

HP OpenView Storage Data Protector Zero Downtime Backup Integration Guide

for

**Oracle
SAP R/3**

Microsoft SQL Server

Microsoft Exchange Server

Microsoft Volume Shadow Copy Service

Manual Edition: February 2006 (build label 249)



Manufacturing Part Number: B6960-90014

Release A.06.00

© Copyright Hewlett-Packard Development Company, L.P.2006.

Legal Notices

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

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

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

UNIX® is a registered trademark of The Open Group.

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

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

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

ARM® is a registered trademark of ARM Limited.

1. Data Protector Oracle ZDB Integration

In This Chapter	2
Introduction	3
Backup and Restore Types	4
Integration Concept	9
Oracle Backup Set ZDB Concept	14
Backup Process for Oracle Backup Set ZDB	17
Oracle Proxy-Copy ZDB Concept	21
Backup Process—Oracle Proxy-Copy	24
Configuring the Integration	28
Prerequisites	28
Limitations	30
Before You Begin	30
Cluster-Aware Clients	34
Linking Oracle with the Data Protector Oracle Integration Media Management Library (MML) on UNIX	35
Configuring Oracle Users on UNIX	37
Configuring Oracle Databases	39
Checking the Configuration	49
Switching Between Oracle Backup Methods	51
Configuring an Oracle ZDB	53
Creating a Data Protector Oracle ZDB Backup Specification	53
Editing the Oracle RMAN Script	69
Backing Up an Oracle Database	75
Scheduling a Backup	76
Running an Interactive Backup	77
Restoring an Oracle Database	80
Prerequisites	82
Limitations	82
Restoring from Backup Media to the Application System on LAN	82
Restoring Oracle Using the Data Protector GUI	83
Duplicating an Oracle Database	93
Restore, Recovery, and Duplicate Options	96
Restoring Oracle Using RMAN	102
Restoring Oracle Using CLI	116
Restoring Using Another Device	117
Instant Recovery and Database Recovery	117
Using Oracle After Removing the Data Protector Oracle Integration on UNIX Systems	126

Contents

Removing the Data Protector Oracle Integration Link on HP-UX Systems	126
Removing the Data Protector Oracle Integration Link on Solaris	127
TBD Troubleshooting	129
Before You Begin	129
Verifying the Prerequisites (Oracle Side)	129
Verifying the Configuration	132
Verifying the Backup	133
Verifying Restore	134
Configuration and Backup Problems	135
Restore Problems	139

2. Data Protector SAP R/3 ZDB Integration

In This Chapter	144
Introduction	145
Backup and Restore Types	146
Prerequisites and Limitations	150
SAP R/3 Integration Concept	154
SAP R/3 ZDB Concept	161
Backup Process for SAP R/3	163
Data Protector SAP R/3 Configuration File	166
Setting, Retrieving, Listing, and Deleting Data Protector SAP R/3 Configuration File Parameters Using the CLI	169
Configuring the Integration	173
Configuring an SAP R/3 Backup Owner in Data Protector on UNIX Systems	173
Configuring the Data Protector SAP R/3 Integration	174
Configuring the SAP R/3 Oracle Database on the Application System	174
Mounting the SAPBACKUP Directory on the Backup System	181
Configuring a Data Protector SAP R/3 Client on the Application System	182
Configuring an SAP R/3 ZDB	195
Creating a Data Protector SAP R/3 ZDB Backup Specification	195
SAP R/3 Specific Backup Options	207
Creating or Modifying the Parameter File on the SAP R/3 Database Server	211
Backing Up an SAP R/3 Database	213
Scheduling a Backup	214
Running an Interactive Backup	215
Restoring an SAP R/3 Database	218
Considerations	218
Finding Information Needed for Restore	219

Restoring Using the Data Protector GUI	219
Instant Recovery and Database Recovery	224
Troubleshooting	229
Before You Begin	229
General Troubleshooting	229
Verifying the Prerequisites (Oracle Side)	230
Verifying the Prerequisites (SAP R/3 Side)	231
Verifying the Configuration	232
Verifying the Backup	233
Verifying Restore	234
Configuration and Backup Problems	237
Examples of SAP R/3 Database Restore	240
Preparing the SAP R/3 Database for Restore	240
Example of Full Database Restore and Recovery	242
Example of Partial Restore	246
Example of Lost Files Restore	247
Example of Archive Log Files Restore	248

3. Data Protector Microsoft SQL Server ZDB Integration

In This Chapter	252
Introduction	253
Backup and Restore Types	253
Prerequisites and Limitations	256
Integration Concept	258
Data Protector Microsoft SQL Server Configuration File	261
Configuring the Integration	263
Configuring Microsoft SQL Server	263
Configuring a Microsoft SQL Server ZDB	267
Backing Up Microsoft SQL Server Databases	279
Scheduling a Backup	279
Running an Interactive Backup	281
Restoring a Microsoft SQL Server Database	283
Restoring from Backup Media to the Application System on LAN	283
Restore Options	287
Instant Recovery	290
Troubleshooting	293
Before You Begin	293
Configuration Problems	293

Contents

Backup Problems	296
Restore Problems	296
Before You Call Support.....	300

4. Data Protector Microsoft Exchange Server ZDB Integration

In This Chapter	302
Introduction	303
Backup and Restore Types.....	303
Prerequisites and Limitations.....	306
Integration Concept	310
Configuring a Microsoft Exchange ZDB Backup Specification	312
Creating a Microsoft Exchange ZDB Backup Specification	313
Modifying a Microsoft Exchange ZDB Backup Specification.....	328
Checking Microsoft Exchange Files for Consistency	329
Backing Up Microsoft Exchange Server.....	332
Scheduling a Backup	332
Running an Interactive Backup.....	334
Restoring a Microsoft Exchange Database.....	336
Restoring from the Backup Media to the Application System on LAN.....	336
Instant Recovery.....	346
Troubleshooting	350
Before You Begin	350
Backup Problems	350

5. Integrating the Data Protector ZDB Integrations and Microsoft Volume Shadow Copy Service

In This Chapter	354
Introduction	355
VSS Backup Types	357
VSS Restore.....	357
Prerequisites and Limitations.....	359
Integration Concepts	361
Backup.....	362
Restore.....	365
Configuring the Integration	367
Configuring the Data Protector VSS Cluster-Aware Integration	367
Microsoft Exchange Server 2003 Writer Specific Configuration	369
Configuration Check.....	369

Contents

Writers Specifics	371
Backing Up Writers Data	376
Creating Backup Specification Using GUI	376
Scheduling the Backup	385
Running an Interactive Backup	386
Restoring Writers Data	387
Restore Procedure	388
MSDE Writer Restore Specifics	391
Microsoft Exchange Server 2003 Writer Restore Specifics	392
Microsoft Data Protection Manager 2006 Writer Restore Specifics	395
Instant Recovery	397
Troubleshooting	400
Before You Begin	400
Backup Problems	400
Restore Problems	403
User Scenario For Microsoft Exchange Server 2003 Backup and Restore	405
Example - VSS Transportable Backup	405
Example Restore Scenario for Microsoft Exchange Server 2003	407

A. Appendix

In This Appendix	A-2
Reconfiguring an Oracle Instance for Instant Recovery	A-3
Examples for Moving the Control Files and Redo Logs to Different Locations	A-5
ZDB Integrations Omnirc Variables	A-9

Glossary

Index

Contents

Printing History

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

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

Table 1

Edition History

Part Number	Manual Edition	Product
B6960-90114	October 2004	Data Protector Release A.05.50
B6960-90013	April 2006	Data Protector Release A.06.00

Conventions


The following typographical conventions are used in this manual.

Table 2

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

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

Figure 1 Data Protector Graphical User Interface



Contact Information

General Information

General information about Data Protector can be found at

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

Technical Support

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

<http://support.openview.hp.com/support.jsp>

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

Information about the latest Data Protector patches can be found at

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

For information on the Data Protector required patches, refer to the *HP OpenView Storage Data Protector Product Announcements, Software Notes, and References*.

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

Documentation Feedback

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

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

Training Information

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

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

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

Data Protector Documentation

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

Manuals

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

HP OpenView Storage Data Protector Concepts Guide

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

HP OpenView Storage Data Protector Installation and Licensing Guide

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

HP OpenView Storage Data Protector Troubleshooting Guide

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

HP OpenView Storage Data Protector Disaster Recovery Guide

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

HP OpenView Storage Data Protector Integration Guide

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

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

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

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

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

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

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

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

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

HP OpenView Storage Data Protector Integration Guide for HP OpenView

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

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

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

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

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

There are two versions of the manual:

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

HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide

This manual describes Data Protector zero downtime backup and instant recovery concepts and provides background information on how Data Protector works in a zero downtime backup environment. It is intended to be used with the task-oriented *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide* and the *HP OpenView Storage Data Protector Zero Downtime Backup Integration Guide*.

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

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

HP OpenView Storage Data Protector Zero Downtime Backup Integration Guide

This manual describes how to configure and use Data Protector to perform zero downtime backup, instant recovery, and standard restore of Oracle, SAP R/3, Microsoft Exchange Server 2000/2003, and Microsoft

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

HP OpenView Storage Data Protector MPE/iX System User Guide

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

HP OpenView Storage Data Protector Media Operations User's Guide

This manual provides tracking and management of offline storage media. It is intended for network administrators responsible for maintaining and backing up systems. It describes the tasks of installing and configuring the application, performing daily media operations and producing reports.

HP OpenView Storage Data Protector Product Announcements, Software Notes, and References

This manual gives a description of new features of HP OpenView Storage Data Protector A.06.00. It also provides information on supported configurations (devices, platforms and online database integrations, SAN, and ZDB), required patches, and limitations, as well as known problems and workarounds. An updated version of the supported configurations is available at http://www.openview.hp.com/products/datapro/spec_0001.html

There are also four other *Product Announcements, Software Notes and References*, which serve a similar purpose for the following:

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

Online Help

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

Documentation Map

Abbreviations

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

Abbreviation	Manual
CLI	Command Line Interface Reference Guide
Concepts	Concepts Guide
DR	Disaster Recovery Guide
GS	Getting Started Guide
Help	Online Help
IG-IBM	Integration Guide—IBM Applications
IG-MS	Integration Guide—Microsoft Applications
IG-O/S	Integration Guide—Oracle & SAP
IG-OV	Integration Guide—HP OpenView Service Information Portal/OpenView Reporter
IG-OVOU	Integration Guide—HP OpenView Operations, UNIX
IG-OVOW	Integration Guide—HP OpenView Operations 7.1x, 7.2x, Windows
IG-OVOW	Integration Guide—HP OpenView Operations 7.5, Windows
IG-Var	Integration Guide—Sybase, Network Node Manager & NDMP
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

Abbreviation	Manual
Trouble	Troubleshooting Guide
ZDB Admin	ZDB Administrator's Guide
ZDB Concpt	ZDB Concepts Guide
ZDB IG	ZDB Integration Guide

Map

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

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

Integrations

Look in these manuals for details of the following integrations:

Integration	Guide
HP OpenView Operations (OVO)	IG-OVOU, IG-OVOW
HP OpenView Reporter (OVR)	IG-OV
HP OpenView Reporter Light	IG-OVOW
HP OpenView Service Information Portal (OVSIP)	IG-OV
HP StorageWorks Disk Array XP	all ZDB
HP StorageWorks Enterprise Virtual Array (EVA)	all ZDB
HP StorageWorks Virtual Array (VA)	all ZDB
IBM DB2 UDB	IG-IBM
Informix	IG-IBM
Lotus Notes/Domino	IG-IBM
Media Operations	MO User
MPE/iX System	MPE/iX
Microsoft Exchange Servers	IG-MS, ZDB IG
Microsoft Exchange Single Mailbox	IG-MS
Microsoft SQL Servers	IG-MS, ZDB IG
Microsoft Volume Shadow Copy Service (VSS)	IG-MS, ZDB IG
NDMP Server	IG-Var
Network Node Manager (NNM)	IG-Var
Oracle	IG-O/S
Oracle ZDB	ZDB IG
SAP DB	IG-O/S
SAP R/3	IG-O/S, ZDB IG
Sybase	IG-Var
Symmetrix (EMC)	all ZDB

In This Book

The *HP OpenView Storage Data Protector Zero Downtime Backup Integration Guide* describes how to configure and use Data Protector disk array integrations with other software products.

Audience

This manual is intended for backup administrators who are responsible for the planning, setup, and maintenance of network backups. It assumes that you are familiar with:

- Basic Data Protector functionality
- Database administration

Conceptual information can be found in the *HP OpenView Storage Data Protector Concepts Guide*, which is recommended to fully understand the fundamentals and the model of Data Protector.

It is also recommended to read the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide* for fundamentals of Data Protector integrations with disk arrays.

Organization

The manual is organized as follows:

- Chapter 1** “Data Protector Oracle ZDB Integration” on page 1
- Chapter 2** “Data Protector SAP R/3 ZDB Integration” on page 143
- Chapter 3** “Data Protector Microsoft SQL Server ZDB Integration” on page 251
- Chapter 5** “Data Protector Microsoft Exchange Server ZDB Integration” on page 301
- Chapter 6** “Integrating the Data Protector ZDB Integrations and Microsoft Volume Shadow Copy Service” on page 353
- Glossary** Definition of terms used in this manual.

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

- Microsoft SQL Server
- Microsoft Exchange Server
- Microsoft Volume Shadow Copy Service

The integrations of Data Protector with the following database applications are described in the *HP OpenView Storage Data Protector Integration Guide for Oracle and SAP*:

- Oracle
- SAP R/3
- SAP DB

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

- Informix Server
- IBM DB2 UDB
- Lotus Notes/Domino Server

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

- Sybase
- Network Node Manager
- Network Data Management Protocol

1 Data Protector Oracle ZDB Integration

In This Chapter

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

The chapter is organized into the following sections:

“Introduction” on page 3

“Integration Concept” on page 9

“Oracle Backup Set ZDB Concept” on page 14

“Oracle Proxy-Copy ZDB Concept” on page 21

“Configuring the Integration” on page 28

“Configuring an Oracle ZDB” on page 53

“Backing Up an Oracle Database” on page 75

“Restoring an Oracle Database” on page 80

“Using Oracle After Removing the Data Protector Oracle Integration on UNIX Systems” on page 126

“TBD Troubleshooting” on page 129

Introduction

You can employ a variety of backup strategies in order 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 in order 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 the following:

- 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 time to create a replica is significantly reduced. Following its creation, tape backup can be created 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.

NOTE

Backup of the recovery catalog and control file is not performed with ZDB to disk.

The Oracle Recovery Manager utility (RMAN) is not aware of ZDB-to-disk sessions.

IMPORTANT

Backup of archived logs cannot be done with the Data Protector Oracle ZDB integration. Backup of archived 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 OpenView Storage 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 a recovery from incremental backup (for ZDB to tape or ZDB to disk+tape)
 - ✓ perform a 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
 - ✓ restore a catalog before restoring the database

Table 1-1 on page 6 provides an overview of recovery methods, depending on the type of backup that was performed and type of recovery required.

Table 1-1 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

Table 1-1 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	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	<ul style="list-style-type: none"> • Restore or • Instant recovery+ database recovery 	<ul style="list-style-type: none"> • Restore or • Instant recovery+ database recovery 	Restore
	ZDB to disk+tape - offline	<ul style="list-style-type: none"> • Restore or • Instant recovery 	<ul style="list-style-type: none"> • Restore or • Instant recovery+ database recovery¹ 	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.

Introduction

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 OpenView Storage Data Protector Zero Downtime Backup Concepts Guide* for an overview of ZDB concepts and terminology.

Integration Concept

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

Components

The software components involved in backup and restore processes are:

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

Integration Functionality Overview

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

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

How Does the Integration Work?

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

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

The Data Protector Oracle Integration agent for Oracle gets the following information about Oracle backup objects either from the control files in the Oracle target database or from the Oracle recovery catalog:

- The physical schema of the Oracle target databases

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

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

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 data files that belong to an Oracle Server instance is known as a whole database backup.

These features can be used for online or offline backup of the Oracle target database. However, you must ensure that the backup objects (such as tablespaces) are switched into the appropriate state before and after a backup session. 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 a specific user. This user must be defined as the owner of the Data

Protector Oracle backup specification. Further on, `ob2rman.pl` prepares the environment to start the backup, and issues the RMAN backup command. RMAN instructs the Oracle Server processes to perform the specified command.

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

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

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

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

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

Restore Flow

A restore session can be started from the Data Protector GUI, or by issuing the RMAN restore command from the RMAN command line. 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.

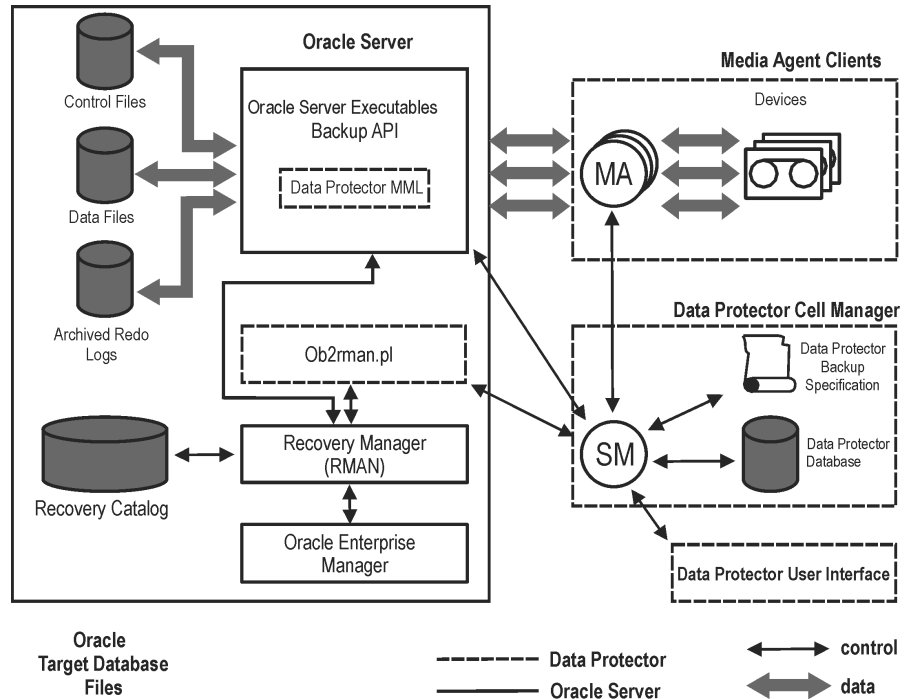
The Restore Session Manager checks the IDB to find the appropriate devices and media, starts the General Media Agent, establishes a connection between MML and the General Media Agent, and then monitors the restore and writes messages and information regarding the restore to the IDB.

Data Protector Oracle ZDB Integration
Integration Concept

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

The concept of Oracle integration, data and the control flow are shown in Figure 1-1 on page 12, and the related terms are explained in the following table.

Figure 1-1 Data Protector Oracle Integration Concept



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

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.

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

Oracle Backup Set ZDB Concept

See the *HP OpenView Storage 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 MML.

Depending on the location of the Oracle control file, online redo log files, and Oracle 9i/10g 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 A-9.

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. In such case, you must restore the overwritten files from a separate backup before you can perform database recovery beyond the replica creation time. Therefore, you may want to move the control files and redo logs to different locations. See “Reconfiguring an Oracle Instance for Instant Recovery” on page A-3 for more details.

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

Figure 1-2 Oracle Backup Set ZDB and Restore Concept

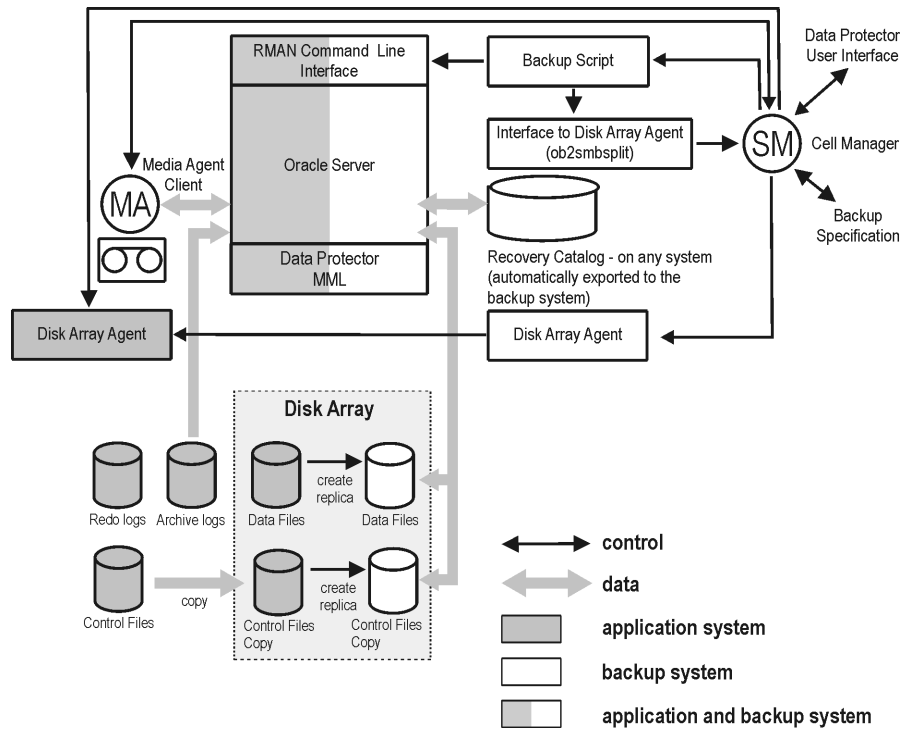


Figure 1-2 presents only the default integration behavior, where Oracle control file, online redo log files, and SPFILE reside on a different volume group (if LVM is used) or source volume than Oracle datafiles. Oracle 10g database files can also be part of ASM configuration, but note that only ZDB to tape is possible with ASM files.

For more information on an alternative Oracle backup and restore concept, see “ZDB Integrations Omnirc Variables” on page A-9.

Legend:

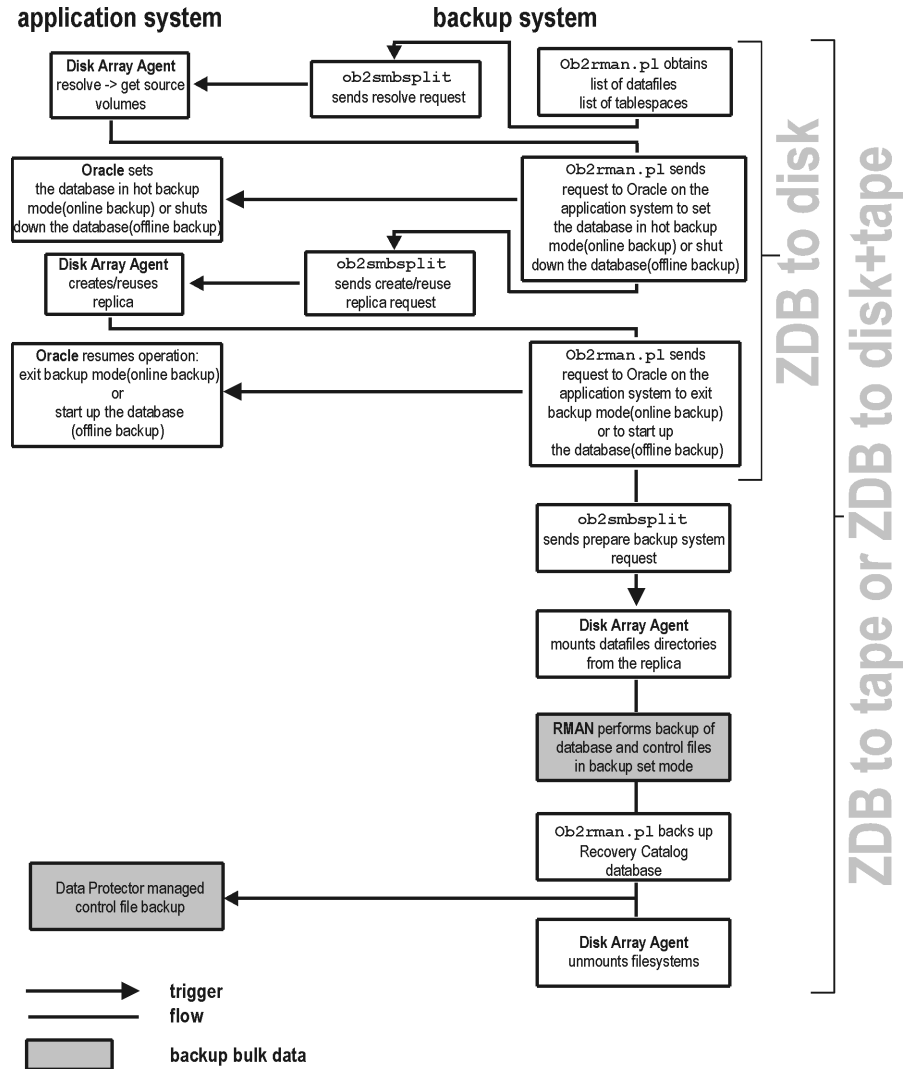
- MA** The General Media Agent writes data from a replica to backup media. The General Media Agent typically resides on the backup system.
- SM** The session manager controls backup and restore sessions and writes session information to the IDB.

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 for Oracle Backup Set ZDB

Figure 1-3 Oracle Backup Set ZDB Flow



NOTE

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

See the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide* for a general description of ZDB and instant recovery concepts.

See the *HP OpenView Storage 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/triggered by ZDB options. See the *HP OpenView Storage 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 disk array source volumes.
- 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 mountpoints 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 standard input of the RMAN command.
 - 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 Concept

See the *HP OpenView Storage 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.

The Oracle interface provides the Oracle Proxy Copy functionality which is utilized by Data Protector. 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 Oracle 9i/10g 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.

- 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.

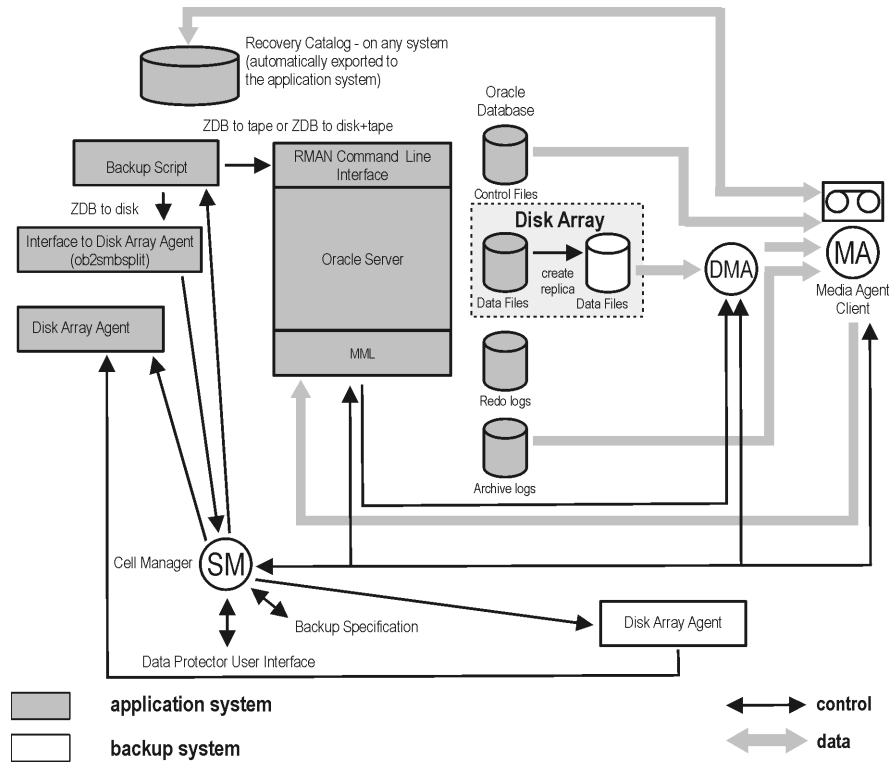
By default, instant recovery is *not* enabled. You can enable instant recovery by setting the `ZDB_ORA_INCLUDE_CF_OLF`, `ZDB_ORA_INCLUDE_SPF`, and `ZDB_ORA_NO_CHECKCONF_IR` omnirc variables to 1. See “ZDB Integrations Omnirc Variables” on page A-9.

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. In such case, you must restore the overwritten files from a separate backup before you can perform database recovery beyond the replica creation time. Therefore, you may want to move the control files and redo logs to different locations. See “Reconfiguring an Oracle Instance for Instant Recovery” on page A-3 for more details.

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

Figure 1-4 Oracle Proxy-copy ZDB and Restore Concept



NOTE

Figure 1-4 presents only the default integration behavior, where Oracle control file, online redo log files, and SPFILE reside on a different disk array source volume group than Oracle datafiles. For more information on alternative Oracle backup and restore concept, see “ZDB Integrations Omnirc Variables” on page A-9.

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

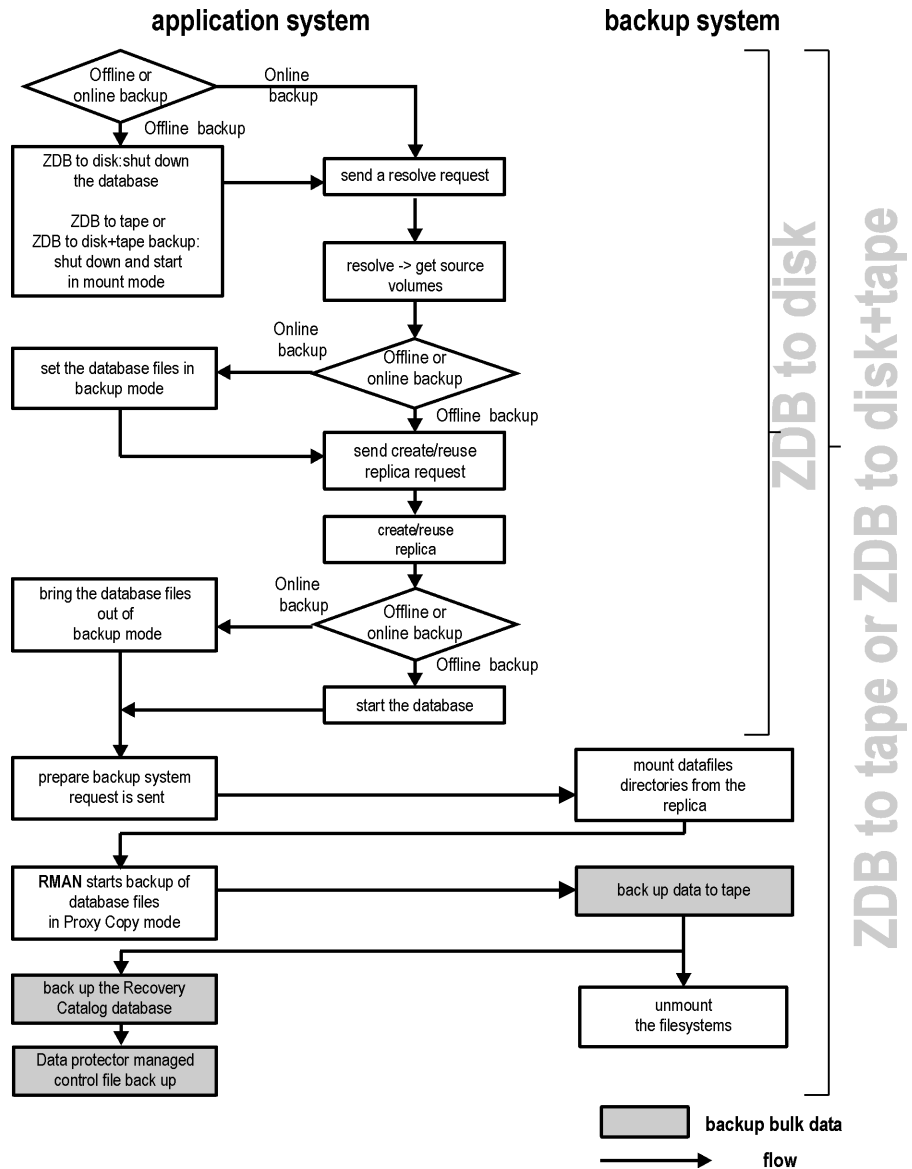
EMC), SSEA (on XP), and SNAPA (on VA).

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—Oracle Proxy-Copy

Figure 1-5 Oracle Proxy-Copy ZDB Flow



NOTE

The ZDB Agent are SYMA on EMC, SSEA on XP, SNAPA on VA, and either EVAA or SMISA on EVA.

See the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide* for a general description of ZDB and instant recovery concepts.

See the *HP OpenView Storage 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/triggered by ZDB options. See the *HP OpenView Storage 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 rawdisks to be backed up 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, `ob2rman.pl` makes a copy of the control file to the specified directory on the application system. This directory has to reside on disk array source volumes.

- In the case of an *online* ZDB, the Oracle target database instance files are switched into 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 backup mode after the replica has been created.

In the case of an *offline* backup, the Oracle database is started by `ob2rman.pl`.

- The `ob2rman.pl` command (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 disk images, the mountpoints containing the Oracle database files are mounted.

NOTE

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.

-
- A ZDB Agent then mounts the database on the backup system to a temporary directory on the application system.
 - 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

- 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 OpenView Storage Data Protector Installation and Licensing Guide* for information about licensing.
- Before you begin, ensure that you have correctly installed and configured the Oracle Server and Data Protector systems. See the:
 - *HP OpenView Storage Data Protector Product Announcements, Software Notes, and References* or http://www.openview.hp.com/products/datapro/spec_0001.html for an up-to-date list of supported versions, platforms, devices, and other information.
 - *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions on how to install Data Protector on various architectures and how to install 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.
- It is assumed that you are familiar with the Oracle database administration and the basic Data Protector functionality.
- A Data Protector ZDB integration (XP, VA, EVA, or EMC) must be correctly installed and configured. For installation, see the *HP OpenView Storage Data Protector Installation and Licensing Guide*. For configuration, see the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide*.

- The Oracle Server software must be installed and the Oracle target database must be open or mounted on the application system.
- 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.

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

See “TBD Troubleshooting” on page 129 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.

- The Oracle Net software must be correctly configured and running on the application system. The Oracle Net software is needed because during a backup, the Data Protector Oracle agent, has to be connected to the Oracle Database on the application system through Oracle TNS.
- On Windows, if the Oracle target database and the Oracle recovery catalog are installed on two different systems, the Data Protector Inet service account on the system with the Oracle target database installed must be configured as a *domain* account that is a member of the Administrators group on both systems. For information on how to change the Data Protector Inet service account, see the online Help index: “changing Data Protector Inet account”.
- 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 “Configuration and Backup Problems” on page 135.

Limitations

- 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 39.
- The Oracle database identifier (DBID) must be unique 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 data files 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 only) cannot be backed up using ZDB.
- The ASM files can be backed up using ZDB to tape only, because instant recovery is not supported for ASM.
- For backing up the ASM files, only backup set ZDB method is supported.
- **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.

To grant the SYSDBA privilege, start the sqlplus prompt and enter:

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

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

- ✓ If the Oracle database is installed on symbolic links, create these symbolic links on the backup system too.
- ✓ If an Oracle 10g **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 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. Enable Oracle automatic log archiving by setting:

```
log_archive_start=true  
  
log_archive_dest=<path_to_archive_logs>
```

in the `init<ORACLE_SID>.ora` file, where `<ORACLE_SID>` is the name of the Oracle database instance.

The default path of the file is:

Windows: <ORACLE_HOME>\database\init<ORACLE_SID>.ora

UNIX: <ORACLE_HOME>/dbs/init<ORACLE_SID>.ora

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:

— tnsnames.ora

Default path: <ORACLE_HOME>/network/admin/tnsnames.ora

— init<ORACLE_SID>.ora

Default path: <ORACLE_HOME>/dbs/init<ORACLE_SID>.ora.

— orapw<ORACLE_SID>

Default path: <ORACLE_HOME>/dbs/orapw<ORACLE_SID>

— 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`, `ORACLE_SID`, 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 Oracle8 database administrator with the `SYSDBA` role.

Example

If the `ORACLE_SID` of the target database is `PROD`, the `ORACLE_SID` 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 ORACLE_SID=PROD
oracle/PROD/bin/sqlplus
SQL> connect rman/rman@RMANCAT
Connected.
```

Data Protector Oracle ZDB Integration

Configuring the Integration

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

- Test whether both 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 `ORACLE_SID` of the target database is `PROD`, the `ORACLE_SID` 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 ORACLE_SID=PROD

rman target system/manager rcvcat rman/rman
Recovery Manager: Release 8.0.5.0.0 - Production
RMAN-06005: connected to target database: PROD
RMAN-06008: connected to recovery catalog database
RMAN> exit
Recovery Manager completed.
```

Cluster-Aware Clients

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

Windows: set `OB2BARHOSTNAME=<virtual_server_name>`

UNIX: export `OB2BARHOSTNAME=<virtual_server_name>`

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

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

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

IMPORTANT

After uninstalling the Data Protector Oracle integration on an Oracle server system, the Oracle server software is still linked to MML. You must re-link the Oracle binary to remove this link. If this is not done, the Oracle server cannot be started after the integration has been removed. See “Using Oracle After Removing the Data Protector Oracle Integration on UNIX Systems” on page 126 for information on removing the integration link.

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

On Oracle Server systems, MML is located in the `/opt/omni/lib` directory.

The filename for MML depends on the platform:

Table 1-2

Filenames for the MML on Different Platforms

Platforms	32-bit	64-bit
HP-UX	<code>libob2oracle8.sl</code>	<code>libob2oracle8_64bit.sl</code>
HP-UX on IA-64 architecture	<code>libob2oracle8.so</code>	<code>libob2oracle8_64bit.so</code>
Solaris	<code>libob2oracle8.so</code>	<code>libob2oracle8_64bit.so</code>

Data Protector Oracle ZDB Integration

Configuring the Integration

Proceed as follows:

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

32-bit Oracle: `cd <ORACLE_HOME>/lib`

64-bit Oracle 8i: `cd <ORACLE_HOME>/lib64`

64-bit Oracle 9i/10g: `cd <ORACLE_HOME>/lib`

2. Perform this step only if the `libobk.sl` (HP-UX) or `libobk.so` (Solaris) file is already created in the `<ORACLE_HOME>/lib` directory. Otherwise, skip this step.

Run:

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

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

IMPORTANT

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

3. Run:

HP-UX:

- 32-bit:

```
ln -s /opt/omni/lib/libob2oracle8.sl libobk.sl
```

- 64-bit:

```
ln -s /opt/omni/lib/libob2oracle8_64bit.sl libobk.sl
```

Solaris:

- 32-bit:

```
ln -s /optS/omni/lib/libob2oracle8.so libobk.so
```

- 64-bit:

```
ln -s /opt/omni/lib/libob2oracle8_64bit.so libobk.so
```

Configuring Oracle Users on UNIX

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

In addition, this operating system user must be registered in the Oracle database and identified by Oracle through the operating system identification.

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

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

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

As the owner of the backup specification, the Oracle user must be added to the Data Protector admin or operator user group.

You can identify this user by running the following command on the Oracle Server system:

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

Figure 1-6

Finding the Oracle User



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

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

Data Protector. For the Oracle backup set ZDB method, an Oracle user with the same user ID, group ID, and other privileges must be configured also on the backup system.

IMPORTANT

Additionally, the user `root` (UNIX only) on the Oracle Server has to be added to the Data Protector `admin` or `operator` user group for both, the application and backup system.

The user account for the added Oracle user must have the same numerical `userID` and `groupID` on the backup system and on the application system (for example, UNIX user `ora` and UNIX group `dba`). The Oracle user must be identical on both systems.

To check the `userID` and `groupID`, switch the user name:

```
#su - ora
```

and run the `id` command:

```
#id  
uid=101(ora) gid=101(dba)
```

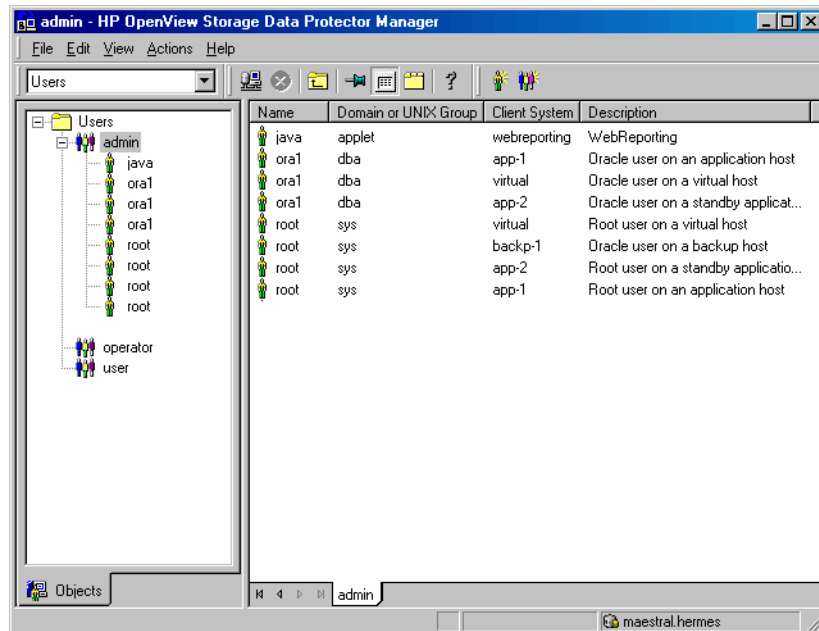
For information on how to add a user to a user group, see the online Help index: “adding users”.

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

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

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

Figure 1-7 Example User Configuration in a Cluster Environment



Configuring Oracle Databases

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

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

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

Using the Data Protector GUI

Configure an Oracle database when you create first ZDB backup specification for the database. Start with the procedure “Creating a Data Protector Oracle ZDB Backup Specification” on page 53 and at step 6 proceed as follows:

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

Figure 1-8 **Configuring Oracle - General (Windows)**

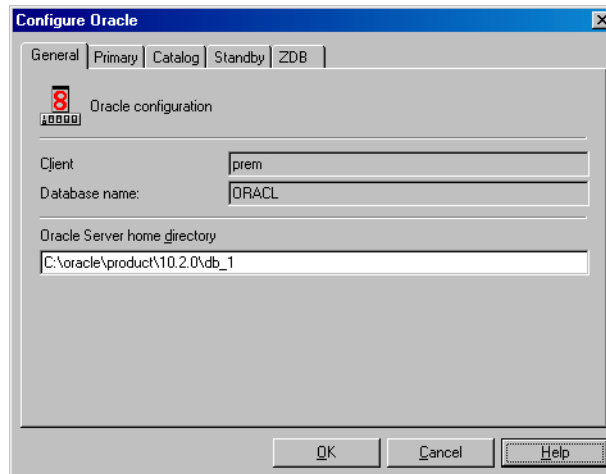
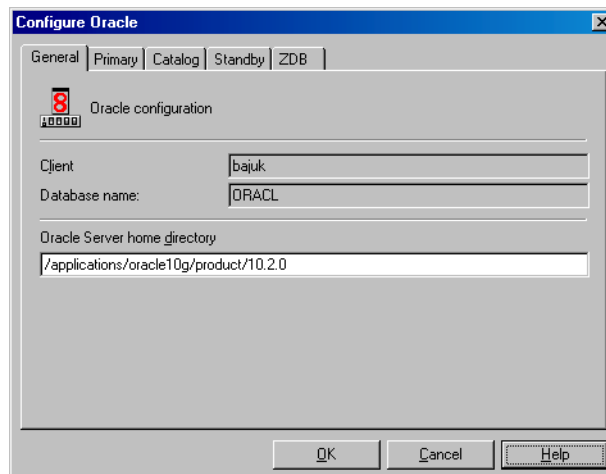


Figure 1-9 **Configuring Oracle - General (UNIX)**



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

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

Data Protector Oracle ZDB Integration

Configuring the Integration

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

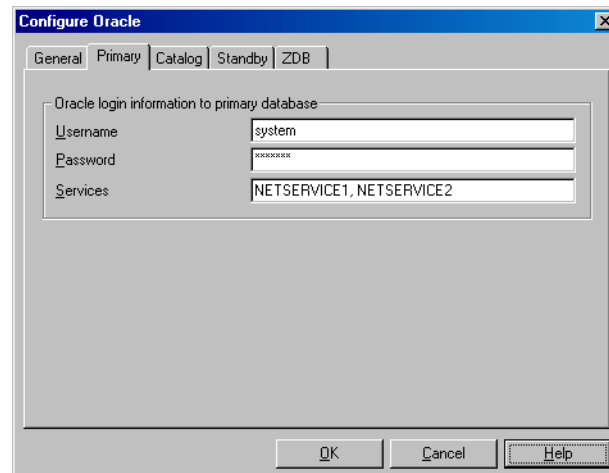
Note that each database instance for which the net service name is provided must be linked with the Data Protector MML. This means that the Data Protector Oracle integration software component must be installed on each system where the specific database instance is running.

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

NOTE

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

Figure 1-10 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.

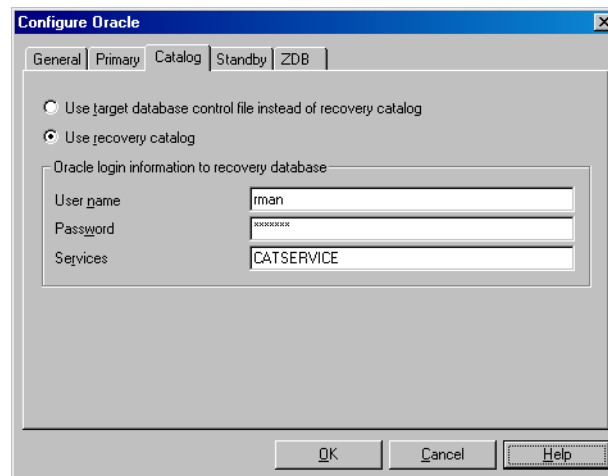
To use the recovery database catalog as an RMAN repository for backup history, select `Use recovery catalog` and specify the login information to the recovery catalog.

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.

Figure 1-11 **Configuring Oracle - Catalog**



- a. 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.

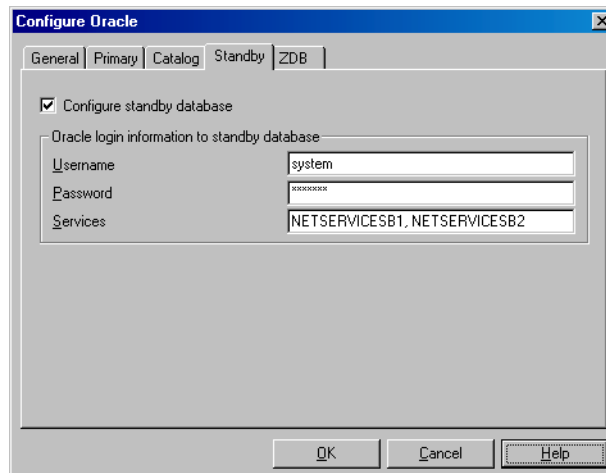
Note that each database instance for which the net service name is provided must be linked with the Data Protector MML. This means that the Data Protector Oracle integration software component must be installed on each system where the specific database instance is running.

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

NOTE

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

Figure 1-12 Configuring Oracle - Standby




4. In the ZDB property 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.

Figure 1-13 **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 1-62.

Using the Data Protector CLI

1. **UNIX only:** Log in to the Oracle Server system as user root or as the Oracle user that is identified as described in “Configuring Oracle Users on UNIX” on page 37.
2. On the Oracle Server system, from the directory:

Windows: <Data_Protector_home>\bin

HP-UX and Solaris: /opt/omni/1bin

run:

Data Protector Oracle ZDB Integration

Configuring the Integration

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

where:

*PRIMARY_DB_LOGIN*is:

```
-prouser <PRIMARY_USERNAME>  
-prpasswd <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_SET}  
[-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. In a cluster environment, name of the virtual server. This is optional for non-cluster environment and is used when the configuration is to be performed on different system than the one on which the `util_oracle8.pl` command is run.

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

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

<DB_NAME> Name of the database to be configured.

<ORACLE_HOME> Pathname of the Oracle Server home directory.

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

<primary_net_service_name_1>
[<primary_net_service_name_2>, ...] Net services names for the primary database.

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

<CATALOG_USERNAME> <CATALOG_PASSWORD> Username and password for login to the recovery catalog. This is optional and is used only if you are using the recovery catalog database catalog as an `RAMN` 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>

Data Protector Oracle ZDB Integration

Configuring the Integration

[*<standby_net_service_name_2>*, ...] Net services names for the standby database.

<BACKUP_CONTROL_FILE_COPY_LOCATION> The location on the source volumes where a copy of the current control file is made during 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.

If you use a raw logical volume as the *<BACKUP_CONTROL_FILE_COPY_LOCATION>*, the raw logical volume must reside on a volume group that will be replicated. If there is no such raw logical volume available, create a new shared disk (volume group) residing on the disk that will be replicated and configure a raw logical volume on it. If you use a raw logical volume, in case of an ZDB to disk, you need to ensure enough free space in the */var/opt/omni/tmp* directory on the backup host to hold the copy of the raw logical volume.

<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.

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: *orac1*
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 -prmsuser 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, TNS listener, or RMAN, these variables must be defined in the Environment section of the Data Protector Oracle global configuration file. See TBD for information on how to define them.

What Happens After the Configuration?

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

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

Checking the Configuration

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

Using the Data Protector GUI

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

IMPORTANT

On UNIX, it is possible that although the GUI check returns a successful result, you may still receive the error 12:8300 when trying to start a backup session. Such a backup session will not start. For more information, see “TBD Troubleshooting” on page 129.

Using the Data Protector CLI

1. **UNIX only:** Log in to the application system as the Oracle user or as user root.
2. From the directory:

Windows: <Data_Protector_home>\bin

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

run:

```
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:

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

HP-UX and Solaris: Run:

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

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 and Solaris: /opt/omni/lbin

run:

```
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 A-3 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 and you must restore the overwritten files from a separate backup before you can perform database recovery beyond the replica creation time. See “ZDB Integrations Omnirc Variables” on page A-9 on how to set the `omnirc` variables.

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 below 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<ORACLE_SID>_bckp.ora
```

UNIX: /etc/opt/omni/server/integ/config/Oracle8/
<client_name>%init<ORACLE_SID>_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 “Restore Problems” on page 139 on how to restore such a backup.

Configuring an Oracle ZDB

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 OpenView Storage 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 a Data Protector Oracle ZDB Backup Specification” on page 53.

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.

Creating a Data Protector Oracle ZDB Backup Specification

Online ZDB

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

You cannot perform a ZDB of the archived redo log files. Therefore, you need to 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

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

Creating an Oracle ZDB Backup Specification To create an Oracle ZDB backup specification, proceed as follows:

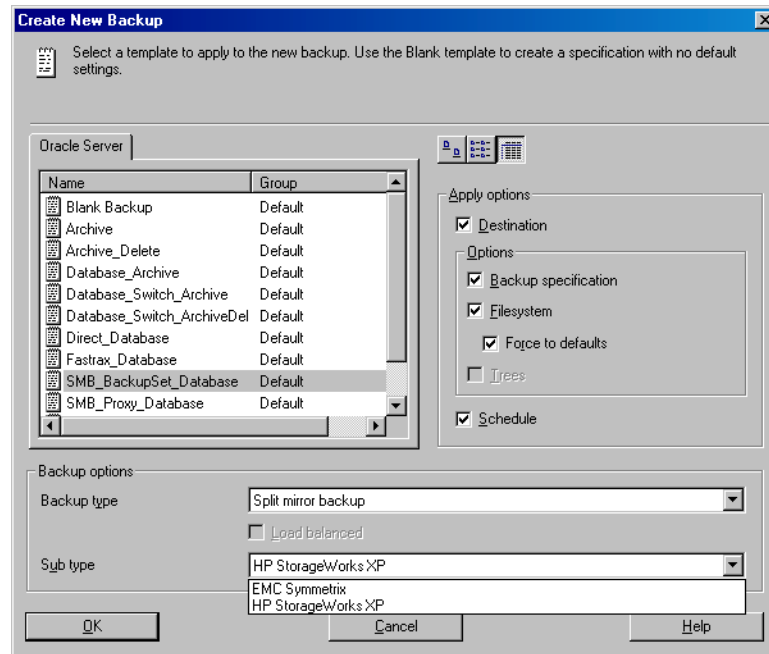
1. In the Context List, click Backup.
2. In the Scoping Pane, expand Backup Specifications, right-click Oracle Server, and click Add Backup.
3. In the Create New Backup dialog box, select the following:

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 1-14.

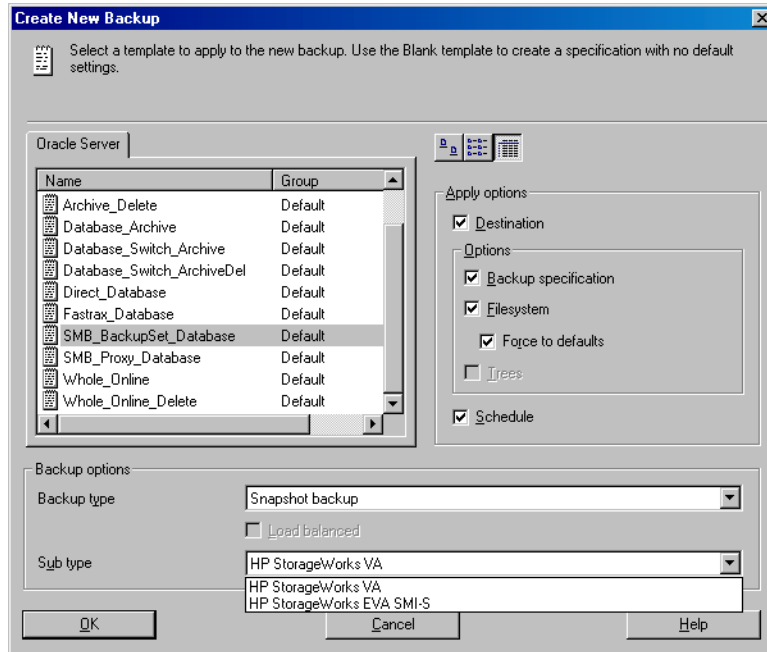
Figure 1-14 Selecting an Oracle ZDB Template and a Split Mirror Backup



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 1-15.

Figure 1-15 Selecting an Oracle ZDB Template and a Snapshot Backup



Click OK.

4. In the Application system drop-down list, select the Data Protector Oracle integration client. In a cluster environment, select the virtual server.

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

In the Backup system drop-down list, select the backup system.

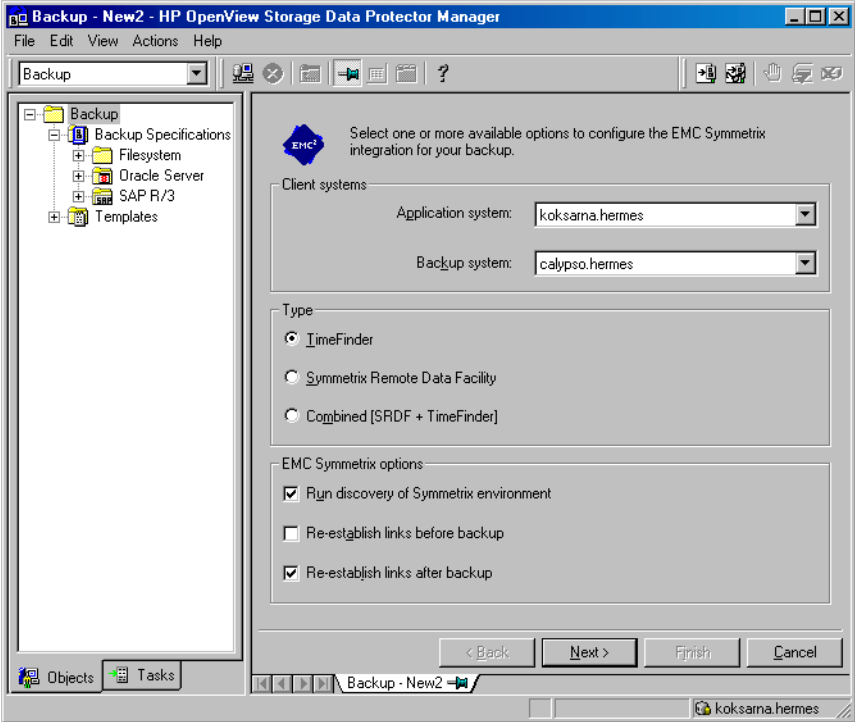
Select other disk array specific backup options. (see Figure 1-16 for EMC, Figure 1-17 for XP, Figure 1-18 for VA, or Figure 1-19 for EVA backup options). For detailed information on the options, press **F1**.

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.

