array XP VSS provider specific term)* One of two XP VSS hardware provider operation modes. When the XP provider is in the resync mode, the source volume (P-VOL) and its replica (S-VOL) are in the suspended mirror relationship after a backup. The maximum number of replicas (S-VOLs per a P-VOL) rotated is three provided that MU range is 0-2 or 0, 1, 2. Restore from a backup in such a configuration is possible only by re-synchronization of an S-VOL with its P-VOL.
*See also* VSS compliant mode, source volume, primary volume (P-VOL), replica, secondary volume (S-VOL), MU number, and replica set rotation.

**RMAN *(Oracle specific term)***
*See* Recovery Manager.

**RSM**
The Data Protector Restore Session Manager controls 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.

**scan**
A function that identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library).

**scanning**
A function which identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library). It is

useful to perform a scan and check the actual media in the device if someone has manually manipulated media without using Data Protector to eject or enter, for example.

**Scheduler**

A function that controls when and how often automatic backups occur. By setting up a schedule, you automate the start of backups.

**secondary volume (S-VOL)**

*(HP StorageWorks Disk Array XP specific term)* secondary volumes, or S-VOLs, are XP LDEVs that act as a secondary CA or BC mirror of another LDEV (a P-VOL). In the case of CA, S-VOLs can be used as failover devices in a MetroCluster configuration. The S-VOLs are assigned separate SCSI addresses, different from the addresses used by the P-VOLs.
*See also* primary volume (P-VOL) 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-exec and post-exec script is a Data Protector unique identification of any session, including preview sessions. The session key is not recorded in the database, and it is used for specifying options for the `omnimnt`, `omnistat`, and `omniabort` commands.

**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.

**SIBF**

The Serverless Integrations Binary Files (SIBF) is a part of the IDB that stores raw NDMP meta data. This data is necessary to perform restore of NDMP objects.

**Site Replication Service**

*(Microsoft Exchange Server specific term)* The Microsoft Exchange Server 2000/2003 service that permits compatibility with Microsoft Exchange Server 5.5 by emulating the Exchange Server 5.5 directory service.
*See also* Information Store and Key Management Service.

**slot**

A mechanical position in a library. Each slot can hold a medium, such as a DLT tape. Data Protector references each slot with a number. To read a medium, a robotic mechanism moves the medium from a slot into the drive.

**SMB**

*See* split mirror backup.

**smart copy**

*(VLS specific term)* A copy of the backed up data created from the virtual tape to the physical tape library. The smart copy process allows Data Protector to distinguish between the source and the target medium thus enabling media management.
*See also* Virtual Library System (VLS).

**smart copy pool**

*(VLS specific term)* A pool that defines which destination library slots are available as smart copy targets for a specified source virtual library.
*See also* Virtual Library System (VLS) and smart copy.

**SMBF**

The Session Messages Binary Files (SMBF) part of the IDB stores session messages generated during backup, restore, object copy, object consolidation, object verification, and media management sessions. One binary file is created per session. The files are grouped by year and month.

**snapshot**

*(HP StorageWorks VA and HP StorageWorks EVA specific term)* A form of replica produced using snapshot creation techniques.

A range of snapshot types is available, with different characteristics, depending on the arrays/techniques used. Such replicas are dynamic and may be either virtual copies, still reliant upon the contents of the source volumes, or independent exact duplicates (clones), depending on the snapshot type and the time since creation.
*See also* replica and snapshot creation.

**snapshot backup *(HP StorageWorks VA and HP StorageWorks EVA specific term)***

*See* ZDB to tape, ZDB to disk, and ZDB to disk+tape.

**snapshot creation**

*(HP StorageWorks VA and HP StorageWorks EVA specific term)* A replica creation technique, in which copies of source volumes are created using storage virtualization techniques. The replicas are considered to be created at one particular point in time, without pre-configuration, and are immediately available for use. However background copying processes normally continue after creation.
*See also* snapshot.

**source (R1) device**

*(EMC Symmetrix specific term)* An EMC Symmetrix device that participates in SRDF operations with a target (R2) device. All writes to this device are mirrored to a target (R2) device in a remote EMC Symmetrix unit. An R1 device must be assigned to an RDF1 group type.
*See also* target (R2) device.

**source volume**

*(ZDB specific term)* A storage volume containing data to be replicated.

**sparse file**

A file that contains data with portions of empty blocks. Examples are: a matrix in which some or much of the data contains zeros, files from image applications, and high-speed databases. If sparse file processing is not enabled during restore, it might be impossible to restore this file.

**split mirror**

*(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)* A replica created using split mirror techniques. Such a replica provides an independent, exact duplicate, or clone of the contents of the source volumes.
*See also* replica and split mirror creation.