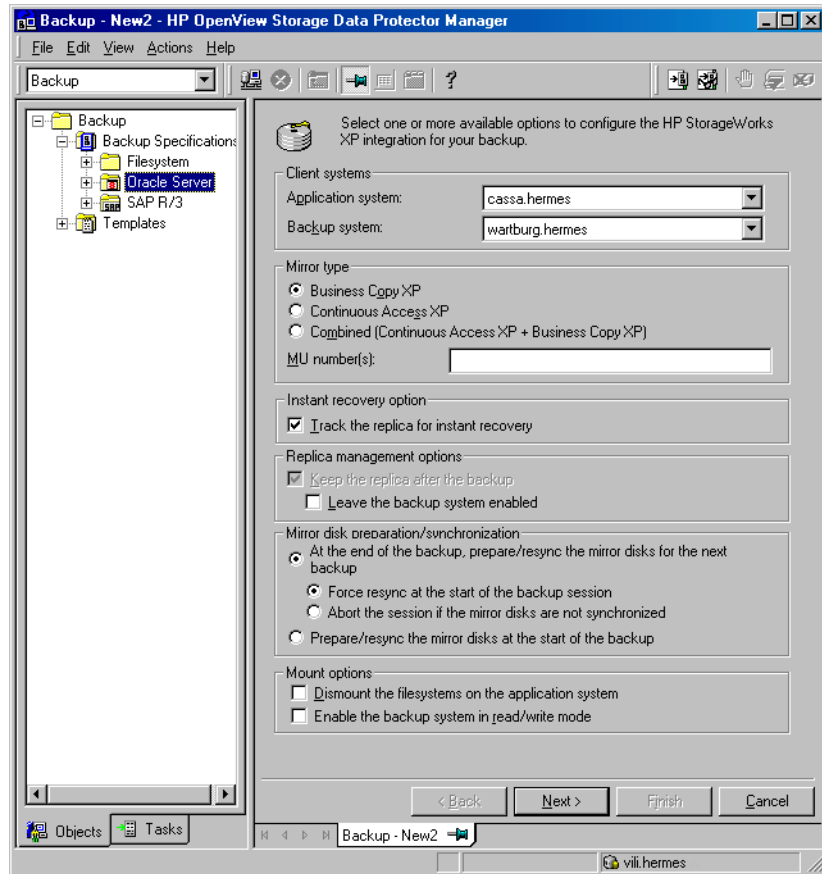
Figure 1-16 EMC Backup Options



On XP

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

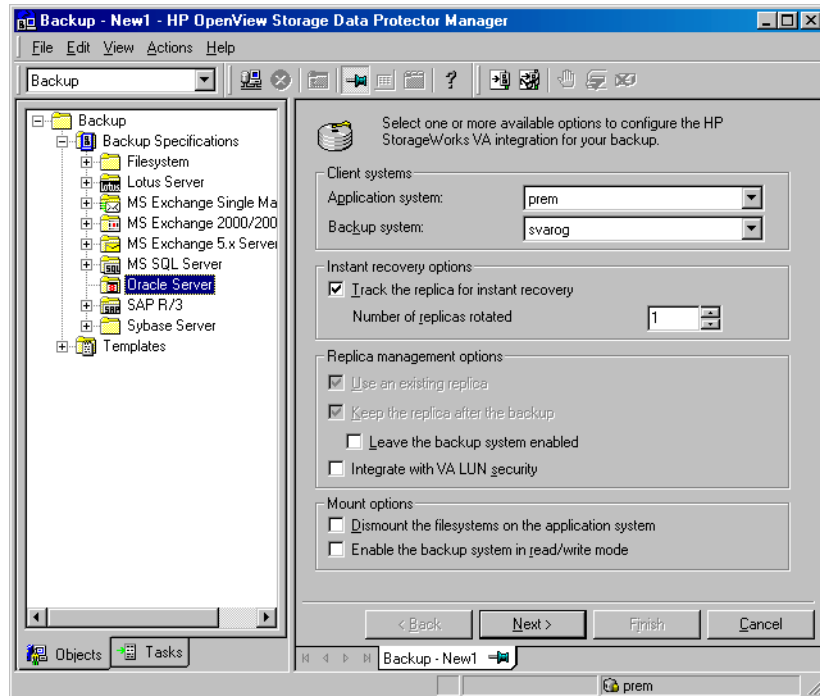
Figure 1-17 XP Backup Options



On VA

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

Figure 1-18 VA Backup Options

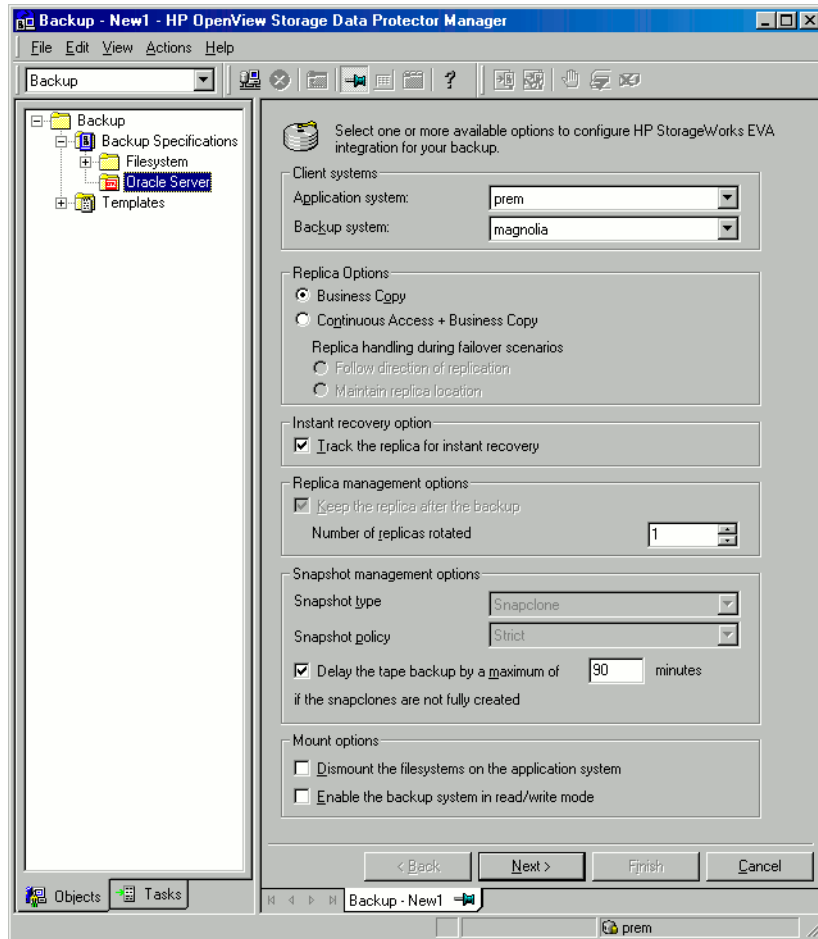


On EVA

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

Figure 1-19

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 as follows:

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

NOTE

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

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

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

UNIX only: Type the username and user group of the Oracle user. See “Configuring Oracle Users on UNIX” on page 37 for information on how to identify that user.

Figure 1-20 Specifying an Oracle Server System (Windows)

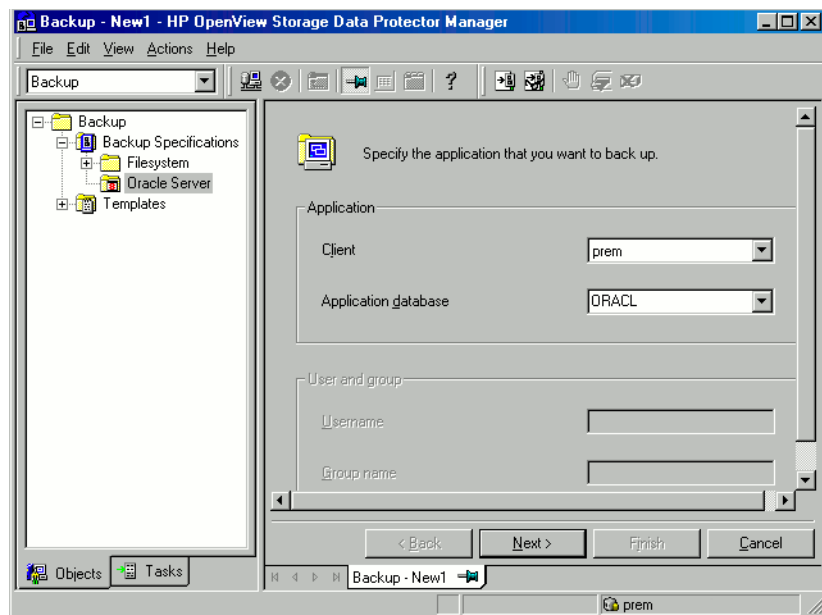
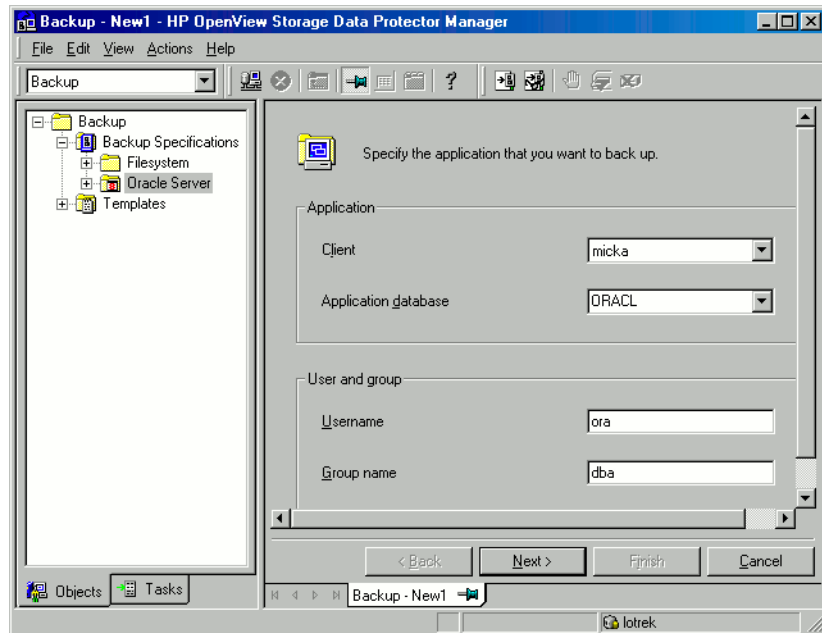


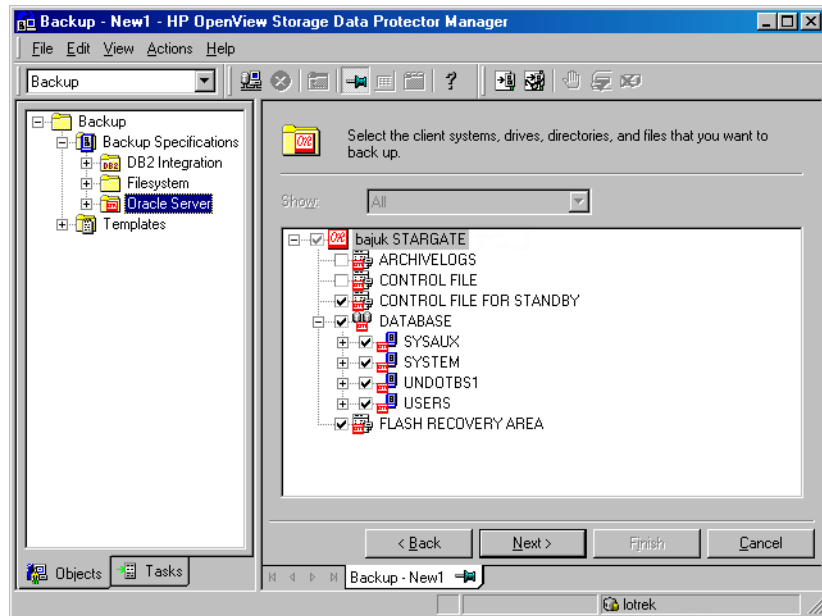
Figure 1-21 Specifying an Oracle Server System (UNIX)



Click Next.

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 39.
7. Select the Oracle database objects to be backed up.

Figure 1-22 **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.

Data Protector Oracle ZDB Integration Configuring an Oracle ZDB

Click Next to proceed.

9. Set the backup options.

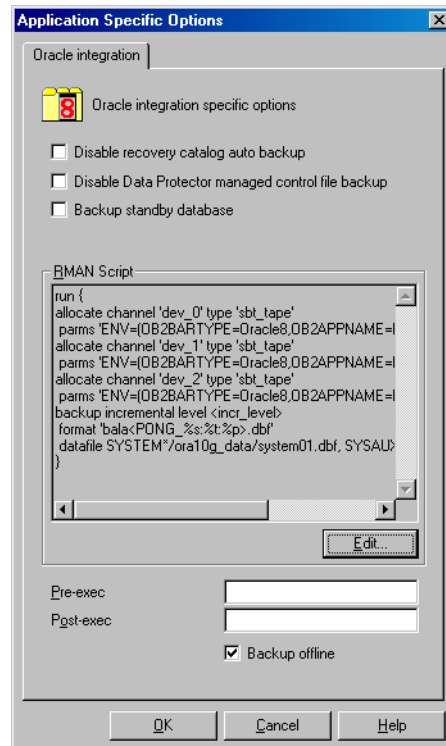
For information on the Backup Specification Options and Common Application Options, see the online Help.

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 1-23.

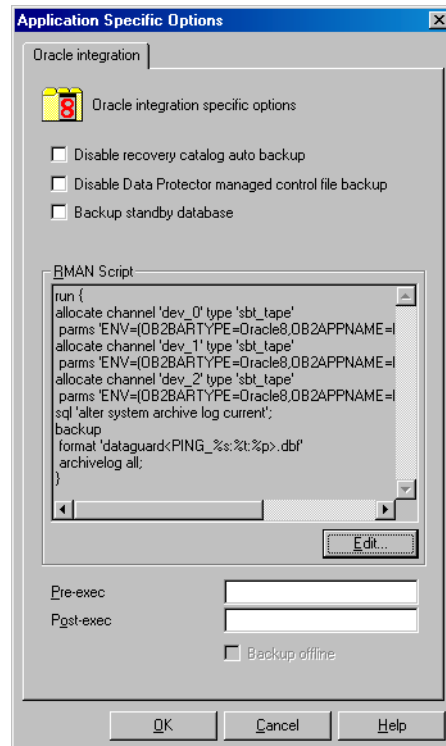
Figure 1-23

Backup Offline Option



For information on other Application Specific Options (Figure 1-24), see Table 1-3 on page 67 or online Help.

Figure 1-24 Oracle Specific Options



Click Next.

10. Optionally, schedule the backup. For more details, see “Scheduling a Backup” on page 76.

Note that only Full backup type is performed.


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. Oracle does not allow full stops in backup piece. Therefore, do not use a punctuation in the names of backup specifications, since the Oracle channel format is created from the backup specification name.

Figure 1-25 Saving the Backup Specification



Click OK.

To start the backup, see “Backing Up an Oracle Database” on page 75.

12. On UNIX, after the backup specification is saved, verify that the owner of the backup specification is the specified Oracle user. See “Configuring Oracle Users on UNIX” on page 37 for details about this user.

Online Backup

13. 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 OpenView Storage 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 1-3 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 69.
Pre-exec, Post-exec	Specify a command or RMAN script that will be started by <code>ob2rman.pl</code> on the Oracle server system before the backup (<code>pre-exec</code>) or after it (<code>post-exec</code>). Do not use double quotes. For example, you can provide scripts to shut-down and start an Oracle instance. For UNIX, see “Examples of Pre-Exec and Post-Exec Scripts on UNIX” on page 68. 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 1-23.

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 \"$ORACLE_SID\" shut down."
exit 0
else
echo "Cannot find Oracle SQLPLUS
($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 \"$ORACLE_SID\" started."
exit 0
else
echo "Cannot find Oracle SQLPLUS
```

```
($ORACLE_HOME/bin/sqlplus) ."  
exit 1  
fi
```

Editing the Oracle RMAN Script

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

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

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

Limitations

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

- The Oracle manual configuration convention must be used and not the Oracle automatic configuration convention (introduced by Oracle 9i).
- Double quotes (") must not be used - single quotes should be used instead.
- By default, RMAN scripts created by Data Protector contain instructions for backing up one or more of the following objects:
 - Databases, tablespaces, or data files (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.

IMPORTANT

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

To edit an Oracle RMAN script, click **Edit** in the **Application Specific Options** window (see Figure 1-24 on page 65), 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 (see also “Example of the RMAN Script” on page 73):

- **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 16 channels can be allocated. If the calculated number exceeds this limitation, you have to manually edit the RMAN script and reduce the number of allocated channels.

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


```
parms 'ENV=(OB2BARTYPE=Oracle8, OB2APPNAME=<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><<ORACLE_SID>_%s:%t:%p>.db
f'
```

NOTE

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

- 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.**

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

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

The backup statement consists of the following:

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

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

NOTE

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

- The RMAN archivelog all command.

If an appropriate template was selected, or if the statement was manually added, the RMAN statement to delete the Archived Redo Logs after they are backed up:

```
archivelog all delete input;
```

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

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

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

NOTE

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

— The RMAN current controlfile command.

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 RMAN Script

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

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

```
;
backup format 'New1<DIPSI_%s:%t:%p>.dbf' archivelog all;
backup format 'New1<DIPSI_%s:%t:%p>.dbf' current controlfile
;}
```

**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;
}
```

Backing Up an Oracle Database

To run a ZDB-to-disk, ZDB-to-tape, or ZDB-to-disk+tape 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 a Backup” on page 76.
- 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 77.

Considerations

Before running an Oracle ZDB session, note the following:

- It is not possible to start 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.
- 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, a warning message is issued in the Data Protector monitor during the backup, since the volume group source volumes on the application system do not have their BC pairs assigned. This warning message should be ignored.
- If the control file, SPFILE, 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 A-3 or to “ZDB Integrations Omnirc Variables” on page A-9.
- 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 a Backup

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

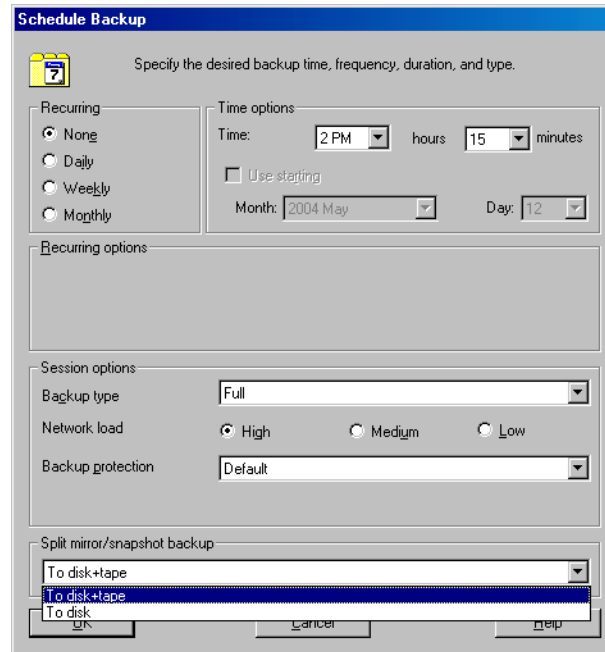
Note that the backup type is ignored for ZDB sessions. It is set to Full.

In the case of a ZDB-to-disk or a ZDB-to-disk+tape session, specify the Split mirror/snapshot backup option. See Figure 1-26.

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.

Figure 1-26 **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.

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

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


Note that the backup type is ignored for ZDB sessions. It 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. See Figure 1-27.

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.

Figure 1-27 **Selecting ZDB-to-Disk or ZDB-to-Disk+Tape Session When Starting an Interactive Backup**



Click OK.

An interactive backup can also be started from the CLI.

Starting a Backup Using the CLI

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

```
omnib -oracle8_list <Name>
```

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

```
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.

NOTE

It is not possible to run a 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.

Restoring an Oracle Database

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

- Control files
- Datafiles
- Tablespaces
- Databases

Using the Data Protector GUI, you can also **duplicate** a production database. See “Duplicating an Oracle Database” on page 93.

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

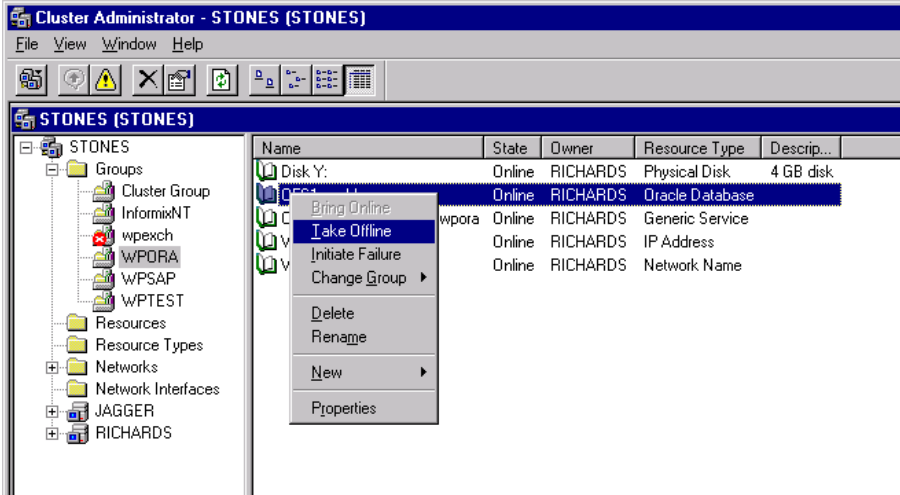
- Standard restore from backup media to the application system on LAN. See “Restoring from Backup Media to the Application System on LAN” on page 82.
- Instant recovery. See “Instant Recovery and Database Recovery” on page 117.

See also Table 1-1 on page 6 for an overview of recovery methods depending on the backup type and type of recovery.

MS Cluster Server Clients

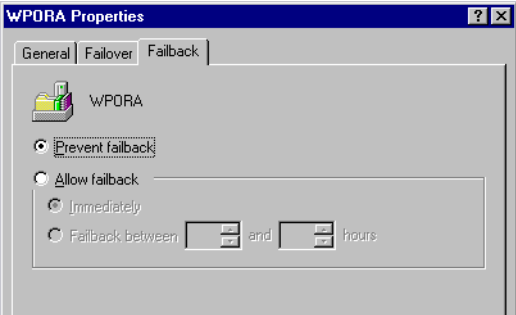
Before you start restoring a cluster-aware Oracle server, take the Oracle Database resource offline using, for example, the Cluster Administrator utility. See Figure 1-28.

Figure 1-28 Taking the Oracle Resource Group Offline



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

Figure 1-29 Checking Properties



MC/ServiceGuard Clients

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

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

Prerequisites

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

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.
- Objects backed up with *previous* versions of Data Protector or OmniBack II, using the proxy-copy ZDB method, will be restored with parallelism set to 1.

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 83.
- RMAN. See “Restoring Oracle Using RMAN” on page 102.

Restoring Oracle Using the Data Protector GUI

For restore, RMAN scripts are generated with necessary commands, depending on selections made in the GUI. If you want to perform additional actions, you cannot edit the RMAN restore script, but you can perform them manually from RMAN itself.

Restoring Database Items in a Disaster Recovery

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

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 object or you perform a duplication of a database, ensure that the database is in the correct state:

Table 1-4

Required Database States

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

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

1. Ensure the recovery catalog database exists and is empty. To check if the recovery catalog database was used as a repository during backup execute the following from the command line on the client system:

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

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

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

If the RCVLogin entry is present in the output, the recovery catalog database was used.

2. Identify the recovery catalog database owner and the instance name of the recovery catalog database using the Data Protector GUI.
3. Ensure that the recovery catalog database is in the Open state. Run the following command:

```
sqlplus /nolog
```

In the SQL> prompt, enter

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

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

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

- If the database is shut down, start it:

```
SQL>startup;
```
 - If the database is in nomount state, mount the database and then open it:

```
SQL>alter database mount;  
SQL>alter database open;
```
 - If the database is mounted, open it:


```
SQL>alter database open;
```
4. In the Data Protector GUI, switch to the `Restore` context. 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.

Figure 1-30

Recovery Catalog and Recovery Catalog Settings Dialog Box



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 96.

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.

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)

By default, the control file was backed up by `ob2rman.pl` at the end of a backup session. If the option `Disable Data Protector managed control file backup` was selected, the control file was not backed up.

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

The control files (`ctrl<ORACLE_SID>.dbf`) are restored in:

Windows: `<Data_Protector_home>\tmp`

HP-UX: `/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 83.
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 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 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 96 for information.
6. Click `Restore`.

Proceed with recovering the Oracle database or items within the database.

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 84 and “Restoring the Control File” on page 87.


To restore Oracle database objects:

1. Put the database in the mount state. See “Changing The Database State” on page 83.
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 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 96.

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 102.

Figure 1-31 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.

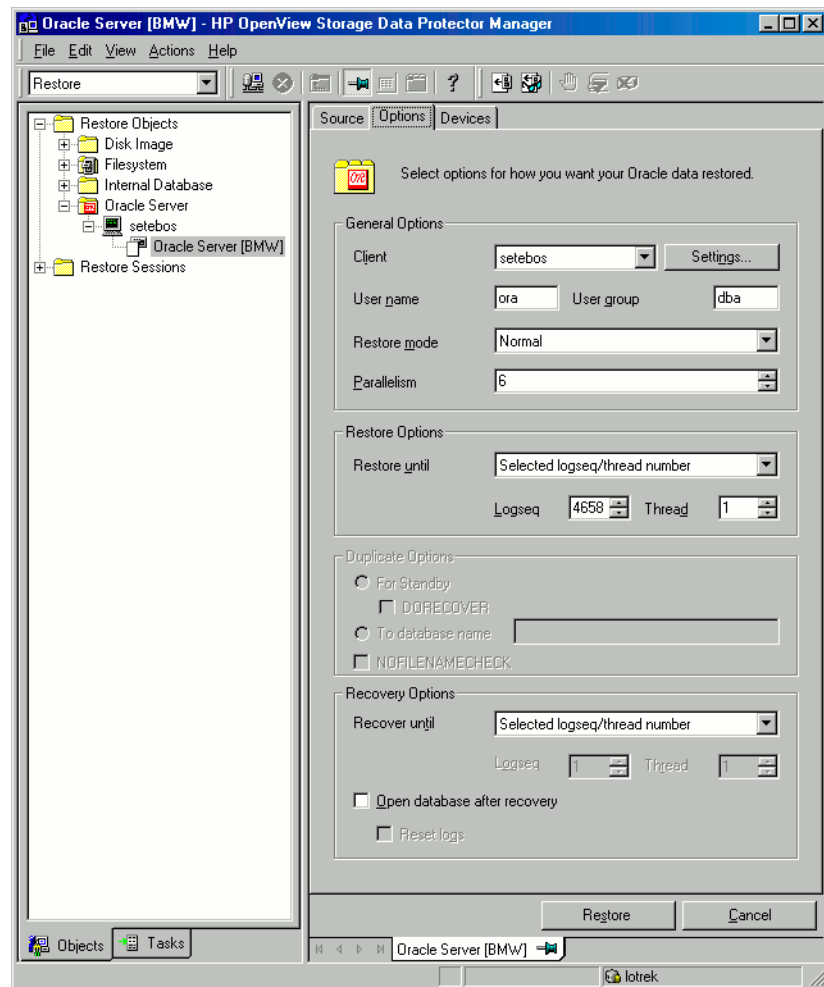
If you select to restore the datafile to another location or with a different name, the datafile will be restored to the selected location. If you want Oracle to use the datafile under the new name, you must issue a switch statement afterwards using Oracle tools. For more details, see the *Oracle Recovery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

Data Protector Oracle ZDB Integration
Restoring an Oracle Database

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 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 96 for information.

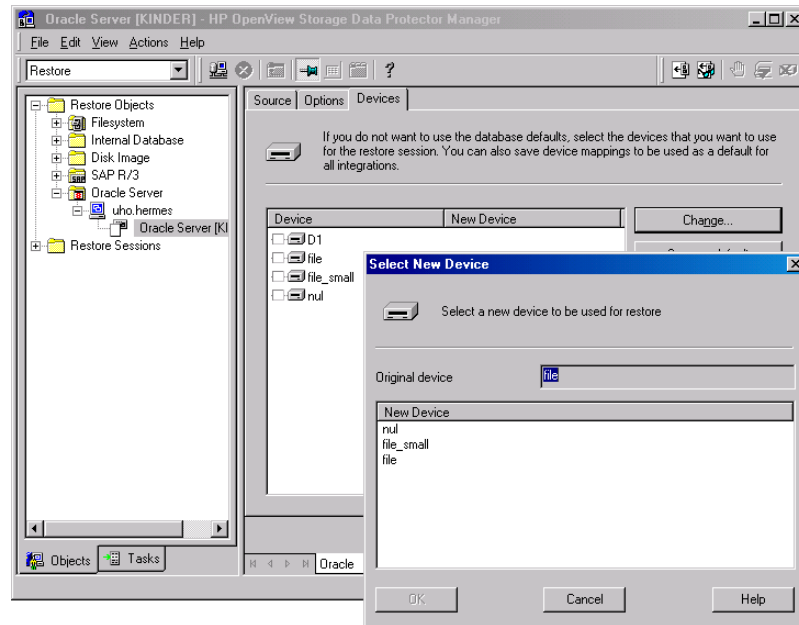
Figure 1-32 Options Page



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**.

Figure 1-33 **Devices Page**



8. Click `Restore`.

After the restore:

1. Put the database in the correct state.

If you used one of the options in the `Source` page containing the word “Recovery” 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 in order to register the new incarnation of database in the recovery catalog.

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

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

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> 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.

Limitation

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

Prerequisites

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

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 83.
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 (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 98 or press **F1**.

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

Figure 1-34 Oracle Duplicate Options



6. Click Restore.

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

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

Restore, Recovery, and Duplicate Options

Restore Action Options

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

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

Perform Restore

Use this option to only restore (but not recover) the database objects using Data Protector. After restore, recover the database manually using RMAN. For information on recovering the database using RMAN, see “Restoring Oracle Using RMAN” on page 102.

Perform Restore and Recovery

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

Perform Recovery Only

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

Perform RMAN Repository Restore

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

Perform Duplication

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

General Options

Client

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

Settings

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

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

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

User name (UNIX systems only)

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

User group (UNIX systems only)

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

NOTE

The user name and the user group must be the same as defined in the backup ownership. See “Configuring Oracle Users on UNIX” on page 37 for more information on this user and on how to identify it.

Restore mode

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

- Normal

This option should be used when a conventional backup or ZDB using the backup set method was performed with version of Data Protector older than A.05.00.

- Proxy copy

This option should be used when the original Oracle backup was made using the Oracle RMAN proxy-copy method, such as ZDB of Oracle 8i/9i using Data Protector version A.05.10.

This option is disabled when you are performing a restore after instant recovery.

Parallelism

This field is used to specify the number of concurrent data streams that can read from the backup device. If you do not enter a value, the number of parallel streams defaults to one.

To optimize restore performance, specify the same number of data streams as were used during the backup. 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.

Recovery Options

Restore until

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

This option is disabled when you are performing a restore after instant recovery.

- Now

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

- Selected time

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

- Selected logseq/thread number

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

- Selected SCN number

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

Recover until

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

- Now

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

- Selected time

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

- Selected logseq/thread number

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

- Selected SCN number

This option allows you to 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 the following:

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

```
RMAN> RESET DATABASE;
```

```
RMAN> exit
```

Open database after recovery

Opens the database after a recovery is performed.

Reset logs

Resets the archive logs after the database is opened.

Always reset the logs:

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

Do not reset the logs:

- After a complete recovery when the backup of a control file was not used in recovery.
- On the primary database, if the archive logs are used for a standby database. However, if you must reset the archive logs, you will need to recreate the standby database.

If you reset the logs when the `Restore until` option is set to `Now`, a warning is displayed, stating that you should reset the logs only if you use a backup of the control file for restore.

NOTE

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

Target DB login

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

Restoring Oracle Using RMAN

Data Protector acts as a media management 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” on page 105
- “Example of Point-in-Time Restore” on page 106
- “Example of Tablespace Restore” on page 108
- “Example of Datafile Restore” on page 110
- “Example of Archive Log Restore” on page 115

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 “Restoring an Oracle Database” on page 80 for details on how to restore the recovery catalog database.
- Check which ZDB method (proxy-copy or backup set) was used for backup that you plan the restore.
- If you are using control files, they must be accessible in order to perform the restore. If the control files are not available, you will probably need to restore them. 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 and the listener.

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

Example of Environment Variables on UNIX

```
ORACLE_BASE=/opt/oracle
ORACLE_HOME=/opt/oracle/product/8.1.6
ORACLE_TERM=hp
ORACLE_SID=PROD
PATH=$PATH:/opt/oracle/product/8.1.6/bin
NLS_LANG=american
```

**Example of
Environment
Variables on
Windows**

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

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

— On UNIX: `PROD:/opt/oracle/product/8.1.6:N`

— On Windows: `PROD:<Oracle_home>\product\8.1.6:N`

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

Connection Strings Used in the Examples

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

- Target connection string for target database:

```
sys/manager@PROD
```

where `sys` is the username, `manager` is the password and `PROD` is the name of the Oracle database.

- Recovery catalog connection string for recovery catalog database:

```
rman/rman@CATAL
```

where `rman` is the username and password and `CATAL` is the name of the Oracle database.

Example of Full Database Restore

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

1. Log in to the Oracle RMAN:

Run the following command:

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

2. Start the full database restore:

For a non-ZDB or ZDB backup set session:

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

For a ZDB proxy-copy session:

```
run{
allocate channel 'dev1' type 'sbt_tape' parms
'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.

Run the following command:

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

Example of Point-in-Time Restore

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

1. Log in to the Oracle RMAN:

Run the following command:

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

2. Start the point-in-time restore:

For a non-ZDB or ZDB backup set session:

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

```
release channel 'dev1';  
}
```

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  
  'ENV=(OB2BARTYPE=Oracle8,OB2PROXYCOPY=1,  
  OB2APPNAME=<DB_NAME>)' ;  
  allocate channel 'dev2' type 'sbt_tape' parms  
  'ENV=(OB2BARTYPE=Oracle8, OB2APPNAME=<DB_NAME>)' ;  
  set until time 'Mar 14 2004 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. Follow the steps below:

1. Create a file `restore_PIT` in the `/var/opt/omni/tmp` or `<Data_Protector_home>\tmp` directory.
2. Start the point-in-time restore.

Run the following command:

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

Example of Tablespace Restore

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

1. Log in to the Oracle RMAN:

Run the following command:

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

2. Start the tablespace restore.

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

For a non-ZDB or ZDB backup set session:

```
run{
allocate channel <dev1> type 'sbt_tape' parms
'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<DB_NAME>)';
sql 'alter tablespace "TEMP" offline immediate';
restore tablespace 'TEMP';
recover tablespace 'TEMP';
sql 'alter tablespace "TEMP" online';
release channel dev1;
}
```

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
'ENV=(OB2BARTYPE=Oracle8,OB2PROXYCOPY=1,
OB2APPNAME=<DB_NAME>)';
```

```
allocate channel 'dev2' type 'sbt_tape' parms
'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 the tablespace should have the following format:

For a non-ZDB or ZDB backup set session:

```
run{
allocate channel <dev1> type 'sbt_tape' parms
'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
'ENV=(OB2BARTYPE=Oracle8,OB2PROXYCOPY=1,
OB2APPNAME=<DB_NAME>)';
allocate channel 'dev2' type 'sbt_tape' parms
'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. The procedure in such cases is as follows:

1. Create a file `restore_TAB` in the `/var/opt/omni/tmp` (UNIX systems) or `<Data_Protector_home>\tmp` (Windows systems) directory.
2. Start the tablespace restore.

Run the following command:

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

Example of Datafile Restore

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

1. Log in to the Oracle RMAN:

Run the following command:

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

2. Start the datafile restore:

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

For a non-ZDB or ZDB backup set session:

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

UNIX


```
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
'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<DB_NAME>)';
allocate channel dev2 type 'sbt_tape' parms
'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<DB_NAME>,
OB2PORXYCOPY=1)';
sql "alter database datafile
'/opt/oracle/data/oradata/DATA/temp01.dbf' offline";
restore datafile
'/opt/oracle/data/oradata/DATA/temp01.dbf';
release channel dev1;
recover datafile
'/opt/oracle/data/oradata/DATA/temp01.dbf';
sql "alter database datafile
'/opt/oracle/data/oradata/DATA/temp01.dbf' online";
release channel dev2;
}
```

Windows

For a non-ZDB or ZDB backup set session:

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

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

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

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

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

release channel dev1;

}
```

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
'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<DB_NAME>)';

allocate channel dev2 type 'sbt_tape' parms
'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 dev1;

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

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

release channel dev2;

}
```

- If the database is in a mount 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
'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
'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<DB_NAME>)' ;
allocate channel dev2 type 'sbt_tape' parms
'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<DB_NAME>,
OB2PROXYCOPY=1)' ;
restore datafile
'/opt/oracle/data/oradata/DATA/temp01.dbf' ;
release channel dev1;
recover datafile
'/opt/oracle/data/oradata/DATA/temp01.dbf' ;
release channel dev2;
}
```

Windows

For a non-ZDB or ZDB backup set session:

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

Data Protector Oracle ZDB Integration

Restoring an Oracle Database

```
'<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  
'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<DB_NAME>)' ;  
allocate channel dev2 type 'sbt_tape' parms  
'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<DB_NAME>,  
OB2PROXYCOPY=1)' ;  
restore datafile  
'<Oracle_home>\data\oradata\DATA\temp01.dbf';  
release channel dev1;  
recover datafile  
'<Oracle_home>\data\oradata\DATA\temp01.dbf';  
release channel dev2;  
}
```

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

1. Create a file `restore_dbf` the `/var/opt/omni/tmp` or `<Data_Protector_home>\tmp` (Windows systems) directory.
2. Start the datafile restore.

Run the following command:

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

Example of Archive Log Restore

To restore an archive log, follow the steps below:

1. Login to the Oracle RMAN:

Run the following command:

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

2. Start the archive log restore:

```
run{
allocate channel dev1 type 'sbt_tape' parms
'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<DB_NAME>)';
restore archivelog all;
release channel dev1;
}
```

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

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

2. Start the archive log restore.

Run the following command:

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

Restoring Oracle Using CLI

Restoring the Recovery Catalog

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

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

To restore the Oracle recovery catalog, proceed as follows:

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

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

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

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

UNIX

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

```
OB2APPNAME="<DB_NAME>"  
export OB2APPNAME
```

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

```
setenv OB2APPNAME "<DB_NAME>"
```

Windows

```
set OB2APPNAME=<DB_NAME>
```

3. Run the following command:

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

- On Windows: `<Data_Protector_home>\bin\ob2rman.pl -restore_catalog -session <Session_ID> -apphost <application_hostname>`

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

Restoring Using Another Device

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

Specify these devices in the `/etc/opt/omni/server/cell/restoreddev` (UNIX systems) or

`<Data_Protector_home>\Config\server\Cell\restoreddev` (Windows systems) file in the following format:

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

where

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

On Windows, this file must be in UNICODE format.

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

Example

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

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

Instant Recovery and Database Recovery

See the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide* and *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide* for general information on instant recovery.

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 the instant recovery procedure 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 disk are restored.

IMPORTANT

If the Oracle control file, online redo logs, and the Oracle 9i/10g SPFILE are on the same source volumes as datafiles and you enable instant recovery by setting the `omnirc` variables, the control file, SPFILE, and online redo logs are overwritten during the instant recovery. As a consequence, you must restore the overwritten files from a separate backup before you can perform database recovery beyond the replica creation time.

Prerequisites

- The control file that reflects the internal database structure at the time of backup must be available on the application client. 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.
- On Windows, instant recovery is not possible from the replicas that were created with the Data Protector versions older than A.05.50.
- 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, edit the `omnirc` file and set the following variable:

```
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. The environment variable `OB2BARHOSTNAME` must be defined as the virtual hostname before running the configuration from the command line. The `OB2BARHOSTNAME` variable is set as follows:

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


3. Reconfigure the Oracle instance on the target node by running the following command:

```
/opt/omni/sbin/util_oracle8.pl -CONFIG_SMB_PROXY  
<ORACLE_GLOBAL_INSTANCE_SID> <ORACLE_HOME>  
<TARGET_PFILE_LOCATION> <CONTROL_FILE_LOCATION>  
<TARGET_DATABASE_LOGIN> <RECOVERY_CATALOG_LOGIN>
```

Note that all the parameters (with the exception of the `<ORACLE_GLOBAL_INSTANCE_SID>` and the `<RECOVERY_CATALOG_LOGIN>`) must be specified for the target node and not for the node that was backed up. The `<TARGET_DATABASE_LOGIN>` connection string must be specified for the Oracle instance running on the target node.

The backup method must be the same as on the original node, when the backup was made.

Instant Recovery Procedure

To perform an instant recovery, proceed as follows:

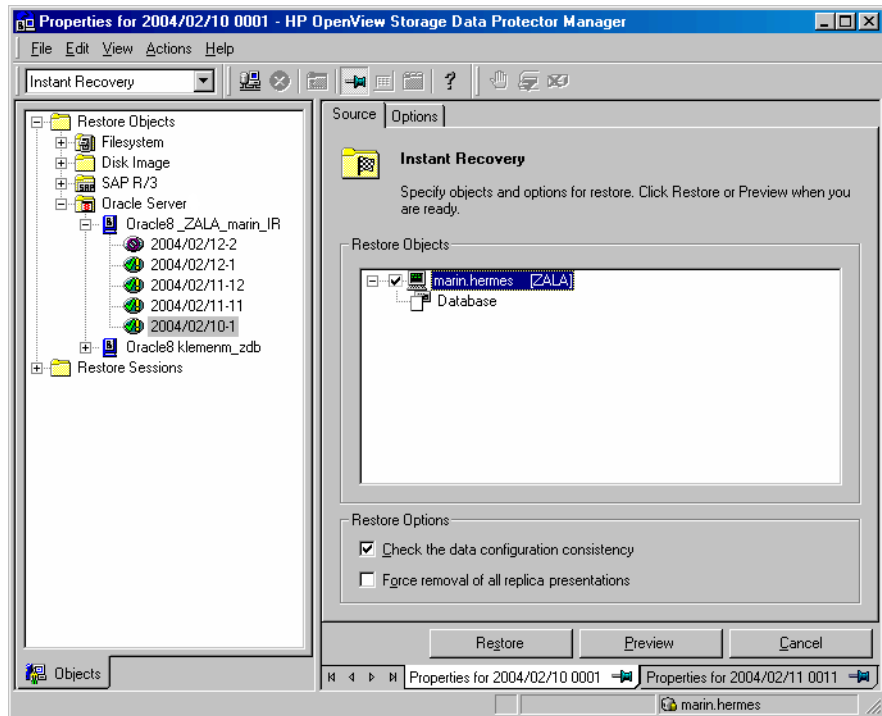
1. Shut down the Oracle database using sqlplus.

For example:

```
sqlplus  
  
sql> shutdown immediate  
  
sql> exit
```

2. In the Data Protector Manager Context List, select 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 StorageWorks Disk Array XP, it is recommended to set the Keep the replica after instant recovery option.

Figure 1-35 Selecting Backup Sessions (EVA example)



- 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 122.

- 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 1-36 and “Restore, Recovery, and Duplicate Options” on page 96 for details on available options.

Click **Restore**.

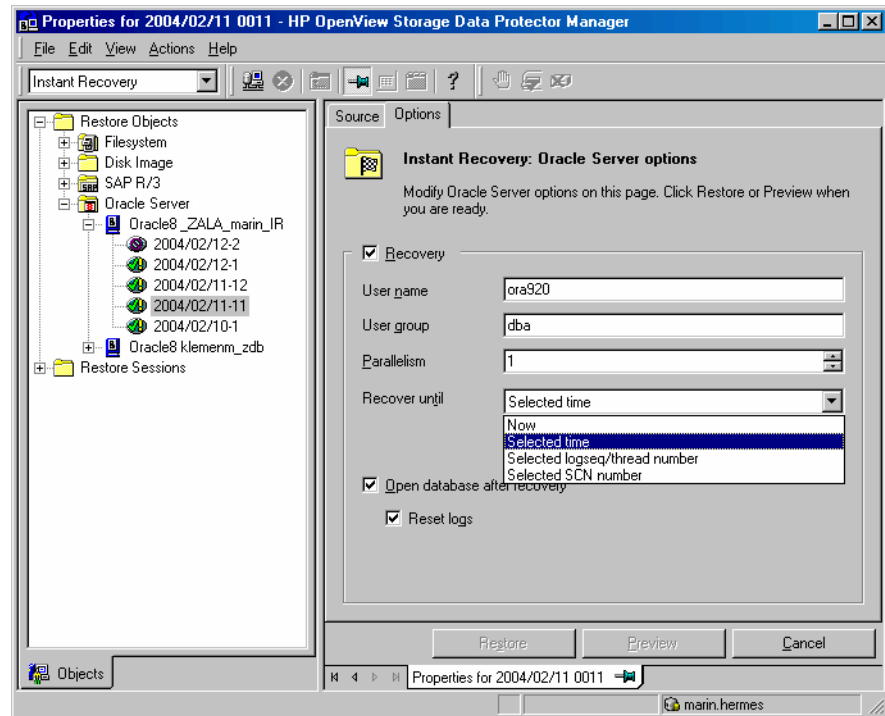
Data Protector recovers the database after performing instant recovery by switching the database to a mount state, restoring the necessary incremental backups and archived redo logs from tape, and applying the redo logs.

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 the following:

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

Figure 1-36

Oracle Recovery Options



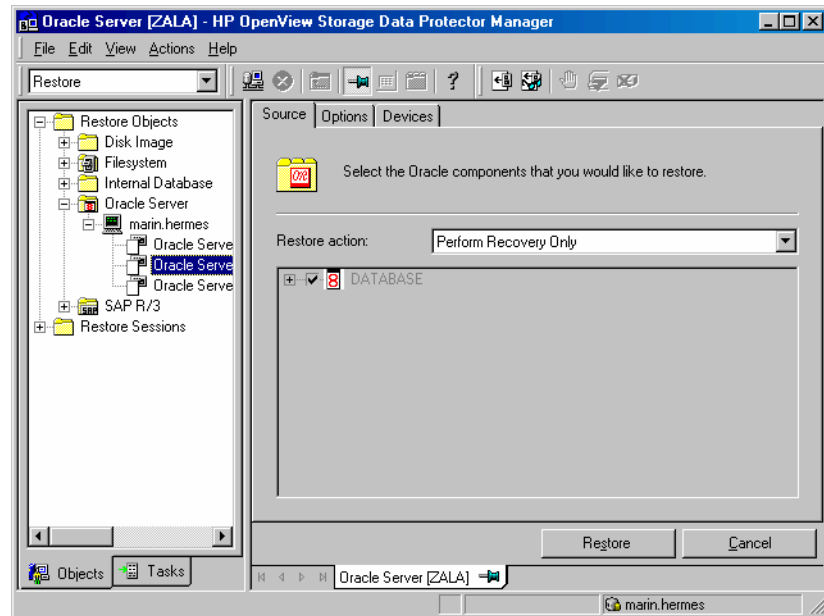
Oracle Database Recovery After the Instant Recovery

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.

Figure 1-37 Selecting the Database for Recovery



- b. In the Options tab, select the recovery options. For details, see “Restore, Recovery, and Duplicate Options” on page 96.
- 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  
        '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 102.

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 122 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.

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

After uninstalling the Data Protector Oracle integration on an Oracle server system, the Oracle server software is still linked to MML. You must rebuild (Oracle 8) or re-link (Oracle 8i/9i/10g) the Oracle binary to remove this link. If this is not done, the Oracle server cannot be started after the integration has been removed.

After you have uninstalled the Data Protector Oracle integration on the Oracle server system, proceed as described in the sections “Removing the Data Protector Oracle Integration Link on HP-UX Systems” on page 126 or “Removing the Data Protector Oracle Integration Link on Solaris” on page 127.

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

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

Oracle8

2. For Oracle 8, perform the following:

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

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

- b. Execute the following command:

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

IMPORTANT

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

Oracle 8i/9i/10g

3. For Oracle 8i/9i/10g, perform the following:
 - a. Change to the `<ORACLE_HOME>/lib` directory:

```
cd <ORACLE_HOME>/lib (32-bit Oracle),  
cd <ORACLE_HOME>/lib64 (64-bit Oracle 8i) or  
cd <ORACLE_HOME>/lib (64-bit Oracle 9i/10g).
```
 - b. If the `libobk.sl.orig` file exists in the `<ORACLE_HOME>/lib` directory, execute the following command:

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

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

Removing the Data Protector Oracle Integration Link on Solaris

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

Oracle 8

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

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

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

IMPORTANT

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

Oracle 8i/9i/10g

3. For Oracle 8i/9i/10g, perform the following:
 - a. Change to the `<ORACLE_HOME>/lib` directory:

```
cd <ORACLE_HOME>/lib (32-bit Oracle),  
cd <ORACLE_HOME>/lib64 (64-bit Oracle 8i) or
```

```
cd <ORACLE_HOME>/lib (64-bit Oracle 9i/10g).
```

- b. If the `libobk.so.orig` file exists in the `<ORACLE_HOME>/lib` directory, execute the following command:

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

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

4. Start all Oracle instances.

TBD 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.

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

For general ZDB, restore, and instant recovery related troubleshooting, see the troubleshooting sections in the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide*.

See also the troubleshooting section in the Oracle chapter of the *HP OpenView Storage 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 OpenView Storage Data Protector Product Announcements, Software Notes, and References* for general Data Protector limitations, as well as known problems and workarounds.
- ✓ See http://www.openview.hp.com/products/datapro/spec_0001.html for an up-to-date list of supported versions, platforms, and other information.

Verifying the Prerequisites (Oracle Side)

1. On the application system, verify that the target database is open as follows:

```
Export the <ORACLE_HOME> variable
```

```
Export the <ORACLE_SID> variable
```

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

```
select * from dba_tablespaces;
```

```
exit;
```

If it fails,

- Start or open the target database
2. From the application system, verify that the recovery catalog database is created and opened as follows:

Export the `<ORACLE_HOME>` variable

```
SQL > connect <login_information_to_recovery_catalog>
```

```
select * from rcver;
```

```
exit;
```

If it fails,

- Start or open the recovery catalog database
3. Verify that the Oracle Net software is correctly configured for the target database and the recovery catalog database in order to establish TNS network connections by doing the following:

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

See the Oracle documentation for information on how to create a TNS configuration file (LISTENER.ORA).

4. From the application system, verify that the target database and recovery catalog database are configured to allow remote connections with system privileges and to allow backup as follows:

- If you use the recovery catalog database:

Export the `<ORACLE_HOME>` variable

Export the `<ORACLE_SID>` variable

```
SQL> connect
```

```
<login_information_to_recovery_catalog_or_to_target  
database> as sysdba;
```

```
> exit
```

```
<ORACLE_HOME>/bin/rman target
```

```
<login_information_to_target database> rcvcat
```

```
<login_information_to_Recovery_Catalog>
```

- If you are not using the recovery catalog database:

Export the `<ORACLE_HOME>` variable

Export the `<ORACLE_SID>` variable

```
<ORACLE_HOME>/bin/rman target  
<login_information_to_target_database> nocatalog
```

See the Oracle documentation for how to set up a password file and parameters in the `init<ORACLE_SID>.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.

5. Verify that the target database is registered in the recovery catalog database as follows:

Export the `<ORACLE_HOME>` variable

```
<ORACLE_HOME>/bin/sqlplus> connect  
<Recovery_Catalog_Login>;  
> select * from rc_database;
```

Start the configuration from the Data Protector user interface for the application system, or see the Oracle documentation for information on how the target database is registered in the recovery catalog database.

6. On the application system, verify the backup and restore directly to disk using an RMAN channel type disk as follows:

If you use the recovery catalog database:

```
<ORACLE_HOME>/bin/rman target <Target_Database_Login>  
rcvcat <Recovery_Catalog_Login>  
  
run{  
  allocate channel dev1 type disk;  
  backup format <filename_for_backup> tablespace  
  <test_tablespace>;  
  sql 'alter tablespace <test_tablespace> offline';  
  restore tablespace <test_tablespace>;  
  recover tablespace <test_tablespace>;  
  sql 'alter tablespace <test_tablespace> online';  
}
```

If you are not using the recovery catalog database:

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

```
run{
allocate channel dev1 type disk;
backup format <filename_for_backup> tablespace
<test_tablespace>;
sql 'alter tablespace <test_tablespace> offline';
restore tablespace <test_tablespace>;
recover tablespace <test_tablespace>;
sql 'alter tablespace <test_tablespace> online';
}
```

Verifying the Configuration

All the verify steps from the “Verifying the Prerequisites (Oracle Side)” on page 129 must be performed before you begin this section. See the Oracle documentation for additional information on the actions that relate to Oracle.

1. Verify on the application system (Oracle proxy-copy ZDB method) or backup system (Oracle backup set ZDB) that the `libob2oracle8.sl` file is linked with the following Oracle executable:

```
/usr/bin/ldd -s <ORACLE_HOME>/bin/oracle (HP-UX 11.0 and 64 bit Oracle), or
```

```
/usr/bin/chatr <ORACLE_HOME>/bin/oracle (32 bit Oracle on HP-UX 11.0).
```

The `/usr/omni/lib/libob2oracle8.sl` Media Management Library (or the soft link to this library) must be listed as required by the Oracle executable, and the `SHLIB_PATH` parameter for dynamic path searching must be enabled.

Run a Data Protector filesystem backup of an Oracle client to check if the client and Cell Manager communicate properly.

2. Examine system errors reported in `/var/opt/omni/log/debug.log` on the application system (Oracle proxy-copy ZDB method) or backup system (Oracle backup set ZDB method).

If you have special Oracle environment settings, use the following command to register them:

```
/opt/omni/lbin/util_cmd -putopt Oracle8 <ORACLE_SID>
<VAR_NAME> <VAR_VALUE> -sublist Environment
```

Verifying the Backup

All steps from the “Verifying the Prerequisites (Oracle Side)” on page 129 and “Verifying the Configuration” on page 132 must be performed before you begin the steps in this section.

1. Login to the backup system (Oracle backup set ZDB method) or to the application system (Oracle proxy-copy ZDB method) as the user `root` or as the Oracle user that is identified as described in “Configuring Oracle Users on UNIX” on page 37. The identified Oracle user and the user `root` must also be added to the Data Protector admin or operator group. Then check the output of the command:

```
/opt/omni/lbin/util_oracle8.pl -CHKCONF_SMB -dbname  
<ORACLE_SID>
```

If an error occurs, use the following command to obtain an explanation:

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

where `<error_number>` is the number returned by the `*RETVAl*``<error_number>` line reported by `util_oracle8.pl`.

2. Verify that the user specified for the restore session is the Oracle backup owner and that they belong to the Data Protector operator or admin group. Check that the respective user group has the See Private Objects user rights.
3. Verify Data Protector backup using the `testbar2` utility.

Before you run `testbar2`, verify that the Cell Manager name is correctly defined on the Oracle client. Check the `/etc/opt/omni/client/cell_server` file, which contains the name of the Cell Manager system.

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

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

Examine errors reported by the `testbar2` utility by viewing the Data Protector troubleshooting file `/opt/omni/gui/help/Trouble.txt`.

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

- Check that the owner of the backup specification is the Oracle backup owner, and that this user belongs to the Data Protector operator or admin group.
- Check that the respective Data Protector user group has the `See private objects` user right enabled.
- Create an Oracle backup specification to back up to a null or file device. If the backup succeeds, the problem may be related to the backup devices. See the *HP OpenView Storage Data Protector Troubleshooting Guide* for instructions on troubleshooting devices.

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

Verifying Restore

All steps from “Verifying the Prerequisites (Oracle Side)” on page 129 and “Verifying the Configuration” on page 132 must be checked. See the Oracle documentation for additional details on actions related to Oracle.

Using `testbar2` utility, verify Data Protector restore.

Before you run `testbar2`, verify that the Cell Manager name is correctly defined on the Oracle client. Check the `/etc/opt/omni/client/cell_server` file, which contains the name of the Cell Manager system.

1. Test Data Protector’s internal data transfer using the `testbar2` utility:

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

2. Examine errors reported by the `testbar2` utility by viewing the Data Protector troubleshooting file `/opt/omni/gui/help/Trouble.txt`.

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

- a. Check if the owner of the restore session is the Oracle backup owner, and that this user belongs to the Data Protector operator or admin group.

- b. Check that the respective Data Protector user group has the `See private objects` user right enabled.
- c. As the owner of the restore session, run the `omnidb` command to see objects in the database.

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

Configuration and Backup Problems

The following list gives a description of problems and actions to be taken to resolve them.

- **SQL*Plus is unable to connect to destination.**

Check if the Oracle TNS 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 section “Data Protector Configuration Files” the *HP OpenView Storage Data Protector Integration Guide* for more information on setting the variables in the Data Protector Oracle configuration files.

- **The configuration procedure fails.**

Check whether the Oracle Server is up and running.

Check whether the Oracle Server with the recovery catalog database is up and running.

Check the login information to the target and the recovery catalog database on the application system. Use SQL*Plus. If you cannot start it, then:

1. Check if the `SYSDBA` roles are set to the Oracle administrator.
2. Examine system errors reported in:
 - ✓ UNIX systems: `/var/opt/omni/log/debug.log` on the application system.
 - ✓ Windows systems:
`<Data_Protector_home>\log\debug.log` on the application system.

3. If you have special Oracle environment settings, then ensure that they are entered in the Data Protector Oracle configuration files on the Cell Manager. See the *HP OpenView Storage Data Protector Integration Guide* for more information on setting the variables in the Data Protector Oracle configuration files.
- **When you start the backup, the configuration check fails.**

UNIX

Log in to the backup system (Oracle backup set ZDB method) or to the application system (Oracle proxy-copy ZDB method) as the user `root` or as the Oracle user that is identified as described in the section “Configuring an Oracle User in Data Protector” of the *HP OpenView Storage Data Protector Integration Guide*. The identified Oracle user and the user `root` must also be added to Data Protector `admin` or `operator` group. Then check the output of the command:

```
/opt/omni/sbin/util_oracle8.pl -CHKCONF_SMB -dbname  
<ORACLE_SID>
```

Windows

On the application system, check the output of the command:

```
<Data_Protector_home>\bin\util_oracle8.pl
```

- **Backup does not work.**

Test whether the backup works by starting a non-ZDB Oracle session on the application system.

Check whether the Cell Manager is correctly set on the backup system. The name of the Cell Manager has to be listed in the `/etc/opt/omni/client/cell_server` file (UNIX systems) or in the `HKEY_LOCAL_MACHINE\SOFTWARE\Hewlett-Packard\OpenView\OmniBackII\Site\CellServer` registry entry (Windows systems) on the Data Protector client.

Check whether the users are properly configured in user groups. The UNIX Oracle administrator from the application system has to be a member of the Data Protector operator class.

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

- **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.

In order to catalog the *controlfilecopy* to the RMAN catalog, follow the procedure below on the application system:

1. Connect to RMAN.
2. Run the following command:

```
RMAN> catalog controlfilecopy  
'<CONTROL_FILE_LOCATION>/ctrl<ORACLE_SID>.ctl'
```

- **Oracle related problems.**

When the following error:

```
ORA-12532: TNS: invalid argument.
```

is reported by SQL*Plus in the Data Protector monitor, the application system may be low on resources (CPU, memory, etc.).

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.

- **Backup Fails After a Point in Time Restore and Recovery**

Problem

Backup fails after a Point in time restore and recovery was performed and the following error is displayed:

```
RMAN-06004: ORACLE error from recovery catalog database:  
RMAN-20003: target database incarnation not found in  
recovery catalog
```

Action

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

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

```
RMAN> RESET DATABASE;
```

```
RMAN> exit
```

- **Backup Set ZDB is Aborted After 10 Minutes**

- Problem** 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 session 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 II or Data Protector:

```
SmWaitForFirstClient=<minutes>
```
 - If you performed a clean installation:

```
SmWaitForFirstBackupClient=<minutes>
```
- See the *HP OpenView Storage Data Protector Troubleshooting Guide* for more information on the global options file.
- **Backup of Archive Logs on RAC Cannot be Performed**

Problem On RAC, the archive logs are not installed on a NFS mounted disk. Backup of archive logs cannot be performed.

Action Edit the archive logs backup specification:

- Add an additional `allocate channel` command for *each* node.
- Replace the `<ORACLE_SID>` in the `OB2APPNAME` parameter with `<GLOBAL_DB_NAME>`.
- Add a command to connect to each instance. The connection parameters should be given as

```
<username>/<passwd>@<INSTANCE>
```

For example, if you are using two nodes, the backup specification might look as follows:

```
run {  
allocate channel 'dev_0' type 'sbt_tape'
```

```
parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<GLOBAL DB
NAME>,OB2BARLIST=RAC_arch)' connect
<username>/<passwd>@<INSTANCE 1>;

allocate channel 'dev_2' type 'sbt_tape'

parms 'ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<GLOBAL DB
NAME>,OB2BARLIST=RAC_arch)' connect
<username>/<passwd>@<INSTANCE 2>;

backup

format 'RAC_arch<QU_%s:%t:%p>.dbf'

archivelog all;

}
```

- **Backup Set ZDB Fails After Changing the Physical Schema of the Database**

Problem

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.

Restore Problems

- Proxy copy restore fails with the following error:

```
RMAN-10035: exception raised in RPC: ORA-27197: skgfprs:
sbtprcrestore returned error
```

```
RMAN-10031: ORA-27197 occurred during call to
DBMS_BACKUP_RESTORE.PROXYRESTOREDATAFILE
```

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> rcvcat  
<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.

- **Restore After a Switch Between ZDB Methods Fails**

Problem

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 1-38 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 {
  ALLOCATE CHANNEL 'dev_0' TYPE 'sbt_tape'
  PARAMS ENV(OB2BARTYPE=Oracle8,OB2APPNAME=<ORACLE_SID>);
  ALLOCATE CHANNEL 'dev_1' TYPE 'sbt_tape'
  PARAMS ENV(OB2BARTYPE=Oracle8,OB2APPNAME=<ORACLE_SID>,
  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.

TBD Troubleshooting

<list of backup_set backup> 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>*.



2

**Data Protector SAP R/3 ZDB
Integration**

In This Chapter

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

The chapter is organized into the following sections:

“Introduction” on page 145

“Prerequisites and Limitations” on page 150

“SAP R/3 Integration Concept” on page 154

“SAP R/3 ZDB Concept” on page 161

“Data Protector SAP R/3 Configuration File” on page 166

“Configuring the Integration” on page 173

“Configuring an SAP R/3 ZDB” on page 195

“Backing Up an SAP R/3 Database” on page 213

“Restoring an SAP R/3 Database” on page 218

“Troubleshooting” on page 229

“Examples of SAP R/3 Database Restore” on page 240

Introduction

You can employ a variety of backup strategies in order 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 in order 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 SAP R/3:

- 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 SAP R/3 ZDB integration are the following:

- 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 Data Protector SAP R/3 ZDB integration offers online and offline backup of your SAP R/3 Database 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.

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.

On EMC

You can perform ZDB (ZDB to tape only) on the following mirror types of the EMC configurations:

- SRDF
- Time Finder
- Combined SRDF + Time Finder

On XP

You can perform ZDB on the following mirror types of the XP configurations:

- BC
- CA
- Combined BC+CA

Restore

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

- Restoring from backup media to the application system on LAN (standard Data Protector restore) and using the SAP R/3 restore utility (SAPDBA) 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 instant recovery and database recovery or instant recovery and the SAP R/3 restore utility (SAPDBA) you can:
 - ✓ recover a whole database
 - ✓ recover a whole database as it was at a specific point in time

NOTE

It is not possible to perform an instant recovery from data backed up using a Data Protector SAP R/3 offline backup a session.

Table 2-1 provides an overview of recovery methods, depending on the type of backup that was performed and type of recovery required:

Table 2-1 SAP R/3 Recovery Methods

Disk arrays	Backup type	Recover a whole database		Recover part of a 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

Table 2-1 SAP R/3 Recovery Methods

Disk arrays	Backup type	Recover a whole database		Recover part of a database
		Until now	To a point in time, logseq/thread or SCN number	
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	<ul style="list-style-type: none"> • Instant recovery + database recovery or • Restore 	<ul style="list-style-type: none"> • Instant recovery + database recovery or • Restore 	Restore
	ZDB to disk+tape - offline	Restore	Restore	Restore

Legend:

Restore

The following two options are possible:

- Use SAPDBA to restore from backup media to the application system on LAN or
- Use the Data Protector GUI to restore the database from backup media to the application system on LAN and then use SAPDBA to recover the database.

“Instant recovery + database recovery” The following two options are possible:

- Perform instant recovery with a database recovery
or
- Perform an instant recovery first and then use SAPDBA to recover the database.

Prerequisites and Limitations

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

Prerequisites

- The database used by SAP R/3 must be an Oracle database. If any other database is used by SAP, then the corresponding Data Protector integration of that database must be used instead.
- You need a license to use the Data Protector SAP R/3 ZDB integration. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for information about licensing.
- Refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide* for an overview of ZDB concepts and terminology.
- Before you begin, make sure that you have correctly installed and configured the SAP R/3 Database Server and Data Protector systems. Refer to the:
 - *HP OpenView Storage Data Protector Product Announcements, Software Notes, and References* or http://www.openview.hp.com/products/datapro/spec_0001.html for an up-to-date list of supported versions, platforms, devices, and other information.
 - *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions on how to install Data Protector on various architectures and how to install a Data Protector disk array integration (EMC, XP, VA, or EVA) with SAP R/3.
 - *SAP R/3 System Online Documentation* for instructions on how to install and configure the SAP R/3 database and the SAP R/3 backup and restore tools (BRBACKUP, BRRESTORE, and BRARCHIVE).
- A Data Protector ZDB integration (EMC, XP, VA, or EVA) must be correctly installed and configured. For installation, refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*. For configuration, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide*.

- On UNIX, if the database is installed on raw partitions (rawdisk or raw logical volumes), then the volume group (HP-UX systems) or the disk group (Solaris systems) name on the application system and backup system must be identical.
- The SAP R/3 database *user* used by ZDB integrations to connect to the target SAP R/3 database during the backup must have the SYSDBA privilege granted. Refer to the Oracle documentation for more information on user privileges in Oracle.

The operating system *root* user on the Oracle Server also has to be added to either the Data Protector *admin* or *operator* user group on both, the application and backup system.

- The following Oracle software must be installed and configured on the application system:

- ✓ Oracle RDBMS Server
- ✓ SQL*Plus
- ✓ Oracle Net8 software

- The database on the application system can be installed on disk images, logical volumes, or filesystems. The Oracle datafiles *must* reside on disk array source volumes.

For online backup, the control files, online redo logs, and archived redo log files do not have to reside on disk array.

For offline backup, the control files, online redo logs, and archived redo log files *must* reside on disk array. However, note the following:

- If Oracle8i/9i control file, online redo logs, and Oracle9i SPFILE reside on a **different** volume group (if LVM is used) or source volume than Oracle8i/9i datafiles., instant recovery *is* enabled.
- If Oracle8i/9i control file, online redo logs, and Oracle9i SPFILE reside on the **same** volume group (if LVM is used) or source volume as Oracle8i/9i datafiles, instant recovery 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. Refer to “ZDB Integrations Omnirc Variables” on page A-9.

IMPORTANT

If you enable instant recovery by setting the above mentioned variables, note that the control file, online redo logs, and Oracle9i SPFILE are overwritten during instant recovery. In such case, you must restore the overwritten files from a separate backup before you can perform database recovery beyond the replica creation time. Therefore, you may want to move the control files and redo logs to different locations. Refer to “Reconfiguring an Oracle Instance for Instant Recovery” on page A-3 for more details.

For online backup, the control files and online redo logs as well as the archived redo log files do not have to reside on disk array.

For offline backup, the control files and online redo logs as well as the archived redo log files *must* reside on disk array.

- The SAP R/3 directory `sapbackup` must be NFS mounted (UNIX systems) or shared and mounted (Windows systems) on the backup system with the same mount path as on the application system. On UNIX, the directory must be shared with root permissions. The `sapbackup`, `saparch`, and `sapreorg` directories must not reside on the same source volumes as the database files.
- If the SAP R/3 database is installed on symbolic links, then these links have to be created on the backup system too.

Limitations

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

- ZDB using the Oracle Recovery Manager is not possible TBD.
- Do not use double quotes (" ") in object-specific pre-exec and post-exec commands.
- For this integration, Oracle8, Oracle8i, and Oracle9i are supported.
- Preview is not possible for SAP R/3 ZDB and restore sessions.
- Automatic database recovery is supported within the Data Protector Instant Recovery context and CLI.
- Incremental backups made with RMAN are not supported for instant recovery.

- Object copying and object mirroring is not supported for ZDB to disk.
- The `splitint` interface is supported only with SAP R/3 BRTOOLS 6.4 or newer.

SAP R/3 Integration Concept

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

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

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

Version 4.5 and higher of the SAP R/3 backup utilities allows Oracle data files to be backed up directly using the Oracle Recovery Manager (**RMAN mode**), as well as using the Data Protector Oracle Integration (hereafter referred to as the **backint mode**).

IMPORTANT

With ZDB, only backint mode is supported.

SAP R/3 Backup Utilities

SAP R/3 backup utilities are the following:

- BRBACKUP

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

- BRARCHIVE

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

- BRRESTORE

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

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

NOTE

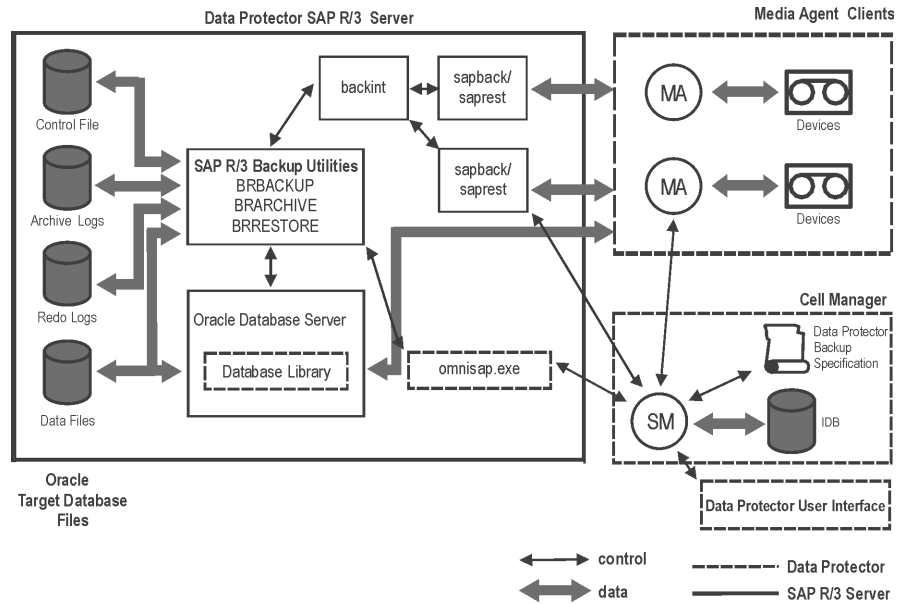
Data Protector supports all SAP R/3 backup utilities options, except for the -a and -b options. In order for Data Protector to support also the -a and -b options, set the OB2BRTNOSECU omnirc variable to 1. For more information about the omnirc file, refer to the *HP OpenView Storage Data Protector Zero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

**Data Protector
Integration
Software**

The Data Protector integration software consists of the following components, as depicted in Figure 2-1 on page 156.

- The `backint` program is a backup interface between the Data Protector software and the SAP R/3 backup and restore tools.
It is started using BRBACKUP or BRARCHIVE during a backup session, and BRRESTORE during a restore session.
- The `sapback` program performs the actual backup of files.
- The `saprest` program performs the actual restore of files.
- The `omnisap.exe` program is used by Data Protector to start the SAP R/3 backup tools.
- The `testbar2` utility checks the Data Protector part of the integration.
- The `util_sap.exe` program is used by Data Protector to configure the integration.
- The configuration file on the Cell Manager system contains data needed by Data Protector to run backups and restores.

Figure 2-1 SAP R/3 Backup Concept



Legend

- SM** The Data Protector Session Manager, which is the Data Protector Backup Session Manager during backup or the Data Protector Restore Session Manager during restore.
- Database Library** The interface between SAP R/3 Server processes and Data Protector.
- IDB** The IDB, which stores information about Data Protector sessions, such as session messages, and information about objects, data, used devices, and media.
- MA** The Data Protector General Media Agent.

**Backup Flow
Using Backint**

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

NOTE

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

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



Legend

- | | |
|-----|--|
| BSM | Data Protector Backup Session Manager |
| RSM | Data Protector Restore Session Manager |
| BMA | Data Protector Backup Media Agent |

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

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

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

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

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

2. `BRBACKUP` does the following:
 - Automatically changes the state of the Oracle Target Database (opened or closed), according to the backup type (online or offline).
 - Switches the Oracle Target Database to the `ARCHIVELOG` mode before the backup.

The archived redo log files are written to the archiving directory by Oracle and are backed up later using `BRARCHIVE`.

- Writes the `BRBACKUP` log during the backup session, with information about the backup file and the backup ID. These logs must be available in order to determine the location of the database files and archived redo log files during restores.
- Sets the tablespace mode (`BEGIN / END BACKUP`) in the case of online backup using `backint`.

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

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

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

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

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

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

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

NOTE

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

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

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

Restore Flow Using Backint

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

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

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

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

The RSM checks the availability of the restore devices, starts General Media Agents and establishes a connection between the `saprest` processes and General Media Agents. Data transfer begins at this stage. Data is sent from the media to the target disks. The General Media Agent finishes as soon as all `saprest` processes connected to it are completed.

3. When all the General Media Agents have finished, the RSM waits for a timeout (`SmWaitForNewClient` global variable) and completes the restore session, if no backint is started within this time frame.

SAP R/3 ZDB Concept

Refer to the *HP OpenView Storage 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.

The database on the application system can be installed on disk images, logical volumes, or filesystems. The Oracle datafiles *must* reside on disk array source volumes.

Depending on the location of the Oracle8i/9i control file, online redo logs, and Oracle9i SPFILE, the following two options are possible:

- Oracle8i/9i control file, online redo logs, and Oracle9i SPFILE reside on a **different** volume group (if LVM is used) or source volume than Oracle8i/9i datafiles.

By default, instant recovery *is* enabled.

- Oracle8i/9i control file, online redo logs, and Oracle9i SPFILE reside on the **same** volume group (if LVM is used) or source volume as Oracle8i/9i datafiles.

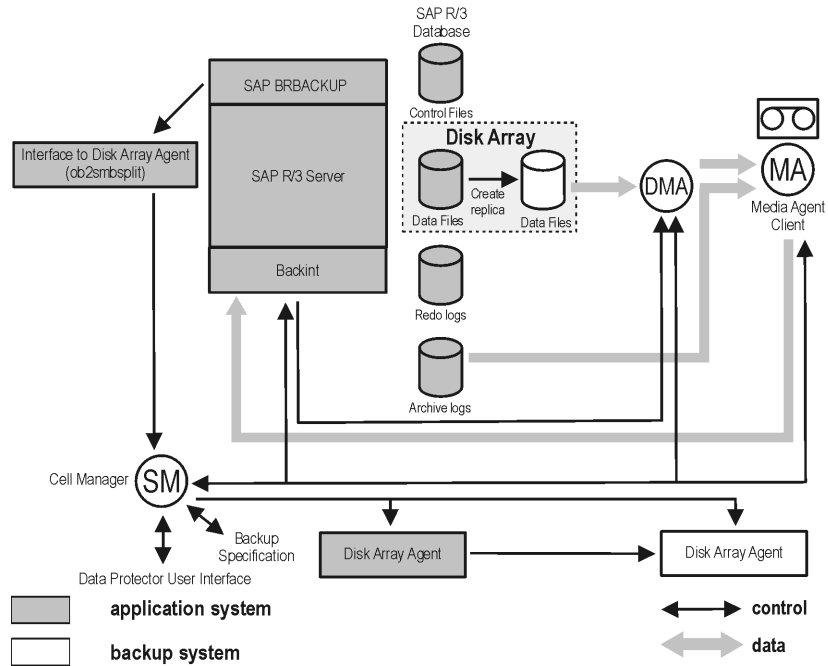
By default, instant recovery is *not* enabled. You can enable instant recovery by setting the ZDB_ORA_INCLUDE_CF_OLF, ZDB_ORA_INCLUDE_SPF, and ZDB_ORA_NO_CHECKCONF_IR omnirc variables. Refer to “ZDB Integrations Omnirc Variables” on page A-9.

IMPORTANT

If you enable instant recovery by setting the above mentioned variables, note that the control file, online redo logs, and Oracle9i SPFILE are overwritten during instant recovery. In such case, you must restore the overwritten files from a separate backup before you can perform database recovery beyond the replica creation time. Therefore, you may want to move the control files and redo logs to different locations. Refer to “Reconfiguring an Oracle Instance for Instant Recovery” on page A-3 for more details.

Note that the control files and online redo logs as well as the archived redo log files do not have to reside on disk array at all.

Figure 2-3 SAP R/3 Online Backup and Restore Concept

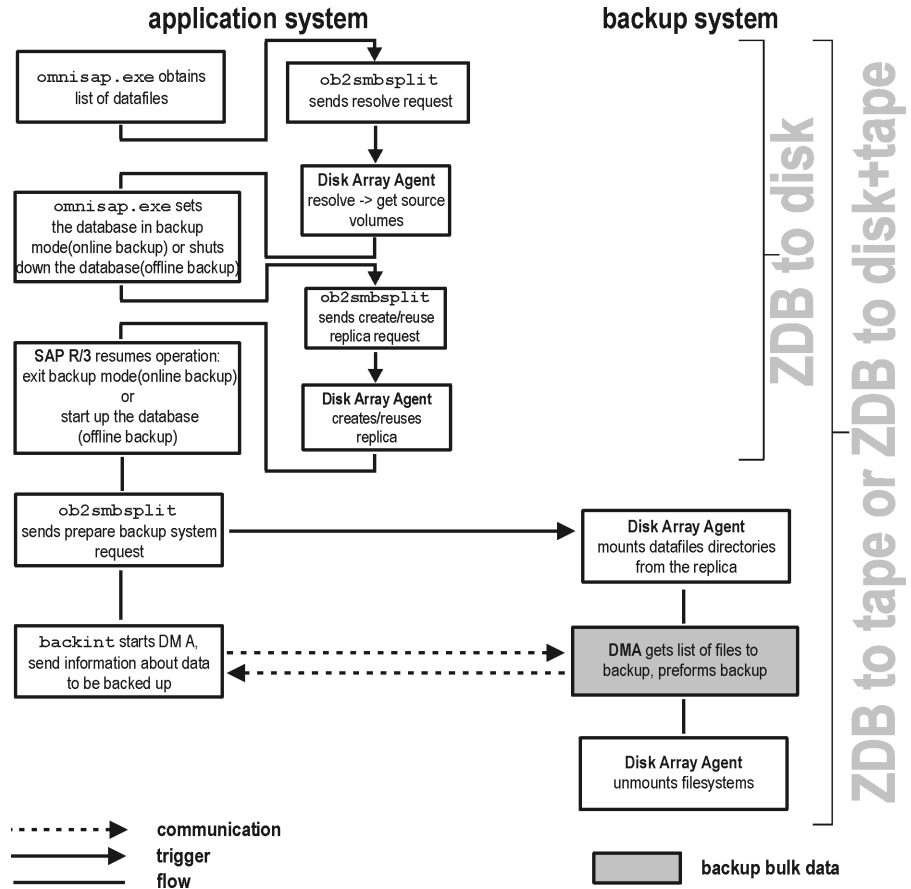


NOTE

The Figure 2-3 presents only the default integration behavior, where Oracle control file, online redo log files, and Oracle9i SPFILE reside on a different volume group than Oracle8i/9i datafiles. For more information on an alternative offline SAP R/3 backup and restore concept, refer to Refer to “ZDB Integrations Omnirc Variables” on page A-9 for more details.

Backup Process for SAP R/3

Figure 2-4 SAP R/3 ZDB Session Flow



Refer to the *HP OpenView Storage 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 information only about that part of the process that is relevant for SAP R/3.

Operations on source volumes (mounting, activating volume/disk groups...) in the process described below are dependent on/triggered by ZDB options. Refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide* for more information on these options.

ZDB Flow

The Data Protector SAP R/3 ZDB integration backup flow can be described as follows:

- The SAP R/3 backup specification is read and the Data Protector `omnisap.exe` program is started on the application system.
- The Data Protector `omnisap.exe` program starts the SAP R/3 BRBACKUP program on the application system. This program connects to the SAP R/3 database on the application system, puts the database into 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 backup mode (online backup) or is restarted (offline backup) after the replica is created.

If supported by BRTOOLS, Data Protector can use the `splitint` interface to reduce the time during which the database is in backup mode. By default, this option is disabled. On how to configure the `online_mirror` and `offline_mirror` options, see “SAP R/3 Specific Backup Options” on page 207.

- The `ob2smbsplit` command starts the process of resolving the backup configuration, producing a replica, and preparing the replica for backup.

The mountpoints for the backed up object are created on the backup system.

The backup volume groups (HP-UX systems) or backup disk groups (Solaris systems) are activated and the filesystems are mounted on the backup 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.

**ZDB to Tape or
ZDB to Disk+Tape**

If a ZDB-to-tape or a ZDB-to-disk+tape session is being performed, the processing continues as follows:

- On the application system, the SAP BRBACKUP program invokes the Data Protector `backint` program, which then starts one or more Data Protector **Data Movement Agents** (DMA) on the backup system and sends them information about the data to be backed up.
- DMAs send a request to the Data Protector Backup Session Manager (BSM) on the Cell Manager to start backup to tape.
- BSM reads the backup specification and then starts the General Media Agents (MA), which control and manage the backup devices.
- BSM sends to the DMAs the data for establishing the connection between the DMAs and MAs.
- DMAs connect to MAs.
- DMAs read data from the backup system and send it to MAs to back up the SAP R/3 client in parallel (to write the data to backup devices).
- At the end of data transfer, the backup system is disabled (filesystems are unmounted on all platforms and volume/disk groups deactivated on UNIX).
- With the EMC and XP integrations, the mirror disks are disabled; filesystems are dismounted and volume groups (HP-UX systems) or disk groups (Solaris systems) are deactivated and, depending on ZDB options, links are re-established.

Data Protector SAP R/3 Configuration File

Data Protector stores the SAP R/3 integration parameters for every configured SAP R/3 instance in the following file on the Cell Manager:

- On UNIX:
`/etc/opt/omni/server/integ/config/SAP/<client_name>%<ORACLE_SID>`
- On Windows:
`<Data_Protector_home>\Config\Server\Integ\Config\Sap\<client_name>%<ORACLE_SID>`

The parameters stored are:

- Oracle home directory
- encoded connection string to the target database
- BRTOOLS home directory
- the variables which need to be exported prior to starting a backup
- speed parameters (time needed for a specific file to back up - in seconds)
- manual balancing parameters

The configuration parameters are written to the Data Protector SAP R/3 configuration file:

- during configuration of the integration
- during creation of a backup specification
- when the configuration parameters are changed

IMPORTANT

To avoid problems with your backups, take extra care to ensure the syntax and punctuation of your configuration file match the examples.

NOTE

You can set up the parameters in the Environment section (sublist) of the file by referring to other environment variables in the following way:

```
SAPDATA_HOME=${ORACLE_HOME}/data
```

Syntax

The syntax of the Data Protector SAP R/3 configuration file is as follows:

```
ORACLE_HOME='<ORACLE_HOME>';
ConnStr='<ENCODED_CONNECTION_STRING_TO_THE_TARGET_DATABASE>';
BR_directory='<BRTOOLS_HOME>;
SAPDATA_HOME='<SAPDATA_HOME>';
Environment={
  [<ENV var1>='<value1>';]
  [<ENV var2>='<value2>';
  ...]
}
SAP_Parameters={<bckup_spec_name>=('-concurrency <# of
concurrency>' | '-time_balance' | '-load_balance' |
'-manual_balance');
}
speed={
  AVERAGE=1;
  '<filename>'=# of seconds needed to backup this file;
}
compression={'<filename>'=<size of the file in bytes after the
compression>;
}
manual_balance={<backup_specification_name>={'<filename>'=<device_
number>;
}
}
```

Example

This is an example of the file:

```
ORACLE_HOME='/app/oracle805/product';
ConnStr='EIBBKIBBEEIBBFIBBGHBBQDBBOFBBCFBBPFBBCFBBIFFBBGFBBDBBB
FBBCFBBDFBBBCFBB';

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

Environment={ }

SAP_Parameters={
    sap_weekly_offline=('-concurrency 1','-no_balance');
    sap_daily_online=('-concurrency 3','-load_balance');
    sap_daily_manual=('-concurrency 3','-manual_balance');
}

speed={
    AVERAGE=203971;
    '/file1'=138186;
    '/file2'=269756;
}

compression={
    '/file1'=1234;
    '/file2'=5678;
}

manual_balance={
    sap_daily_manual={
    '/file1'=1; /* file 1 is backed up by the first sapback */
    '/file2'=2; /* file 2 is backed up by the second sapback */
    '/file3'=1; /* file 3 is backed up by the first sapback */
    '/file4'=1;
    }
}
}
```

Setting, Retrieving, Listing, and Deleting Data Protector SAP R/3 Configuration File Parameters Using the CLI

The Data Protector SAP R/3 configuration file parameters are normally written to the Data Protector SAP R/3 configuration file after:

- the configuration of the SAP R/3 instance in Data Protector is completed.
- a new backup specification is created.
- a backup that uses balancing by time algorithm is completed.

The `util_cmd` Command

You can set, retrieve, list, or delete the Data Protector SAP R/3 configuration file parameters using the `util_cmd -putopt` (setting a parameter), `util_cmd -getopt` (retrieving a parameter), or `util_cmd -getconf` (listing all parameters) command on the Data Protector SAP R/3 client. The command resides in the `<Data_Protector_home>\bin` (Windows systems) or `/opt/omni/lbin` (HP-UX and 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 <sap_instance> [-local <filename>]
util_cmd -getopt[ion] [SAP <sap_instance>] <option_name>
[-sub[list] <sublist_name>] [-local <filename>]
util_cmd -putopt[ion] [SAP <sap_instance>] <option_name>
[<option_value>] [-sub[list] <sublist_name>] [-local
<filename>]
```

where:

`<option_name>` is the name of the parameter

`<option_value>` is the value for the parameter

`[-sub[list] <sublist_name>]` specifies the sublist in the configuration file to which a parameter is written to or taken from.

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

- When it is used with the `-getconf [ig]` option, it specifies the filename for the output of the command to be written to. If the `-local` option is not specified, the output is written to the standard output.
- When it is used with the `-getopt [ion]`, it specifies the filename of the file from which the parameter and its value are to be taken and then written to the standard output. If the `-local` option is not specified, the parameter and its value are taken from the Data Protector SAP R/3 configuration file and then written to the standard output.
- When it is used with the `-putopt [ion]` option, it specifies the filename for the output of the command to be written to. If the `-local` option is not specified, the output is written to the Data Protector SAP R/3 configuration file.

NOTE

If you are setting the `option_value` parameter as a number, the number must be put in single quotes, surrounded by double quotes.

Return Values

The `util_cmd` command displays a short status message after each operation (writes it to the standard error):

- Configuration read/write operation successful.
This message is displayed when all the requested operations have been completed successfully.
- Configuration option/file not found.
This message is displayed when either an option with the specified name does not exist in the configuration, or the file specified as the `-local` parameter does not exist.
- Configuration read/write operation failed.
This message is displayed if any fatal errors occurred, for example: the Cell Manager is unavailable, the Data Protector SAP R/3 configuration file is missing on the Cell Manager, etc.

Setting Parameters To set the Data Protector OB2OPTS and the Oracle NLS_LANG parameters for the SAP R/3 instance ICE, use the following commands on the Data Protector SAP R/3 client:

Windows

```
<Data_Protector_home>\bin\util_cmd -putopt SAP ICE OB2OPTS  
'-debug 1-200 INSTANCE.txt' -sublist Environment  
  
<Data_Protector_home>\bin\util_cmd -putopt SAP ICE  
NLS_LANG 'AMERICAN_AMERICA.US7ASCII' -sublist Environment  
  
<Data_Protector_home>\bin\util_cmd -putopt SAP ICE  
NLS_LANG "'10'" -sublist Environment
```

HP-UX and Solaris

```
/opt/omni/lbin/util_cmd -putopt SAP ICE OB2OPTS '-debug \  
1-200 INSTANCE.txt' -sublist Environment  
  
/opt/omni/lbin/util_cmd -putopt SAP ICE NLS_LANG \  
'AMERICAN_AMERICA.US7ASCII' -sublist Environment  
  
/opt/omni/lbin/util_cmd -putopt SAP ICE BR_TRACE "'10'"  
-sublist Environment  
  
/usr/omni/bin/util_cmd -putopt SAP ICE OB2OPTS '-debug \  
1-200 INSTANCE.txt' -sublist Environment
```

Retrieving Parameters To retrieve the value of the OB2OPTS parameter for the SAP R/3 instance ICE, use the following command on the Data Protector SAP R/3 client:

- On Windows: `<Data_Protector_home>\bin\util_cmd -getopt SAP ICE OB2OPTS -sublist Environment`
- On HP-UX and Solaris: `/opt/omni/lbin/util_cmd -getopt SAP ICE OB2OPTS -sublist \ Environment`

Listing Parameters To list all the Data Protector SAP R/3 configuration file parameters for the SAP R/3 instance ICE, use the following command on the Data Protector SAP R/3 client:

- On Windows: `<Data_Protector_home>\bin\util_cmd -getconf SAP ICE`
- On HP-UX and Solaris: `/opt/omni/lbin/util_cmd -getconf SAP ICE`

Deleting Parameters

To remove the value of the OB2OPTS parameter for the SAP R/3 instance ICE, use the following command on the Data Protector SAP R/3 client:

- **On Windows:** `<Data_Protector_home>\bin\util_cmd -putopt SAP ICE OB2PTS -sublist Environment`
- **On HP-UX and Solaris:** `/opt/omni/1bin/util_cmd -putopt SAP ICE OB2OPTS -sublist Environment`

Configuring the Integration

Configuration Overview

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

1. Configure an SAP R/3 backup owner (UNIX systems only).
2. Configure the Data Protector SAP R/3 integration.
3. Configure the SAP R/3 Oracle Database on the application system.
4. Configure the Data Protector SAP R/3 client on the application system.

Configuring an SAP R/3 Backup Owner in Data Protector on UNIX Systems

NOTE

This section is relevant only if you are configuring a UNIX environment.

To back up an SAP R/3 database using the Data Protector ZDB integration, add the following UNIX users to the Data Protector admin or operator user group for the *application system*:

- UNIX user `<ORACLE_SID>adm` in UNIX group `sapsys` - UNIX SAP administrator
- UNIX user `ora<ORACLE_SID>` in UNIX group `dba` - UNIX Oracle administrator
- UNIX user `root`

Add the user `root` to the Data Protector admin or operator user group also for the *backup system*.

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

Configuring the Data Protector SAP R/3 Integration

The environment must be configured for the following list of files:

- Oracle datafiles
- Oracle control file
- Oracle online redo log files

On UNIX systems, the owner of the filesystem or of the raw logical volume on which the database is mounted should be the UNIX user `ora<ORACLE_SID>` in the UNIX group `dba` (SAP R/3 backup owner). The minimum permissions should be `740`.

Configuring the SAP R/3 Oracle Database on the Application System

This section contains detailed steps needed to prepare the SAP R/3 software and the SAP R/3 Oracle database on the application system for a ZDB. To configure the SAP R/3 software on the application system, proceed as follows:

Perform the first two steps only if the Oracle SQL*Net V2 or Net8 TNS listener for the target database is not configured.

1. Configure the SQL*Net V2 or Net8 TNS listener, `listener.ora`, for the target database as described in the examples below, and start it using the following command on the system where the Oracle Server is installed:

```
<ORACLE_HOME>/bin/lsnrctl start (UNIX systems) or
```

```
<ORACLE_HOME>\bin\lsnrctl start (Windows systems)
```

Example - Oracle8

The following is an example of the
`<ORACLE_HOME>/network/admin/listener.ora` (UNIX systems) or
`<ORACLE_HOME>\network\admin\listener.ora` (Windows systems)
file for *Oracle8* SAP R/3 Database PROD on the application system
with hostname `alpha.hp.com`:

```
LISTENER =  
  (ADDRESS_LIST =  
    (ADDRESS=  
      (PROTOCOL= TCP)
```



```

        (Host= alpha.hp.com)
        (Port= 1521)
    )
)
SID_LIST_LISTENER =
  (SID_LIST =
    (SID_DESC =
      (GLOBAL_DBNAME= PROD)
      (ORACLE_HOME= /app/oracle805/product)
      (SID_NAME = PROD)
    )
  )
STARTUP_WAIT_TIME_LISTENER = 0
CONNECT_TIMEOUT_LISTENER = 10
TRACE_LEVEL_LISTENER = OFF

```

Example - Oracle8i

The following is an example of the
 <ORACLE_HOME>/network/admin/listener.ora (UNIX systems) or
 <ORACLE_HOME>\network\ad7min\listener.ora (Windows systems)
 file for the *Oracle8i* SAP R/3 Database PROD on the application system
 with hostname alpha.hp.com:

```

LISTENER =
  (DESCRIPTION_LIST =
    (DESCRIPTION =
      (ADDRESS =
        (PROTOCOL = TCP) (HOST = alpha.hp.com) (PORT =
1522)
      )
    )
  )
SID_LIST_LISTENER =

```

Data Protector SAP R/3 ZDB Integration

Configuring the Integration

```
(SID_LIST =
  (SID_DESC =
    (GLOBAL_DBNAME = PROD)
    (SID_NAME = PROD)
    (ORACLE_HOME = /app/oracle815/product)
  )
)
```

2. Configure the file `tnsnames.ora`, on the application system as described in the examples below in order to specify the SQL*Net V2 or Net8 service name of the target database:

Example - Oracle8

Here is an example for the
<ORACLE_HOME>/network/admin/tnsnames.ora (UNIX systems) or
<ORACLE_HOME>\network\admin\tnsnames.ora (Windows systems)
file for the *Oracle8* SAP R/3 Database PROD on the application system
with hostname `alpha.hp.com`:

```
PROD =
  (DESCRIPTION =
    (ADDRESS =
      (PROTOCOL = TCP)
      (Host= alpha.hp.com)
      (Port= 1521)
    )
    (CONNECT_DATA =
      (SID = PROD)
    )
  )
```

Example - Oracle8i

Here is an example for the
<ORACLE_HOME>/network/admin/tnsnames.ora (UNIX systems) or
<ORACLE_HOME>\network\admin\tnsnames.ora (Windows systems)
file for the *Oracle8i* SAP R/3 Database PROD on the application system
with hostname `alpha.hp.com`:

```
PROD =
```

```
(DESCRIPTION =  
  (ADDRESS_LIST =  
    (ADDRESS = (PROTOCOL = TCP) (HOST =  
alpha.hp.com) (PORT = 1522))  
  )  
  (CONNECT_DATA =  
    (SERVICE_NAME = PROD)  
  )  
)
```

3. Check if you can connect to the SAP R/3 Oracle target database from the application system:

Example

If the TNS service name for target database is PROD (see the previous example) and the Oracle system user has the password manager, execute the following `svrmgr1` (Oracle8/8i) or `sqlplus` (Oracle9i) command on the application system:

```
SVRMGR> connect system/manager@PROD  
Connected.
```

4. Shut down the SAP R/3 Oracle target database on the application system.
5. Enable the use of the Authentication Password File for the database administrator, by specifying the following in the `init<ORACLE_SID>.ora` file:

```
remote_login_passwordfile = exclusive
```

Refer to the Oracle documentation for instructions on how to set up the password file.

6. Enable automatic log archiving by setting `log_archive_start = true` in the `init<ORACLE_SID>.ora` file, where `<ORACLE_SID>` is the Oracle SID of the target database. Also specify the `log_archive_dest` option.

Example

This is an example of the `init<ORACLE_SID>.ora` file for the SAP R/3 Oracle target database (ORACLE_SID is PROD):

```
# @(#)initSID.ora 20.4.6.1 SAP 98/03/30  
#####  
# (c)Copyright SAP AG, Walldorf
```

Data Protector SAP R/3 ZDB Integration

Configuring the Integration

```
#####  
. . . .  
. . . . . . . . . .  
. . . . . . . . . .  
. . . . . . . . . .  
### ORACLE Authentication Password File  
remote_login_passwordfile = exclusive  
### ORACLE archiving  
log_archive_dest = /oracle/PROD/saparch/PRODarch  
log_archive_start = true  
. . . .
```

The changes in the `init<ORACLE_SID>.ora` file will become effective after the next start of the database.

7. Mount the Oracle database using the Oracle Server Manager and run the following commands in the Oracle Server Manager to enable the archive log mode:

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

Example

For example, if the SAP R/3 target database instance name is PROD, the `svrmgr1` (Oracle8/8i) command will look like:

```
export ORACLE_SID=PROD (UNIX systems) or  
set ORACLE_SID=PROD (Windows systems)  
  
svrmgr1  
SVRMGR> connect <user>/<passwd>@PROD;  
Connected.  
SVRMGR> 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.  
SVRMGR> alter database archivelog;  
Statement processed.  
SVRMGR> archive log start;  
Statement processed.  
SVRMGR> alter database open;
```

Note that with Oracle9i, the `sqlplus` command is used.

8. Ensure that the password of the Oracle user who is the database administrator is the same as the one specified by the `orapwd` command. If necessary, change the password by running the following Oracle Server Manager command:

```
alter user <user_name> identified by <password>;
```

Example

In this example, if the SAP R/3 target database instance name is PROD, a password file was created with the password manager, and the Oracle administrator user is `system`, the `svrmgrl` (Oracle8/8i) command will look like:

```
export ORACLE_SID=PROD (UNIX systems) or  
set ORACLE_SID=PROD (Windows systems)  
  
svrmgrl  
SVRMGR> connect internal;  
Connected.  
  
SVRMGR> alter user system identified by manager;
```

Note that with Oracle9i, the `sqlplus` command is used.

9. Grant `sysdba` and `sysoper` roles to the Oracle administrator user, `<user_name>`, by running the following Oracle Server Manager commands:

```
grant sysoper to <user_name>;  
grant sysdba to <user_name>;
```

Example

In this example, if the SAP R/3 target database instance name is PROD, a password file was created with the password manager, and the Oracle administrator user is `system`, the `svrmgrl` (Oracle8/8i) command will look like:

```
export ORACLE_SID=PROD (UNIX systems) or
```

Data Protector SAP R/3 ZDB Integration

Configuring the Integration

```
set ORACLE_SID=PROD (Windows systems)

$ svrmgrl
SVRMGR> connect system/manager@PROD;
Connected.
SVRMGR> grant sysoper to system;
Statement processed.
SVRMGR> grant sysdba to system;
Statement processed.
```

Note that with Oracle9i, the sqlplus command is used.

10. Test whether the Oracle user can log on to the SAP R/3 Oracle target database as the Oracle database administrator from the application system:
 - a. Login as an Oracle user on the application system.
 - b. On UNIX systems, export the `<ORACLE_HOME>` and `<ORACLE_SID>` variables.

On Windows systems, set the `<ORACLE_HOME>` and `<ORACLE_SID>` variables.
 - c. Start `svrmgrl` (Oracle8/8i) or `sqlplus` (Oracle9i).
 - d. Connect to the Oracle target database as an Oracle database administrator with the `SYSDBA` role.
 - e. Connect to the Oracle target database as an Oracle database administrator with the `SYSOPER` role.

Example - UNIX

In this UNIX example, if `ORACLE_SID` of the target database is `PROD`, and `ORACLE_HOME` is `/app/oracle816/product`, the commands on Oracle8/8i will look like:

```
id
uid=102(oracle) gid=101(dba)
export ORACLE_SID=PROD
export ORACLE_HOME=/app/oracle816/product
export
SHLIB_PATH=/app/oracle816/product/lib:/opt/omni/libin
svrmgrl
```

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

Note that with Oracle9i, the `sqlplus` command is used.

Example - Windows

In this Windows example, if `ORACLE_SID` of the target database is `PROD`, and `ORACLE_HOME` is `c:\oracle`, the commands on Oracle8/8i will look like:

```
set ORACLE_SID=PROD  
set ORACLE_HOME=c:\oracle  
svrmgr1
```

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

Note that with Oracle9i, the `sqlplus` command is used.

Mounting the SAPBACKUP Directory on the Backup System

UNIX Systems

On UNIX systems, the `SAPBACKUP` directory must be shared through NFS with root permissions:

- On HP-UX systems, add the `SAPBACKUP` directory to the `/etc/exports` file on the application system. For example, if your backup client is `backup.company.com` and `SAPBACKUP` points to `/opt/oracle/8.1.7/sapbackup`, add the following line:

```
/opt/oracle/8.1.7/sapbackup -root=backup.company.com
```
- On Solaris systems, add the `SAPBACKUP` directory to the `/etc/dfs/dfstab` file on the application system. For example, if your backup client is `backup.company.com` and `SAPBACKUP` points to `/opt/oracle/8.1.7/sapbackup`, add the following line:

```
share -F nfs -o root=backup.company.com  
/opt/oracle/8.1.7/sapbackup
```

Mount the directory on the backup host with the same name as on the application system.

Data Protector SAP R/3 ZDB Integration

Configuring the Integration

For example, on HP-UX, if your application client is `app.company.com`, you can add the following line to `/etc/fstab`:

```
app.company.com:/opt/oracle/8.1.7/sapbackup  
/opt/oracle/8.1.7/sapbackup nfs defaults 0 0
```

Windows Systems On Windows systems, share the SAP R/3 `sapbackup` directory with write permissions and name the share `sapmnt`.

Configuring a Data Protector SAP R/3 Client on the Application System

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

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

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

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

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

Cluster-Aware Clients on Windows/UNIX

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


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

NOTE

Make sure that the SAPLOCALHOST variable is not defined in the Environment section of the Data Protector SAP R/3 configuration file. Refer to “Data Protector SAP R/3 Configuration File” on page 166 for information on how to do that.

Cluster-Aware Clients on UNIX

Configure the Data Protector SAP R/3 integration on only one cluster node, since the Data Protector SAP R/3 configuration file resides on the Cell Manager. Use the virtual hostname when configuring the integration. However, you need to create a link to the Data Protector backint and splitint interface on all other nodes. Enter the following command on all other nodes:

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

If splitint is supported by BRTOOLS, enter also

```
ln -s /opt/omni/sbin/ob2smbsplit \
/usr/sap/<ORACLE_SID>/sys/exe/run/splitint
```

In a cluster environment, the environment variable OB2BARHOSTNAME must be defined as the virtual hostname before running the configuration from the command line (on the client). The OB2BARHOSTNAME variable is set as follows:

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

Add the SAP R/3 group dba user to Data Protector for the virtual server and for every node in the cluster. Refer to *HP OpenView Storage Data Protector Administrator's Guide* for information on how to add a user to Data Protector.

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

Cluster-Aware Clients on Windows

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

However, the Data Protector `backint` and `splitint` programs need to be manually copied to the correct location on all other nodes. On every other node, copy the `<Data_Protector_home>\bin\backint.exe` and `<Data_Protector_home>\bin\ob2smbsplit.exe` (if `splitint` is supported by BRTOOLS) to the directory where the SAP R/3 backup utilities reside and rename `ob2smbsplit.exe` to `splitint.exe`.

NOTE

Each SAP R/3 instance must be configured separately.

NOTE

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

Data Protector Inet User Account on Windows

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

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

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

The util_sap.exe Command

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

- List all Oracle instances on a particular system.

```
util_sap.exe -APP
```

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

```
util_sap.exe -OBS0 <ORACLE_SID>
```

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

```
util_sap.exe -OBS1 <ORACLE_SID> <TABLESPACE>
```

Using the CLI - UNIX Systems Only

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

NOTE

Each instance must be configured separately.

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

- <ORACLE_SID>

is the name of the Oracle database instance to be configured

- <ORACLE_HOME>

is the directory in which Oracle binaries are installed

- <targetdb_connection_string>

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

The <user_name> is the name by which a user is known to Oracle Server and to other users. Every user is identified by a password, and both must be entered to connect to an Oracle database. This user is, by default, used by brbackup and brarchive during backup. To define a different user when backing up, use the -u <user_name> as a BR Backup SAP R/3 backup option. See “SAP R/3 Specific Backup Options” on page 207.

NOTE

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

- `<SAPTOOLS_DIR>`

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

- `<SAPDATA_HOME>`

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

Using the GUI

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

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

The properties of a particular backup template can be seen in the corresponding pop-up window.

4. In the Results Area of the next page of the wizard, enter the following information:
 - Name of the SAP R/3 Database Server you want to configure. If the application is cluster-aware, select the virtual server of the SAP R/3 resource group.
 - Name of the Oracle Server instance (`ORACLE_SID`) on which the SAP R/3 Database Server is running.
 - On UNIX, enter also the UNIX user name and user group of the SAP R/3 user, as described in “Configuring an SAP R/3 Backup Owner in Data Protector on UNIX Systems” on page 173.

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

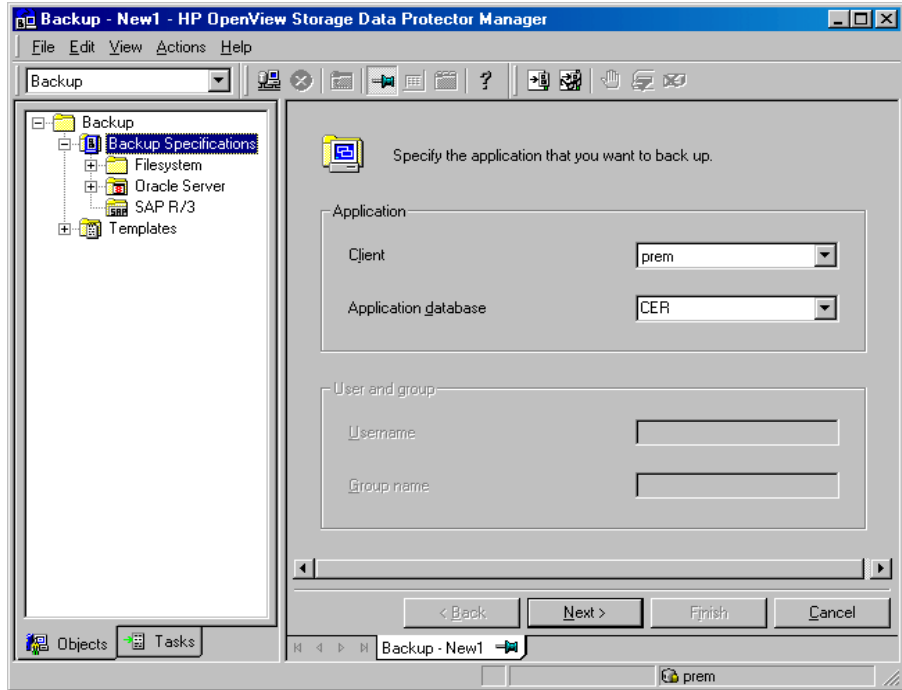
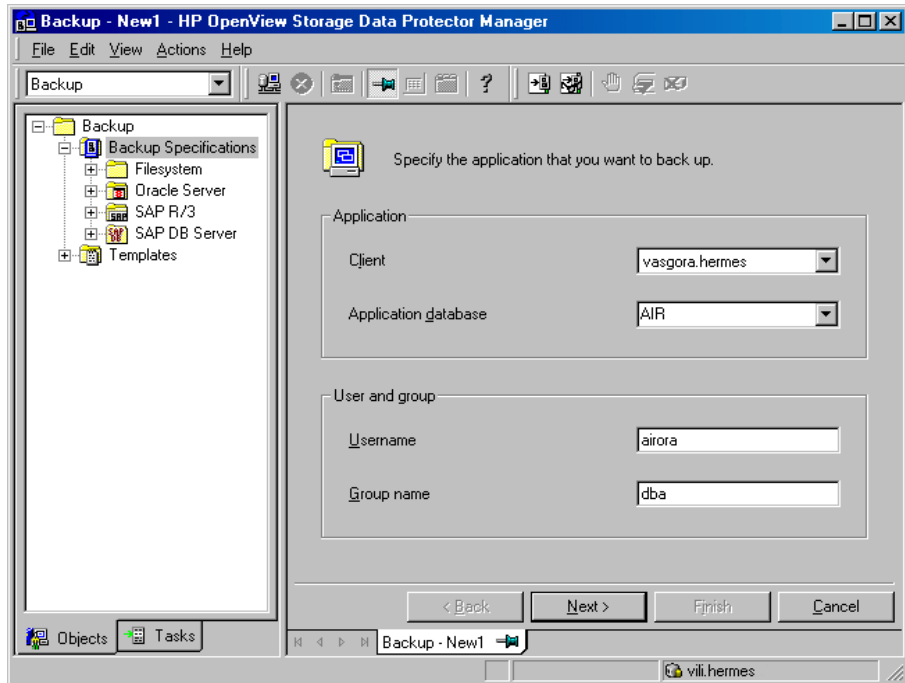


Figure 2-6 Specifying the SAP R/3 Database Server and the Oracle SID on UNIX



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

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

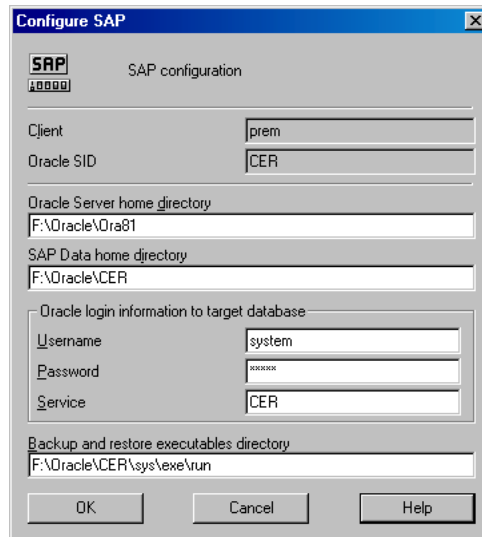
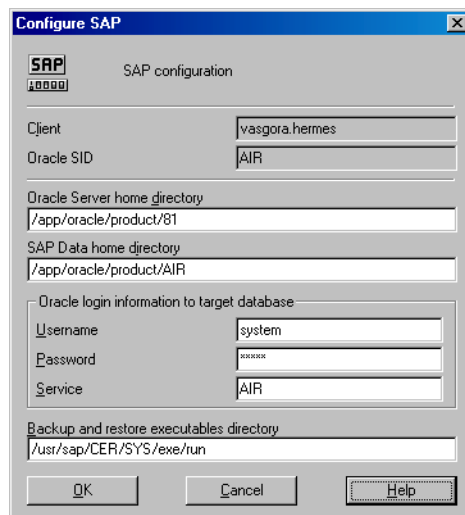


Figure 2-8 **Configuring an SAP R/3 Database Server on UNIX**



Data Protector SAP R/3 ZDB Integration

Configuring the Integration

5. Enter the following information in the Configure SAP dialog box:

- The Oracle Server home directory. If not specified, this is set to the default Oracle home directory.
- SAP data home directory (if not entered, this is set to `<ORACLE_HOME>`)
- The connection string to the Oracle Target Database.
See “Glossary” on page G-1 for more information on login connection strings.
- The directory where the SAP R/3 backup utilities are stored. By default, the utilities reside in the `\\<SAP_system>\sapmnt\<ORACLE_SID>\sys\exe\run` (Windows systems) or `/usr/sap/<ORACLE_SID>/SYS/exe/run` (UNIX systems) directory.

What Happens?

The following happens after saving the configuration.

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

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

Checking the SAP R/3 Configuration - Data Protector GUI

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

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

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

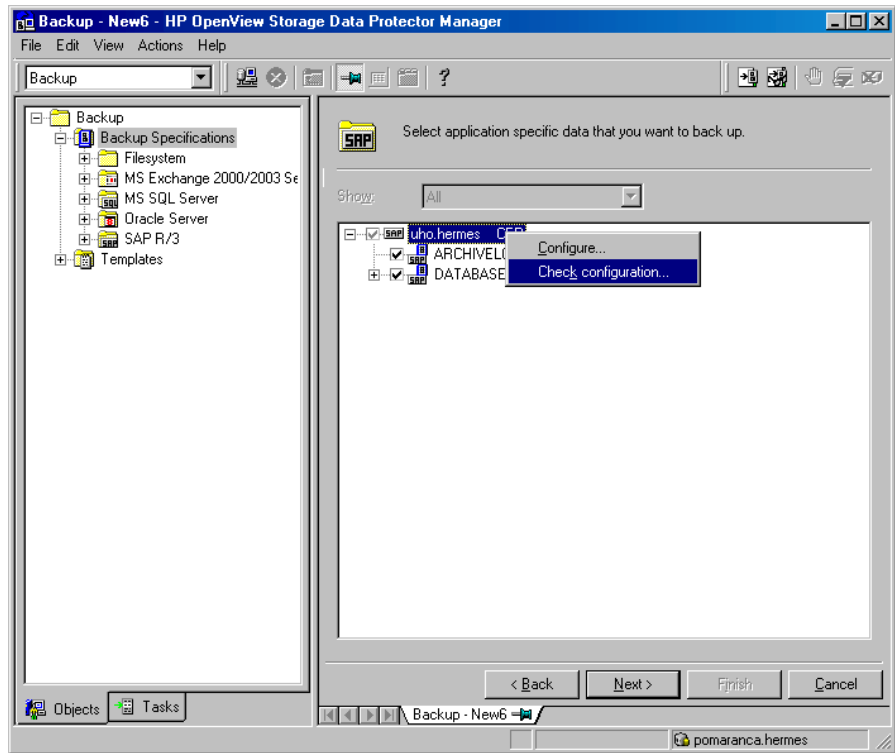
NOTE

The target database must be online during the check.

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

1. In the Data Protector Manager, switch to the Backup context.
In the Scoping Pane, expand Backup, Backup Specification, then SAP R/3.
2. In the Results Area, double-click the backup specification, then select Properties.
3. In the Source property page, right-click the name of the SAP R/3 Database Server, then click Check Configuration.

Figure 2-9 Checking the SAP R/3 Configuration



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

Checking the SAP R/3 Configuration - Data Protector CLI

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

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

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

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

On UNIX, to get the error description, start the
`/opt/omni/lbin/omnigetmsg 12 <error_number> command.`

Testing a ZDB Configuration using the CLI

UNIX

Log in to the application system as the user `root` or as the SAP R/3 user that is identified as described in “Configuring an SAP R/3 Backup Owner in Data Protector on UNIX Systems” on page 173. The identified SAP R/3 user and the user `root` must also be added to Data Protector `admin` or `operator` group. Then run the following command:

```
/opt/omni/lbin/util_sap.exe -CHKCONF <ORACLE_SID>
```

If an error occurs, use the following command to obtain an explanation:

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

where `<error_number>` is the number returned by the
`*RETVAL* <error_number>` line reported by the `util_sap.exe` script.

To check if the Oracle configuration is suitable for instant recovery run:

```
/opt/omni/lbin/util_sap.exe -CHKCONF_IR <ORACLE_SID>
```

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 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 instance. Refer to “Reconfiguring an Oracle Instance for Instant Recovery” on page A-3, 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 and database recovery beyond the replica creation time requires additional steps. Refer to “ZDB Integrations Omnirc Variables” on page A-9 on how to set the omnirc variables.

Windows

Run the following command on the application system:

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

If an error occurs, it is explained.

To check if the Oracle configuration is suitable for instant recovery, that is, if the recovery catalog or control files are not on source volumes that will be replicated during instant recovery, run:

```
<Data_Protector_home>\bin\util_sap.exe -CHKCONF_IR  
<ORACLE_SID>
```

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 the control files, Oracle9i 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 instance. Refer to “Reconfiguring an Oracle Instance for Instant Recovery” on page A-3, 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, Oracle9i SPFILE and online redo logs are overwritten and database recovery beyond the replica creation time requires additional steps. Refer to “ZDB Integrations Omnirc Variables” on page A-9 on how to set the `omnirc` variables.

Configuring an SAP R/3 ZDB

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

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

For a ZDB to disk, you also need to configure a backup device (for example, a standalone file device or null 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. For information on configuring a standalone device, refer to the online Help index keyword “standalone devices”.
2. Configure media pools and media for a backup. Refer to the online Help index keyword “creating media pools” for instructions.
3. Create a Data Protector SAP R/3 ZDB backup specification. See “Creating a Data Protector SAP R/3 ZDB Backup Specification” on page 195.
4. Create or modify the parameter file on the SAP R/3 Database Server. See “Creating or Modifying the Parameter File on the SAP R/3 Database Server” on page 211.

Creating a Data Protector SAP R/3 ZDB Backup Specification

Depending on whether you create an online or offline ZDB specification note the following:

- To perform an **online ZDB** of an SAP R/3 database, the database has to run in the ARCHIVELOG mode.

The archived redo log files, which are needed for recovery, are not backed up in a ZDB. Therefore, it is recommended to create two backup specifications:

- ZDB backup specification for backing up data files
- standard Data Protector SAP R/3 integration backup specification for backing up the application system archived redo log files

Configuring an SAP R/3 ZDB

In the case of ZDB to tape and ZDB to disk+tape, you can create one ZDB backup specification, in which you also select archived redo log files to be backed up.

In the case of a ZDB to disk, you *must* create two backup specifications.

- To perform an **offline ZDB**, create only a ZDB backup specification as the database is already in a consistent state.

UNIX

On UNIX systems, the owner of the backup specification should be the user `adm<ORACLE_SID>` in the group `sapsys`. Ensure that the user `adm<ORACLE_SID>` is a valid Data Protector user. The owner of `BRBACKUP` (`ora<ORACLE_SID>` by default) must be a valid Data Protector user as well.

Creating an SAP R/3 ZDB Backup Specification

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

1. In the HP OpenView Storage Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, and then Backup Specifications. Right-click SAP R/3 and then click Add Backup. The Create New Backup dialog box is displayed.
3. In the Create New Backup dialog box, select the appropriate template:

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 time during which the database is offline is shorter than with `SMB_offline`, but `splitint` must be supported by `BRTOOLS`.

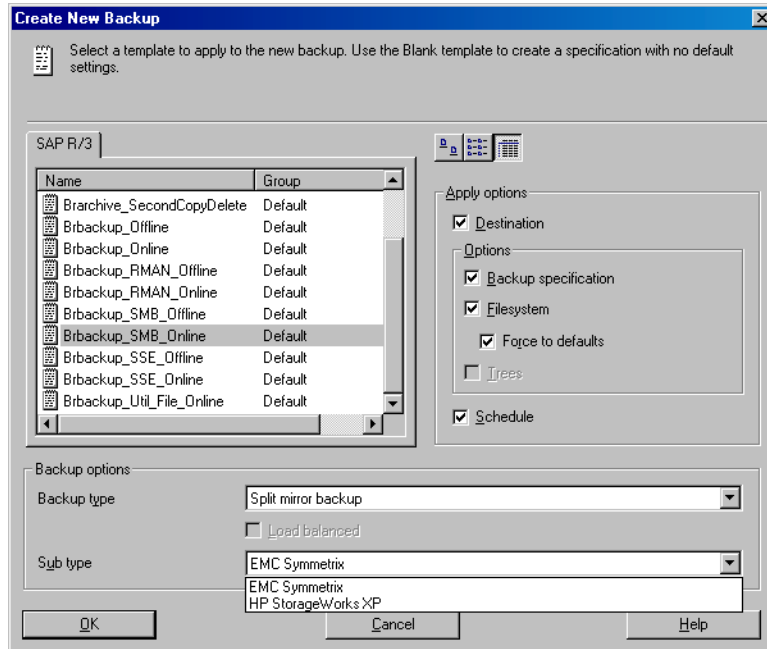
Brbackup_SPLITINT_Online Used for an online ZDB (split mirror or snapshot backup). The database is active during the creation of a replica. The time during which the database is in backup mode is shorter than with SMB_online, but splitint must be supported by BRTOOLS.

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 2-10.

Figure 2-10

**Selecting an Online ZDB Template and a Split Mirror Backup
TBD**

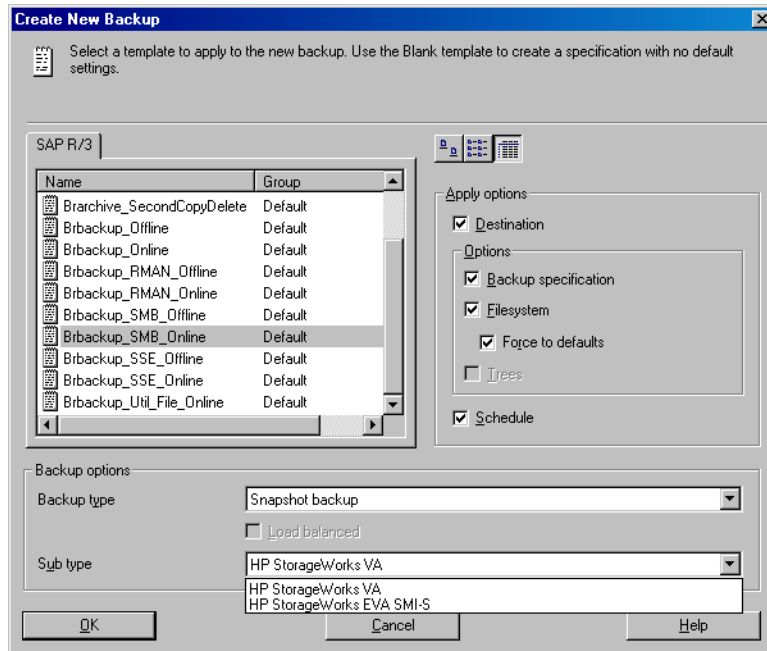


On VA and EVA

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 SMIS-S). See Figure 2-11.