**split mirror backup *(EMC Symmetrix specific term)***

*See* ZDB to tape.

**split mirror backup *(HP StorageWorks Disk Array XP specific term)***

*See* ZDB to tape, ZDB to disk, and ZDB to disk+tape.

**split mirror creation**

*(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)* A replica creation technique, in which a pre-configured set of target volumes (a mirror) is kept synchronized with a set of source volumes until the time at which a replica of the contents of the source volumes is required. Then, the synchronization is stopped (the mirror is split) and a split mirror replica of the source volumes at the time of the split remains in the target volumes.
*See also* split mirror.

**split mirror restore**

*(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)* A process in which data backed up in a ZDB-to-tape or a ZDB-to-disk+tape session is restored from tape media to a split mirror replica, which is then synchronized to the source volumes. Individual backup objects or complete sessions can be restored using this method.
*See also* ZDB to tape, ZDB to disk+tape, and replica.

**sqlhosts file**

*(Informix Server specific term)* An Informix Server connectivity information file (on UNIX) or registry (on Windows) that contains the names of each of the database servers and any aliases to which the clients on a host computer can connect.

**SRD file**

*(disaster recovery specific term)* 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**

*(HP StorageWorks Disk Array XP specific term)* A Data Protector software module that executes all tasks required for a split mirror backup integration. It communicates with the HP StorageWorks Disk Array XP storing system using the RAID Manager XP utility (HP-UX and Windows systems) or RAID Manager Library (Solaris systems).

**sst.conf file**

The file `/usr/kernel/drv/sst.conf` is required on each Data Protector Sun Solaris client to which a multi-drive library device is connected. It must contain an entry for the SCSI address of the robotic mechanism of each library device connected to the client.

**st.conf file**

The file `/kernel/drv/st.conf` is required on each Data Protector Solaris client with a backup device connected. It must contain device information and a SCSI address for each backup drive connected to that client. A single SCSI entry is required for a single-drive device and multiple SCSI entries are required for a multi-drive library device.

**stackers**

Devices with multiple slots for media storage usually with only one drive. A stacker selects media from the stack sequentially. In contrast, a library can randomly select media from its repository.

**standalone file device**

A file device is a file in a specified directory to which you back up data.

**Storage Group**

*(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.

**StorageTek ACS library**

*(StorageTek specific term)* Automated Cartridge System is a library system (also known as Silo) consisting of one Library Management Unit (LMU) and one to 24 Library Storage Modules (LSM) connected to the unit.

**storage volume**

*(ZDB specific term)* A storage volume represents an object that may be presented to an operating system or some other entity (for example, a virtualization mechanism) upon which volume

management systems, file systems, or other objects may exist. The volume management systems, file systems are built on this storage. Typically, these can be created or exist within a storage system such as a disk array.

**switchover**

*See* failover.

**Sybase Backup Server API**

*(Sybase specific term)* An industry-standard interface developed for the exchange of backup and recovery information between a Sybase SQL Server and a backup solution like Data Protector.

**Sybase SQL Server**

*(Sybase specific term)* The server in the Sybase "client-server" architecture. Sybase SQL Server manages multiple databases and multiple users, keeps track of the actual location of data on disks, maintains mapping of logical data description to physical data storage, and maintains data and procedure caches in memory.

**Symmetrix Agent (SYMA)**

*(EMC Symmetrix specific term)* The Data Protector software module that prepares the EMC Symmetrix environment for backup and restore operations.

**synthetic backup**

A backup solution that produces a synthetic full backup, an equivalent to a conventional full backup in terms of data, without putting stress on the production servers or the network. A synthetic full backup is created from a previous full backup and any number of incremental backups.

**synthetic full backup**

The result of an object consolidation operation, where a restore chain of a backup objects is merged into a new, synthetic full version of this object. A synthetic full backup is equivalent to a conventional full backup in terms of restore speed.

**System Backup to Tape**

*(Oracle specific term)* An Oracle interface that handles the actions required to load, label, and unload correct backup devices when Oracle issues a backup or restore request.

**system databases**

*(Sybase specific term)* The four system databases on a newly installed Sybase SQL Server are the:

- master database (master)
- temporary database (tempdb)
- system procedure database (sybsystemprocs)
- model database (model).

**System 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.

**tablespace**
A part of a database structure. Each database is logically divided into one or more tablespaces. Each tablespace has data files or raw volumes exclusively associated with it.

**tapeless backup *(ZDB specific term)***
*See* ZDB to disk.

**target database**
*(Oracle specific term)* In RMAN, the target database is the database that you are backing up or restoring.