Figure 2-11

Selecting an Online ZDB Template and a Snapshot Backup TBD



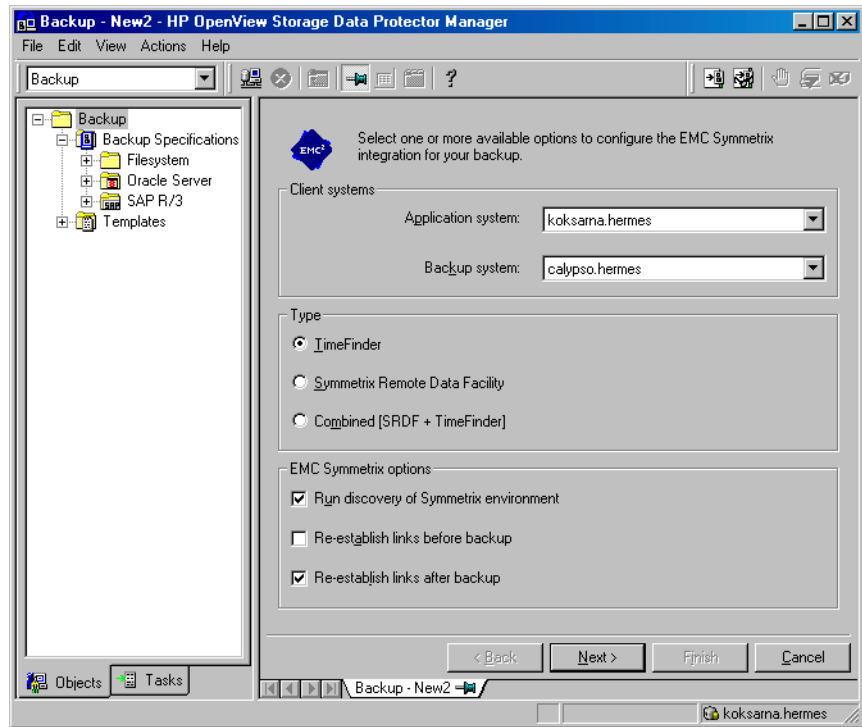
Click OK.

4. Under Client systems, in the Application system drop-down list, select the SAP R/3 Database server that you want to back up. If the application is cluster-aware, select the virtual server of the SAP R/3 package (UNIX systems) or resource group (Windows systems).

In the Backup system drop-down list, select the backup system.

Select other disk array specific backup options (see Figure 2-12 for EMC, Figure 2-13 for XP, Figure 2-14 for VA, or Figure 2-15 for EVA backup options). For detailed information on the options, press F1.

Figure 2-12 EMC Backup Options

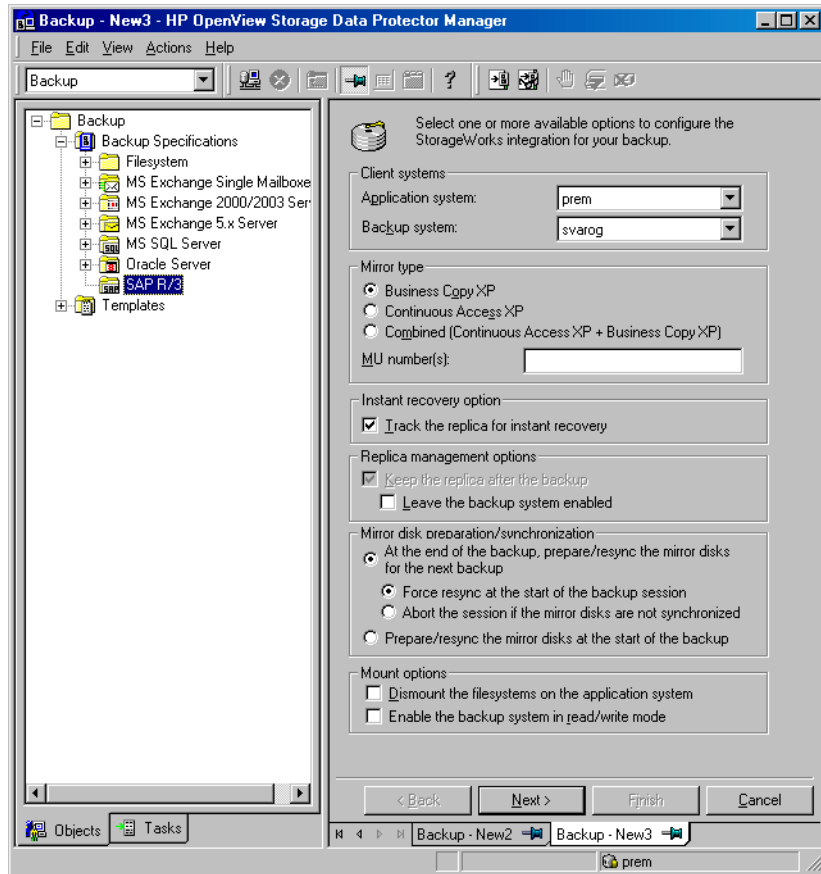


On XP

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

Figure 2-13

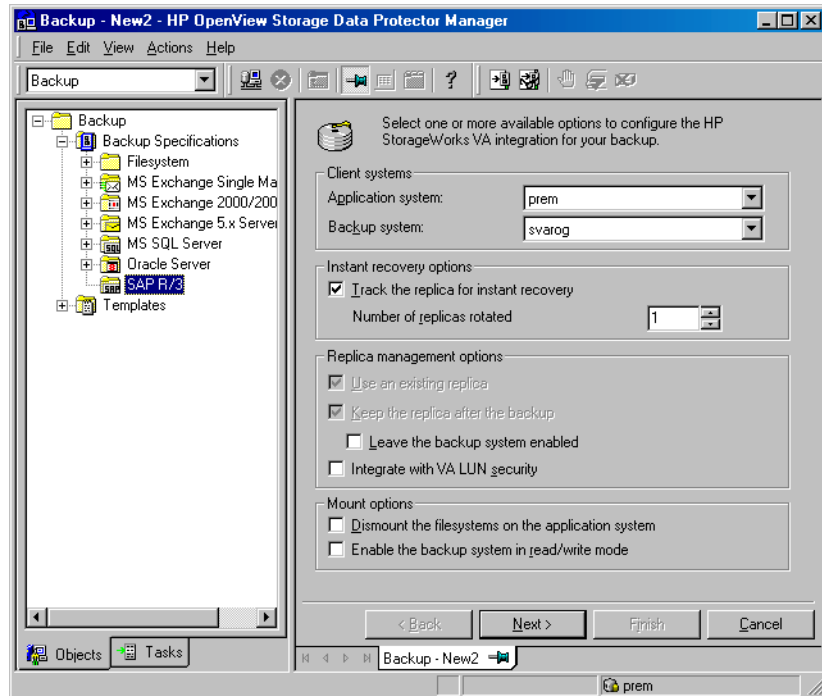
XP Backup Options



On VA

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

Figure 2-14 VA Backup Options

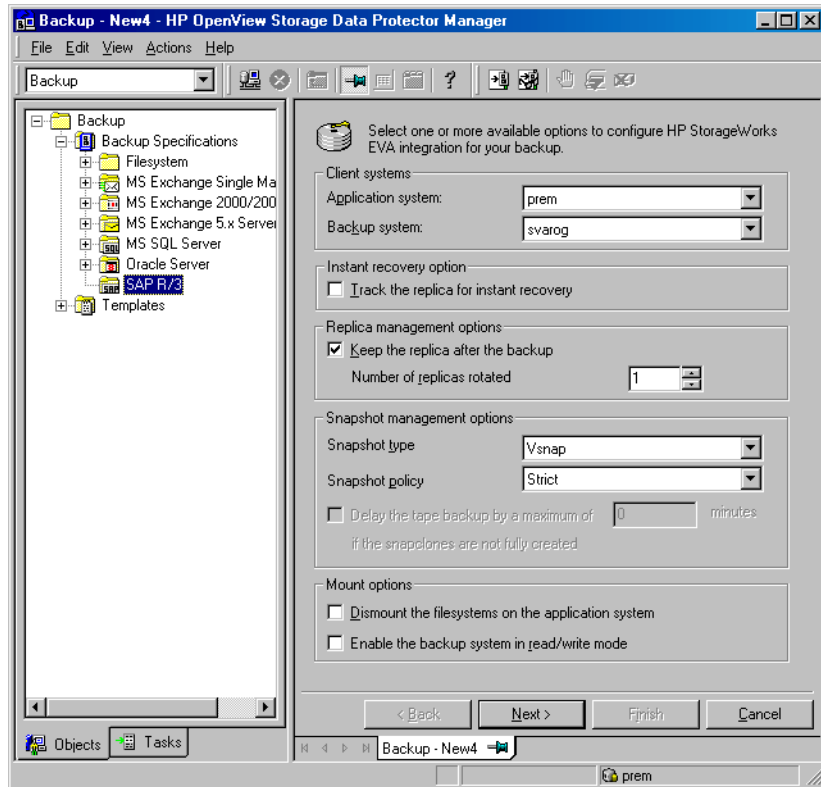


On EVA

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

Figure 2-15

EVA Backup Options



Click Next.

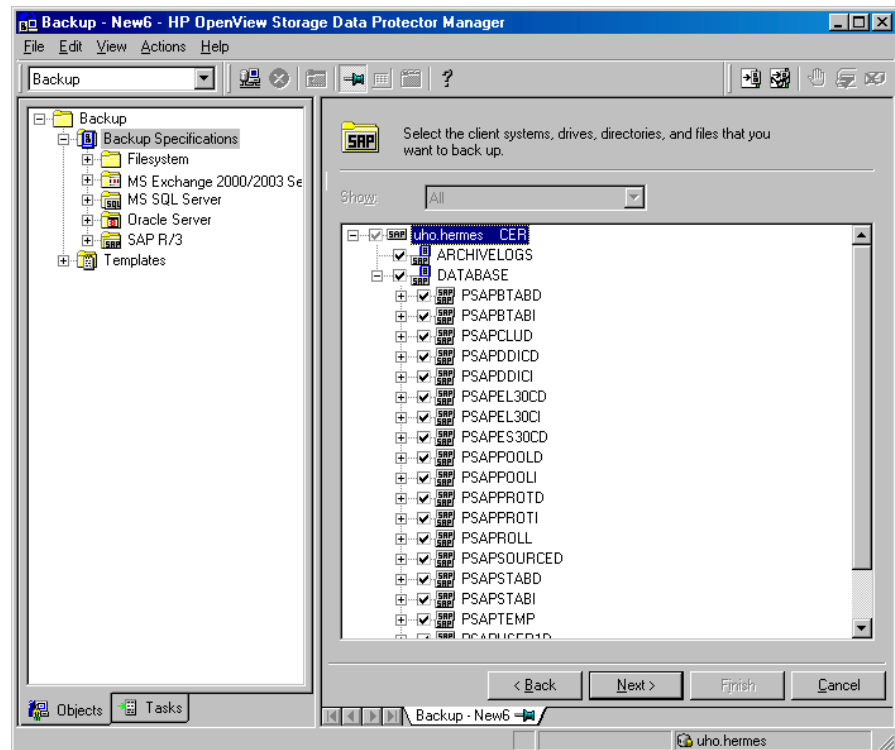
5. In the Results Area, provide the following information:

- In the Application database drop-down list, select the name of the Oracle Server instance (ORACLE_SID) on which the SAP R/3 Database Server is running.
- On UNIX, enter also the SAP R/3 user name and its group name as described in “Configuring an SAP R/3 Backup Owner in Data Protector on UNIX Systems” on page 173.

Click Next.

6. If the SAP R/3 Database Server is already configured, the Source dialog box is displayed. Otherwise, you are prompted to configure it. See “Configuring a Data Protector SAP R/3 Client on the Application System” on page 182 for details.
7. In the Source property page, if you have selected the Track the replica for instant recovery option, leave the whole DATABASE item selected; otherwise, select the database objects you want to back up. Database objects include archive logs, tablespaces, and data files. In the case of an online ZDB to tape or a ZDB to disk+tape, you can also select ARCHIVELOGS to be backed up.

Figure 2-16 Selecting Backup Objects



Click Next.

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**.

9. Select the backup options.

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

For information on the **Application Specific Option (SAP R/3 specific backup options)**, see “SAP R/3 Specific Backup Options” on page 207 or online Help.

Figure 2-17 SAP R/3 Application Specific Options - Online Backup

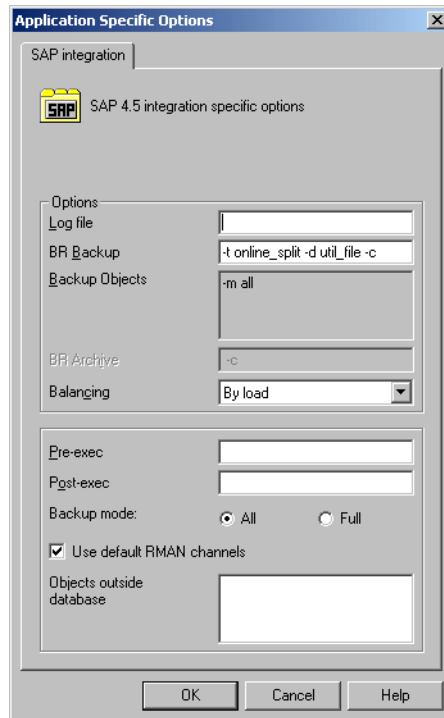
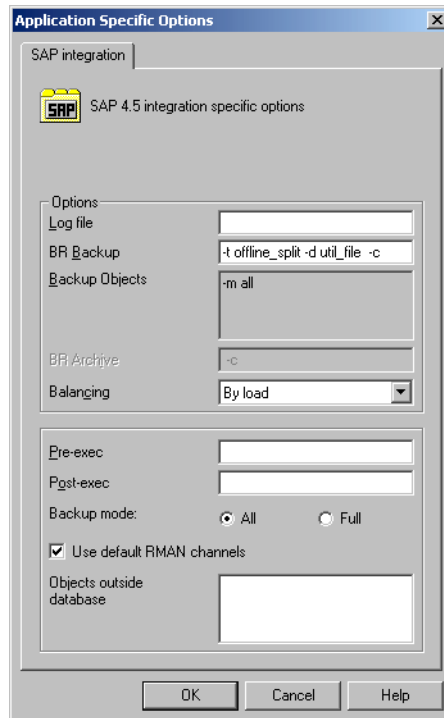


Figure 2-18

SAP R/3 Application Specific Options - Offline Backup



Click Next.


10. Optionally, schedule the backup. For more information, refer to “Scheduling a Backup” on page 214.

Note that only Full backup type is performed.

Click Next.

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

Figure 2-19 Saving the Backup Specification



Click OK.

To start the backup, see “Backing Up an SAP R/3 Database” on page 213.

12. On UNIX, after the backup specification is saved, verify that the owner of the backup specification is the specified SAP R/3 user. See “Configuring an SAP R/3 Backup Owner in Data Protector on UNIX Systems” on page 173 for details about this user.
13. If you have selected the ARCHIVELOGS item as shown in Figure 2-16 on page 203, you do not need to create another backup specification. Otherwise, create the standard SAP R/3 backup specification, as described in the *HP OpenView Storage Data Protector Integration Guide*, and select only ARCHIVELOGS to be backed up.

SAP R/3 Specific Backup Options

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

Log file

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

BR Backup

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

For example, for online or offline backup using the splitint interface, type `-t online_mirror` or `-t offline_mirror`.

If splitint is not supported by BRTOOLS, type `-t online_split` or `-t offline_split`.

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

Backup Objects

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

BR Archive

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

This option is not relevant for ZDB to disk.

Balancing: By Load

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

If this option is set and your backup device uses hardware compression, the size of the backed up file on the medium will not be the same as on the disk. To make Data Protector aware of this, make sure that you specify the size of the backed up file

on the medium in the `compression` section of the Data Protector SAP R/3 configuration file. See “Data Protector SAP R/3 Configuration File” on page 166 for information on how to do this.

Balancing: By Time

Groups files in subsets so that backup to all backup devices takes approximately the same time. This depends on the file types, the speed of the backup devices, and external influences (such as mount prompts), and is therefore best for environments with large libraries of the same quality. Each subset is backed up by one Data Protector `sapback` program, thus allowing concurrent backup of all subsets of the same type. Data Protector automatically stores backup speed information in the `speed` section of the Data Protector integration configuration file on the Cell Manager. It uses this information to optimize backup time.

This type of balancing may lead to non-optimal grouping of files in the case of online backup, or if the speed of backup devices varies significantly among devices.

Balancing: Manual

Manual balancing optimizes backups by allowing you to group files into subsets and back up these subsets using specific devices.

This option is not relevant 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

	<p>order use the Oracle Server Manager SQL command: <code>select * from dba_data_files</code></p>
Pre-exec	<p>Specifies an object pre-exec command with options that will be started on the SAP R/3 Database Server before backup. The command/script is started by Data Protector <code>omnisap.exe</code> and has to reside in the <code><Data_Protector_home>\bin</code> (Windows systems), or <code>/opt/omni/bin</code> (HP-UX and Solaris systems) directory. Only the filename must be provided in the backup specification.</p>
Post-exec	<p>Specifies an object post-exec command with options that will be started on the SAP R/3 Database Server after backup. The command/script is started by Data Protector <code>omnisap.exe</code> and has to reside in the <code><Data_Protector_home>\bin</code> (Windows systems), or <code>/opt/omni/bin</code> (HP-UX and Solaris systems) directory. Only the filename must be provided in the backup specification.</p>
Backup mode	<p>This option is not relevant for ZDB.</p>
Use default RMAN channels	<p>This option is not relevant for ZDB.</p>
Objects outside database	<p>With this option, you save all non-database files of the SAP R/3 and Oracle environments. This means that the following directory trees can be saved:</p> <pre>/sapmnt/<ORACLE_SID> /usr/sap/<ORACLE_SID>, /usr/sap/trans/<ORACLE_HOME></pre>

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

NOTE

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

See online Help and the *HP OpenView Storage Data Protector Administrator's Guide* for details on other specific Data Protector backup options.

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

The parameter file is used by SAP R/3 to set specific SAP R/3 backup options. A template for the parameter file is located on the application system as:

- On UNIX: `<ORACLE_HOME>/dbs/init<ORACLE_SID>.sap`
where `<ORACLE_SID>` represents the identifier for the instance.
- On Windows: `<ORACLE_HOME>\database\init<ORACLE_SID>.sap`

Perform the following to configure the parameter file for this integration:

- BRBACKUP calls the command specified in the parameter `split_cmd` to trigger the creating of a replica.

Define the `split_cmd` parameter as follows:

— On UNIX: `split_cmd = "/opt/omni/lbin/ob2smbsplit $"`

— On Windows: `split_cmd = "<Data_Protector_home>\bin\ob2smbsplit $"`

NOTE

On Windows systems, if the path specified in the `split_cmd` parameter contains space(s), use Windows short name(s) to set the path.

At run time BRBACKUP will replace the optional sign “\$” with the name of the text file containing the names of files to be backed up.

Configuring an SAP R/3 ZDB

- Define the service name of the target database to link the backup system to the application system:

```
primary_db = LOCAL
```

Backing Up an SAP R/3 Database

To run a ZDB-to-disk, ZDB-to-tape, or ZDB-to-disk+tape session of an SAP R/3 database, use any of the following methods:

Backup Methods

- Schedule a backup of an existing SAP R/3 ZDB backup specification using the Data Protector Scheduler. See “Scheduling a Backup” on page 214.
- Start an interactive backup of an existing SAP R/3 ZDB backup specification using the Data Protector GUI or the Data Protector CLI. See “Running an Interactive Backup” on page 215.

Considerations

Before running an SAP R/3 ZDB session, note the following:

- It is not possible to start 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.
- If the ARCHIVELOGS item as shown in Figure 2-16 on page 203 was selected during the configuration of the backup specification, you cannot run a ZDB-to-disk session until the archive log backup is completed.
- In the case of a ZDB-to-disk session, BRBACKUP returns an error in the Data Protector monitor, which should be ignored:

```
Ignoring BRBACKUP return value 5 because it is disk-only backup
```
- With the XP integration, if the LVM Mirroring configuration is used, a warning message is issued in the Data Protector monitor during the backup, since the volume group source volumes on the application system do not have their BC pairs assigned. This warning message should be ignored.
- If the control files, Oracle9i 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 instance. Refer to “Reconfiguring an Oracle Instance for Instant Recovery” on page A-3, 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, Oracle9i SPFILE and online redo logs are overwritten during instant recovery and you must restore the overwritten files from a separate backup before you can perform database recovery beyond the replica creation time. Refer to “ZDB Integrations Omnirc Variables” on page A-9 on how to set the `omnirc` variables.

Scheduling a Backup

Scheduling a backup specification means setting 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, refer to the online Help index keyword “scheduled backups”.

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

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

Note that the backup type is ignored for ZDB sessions. It is set to Full.

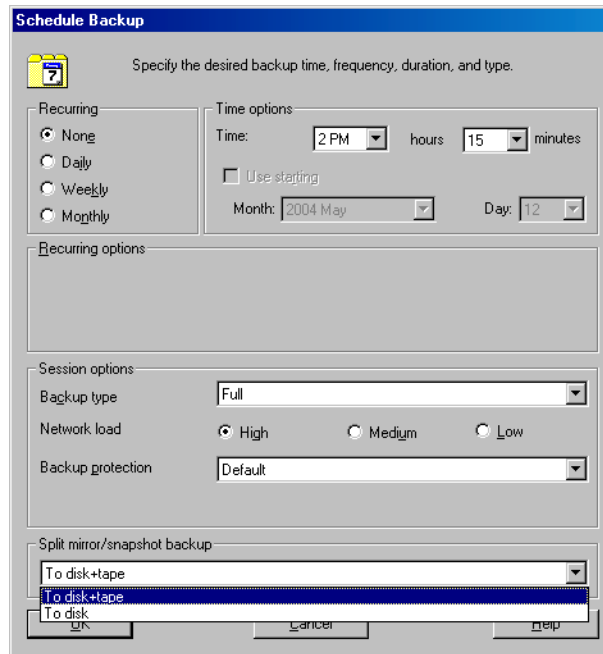
In the case of a ZDB-to-disk or a ZDB-to-disk+tape session, specify the Split mirror/snapshot backup option. See Figure 2-20.

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 option is not selected in the backup specification.

Figure 2-20

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.

Starting a Backup Using the GUI

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

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

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

Note that the backup type is ignored for ZDB sessions. It 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. See Figure 2-21.

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 option is not selected in the backup specification.

Figure 2-21

Selecting ZDB-to-Disk or ZDB-to-Disk+Tape Session When Starting an Interactive Backup



Click OK.

An interactive backup can also be started from the CLI.

Cluster-Aware Clients

In a cluster environment, the environment variable `OB2BARHOSTNAME` must be defined as the virtual hostname before running a backup from the command line (on the client). The `OB2BARHOSTNAME` variable is set as follows:

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

Tru64 Cluster

Before starting a backup on the Tru64 Cluster, create the following links:

```
ln -s /sapfiles/admin/dbs/initsap.dba initSAP.dba
ln -s /sapfiles/admin/dbs/initsap.ora initSAP.ora
ln -s /sapfiles/admin/dbs/initsap.sap init SAP.sap
```

Starting a Backup Using the CLI

To start an SAP R/3 **ZDB-to-tape** or **ZDB-to-disk+tape** session using the Data Protector CLI, use the following command:

```
omnib -sap_list <ListName>
```

To start an SAP R/3 **ZDB-to-disk** session using the Data Protector CLI, use the following command:

```
omnib -sap_list <ListName> -disk_only
```

where *<ListName>* is the name of the backup specification. For more information on the `omnib` command, refer to its man page.

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.

Restoring an SAP R/3 Database

An SAP R/3 database can be restored using one of the following two methods:

- Standard restore from backup media to the application system on LAN.

To utilize this method, a backup copy on media as a result of a ZDB-to-tape or ZDB-to-disk+tape session must exist. Using this method, the following can be achieved:

- partial restore and database recovery
- full database restore and database recovery
- restore of lost files and archive log restore, and database recovery

For information, see “Restoring Using the Data Protector GUI” on page 219, “Restoring Using the Data Protector CLI” on page 239, or “Restoring Using the SAP R/3 Commands” on page 240.

- Instant recovery

To utilize this method, a replica as a result of a ZDB-to-disk or ZDB-to-disk+tape session must exist and must be marked for instant recovery. Using this method, a recovery of an SAP R/3 database can be performed.

For information, see “Instant Recovery and Database Recovery” on page 224.

See also Table 2-1 on page 147 for an overview of recovery methods depending on the backup type and type of recovery.

Considerations

Use the Data Protector GUI to restore an SAP R/3 database from backup media to the application system on LAN.

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

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

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

Localized SAP R/3 Object Names When selecting objects for restore, Data Protector displays the actual names of the files as they are written to the filesystem and not SAP R/3 names, which are displayed when selecting objects for backing up. As a result, if the names contain non-ASCII characters, some of the characters may display different as in the backup specification, depending on your system settings (code pages or locale). This does not impact restore, which is still completed successfully.

Finding Information Needed for Restore

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

Execute the following commands:

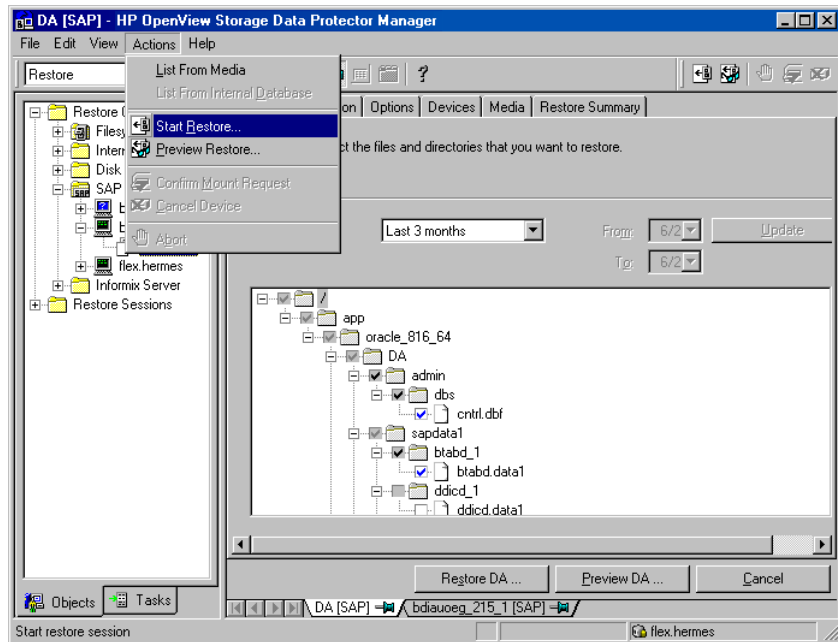
- `omnidb -sap`
to get a list of SAP R/3 objects.
- `omnidb -sap <object_name>`
to get details on a specific object, including the SessionID.

Restoring Using the Data Protector GUI

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

1. In the HP OpenView Storage Data Protector Manager, switch to the Restore context.
2. In the Scoping Pane, expand Restore Objects, SAP R/3, and then select the client (backup system) from which you want to restore. A list of backed up objects is displayed in the Results Area. See Figure 2-22.

Figure 2-22 Restoring SAP R/3 Database Objects



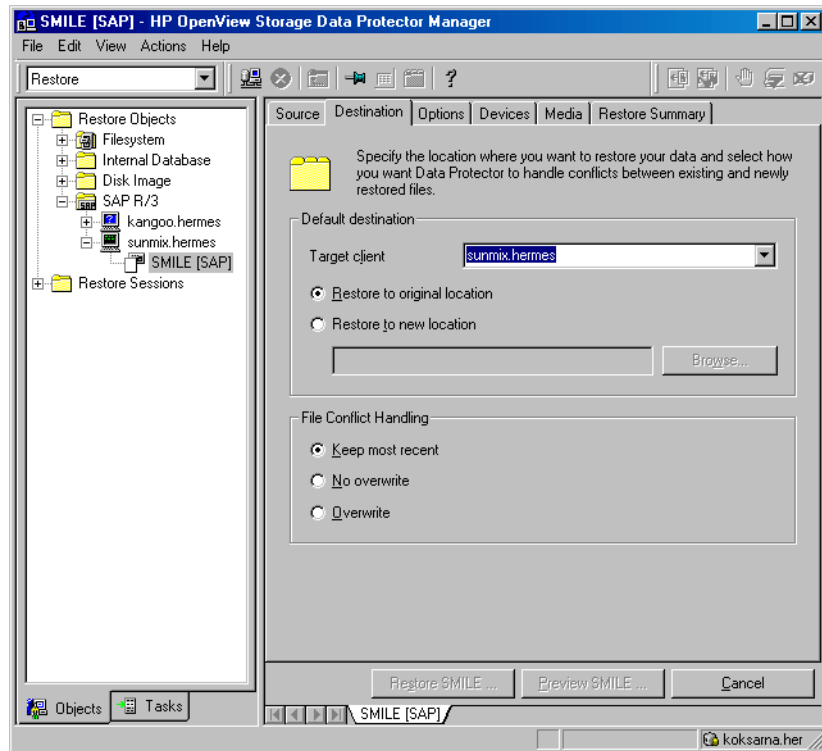
3. Select the backed up SAP R/3 object you want to restore.

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

Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for more information on various restore options available.

4. Under the Destination tab, select the application system as Target client. See Figure 2-23.

Figure 2-23 Selecting the Application System



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

Note that you can use a different device for the restore than the one used for the backup. Refer to the “Restoring Under Another Device” section in the *HP OpenView Storage Data Protector Administrator’s Guide* for more information on how to perform a restore using another device.

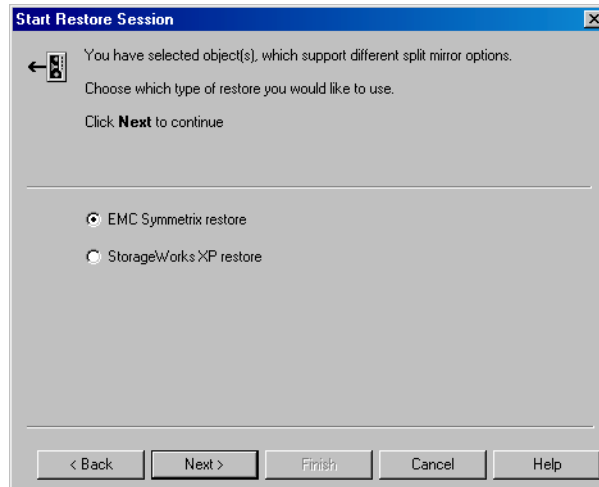
6. After you have set the restore options, click Restore. The Start Restore Session dialog box is displayed.
7. Click Next to specify the report level and network load.

On EMC and XP

8. This step is relevant only if you have both the EMC Symmetrix Agent and HP StorageWorks XP Agent components installed on the application system.

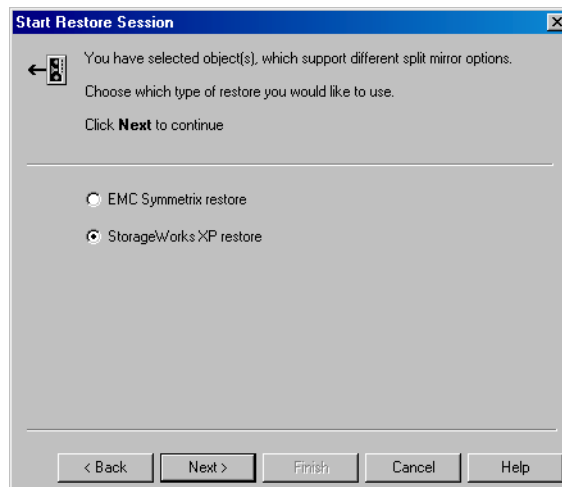
On EMC, leave the EMC Symmetrix restore option selected. See Figure 2-24.

Figure 2-24 **Selecting the EMC Restore**



On XP, select StorageWorks XP restore. See Figure 2-25.

Figure 2-25 **Selecting the XP Restore**



Click Next.

On EMC or XP

9. In the EMC Symmetrix mode or Mirror mode drop-down list, leave the Disabled option selected. This sets the restore from backup media to the application system directly.

Figure 2-26

EMC - Selecting Restore to the Application System Directly

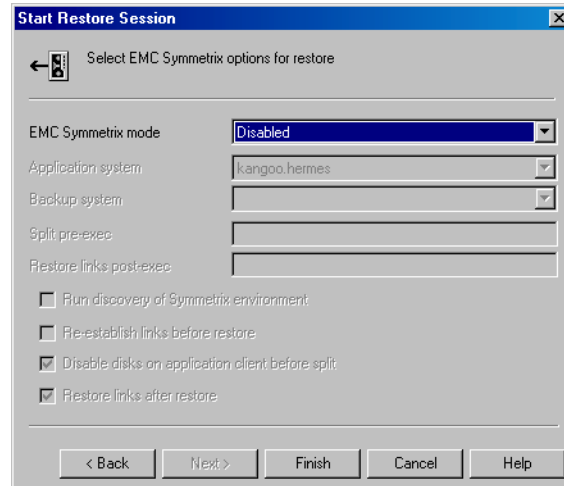
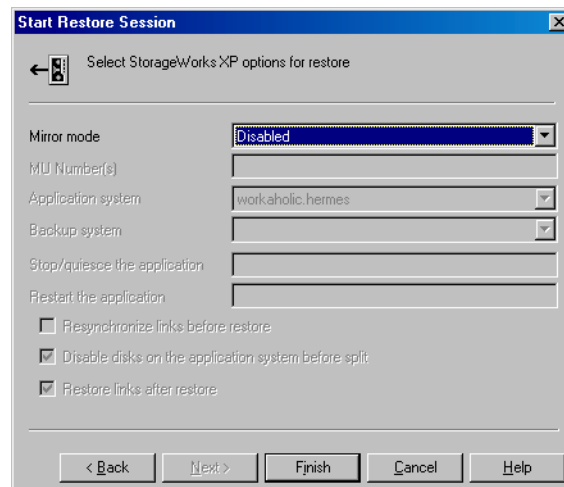


Figure 2-27

XP - Selecting Restore to the Application System Directly



10. Click Finish to start the restore session.

11. To recover the database, additional database or application recovery-related steps must be performed. See “Examples of SAP R/3 Database Restore” on page 240 for examples on how to perform database or application recovery-related steps.

Instant Recovery and Database Recovery

Refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide* and *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide* for general information on instant recovery.

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 the instant recovery procedure. During database recovery, archive log backups performed after the ZDB are restored from tape by the SAP BRTOOLS utility. If selected, the logs are reset and the database is opened.

IMPORTANT

If the control file, online redo logs, and Oracle9i SPFILE are on the same source volumes as datafiles and you enable instant recovery by setting the `omnirc` variables as described in “ZDB Integrations Omnirc Variables” on page A-9, the control file is overwritten during the instant recovery. As a consequence, you must restore the overwritten files from a separate backup before you can perform database recovery beyond the replica creation time.

If data is backed up using the Data Protector SAP R/3 ZDB offline backup, it is not possible to perform instant recovery from such a session.

Instant Recovery Procedure

To perform an instant recovery, proceed as follows:

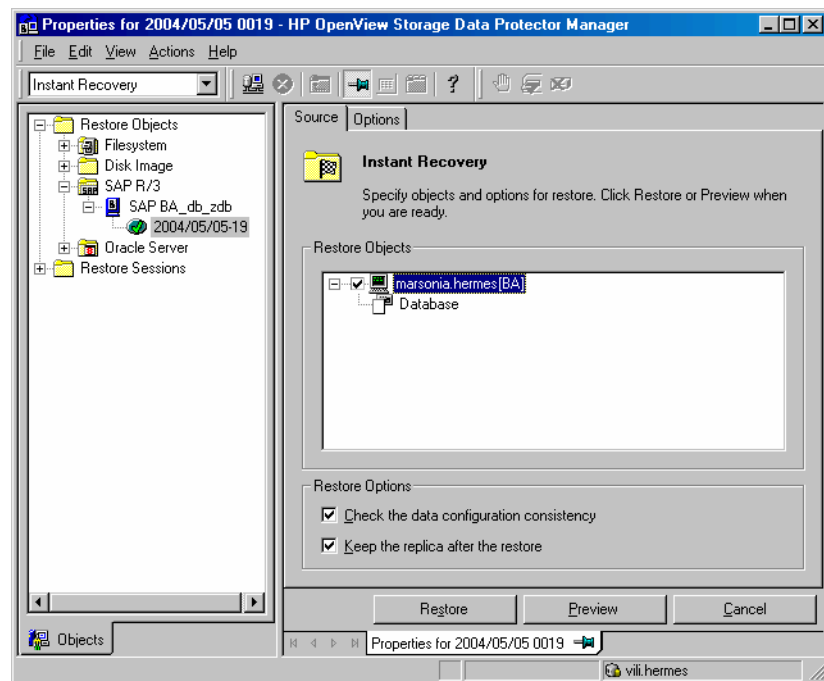
1. Shut down the SAP R/3 database using `svrmgrl` (Oracle8i) or `sqlplus` (Oracle9i):

For example:

```
svrmgrl
svrmgrl> shutdown immediate
svrmgrl> 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. With HP StorageWorks Virtual Array or StorageWorks Disk Array XP, it is recommended to set the Keep the replica after the restore option to enable a restart of an instant recovery session.

Figure 2-28 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.

NOTE

If you intend to manually perform a database recovery after instant recovery using SAPDBA, see “Examples of SAP R/3 Database Restore” on page 240 for examples on how to perform database or application recovery-related steps.

- To automatically perform a database recovery after an instant recovery, select the recovery options. For details refer to “Database Recovery Options” on page 226.

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

Recovery Enables database recovery after instant recovery.

Recover until The options in this drop-down list enable 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.

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. See “Configuring an SAP R/3 Backup Owner in Data Protector on UNIX Systems” on page 173 for more information on this user and on how to identify it.

`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 not available if the `Recover until` option is set to `Now`.

The following are Oracle recommendations on when to reset the logs:

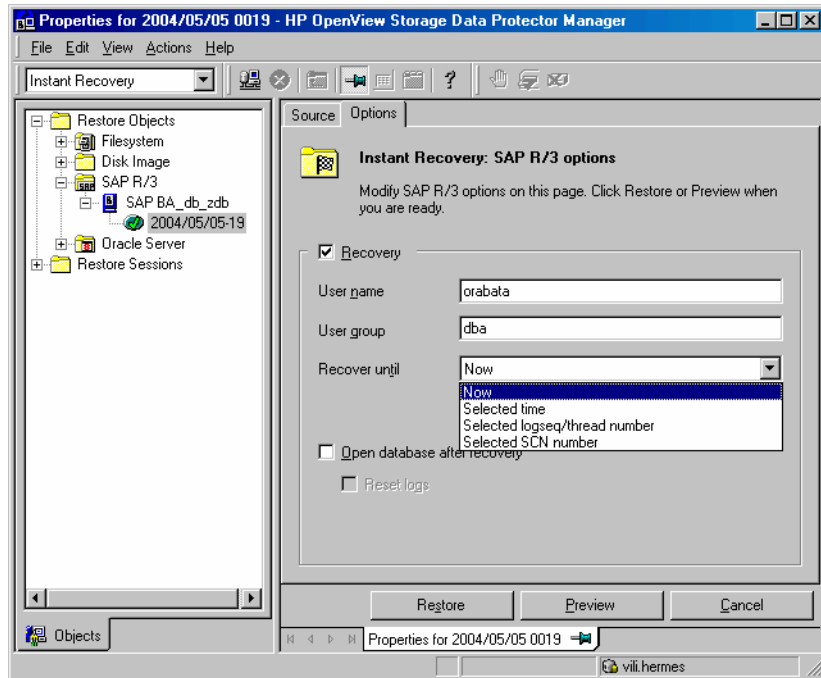
Always reset the logs:

- 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.

Figure 2-29 SAP R/3 Recovery Options



Troubleshooting

This section contains general checks and verifications and a list of problems you might encounter when using the Data Protector SAP R/3 integration.

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

For general ZDB, restore, and instant recovery related troubleshooting, see the troubleshooting sections in the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide*.

See also the troubleshooting section in the SAP R/3 chapter of the *HP OpenView Storage 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 OpenView Storage Data Protector Product Announcements, Software Notes, and References* for general Data Protector limitations, as well as known problems and workarounds.
- ✓ See http://www.openview.hp.com/products/datapro/spec_0001.html for an up-to-date list of supported versions, platforms, and other information.

General Troubleshooting

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

Problem

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

```
*RETVAl*8422
```

Action

Rename the Oracle8i `svrmgr1` binary to something else so that Data Protector will not find it. The Oracle upgrade process from Oracle8i to Oracle9i does not remove the Oracle8i `svrmgr1` binary, rather it changes

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

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. For Oracle9i, use `sqlplus` (and not `svrmgrl`):

```
export ORACLE_SID
export ORACLE_HOME
<ORACLE_HOME>/bin/svrmgrl
```

At the `Svrmgrl` or `SQLPlus` prompt, type:

```
connect <user>/<passwd>@<service>
select * from dba_tablespaces
exit;
```

If it fails, start 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:

```
<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. For Oracle9i, use `sqlplus` (and not `svrmgrl`):

```
export ORACLE_SID
export ORACLE_HOME
<ORACLE_HOME>/bin/svrmgrl
```

At the `Svrmgrl` or `SQLPlus` prompt, type:

```
connect <login information to target database>;
```



```
exit;
```

If it fails, refer to the Oracle documentation on how to create the TNS configuration file (TSNAMES.ORA).

Verifying the Prerequisites (SAP R/3 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 R/3 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 R/3 Online help for instructions on how to execute a backup to disk using the SAP R/3 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 R/3 Online help for instructions on how to execute a restore to disk using the SAP R/3 restore utility.

3. On the application system, verify that SAP R/3 is configured properly, as follows:

Move the original backint. Create a test script with the name backint in the directory with the SAP R/3 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 R/3 (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 R/3 is properly configured for backup using `backint`. Otherwise, you should reconfigure SAP R/3.

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 R/3 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 R/3 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.

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

If it fails, refer to the *HP OpenView Storage 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 R/3 backup utilities, Oracle Server Manager, or the TNS listener, then define the variables in the `Environment` section of the configuration file on the Cell Manager:

```
/etc/opt/omni/server/integ/config/SAP/<hostname>%<ORACLE_  
SID>
```

3. Verify the permissions of the SAP R/3 user on application system:

SAP R/3 user permissions must be set to enable you to perform an SAP R/3 backup or restore with Data Protector. Refer to “Configuring an SAP User in Data Protector” in the *HP OpenView Storage Data Protector Integration Guide* for information. Use the `testbar2` to check the permissions:

- Login in as an SAP R/3 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

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 R/3 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 R/3 ZDB configuration on the application system:

Execute the following command:

```
/opt/omni/lbin/util_sap.exe -CHKCONF <ORACLE_SID>
```

If an error occurs, use the following command to obtain an explanation:

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

where `<error_number>` is the number returned by the `RETVAL<error_number>` line reported by the `util_sap.exe` script.

2. Verify the SAP R/3 user.

Check that the respective user group has the See Private Objects user right selected. Also refer to “Configuring an SAP R/3 User in Data Protector” in the *HP OpenView Storage Data Protector Integration Guide* for more information.

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 R/3 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`;

SAPDBA also uses 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 R/3 user account?
- Call a support representative for assistance.

3. Verify the restore using Data Protector or CLI:

Refer to “Restoring an SAP R/3 Database” in the *HP OpenView Storage Data Protector Integration Guide* for more information.

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 R/3 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 R/3 user account
 - Call a support representative for assistance.
5. Verify the restore using `backint`:

`backint` is the same command used by the SAP R/3 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 R/3 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>`. Enter these variables in the Environment sublist of the Data Protector SAP R/3 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 the *HP OpenView Storage Data Protector Integration Guide* for more information on the Data Protector SAP R/3 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 R/3 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 the *HP OpenView Storage Data Protector Integration Guide* for more information on the Data Protector SAP R/3 configuration file.

- **When you start the backup, the configuration fails.**

On UNIX systems, check the output of the following command on the application system:

```
/opt/omni/sbin/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/sbin/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 R/3 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.

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.

Examples of SAP R/3 Database Restore

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

- “Example of Full Database Restore and Recovery” on page 242
- “Example of Partial Restore” on page 246
- “Example of Lost Files Restore” on page 247
- “Example of Archive Log Files Restore” on page 248

IMPORTANT

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

Preparing the SAP R/3 Database for Restore

If you are performing a full database restore, you need to know how the backup was performed; whether you have used the Oracle RMAN channels (during the non-ZDB sessions) or only BRBACKUP tools. If you used RMAN, use `svrmgr1` (Oracle8/8i) or `sqlplus` (Oracle9i), and RMAN commands to perform the restore. If you have used BRBACKUP utility, use SAPDBA to perform the restore.

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

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

- ORACLE_SID: system ID of the database instance
Example: P01

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

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

IMPORTANT

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

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

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

It is also the standard directory for export and unload dump files, if the parameter `exireo_dumpdir` in the profile `init<DBSID>.dba` is not set.

- **SAPTRACE:** directory for Oracle trace files and the alert file is `<SAPDATA_HOME>/saptrace` (UNIX systems) or `<SAPDATA_HOME>\saptrace` (Windows systems).
- **SAPDATA1:** directory of the database data files is by default `<SAPDATA_HOME>/sapdata1` (UNIX systems) or `<SAPDATA_HOME>\sapdata1` (Windows systems).

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

- `TWO_TASK`: identification of a remote database system
This environment variable must not be set.

Other optional environment variables that can be set:

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

Example of Full Database Restore and Recovery

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

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

Figure 2-30

Starting the SAPDBA in Expert Mode

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

ORACLE version: 8.0.5.0.0
ORACLE_SID      : ABA
ORACLE_HOME     : /app/oracle005/product
DATABASE        : open
SAPR3           : not connected

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

q - Quit

Please select ==> m

User and Security

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

q - Return

Please select ==> a
```

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

Figure 2-31

Selecting the Restore/Recovery Option

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

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

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

q - Quit

Please select ==> j
```

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

Figure 2-32

Selecting Full Restore and Recovery

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

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

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

h - Help
q - Return

Please select ==> b
```

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

Figure 2-33

Selecting the backup type and version for restore

```

a0y0C
-----
                                Full Restore and Recovery (2001-10-09)
-----
DATABASE STATE      : open
RESTORE / RECOVER: disallowed (see status)

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

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

S - Start restore and recover
q - Return

Please select ==> █
  
```

5. Afterwards, enter the full pathname name for the backup tool parameter file.
6. Select the Start restore and recover option to start the restore session.

Figure 2-34

Starting the Restore Session

```

a0y0C
-----
                                Full Restore and Recovery (2001-10-09)
-----
DATABASE STATE      : open
RESTORE / RECOVER: allowed

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

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

S - Start restore and recover
q - Return

Please select ==> $█
  
```

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

Figure 2-35

Selecting Return to restore process and continue option

```
a0y0c
Specify Restore Parameters for Backup Files

Selected bdgjp1a.anf 2001-10-08 14.51.06

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

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

Please select ==> q
```

IMPORTANT

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

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

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

Example of Partial Restore

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

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

- SAP tablespaces data file (PSAP<name>D/I)
- System tablespace files (SYSTEM)

- Rollback tablespace files (PSAPROLL)
- Temporary tablespace files (PSAPTEMP)

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

Example of Lost Files Restore

To restore the lost files, follow the steps below:

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

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

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

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

- Backup file: *<name> on <tape/disk>*
Backed up by <name of the external backup program>
This means that the file was backed up using the specific program. This comment appears when the parameter `backup_util_name` of the profile `init<DBSID>.dba` contains the name of the external backup program. Otherwise, the comment is displayed as, for example: `ext. backup utility`.
 - No restore of a backup file required
This means that the existing file can be used.
 - No backup file found
This means that no backup was found for this file in the specified period of time.
3. Select the Show the list of backup files option to specify the lost files for which you would like to see the available backup files. Each file that has been lost can have several backup files.

4. Select the `Select a backup file for restore` if you would like to change the proposed backup file, that should be restored. The file that is selected for the restore is flagged with `Selected for restore`.
5. Select the `Select a BRBACKUP run for restore` if you want to change the newest found backup file for each individual file from which the requested files can be restored. You can change this setting, for example, if all the files for restore were backed up in the same backup session and you want to specify only that backup session. The following information is listed:
 - Sequential number of the backup file found
 - Coded timestamp, date and time of the backup
 - The medium on which the backup was performed
 - The number of files found in this backup which are to be restored
6. Select the `Return` option to continue with the recovery process.
The lost files are restored using the SAP utility `BRRESTORE`.
7. Select the `Start restore of backup files`.

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

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

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

Example of Archive Log Files Restore

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

1. Select the `Restore archive files` option.

Archive log files are restored using the SAP `BRRESTORE` utility. If SAPDBA determines that the archiving directory

`<Oracle_home>/saparch` (UNIX systems) or

`<Oracle_home>\saparch` (Windows systems) does not have enough

space to restore all the necessary redo log files, the redo log files that have already been used will be deleted and the next required redo logs are restored during the subsequent recovery.

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

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

SAPDBA displays the following information on the screen:

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

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

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

3

Data Protector Microsoft SQL Server ZDB Integration

In This Chapter

This chapter explains how to configure and use the Data Protector Microsoft SQL Server ZDB integration.

The chapter is organized into the following sections:

“Introduction” on page 253

“Prerequisites and Limitations” on page 256

“Integration Concept” on page 258

“Data Protector Microsoft SQL Server Configuration File” on page 261

“Configuring the Integration” on page 263

“Backing Up Microsoft SQL Server Databases” on page 279

“Restoring a Microsoft SQL Server Database” on page 283

“Troubleshooting” on page 293

Introduction

The Data Protector Microsoft SQL Server ZDB integration provides a way to protect the Microsoft SQL Server data. The Data Protector zero downtime backup (ZDB) functionality offers a solution for eliminating performance degradation on the application system. During the backup, a Microsoft 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 that a Microsoft SQL Server snapshot is a Microsoft SQL Server related term and does not mean the same as a disk array snapshot.

Supported Disk Arrays

The following disk arrays can be used for ZDB of Microsoft SQL Server:

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

Advantages

Advantages of using the Data Protector Microsoft SQL Server ZDB integration are the following:

- The database runs on the application system with almost no performance degradation.
- The database is in snapshot mode only during the time it takes to create a replica.

Backup and Restore Types

Backup

All types of ZDB sessions (ZDB to tape, ZDB to disk, and ZDB to disk+tape) are possible with the Data Protector Microsoft SQL Server ZDB integration.

During a ZDB session, before the replica is created, the application system's database files are frozen and any transactions to them are cached, so that the database files are in a consistent state. The I/O to the

Introduction

database is suspended in this way only during the time it takes to create the replica, reducing disruption to the application system operation to a minimum.

When the replica is created, the application system database can return to the normal operation.

If a ZDB-to-tape session is being performed, the replica is then mounted on the backup system and the backup performed without impacting the application system further.

A separate transaction log backup to tape must be performed to enable rollforward recovery; transaction logs on disk cannot be used.

On EMC

You can perform ZDB (ZDB to tape only) on the following mirror types of the EMC configurations:

- SRDF
- Time Finder
- Combined SRDF + Time Finder

On XP

You can perform a ZDB session on the following mirror types of the XP configurations:

- BC
- CA
- Combined BC+CA

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 the instant recovery, the data in the specified replica (left unchanged for the purpose of instant recovery) is restored to the application system source volumes. Only the needed differential and transaction log backups are restored from the backup medium.

Table 3-1

Microsoft SQL Server Recovery Methods

	Recovery method
ZDB to tape	Standard restore

Table 3-1 Microsoft SQL Server Recovery Methods

	Recovery method
ZDB to disk	Instant recovery
ZDB to disk+tape	Standard restore, instant recovery

Prerequisites and Limitations

Prerequisites

- You need a license to use the Data Protector Microsoft SQL Server ZDB integration. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for information about licensing.
- Refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide* for an overview of ZDB concepts and terminology.
- Before you begin, make sure that you have correctly installed and configured the Microsoft SQL Server and Data Protector systems. Refer to the:
 - *HP OpenView Storage Data Protector Product Announcements, Software Notes, and References* for an up-to-date list of supported versions, platforms, devices, limitations, and other information.
 - *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions on how to install Data Protector on various architectures and how to install a Data Protector disk array integration (EMC, XP, VA, or EVA) with Microsoft SQL Server.
 - *SQL Server Books Online* for online information on Microsoft SQL Server.
- Microsoft SQL Server has to be installed on the application system. User databases *must* reside on the disk array source volumes, while system databases can be installed anywhere. However, if the system databases are also installed on the disk array, they *must* be installed on *different* source volumes than user databases.
- A Data Protector ZDB integration (EMC, XP, VA, or EVA) must be correctly installed and configured. For installation, refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*. For configuration, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide*.
- Do not use double quotes (" ") in object-specific pre-exec and post-exec commands.

Limitations

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

- If Microsoft SQL Server is installed on the backup system as well, its databases must be installed on source volumes that are different from the source volumes used for this integration. The drive letters or mount points assigned to those volumes must also be different from the drive letters or mount points assigned to the volumes used for this integration.
- It is not possible to *selectively* restore objects (Microsoft SQL Server databases) using instant recovery, but only the whole ZDB-to-disk or ZDB-to-disk+tape session that produced the replica.
- Preview is not possible for SQL ZDB and restore sessions.
- Object copying and object mirroring is not supported for ZDB to disk.

It is assumed that you are familiar with the Microsoft SQL Server database administration and the basic Data Protector ZDB functionality.

Integration Concept

Refer to the or *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide* for a general description of ZDB (split mirror or snapshot backup) and instant recovery concepts.

This section provides only the information relevant to the Data Protector Microsoft SQL Server ZDB integration.

Operations on source volumes (mounting a filesystem, creating a replica...) in the process described below are dependent on/triggered by ZDB options.

Microsoft SQL Server on the application system has to be running for a backup to start. Otherwise, data consistency cannot be guaranteed and the backup will not be performed.

The Data Protector Microsoft SQL Server ZDB integration backs up the Microsoft SQL Server database(s), which are replicated within a disk array.

ZDB Flow

The Data Protector Microsoft SQL Server ZDB integration backup flow can be described as follows:

- The Data Protector Session Manager starts the `sql_bar.exe` on the application system. The `sql_bar.exe` connects to the Microsoft SQL Server to find the locations of the database files to be backed up. The backup fails if the Microsoft SQL Server services are not running at this point.
- The `sql_bar.exe` initiates the resolving of backup objects on the application system and makes a Microsoft SQL Server snapshot of the Microsoft SQL Server databases on the application system. This ensures database consistency.
- A disk array agent initiates the creation of a replica.
- When a replica is created, the `sql_bar.exe` informs the Microsoft SQL Server on the application system that the Microsoft SQL Server snapshot is no longer needed—the database resumes normal operation.
- Data Protector assigns the same drive letter or mount point as on the application system to the target volumes on the backup system.

NOTE

In the case of a ZDB-to-disk session, at this point the remaining ZDB options are processed and the session finishes. The description that follows occurs only with a ZDB-to-tape and a ZDB-to-disk+tape sessions.

- The `sql_bar.exe` informs `sql_smb.exe` what is to be backed up. The `sql_smb.exe` passes database data to the General Media Agent for performing the backup of Microsoft SQL Server.

Restore Flow

The Data Protector Microsoft SQL Server integration restore flow is as follows:

- The objects and object versions which are to be restored are defined using the Data Protector GUI.

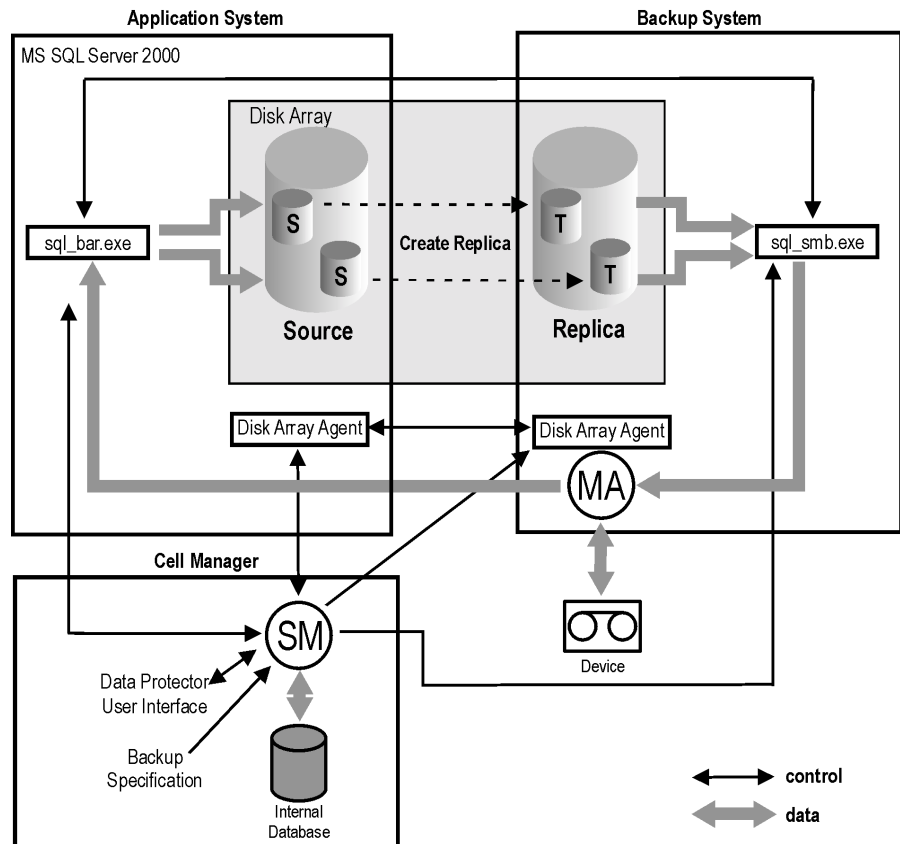
If a restore from backup media to the application system on LAN (standard restore) is being performed:

- A restore session is started by the Restore Session Manager (RSM), which starts `sql_bar.exe` and the Data Protector General Media Agents. The `sql_bar.exe` connects to Microsoft SQL Server and receives data from the General Media Agents. The Microsoft SQL Server then writes the data restored by Data Protector to the disks.

If an instant recovery is being performed:

- An instant recovery session is started and the files belonging to the database are restored. For detailed information on instant recovery, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide* and *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide*.
- `sql_bar.exe` puts the database in non-recovered state. Differential backups and transaction logs are restored from the tape and transactions are applied. At the end of the process, the SQL server recovers the database.

Figure 3-1 Backup and Restore Concept



Data Protector Microsoft SQL Server Configuration File

Data Protector stores the Microsoft SQL Server integration parameters for every configured Microsoft SQL Server in the `/etc/opt/omni/server/integ/config/MSSQL/<client_name>%<instance_name>` file (HP-UX and Solaris systems), or in the `<Data_Protector_home>\Config\Server\Integ\Config\MSSQL\<client_name>%<instance_name>` file (Windows systems) on the Cell Manager. The parameters stored are the user name and password for the Microsoft SQL Server user, who must have permissions to run backup and restore within Microsoft SQL Server (assuming the standard security is used during the configuration of the integration).

The configuration parameters are written to the Data Protector Microsoft SQL Server configuration file:

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

Syntax

The syntax of the Data Protector Microsoft SQL Server configuration file is as follows:

```
Login='<user>';  
Password='<encoded_password>';
```

IMPORTANT

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

Example

This is an example of the file:

- if standard security is used:

```
Login='TROLL\Administrator';  
Password='dsjf08m80fh43kdf';
```
- if integrated security is used:

```
Login='';
```

Data Protector Microsoft SQL Server ZDB Integration
Data Protector Microsoft SQL Server Configuration File

Password='dsjf08m80fh43kdf';

Configuring the Integration

The configuration of the Data Protector Microsoft SQL Server ZDB integration consists of the following:

1. “Configuring Microsoft SQL Server” on page 263
2. “Configuring a Microsoft SQL Server ZDB” on page 267

Before You Begin

It is recommended that you configure and run some test filesystem backups using Data Protector.