**target (R2) device**
*(EMC Symmetrix specific term)* An EMC Symmetrix device that participates in SRDF operations with a source (R1) device. It resides in the remote EMC Symmetrix unit. It is paired with a source (R1) device in the local EMC Symmetrix unit and receives all write data from its mirrored pair. This device is not accessed by user applications during normal I/O operations. An R2 device must be assigned to an RDF2 group type.
*See also* source (R1) device.

**target system**
*(disaster recovery specific term)* A system after a computer disaster has occurred. The target system is typically in a non-bootable state and the goal of disaster recovery is to restore

this system to the original system configuration. The difference between a 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 Symmetrix Logical Devices (SLDs). The instant copy is created on specially preconfigured SLDs called BCVs and is accessible via a separate device address to the system(s).

**TLU** Tape Library Unit.

**TNSNAMES.ORA** *(Oracle and SAP R/3 specific term)* A network configuration file that contains connect descriptors mapped to service names. The file may be maintained centrally or locally, for use by all or individual clients.

**transaction** A mechanism for ensuring that a set of actions is treated as a single unit of work. Databases use transactions to keep track of database changes.

**transaction backup** Transaction backups generally use fewer resources than database backups, so they can be created more frequently than database backups. By applying transaction backups, you can recover the database to a specific point in time prior to when a problem occurred.

**transaction backup** *(Sybase and SQL specific term)* A backup of the transaction log providing a record of changes made since the last full or transaction backup.

**transaction log backup** Transaction log backups generally use fewer resources than database backups so they can be created more frequently than database backups. By applying transaction log backups, you can recover the database to a specific point in time.

**transaction log files**
Files that record transactions of the database modifications, and provide fault tolerance in case of a database disaster.

**transaction logs**
*(Data Protector specific term)* Keep track of IDB changes. The archiving of transaction logs should be enabled to prevent you from losing the transaction log files that are created after the last IDB backup and are necessary for IDB recovery.

**transaction log table**
*(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 where a backup can be performed.
*See also* Microsoft Volume Shadow Copy Service (VSS).

**TSANDS.CFG file**
*(Novell NetWare specific term)* A file that allows you to specify the names of containers where you want backups to begin. It is a text file located in the `SYS:SYSTEM\TSA` directory on the server where `TSANDS.NLM` is loaded.

**UIProxy**
The Java GUI Server (`UIProxy` service) runs on the Data Protector Cell Manager. It is responsible for communication between the Java GUI Client and the Cell Manager, moreover, it performs business logic operations and sends only important information to the client. The service is started as soon as Data Protector is installed on the Cell Manager.

**unattended operation**
*See* lights-out operation.

**user account (Data Protector user account)**
You can use Data Protector only if you have a Data Protector user account, which restricts unauthorized access to Data Protector and to backed up data. Data Protector administrators create this account specifying a user logon name, the systems from which the user can log on, and a Data Protector user group membership. This is checked whenever the user starts the Data Protector user interface or performs specific tasks.

**User 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.

**vaulting media** The process of storing media to a safe and remote place. The media are brought back to the data center when they are needed for restore or are ready for reuse in subsequent backups. The vaulting procedure depends on your company's backup strategy and policies for data protection/reliability.

**verify** A function that lets you check whether the Data Protector data on a specified medium is readable. Additionally, consistency within each block can be checked if the backup was performed with the cyclic redundancy check (CRC) option ON.

**Virtual Controller Software (VCS)** *(HP StorageWorks EVA specific term)* The firmware that manages all aspects of storage system operation, including communication with Command View EVA through the HSV controllers.
*See also* Command View (CV) EVA.

**Virtual Device Interface**

*(Microsoft SQL Server specific term)* This is a SQL Server programming interface that allows fast backup and restore of large databases.

**virtual disk**

*(HP StorageWorks EVA specific term)* A unit of storage allocated from an HP StorageWorks Enterprise Virtual Array storage pool. Virtual disks are the entities that are replicated using the HP StorageWorks Enterprise Virtual Array snapshot functionality.
*See also* source volume and target volume.

**virtual full backup**

An efficient type of synthetic backup where data is consolidated using pointers instead of being copied. It is performed if all the backups (the full backup, incremental backups, and the resulting virtual full backup) are written to a single file library that uses distributed file medium format.

**Virtual 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**

*(VLS specific term)* An archival storage technology that backs up data to disk drives in the same way as if it were being stored on tape. Benefits of virtual tape systems include improved backup and recovery speed and lower operating costs.
*See also* Virtual Library System (VLS) and Virtual Tape Library.

**Virtual Tape Library (VTL)**

*(VLS specific term)* An emulated tape library that provides the functionality of traditional tape-based storage.
*See also* Virtual Library System (VLS).

**VMware management client**

*(VMware integration specific term)* The client that Data Protector uses to communicate with VMware Virtual Infrastructure. This can be a VirtualCenter Server system (VirtualCenter environment) or an ESX Server system (standalone ESX Server environment).