This includes installing the Disk Agent on the Microsoft SQL Server system. Any device can be used for this test. Configure a standard filesystem backup, which can include one directory only.

Thus, you can check whether the Microsoft SQL Server client system and the Data Protector Cell Manager are communicating properly.

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

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

Furthermore, it is recommended to run some test filesystem ZDB sessions. Configure a backup specification for ZDB and run the ZDB session as described in the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide*.

In case of any difficulties with the filesystem backups (non ZDB or ZDB), do not continue configuring the integration until you have solved the problems.

Configuring Microsoft SQL Server

The configuration is performed during the creation of the first backup specification for Microsoft SQL Server databases. For creating a backup specification, see “Configuring a Microsoft SQL Server ZDB” on page 267.

However, you can change the configuration any time after you have created at least one backup specification. For information on changing the configuration, see “Changing and Checking the Microsoft SQL Server Configuration” on page 266.

Prerequisites

- Microsoft SQL Server must be online during the configuration procedure.
- Configuration must be performed for every single Microsoft SQL Server system.

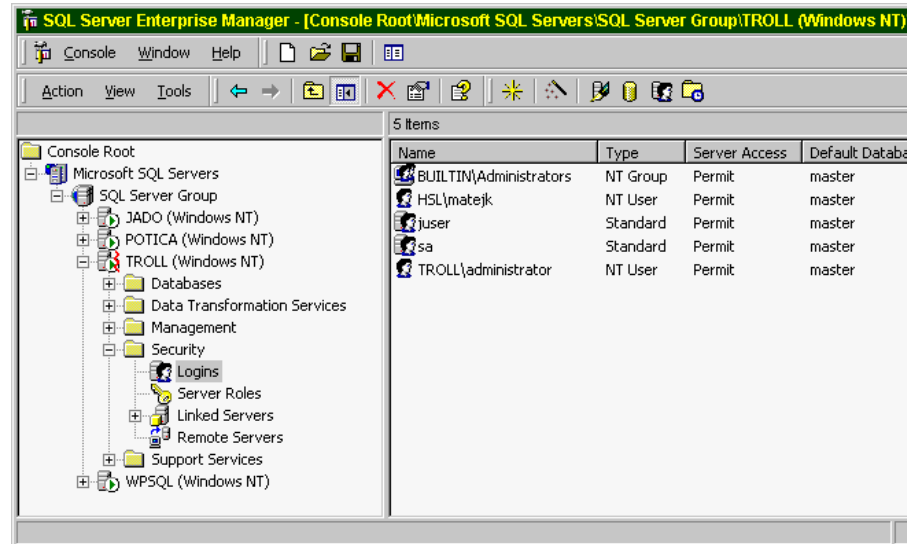
The configuration consists of setting the user name and password for the Data Protector services. Thereafter, the services are able to connect to the Microsoft SQL Server and operate under the specified account.

The user must have appropriate permissions to run backup and restore on the respective Microsoft SQL Server.

You can check this using the Microsoft SQL Server Enterprise Manager.

Figure 3-2

Microsoft SQL Server Users



You need to define the way in which the Data Protector sessions will run on the Microsoft SQL Server system, using either Data Protector Inet account (in most cases the system account) or a specified user account (preferred option).

Configuration Procedure

To configure the Microsoft SQL Server while creating the first backup specification or while changing the configuration, proceed as follows:

In the Configure MS SQL dialog box, select either Integrated Security or Standard Security. See Figure 3-3.

Figure 3-3 Configuring the Microsoft SQL Server



NOTE

It is recommended that the Microsoft SQL Server system administrator configures the Data Protector Microsoft SQL Server integration.

- If you use Standard Security, provide a user name in the format <DOMAIN>\<user_name> and a password for a Microsoft SQL Server user, who must have permissions to run backup and restore of the Microsoft SQL Server.
- If you use Integrated Security, the Data Protector SQL Server integration will use the Data Protector Inet account to connect to the Microsoft SQL Server.

See Microsoft SQL Server documentation for more detailed information about security and for a description of the two connection types.

Click OK to confirm the configuration.

What Happens?

The login information is written to the Data Protector Microsoft SQL Server configuration file on the Cell Manager:

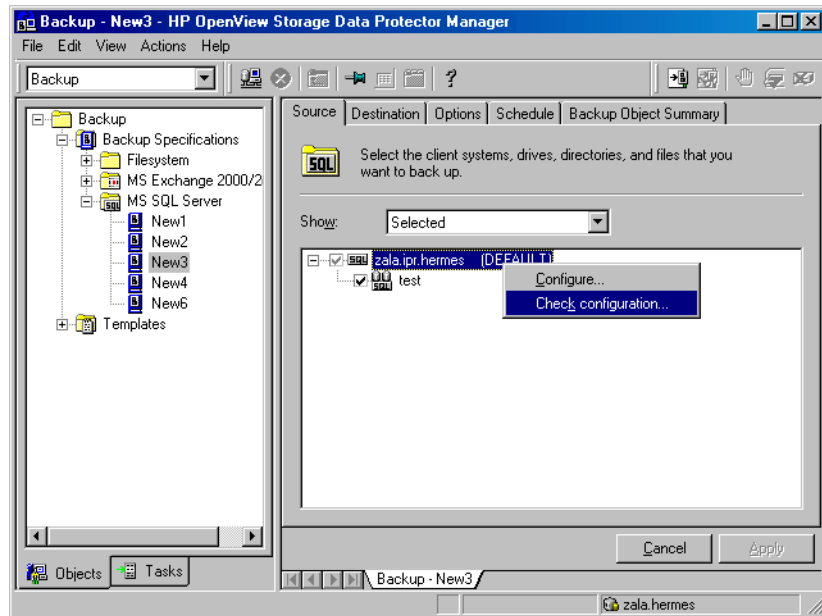
```
<Data_Protector_home>\Config\server\Integ\Config\MSSQL\<hostname>%<instance_name> (Windows Cell Manager) or  
/etc/opt/omni/server/integ/config/MSSQL/<hostname>%<instance_name> (UNIX Cell Manager).
```

Changing and Checking the Microsoft SQL Server Configuration

You can change the configuration of a specific Microsoft SQL Server and its instance any time after you have created at least one backup specification for this Microsoft SQL Server. To change the configuration, proceed as follows:

1. In the HP OpenView Storage Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, Backup Specifications, and then MS SQL Server. Click an existing backup specification for the Microsoft SQL Server for which you want to change the configuration.
3. In the Source property page, right-click the name of the Microsoft SQL Server and select Configure.
4. Configure the Microsoft SQL Server as described in “Configuring Microsoft SQL Server” on page 263.
5. Right-click the name of the Microsoft SQL Server and select Check Configuration. See Figure 3-4.

Figure 3-4 Checking Configuration



Once you start checking the configuration procedure, the Data Protector service reads the login information from the configuration file.

Configuring a Microsoft SQL Server ZDB

To configure a Microsoft SQL Server ZDB, perform the following steps:

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

For a ZDB to disk, you also need to configure a backup device (for example, a standalone file device), as you will have to select it while configuring a backup specification. Otherwise, you cannot create a backup specification for a ZDB to disk. For information on configuring a standalone device, refer to the online Help index keyword “standalone devices”.

2. Configure media pools and media for a backup. Refer to the online Help index keyword “creating media pools” for instructions.

3. Create a Data Protector Microsoft SQL Server ZDB backup specification.

Creating a ZDB Backup Specification

If you intend to perform instant recovery, create two different backup specifications: one for user databases and another for system databases.

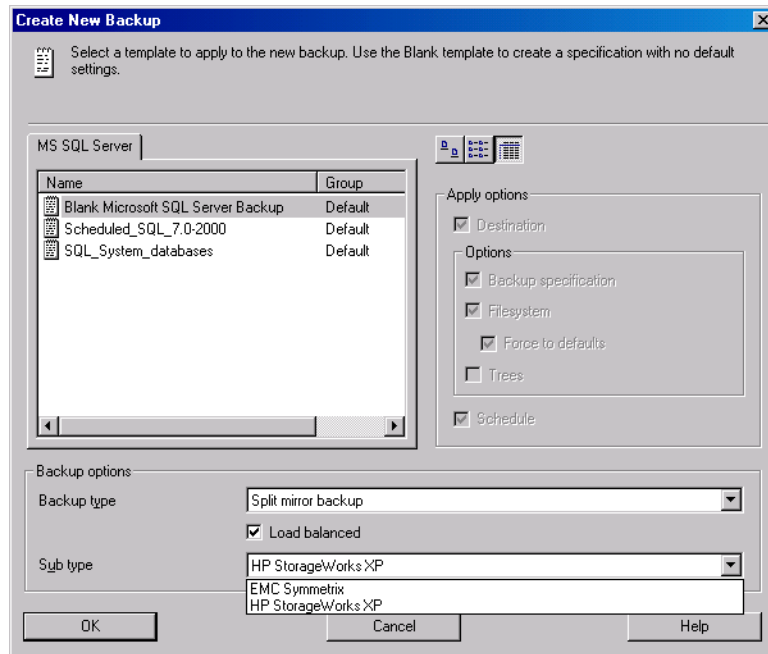
To create a Microsoft SQL Server ZDB backup specification, perform the following steps:

1. In the HP OpenView Storage Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, and then Backup Specifications. Right-click MS SQL Server and click Add Backup.
3. In the Create New Backup dialog box, select the Blank Microsoft SQL Server Backup 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 3-5.

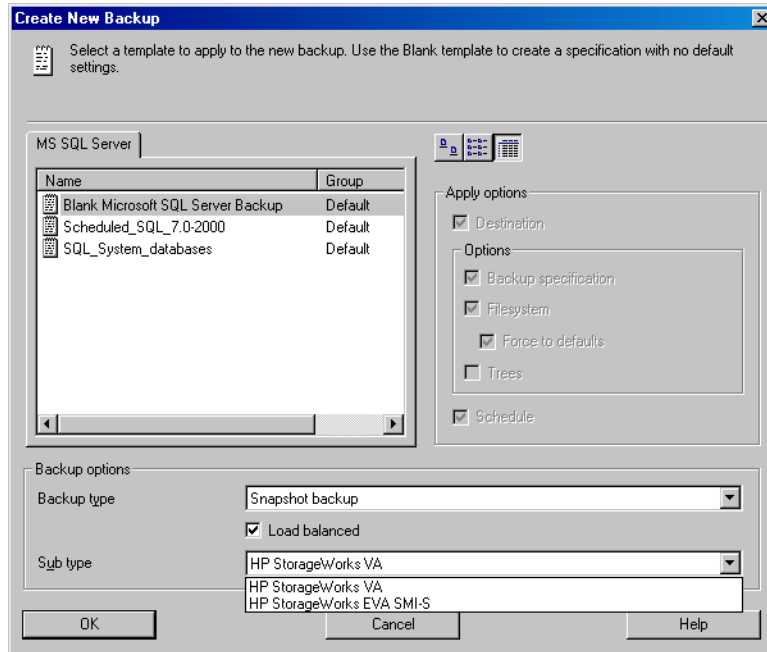
Figure 3-5 **Selecting a Blank Template and a Split Mirror Backup**



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, HP StorageWorks EVA (legacy), or HP StorageWorks EVA SMIS-S). See Figure 3-6.

Figure 3-6 Selecting a Blank Template and a Snapshot Backup



Click OK.

- Under Client systems, in the Application system drop-down list, select the system on which the Microsoft SQL Server database runs. If the application is cluster-aware, select the virtual server of the Microsoft SQL Server resource group.

In the Backup system drop-down list, select the backup system.

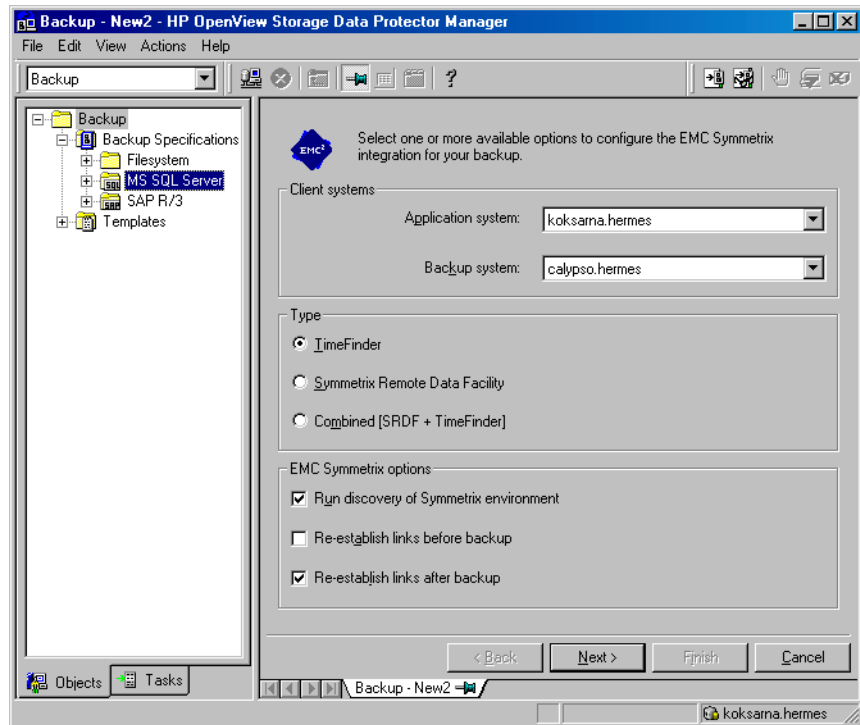
Select other disk array specific backup options (see Figure 3-8 for EMC, Figure 3-9 for XP, Figure 3-10 for VA, or Figure 3-10 for EVA backup options). For detailed information on the options, press **F1**.

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.

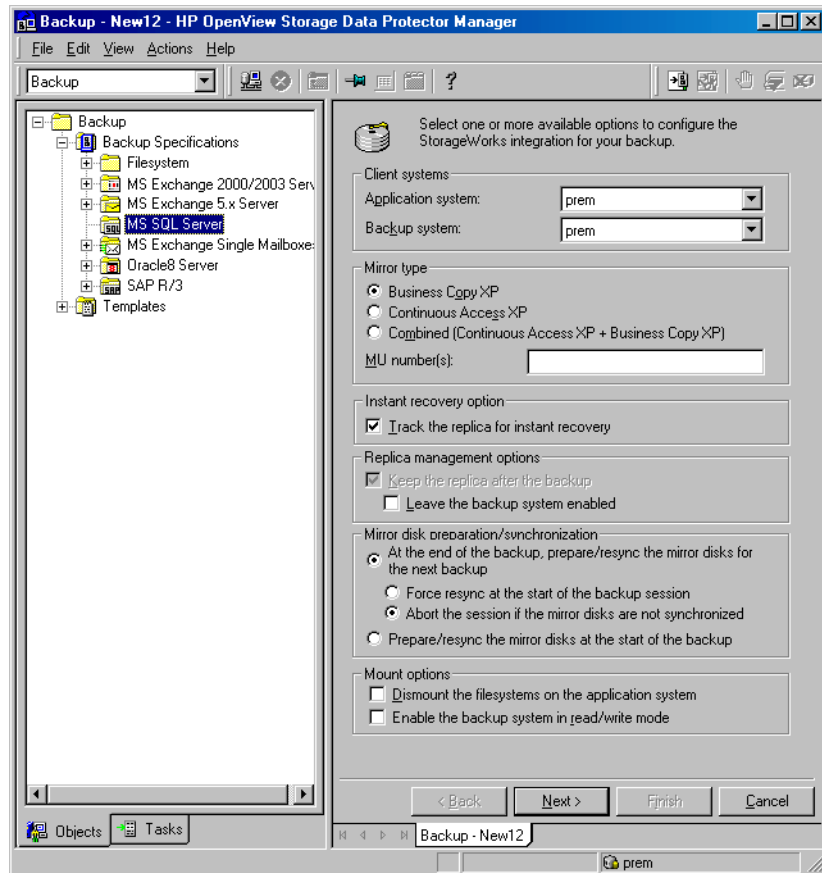
Figure 3-7 EMC Backup Options



On XP

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

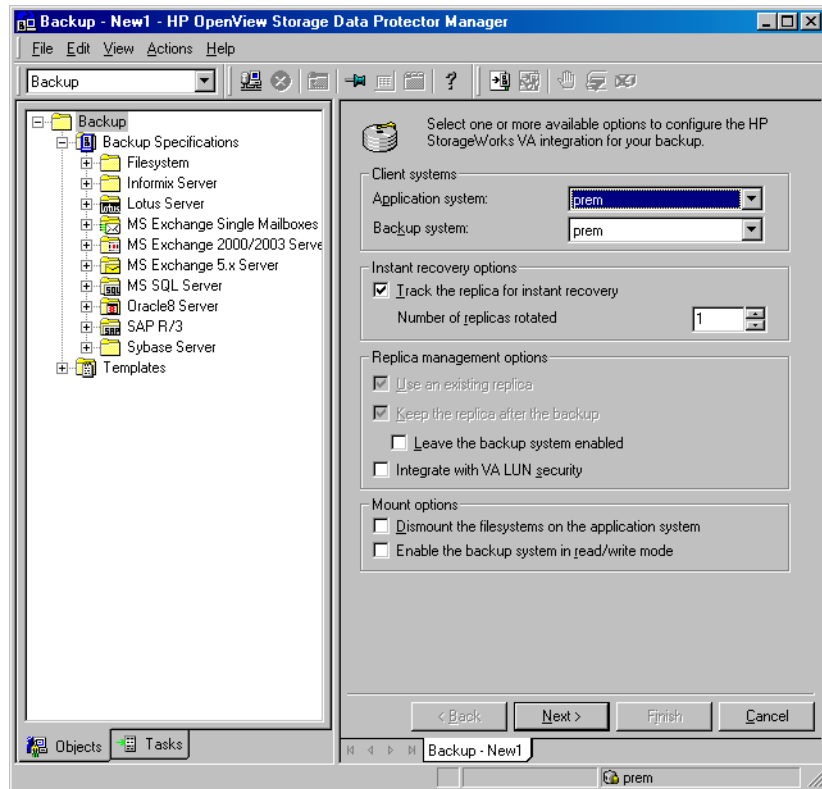
Figure 3-8 XP Backup Options



On VA

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

Figure 3-9 VA Backup Options

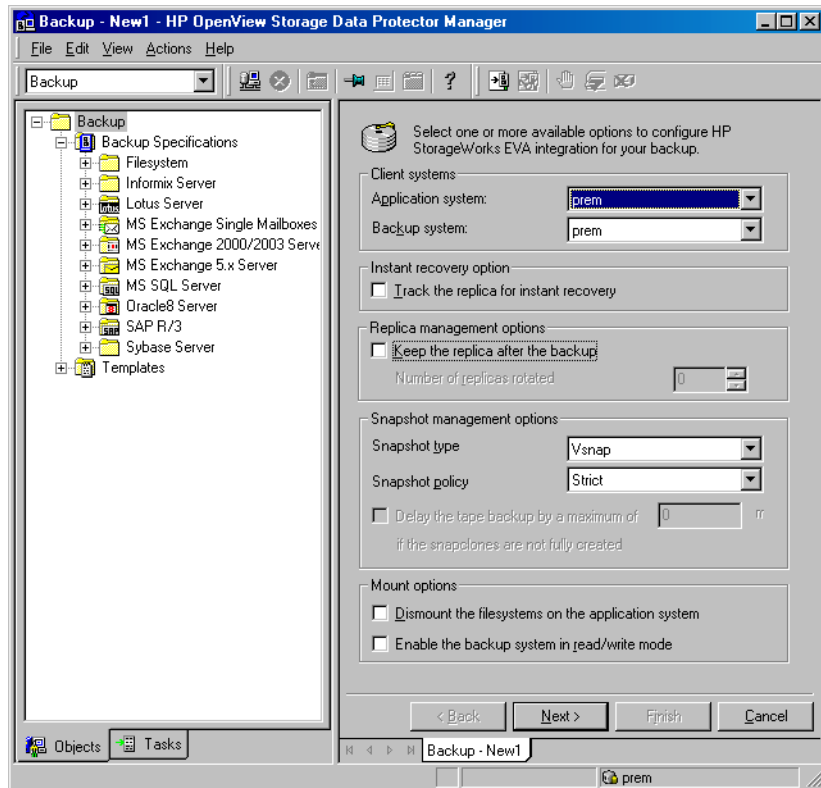


On EVA

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

Figure 3-10

EVA Backup Options



Click Next.

5. In the Application database drop-down list, leave the instance name for Microsoft SQL Server.

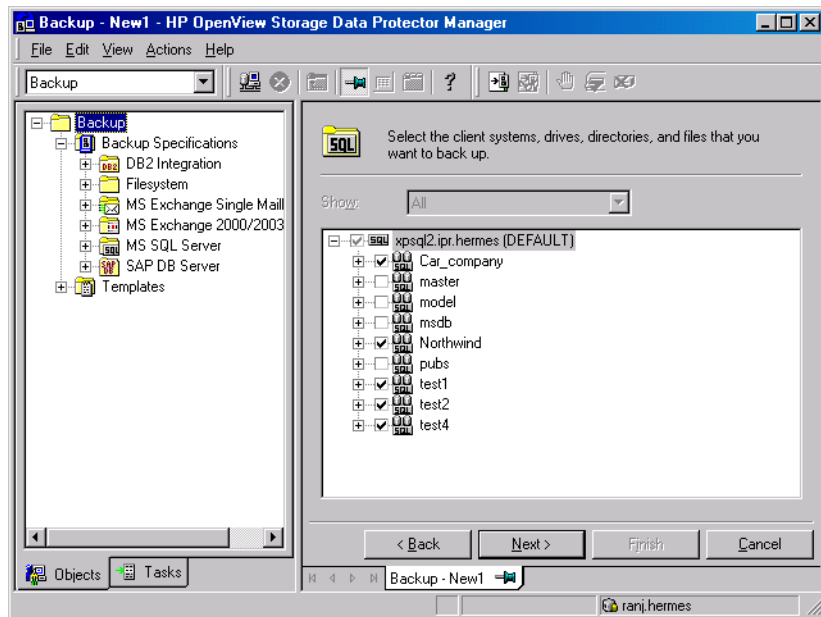
Click Next.

6. If the client has not been configured yet, the Configure Microsoft SQL dialog box appears. See “Configuring Microsoft SQL Server” on page 263 for detailed steps.
7. Select the Microsoft SQL Server databases you want to back up.

IMPORTANT

If you intend to perform instant recovery, do not select user databases and system databases in the same backup specification.

Figure 3-11 **Selecting User Databases**



Click Next.

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. The minimum number of devices required for mirroring Microsoft SQL Server integration objects equals the number of devices used for backup.

For detailed information on the object mirror functionality, see *HP OpenView Storage Data Protector Administrator's Guide*.

NOTE

Object mirroring is not supported for ZDB to disk.

Click Next.

9. Select the backup options.

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

For information on Application Specific Option, see “Microsoft SQL Server Specific Backup Options” on page 277 or online Help.

Click Next.


10. Optionally, schedule the backup. For information on scheduler, press **F1**.

Note that only Full backup type is performed.

11. Save the backup specification. It is recommended that you save all Microsoft SQL Server backup specifications in the MSSQL group. See the following figure.

Figure 3-12

Saving a Backup Specification



12. Once saved, the backup specification can be started by clicking Start Backup.

See “Backing Up Microsoft SQL Server Databases” on page 279 for information on starting a backup.

Microsoft SQL Server Specific Backup Options

The Microsoft SQL Server specific backup options are specified using the Data Protector GUI by clicking the Advanced tab in the Application Specific Options group box.

The following are the Microsoft SQL Server application specific backup options:

Concurrent streams This option is ignored for ZDB sessions.

Fast Direct Mode This option is ignored for ZDB sessions.

Check Database Integrity By selecting this option the structure of the MS SQL data is verified before the backup. In other words, the MS SQL data integrity validation is performed.

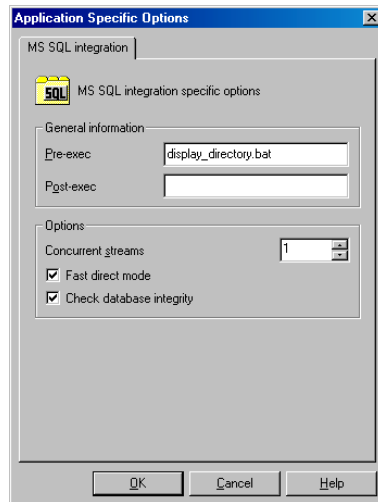
NOTE

If the check fails, the backup session is still completed (with warnings).

Pre-exec Specifies a command with arguments or a script that will be started on the Microsoft SQL Server before the backup starts. The command/script is started by the Data Protector `sql_bar.exe` and must reside in the `<Data_Protector_home>\bin` directory of the Microsoft SQL Server system. Only the filename must be provided in the backup specification.

Post-exec Specifies a command with arguments or a script that will be started on the Microsoft SQL Server after the backup. The command/script is started by the Data Protector `sql_bar.exe` and must reside in the `<Data_Protector_home>\bin` directory of the Microsoft SQL Server system. Only the filename must be provided in the backup specification.

Figure 3-13 **Application Specific Options**



Backing Up Microsoft SQL Server Databases

To run a ZDB-to-disk, ZDB-to-tape, or ZDB-to-disk+tape session for a Microsoft SQL Server database, use any of the following methods:

Backup Methods

- Schedule a backup of an existing Microsoft SQL Server ZDB backup specification using the Data Protector Scheduler.
- Start an interactive backup of an existing Microsoft SQL Server ZDB backup specification using the Data Protector GUI or the Data Protector CLI.

Consideration

It is not possible to start 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 or restore session; otherwise, the session will fail.

Scheduling a Backup

Scheduling a backup specification means setting 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, refer to the online Help index keyword “scheduled backups”.

To schedule a Microsoft SQL Server ZDB, proceed as follows:

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

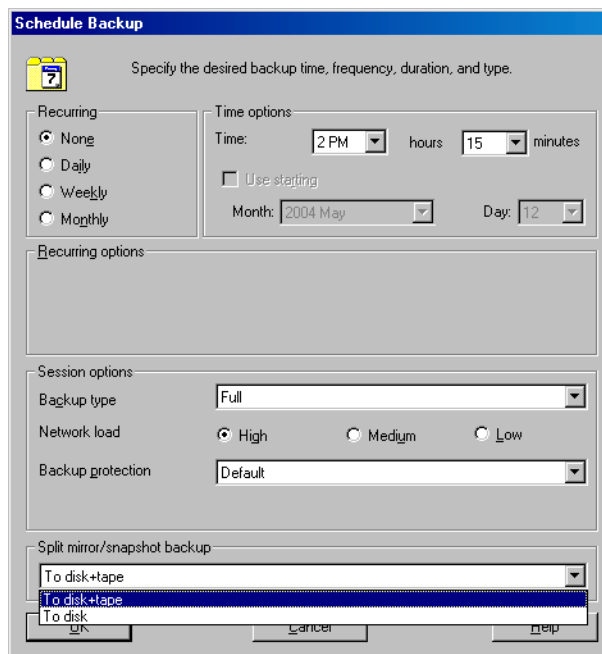
Note that the backup type is ignored for ZDB sessions. It is set to Full.

In the case of a ZDB-to-disk or a ZDB-to-disk+tape session, specify the Split mirror/snapshot backup option. See Figure 3-14.

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 option is not selected in the backup specification.

Figure 3-14 **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

Starting Backup Using the GUI

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

To start an interactive ZDB session of a Microsoft SQL Server database, proceed as follows:

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

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

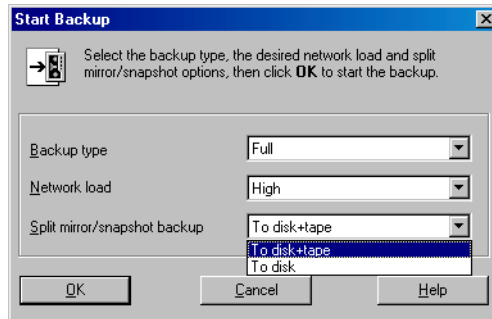
Note that the backup type is ignored for ZDB sessions. It is set to Full.

In the case of a ZDB-to-disk or a ZDB-to-disk+tape session, specify the Split mirror/snapshot backup option. See Figure 3-14.

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 option is not selected in the backup specification.

Figure 3-15 **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 a Microsoft SQL Server **ZDB-to-tape** or **ZDB-to-disk+tape** session using the Data Protector CLI, use the following command:

```
omnib -mssql_list <ListName>
```

To start a Microsoft SQL Server **ZDB-to-disk** session using the Data Protector CLI, use the following command:

```
omnib -mssql_list <ListName> -disk_only
```

where *<ListName>* is the name of the backup specification. For more information on the omnib command, refer to its man page.

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.

Restoring a Microsoft SQL Server Database

There are two ways of restoring a Microsoft SQL Server database:

- Restoring from backup media to the application system on LAN (standard restore)
- Restoring using the instant recovery functionality

Restoring from Backup Media to the Application System on LAN

When restoring from backup media to the application system on LAN, Microsoft SQL Server databases are restored online directly to the application system.

When restoring from backup media to the application system on LAN, you can restore Microsoft SQL Server databases to any Microsoft SQL Server with the same configuration as the original system within the Data Protector cell.

There is no need to create an empty database before restoring a database, because the database and its files are generated automatically by the Microsoft SQL Server.

If the database already exists and has a different structure, the restore will fail unless you select the `Force restore over existing database` option.

See “Restore Options” on page 287 for a detailed description.

General restore options that apply to all objects within a restore session, such as `Restore database to another Microsoft SQL Server` and `Restore using a different device` can be combined with the object-specific restore options, which are the following:

- `Point-in-time restore`
- `Recovery completion state`
- `Force restore over existing database`

In this way, you can choose among several restore scenarios.

Prerequisite

Before you start a restore session, verify that the database is not being used by any user.

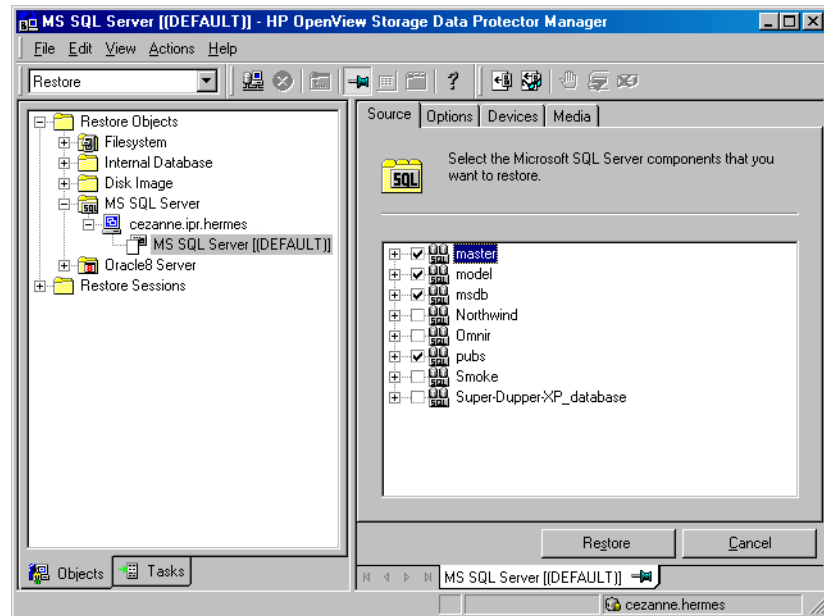
Restore Procedure

To restore the Microsoft SQL Server databases using the Data Protector GUI, proceed as follows:

1. In the HP OpenView Storage Data Protector Manager, switch to the Restore context.
2. In the Scoping Pane, expand Restore Objects, 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 backed up Microsoft SQL Server databases you want to restore. See Figure 3-16.

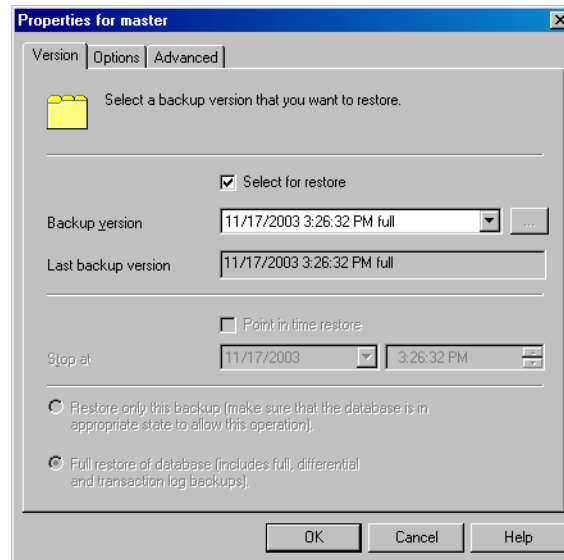
Figure 3-16

Restore Objects



To select the backup object specific options, right-click the object and select Properties.

Figure 3-17 **Selecting the Object Specific Options**



Here you can select the version (date of a backup) from which you want to restore and choose the Microsoft SQL Server specific restore options. See “Restore Options” on page 287 for details about these options. Click OK.

4. In the Options property page, specify whether you want to restore your data to another client or instance. In this case, you have to specify new locations for the databases you want to restore. See “Restore Options” on page 287.

NOTE


When you click the Options tab, Data Protector browses the cell for the running Microsoft SQL Server instances that can be selected as target instances for restore. If no instances are found, the Restore to another instance option is automatically 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 this action to restore the whole database.

- Restore and display file list only. Use this action if you do not know the original file names. In this case, the list of files backed up in a particular backup session is displayed.
- Restore and display headers only. Select this action if you need specific details about the database backup. The SQL Server's header information is displayed.

Figure 3-18 Restore Options



5. Click **Devices** and then **Media** to select the devices and media to be used for the restore.

Note that you can use a different device for the restore than the one used for the backup. Refer to the “Restoring Under Another Device” section in the *HP OpenView Storage Data Protector Administrator’s Guide* for more information on how to perform a restore using another device.

IMPORTANT

If the devices used for the restore are not those used for the backup, select the same number of devices in the `Devices` property page as you used when you backed up the databases.

6. Click `Restore MS SQL Server` and then `Next` to select the `Report level` and `Network load`.

Click `Finish` to start the restore session.

Restore Options

Backup Version

Specify the backup session from which the selected objects will be restored.

Point-in-Time Restore

Point-in-time restore means that a user can specify a point in time to which the database state must be restored. After recovery, the database is recovered in the state it was at the specified date and time.

Only transaction log records written before the specified date and time are applied to the database.

Point-in-time restore is specified by selecting a backup version and by setting the `Stop at` option.

Stop at

The `Stop at` option specifies the exact time when the rollforward of transactions will be stopped. Therefore, the backup you restore from must include transaction log backups so that the Microsoft SQL Server can recover the database to a particular point in time.

This option cannot be used with `NORECOVERY` or `STANDBY`. If you specify a `Stop at` time that is after the end of the `RESTORE LOG` operation, the database is left in a non-recovered state, just as if `RESTORE LOG` had been run with `NORECOVERY`.

Restore only this backup

Restoring a Microsoft SQL Server Database

If you have restored a version of the database and left it in a non-operational or standby state, then 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

The existing database residing on the target restore server system will be overwritten.

If a database with the same name as the one that you want to restore already exists on the server, and it has a different internal structure, then the Microsoft SQL Server does not let you rewrite the database without the Force Restore over existing database option turned on.

Recovery Completion State

These options let you select the state of the database after the recovery. You may select among the following:

- Leave the database operational. Once the last transaction log has been restored and the recovery has completed, the database is already operational.
- Leave the database non-operational after the last transaction log has been restored. You may further restore additional transaction logs one by one.
- Leave the database as read-only. You may further restore transaction logs before the database is set to read-write mode.

Restore database with a new name

This option lets you restore your database under a different name. You have to specify the database's logical file name and the destination file name (suboptions of the **Restore files to new locations** option) when selecting this option.

Restore files to new locations

This option allows you to restore files to a new location. You need to specify the database's logical file name, and a destination target file name for the specified logical file name. Use this option if you are restoring data to another server, instance or making a copy of the database on the same server.

Restoring to Another Microsoft SQL Server Instance or (and) to Another Microsoft SQL Server

To restore databases to a different Microsoft SQL Server system or (and) to a different Microsoft SQL Server instance, check the prerequisites below:

Prerequisites

- Both Microsoft SQL Servers must have the same local settings, such as code page and sort order. This information is displayed in the session monitor for each backup.
- The Target Microsoft SQL Server must be in the same Data Protector cell as the original Microsoft SQL Server and it must be configured.

Proceed as follows:

1. If the target Microsoft SQL Server is not yet configured, create a backup specification and configure the server.
See "Configuring a Microsoft SQL Server ZDB" on page 267.
2. Select the databases you want to restore and their versions.
3. Select whether you want to restore the data to another Microsoft SQL Server client or (and) to another Microsoft SQL Server instance:
 - To restore the data to another Microsoft SQL Server client, select the Restore to another client option and then select the target client from the drop-down list.

Restoring a Microsoft SQL Server Database

- To restore the data to another Microsoft SQL Server instance, select the `Restore to another instance` option. If you do not see the list of instances in the drop-down list, enter the instance name by yourself.
- To restore the data to another Microsoft SQL Server client and to another Microsoft SQL Server instance, make sure you entered the name of the instance that exists on the target client. Otherwise, the restore will fail.

Also, specify the new locations for the databases you want to restore.

4. Start restore.

See “Restoring a Microsoft SQL Server Database” on page 283

Instant Recovery

Refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide* and *HP OpenView Storage Data Protector Zero Downtime Backup Administrator’s Guide* for general information on instant recovery.

To recover the user and system databases, perform the following steps:

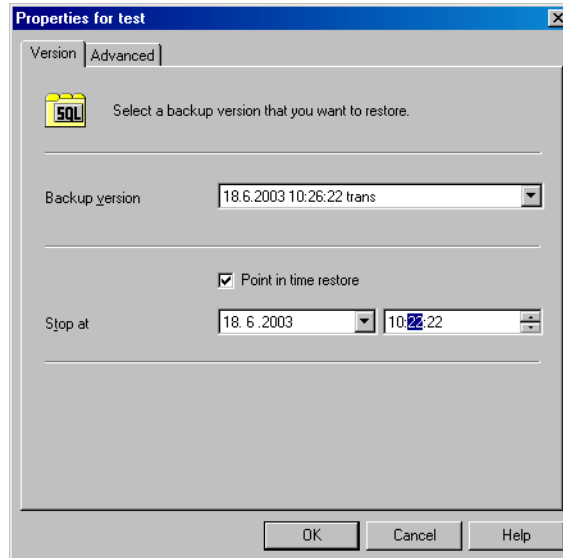
1. If you are restoring a user database, put the *database* offline: start the SQL Server Enterprise Manager, select the database, and click `Action`. Select all the tasks, and take them offline.

If you are restoring a system database, put the SQL *server* offline: start the SQL Server Enterprise Manager, right click the server you would like to restore and click `Stop`.

2. In the Data Protector GUI Context List, select `Instant Recovery`.
3. Expand `MS SQL Server` and select the backup session (replica) from which you want to perform the restore. By default, the database will be recovered until the last backed up transaction.
4. To recover the *user* databases to a specific point in time, perform the following steps:
 - a. In the `Source` property page, under `Restore Objects`, right click any database name and click `Properties`.
 - b. In the `Properties` dialog box, select the replica from which you want to perform the restore from the `Backup version` drop-down list. The latest version is selected by default.

Select Point in time restore. From the Stop at drop down lists, select the point in time to which the transactions should be applied and click OK. If there are no transaction logs available, this option is shaded. See Figure 3-19.

Figure 3-19 Point in Time Restore

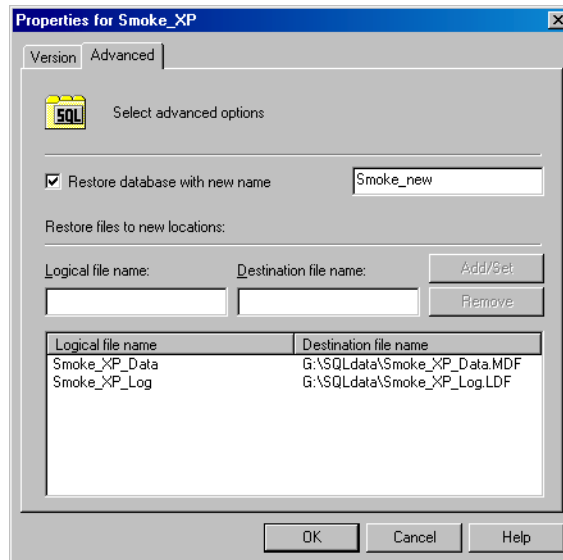


To restore the database under a different name, click on the Advanced tab and select the Restore database with new name option. See Figure 3-20.

IMPORTANT

If the logical filename and physical filename are not listed, add them to the list. Specify the same names that were used at the time of the ZDB session; otherwise, the instant recovery will fail.

Figure 3-20 Restoring the Database With a New Name



5. Click Restore MS SQL Server.
6. If you restored a system database, bring the SQL server online: in the SQL Server Enterprise Manager, right click the server you restored and click Start.

Troubleshooting

This section contains general checks and verifications and a list of problems you might encounter when using the Data Protector Microsoft SQL integration.

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

For general ZDB, restore, and instant recovery related troubleshooting, see the troubleshooting sections in the *HP OpenView Storage 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 OpenView Storage Data Protector Product Announcements, Software Notes, and References* for general Data Protector limitations, as well as known problems and workarounds.
- ✓ See http://www.openview.hp.com/products/datapro/spec_0001.html for an up-to-date list of supported versions, platforms, and other information.

Configuration Problems

If a configuration procedure does not work:

- Check that the Microsoft SQL Server services are running.
- Examine system errors reported in `<Data_Protector_home>\log\debug.log` on the Microsoft SQL Server, functioning as a Data Protector client.
- Connect to the Microsoft SQL Server via Microsoft SQL Server Enterprise Manager using the same login ID as the one specified in the Data Protector Configuration dialog box.
- Perform a backup of Microsoft SQL Server databases using the Microsoft SQL Server Enterprise Manager.

If the backup fails, fix any Microsoft SQL Server problems, and then perform a backup using Data Protector.

Using the Data Protector Command Line to Check Configuration

The configuration can also be performed from the command line. Enter the following string at the `<Data_Protector_home>\bin` directory on the Microsoft SQL Server system:

```
sql_bar config -dbuser:<dbuser> -password:<password>  
-appsrv:<appsrv> [-instance:<instance name>]
```

Enter the same information as if using the Data Protector GUI:

- The username and password of the SQL Server user who has permissions to back up and restore the SQL Server backup objects.
- The name of the SQL Server system.

To check configuration using the command line, enter the following string from the `<Data_Protector_home>\bin` directory on the Microsoft SQL Server computer:

```
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 the following output:

```
*RETVAl*8523
```

If you want the information about the existing configuration, enter the following string:

```
sql_bar getconf [-instance:<instance_name>]
```

If the optional parameter `-instance:<instance_name>` is not specified, the configuration for the default instance is returned.

What Happens?

Once you start the configuration from the command line, the login information will be written in the `<hostname>%<instance name>` file in the `<Data_Protector_home>\Config\Server\Integ\Config\MSSQL` directory.

Once you start checking the configuration procedure, the Data Protector service reads login information from the `<hostname>%<instance name>` file on the Data Protector Cell Manager and tries to connect to the server using this account.

Miscellaneous Problems

Problems

- The integration is properly configured and the backup of all databases fails after a timeout, with an error message similar to the following:

```
[Warning] From: OB2BAR@paradajz.hermes.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@paradajz.hermes.com "MSSQL70" Time:
3/14/00 8:19:24 PM
```

```
Received ABORT request from SM => aborting.
```

- The SQL Server's error log contains an entry similar to the following:

```
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.).
```

- The SQL Server's VDI.LOG file contains an entry similar to the following:

```
2000/03/15 13:19:31 pid(2112)
Error at BuildSecurityAttributes: SetSecurityDescriptorDacl
Status Code: 1338, x53A Explanation: The security descriptor
structure is invalid.
```

Cause

The SQL Server service and the Data Protector Inet service are running under different accounts. The SQL Server integration cannot access the SQL Server's data for backup due to security problems.

Solution

Restart the Data Protector Inet service under the same account as the SQL Server service is running.

Backup Problems

1. If a backup does not work:
 - Verify the configuration file to check if the Cell Manager is correctly set on the Microsoft SQL Server that is functioning as a Data Protector client.
 - Check that the Microsoft SQL Server services are running.
 - Check that `sql_bar.exe` is installed on the system.
 - Examine system errors reported in `<Data_Protector_home>\log\debug.log` on the Microsoft SQL Server system.

Check also the `errorlog` and `VDI.log` files in the `<MSSQL>\log` directory on the server system.
 - Perform a backup of Microsoft SQL Server databases using the Microsoft SQL Server Enterprise Manager.

If the backup fails, fix the Microsoft SQL Server problems and perform a backup using Data Protector.
2. If during the creation of a backup specification you do not see the instance of the Microsoft SQL Server as the application database, enter the instance name by yourself. When the “not-named instance” is not displayed, the `DEFAULT` string must be inserted as an application database.
3. When performing a backup, Microsoft SQL Server reports that the database backup cannot take place because of inappropriate user rights.

If the Data Protector Manager (and `sql_bar.exe`) reports that the integration is properly configured, verify that the Microsoft SQL Server user has appropriate rights to access the databases that cannot be backed up.

It is recommended that the Microsoft SQL Server system administrator (`sa`) configure the Data Protector Microsoft SQL Server integration.

Restore Problems

- If a restore does not work:

- Check whether a filesystem backup of the problematic client works. It is much easier to troubleshoot a filesystem backup.
- Check that the Microsoft SQL Server services are running.
- Check that `sql_bar.exe` is installed on the system.
- Examine the system errors reported in
`<Data_Protector_home>\log\debug.log` on the Microsoft SQL Server that is functioning as a Data Protector client.

Check also the `errorlog` and `VDI.log` files in the `<MSSQL>\log` directory on the same system.

- The following error has occurred when executing an SQL statement:

```
Error message: "Microsoft SQL-DMO (ODBC SQLState:
01000)?15[152:5] 1646 [Microsoft] [ODBC SQL Serevr 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."
```

It is an expected behavior when the master database is restored in a single user mode, so this message should not be treated as an error.

A Restore from an Object Copy Hangs.

Problem

A restore from an object copy hangs.

Action

Before restarting the restore:

- Increase the number of Disk Agent buffers for the device used for the restore.
- If all objects of the backup are recorded in the IDB, perform the following steps:
 1. In the Internal Database context of the Data Protector GUI, search for all objects belonging to the same backup. The objects are identified by the same backup ID.
 2. Copy each object in a separate object copy session to a separate device, for example a file library. For each object, use a separate medium with the non-appendable media policy.
 3. Set the highest media location priority for the newly created copies.

Database is Left in Unrecovered State After the “Invalid value specified for STOPAT parameter” Message is Reported in the Data Protector Monitor

Problem If the Invalid value specified for STOPAT parameter Message is Reported in the Data Protector Monitor, the database remains in the unrecovered state as if the RESTORE LOG operation was run with the Leave the database non-operational option.

Action The database can be recovered to the latest point in time by using the Microsoft SQL Query Analyzer. To recover the database, run the following T-SQL command:

```
RESTORE DATABASE <database_name> WITH RECOVERY
```

After the database is recovered, additional transaction logs cannot be applied.

Database Restore Fails with: “Error Has Occurred While Executing SQL Statement”

Problem 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 the SQL Server Enterprise Manager, go to the database and click Action. Select all the tasks, and then put them offline.

Transaction Logs Cannot be Restored From Tape

Problem After the recovery session completed successfully and the database was put in norecovery state, the transaction logs cannot be restored from tape.

Action The database can be recovered to the state of ZDB to disk by using the Microsoft SQL Query Analyzer.

To recover the database, run the following T-SQL command:

```
RESTORE DATABASE <database name> WITH RECOVERY
```

After the database is recovered, additional transaction logs cannot be applied.

Restoring to Another Client

Problem You want to perform a restore of the Microsoft SQL Server database to another client in the Data Protector cell not configured to use with the Microsoft SQL Server, but the restore does not work.

Action Create the configuration file by configuring the Microsoft SQL integration on this client. See “Configuring the Integration” on page 263.

Database Left in Unrecovered State After the Restore Session Completed Successfully

Problem If you set the time for the `Stop at restore` option beyond the end of the `RESTORE LOG` operation, the database remains in the unrecovered state as if the `RESTORE LOG` operation was run with the `Leave the database non-operational` option.

Action The database can be recovered to the latest point in time by using the Microsoft SQL Query Analyzer. To recover the database, run the following T-SQL command:

```
RESTORE DATABASE <database_name> WITH RECOVERY
```

After the database is recovered, additional transaction logs cannot be applied.

Before You Call Support

If you have performed all the troubleshooting procedures without solving your problem, you should gather the following information for Data Protector support before you make a call:

1. Provide details about your hardware and software configuration, including the official patches you use, the Microsoft SQL Server version, the SP, the Windows version and the SP.
2. Provide a detailed description of the action you failed to perform. If you had backup problems, attach the backup specification.
3. Provide the information from the following files:
 - `<Data_Protector_home>\log\debug.log`
 - `<MSSQL>\log\errorlog`
 - `<MSSQL>\log\vdi.log`

Copy the session output into a file.

4

**Data Protector Microsoft
Exchange Server ZDB
Integration**

In This Chapter

This chapter explains how to configure and use the Data Protector Microsoft Exchange ZDB integration.

The chapter is organized into the following sections:

“Introduction” on page 303

“Prerequisites and Limitations” on page 306

“Integration Concept” on page 310

“Configuring a Microsoft Exchange ZDB Backup Specification” on page 312

“Backing Up Microsoft Exchange Server” on page 332

“Restoring a Microsoft Exchange Database” on page 336

“Troubleshooting” on page 350

Introduction

The Data Protector Microsoft Exchange ZDB integration provides a way to protect the Microsoft Exchange Server data. The zero downtime backup (ZDB) concept offers a solution for minimal performance degradation of Microsoft Exchange Server.

Supported Disk Arrays

The following disk arrays can be used for ZDB of Microsoft Exchange Server:

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

Advantages

Advantages of using the Data Protector Microsoft Exchange Server ZDB integration are the following:

- The database runs on the application system with almost no performance degradation.
- The database is stopped (offline backup) only for the time it takes to create a **replica** (split the mirror disks or create snapshots).

Backup and Restore Types

Backup

All types of ZDB sessions (ZDB to tape, ZDB to disk, and ZDB to disk+tape) are possible with the Data Protector Microsoft Exchange ZDB integration.

During a ZDB session, the database is stopped on the application system only for the few minutes it takes to create a replica. In the case of a ZDB-to-tape session, backup is subsequently performed offline on the backup system.

During the Microsoft Exchange database ZDB-to-tape session, only the database files are backed up from the backup system and the transaction logs are not deleted. Transaction logs can be backed up in a separate backup session. Transaction logs are backed up from the application system. The `omniEx2000.exe` command purges all *backed up* transaction log files up to the checkpoint log file.

Introduction

The Data Protector Microsoft Exchange ZDB integration is based on the `ese_bar.exe` command for retrieving location of files and folders pertaining to a storage group or a store, and `omniEx2000.exe` command for mounting and dismounting databases and purging the transaction log files after the transaction log files backup.

On XP

You can perform an offline ZDB session on the following mirror types of the XP configurations:

- BC
- CA
- Combined BC+CA

Restore

Restore is performed *offline*, using Data Protector. The Microsoft Exchange backup object(s) are stored in the Data Protector database as filesystem objects; therefore, restore can be performed following the standard Data Protector filesystem restore procedure from backup media to the application system on LAN.

The Microsoft Exchange database, installed on a disk array, can also be restored using the instant recovery functionality. During the instant recovery, the data in the specified replica (left unchanged for the purpose of instant recovery) is restored to the application system source volumes without restoring from a backup medium.

Table 4-1 on page 304 provides an overview of recovery methods:

Table 4-1**Microsoft Exchange Recovery Methods**

	Roll forward recovery	Point-in-time recovery
ZDB to tape	Standard restore (database) + standard restore (transaction log files)	Standard restore (database)
ZDB to disk	Instant recovery (database) + standard restore (transaction log files)	Instant recovery (database)

Table 4-1 Microsoft Exchange Recovery Methods

	Roll forward recovery	Point-in-time recovery
ZDB to disk+tape	<ul style="list-style-type: none"> • Instant recovery (database) + standard restore (transaction log files) • Standard restore (database) + standard restore (transaction log files) 	<ul style="list-style-type: none"> • Instant recovery (database) • Standard restore (database)

There are some additional considerations that must be taken into account with the recovery methods described in the Table 4-1 on page 304:

- The instant recovery functionality restores data from a replica on the backup system to the source volumes on the application system. Therefore it is, using instant recovery, not possible to *selectively* restore objects (storage groups, stores or Microsoft Exchange server) if they do not reside on separate source volumes.
- Transaction logs must be backed up to be able to perform the roll forward operation.

NOTE

Backup of log files can be performed only if circular logging option is disabled for the involved on Microsoft 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.

Prerequisites and Limitations

Prerequisites

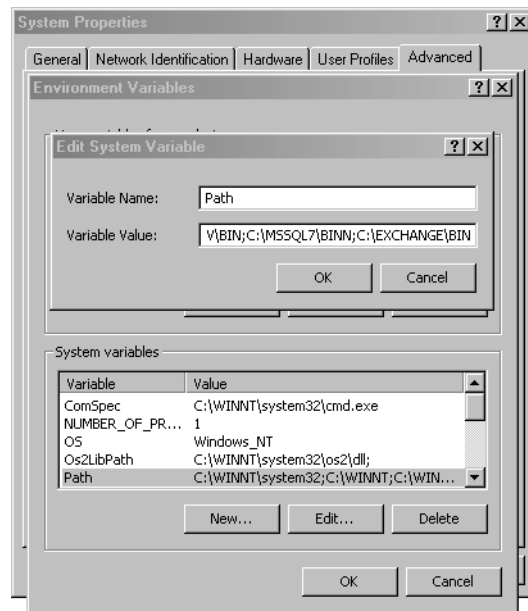
- You need a license to use the Data Protector Microsoft Exchange ZDB integration. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for information about licensing.
- Refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide* for the ZDB and instant recovery concepts and terminology.
- Before you begin, make sure that you have correctly installed and configured the Microsoft Exchange and Data Protector systems. Refer to the:
 - *HP OpenView Storage Data Protector Product Announcements, Software Notes, and References* for an up-to-date list of supported versions, platforms, devices, limitations, and other information.
 - *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions on how to install Data Protector on various architectures and how to install a Data Protector disk array integration (XP, VA, or EVA) with Microsoft Exchange Server.
 - *Microsoft Exchange Server Books Online* for online information on Microsoft Exchange Server.
- A Data Protector ZDB integration (XP, VA, or EVA) must be correctly installed and configured. For installation, refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*. For configuration, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide*.
- The Microsoft Exchange Server has to be installed on the application system. All parts of the Microsoft Exchange database (Information Store (MIS), the Key Management Service (KMS), and the Site Replication Service (SRS)) have to be installed on the disk array source volumes.
- Backup of log files can be performed only if the circular logging option is disabled for the involved storage groups on the Microsoft Exchange Server. Circular logging is a Microsoft Exchange Server mode, where transaction logs are automatically overwritten when the data they contain is committed to the database.

- When backing up Microsoft Exchange log files, either the Log all or the Log files logging level must be selected in the Advanced Filesystem Options dialog box. This enables purging backed up log files from disk. For more information on advanced filesystem options, see online Help index: “filesystem options”.
- Before any operation is performed, add the <Exchange_home>\bin directory to the Windows Path environment variable.
- Add the <Exchange_home>\bin directory to the Windows Path environment variable on all cluster nodes as follows:
 1. In the Microsoft Windows Explorer, right-click My Computer and click Properties.
 2. In the Properties dialog box, click Advanced and then Environment Variables.
 3. In the Environment Variables dialog box, select Path in the System Variables list and click Edit.
 4. Add <Exchange_home>\bin in the Variable Value text box and click OK.

**Cluster-Aware
Clients
Prerequisites**

See Figure 4-1 on page 308.

Figure 4-1 Path System Variable



- In a cluster environment, prior to configuring a ZDB backup specification for the Microsoft Exchange Server database and transaction log files, a new physical resource (LUN in case of an HP StorageWorks Virtual Array, virtual disk in case of an HP StorageWorks Enterprise Virtual Array, or LDEV in case of an HP StorageWorks Disk Array XP) must be created for the cluster, and the Exchange Server database and transaction log files must be moved to it.

For information on how to move the Exchange Server database and transaction log files, refer to the following online 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)*

Limitations

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

- Microsoft Exchange backup/restore API does not have the snapshot capability for suspending database I/O during the creation of a replica. The only way to perform a ZDB is to stop Microsoft Exchange database before a replica creation to ensure data consistency on disks. All updated pages in memory are written to the database files during the dismount process.
- Page level integrity check of Microsoft Exchange database files is not supported on EVA.
- Preview is not possible for Exchange Server ZDB and restore sessions.
- Object copying and object mirroring is not supported for ZDB to disk.

It is assumed that you are familiar with the Microsoft Exchange database administration and the basic Data Protector ZDB functionality.

Integration Concept

Refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide* for a general description of ZDB (split mirror or snapshot backup) and instant recovery concepts.

This section provides only the information relevant to the Microsoft Exchange integration.

Operations on source volumes (mounting a filesystem, creating a replica...) in the process described below are dependent on/triggered by ZDB options. See Table 4-3 on page 315 for the XP options, Table 4-4 on page 318 for the VA options, or Table 4-5 on page 320 for the EVA options.

The Data Protector Microsoft Exchange ZDB integration backs up the Microsoft Information Store (MIS), the Key Management Service (KMS), and the Site Replication Service (SRS), which are replicated within a disk array.

The Microsoft Exchange Server on the application system has to be running for a backup to start. Otherwise, data consistency cannot be guaranteed.

The omnicreated1 Command

The concept of backing up a Microsoft Exchange database installed on a disk array is based on the `omnicreated1` command. The `omnicreated1` command creates a ZDB filesystem backup specification file (`datalist`). The Microsoft Exchange ZDB backup specification, created using the `omnicreated1` command, includes the `Stop/quiesce` the application and `Restart` the application scripts for mounting and dismounting databases that are to be backed up and checking their consistency, and the `post-exec` script for purging the log files after the log file backup. The backup specifications, created using the `omnicreated1` command, can be modified and backup can be started using the Data Protector GUI.

ZDB Flow

The Data Protector Microsoft Exchange Server ZDB integration backup flow can be described as follows:

- The `omnicreated1` command creates a Microsoft Exchange ZDB backup specification file. The backup specification also contains the `Stop/quiesce` the application and `Restart` the application scripts (`omniEx200.exe`).

- The `ese_bar.exe` command is used to resolve the database objects and to retrieve a list of files and folders to be backed up.
- Before creating a replica, the `Stop/quiesce` application script is executed, which dismounts the databases to be backed up and verifies their consistency using the `esefile.exe` (Exchange 2000) or the `eseutil.exe` (Exchange 2003) utility. See “Checking Microsoft Exchange Files for Consistency” on page 329.
- When the replica is created, the `Restart` application script is executed, which mounts the dismounted databases.

NOTE

In the case of a ZDB-to-disk session, at this point the remaining ZDB options are processed and the session finishes. The description that follows occurs only with a ZDB-to-tape and a ZDB-to-disk+tape sessions.

-
- Data Protector backs up the Microsoft Exchange databases on the backup system.

The Microsoft Exchange database on the application system is stopped only while the replica is created.

NOTE

Transaction logs are backed up from the application system to a backup device, using the non-ZDB procedure.

Configuring a Microsoft Exchange ZDB Backup Specification

Backup specification for backing up a Microsoft Exchange database, installed on a disk array, is created using the `omnicreatedl` command. Based on this backup specification, you can perform a filesystem ZDB session using the Data Protector GUI.

The `omnicreatedl` command creates a Microsoft Exchange ZDB backup specification (datalist) with included the `Stop/quiesce` the application and the `Restart` the application scripts (`omniEx2000.exe`) for dismounting/mounting backed up databases and checking their consistency.

Backing Up Transaction Logs

Additionally, the `omnicreatedl` command creates a Microsoft Exchange transaction logs backup specification (datalist) for each storage group specified in a Microsoft Exchange ZDB backup specification (datalist), provided they have circular logging disabled. See Figure 4-18 on page 351. The Microsoft Exchange transaction logs backup specification (datalist) includes the `post-exec` script for purging backed up log files.

TIP

The Microsoft Exchange transaction logs backup specification can be either triggered by the `Post-Exec` (configured on the level of backup specification) option in the backup specification for the backup of database files (recommended), or started manually after the backup specification for the backup of database files has been started. For more information on configuring the `Pre-Exec` and `Post-Exec` commands, see online Help index: “pre- and post-exec commands”.

Creating a Microsoft Exchange ZDB Backup Specification

Using the `omnicreated1` command

A Microsoft Exchange ZDB backup specification (datalist) and Microsoft Exchange transaction logs backup specification are configured by running the `omnicreated1` command on any system in the cell. The following is the synopsis for the `omnicreated1` command:

```
omnicreated1 -ex2000 -datalist <Name> [-device <Name>] {  
DISK_ARRAY_XP_OPTIONS | VIRTUAL_ARRAY_OPTIONS |  
ENTERPRISE_VIRTUAL_ARRAY_OPTIONS } EXCHANGE_OPTIONS  
[-force] [-virtualSrv <Name>]
```

where:

`DISK_ARRAY_XP_OPTIONS` are:

1. For the Business Copy XP (BC) configuration and the ZDB-to-disk or ZDB-to-disk+tape session:

```
-split_mirror -sse -local <app_sys> <bck_sys>  
[-mirrors <MU_Numbers>] -instant_restore  
[-leave_enabled_bs] [ -split | -establish ]
```

2. For the Business Copy XP (BC) configuration and the ZDB-to-tape session:

```
-split_mirror -sse -local <app_sys> <bck_sys>  
[-mirrors <MU_Numbers>] [-keep_version  
[-leave_enabled_bs] ] [ -split | -establish ]
```

3. For the Continuous Access XP (CA) configuration or the Combined (BC+CA) configuration and the ZDB-to-tape session:

```
-split_mirror -sse { -remote <app_sys> <bck_sys> |  
-combined <app_sys> <bck_sys> } [-keep_version  
[-leave_enabled_bs] ] [ -split | -establish ]
```

`VIRTUAL_ARRAY_OPTIONS` are:

1. For the ZDB-to-disk or ZDB-to-disk+tape session:

```
-snapshot -va <app_sys> <bck_sys> -instant_recovery  
[-snapshots <number>] [-leave_enabled_bs]  
[-lun_security]
```

2. For the ZDB-to-tape session:

```
-snapshot -va <app_sys> <bck_sys>
```

Data Protector Microsoft Exchange Server ZDB Integration

Configuring a Microsoft Exchange ZDB Backup Specification

```
[-use_existing_snapshot] [-leave_version  
[-leave_enabled_bs] ] [-lun_security]
```

ENTERPRISE_VIRTUAL_ARRAY_OPTIONS are:

1. For the ZDB-to-disk session:

```
-snapshot { -eva | -smis } <app_sys> <bck_sys>  
-instant_recovery [-snapshots <number>]
```

2. For the ZDB-to-disk+tape session:

```
-snapshot { -eva | -smis } <app_sys> <bck_sys>  
-instant_recovery [-snapshots <number>] [-wait_cloncopy  
<number>]
```

3. For the ZDB-to-tape session:

```
-snapshot { -eva | -smis } <app_sys> <bck_sys>  
-snapshot_type { standard | vsnap | clone  
[-wait_cloncopy <number>] } -snapshot_policy { strict  
| loose } [-snapshots <number>]
```

NOTE

The `-eva` parameter needs to be used when you have the HP StorageWorks EVA Agent (legacy) installed on your application and backup systems. If you have the HP StorageWorks EVA SMI-S Agent installed, run the `omnicreated1` command with the `-smis` parameter.

EXCHANGE_OPTIONS are:

```
-annotation { MIS | SRS | KMS }  
{ -all_storage_groups | -storage_group  
<Storage_Group_Name1>[ -store <Store1>[ <Store2>... ] ] [  
-storage_group <Storage_Group_Name2>[ -store  
<Store1>[<Store2>... ] ]... ] }
```

For parameters and their descriptions see Table 4-2 on page 4-315 through Table 4-6 on page 4-323.

IMPORTANT

If the parameters contain spaces, use double quotes when specifying them in the `omnicreatedl` command. For example, `-storage_group "First Storage Group"`.

Table 4-2 **General Options**

Parameter	Description
<code>-ex2000</code>	Instructs the <code>omnicreatedl</code> command to create a Microsoft Exchange ZDB backup specification and Microsoft Exchange transaction logs backup specification for all specified storage groups with disabled circular logging.
<code>-datalist <Name></code>	Specifies the name of the Microsoft Exchange ZDB backup specification file (datalist) for the Microsoft Exchange ZDB. The datalist is created in the <code><Data_Protector_home>\config\server\datalists</code> directory (on Windows) or in the <code>/etc/opt/omni/server/datalists</code> (on UNIX) directory on the Cell Manager.
<code>-device <Name></code>	Specifies the backup device to be used for the backup. If this option is not specified, the backup device must be specified using the Data Protector GUI.
<code>-force</code>	Forces overwriting of an existing backup specification file with the same name.
<code>-virtualSrv <Name></code>	The name of the Microsoft Exchange Server virtual server. This option is obligatory and used only in cluster configurations.

Table 4-3 **XP Options**

Parameter	Description
<code>-split_mirror -sse</code>	Instructs the <code>omnicreatedl</code> command to create a Microsoft Exchange split mirror backup specification file for XP.

Table 4-3 XP Options

Parameter	Description
<pre>(-local <app_sys> <bck_sys> -remote <app_sys> <bck_sys> -combined <app_sys> <bck_sys>)</pre>	<p>Specify one of the three XP split mirror configurations:</p> <p>-local selects the Business Copy XP (BC) configuration, with the application system <app_sys> and the backup system <bck_sys>. In a cluster environment, specify the virtual server hostname (rather than the physical node hostname).</p> <p>-remote selects the Continuous Access XP (CA) configuration, with the application system <app_sys> and the backup system <bck_sys>. In a cluster environment, specify the virtual server hostname (rather than the physical node hostname).</p> <p>-combined selects the Combined (BC+CA) configuration, with the application system <app_sys> and the backup system <bck_sys>. In a cluster environment, specify the virtual server hostname (rather than the physical node hostname).</p>
<pre>-mirrors <MU_Numbers></pre>	<p>This parameter is optional.</p> <p>Specify a specific replica or a replica set to be used in the backup session to define a replica set from which the integration, according to the replica set rotation, selects one replica to be used in the backup session. If this option is not specified, the MU# 0 is set.</p> <p>Enter an integer number from 0 to 2, any range of integer numbers from 0 to 2, or any combination of integer numbers from 0 to 2 separated by a comma. For example:</p> <pre>1 1-2 2, 0, 1</pre> <p>If the sequence is specified, it does not set the order in which the replicas are used. They are used according to the replica set rotation.</p> <p>If a range is entered, it must be specified in ascending order.</p>

Table 4-3 XP Options

Parameter	Description
[-split -establish]	<p>This parameter is optional. If it is not specified, the <code>-establish</code> option is set by default.</p> <p>If the <code>-split</code> option is specified, the mirrored disks in the replica selected for the current backup session are resynchronized with the P-VOLs at the start of the current backup session.</p> <p>If the <code>-establish</code> option is specified and the replica for the next backup is not synchronized, a resync will be initiated before the next backup.</p>
-leave_enabled_bs	<p>This parameter is optional.</p> <p>To specify this option, the <code>-keep_version</code> option has to be specified.</p> <p>By default, Data Protector dismounts the filesystems on the backup system after each backup. If this option is specified, the filesystems remain mounted after the backup.</p> <p>Thus, you can use the backup system for some data warehouse activity afterwards, but not for instant recovery.</p>
-instant_restore	<p>This parameter is optional.</p> <p>When this option is specified, the <code>omnicreated1</code> command automatically sets the <code>-keep_version</code> option.</p> <p>Specify the <code>-instant_restore</code> option to enable ZDB to disk or ZDB to disk+tape and instant recovery from the replica. If this option is not specified, it is not possible to perform a ZDB to disk or a ZDB to disk+tape and instant recovery from the replica. However, this option does not influence the replica set rotation.</p>

Table 4-3 XP Options

Parameter	Description
-keep_version	<p>This parameter is optional.</p> <p>If this option is specified, the pairs involved in the backup session will remain split after the backup session, enabling you to restore from the replica if an instant recovery is needed.</p> <p>If this option is not specified, the disks involved in the backup session are resynchronized after the backup session, only if one or no replica is set by the -mirrors option. If more than one replica is set by the -mirrors option, the disks involved in the backup session will remain split after the backup session.</p> <p>If this option is not specified, it is not possible to specify the -leave_enabled_bs option.</p>

Table 4-4 VA Options

Parameter	Description
-snapshot	Instructs the omncreatedl command to create an Exchange snapshot backup specification file.
-va <app_sys> <bck_sys>	Specifies the application system <app_sys> and the backup system <bck_sys> for the VA Exchange backup.
-use_existing_snapshot	<p>This parameter is optional.</p> <p>By default, Data Protector automatically sets this option if the -instant_recovery option is specified.</p> <p>If configuring a ZDB to tape, select this option to reuse an existing replica.</p> <p>Data Protector can reuse a replica only if the following condition is met, otherwise the backup session will fail:</p> <p>On a disk array, there must already exist a replica that can be reused. Only replicas that are not marked for instant recovery (are not part of the replica set) or include no snapshots that are listed in the VA LUN exclude file can be reused. Any such replica can be reused.</p>

Table 4-4 VA Options

Parameter	Description
-instant_recovery	<p>This parameter is optional.</p> <p>Specify this option to perform either a ZDB-to-disk or a ZDB-to-disk+tape session and leave the replica on a disk array (after the backup session) to use it in the future for instant recovery. If this option is not set, it is not possible to perform instant recovery from the replica created or reused in this backup session.</p> <p>If this option is specified, you should also specify the <code>-snapshots number</code> option.</p> <p>Note that when this option is selected, the options <code>-use_existing_snapshot</code> and <code>-leave_version</code> are automatically set by Data Protector.</p>
-snapshots <number>	<p>This parameter is optional.</p> <p>With <number>, specify how many replicas you want to keep on a disk array. During every backup session, Data Protector creates a new replica and leaves it on a disk array as long as the specified number is not reached. When the specified number is reached, Data Protector reuses the oldest replica.</p> <p>Note that this option sets the number of replicas in the replica set for a backup specification.</p> <p>You need to specify this number if you have selected the <code>-instant_recovery</code> option.</p> <p>If the option is not specified, it is set to 1. The maximum is 1024.</p>
-leave_version	<p>This parameter is optional.</p> <p>By default, Data Protector automatically sets this option if the <code>-instant_recovery</code> option is specified.</p> <p>If configuring a ZDB to tape (the <code>-instant_recovery</code> option is not specified), specify this option to keep the replica on a disk array after the ZDB-to-tape session. In this case, the replica will not be available for instant recovery, but can be reused in future backup sessions using the same backup specification with the option <code>-use_existing_snapshot</code> specified.</p> <p>If this option is not specified, the replica is deleted after the backup session.</p>

Table 4-4 VA Options

Parameter	Description
-leave_enabled_bs	<p>To specify this option, the <code>-leave_version</code> option has to be specified.</p> <p>By default, Data Protector dismounts the filesystems on the backup system after each backup. If this option is specified, the filesystems remain mounted after the backup.</p> <p>Thus, you can use the backup system for some data warehouse activity afterwards, but not for instant recovery.</p>
-lun_security	<p>Specify this option to apply the LUN security to the child LUNs (target volumes or snapshots) that the integration creates.</p> <p><i>If Secure Manager is activated on VA, specify this option and configure passwords correctly; otherwise, the backup sessions will fail.</i></p> <p>The LUN security is set using the <code>omnidbva</code> command. Refer to the <i>HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide</i> for more information.</p>

Table 4-5 EVA Options

Parameter	Description
-snapshot	Instructs the <code>omnicreated1</code> command to create an Exchange snapshot backup specification file.
-eva <app_sys> <bck_sys>	Specifies the application system <app_sys> and the backup system <bck_sys> for the EVA Exchange ZDB in case the HP StorageWorks EVA Agent (legacy) is installed on the systems.
-smis <app_sys> <bck_sys>	Specifies the application system <app_sys> and the backup system <bck_sys> for the EVA Exchange ZDB in case the HP StorageWorks SMI-S Agent is installed on the systems.

Table 4-5 **EVA Options**

Parameter	Description
-instant_recovery	<p>This parameter is optional.</p> <p>Specify this option to perform either a ZDB-to-disk or a ZDB-to-disk+tape session and leave the replica on a disk array (after the backup session) to use it in the future for instant recovery. If this option is not set, it is not possible to perform instant recovery from the replica created in this backup session.</p> <p>Note that when this option is selected, the options <code>-snapshot_type clone</code> and <code>-snapshot_policy strict</code> are automatically set by Data Protector. If the option <code>-snapshots number</code> is not specified, it is set to 1.</p>
-snapshots <number>	<p>This parameter is optional. By default, Data Protector automatically sets this option to 1 if the <code>-instant_recovery</code> option is specified.</p> <p>With <number>, specify how many replicas you want to keep on the disk array. During every backup session, Data Protector creates a new replica and leaves it on a disk array as long as the specified number is not reached. When the specified number is reached, Data Protector deletes the oldest replica and creates a new one.</p> <p>The maximum number for vsnaps and standard snapshots is 7. Data Protector does not limit the number of replicas rotated, but the session will fail if the limit is exceeded.</p> <p>Note that this option sets the number of replicas in the replica set for a backup specification.</p>
-snapshot_type {standard vsnap clone}	<p>This option instructs Data Protector to create one of the three types of EVA snapshots during the backup session.</p> <p>Setting <code>standard</code> creates snapshots with the pre-allocation of disk space.</p> <p>Setting <code>vsnap</code> creates snapshots without the pre-allocation of disk space.</p> <p>Setting <code>clone</code> creates a clone of a source volume (original virtual disk).</p>

Table 4-5 EVA Options

Parameter	Description
<p><code>-snapshot_policy {strict loose}</code></p>	<p>Specifies how Data Protector creates snapshots with regard to types of already existing snapshots for the same source volume (original virtual disk).</p> <p>When <code>strict</code> is set, Data Protector attempts to create snapshots of the type selected by the <code>-snapshot_type</code> option. If some of the source volumes (original virtual disks) used in the backup session already have existing snapshots of different type, the selected type of snapshots cannot be used. Such a backup session will be aborted.</p> <p>When <code>loose</code> is set, Data Protector creates snapshots of a different type than specified by the <code>-snapshot_type</code> option, when this would help to make a successful session.</p> <p>For example, if you select standard snapshots to be created, but Data Protector detects that standard snapshots cannot be created because some vsnaps or snapclones of the source volumes already exist in a replica set, the following happens: with the <code>loose</code> option selected, Data Protector creates either vsnaps (if vsnaps already exist) or snapclones (if snapclones already exist) instead of standard snapshots.</p> <p>Note that Data Protector can use only one type of snapshots in the backup session. In case when some of the source volumes used in the backup session have existing standard snapshots and some of them existing vsnaps, the backup session will be aborted.</p>
<p><code>-wait_cloncopy <number></code></p>	<p>This parameter is optional and can be specified only if the <code>-snapshot_type clone</code> option is specified.</p> <p>In the case of a ZDB-to-tape or a ZDB-to-disk+tape session, specify this option to delay moving data to tape media until the cloning process is completed. By <code><number></code>, specify the maximum waiting time in minutes. After the specified number of minutes, the backup to tape will start, even if the cloning process is not finished yet.</p> <p>With this option, you prevent degradation of the application data access times during the phase of backup to tape.</p>

Table 4-6 Microsoft Exchange Server Options

Parameter	Description
-annotation { MIS SRS KMS }	This option specifies the possible Microsoft Exchange Server annotations: Microsoft Information Store (MIS), Site Replication Service (SRS), and Key Management Service (KMS). MIS is the default setting and does not need to be specified in case when the MIS will be backed up.
-all_storage_groups	This option creates a backup specification for all databases relating to Microsoft Exchange Microsoft Information Store. It must be specified by the -annotation MIS parameter.
-storage_group <Storage_Group_Name>	This option creates a backup specification for all stores relating to the specified. Multiple declarations of the -storage_group parameter are possible to create a backup specification for the selected storage groups. Logical names can be obtained by using the Exchange System Administrator tool, which is a part of Microsoft Exchange Server.
-store <Store>	When the -store parameter is specified, backup specification is created only for specified store(s) inside the storage group. List of stores can be specified after the -store parameter to create a backup specification for many stores. Store names can be obtained by using Exchange System Administrator tool, which is a part of Microsoft Exchange Server.

For additional information, refer to the omncreatedl man page.

Using the omniex2000SM.bat File

You can also use the omniex2000SM.bat file, located in the <Data_Protector_home>\bin directory on the application system to create a Microsoft Exchange ZDB backup specification and a Microsoft Exchange transaction logs backup specification. The script provides templates and examples of the omncreatedl usage. You can modify the omniex2000SM.bat file using any text editor. Uncomment (delete the @REM before the command) the line with the appropriate command and edit the parameters; for example, specify the backup specification name, storage group, application, and backup systems. See Figure 4-2 on page 324.

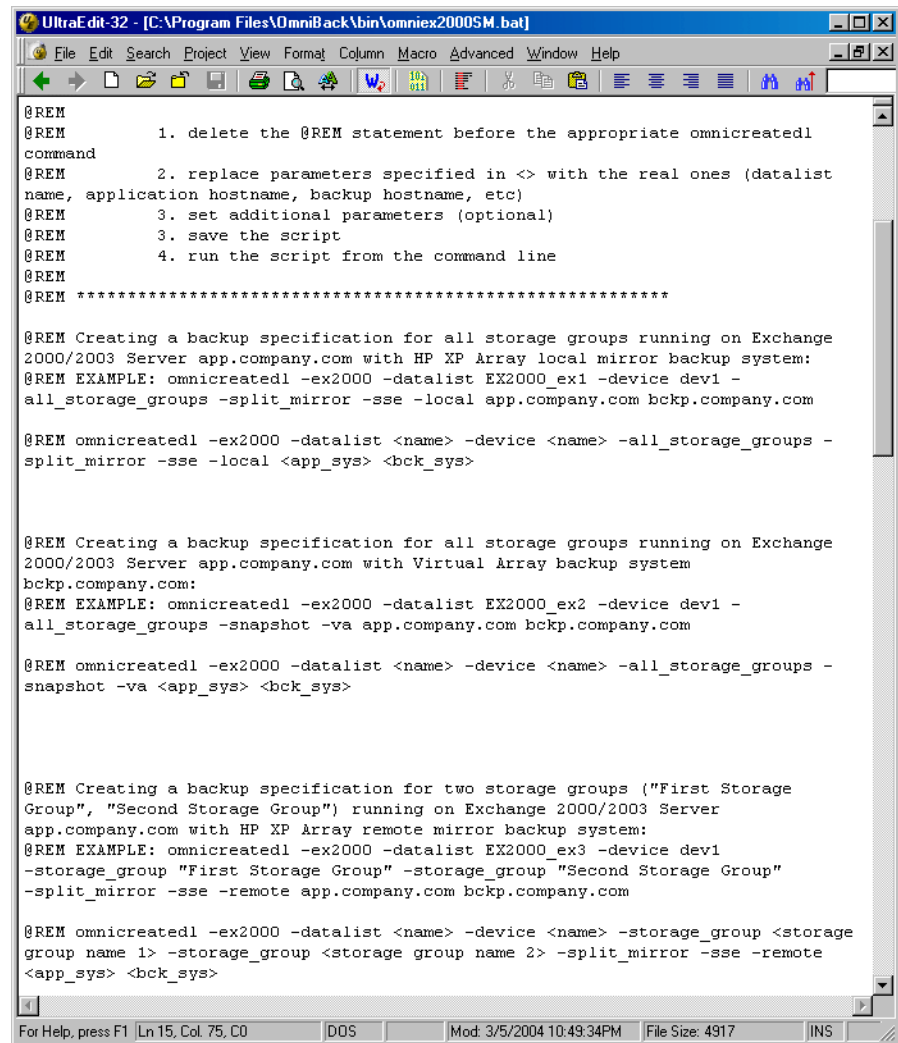
Data Protector Microsoft Exchange Server ZDB Integration Configuring a Microsoft Exchange ZDB Backup Specification

To create a Microsoft Exchange ZDB backup specification and a Microsoft Exchange transaction logs backup specification using the `omniex2000SM.bat` file, make the necessary modifications and run the following command:

```
<Data_Protector_home>\bin\omniex2000SM.bat
```

Figure 4-2

omniex2000SM.bat File



```
@REM
@REM      1. delete the @REM statement before the appropriate omnicreated1
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: omnicreated1 -ex2000 -datalist EX2000_ex1 -device dev1 -
all_storage_groups -split_mirror -sse -local app.company.com bckp.company.com

@REM omnicreated1 -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: omnicreated1 -ex2000 -datalist EX2000_ex2 -device dev1 -
all_storage_groups -snapshot -va app.company.com bckp.company.com

@REM omnicreated1 -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: omnicreated1 -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 omnicreated1 -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>
```

The following are some examples of using the `omnicreatedl` command to create a Microsoft Exchange ZDB backup specification (`datalist`) and a Microsoft Exchange transaction logs backup specification:

Example 1 - XP

Use the following command to create a backup specification, named `BS1`, using device `dev1`, for all storage groups running on Microsoft Exchange Server on application system `computer_app.company.com` and backup system `computer_bck.company.com`, using the HP StorageWorks XP BC configuration:

```
omnicreatedl -ex2000 -datalist BS1 -device dev1  
-all_storage_groups -split_mirror -sse -local  
computer_app.company.com computer_bck.company.com
```

Example 2 - XP

Use the following command to create a backup specification, named `BS2`, using device `dev1`, for two storage groups (First Storage Group and Second Storage Group) running on Microsoft Exchange Server on application system `computer_app.company.com` and backup system `computer_bck.company.com`, using the **XP CA configuration**:

```
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

Use the following command to create a backup specification, named `BS3`, using device `dev1`, for two stores (Public and Mailbox, part of First Storage Group) running on Microsoft Exchange Server on application system `computer_app.company.com` and backup system `computer_bck.company.com`, with the **XP BC/CA configuration**:

```
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

The following command will create a VA Exchange **ZDB-to-tape** backup specification, named `BS1`, using the backup device `dev1`, for all storage groups running on Microsoft Exchange Server on the application system `computer1.company.com` and the backup system `computer2.company.com`. If the backup specification already exists, it will be overwritten. In case it does not already exist, the `omnicreatedl` command creates an HP StorageWorks VA Exchange transaction logs backup specification file `First Storage Group (LOGS)`

Data Protector Microsoft Exchange Server ZDB Integration

Configuring a Microsoft Exchange ZDB Backup Specification

computer1.company.com for First Storage Group log files backup. The replica is kept for future (non-instant recovery) use if the backup session is successful.

```
omnicreatedl -ex2000 -datalist BS1 -device dev1 -snapshot  
-va computer1.company.com computer2.company.com  
-leave_version -storage_group "First Storage Group" -force
```

Example 5 - VA

The following command will create a VA Exchange **ZDB-to-disk+tape** or **ZDB-to-disk** backup specification, named Exchange_example, to back up Site Replication Service on the backup device dev1, using the replica set with 5 replicas. In case it does not already exist, the omnicreatedl command creates a VA Exchange transaction logs backup specification file SRS (LOGS) computer1.company.com for Site Replication Service log files backup (the circular logging must be disabled). When the omnib command or Data Protector GUI is used to start the backup using the created backup specification, choose between the ZDB-to-disk or ZDB-to-disk+tape session:

```
omnicreatedl -ex2000 -datalist Exchange_example -device dev1  
-snapshot -va computer1.company.com computer2.company.com  
-instant_recovery -snapshots 5 -annotation SRS
```

Example 6 - EVA

The following commands (the first command is executed on the EVA that uses the HP StorageWorks EVA Agent (legacy), and the second one - on the EVA that uses the HP StorageWorks EVA SMI-S Agent) will create a **ZDB-to-tape** backup specification, named BS1, using the backup device dev1, for all storage groups running on Microsoft Exchange Server on the application system computer1.company.com and the backup system computer2.company.com. In case it does not already exist, the omnicreatedl command creates an EVA Exchange transaction logs backup specification file First Storage Group (LOGS) computer1.company.com for First Storage Group log files backup. Data Protector will try to create the vsnap type of snapshots. If they cannot be created, the session aborts.

```
omnicreatedl -ex2000 -datalist BS1 -device dev1 -snapshot  
-eva computer_app.company.com computer_bck.company.com  
-snapshot_type vsnap -snapshot_policy strict -storage_group  
"First Storage Group"
```

if you have the HP StorageWorks EVA Agent (legacy) installed, or:


```
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"
```

if you have the HP StorageWorks EVA SMI-S Agent installed.

Example 7 - EVA

The following commands (the first command is executed on the EVA that uses the HP StorageWorks EVA Agent (legacy), and the second one - on the EVA that uses the HP StorageWorks EVA SMI-S Agent) will create an EVA Exchange **ZDB-to-disk** backup specification, named Exchange_example, to back up Site Replication Service using the backup device dev1, using the replica set with 5 replicas. In case it does not already exist, the omnicreatedl command creates an EVA Exchange transaction logs backup specification file SRS (LOGS) computer1.company.com for Site Replication Service log files backup (the circular logging must be disabled). When the omnib command or Data Protector GUI is used to start the backup using the created backup specification, choose the ZDB-to-disk session:

```
omnicreatedl -ex2000 -datalist Exchange_example -device dev1  
-snapshot -eva computer1.company.com computer2.company.com  
-instant_recovery -snapshots 5 -annotation SRS
```

if you have the HP StorageWorks EVA Agent (legacy) installed, or:

```
omnicreatedl -ex2000 -datalist Exchange_example -device dev1  
-snapshot -smis computer1.company.com computer2.company.com  
-instant_recovery -snapshots 5 -annotation SRS
```

if you have the HP StorageWorks EVA SMI-S Agent installed.

Example 8 - EVA

The following commands (the first command is executed on the EVA that uses the HP StorageWorks EVA Agent (legacy), and the second one - on the EVA that uses the HP StorageWorks EVA SMI-S Agent) will create an EVA Exchange **ZDB-to-disk+tape** backup specification, named Exchange_example, to back up Site Replication Service to the backup device dev1, using the replica set with 3 replicas and to delay the backup to tape for the maximum of 50 minutes. In case it does not already exist, the omnicreatedl command creates an EVA Exchange transaction logs backup specification file SRS (LOGS) computer1.company.com for Site Replication Service log files backup (the circular logging must be

disabled). When the `omnib` command or Data Protector GUI is used to start the backup using the created backup specification, choose the ZDB-to-disk+tape session:

```
omnicreatedl -ex2000 -datalist Exchange_example -device dev1  
-snapshot -eva computer1.company.com computer2.company.com  
-instant_recovery -snapshots 3 -wait_cloncopy 50  
-annotation SRS
```

if you have the HP StorageWorks EVA Agent (legacy) installed, or:

```
omnicreatedl -ex2000 -datalist Exchange_example -device dev1  
-snapshot -smis computer1.company.com computer2.company.com  
-instant_recovery -snapshots 3 -wait_cloncopy 50  
-annotation SRS
```

if you have the HP StorageWorks EVA SMI-S Agent installed.

Modifying a Microsoft Exchange ZDB Backup Specification

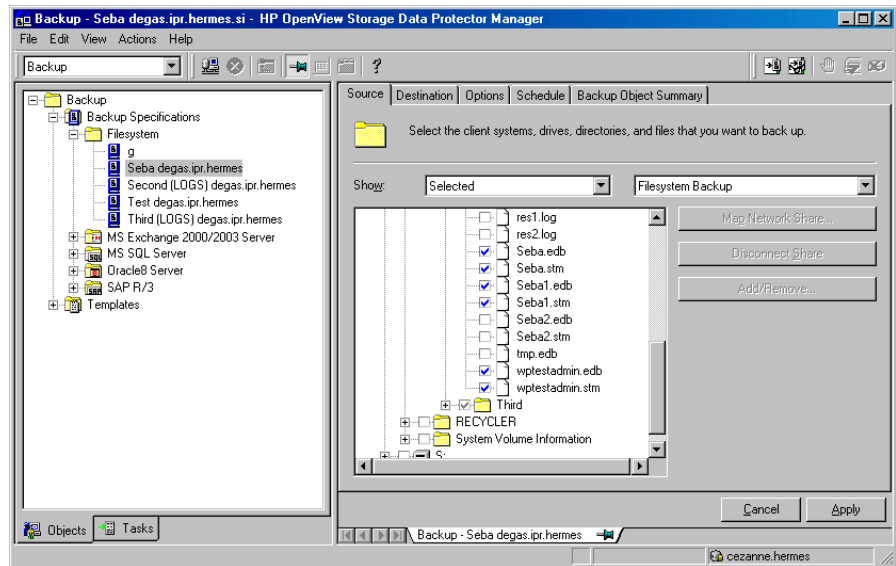
After you have created the Microsoft Exchange ZDB backup specification and a Microsoft Exchange transaction logs backup specification using the `omnicreatedl` command, they can be modified using the Data Protector GUI.

When you want to add or remove a Microsoft Exchange storage group from a Microsoft Exchange ZDB backup specification, it is recommended to create new backup specification file rather than modify the existing one. Changing the saved Microsoft Exchange ZDB backup specification file manually could impact the `Stop/quiesce` the application and `Restart` the application scripts' parameters, which are automatically defined when a new backup specification is created using the `omnicreatedl` command.

Follow the procedure to modify the Microsoft Exchange ZDB backup specification or a Microsoft Exchange transaction logs backup specification:

1. In the Data Protector GUI, select **Backup** in the Context List.
2. In the Scoping Pane expand **Backup Specifications** and then **Filesystem**. Click the Microsoft Exchange ZDB backup specification, which you created using the `omnicreatedl` command. See Figure 4-3 on page 329.

Figure 4-3 Microsoft Exchange ZDB Backup Specification (Datalist)



3. You can modify the backup devices, set the ZDB options, schedule, and start backup. For a particular database, both the .edb and .stm files must be backed up.

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 under the Destination tag. Select separate devices for the backup and for each mirror.

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

NOTE

Object mirroring is not supported for ZDB to disk.

Checking Microsoft Exchange Files for Consistency

Microsoft Exchange Server comes with the utility that allows you to perform a page level integrity check of the database files. This consistency check can be run on the Microsoft Exchange offline database

from the Data Protector during Data Protector ZBD Exchange backup. For more information on the `esefile` (Exchange 2000) and the `eseutil` (Exchange 2003) utilities, refer to the *Esefile Support Utility for Exchange Server 5.5 and Exchange 2000 Server (248406)* and to the *How To: Use the Eseutil Utility to Detect File Header Damage in Exchange 2003 (825088)* documents respectively, available at <http://support.microsoft.com>.

IMPORTANT

Page level integrity check of the Microsoft Exchange database is not supported with EVA.

The procedure below gives instructions on how to use the utility:

1. Copy the utility to the backup system.

The `esefile` utility can be found on the Microsoft Exchange 2000 Installation CD at the following location:

```
<CD>:\ENGLISH\EXCH2000\ENT\SUPPORT\UTILS\I386\ESEFILE
```

The `eseutil` utility can be found in the

```
<Exchange_Server_home>\bin
```

directory once the Microsoft Exchange Server 2003 is installed.

NOTE

On Disk Array XP, location on which `esefile.exe` or `eseutil.exe` is copied should not be mirrored, so that the file will not be overwritten when the mirror is synchronized.

2. Write a script that will run the `esefile.exe` or `eseutil.exe` utility and save it, for example, as `integritycheck.bat`. Basically, the command must look 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 the `eseutil.exe` utility with the `/k` option tests the checksums on the Exchange Server database.

Example

In the following example, the `esefile.exe` resides in `C:\` directory on the backup system. Data Protector mounts data to be backed up to the `C:\Program Files\Omniback\tmp` directory 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`.

Figure 4-4

Example Script for MS Exchange Server 2000

```
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"
```

Refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide* or for information on how Data Protector creates mountpoints on the backup system. The commands in the script will ensure that `*.edb` database files specified to be backed up in the backup specification will be checked for consistency.

3. Modify the Microsoft Exchange ZDB backup specification to include the `integritycheck.bat` as Post-Exec script on the backup specification level. The script must be executed on the backup system. For more information on configuring Pre-Exec and Post-Exec commands, see online Help index: "pre- and post-exec commands".
4. In the backup specification, select the Leave the backup system enabled option. You can select this option only if the Keep the replica after the backup option is selected. This option ensures that the database will be available on the backup system at the time when the integrity check is started.

After the backup is done, the `esefile` or the `esetool` utility will perform page level integrity check against database files specified in your `integritycheck.bat` script.

Backing Up Microsoft Exchange Server

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

Backup Methods

- Schedule the backup of an existing Microsoft Exchange ZDB backup specification using the Data Protector Scheduler.
- Start an interactive backup of an existing Microsoft Exchange ZDB backup specification using the Data Protector GUI or the Data Protector command-line interface.

Considerations

Before running a Microsoft Exchange ZDB session, note the following:

- It is not possible to start 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 or restore session; otherwise, the session will fail.
- Transaction logs must be backed up using the common filesystem backup functionality.

Scheduling a Backup

Scheduling a backup specification means setting 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, refer to the online Help index keyword “scheduled backups”.

To schedule a Microsoft Exchange ZDB session, proceed as follows:

1. In the HP OpenView Storage Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, Backup Specifications, and then Filesystem.
3. Double-click the backup specification you want to schedule and click the Schedule tab.

4. In the Schedule page, select a date in the calendar and click Add to open the Schedule Backup dialog box.
5. Specify Recurring, Time options, Recurring options, and Session options.

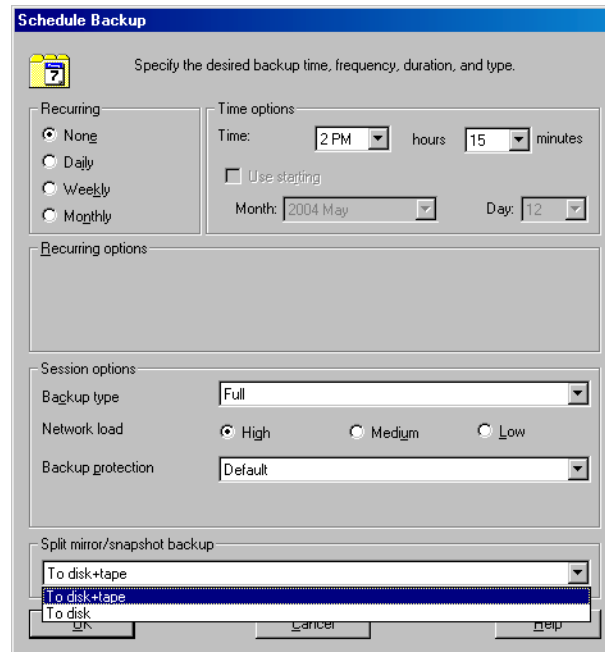
Note that the backup type is ignored for ZDB sessions. It is set to Full.

In the case of a ZDB-to-disk or a ZDB-to-disk+tape session, specify the Split mirror/snapshot backup option. See Figure 4-5.

NOTE

It is not possible to run a ZDB-to-disk or a ZDB-to-disk+tape session if the `-instant_restore` (for XP) or `-instant_recovery` (for VA and EVA) option is not specified when creating the backup specification.

Figure 4-5 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

Starting Backup Using the GUI

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

To start an interactive ZDB session of a Microsoft Exchange database, proceed as follows:

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

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

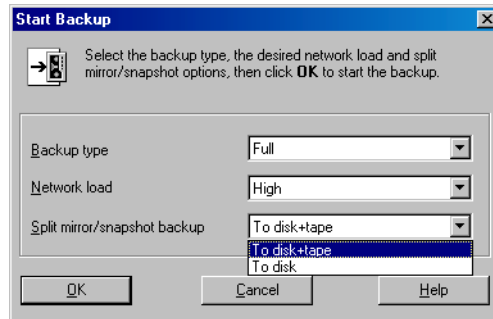
Note that the backup type is ignored for ZDB sessions. It is set to Full.

In the case of a ZDB-to-disk or a ZDB-to-disk+tape session, specify the Split mirror/snapshot backup option. See Figure 4-5.

NOTE

It is not possible to run a ZDB-to-disk or a ZDB-to-disk+tape session if the `-instant_restore` (for XP) or `-instant_recovery` (for VA and EVA) option is not specified when creating the backup specification.

Figure 4-6 **Selecting ZDB-to-Disk or ZDB-to-Disk+Tape Session When Starting an Interactive Backup**



Click OK.

Starting Backup Using the CLI

To start a Microsoft Exchange ZDB-to-tape or ZDB-to-disk+tape session using the Data Protector CLI, use the following command:

```
omnib -datalist <Name>
```

To start a Microsoft Exchange **ZDB-to-disk** session using the Data Protector CLI, use the following command:

```
omnib -datalist <Name> -disk_only
```

where <Name> is the name of the backup specification. For more information on the omnib command, refer to its man page.

NOTE

It is not possible to run a ZDB-to-disk or a ZDB-to-disk+tape session if the `-instant_restore` (for XP) or `-instant_recovery` (for VA and EVA) option is not selected during the creation of the backup specification using the `omnicreatedl` command.

Restoring a Microsoft Exchange Database

There are two ways of restoring a Microsoft Exchange database:

- restoring from backup media to the application system on LAN
- restoring using the instant recovery functionality

Restoring from the Backup Media to the Application System on LAN

When restoring from backup media to the application system on LAN, Microsoft Exchange backup objects are restored online directly to the application system.

When restoring from backup media to the application system on LAN, you can restore Microsoft Exchange backup objects to any Microsoft Exchange Server with the same configuration as the original system within the Data Protector cell.

A Microsoft Exchange database can be recovered using one of the following two methods:

- Point-in-time recovery
Only data (.edb and .stm files) is restored, without the log files. Database is restored to the state when the backup was performed, and all data created after the backup is lost.
- Rollforward recovery
This type of recovery consists of restoring the Microsoft Exchange database files and transaction logs, and then replaying the transaction logs. This action recovers the Microsoft Exchange database to the last consistent state.

TIP

You can improve the data transfer rate when restoring by connecting the backup device to the application system and configuring this backup device on the application system using the Data Protector GUI. See online Help index: “configuring, backup devices”, for more information on configuring the backup devices. You can perform a restore using another device, see online Help index “select, devices for restore”.

Point-in-Time Recovery

Using Data Protector GUI you can perform filesystem restore of Microsoft Exchange database (the .edb and .stm files), but to recover the database, additional database recovery related steps must be performed. For detailed procedure on how to recover the Microsoft Exchange Server, refer to:

Offline Backup and Restoration Procedures for Exchange (296788) document, available at <http://support.microsoft.com>.

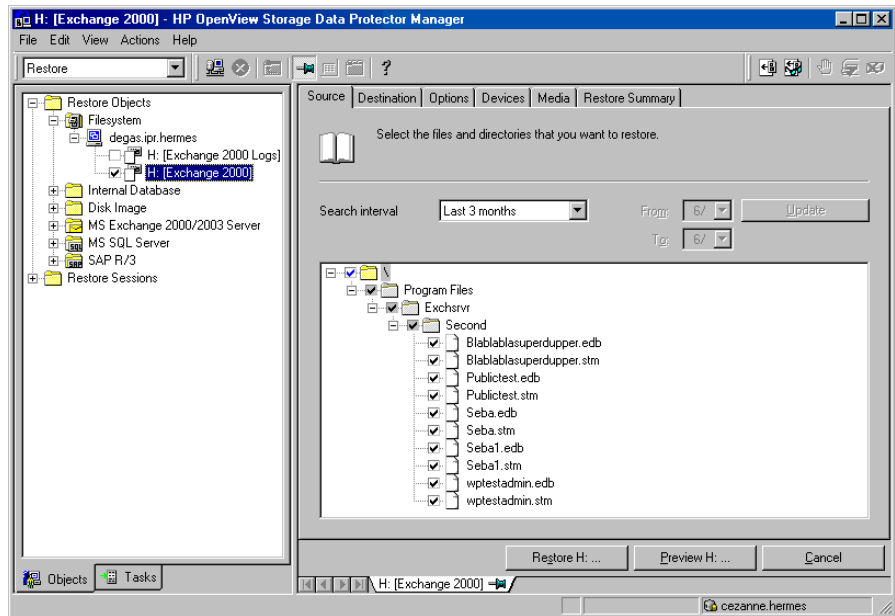
Microsoft Exchange databases are backed up as filesystem objects and can be restored using Data Protector GUI from backup media to the application system on LAN. For a detailed procedure on how to perform filesystem restore, see online Help index: “standard restore procedure”. The restore procedure described here provides only a general description of how to restore Microsoft Exchange databases backed up using the ZDB functionality from backup media to the application system on LAN.

IMPORTANT

The restore procedure described here gives instructions on how to copy the .edb and .stm files to the appropriate database location and is only part of the procedure described in the *Offline Backup and Restoration Procedures for Exchange (296788)* document, available at <http://support.microsoft.com>.

1. In the Context List, select Restore.
2. Expand Restore Objects, Filesystem, and the name of the backed up server. Then mark the drive letter on which the database resides/resided as shown in the Figure 4-7 on page 338.

Figure 4-7 Restoring Microsoft Exchange Database

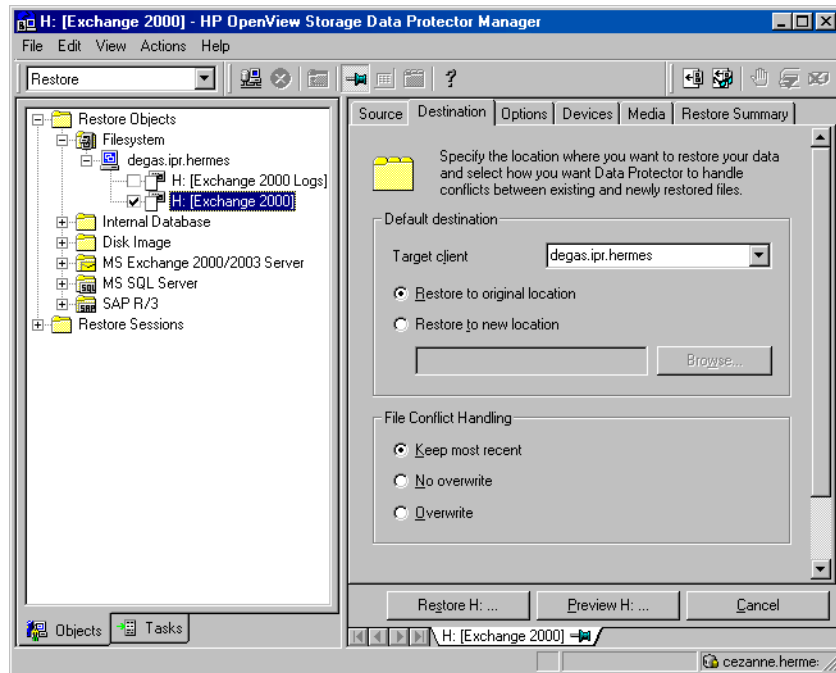


The Microsoft Exchange database consists of two files: `<name>.edb` and `<name>.stm`. For a particular database, both files must be selected for restore. Restoring a storage group by selecting the storage group folder is possible when the entire storage group was backed up and all databases reside in the same directory on disk.

For more information on available restore options, see online Help index: "restore, options".

In order to restore a Microsoft Exchange database object backed up using the ZDB functionality from backup media to the application system on LAN, make sure to select the application system as the Target client under the Destination tab.


Figure 4-8 **Selecting the Application System**



3. After you have set the restore options, click the Restore button. The Start Restore Session dialog box is displayed.
4. Click Next to specify the Report level and Network load. Click Next.
5. This step and Figure 4-9 on page 340 are relevant only for XP and if you have the EMC Symmetrix Agent software component installed on the application system.

Select StorageWorks XP restore. Click Next to display the Start Restore Session dialog box.

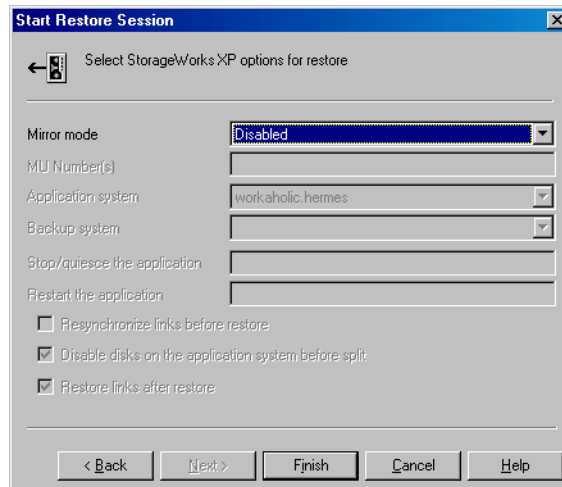
Figure 4-9 **Selecting the XP Restore**



6. This step and Figure 4-10 on page 341 are relevant only for 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.

Figure 4-10 **Selecting the XP Restore Option**



7. Click **Finish** to start the restore session.

Disaster Recovery In case of a disaster, Microsoft Exchange Server must be installed and configured with the same database names and locations as before the disaster.

Rollforward Recovery

Using Data Protector GUI you can perform filesystem restore of Microsoft Exchange database (the `.edb` and `.stm` files) and transaction logs (the `.log` files); but to recover the database, additional database recovery related steps must be performed. For the detailed procedure on how to recover the Microsoft Exchange Server, refer to:

Offline Backup and Restoration Procedures for Exchange (296788) document, available at <http://support.microsoft.com>.

Microsoft Exchange databases are backed up as filesystem objects and can be restored using Data Protector GUI from backup media to the application system on LAN. For a detailed procedure on how to perform a filesystem restore, see online Help index: “standard restore procedure“. The restore procedure described here provides only a general description of how to restore Microsoft Exchange databases backed up using the ZDB functionality from backup media to the application system on LAN.

IMPORTANT

The restore procedure described here gives instructions on how to copy the .edb, .stm and .log files to the appropriate database location and is just a part of the procedure described in the *Offline Backup and Restoration Procedures for Exchange (296788)* document.

When following the Microsoft procedure, use the procedure below to copy the .edb, .stm and .log files to the appropriate database location.

The procedure below needs to be utilized twice: for the restore of the database and for the restore of the transaction logs.

-
1. In the Context List, select Restore.
 2. Perform one of the following:
 - To restore the database (.stm and .edb files), expand Restore Objects, Filesystem, and the name of the backed up server. Then select the [Exchange 2000] object as shown in Figure 4-11 on page 343.
 - To restore the transaction logs (.log files), expand Restore Objects, Filesystem, and the name of the backed up server. Then select the [Exchange 2000 Logs] object as shown in Figure 4-12 on page 343.

Figure 4-11 Restoring Microsoft Exchange Database

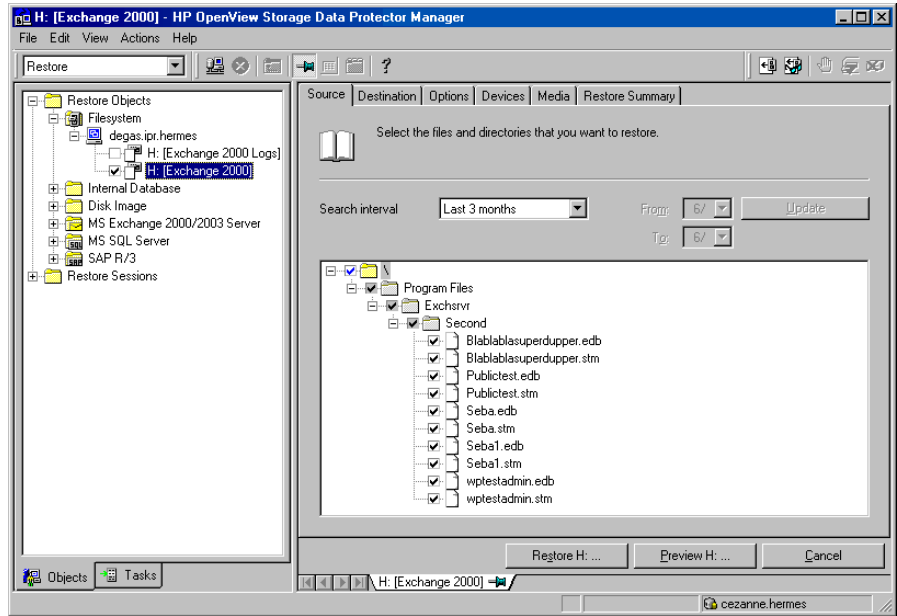
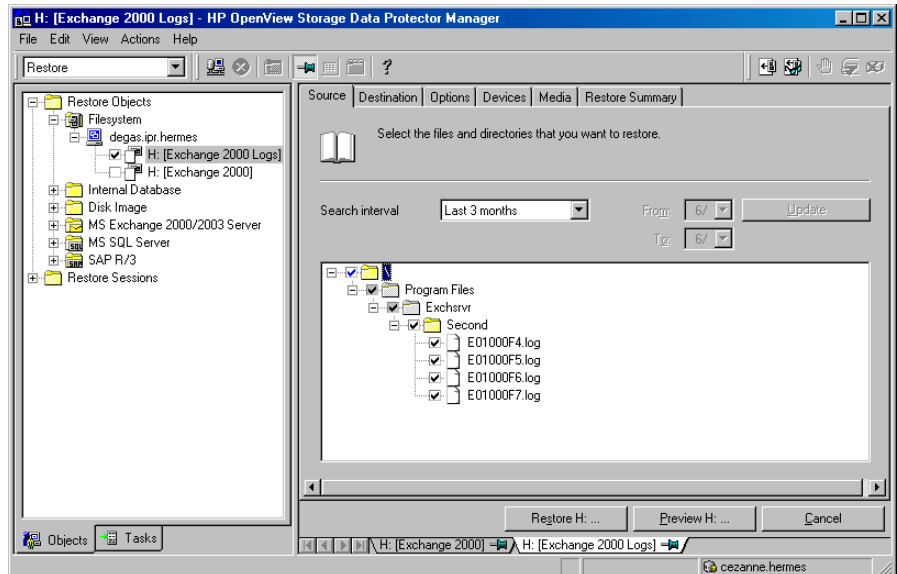


Figure 4-12 Restoring Microsoft Exchange Log Files



3. Make the following selections in the Results pane:

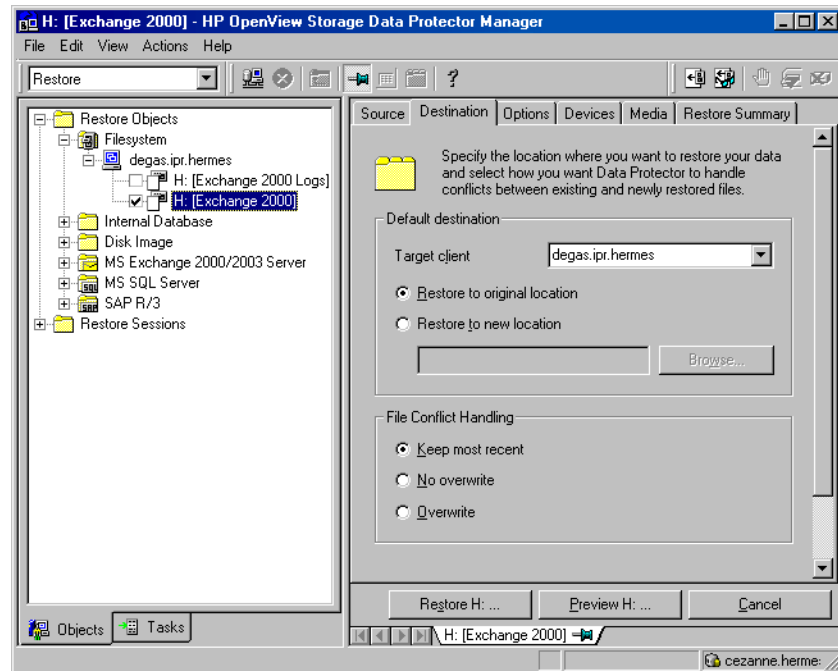
- To restore the Microsoft Exchange database, select `.edb` and `.stm` files. For a particular database, both files must be selected for restore. Restoring a storage group by selecting the storage group folder is possible when the entire storage group was backed up and all databases reside in the same directory on disk.
- To restore the transaction log files, select the `.log` files.

For more information on available restore options, see online Help index: “restore, options“. Make sure to make the following selections in the scoping pane in order to restore a Microsoft Exchange database backed up using the ZDB functionality from backup media to the application system on LAN.

4. Select the application system as the Target client under the Destination tab.

Figure 4-13

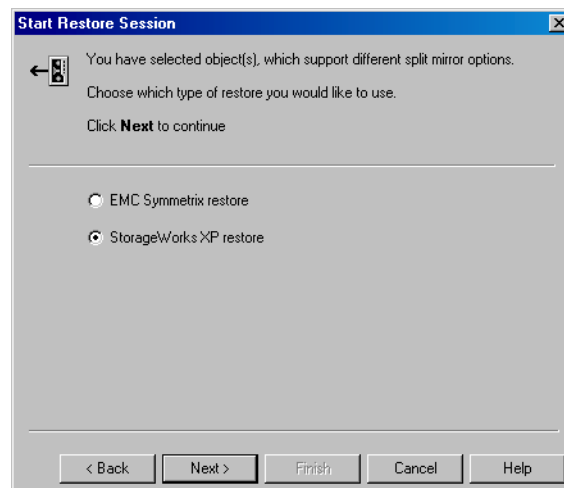
Selecting the Application System



5. After you have set the restore options, click the Restore button. The Start Restore Session dialog box is displayed.
6. Click Next to specify the Report level and Network load. Click Next.
7. This step and Figure 4-14 on page 345 are relevant only for XP and if you have the EMC Symmetrix Agent software component installed on the application system.

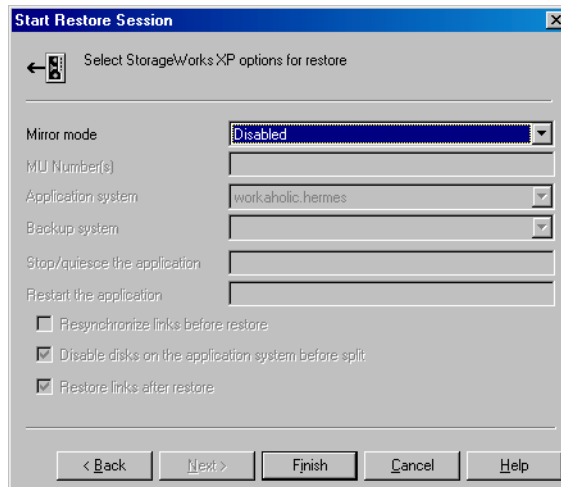
Select StorageWorks XP restore. Click Next to display the Start Restore Session dialog box.

Figure 4-14 **Selecting XP Restore**



8. This step and Figure 4-15 on page 346 are relevant only for 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.

Figure 4-15 **Selecting the XP Restore Option**



9. Click **Finish** to start the restore session.

Disaster Recovery In case of a disaster, Microsoft Exchange Server must be installed and configured with the same database names and locations as before the disaster.

Instant Recovery

IMPORTANT

The Data Protector instant recovery functionality does not provide recovery of a database or an application. It moves the data from a replica to the source volumes. To recover a database or an application, additional database or application recovery related steps must be performed.

A Microsoft Exchange database can be recovered using one of the following two methods:

- Point-in-time recovery
Database is restored to the state when the backup was performed, and all data created after the backup is lost.

- Roll forward recovery

This type of recovery consists of restoring the Microsoft Exchange database files and transaction logs, and then replaying the transaction logs. This action recovers the Microsoft Exchange database to the last consistent state.

Point-in-Time Recovery

For detailed procedure on how to recover the Microsoft Exchange Server, refer to:

Offline Backup and Restoration Procedures for Exchange (296788), available at <http://support.microsoft.com>.

When following the Microsoft procedure, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide* for information on how to copy the backed up .edb and .stm files to the appropriate database using the Data Protector instant recovery functionality.

Roll Forward Recovery

For detailed procedure on how to recover the Microsoft Exchange Server, refer to:

Offline Backup and Restoration Procedures for Exchange (296788), available at <http://support.microsoft.com>.

When following the Microsoft procedure, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide* for information on how to copy the backed up .edb and .stm files to the appropriate database using the Data Protector instant recovery functionality.

The procedure below gives instructions on how to restore Microsoft Exchange transaction logs using the Data Protector GUI when following the Microsoft procedure:

1. In the Context List, select Restore.
2. Expand Restore Objects, Filesystem, and the name of the backed up server. Then select the [Exchange 2000 Logs] object as shown in Figure 4-12 on page 343.

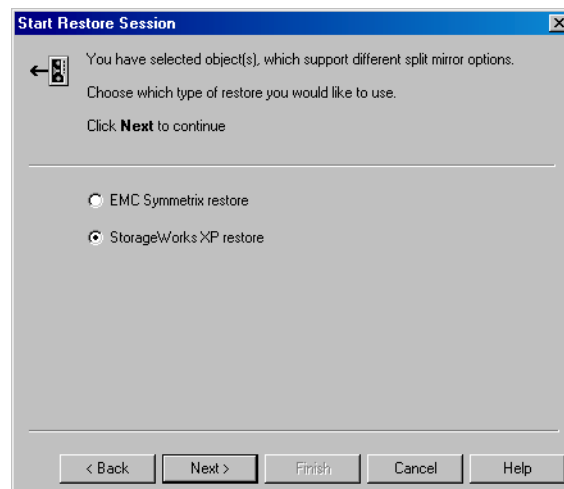
For more information on various restore options available, see online Help index: “restore, options“. Make sure to make the following selections in the scoping pane in order to restore a Microsoft Exchange database backed up using the ZDB functionality from backup media to the application system on LAN:

- Select the application system as the Target client under the Destination tag as shown in the Figure 4-13 on page 344.
3. After you have set the restore options, click the Restore button. The Start Restore Session dialog box is displayed.
 4. Click Next to specify the Report level and Network load. Click Next.
 5. This step and Figure 4-9 on page 340 are relevant only for XP and if you have the EMC Symmetrix Agent software component installed on the application system.

Select StorageWorks XP restore. Click Next to display the Start Restore Session dialog box.

Figure 4-16

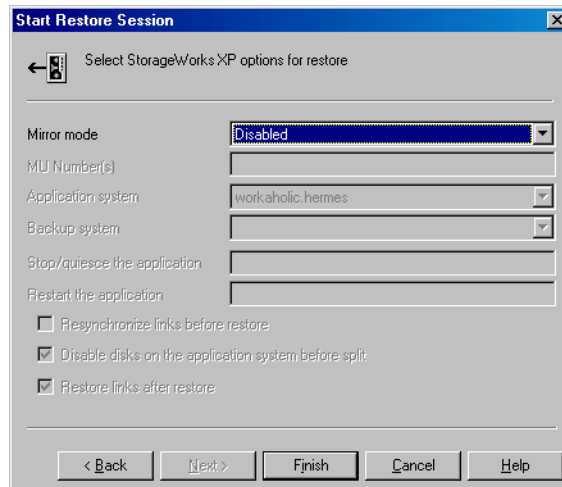
Selecting the XP Restore



6. This step and Figure 4-10 on page 341 are relevant only for 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.

Figure 4-17 **Selecting the XP Restore Option**



7. Click **Finish** to start the restore session.

Disaster Recovery In case of a disaster, Microsoft Exchange Server must be installed and configured with the same database names and locations as before the disaster.

Troubleshooting

This section contains a list of problems you might encounter when using the Data Protector Microsoft Exchange integration.

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

For general ZDB, restore, and instant recovery related troubleshooting, see the troubleshooting sections in the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide*.

See also the Microsoft Exchange troubleshooting section in the *HP OpenView Storage 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 OpenView Storage Data Protector Product Announcements, Software Notes, and References* for general Data Protector limitations, as well as known problems and workarounds.
- ✓ See http://www.openview.hp.com/products/datapro/spec_0001.html for an up-to-date list of supported versions, platforms, and other information.

Backup Problems

Unable to initiate clsEx2000 class

Microsoft Exchange integration reports the following error in the debug.log file on the application system:

```
Unable to initiate clsEx2000 class. Error: -2147221164
```

You need to register the omniex2000.dll. Run the following command in the <Data_Protector_home>\bin directory:

```
regsvr32 omniex2000.dll
```

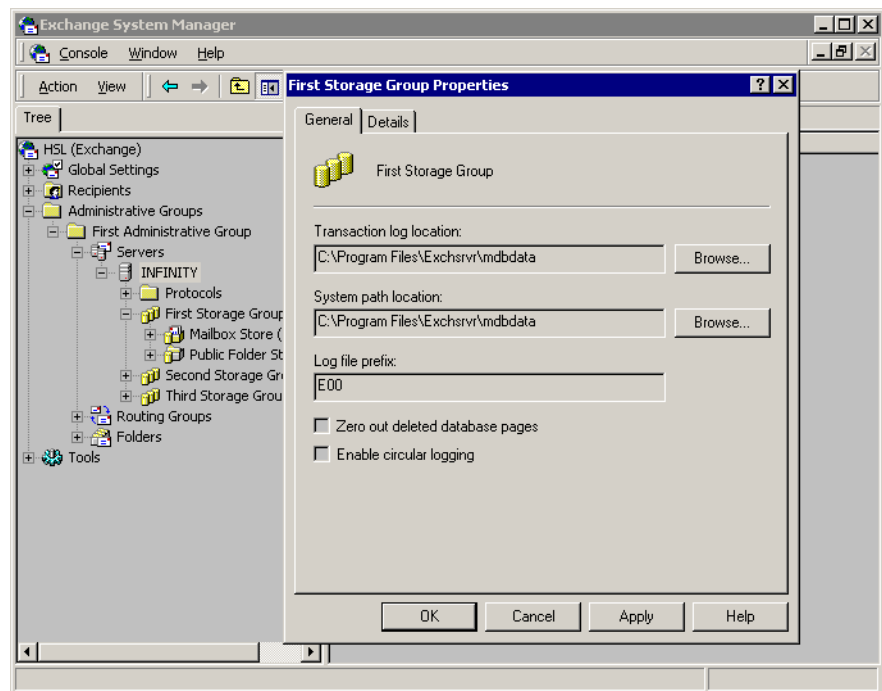

Cannot create log datalist for < name>

The following error is reported if the `omnicreatedl` command cannot create a backup specification for log file backup:

```
Cannot create log datalist for < name>.
```

Check if circular logging is enabled on the Microsoft Exchange server for the particular storage group and disable it. See the following figure:

Figure 4-18 **Disabling Circular Logging**



Could not obtain any data for backup from host

The `omnicreatedl` command 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.

Troubleshooting

- If the application is cluster-aware, specify the `-virtualSrv` parameter.
- Check the user rights on the application system client. You should have the `Save backup specification` user right to be able to create and save a backup specification.
- Check if the `<Exchange_home>\bin` directory is listed in the `Path` system variable. For instructions, see “Prerequisites and Limitations” on page 306.

Error starting service

If KMS is configured in such a way that the KMS service password is kept on a secured piece of paper and not on a disk, the following error is reported during the backup:

```
[Major] From: OB2BAR_main@machine.com "" Time: 10/20/2003  
5:17:24 PM
```

Error starting service.

Reconfigure the KMS to store the KMS service password on a disk. Refer to *XADM: How to Change the KMS Service Password Startup Location (196129)* for information on how to do this. Then restart the backup.

5**Integrating the Data Protector
ZDB Integrations and Microsoft
Volume Shadow Copy Service**

In This Chapter

This chapter explains how to configure and use the Data Protector Microsoft Volume Shadow Copy integration.

The chapter is organized into the following chapters:

“Introduction” on page 355

“Prerequisites and Limitations” on page 359

“Integration Concepts” on page 361

“Configuring the Integration” on page 367

“Writers Specifics” on page 371

“Backing Up Writers Data” on page 376

“Restoring Writers Data” on page 387

“Troubleshooting” on page 400

“User Scenario For Microsoft Exchange Server 2003 Backup and Restore” on page 405

Introduction

A traditional backup process is based on the direct communication between the backup application and the application to be 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. This service collaborates with the backup application, applications to be backed up, shadow copy providers, and the operating system kernel to implement the management of volume shadow copies and shadow copy sets.

HP OpenView Storage Data Protector supports the integration with the Microsoft Volume Shadow Copy service (VSS).

The Data Protector Volume Shadow Copy integration provides a unified communication interface that can coordinate backup and restore of any application regardless of their specific features. With this approach, 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 5-1 and Figure 5-2 show the differences between the traditional backup model and the model with the Data Protector MS Volume Shadow Copy integration.

Figure 5-1

Actors of the Traditional Backup Model

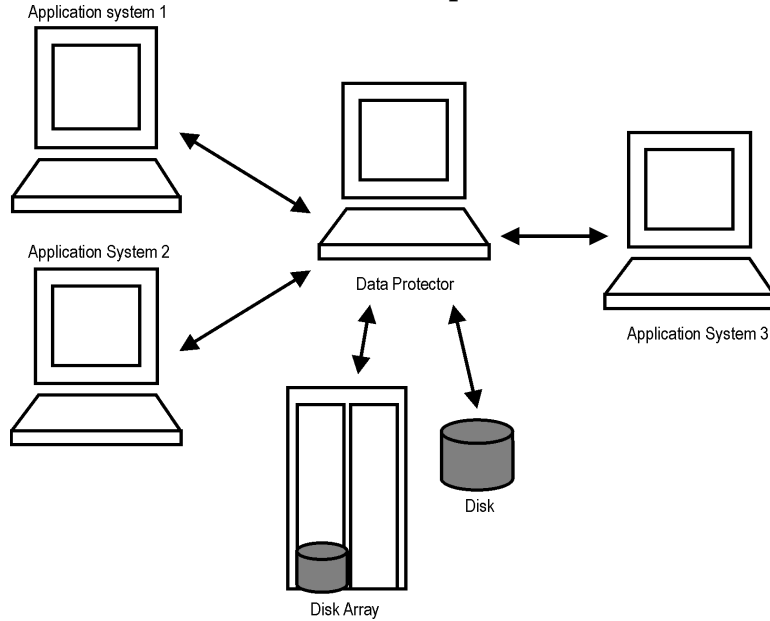



Figure 5-2

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 or network VSS backup

The shadow copy is backed up to tape or to disc from the application client.

- VSS transportable backup

Instead of presenting the volume shadow copies to the application client they are presented to an alternate system, the backup client, which then performs the backup, thus reducing the impact of backup on the application server to the creation of volume shadow copies. A VSS hardware provider that supports transportable snapshots is required.

VSS Transportable Backup Advantages

VSS transportable snapshot backup offers additional advantages over local or network VSS backup:

- The impact on performance on the application server is reduced.
- Following the shadow copy creation, the tape backup can be performed on the backup client at leisure pace.

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 with the Microsoft Exchange 2003 and Microsoft SQL Server ZDB integrations. During the instant recovery, the data in the specified replica (left unchanged for the

purpose of instant recovery) is restored to the application system source volumes. 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.

Prerequisites and Limitations

This is a list of prerequisites and limitations for the Data Protector MS Volume Shadow Copy integration:

Prerequisites

- Before you begin, ensure that you have correctly installed and configured Data Protector, writers and shadow copy providers. Refer to the:
 - *HP OpenView Storage Data Protector Product Announcements, Software Notes, and References* for an up-to-date list of supported versions, platforms, devices, limitations, and other information.
 - *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions on how to install Data Protector on various architectures and how to install the Data Protector MS Volume Shadow Copy integration.
 - Writers and shadow copy providers documentation for instructions on how to install and configure writers and providers on your system.

VSS Transportable Snapshot Prerequisites

For VSS transportable snapshots there are additional prerequisites:

- The backup client must be configured to accept connections from the application client.
- The VSS hardware provider must be installed and configured on the application and backup client. Refer to the provider documentation for details.

Instant Recovery Prerequisite

To be able to perform an instant recovery, the VDS hardware provider must be installed and configured on the application and backup client. Refer to the provider documentation for details.

Limitations

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

- Maximum 64 volumes in a single volume shadow copy set is allowed. The number of shadow copy sets per volume is limited by system resources.

- To run a VSS integration backup, the writer's data must be on an NTFS filesystem.
- The VSS integration backup of writers which store their data on network shared volumes is not supported.
- The Data Protector MS 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 Transportable Snapshot Limitations

- The restore of VSS transportable snapshots is possible only from backup media to the application client.
- Transportable hardware snapshots must be transported out of the cluster before a backup can be performed.
- Only writers that have data on volumes managed by VSS hardware providers, can be included in VSS transportable backups.

Instant Recovery Limitations

- Only writers that have data on volumes managed by VDS hardware providers can be restored using the instant recovery functionality.
- (TBD) 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, refer to the latest support matrices at http://www.openview.hp.com/products/datapro/spec_0001.html.

Integration Concepts

The Data Protector integration with the MS Volume Shadow Copy service provides full support for certified VSS writers. This includes automatic detection of the VSS writers and backup and restore functionality.

Refer to the or *HP OpenView Storage 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 providers refer to the latest support matrices at

http://www.openview.hp.com/products/datapro/spec_0001.html.

Benefits of Using the Integration

The 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 MS Volume Shadow Copy service. Data Protector MS 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)

NOTE

A **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 provided on the level of writer and not dependent on Data Protector functionality. The backup process of the VSS-aware writers consists of the following phases:

1. When you selected writers and components you want to back up and started 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 the Data Protector.
8. Depending on the backup type, in the next step the following two options are possible:
 - With a local or network backup, Data Protector backs up the data from the shadow copy to media and then notifies the VSS service that the shadow copy can be released. VSS issues a command to the shadow copy provider to destroy the shadow copy that has been already backed up. Figure 5-3 shows the relations between the actors of a local or network VSS backup.
 - With a transportable VSS backup, the shadow copy is presented to the backup client. Data Protector backs up the data from the shadow copy on the backup client to the backup media.

After the backup is completed, the vssbar agent on the backup client deletes the shadow copy and disconnects from the application client. Figure 5-4 shows the relations between the actors of a VSS transportable backups.

Figure 5-3

Local or Network VSS Backup

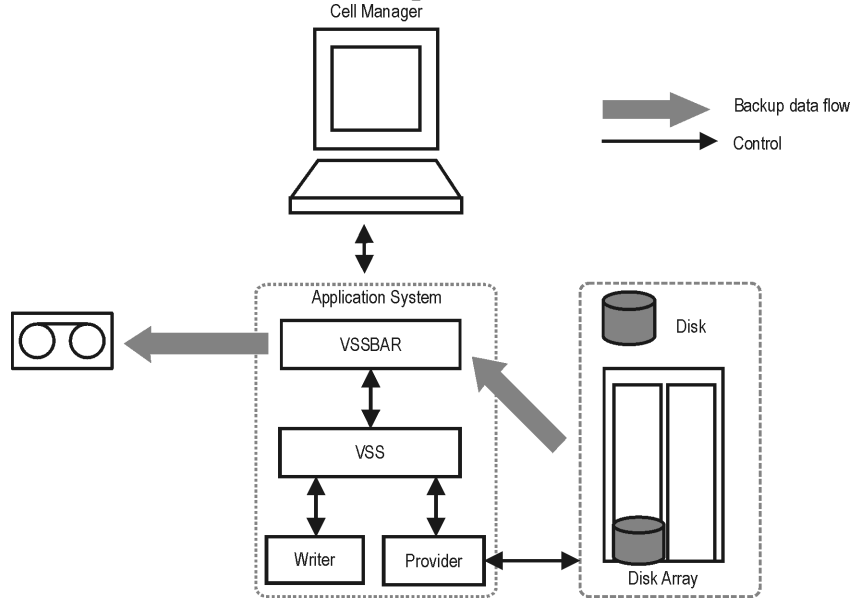
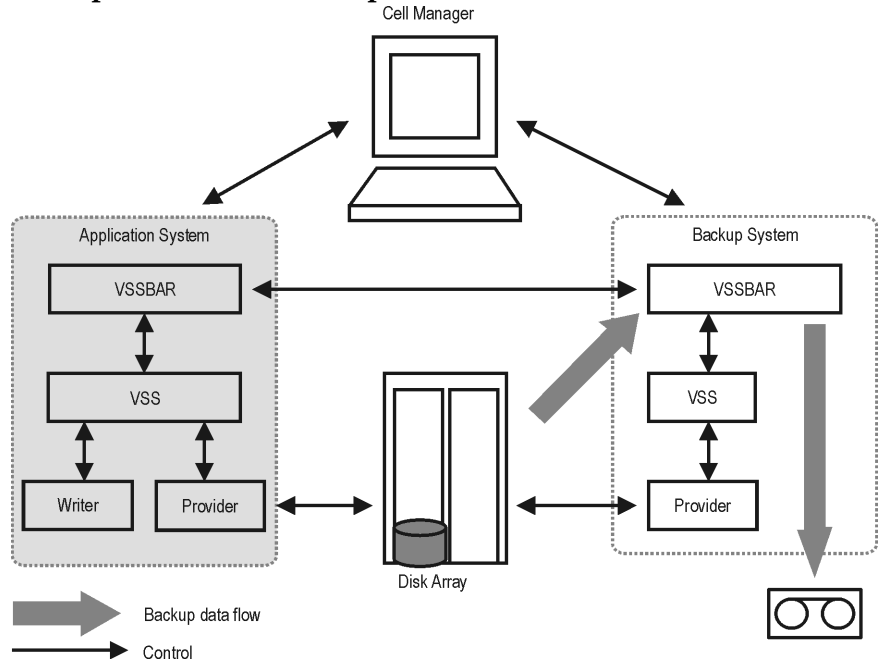


Figure 5-4

Transportable VSS Backup



Restore

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 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.

Instant Recovery

Instant recovery consists of two phases:

- An instant recovery session is started and the relevant files are restored back to the system.
- A writer performs the instant recovery, which is based on the restored data. In case of MSDE and Microsoft Exchange Server 2003 Writers, database recovery is performed automatically by the database engine.

For detailed information on instant recovery, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide* and *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide*.

Configuring the Integration

The Data Protector MS Volume Shadow Copy integration does not require any configuration steps neither on the Data Protector nor on the application side, unless you are configuring a cluster-aware Data Protector VSS integration.

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`
- VDS hardware providers should be present in the list of installed software. Check Control Panel -> Add/Remove Programs.

Configuring the Data Protector VSS Cluster-Aware Integration

The configuration of the Data Protector VSS cluster-aware integration consists of:

1. Configuring an VSS cluster-aware client. Refer to “Configuring a VSS Cluster-Aware Client” on page 367.
2. Configuring a cluster-aware VSS integration backup. Refer to “Configuring Backups for a Cluster-Aware VSS Client” on page 368

Find below an overview of global configuration tasks with cluster-specific steps.

Configuring a VSS Cluster-Aware Client

The client configuration must be performed on one cluster node per one VSS client, since the Data Protector VSS configuration file resides on the Cell Manager.

Configuring Backups for a Cluster-Aware VSS Client

To configure backups for a cluster-aware VSS client, create a Data Protector VSS backup specification, as explained in “Creating Backup Specification Using GUI” on page 376 taking into account the VSS infrastructure specifics described below.

The MS VSS infrastructure does not identify writers that run as cluster resources (for example, cluster-aware writers). Therefore, the MS VSS integration agent cannot distinguish between the cluster-aware and non-cluster-aware writers when creating a backup specification. This means you need to configure different backup specifications for cluster-aware and non-cluster-aware writers.

When backing up cluster-aware writers (such as SQL Server via the MSDE Writer), specify the name of the VSS client system as the virtual server name given in the particular writer resource group.

When backing up writers that are not cluster-aware (such as System Writer or Event Log Writer), specify the name of the VSS client system as the physical node.

Example 5-1

VSS Cluster Specifics

The example below shows why it is necessary to create different backup specifications for cluster-aware and non-cluster-aware writers.

You have `node_A` and `node_B`, and MS Exchange Server 2003 running on a virtual host `exchsvr`. When creating a backup specification, you can select, among others, MS Exchange Writer and Event Log Writer. Suppose, at the time of a backup, Exchange is running on `node_A`. If you create just one backup specification for both writers, the following problems will occur:

- If you select `node_A` as your source host, you have Event Log Writer and Exchange Server 2003 associated with `node_A`. While it is true, that Event Log Writer is a property of the physical node, it is wrong to associate Exchange Server with it, as it is a property of the virtual server.

Suppose that after a failover, MS Exchange Server 2003 is running on `node_B`. When you try to restore the data to `node_A`, the restore will fail because Exchange disks are now owned by `node_B` and you cannot write to them. However, the restore of Event Log Writer will succeed.

- If you select `exchsvr` as your source host, you have Event Log Writer and Exchange Server associated with a virtual server `exchsvr`. While it is true, that Exchange Server 2003 is a property of the virtual server, it is wrong to associate Event Log Writer with it, as it is a property of the physical node.

Suppose that after a failover, MS Exchange Server 2003 is running on `node_B`. When you try to restore Event Log Writer data to `exchsvr`, it will overwrite (or try to overwrite) the data in the Event Log of `node_B` with the data from the Event Log of `node_A`. The restore of Event Log Writer will fail, while Exchange Server will be restored successfully.

Microsoft Exchange Server 2003 Writer Specific Configuration

The Microsoft Exchange database can be successfully backed up only if the consistency check of the replicated datafiles succeeded. To be able to perform the consistency check for transportable snapshots, install the Exchange Management Tool on the backup system.

Configuration Check

To use the instant recovery functionality, consider that 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.

At backup time, Data Protector checks whether the individual components can be selectively restored using the instant recovery functionality. At instant recovery time, Data Protector checks, whether the data that is required for the components restore is available.

In case of the Microsoft Exchange Server 2003 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 the whole writer. 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
- Non-strict
- Disabled

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 applied to the MSDE Writer and the Microsoft Exchange Server 2003 Writer) may create temporary files on components volumes during backup and instant recovery causing the check failure. In such cases, instant recovery will not be possible, though there is no reason for this.

IMPORTANT

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.

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 refer to the latest support matrices at http://www.openview.hp.com/products/datapro/spec_0001.html.

The Data Protector MS 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 `omnic` variable `OB2VSS_SHOWALLWRITERS` must be set to 1 for all writers to be displayed.

Table 4-1 provides a description of VSS writers.

Table 5-1 **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.

Table 5-1 **Writer description**

Writer Name	Description	Restore Method
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
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.

Table 5-1 **Writer description**

Writer Name	Description	Restore Method
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. 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 391.
Microsoft Data Protection Manager 2006 Writer	This is a writer used to back up and restore Microsoft Data Protection Manager 2006. Microsoft Data Protection Manager is a server that creates and stores replicas of clients and uses them for recovering the data on clients	Refer to “Microsoft Data Protection Manager 2006 Writer Restore Specifics” on page 395.

Table 5-1 **Writer description**

Writer Name	Description	Restore Method
Microsoft Exchange 2003 Server Writer	This is a writer used to back up and restore Microsoft Exchange Server 2003. Microsoft Exchange Server 2003 is a mail and groupware server.	Refer to “Microsoft Exchange Server 2003 Writer Restore Specifics” on page 392.
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.
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.

Table 5-1 **Writer description**

Writer Name	Description	Restore Method
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.
TermServLicensing 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

To run backups and restores of the VSS-aware writers, you need to configure the Data Protector MS Volume Shadow Copy integration backup specifications.

To configure the backup using the VSS integration, perform the following steps:

Configuration Steps

1. Configure 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 2003 Writer, the Microsoft SQL Server or Microsoft Exchange Server have to be running on the application system for a backup to start. For more information, see “Data Protector Microsoft SQL Server ZDB Integration” on page 251 and “Data Protector Microsoft Exchange Server ZDB Integration” on page 301 respectively.

Creating Backup Specification Using GUI

The procedure below shows how to back up MS VSS objects using the Data Protector GUI. Some writers have specific limitations. For writers specific limitations, refer to the appropriate sections:

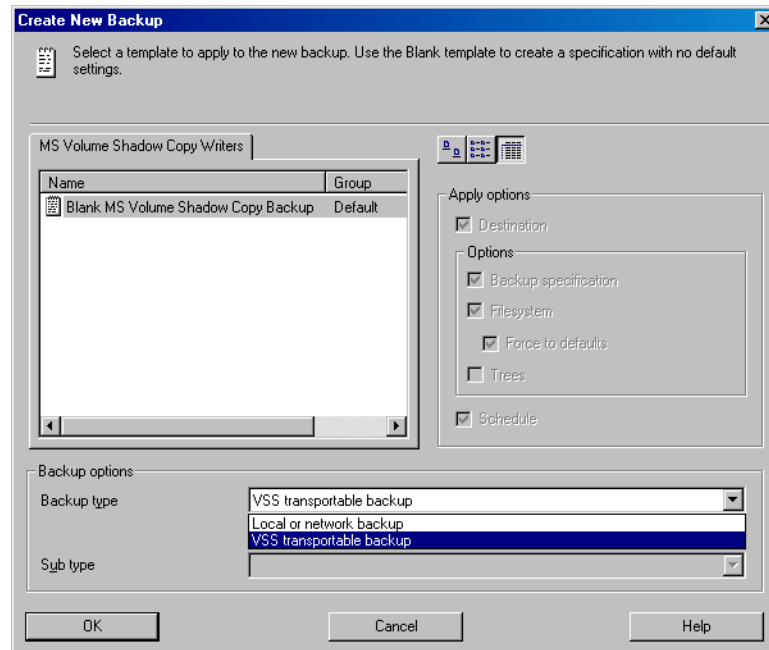
- For Microsoft Exchange Server 2003 specifics, see “Microsoft Exchange Server 2003 Writer Specifics” on page 381.
- For Microsoft Data Protection Manager 2006 specifics, see “Microsoft Data Protection Manager 2006 Writer Specifics” on page 383.

You need to configure different backup specifications for cluster-aware and non-cluster-aware writers. Refer to “Configuring the Data Protector VSS Cluster-Aware Integration” on page 367.

To create a new backup specification for the VSS integration, proceed as follows:

1. In the HP OpenView Storage Data Protector Manager, switch to the Backup context.
2. In the Scoping Pane, expand Backup, and then Backup Specifications.
3. Right-click MS Volume Shadow Copy Writers and then click Add Backup. The Create New Backup dialog box is displayed.
4. In the Create New Backup dialog box, click Blank Microsoft Volume Shadow Copy Backup to select a template.

Figure 5-5 Selecting a Blank Template and VSS Transportable Backup



Select the backup type. You can choose between the following types:

- Local or network backup

This type is used for single host VSS backup. If you want to perform a ZDB backup for instant recovery purposes, you need a hardware provider. Otherwise, no hardware provider is required.

- VSS transportable backup

This option selects the VSS ZDB backup. A hardware provider is required for this type of backup.

5. Specify the following options:

- The name of the application client.

If you are creating a VSS transportable backup specification, the application system drop-down list is disabled.

When backing up cluster-aware writers (such as SQL Server via the MSDE Writer), specify the name of the VSS client system as the virtual server name given in the particular writer resource group.

- For transportable backup, the name of the backup client from where the shadow copy will be backed up to tape.
- For local or network backup, the type of provider: select `Use Hardware Provider` if you want to perform instant recovery.
- If you select the `Track replica for instant recovery` option, 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:

<code>Strict</code>	If any file or folder on the volume does not belong to the component, the disk backup or instant recovery fails.
<code>Non-strict</code>	If any folder on the volume does not belong to the component, the disk backup or instant recovery fails.
<code>Disabled</code>	A check will detect, whether there are more than one component on the volume and whether there is any data besides component's data, but the session will not fail. In case of an instant recovery session, you will lose the data that does not belong to a component, but resides on the same volume as the component.

- The replica type:

<code>Default</code>	The provider decides which replica type is used.
----------------------	--

Mirror/Clone (Plex) A target volume, independent from its source volume, for example snapclone for HP StorageWorks Enterprise Virtual Array.

Snapshot (Differential) A target volume, dependent on its source volume, for example vsnap for HP StorageWorks Enterprise Virtual Array.

A provider may offer one or both types. If you select an unsupported replica type, the backup will fail.

A disk array may also provide more than one snapshot type. In this case use the VSS hardware provider's tools to configure the provider. Refer to the VSS hardware provider's documentation for detailed information on how to configure the provider.

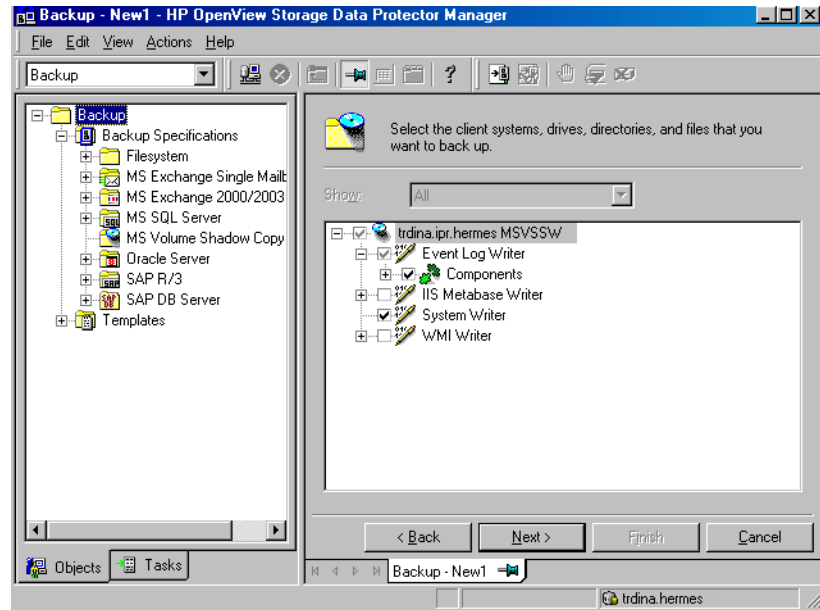
Some providers ignore this selection; in these cases, make sure the provider is configured to create the intended type of copy.

Click Next.

Figure 5-6 VSS Local and Network Backup Options (TBD)

6. Select the backup objects you want to back up.

Figure 5-7 Selecting Backup Objects



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.

In case a writer requires all of its components to be backed up, lower-level items are disabled and you cannot select them. 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.

7. Following the wizard, select the devices, backup options, and schedule your backup.

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`.

TIP

If you are not sure about selecting the backup options, keep the default values.

Refer to online `Help` for details about the options common to all Data Protector backup specifications.

8. 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 MS Volume Shadow Copy Writers backup specification.

9. You can review the newly-created and saved backup specification in the `Backup` context, under the specified group of backup specifications.
10. You can run backup using one of the following methods:
 - Schedule the backup of an existing MS Volume Shadow Copy Writers backup specification using the Data Protector Scheduler.
 - Start an interactive backup of an existing MS Volume Shadow Copy Writers backup specification.

Microsoft Exchange Server 2003 Writer Specifics

The Microsoft Exchange Server 2003 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. Requires Service Pack 1.
- **Copy** - a Full backup, but the logs are not truncated. This type of backup is not intended for use in recovering failed systems. Requires Service Pack 1.

Limitations

- A combination of VSS snapshot backups and incremental stream backups is not possible.
- You can back up only the whole server or full storage groups. Single stores cannot be backed up.
- Incremental and differential backups cannot be mixed in one restore chain.
- Circular logging must be disabled; otherwise, only full backup recovery is possible.
- Only one VSS backup session of the Microsoft Exchange Server 2003 Writer can be running at once on the application client.

Rollforward Recovery

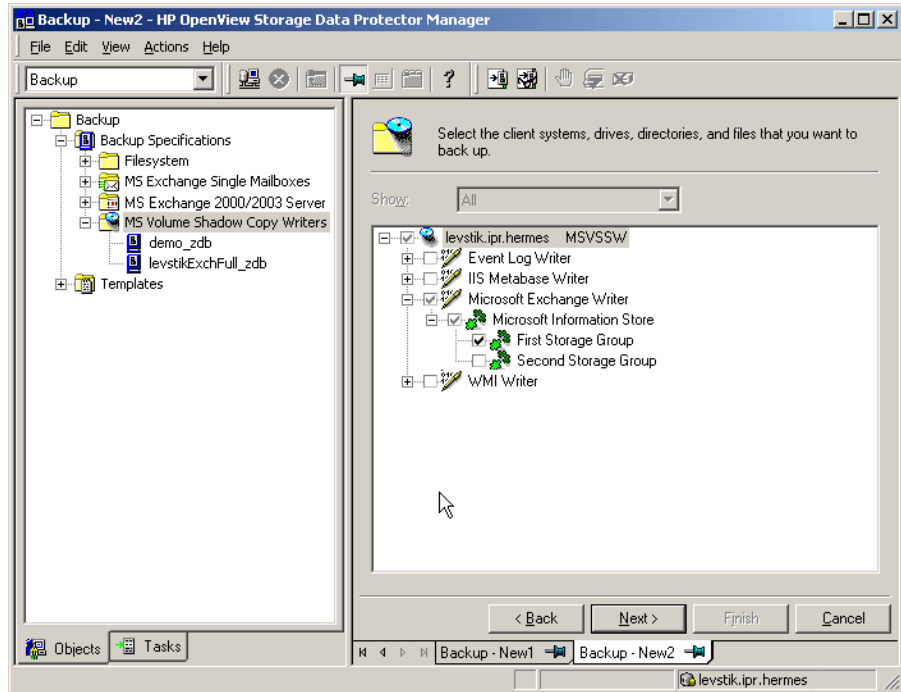
Transaction logs must be backed up to be able to perform the rollforward operation. If you want to combine rollforward recovery with instant recovery, you need to create a separate backup specification for **Incremental** or **Differential** backup with the same object and description as you have in the backup specification for instant recovery.

Consistency Check

The database can be successfully backed up only if the consistency check of the replicated datafiles succeeded.

To disable consistency checking, set the `OB2VSS_EXCHANGE_DISABLE_CONSISTENCY_CHECK` omnirc variable to 1.

Figure 5-8 **Selecting Microsoft Exchange Server 2003 Storage Groups TBD**



Microsoft Data Protection Manager 2006 Writer Specifics

Microsoft Data Protection Manager 2006 (DPM) is a server application that creates replicas of the clients, synchronizes them through LAN, and stores these replicas as snapshots.

The Data Protection Manager writer is used to back up:

- the Data Protection Manager database and the Data Protection Manager Report database
- the *latest version* of the DPM replicas.

IMPORTANT

The DPM uses DPM snapshots for restore. These snapshots are *not* backed up. To be able to recreate DPM snapshots you must manually schedule a backup of the replica each time after the DPM creates a new replica.

Two backup types are supported:

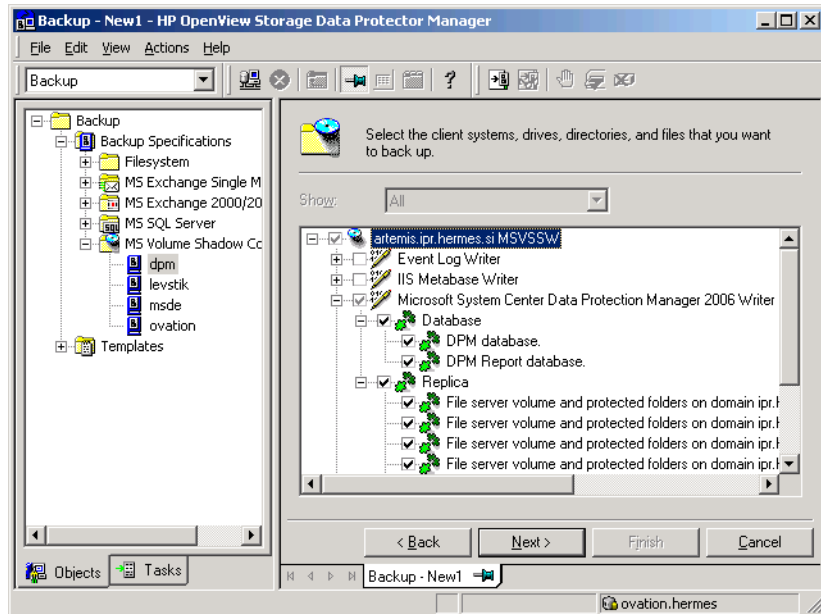
- Full (for the DPM databases and replicas)
- Incremental (replicas only).

If you select unsupported backup types (Copy or Differential) when scheduling the backup, Data Protector will abort the backup and display an error message.

Prerequisite

The MSDE writer (used for backing up the DPM databases) must be installed.

Figure 5-9 Selecting Microsoft Data Protection Manager Database and Replicas



Scheduling the Backup

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

To schedule a MS Volume Shadow Copy Writers backup specification, perform the following steps in the Data Protector GUI:

1. In the HP OpenView Storage 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.

Figure 5-10 Scheduling a Backup (TBD)

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 OpenView Storage 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|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 MS Volume Shadow Copy integration 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 388.
- Using the instant recovery functionality with the MS Exchange and MS SQL Server ZDB integrations. Refer to “Instant Recovery” on page 397.

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 MS 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.

NOTE

Writers requiring custom restore methods are by default not shown by Data Protector. The `omniirc` variable `OB2VSS_SHOWALLWRITERS` must be set to `1` for all writers to be displayed.

Restore Procedure

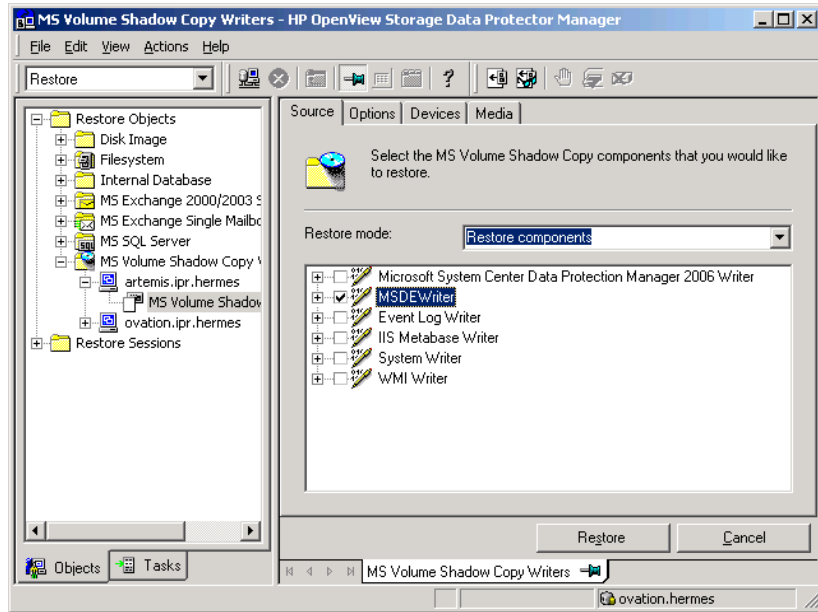
The procedure below shows how to restore MS VSS components using the Data Protector GUI. Some writers require custom restore procedures and/or have specific limitations. See also the appropriate sections:

- For Microsoft Exchange Server 2003 Writer specifics see “Microsoft Exchange Server 2003 Writer Restore Specifics” on page 392.
- For MSDE Writer specifics see “MSDE Writer Restore Specifics” on page 391.
- For Microsoft Data Protection Manager 2006 Writer specifics see “Microsoft Data Protection Manager 2006 Writer Restore Specifics” on page 395.

To restore MS VSS objects using the Data Protector GUI, proceed as follows:

1. In the HP OpenView Storage Data Protector Manager, switch to the Restore context.
2. Expand Restore and Microsoft Volume Shadow Copy Writers and select the client from which you want to restore the data. In the Results Area, a list of writers, which were backed up on this client, is displayed.
3. Select the restore mode:
 - To restore components using the Volume Shadow Copy Service, select Restore Components.
 - To restore individual files or a group of files without using the Volume Shadow Copy service, select Restore files.
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).

Figure 5-11 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 backup version is selected by default, however, you can select a different version from the drop-down list.

5. In the **Options** property page, select the MS Volume Shadow Copy specific restore options. Refer to “Restore Options” on page 390.
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. Refer to the online Help Index: TBD for more information on how to perform a restore using another device.

7. Click the `Restore MS Volume...` button. 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.

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 MS 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.

Figure 5-12 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 unlock user databases before the restore, while SQL service will have to be stopped manually in order to restore the system databases.

Microsoft Exchange Server 2003 Writer Restore Specifics

Microsoft Exchange Server 2003 Writer is used to restore Microsoft Exchange Server 2003 database files.

When restoring from a Microsoft Exchange 2003 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 393.
- 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 394.

Limitations

The following limitations apply when restoring Microsoft Exchange Server 2003:

- Shadow copies cannot be restored to alternate locations on the backup client.
- You cannot restore the shadow copy to the Recovery Storage Group.
- Rollforward recovery cannot be performed after a point-in-time restore.

Rollforward Recovery from the Loss of One or More Databases

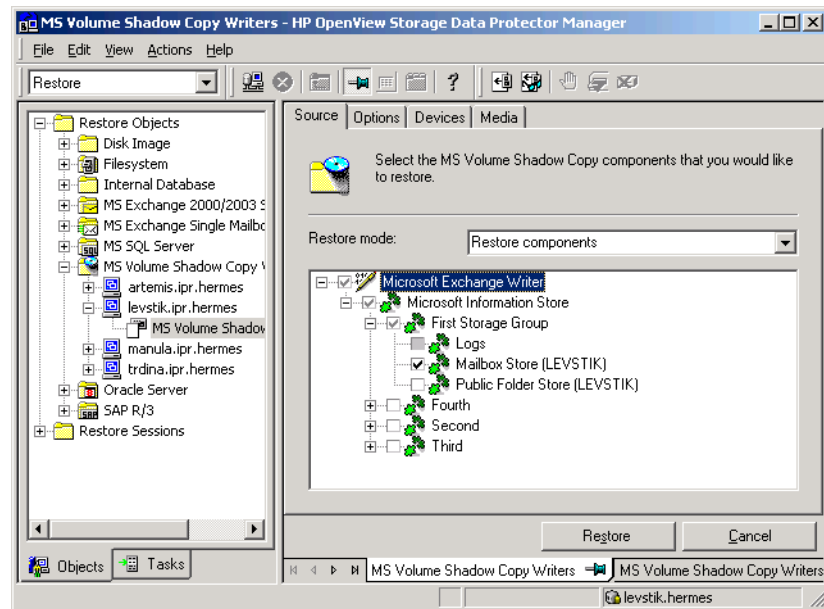
For a rollforward recovery:

1. 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 and Microsoft Volume Shadow Copy Writers and select the client from which you want to restore the data

In the Results Area, expand the Microsoft Exchange Server 2003 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.

Figure 5-13

Selecting Microsoft Exchange Server 2003 Stores for Rollforward Recovery



3. Proceed as with general VSS writer restore. See “Restore Procedure” on page 388 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. Start Exchange System Manager and check if the storage group is already unmounted. If not, unmount the whole group.
2. Switch to the Restore context. Expand Restore and Microsoft Volume Shadow Copy Writers and select the client from which you want to restore the data.

In the Results Area, expand the Microsoft Exchange 2003 writer and select the whole storage group. Do not select individual stores.

Figure 5-14

Selecting Microsoft Exchange Server 2003 Stores for Point-in-Time Restore



3. Proceed as with general VSS writer restore. See “Restore Procedure” on page 388 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 Data Protection Manager 2006 Writer Restore Specifics

When restoring the DPM writer, you can:

- Restore the DPM *server* first and then use the DPM to restore clients (for example, if only a DPM DB or individual replicas are lost).
- Restore individual DPM *clients* directly, without using the DPM server (for example, if you cannot restore the DPM server or if you want to avoid the additional step of recreating the DPM snapshot). When restoring the DPM clients directly you can select between component restore and file restore modes.

NOTE

Although the Data Protection Manager database can also be restored using the MSDE writer, this method is not recommended, because DPM is *not* shut down automatically as with the DPM writer. If you really need to use this writer, shut down the DPM server manually.

Limitations

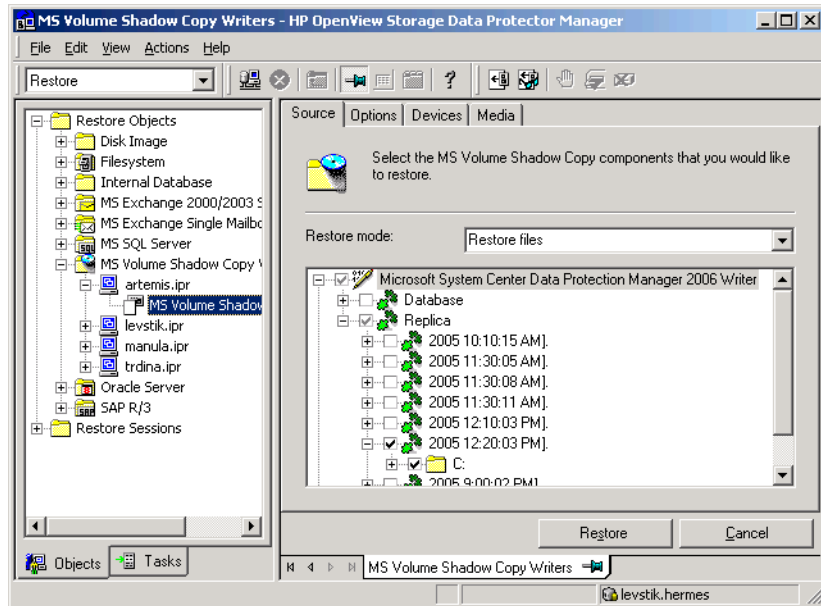
- Restore to another server is not supported by the Data Protection Manager writer.
- Parallel restore to different clients is not supported.

Restore the DPM Server First

1. Switch to the Restore context. Expand Restore and Microsoft Volume Shadow Copy Writers and select the client from which you want to restore the data.
2. In the Results Area, expand the DPM writer and select the components for restore:
 - If the whole DPM server was lost, select both, the Data Protection Manager databases and the replicas.

- If only one or more replicas were lost, select only the necessary replicas.

Figure 5-15 Restoring the Microsoft Data Protection Manager 2006 Client



Proceed as with general VSS writer restore. See “Restore Procedure” on page 388 for the general VSS writer restore procedure.

3. Use the DPM to restore individual clients.

IMPORTANT

The DPM console does not automatically check for new or restored snapshots. Before you can start the restore of clients, you must use the Data Protection Manager to recreate a DPM snapshot.

- a. In the DPM console, open the Recovery context. Under the Browse tab, select the server, right click on the restored replica, and select Create shadow copy now.
- b. Select and restore the new snapshot to the client.

Restore the DPM Clients Directly

1. Switch to the Restore context. Expand Restore and Microsoft Volume Shadow Copy Writers and select the client from which you want to restore the data.
2. Select the restore modes:
 - Restore Components
Use this mode *only* if the client to which you want to restore supports VSS, for example if you restore to Windows 2003 clients. You can restore only entire replicas.
 - Restore Files
The client does not need to support VSS and you can restore individual folders or files.
3. When selecting the DPM writer for restore, select *only* the Replica components. Do not select the DPM database.
4. Click the Options tab, and under Restore to another client enter the name of the target client. Click Next.
5. Proceed as with general VSS writer restore. See “Restore Procedure” on page 388 for the general VSS writer restore procedure.

Instant Recovery

Refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide* and *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide* for general information on instant recovery.

The Data Protector VSS integration provides instant recovery of the MSDE Writer and the Microsoft Exchange Server 2003 Writer.

Instant Recovery Procedure

You can *selectively* restore separate writer's components only 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.

Before starting an instant recovery, Data Protector checks, whether the data that is required for the components restore is available. If the check fails, an instant recovery session also fails to ensure, that no data that does not belong to the components is lost.

Limitation

Instant recovery cannot be performed after a point-in-time restore. To be able to perform an instant recovery, you need to run a new backup with the instant recovery options set.

Procedure

1. Manually dismount the databases to be recovered.
2. In the HP OpenView Storage Data Protector Manager, switch to the Instant Recovery context.
3. In the Scoping Pane, expand Restore Objects, MS Volume Shadow Copy Writers, the writer you want to restore and double-click the backup session from which you want to perform a recovery and select a component (replica) for backup.

Figure 5-16

Selecting Writers Components for Instant Recovery (TBD)

4. In the Source property page, specify writers and/or components for recovery. Select the configuration check mode:

Strict	If any file or folder on the volume does not belong to the component, the instant recovery fails.
Non-strict	If any folder on the volume does not belong to the component, the instant recovery fails.
Disabled	A check will detect, whether there are more than one component on the volume and whether there is any data besides component's data, but the session will not fail. In this case, you will lose the data that does not belong to a component, but resides on the same volume.

Click Restore.

5. Re-mount the database stores.

Clusters

To perform an instant recovery in the cluster environment, use the above instant recovery procedure. Use virtual hostnames.

Microsoft Exchange Server 2003 Writer Specifics

This section provides specific information on the instant recovery and rollforward recovery of the Microsoft Exchange Server 2003 Writer.

When recovering Microsoft Exchange Server 2003 Writer, the following configuration scenarios are possible:

- Transaction logs and database stores are on the same volume.
You cannot select database stores for instant recovery in the GUI. If the transaction logs and/or database stores are lost, the whole storage group should be recovered.
- Transaction logs and database stores are on different volumes.
You can select database stores for instant recovery in the GUI. If a data store is lost, it can be recovered separately. If the transaction logs are lost, the whole storage group should be recovered.

Database Recovery

You can run a database recovery from the `Instant Recovery` context of the Data Protector GUI. This possibility 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.

Troubleshooting

This section contains a list of problems you might encounter when using the Data Protector Microsoft Volume Shadow Copy integration.

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

Before You Begin

- ✓ Ensure that the latest official Data Protector patches are installed. See the online Help index: “patches” on how to verify this.
- ✓ See the *HP OpenView Storage Data Protector Product Announcements, Software Notes, and References* for general Data Protector limitations, as well as known problems and workarounds.
- ✓ See http://www.openview.hp.com/products/datapro/spec_0001.html for an up-to-date list of supported versions, platforms, and other information.

Backup Problems

Problem

After a backup to tape was performed, the VSS shadow copy cannot be deleted.

This can happen if the shadow copy was (un)presented or manipulated using a non-VSS API based application, for example, the disk array’s management appliance.

Action

Rebuild the VSS database:

1. Manually delete all shadow copies corresponding to the application server’s disks.
2. Manually delete all shadow copies (snapshots) presented to the backup server.
3. Delete the VSS database on the application and backup host. The location of the files is defined with the following registry keys:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\
BackupRestore\FilesNotToBackup\VSS Service DB
```

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\  
BackupRestore\FilesNotToBackup\VSS Service Alternate DB
```

NOTE The above procedure removes all information about existing snapshots from the VSS database and should be used carefully.

Problem **The presentation time exceeds the BSM time-out and the session is aborted.**

When HP StorageWorks EVA is used as a VSS Hardware Provider, it may take up to 4 or 5 minutes per volume in a snapshot set to be presented to the backup host (or to the application host in case of a local backup). In the case when there are many volumes in a snapshot, the total presentation time may exceed the BSM time-out for waiting new connections. BSM will abort the session with the following error:

```
Bar backup session was started but no client connected in 600  
seconds.
```

Action Set the global variable `SmWaitForFirstClient` to a longer time-out. The default time-out is 10 minutes. The new value should be `5*NumberOfVolumes`, where `NumberOfVolumes` is the number of volumes in a snapshot set being imported. The presentation time is much shorter, when the management appliance version 3.0 is used.

Problem **When HP StorageWorks EVA is used as a VSS Hardware Provider, the option Snapshot Type is ignored by the provider.**

Action Use the HP StorageWorks EVA configuration tool to select the desired type of a shadow copy (for example snapshot, vsnap or snapclone).

Problem **Shadow copies are not deleted at the end of the failed backup session.**

If importing of shadow copies fails on the backup client (for any reason), the created shadow copies cannot be deleted by the VSS service. Therefore, VSSBAR on the application client tries to import the shadow copies in order to delete them. If this import fails, you have to delete the shadow copies manually.

Action To delete the shadow copies proceed as follows:

1. Stop the provider service using the Service Manager.
2. Stop the VSS and VDS services.
3. Delete the VSS Snapshot Database on the backup server and reboot the server.

To locate the VSS Snapshot Database files, use the registry editor to find the value of the following registry keys:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\BackupRestore\FilesNotToBackup\VSS Service DB
```

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\BackupRestore\FilesNotToBackup\VSS Service Alternate DB
```

4. Connect to the management appliance, identify the shadow copies and delete them.
5. Run the backup again. If the same error persists, repeat the procedure and reboot the system.

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 client, 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:

- Run the backup again.
- Limit the number of users that are accessing the management appliance.

- Reduce the number of volumes in a snapshot set. For example, keep different store's data on separate volumes or create a backup specification dedicated to each store instead of one specification for the whole server.

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

Some components are not restored during the restore session.

If a component cannot be restored to the location specified in Writer Metadata Document (for example, if this location is locked or it is not possible to perform regular restore), this component will be skipped during the restore procedure.

Action

Specify a location, where skipped files will be redirected in case of failure, by setting the OB2VSS_DUMPTO environmental variable in the `<Data_Protector_home>\omnirc` file. Restart the Data Protector services to apply the changes in the `omnirc` file.

Example

If you want the files that are skipped during the restore to be copied to the `F:\Restore` directory, set `OB2VSS_DUMPTO=F:\Restore` in the `omnirc` file. In case the SQL component `Company` was skipped during the restore, it will be copied to the specified directory as follows:

```
F:\Restore\2002-12-09-23\G\SQL\Log\Company.ldf  
F:\Restore\2002-12-09-23\G\SQL\Log\Company.mdf
```

The pathname includes the backup session ID and the pathname to the original location.

Problem

Instant recovery of the Microsoft Exchange Server 2003 Writer fails.

This problem may occur, if the Microsoft Exchange 2003 Writer is not in the stable state. Check this by running `vssadmin list writers` from the command prompt.

Action

Bring the Exchange Server 2003 Writer to a stable state by restarting the Microsoft Exchange Information Store.

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:

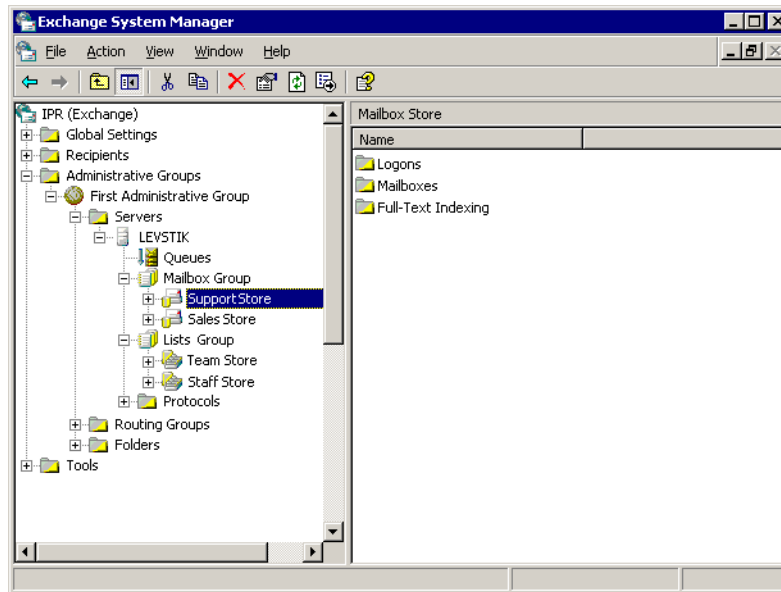
- A Microsoft Exchange Server 2003 is running on the application client and is connected to a HP StorageWorks Enterprise Virtual Array. A tape drive is connected to the application client for restore purposes.
- A separate backup client is connected to a tape library.
- Two storage groups are configured 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 400.

Figure 5-17 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 client 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 client 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 405 and the backup policy given in the example was implemented.

Example 5-2

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 client. 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 5-3

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 client. 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 A-3
- “ZDB Integrations Omnirc Variables” on page A-9

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 Concept” on page 14 and to “Oracle Proxy-Copy ZDB Concept” on page 21 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 5.

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 source volumes to be replicated, to other locations, perform the following steps:

1. For Oracle8i, display the control file information:

```
SQL> select name from v$controlfile
```

For Oracle9i/10g, 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/init<SID>.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 using:

```
$ 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 Oracle8i

To move the Oracle8i *control files* from source volumes (/opt/oracle/oradata) to a non-replicated source volume (/oracle8/oractl):

1. Display the control file information:

```
SQL> select name from v$controlfile
/opt/oracle/oradata/control01.ctl
```



```
/opt/oracle/oradata/control02.ctl
```

```
/opt/oracle/oradata/control02.ctl
```

2. Shut down the database:

```
SQL> shutdown
```

3. Move the control files:

```
mv /opt/oracle/oradata/control01.ctl /oracle8/oractl
```

```
mv /opt/oracle/oradata/control02.ctl /oracle8/oractl
```

```
mv /opt/oracle/oradata/control03.ctl /oracle8/oractl
```

4. Edit the CONTROL_FILES parameter in the database's initialization parameter file:

```
control_files = ("/oracle8/oradata/control01.ctl",
"/oracle8/oractl/control02.ctl",
"/oracle8/oractl/control03.ctl")
```

5. Restart the database:

```
SQL> startup
```

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:

```
SQL> alter system set
control_files='/oracle/logs/RCVCAT/control01.ctl',
'/oracle/logs/RCVCAT/control02',
'/oracle/logs/RCVCAT/control03.ctl' scope=spfile;
```

3. Shut down the database:

```
SQL> shutdown
```

Reconfiguring an Oracle Instance for Instant Recovery

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. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for information on how to use the omnirc file.

For information on Data Protector ZDB agents omnirc file variables, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide*.

This section explains the omnirc file variables that can be set for Data Protector ZDB integrations.

ZDB_ORA_INCLUDE_CF_OLF: An Oracle8i/9i/10g (if the Oracle integration is used) and TBD Oracle8/8i/9i (if the SAP R/3 integration is used) related variable.

NOTE

This variable is not supported on EMC.

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 in the following cases:

- Oracle8.0.x is used with the Data Protector Oracle integration.
- Offline backup is performed using the Data Protector SAP R/3 integration.

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:

- By default (if this variable is set to 0), during a ZDB session, Data Protector creates target volumes only for the source volumes containing Oracle8i/9i/10g (if the Oracle integration is used) or TBD

Oracle8/8i/9i (if the SAP R/3 integration is used) 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 concept when this variable is set to 0, see “Oracle Backup Set ZDB Concept” on page 14 and “Oracle Proxy-Copy ZDB Concept” on page 21. For SAP R/3 backup and restore concept when this variable is set to 0, see “SAP R/3 ZDB Concept” on page 161.

- 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 and the control files and redo logs are on the same source volumes as datafiles, they are overwritten during instant recovery.

Opening the Database on the Backup System

If you want to successfully open the database on the backup system for *other* purposes than Data Protector, note the following:

- 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:

- ✓ If the Oracle integration is used: Oracleg datafiles, Oracle control file, and Oracle online redo logs are all installed on disk array; they must not be installed anywhere on the network.

If the SAP R/3 integration is used: TBD Oracle8/8i/9i datafiles and Oracle8/8i/9i control file are all installed on disk array; they must not be installed anywhere on the network.

- ✓ If the Oracle integration is used: Oracle datafiles on one hand, and Oracle control file and Oracle online redo log files on the other are not installed on the same disk image, logical volume, or filesystem.

If the SAP R/3 integration is used: TBD Oracle8/8i/9i datafiles on one hand, and Oracle8/8i/9i control file on the other are not installed on the same disk image, logical volume, or filesystem.

Refer to Figure A-1 on page A-11 and Figure A-2 on page A-12 for Oracle backup and restore concept when this variable is set to 1. Refer to Figure A-3 on page A-13 for SAP R/3 backup and restore concept when this variable is set to 1.