**volser**

*(ADIC and STK specific term)* A VOLume SERial number is a label on the medium to identify the physical tape used in very

large libraries. A volser is a naming convention specific to ADIC/GRAU and StorageTek devices.

**volume group**
A unit of data storage in an LVM system. A volume group can consist of one or more physical volumes. There can be more than one volume group on the system.

**volume mount point**
*(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**
*See* Microsoft Volume Shadow Copy Service.

**VSS compliant mode**
*(HP StorageWorks Disk Array XP VSS provider specific term)* One of two XP VSS hardware provider operation modes. When the XP provider is in the VSS compliant mode, the source volume (P-VOL) and its replica (S-VOL) are in simplex, unpaired state after a backup. Therefore the number of replicas (S-VOLs per a P-VOL) rotated is not limited. Restore from a backup in such a configuration is possible only by switching of the disks.
*See also* resync mode, source volume, primary volume (P-VOL), replica, secondary volume (S-VOL), and replica set rotation.

**VxFS**
Veritas Journal Filesystem.

**VxVM (Veritas Volume Manager)**
A Veritas Volume Manager is a system for managing disk space on Solaris platforms. A VxVM system consists of an arbitrary group of one or more physical volumes organized into logical disk groups.

**Wake ONLAN**
Remote power-up support for systems running in power-save mode from some other system on the same LAN.

**Web reporting**
The Data Protector functionality that allows you to view reports on backup, object copy, and object consolidation status and Data Protector configuration using the Web interface.

**wildcard character**
A keyboard character that can be used to represent one or many characters. The asterisk (*), for example, typically represents

one or more characters, and the question mark (?) typically represents a single character. Wildcard characters are often used in operating systems as a means of specifying more than one file by name.

**Windows CONFIGURATION backup**
Data Protector allows you to back up Windows CONFIGURATION, including Windows Registry, user profiles, Event Logs, and WINS and DHCP server data (if configured on a system) in one step.

**Windows Registry**
A centralized database used by Windows to store configuration information for the operating system and the installed applications.

**WINS server**
A system running Windows Internet Name Service software that resolves Windows networking computer names to IP addresses. Data Protector can back up WINS server data as part of the Windows configuration.

**writer**
*(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.

**XBSA interface**
*(Informix Server specific term)* ON-Bar and Data Protector communicate with each other through the X/Open Backup Services Application Programmer's Interface (XBSA).

**XCopy engine**
*(direct backup specific term)* A SCSI-3 copy command that allows you to copy data from a storage device having a SCSI source address to a backup device having a SCSI destination address, thus enabling direct backup. The data flows from a source device (either block or streaming, that is, disk or tape) to the destination device (either block or streaming) through XCopy. This releases the controlling server of reading the data from the storage device into memory and then writing the information to the destination device.
*See also* direct backup.

**ZDB**
*See* zero downtime backup (ZDB).

**ZDB database**
*(ZDB specific term)* A part of the IDB, storing ZDB related information such as source volumes, replicas and security

information. The ZDB database is used for ZDB, instant recovery, and split mirror restore.
*See also* zero downtime backup (ZDB).

**ZDB to disk**

*(ZDB specific term)* A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time. Multiple replicas, produced using the same backup specification at different times, can be kept in a replica set. A replica from ZDB to disk can be restored using the instant recovery process.
*See also* zero downtime backup (ZDB), ZDB to tape, ZDB to disk+tape, instant recovery, and replica set rotation.

**ZDB to disk+tape**

*(ZDB specific term)* A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time, in the same way as with ZDB to disk. However, data in the replica is also streamed to a backup medium, as with ZDB to tape. If this backup method is used, data backed up in the same session can be restored using the instant recovery process, the standard Data Protector restore from tape, or on split mirror arrays, split mirror restore.
*See also* zero downtime backup (ZDB), ZDB to disk, ZDB to tape, instant recovery, replica, and replica set rotation.

**ZDB to tape**

*(ZDB specific term)* A form of zero downtime backup where data in the replica produced is streamed to a backup medium, typically tape. Instant recovery is not possible from such a backup, so the replica need not be retained on the disk array after backup completion. The backed-up data can be restored using standard Data Protector restore from tape. On split mirror arrays, split mirror restore can also be used.
*See also* zero downtime backup (ZDB), ZDB to disk, instant recovery, ZDB to disk+tape, and replica.

**zero downtime backup (ZDB)**

A backup approach in which data replication techniques, provided by a disk array, are used to minimize the impact of backup operations on an application system. A replica of the data to be backed up is created first. All subsequent backup operations are performed on the replicated data rather than the original data, while the application system can be returned to normal operation.
*See also* ZDB to disk, ZDB to tape, ZDB to disk+tape, and instant recovery.

# Index