Figure A-1 Oracle Proxy-copy ZDB and Restore Concept when the ZDB_ORA_INCLUDE_CF_OLF variable is set to 1

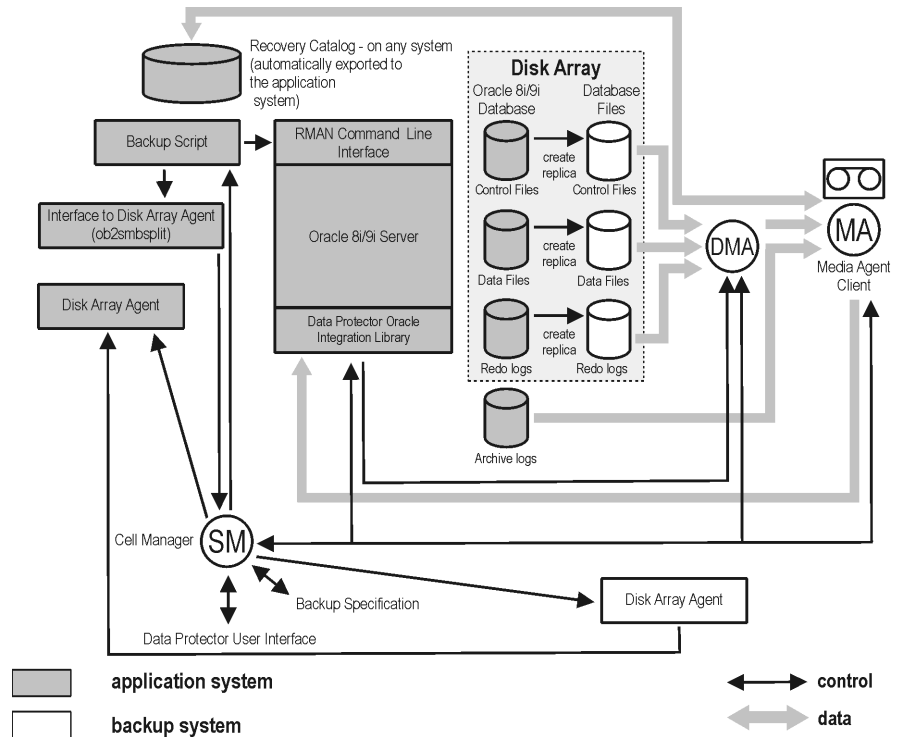


Figure A-2

Oracle Backup Set ZDB and Restore Concept when the ZDB_ORA_INCLUDE_CF_OLF variable is set to 1

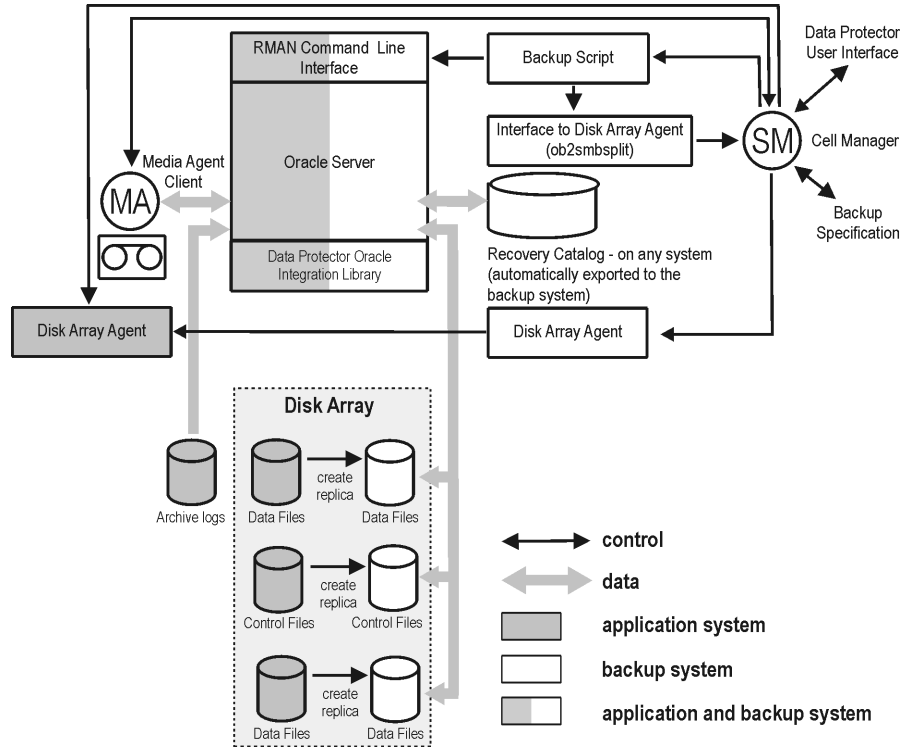
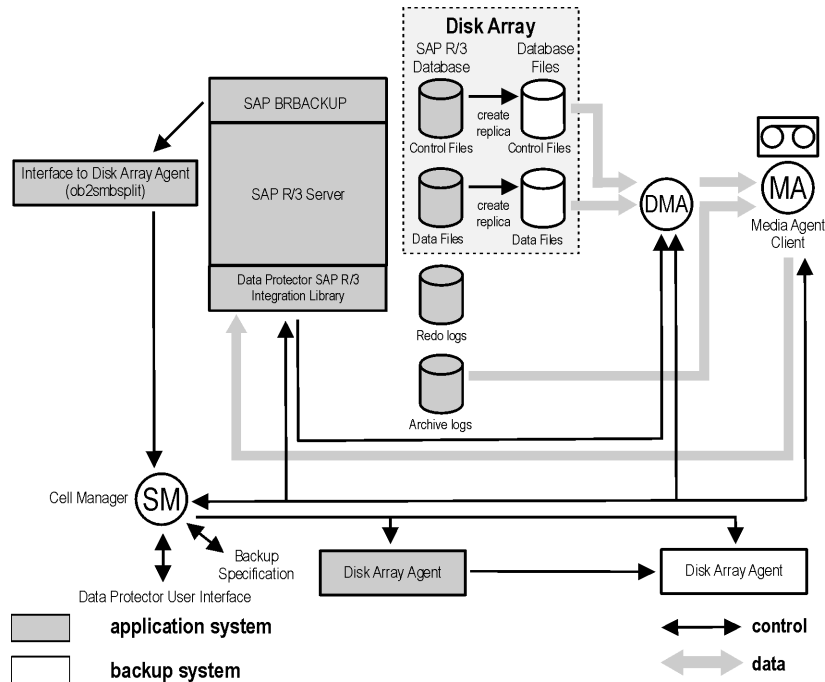


Figure A-3

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: An Oracle9i/10g 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 in the following cases:

- Oracle8/8i is used with the Data Protector Oracle integration.
- 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 Oracle9i/10g 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: An Oracle8i/9i/10g (if the Oracle integration is used) and TBD Oracle8/8i/9i (if the SAP R/3 integration is used) 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 controlfile, 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 and 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.

access rights

See user rights.

ACSLs (*StorageTek specific term*)

The Automated Cartridge System Library Server (ACSLs) software that manages the Automated Cartridge System (ACS).

Active Directory (*Windows specific term*)

The directory service in a Windows network. It contains information about resources on the network and makes them accessible to users and applications. The directory services provide a consistent way to name, describe, locate, access and manage resources regardless of the physical system they reside on.

AML (*EMASS/GRAU specific term*)

Automated Mixed-Media library.

application agent

A component needed on a client to back up or restore online database integrations.

See also Disk Agent.

application system (*ZDB specific term*)

A system the application or database runs on. The application or database data is located on source volumes.

See also backup system and source volume.

archived redo log (*Oracle specific term*)

Also called offline redo log. If the Oracle database operates in the ARCHIVELOG mode, as each online redo log is filled, it is copied to one (or more) archived log destination(s). This copy is the Archived Redo Log. The presence or absence of an Archived Redo Log is determined by the mode that the database is using:

- ARCHIVELOG - The filled online redo log files are archived before they are reused. The database can be recovered from an instance and disk failure. The “hot” backup can be performed only when the database is running in this mode.
- NOARCHIVELOG - The filled online redo log files are not archived.

See also online redo log.

archive logging (*Lotus Domino Server specific term*)

Lotus Domino Server database mode where transaction log files are overwritten only after they have been backed up.

ASR Set

A collection of files stored on several diskettes, required for proper reconfiguration of the replacement disk

Glossary

(disk partitioning and logical volume configuration) and automatic recovery of the original system configuration and user data that was backed up during the full client backup.

These files are stored as an ASR archive file on the Cell Manager (in `<Data_Protector_home>\Config\Server\dr\asr` on a Windows Cell Manager or in `/etc/opt/omni/server/dr/asr/` on a UNIX Cell Manager) as well as on the backup medium. The ASR archive file is extracted to three diskettes for 32-bit Windows systems or four diskettes for 64-bit Windows systems after a disaster occurs. You need these diskettes to perform ASR.

autochanger

See **library**

autoloader

See **library**

BACKINT (*SAP R/3 specific term*)

SAP R/3 backup programs can call the Data Protector `backint` interface program via an open interface, which enables them to communicate with Data Protector software. For backup and restore, SAP R/3 programs issue orders for the Data Protector `backint` interface.

backup API

The Oracle interface between the Oracle `backup/restore` utility and the `backup/restore` media management layer. The

interface defines a set of routines to allow the reading and writing of data to the backup media, the creation, searching and removing the backup files.

backup chain

This relates to a situation where full and incremental backups are performed. Based on the level of the incremental backups used (`Incr`, `Incr 1`, `Incr 2`, and so on), simple or rather complex dependencies of incrementals to previous incrementals can exist. The backup chain are all backups, starting from the full backup plus all the dependent incrementals up to the desired point in time.

backup device

A device configured for use with Data Protector, which can write data to and read data from storage media. This can be, for example, a standalone `DDS/DAT` drive or a library.

backup generation

One backup generation includes one full backup and all incremental backups until the next full backup.

backup ID

An identifier of an integration object that equals the session ID of the backup of this object. The backup ID is preserved when an object is copied, exported, or imported.

Glossary

backup object

A backup unit that contains all items backed up from one disk volume (logical disk or mount point). The backed up items can be any number of files, directories, or the entire disk or mount point. Additionally, a backup object can be a database entity or a disk image (rawdisk).

A backup object is defined by:

- **Client name:** hostname of the Data Protector client where the backup object resides.
- **Mount point:** the access point in a directory structure (drive on Windows and mount point on UNIX) on the client where the backup object is located.
- **Description:** uniquely defines backup objects with identical client name and mount point.
- **Type:** backup object type (for example filesystem or Oracle).

backup owner

Each backup object in the IDB has an owner. The default owner of a backup is the user who starts the backup session.

backup session

A process that creates a copy of data on storage media. The activities are

specified in a backup specification or an interactive session. All clients configured in one backup specification are backed up together in one backup session using the same backup type (full or incremental). The result of a backup session is a set of media, which was written to, also called the backup or media set.

*See also **incremental backup** and **full backup**.*

backup set

A complete set of integration objects associated with a backup.

backup set (*Oracle specific term*)

A logical grouping of backed up files that are created using the RMAN backup command. A backup set is a complete set of files associated with a backup. The files can be multiplexed to increase performance. A backup set contains either datafiles or archived logs, but not both together.

backup specification

A list of objects to be backed up, together with a set of devices or drives to be used, backup options for all objects in the specification, days and time that you want backups to be performed. The objects are entire disks/volumes or parts of them such as files, directories, or even the Windows

Glossary

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 to the backup system. *See also* **HP StorageWorks Disk Array XP LDEV, CA, Main Control Unit, application system, and backup system.**

BC Process (*EMC Symmetrix specific term*)

A protected storage environment solution that has defined specially configured EMC Symmetrix devices as mirrors or Business 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

Glossary

maintain internal copies of HP StorageWorks Virtual Array LUNs for data backup or data duplication within the same virtual array. The copies (child or Business Copy LUNs) can be used for various purposes, such as backup, data analysis or development. When used for backup purposes, the original (parent) LUNs are connected to the application system and the Business Copy (child) LUNs are connected to the backup system.

See also **HP StorageWorks Virtual Array LUN, application system, and backup system.**

BCV (*EMC Symmetrix specific term*) Business Continuance Volumes, or BCV devices, are dedicated SLDs that are pre-configured in the ICDA on which the business continuation operation runs. BCV devices are assigned separate SCSI addresses, differing from the addresses used by the SLDs they mirror. The BCV devices are used as splittable mirrors of the primary EMC Symmetrix SLDs that need to be protected.

See also **BC** and **BC Process.**

Boolean operators

The Boolean operators for the full text search functionality of the online Help system are AND, OR, NOT, and NEAR. Used when searching, they enable you to define your query precisely by creating a relationship between search terms. If no operator is specified in a

multi-word search, AND is used by default. For example, the query manual disaster recovery is equivalent to manual AND disaster AND recovery.

boot volume/disk/partition

A volume/disk/partition with files required for the initial step of the boot process. Microsoft terminology defines the boot volume/disk/partition as a volume/disk/partition containing the operating system files.

BRARCHIVE (*SAP R/3 specific term*)

An SAP R/3 backup tool that allows you to archive redo log files. BRARCHIVE also saves all the logs and profiles of the archiving process.

See also **SAPDBA, BRBACKUP** and **BRRESTORE.**

BRBACKUP (*SAP R/3 specific term*)

An SAP R/3 backup tool that allows an online or offline backup of the control file, of individual data files, or of all tablespaces and, if necessary, of the online redo log files.

See also **SAPDBA, BRARCHIVE** and **BRRESTORE.**

BRRESTORE (*SAP R/3 specific term*)

An SAP R/3 tool that can be used to restore files of the following type:

- Database data files, control files, and online redo log files saved with BRBACKUP

Glossary

- Redo log files archived with BRARCHIVE
- Non-database files saved with BRBACKUP

You can specify files, tablespaces, complete backups, log sequence numbers of redo log files, or the session ID of the backup.

See also **SAPDBA**, **BRBACKUP** and **BRARCHIVE**.

BSM

The Data Protector Backup Session Manager controls the backup session. This process always runs on the Cell Manager system.

CA (*HP StorageWorks Disk Array XP specific term*)

Continuous Access XP allows you to create and maintain remote copies of HP StorageWorks Disk Array XP LDEVs for purposes such as data duplication, backup, and disaster recovery. CA operations involve the main (primary) disk arrays and the remote (secondary) disk arrays. The main disk arrays contain the CA primary volumes (P-VOLs), which contain the original data and are connected to the application system. The remote disk arrays contain the CA secondary volumes (S-VOLs) connected to the backup system.

See also **BC** (*HP StorageWorks Disk*

Array XP specific term), **Main Control Unit** and **HP StorageWorks Disk Array XP LDEV**.

CAP (*StorageTek specific term*)

Cartridge Access Port is a port built into the door panel of a library. The purpose is to enter or eject media.

catalog protection

Defines how long information about backed up data (such as file names and file versions) is kept in the IDB.

See also **data protection**.

CDB

The Catalog Database is a part of the IDB that contains information about backups, object copies, restores, media management sessions, and backed up data. Depending on the selected logging level, it also contains file names and file versions. This part of the database is always local to the cell.

See also **MMDB**.

CDF file (*UNIX specific term*)

A Context Dependent File is a file consisting of several files grouped under the same pathname. The system ordinarily selects one of the files using the context of the process. This mechanism allows machine dependent executables, system data, and device files to work correctly from all hosts in a cluster while using the same pathname.

Glossary

cell

A set of systems that are under the control of a Cell Manager. The cell typically represents the systems on a site or an organizational entity, which are connected to the same LAN. Central control is available to administer the backup and restore policies and tasks.

Cell Manager

The main system in the cell where the essential Data Protector software is installed and from which all backup and restore activities are managed. The GUI used for management tasks can be located on a different system. Each cell has one Cell Manager system.

centralized licensing

Data Protector allows you to configure centralized licensing for the whole enterprise environment consisting of several cells. All Data Protector licenses are installed and kept on the Enterprise Cell Manager system. You can then allocate licenses to specific cells to suit your needs.

See also MoM.

Centralized Media Management Database (CMMDB)

See CMMDB.

channel (*Oracle specific term*)

An Oracle Recovery Manager resource allocation. Every allocated channel starts a new Oracle process, which

performs backup, restore, and recovery actions. The type of channel allocated determines the type of media used:

- type “disk”
- type ‘SBT_TAPE’

If the specified channel is type ‘SBT_TAPE’ and Oracle is integrated with Data Protector, the server process will attempt to read backups from or write data files to Data Protector.

circular logging (*Microsoft Exchange Server and Lotus Domino Server specific term*)

Circular logging is a Microsoft Exchange Server database and Lotus Domino Server database mode, in which the transaction log file contents are periodically overwritten after the corresponding data is committed to the database. Circular logging reduces disk storage space requirements.

client backup

A backup of all filesystems mounted on a client. Filesystems mounted on the client after the backup specification was created are not automatically detected.

client backup with disk discovery

A backup of all filesystems mounted on a client. When the backup starts, Data Protector discovers the disks on the clients. Client backup with disk

Glossary

discovery simplifies backup configuration and improves backup coverage of systems that often mount or dismount disks.

client or client system

Any system configured with any Data Protector functionality and configured in a cell.

cluster-aware application

It is an application that supports the cluster Application Programming Interface. Each cluster-aware application declares its own critical resources (disk volumes (on Microsoft Cluster Server), volume groups (on MC/ServiceGuard), application services, IP names and addresses ...).

CMD Script for OnLine Server

(Informix specific term)

Windows CMD script that is created in INFORMIXDIR when Informix OnLine Server is configured. The CMD script is a set of system commands that export environment variables for OnLine Server.

CMMDB

The Data Protector Centralized Media Management Database (CMMDB) is the result of merging MMDBs from several cells in the MoM environment. It allows you to share high-end devices and media across multiple cells in a MoM environment. One cell can control the

robotics, serving the devices that are connected to systems in other cells. The CMMDB must reside on the MoM Manager. A reliable network connection between the MoM cell and the other Data Protector cells is highly recommended
See also MoM.

COM+ Registration Database

(Windows specific term)

The COM+ Registration Database and the Windows Registry store COM+ application attributes, class attributes, and computer-level attributes. This guarantees consistency among these attributes and provides common operation on top of these attributes.

command-line interface

A set of DOS and UNIX like commands that you can use in shell scripts to perform Data Protector configuration, backup, restore, and management tasks.

Command View (CV) EVA *(HP*

StorageWorks EVA specific term)

The user interface that allows you to configure, manage, and monitor your HP StorageWorks EVA storage system. It is used to perform various storage management tasks, for example, creating virtual disk families, managing storage system hardware, creating snapclones and snapshots of virtual disks. The Command View EVA software runs on the HP OpenView

Glossary

Storage Management Appliance, and is accessed by a Web browser.

See also **HP StorageWorks EVA Agent (legacy)** and **HP StorageWorks EVA SMI-S Agent**.

concurrency

See **Disk Agent concurrency**.

control file (*Oracle and SAP R/3 specific term*)

An Oracle data file that contains entries specifying the physical structure of the database. It provides database consistency information used for recovery.

CRS

The Cell Request Server process (service) runs on the Data Protector Cell Manager. It starts and controls the backup and restore sessions. The service is started as soon as Data Protector is installed on the Cell Manager.

CRS runs under the account root on UNIX systems, and under any Windows account. By default, it runs under the account of the user, specified at installation time.

CSM

The Data Protector Copy Session Manager process controls the object copy session and runs on the Cell Manager system.

data file (*Oracle and SAP R/3 specific term*)

A physical file created by Oracle that contains data structures such as tables and indexes. A data file can only belong to one Oracle database.

data protection

Defines how long the backed up data on media remains protected, that is, Data Protector will not overwrite it. When the protection expires, Data Protector will be able to reuse the media in one of the next backup sessions.

See also **catalog protection**.

Data Protector Event Log

A central repository of all Data Protector related notifications. By default, all notifications are sent to the Event Log. The Event Log is accessible only to Data Protector users in the Admin group and to Data Protector users who are granted the Reporting and notifications user rights. You can view or delete all events in the Event Log.

Data Protector user account

You can use Data Protector only if you have a Data Protector user account, which restricts unauthorized access to Data Protector and to backed up data. Data Protector administrators create this account specifying a user logon name, the systems from which the user can log on, and a Data Protector user group

Glossary

membership. This is checked whenever the user starts the Data Protector user interface or performs specific tasks.

data stream

Sequence of data transferred over the communication channel.

database library

A Data Protector set of routines that enables data transfer between Data Protector and a server of an online database integration, for example, the Oracle Server.

database parallelism

More than one database is backed up at a time if the number of available devices allows you to perform backups in parallel.

database server

A computer with a large database stored on it, such as the SAP R/3 or Microsoft SQL database. A server has a database that can be accessed by clients.

Dboject (*Informix specific term*)

An Informix physical database object. It can be a blob space, db space, or logical-log file.

DC directory

The Detail Catalog (DC) directory consists of DC binary files, which store information about file versions. It represents the DCBF part of the IDB,

which occupies approximately 80% of the IDB. The default DC directory is called the dcbf directory and is located in the `<Data_Protector_home>\db40` directory on a Windows Cell Manager and in the `/var/opt/omni/server/db40` directory on a UNIX Cell Manager. You can create more DC directories and locate them as appropriate to you. Up to 10 DC directories are supported per cell. The default maximum size of a DC directory is 4 GB.

DCBF

The Detail Catalog Binary Files (DCBF) part of the IDB stores information about file versions and attributes. It occupies approximately 80% of the IDB. One DC binary file is created for each Data Protector medium used for backup.

delta backup

A delta backup is a backup containing all the changes made to the database from the last backup of any type.

See also **backup types**

device

A physical unit which contains either just a drive or a more complex unit such as a library.

device chain

A device chain consists of several standalone devices configured for sequential use. When a medium in one

Glossary

device gets full, the backup automatically continues on a medium in the next device in the device chain.

device group (*EMC Symmetrix specific term*)

A logical unit representing several EMC Symmetrix devices. A device cannot belong to more than a single device group. All devices in a device group must be on the same EMC Symmetrix unit. You can use a device group to identify and work with a subset of the available EMC Symmetrix devices.

device streaming

A device is streaming if it can feed enough data to the medium to keep it moving forward continuously. Otherwise, the tape has to be stopped, the device waits for more data, reverses the tape a little and resumes to write to the tape, and so on. In other words, if the data rate written to the tape is less or equal the data rate which can be delivered to the device by the computer system, then the device is streaming. Streaming significantly improves the performance of the device and use of space.

DHCP server

A system running the Dynamic Host Configuration Protocol (DHCP) providing dynamic configuration of IP addresses and related information.

differential backup

An incremental backup (incr) based on any previous Data Protector backup (full or any incremental), which must still be protected.

See **incremental backup**.

differential backup (*MS SQL specific term*)

A database backup that records only the data changes made to the database after the last full database backup.

See also **backup types**.

differential database backup

A differential database backup records only those data changes made to the database after the last full database backup.

direct backup

A SAN-based backup solution in which data movement directly from disk to tape (or to other secondary storage) is facilitated by the SCSI Extended Copy (Xcopy) command. Direct backup lessens the backup I/O load on systems in a SAN environment. The data movement is facilitated directly from disk to tape (or to other secondary storage) by the SCSI Extended Copy (XCOPY) command. The command is provided by any element of the infrastructure including bridges, switches, tape libraries, and disk subsystems.

See also **XCOPY engine**.

Glossary

directory junction (*Windows specific term*)

Directory junctions use the reparse point concept of Windows. An NTFS 5 directory junction allows you to redirect a directory/file request to another location.

disaster recovery

A process to restore a client's main system disk to a state close to the time when a (full) backup was performed.

Disk Agent

A component needed on a client to back it up and restore it. The Disk Agent controls reading from and writing to a disk. During a backup session, the Disk Agent reads data from a disk and sends it to the Media Agent, which then moves it to the device. During a restore session the Disk Agent receives data from the Media Agent and writes it to the disk.

Disk Agent concurrency

The number of Disk Agents that are allowed to send data to one Media Agent concurrently.

disk discovery

The detection of disks during client backup with disk discovery. During this backup, Data Protector discovers (detects) the disks that are present on the client — even though they might not have been present on the system when the backup was configured — and backs

them up. This is particularly useful in dynamic environments, where configurations change rapidly. After the disks are expanded, each inherits all options from its master client object. Even if pre- and post-exec commands are specified once, they are started many times, once per each object.

disk group (*Veritas Volume Manager specific term*)

The basic unit of data storage in VxVM system. A disk group can consist of one or more physical volumes. There can be more than one disk group on the system.

disk image (rawdisk) backup

A high-speed backup where Data Protector backs up files as bitmap images. A disk image (rawdisk) backup does not track the files and directory structure stored on the disk, but stores a disk image structure on byte level. You can perform a disk image backup of either specific disk sections or a complete disk.

disk quota

A concept to manage disk space consumption for all or a subset of users on a computer system. This concept is used by several operating system platforms.

disk staging

The process of backing up data in several phases to improve the

Glossary

performance of backups and restores, reduce costs of storing the backed up data, and increase the data availability and accessibility for restore. The backup stages consist of backing up data to one media type first (for example disk) and later copying it to a different media type (for example tape).

Distributed File System (DFS)

A service that connects file shares into a single namespace. The file shares can reside on the same or on different computers. DFS provides client access to the resources in a location-transparent manner.

DMZ

The Demilitarized Zone (DMZ) is a network inserted as a "neutral zone" between a company's private network (intranet) and the outside public network (Internet). It prevents outside users from getting direct access to company servers in the intranet.

DNS server

In the DNS client-server model, this is the server containing information about a portion of the DNS database that makes computer names available to client resolvers querying for name resolution across the Internet.

domain controller

A server in a network that is responsible for user security and verifying passwords within a group of other servers.

DR image

Data required for temporary disaster recovery operating system (DR OS) installation and configuration.

DR OS

A disaster recovery operating system is an operating system environment in which disaster recovery runs. It provides Data Protector a basic runtime environment (disk, network, tape, and filesystem access). The OS has to be installed and configured before the Data Protector disaster recovery can be performed. DR OS not only hosts the Data Protector disaster recovery process but is also a part of the restored system because it replaces its own configuration data with the original configuration data.

drive

A physical unit that receives data from a computer system and can write it onto a magnetic medium (typically a tape drive). It can also read the data from the medium and send it to the computer system.

Glossary

drive index

A number that identifies the mechanical position of a drive inside a library device. This number is used by the robotic control to access a drive.

dynamic client

See **client backup with disk discovery**.

EMC Symmetrix Agent (SYMA)

(EMC Symmetrix specific term)

See **Symmetrix Agent (SYMA)**

emergency boot file *(Informix specific term)*

An Informix configuration file that resides in the <INFORMIXDIR>\etc directory (on HP-UX) or <INFORMIXDIR>/etc directory (on Windows) and is called ixbar.<server_id>, where <INFORMIXDIR> is the OnLine Server home directory and <server_id> is the value of the SERVERNUM configuration parameter. Each line of the emergency boot file corresponds to one backup object.

Enterprise Backup Environment

Several cells can be grouped together and managed from a central cell. The enterprise backup environment includes all clients located in several Data Protector cells which are managed and

administered from a central cell using the Manager-of-Managers concept. See also **MoM**.

Event Logs

Files in which Windows logs all events, such as the starting or stopping of services and the logging on and off of users. Data Protector can back up Windows Event Logs as part of the Windows configuration backup.

exchanger

Also referred to as SCSI Exchanger. See also **library**.

exporting media

A process that removes all data about backup sessions, such as systems, objects, and file names, which reside on the media from the IDB. Information about the media and their relation to a pool is also removed from the IDB. The data on the media remains unchanged. See also **importing media**.

Extensible Storage Engine (ESE)

(Microsoft Exchange Server 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)

Glossary

or package (on Unix) from one cluster node to another. A failover can occur mostly because of software or hardware failures or maintenance on the primary node.

FC bridge

See **Fibre Channel bridge**

Fibre Channel

An ANSI standard for high-speed computer interconnection. Using either optical or copper cables, it allows the high speed bidirectional transmission of large data files and can be deployed between sites kilometers apart. Fibre Channel connects nodes using three physical topologies: point-to-point, loop, and switched.

Fibre Channel bridge

A Fibre Channel bridge or multiplexer provides the ability to migrate existing parallel SCSI devices, like RAID arrays, solid state disks (SSD), and tape libraries to a Fibre Channel environment. On one side of the bridge or multiplexer there is a Fibre Channel interface while on the other side there are parallel SCSI ports. The bridge or multiplexer enables SCSI packets to be moved between the Fibre Channel and parallel SCSI devices.

file depot

A file containing the data from a backup to a file library device.

file jukebox device

A device residing on disk consisting of multiple slots used to store file media.

file library device

A device which resides on a disk emulating a library with several media, hence containing multiple files, referred to as file depots.

File Replication Service (FRS)

A Windows service that replicates the domain controller store logon scripts and group policies. FRS also enables replication of Distributed File System (DFS) shares between systems and allows any server to perform replication activity.

file version

The same file can be backed up multiple times in case of full backups and incremental backups (if the file changed). If the log level ALL is selected for backup, Data Protector retains one entry in the IDB for the filename itself and one for each version (date/time) of the file.

filesystem

The organization of files on a hard disk. A filesystem is backed up so that the file attributes and the file contents are stored on the backup media.

Glossary

first level mirror (*HP StorageWorks Disk Array XP specific term*)

HP StorageWorks Disk Array XP allows up to three mirror copies of a Primary Volume and each of these copies can have additional two copies. The three mirror copies are called first level mirrors.

See also **Primary Volume**, and **MU numbers**.

fnames.dat

The fnames.dat files of the IDB contain information on the names of the backed up files. Typically, these files occupy about 20% of the IDB, if filenames are stored.

formatting

A process that erases any data contained on a medium and prepares it for use with Data Protector. Information about media (media ID, description, and location) is saved in the IDB as well as on the respective media (media header). Data Protector media with protected data are not formatted until the protection expires or the media are unprotected/recycled.

free pool

An auxiliary source of media for use by media pools when they run out of media. The media pools must be configured to use free pools.

full backup

A backup in which all selected objects are backed up, whether or not they have been recently modified.

See also **backup types**.

full database backup

A backup of all data in a database, not only the data that has been changed after the last (full or incremental) backup of the database. A full database backup does not depend on any other backup.

full mailbox backup

A full mailbox backup is a backup of the entire mailbox content.

full ZDB

A ZDB backup in which all selected objects are backed up, even if there are no changes from the previous backup.

See also **incremental ZDB**.

global options file

A file that allows you to customize Data Protector. It explains the global options, which cover various aspects of Data Protector, typically time-outs and limits, and affect the entire Data Protector cell. The file is located in the /etc/opt/omni/server/options directory on HP-UX and Solaris systems and in the <Data_Protector_home>\Config\Server\Options directory on Windows systems.

Glossary

group (*Microsoft Cluster Server specific term*)

A collection of resources (for example disk volumes, application services, IP names and addresses) that are needed to run a specific cluster-aware applications.

GUI

A cross-platform (HP-UX, Solaris, and Windows) graphical user interface, provided by Data Protector for easy access to all configuration, administration, and operation tasks.

hard recovery (*Microsoft Exchange Server specific term*)

A Microsoft Exchange Server database recovery that is performed after a restore by the database engine, using transaction log files.

heartbeat

A cluster data set with a time stamp carrying information about the operational status of a particular cluster node. This data set or packet is distributed among all cluster nodes.

Hierarchical Storage Management (HSM)

A method for optimizing the use of expensive hard disk storage by migrating less frequently used data to less expensive optical platters. When needed, the data is migrated back to

hard disk storage. This balances the need for fast retrieval from hard disk with the lower cost of optical platters.

Holidays file

A file that contains information about holidays. You can set different holidays by editing the Holidays file: `/etc/opt/omni/server/Holidays` on the UNIX Cell Manager and `<Data_Protector_home>\Config\Server\holidays` on the Windows Cell Manager.

host backup

See **client backup with disk discovery**.

hosting system

A working Data Protector client used for Disk Delivery Disaster Recovery with a Data Protector Disk Agent installed.

HP ITO

See **OVO**.

HP OpC

See **OVO**.

HP OpenView SMART Plug-In (SPI)

A fully integrated, out-of-the-box solution which "plugs into" HP OpenView Operations, extending the managed domain. Through the Data Protector integration, which is implemented as an HP OpenView SMART Plug-In, a user can have an

Glossary

arbitrary number of Data Protector Cell Managers monitored as an extension to HP OpenView Operations (OVO).

HP OVO

See **OVO**.

HP StorageWorks Disk Array XP LDEV

A logical partition of a physical disk within an HP StorageWorks Disk Array XP. LDEVs are entities that can be replicated in the Continuous Access XP (CA) and Business Copy XP (BC) configurations, or can be used as standalone entities.

See also **BC** (*HP StorageWorks Disk Array XP specific term*), **CA** (*HP StorageWorks Disk Array XP specific term*), and **replica**.

HP StorageWorks EVA Agent (legacy)

A Data Protector software module that executes all tasks required for the HP StorageWorks Enterprise Virtual Array integration operating on HP StorageWorks EVA with Command View (CV) EVA software v3.1 or lower, and the EVA VCS firmware v3.01x or lower.

See also **Command View (CV) EVA** and **HP StorageWorks EVA SMI-S Agent**.

HP StorageWorks EVA SMI-S Agent

A Data Protector software module that executes all tasks required for the HP StorageWorks Enterprise Virtual Array integration operating on HP StorageWorks EVA with Command View (CV) EVA software starting with v3.2. With the EVA SMI-S Agent, the control over the array is established through HP StorageWorks SMI-S EVA provider, which directs communication between incoming requests and CV EVA.

See also **Command View (CV) EVA**, **HP StorageWorks SMI-S EVA provider**, and **HP StorageWorks EVA Agent (legacy)**.

HP StorageWorks SMI-S EVA provider

An interface used for controlling HP StorageWorks Enterprise Virtual Array. SMI-S EVA provider runs as a separate service on the HP OpenView Storage Management Appliance system and acts as a gateway between incoming requests and Command View EVA. With the Data Protector HP StorageWorks EVA integration, SMI-S EVA provider accepts standardized requests from the EVA SMI-S Agent, communicates with Command View EVA for information or method invocation, and returns standardized responses.

See also **HP StorageWorks EVA SMI-**

Glossary

S Agent and Command View (CV) EVA.

HP StorageWorks Virtual Array LUN

A logical partition of a physical disk within an HP StorageWorks Virtual Array. LUNs are entities that can be replicated in the HP StorageWorks Business Copy VA configuration, or can be used as standalone entities.
See also **BC VA** and **replica**.

HP VPO
See **OVO**.

ICDA (*EMC Symmetrix specific term*)
EMC's Symmetrix Integrated Cached Disk Arrays (ICDA) is a disk array device that combines a set of physical disks, a number of FWD SCSI channels, an internal cache memory, and control and diagnostic software commonly referred to as the microcode.

IDB

The Data Protector Internal Database is an embedded database located on the Cell Manager that keeps information regarding which data is backed up, on which media it is backed up, how backup and restore sessions are run, and which devices and libraries are configured.

importing media

A process that re-reads all data about backup sessions which are on the medium back into the IDB. This then allows for fast and convenient access to the data on the media.
See also **exporting media**.

incremental backup

A backup that selects only files that have changed since a previous backup. Several levels of incremental backup are available, allowing selective backup of only files that have changed since the last incremental backup.
See also **backup types**.

incremental backup (*Microsoft Exchange Server specific term*)

A backup of the Microsoft Exchange Server data that has changed since the last full or incremental backup. With the incremental backup, only the transaction log files are backed up.
See also **backup types**.

incremental mailbox backup

An incremental mailbox backup backs up all the changes made to the mailbox after the last backup of any type.

incremental1 mailbox backup

An incremental1 mailbox backup backs up all the changes made to the mailbox after the last full backup.

Glossary

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 ZDB to tape or ZDB to disk+tape session in which only changes from the last full or incremental protected 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

Glossary

that are shared among several users.
See also **Key Management Service** and **Site Replication Service**.

initializing

See **formatting**.

Installation Server

A computer system that holds a repository of the Data Protector software packages for a specific architecture. The Installation Server is used for remote installation of Data Protector clients. In mixed environments at least two Installation Servers are needed: one for UNIX systems and one for Windows systems.

instant recovery (*ZDB specific term*)

A process in which a replica, produced by a ZDB-to-disk or a ZDB-to-disk+tape sessions, is used to restore the contents of the source volumes to their states at the time at which the replica was created, avoiding the need to perform a restore from tape. Depending on the application/database concerned, this may be all that is required, or other steps, such as the application of transaction log files, may be required for full recovery.

See also **replica, zero downtime backup (ZDB), ZDB to disk, and ZDB to disk+tape**.

integrated security (*MS SQL specific term*)

Integrated security allows the Microsoft SQL Server to use Windows authentication mechanisms to validate Microsoft SQL Server logins for all connections. Using integrated security means that users have one password for both Windows and Microsoft SQL Server. Integrated security should be used in environments where all clients support trusted connections. Connections validated by Windows Server and accepted by Microsoft SQL Server are referred to as trusted connections. Only trusted connections are allowed.

integration object

A backup object of a Data Protector integration, such as Oracle or SAP DB.

Internet Information Server (IIS)

(Windows specific term)

Microsoft Internet Information Server is a network file and application server that supports multiple protocols. Primarily, IIS transmits information in Hypertext Markup Language (HTML) pages by using the Hypertext Transport Protocol (HTTP).

IP address

Internet Protocol address is a numeric address of a system used to uniquely identify the system on the network. The

Glossary

IP address consists of four groups of numbers separated by periods (full stops).

ISQL (*Sybase specific term*)

A Sybase utility used to perform system administration tasks on Sybase SQL Server.

ITO

See **OVO**.

jukebox

See **library**.

jukebox device

A device consisting of multiple slots used to store either optical or file media. When being used to store file media, the jukebox device is known as the “file jukebox device”.

Key Management Service (*Microsoft Exchange Server specific term*)

The Microsoft Exchange Server service that provides encryption functionality for enhanced security.

See also **Information Store** and **Site Replication Service**.

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

Glossary

used. If you do not want to use load balancing, you can select which device will be used for each object in the backup specification. Data Protector will access the devices in the specified order.

local and remote recovery

Remote recovery is performed if all Media Agent hosts specified in the SRD file are accessible. If any of them fails, the disaster recovery process fails over to the local mode. This means that the target system is searched for locally attached devices. If only one device is found, it is automatically used. Otherwise, Data Protector prompts you to select the device, which will be used for restore.

lock name

You can configure the same physical device several times with different characteristics, by using different device names.

The lock name is a user specified string that is used for locking all such device configurations to prevent collision if several such devices (device names) are used concurrently. Use an identical lock name for all device definitions which use the same physical device.

log_full shell script (*Informix UNIX specific term*)

A script provided by ON-Bar that you

can use to start backing up logical-log files when OnLine Server issues a log-full event alarm. The Informix ALARMPROGRAM configuration parameter defaults to the `<INFORMIXDIR>/etc/log_full.sh`, where `<INFORMIXDIR>` is the OnLine Server home directory. If you do not want logical logs to be backed up continuously, set the ALARMPROGRAM configuration parameter to `<INFORMIXDIR>/etc/no_log.sh`.

logging level

The logging level determines the amount of details on files and directories written to the IDB during backup or object copying. You can always restore your data, regardless of the logging level used during backup. Data Protector provides four logging levels: Log All, Log Directories, Log Files, and No Log. The different logging level settings influence the IDB growth, backup speed, and the convenience of browsing data for restore.

logical-log files

This applies to online database backup. Logical-log files are files in which modified data is first stored before being flushed to disk. In the event of a failure, these logical-log files are used to roll forward all transactions that have been

Glossary

committed as well as roll back any transactions that have not been committed.

login ID (*MS SQL Server specific term*)

The name a user uses to log on to Microsoft SQL Server. A login ID is valid if Microsoft SQL Server has an entry for that user in the system table syslogin.

login information to the Oracle Target Database (*Oracle and SAP R/3 specific term*)

The format of the login information is <user_name>/<password>@<service>, where:

- <user_name> is the name by which a user is known to Oracle Server and to other users. Every user name is associated with a password and both have to be entered to connect to an Oracle Target Database. This user must have been granted Oracle SYSDBA or SYSOPER rights.
- <password> is a string used for data security and known only to its owner. Passwords are entered to connect to an operating system or software application. The password has to be the same as the password specified in the Oracle password file (orapwd), which is used for authentication of users performing database administration.

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

login information to the Recovery Catalog Database (*Oracle specific term*)

The format of the login information to the Recovery (Oracle) Catalog Database is <user_name>/<password>@<service>, where the description of the user name, password, and service name is the same as in the Oracle SQL*Net V2 login information to the Oracle target database. In this case, <service> is the name of the service to the Recovery Catalog Database, not the Oracle target database.

Note that the Oracle user specified here has to be the owner of the Oracle Recovery (Oracle) Catalog.

Lotus C API (*Lotus Domino Server specific term*)

An interface for the exchange of backup and recovery information between Lotus Domino Server and a backup solution, like Data Protector.

LVM

A Logical Volume Manager is a subsystem for structuring and mapping physical disk space to logical volumes on UNIX systems. An LVM system

Glossary

consists of several volume groups, where each volume group has several volumes.

Magic Packet

See **Wake ONLAN**.

mailbox (*Microsoft Exchange Server specific term*)

The location to which e-mail is delivered, which is set up by the administrator for each user. If a set of personal folders is designated as the e-mail delivery location, e-mail is routed from the mailbox to this location.

Mailbox Store (*Microsoft Exchange Server specific term*)

A part of the Information Store that maintains information about user mailboxes. A mailbox store consists of a binary rich-text .edb file and a streaming native internet content .stm file.

Main Control Unit (MCU) (*HP StorageWorks Disk Array XP specific term*)

An HP StorageWorks XP disk array that contains the primary volumes for the Continuous Access configuration and acts as a master device.

See also **BC** (*HP StorageWorks Disk Array XP specific term*), **CA** (*HP StorageWorks Disk Array XP specific term*), and **HP StorageWorks Disk Array XP LDEV**.

Manager-of-Managers (MoM)

See **Enterprise Cell Manager**.

Media Agent

A process that controls reading from and writing to a device, which reads from or writes to a medium (typically a tape).

During a backup session, a Media Agent receives data from the Disk Agent and sends it to the device for writing it to the medium. During a restore session, a Media Agent locates data on the backup medium and sends it to the Disk Agent. The Disk Agent then writes the data to the disk. A Media Agent also manages the robotics control of a library.

MAPI (*Microsoft Exchange specific term*)

The MAPI (Messaging Application Programming Interface) is the programming interface that lets applications and messaging clients interact with messaging and information systems.

media allocation policy

Determines in which sequence media are used for backup. The Strict allocation policy directs Data Protector to prompt for a specific medium. The Loose policy directs Data Protector to prompt for any suitable medium. The Formatted First policy directs Data Protector to give preference to unknown media, even if unprotected media are available in the library.

Glossary

media condition

The quality of a medium as derived from the media condition factors. Heavy usage and age result in an increased number of read and write errors with tape media. Media need to be replaced when they are marked as POOR.

media condition factors

The user-assigned age threshold and overwrite threshold used to determine the state of a medium.

media ID

A unique identifier assigned to a medium by Data Protector.

media label

A user-defined identifier used to describe a medium.

media location

A user-defined physical location of a medium, such as "building 4" or "off-site storage".

media management session

A session performing some action on a medium, such as initializing, scanning the content, verifying data on a medium, or copying a medium.

media pool

A set of media of the same type (such as DDS) used and tracked as a group. Media are formatted and assigned to a media pool.

media set

The result of a backup session is data backed up on a group of media called media set. Depending on the media usage policy, several sessions can share the same media.

media type

The physical type of media, such as DDS or DLT.

media usage policy

The media usage policy controls how new backups are added to the already used media. It can be Appendable, Non-Appendable, or Appendable for incrementals only.

merging

This defines one mode to resolve file conflicts during restore. If the file to be restored already exists at the destination, the one with the more recent modification date is kept. Files not present on the disk are always restored. *See also* **overwrite**.

MFS

The Migrating File System enables a standard JFS filesystem with migration capabilities (on HP-UX 11.00). The MFS is accessed via a standard filesystem interface (DMAPI), it is mounted to a directory the same way as any HP-UX filesystem. In an MFS, only the superblock, the inode and the 'extended attribute' information remain

Glossary

permanently on the hard disk and are never migrated.

See also **VBFS**.

Microsoft Exchange Server

A “client-server” messaging and a workgroup system that offers a transparent connection to many different communication systems. It provides users with an electronic mail system, individual and group scheduling, online forms, and workflow automation tools. It provides a developer with a platform on which to build custom information-sharing and messaging-service applications.

Microsoft Management Console (MMC) *(Windows specific term)*

An administration model for Windows-based environments. It provides a simple, consistent, and integrated administration user interface allowing management of many applications through the same GUI, provided that the applications adhere to the MMC model.

Microsoft SQL Server

A database management system designed to meet the requirements of distributed “client-server” computing.

Microsoft Volume Shadow Copy service (VSS)

A software service that provides a unified communication interface to coordinate backup and restore of a VSS-

aware application regardless of its specific features. This service collaborates with the backup application, writers, shadow copy providers, and the operating system kernel to implement the management of volume shadow copies and shadow copy sets.

See also **shadow copy, shadow copy provider, writer**.

mirror *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)*

See **target volume**.

mirror rotation *(HP StorageWorks Disk Array XP specific term)*

See **replica set rotation**.

MMD

The Media Management Daemon process (service) runs on the Data Protector Cell Manager and controls media management and device operations. The process is started when Data Protector is installed on the Cell Manager.

MMDB

The Media Management Database (MMDB) is a part of the IDB that contains information about media, media pools, devices, libraries, library drives, and slots configured in the cell, as well as the Data Protector media used for backup. In an enterprise backup

Glossary

environment, this part of the database can be common to all cells.
See also **CMMDB, CDB.**

MoM

Several cells can be grouped together and managed from a central cell. The management system of the central cell is the Manager-of-Managers (MoM). The MoM allows you to configure and manage multiple cells from a central point.

mount request

A screen prompt that tells you to insert a specific medium into a device. Once you respond to the mount request by providing the required medium and confirm the mount request, the session continues.

mount point

The access point in a directory structure for a disk or logical volume, for example /opt or d:. On UNIX, the mount points are displayed using the bdf or df command.

MSM

The Data Protector Media Session Manager, which runs on the Cell Manager and controls media sessions, such as copying media.

MU number (*HP StorageWorks Disk Array XP specific term*)

A Mirror Unit number is an integer

number (0, 1 or 2), used to indicate a first level mirror.
See also **first level mirror.**

multi-drive server

A license that allows you to run an unlimited number of Media Agents on a single system. This license, which is bound to the IP address of the Cell Manager, is no longer available.

obdrindex.dat

An IDB file with information about IDB backups, media, and devices used for the backup. This information can significantly simplify IDB recovery. It is recommended to relocate the file, together with IDB transaction logs, on a separate physical disk from other IDB directories, and, additionally, to make a copy of the file and locate it where you want.

OBDR capable device

A device that can emulate a CD-ROM drive loaded with a bootable disk and can thus be used as a backup or boot device for disaster recovery purposes.

object

See **backup object**

object copy

A copy of a specific object version that is created during an object copy session or a backup session with object mirroring.

Glossary

object copy session

A process that creates an additional copy of the backed up data on a different media set. During an object copy session, the selected backed up objects are copied from the source to the target media.

object copying

The process of copying selected object versions to a specific media set. You can select object versions from one or several backup sessions to be copied.

Object ID (*Windows specific term*)

The object IDs (OIDs) enable access to NTFS 5 files no matter where in the system the files reside. Data Protector treats the OIDs as alternate streams of the files.

object mirror

A copy of a backup object created using object mirroring. Object mirrors are often referred to as object copies.

object mirroring

The process of writing the same data to several media sets during a backup session. Data Protector enables you to mirror all or some backup objects to one or more media sets.

offline backup

A backup during which an application database cannot be used by the application.

- For simple backup methods (non ZDB), the database is generally put into a quiescent state that allows use by the backup system, but not the application, for the whole backup period (~minutes/hours). For instance, for backup to tape, until streaming of data to the tape is finished.
- For ZDB methods, the database is also put into the quiescent state, but for the period of the data replication process only (~seconds). Normal database operation can then be resumed for the rest of the backup process.

*See also **zero downtime backup (ZDB)** and **online backup**.*

offline recovery

Offline recovery is performed if the Cell Manager is not accessible, for example, due to network problems. Only standalone and SCSI library devices can be used for offline recovery. Recovery of the Cell Manager is always offline.

offline redo log

*See **archived redo log***

OmniStorage

Software providing transparent migration of less frequently used data to the optical library while keeping more

Glossary

frequently used data on the hard disk. HP OmniStorage runs on HP-UX systems.

On-Bar (*Informix specific term*)

A backup and restore system for OnLine Server. ON-Bar enables you to create a copy of your OnLine Server data and later restore the data. The ON-Bar backup and restore system involves the following components:

- onbar utility
- Data Protector, as the backup solution
- XBSA interface
- ON-Bar catalog tables, which are used to back up dbobjects and track instances of dbobjects through multiple backups.

onbar utility (*Informix specific term*)

The Informix utility that communicates backup and restore requests to OnLine Server. The utility uses XBSA to exchange control data and back up and restore data with Data Protector.

ONCONFIG (*Informix specific term*)

An environment variable that specifies the name of the active ONCONFIG configuration file. If the ONCONFIG environment variable is not present, OnLine uses the configuration values

from the file `<INFORMIXDIR>/etc/onconfig` (on HP-UX) or `<INFORMIXDIR>\etc\onconfig` (on Windows).

online backup

A backup performed while a database application remains available for use. The database is placed into a special backup mode of operation for the time period that the backup application requires access to the original data objects. During this period, the database is fully operational, but there may be a small performance impact and log files may grow very quickly.

- For simple backup methods (non ZDB), backup mode is required for the whole backup period (~minutes/ hours). For instance, for backup to tape, until streaming of data to tape is finished.
- For ZDB methods, backup mode is required for the short period of the data replication process only (~seconds). Normal database operation can then be resumed for the rest of the backup process.

In some cases, transaction logs may also have to be backed up to allow a consistent database to be restored. *See also* **zero downtime backup (ZDB)** and **offline backup**.

Glossary

online redo log (*Oracle specific term*)

Redo logs that have not been archived, but are either available to the instance for recording database activity or are filled and waiting to be archived or reused.

See also **archived redo log**.

OnLine Server (*Informix specific term*)

Refers to INFORMIX-OnLine Dynamic Server.

OpC

See **OVO**.

Oracle instance (*Oracle specific term*)

Each installation of an Oracle database on one or more systems. One computer system can have several instances of a database running.

ORACLE_SID (*Oracle specific term*)

A unique name for an Oracle Server instance. To switch among Oracle Servers, specify the desired `<ORACLE_SID>`. The `<ORACLE_SID>` is included in the `CONNECT DATA` parts of the connect descriptor in a `TNSNAMES.ORA` file and in the definition of the TNS listener in the `LISTENER.ORA` file.

original system

The system configuration backed up by Data Protector before a computer disaster hits the system.

overwrite

An option that defines one mode to resolve file conflicts during restore. All files are restored from a backup even if they are older than existing files.

See also **merging**.

OVO

HP OpenView Operations for Unix provides powerful capabilities for operations management of a large number of systems and applications on in a network. Data Protector provides an integration into this management product. This integration is implemented as a SMART Plug-In for OVO management servers on HP-UX and Solaris. Earlier versions of OVO were called IT/Operation, Operations Center and Vantage Point Operations.

See also **merging**.

ownership

The ownership of a backup determines who can restore from the backup. The user who starts an interactive backup is the session owner. If a user starts an existing backup specification without modifying it, the session is not considered interactive. In that case, if the backup owner has been defined in the backup specification, they remain the session owner. Otherwise, the session owner becomes the user who started the backup in question. For the scheduled backups, by default, the session owner is for the UNIX Cell

Glossary

Manager: root.sys@<Cell Manager>, and for the Windows Cell Manager, the user that was specified during the installation of the Cell Manager. It is possible to modify the ownership, so that the specific user becomes the session owner.

P1S file

P1S file contains information on how to format and partition all disks installed in the system during Enhanced Automated Disaster Recovery (EADR). It is created during full backup and is saved on backup medium and on Cell Manager into

<*Data_Protector_home*>\Config\Server\dr\p1s directory on a Windows Cell Manager or in /etc/opt/omni/server/dr/p1s directory on a UNIX Cell Manager with the filename recovery.p1s.

package (*MC/ServiceGuard and Veritas Cluster specific term*)

A collection of resources (for example volume groups, application services, IP names and addresses) that are needed to run a specific cluster-aware application.

pair status (*HP StorageWorks Disk Array XP specific term*)

A mirrored pair of disks can have various status values depending on the action performed on it. The three most important status values are:

- **COPY** - The mirrored pair is currently resynchronizing. Data is transferred from one disk to the other. The disks do not contain the same data.
- **PAIR** - The mirrored pair is completely synchronized and both disks (the primary volume and the mirrored volume) contain identical data.
- **SUSPENDED** - The link between the mirrored disks is suspended. That means that both disks are accessed and updated independently. However, the mirror relationship is still maintained and the pair can be resynchronized without transferring the complete disk.

parallel restore

Restoring backed up data to multiple disks at the same time (that is, in parallel) by running multiple Disk Agents, that receive data from one Media Agent. For the parallel restore to work, select data that is located on different disks or logical volumes and during backup, the data from the different objects must have been sent to the same device using a concurrency of 2 or more. During a parallel restore, the data for multiple objects selected for restore is read from media at the same time, thereby improving performance.

Glossary

parallelism

The concept of reading multiple data streams from an online database.

physical device

A physical unit that contains either a drive or a more complex unit such as a library.

post-exec

A backup option that executes a command or script after the backup of an object or after the entire session completes. Post-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows and as shell scripts on UNIX.

See also **pre-exec**.

pre- and post-exec commands

Pre- and post-exec commands are used to perform additional action before and after a backup or restore session. They are not supplied by Data Protector. You need to create your own commands. They can be written as executables or batch files on Windows and as shell scripts on UNIX.

prealloc list

A subset of media in a media pool that specifies the order in which media are used for backup.

pre-exec

A backup option that executes a command or script before the backup of an object or before the entire session is started. Pre-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows and as shell scripts on UNIX.

See also **post-exec**.

Primary Volume (P-VOL) (*HP*

StorageWorks Disk Array XP specific term)

Standard HP StorageWorks Disk Array XP LDEVs that act as a primary volume for the CA and BC configurations. The P-VOL is located in the MCU.

See also **Secondary Volume (S-VOL)**.

protection

See **data protection** and also **catalog protection**.

public folder store (*Microsoft Exchange Server specific term*)

The part of the Information Store that maintains information in public folders. A public folder store consists of a binary rich-text .edb file and a streaming native internet content .stm file.

public/private backed up data

When configuring a backup, you can select whether the backed up data will be:

Glossary

- public, that is visible (and accessible for restore) to all Data Protector users
- private, that is, visible (and accessible for restore) only to the owner of the backup and administrators

RAID

Redundant Array of Inexpensive Disks.

RAID Manager Library (*HP StorageWorks Disk Array XP specific term*)

The RAID Manager Library is used internally by Data Protector on Solaris systems to allow access to HP StorageWorks Disk Array XP configuration, status, and performance data and to key HP StorageWorks Disk Array XP features through the use of function calls translated into a sequence of low level SCSI commands.

RAID Manager XP (*HP StorageWorks Disk Array XP specific term*)

The RAID Manager XP application provides an extensive list of commands to report and control the status of the CA and BC applications. The commands communicate through a RAID Manager instance with the HP StorageWorks Disk Array XP Disk Control Unit. This instance translates the commands into a sequence of low level SCSI commands.

rawdisk backup

See disk image backup.

RCU (*HP StorageWorks specific term*)

The Remote Control Unit acts as a slave of an MCU in a CA configuration. In bidirectional configurations, the RCU can act as an MCU.

RDBMS

Relational Database Management System.

RDF1/RDF2 (*EMC Symmetrix specific term*)

A type of SRDF device group. Only RDF devices can be assigned to an RDF group. An RDF1 group type contains source (R1) devices and an RDF2 group type contains target (R2) devices.

RDS

The Raima Database Server process (service) runs on the Data Protector Cell Manager and manages the IDB. The process is started when Data Protector is installed on the Cell Manager.

Recovery Catalog (*Oracle specific term*)

A set of Oracle tables and views that are used by Recovery Manager to store information about Oracle databases. This information is used by Recovery Manager to manage the backup, restore,

Glossary

and recovery of Oracle databases. The recovery catalog contains information about:

- The physical schema of the Oracle target database
- Data file and archived log backup sets
- Data file copies
- Archived Redo Logs
- Stored scripts.

Recovery Catalog Database (*Oracle specific term*)

An Oracle database that contains a recovery catalog schema. You should not store the recovery catalog in your target database.

RecoveryInfo

When backing up Windows configuration files, Data Protector collects the information about the current system configuration (information on disk layout, volume, and network configuration). This information is needed for disaster recovery.

Recovery Manager (RMAN) (*Oracle specific term*)

An Oracle command-line interface that directs an Oracle Server process to back

up, restore, or recover the database it is connected to. RMAN uses either the recovery catalog or the control file to store information about backups. This information can be used later in restore sessions.

recycle

A process that removes the data protection from all backed up data on a medium, allowing Data Protector to overwrite it during one of the next backups. Data that belongs to the same session(s) but resides on other media is also unprotected. Recycling does not actually alter the data on the medium.

redo log (*Oracle specific term*)

Every Oracle database has a set of two or more redo log files. The set of redo log files for the database is known as the database's redo log. Oracle uses the redo log to record all changes made to data.

Remote Control Unit (RCU)

(HP StorageWorks Disk Array XP specific term)

The Remote Control Unit (RCU) acts as a slave of an MCU in a CA configuration. In bidirectional configurations, the RCU can act as an MCU.

Removable Storage Management Database (*Windows specific term*)

A Windows service used for managing removable media (such as tapes and

Glossary

disks) and storage devices (libraries). Removable Storage allows applications to access and share the same media resources.

reparse point (*Windows specific term*)

A system-controlled attribute that can be associated with any directory or file. The value of a reparse attribute can have user-controlled data. The format of the data is understood by the application that stored the data and a filesystem filter that was installed to interpret the data and process such files. Whenever the filesystem encounters a file with a reparse point, it attempts to find the filesystem filter associated with the data format.

replica (*ZDB specific term*)

An image, at a particular point in time, of the data in source volumes that contain user-specified backup objects. Depending on the hardware/software with which it is created, the image may be an independent exact duplicate (clone) of the storage blocks at the physical disk level (for example, a split mirror), or a virtual copy (for example, a snapshot). From a host's perspective, on a basic UNIX or Windows system, the complete physical disk containing a backup object is replicated. However, if a volume manager is used on UNIX, the whole volume/disk group containing a

backup object is replicated.

See also **snapshot**, **snapshot creation**, **split mirror**, and **split mirror creation**.

replica set (*ZDB specific term*)

A group of replicas, all created using the same backup specification.

See also **replica** and **replica set rotation**.

replica set rotation (*ZDB specific term*)

The use of a replica set for regular backup production: Each time the same backup specification requiring the use of a replica set is run, a new replica is created and added to the set, until the maximum number for the set is reached. After that, the oldest replica in the set is replaced, maintaining the maximum number of replicas in the set.

See also **replica** and **replica set**.

restore session

A process that copies data from backup media to a client.

RMAN (*Oracle specific term*)

See **Recovery Manager**.

RSM

The Data Protector Restore Session Manager controls the restore session. This process always runs on the Cell Manager system.

Glossary

RSM (*Windows specific term*)

Removable Storage Manager (RSM) includes a media management service that facilitates communication among applications, robotic changers, and media libraries. It enables multiple applications to share local robotic media libraries and tape or disk drives and to manage removable media.

SAPDBA (*SAP R/3 specific term*)

An SAP R/3 user interface that integrates the BRBACKUP, BRARCHIVE, and BRRESTORE tools.

scan

A function that identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library).

scanning

A function which identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library). It is useful to perform a scan and check the actual media in the device if someone has manually manipulated media without using Data Protector to eject or enter, for example.

Scheduler

A function that controls when and how often automatic backups occur. By setting up a schedule, you automate the start of backups.

Secondary Volume (S-VOL) (*HP StorageWorks Disk Array XP specific term*)

Secondary Volumes, or S-VOLs, are XP LDEVs that act as a secondary CA or BC mirror of another LDEV (a P-VOL). In the case of CA, S-VOLs can be used as failover devices in a MetroCluster configuration. The S-VOLs are assigned separate SCSI addresses, different from the addresses used by the P-VOLs. *See also* **Primary Volume (P-VOL)**.

session

See **backup session**, **media management session**, and **restore session**.

session ID

An identifier of a backup, restore, object copy, or media management session, consisting of the date when the session ran and a unique number.

session key

This environment variable for the Pre- and Post-exec script is a Data Protector unique identification of any session, including preview sessions. The session key is not recorded in the database, and

Glossary

it is used for specifying options for the omnimnt, omnistat and omniabort CLI commands.

shadow copy (*MS VSS specific term*)

A volume that represents a duplicate of the original volume at a certain point in time. The data is then backed up from the shadow copy and not from the original volume. The original volume continues to change as the backup process continues, but the shadow copy of the volume remains constant. |

See also **Microsoft Volume Shadow Copy service**.

shadow copy provider (*MS VSS specific term*)

An entity that performs the work on creating and representing the volume shadow copies. Providers own the shadow copy data and expose the shadow copies. Providers can be software (for example, system providers) or hardware (local disks, disk arrays).

See also **shadow copy**.

shadow copy set (*MS VSS specific term*)

A collection of shadow copies created at the same point in time.

See also **shadow copy**.

shared disks

A Windows disk on another system that has been made available to other users

on the network. Systems with shared disks can be backed up without a Data Protector Disk Agent installed.

SIBF

The Serverless Integrations Binary Files (SIBF) is a part of the IDB that stores raw NDMP meta data. This data is necessary to perform restore of NDMP objects.

Site Replication Service (*Microsoft Exchange Server specific term*)

The Microsoft Exchange Server 2000/2003 service that permits compatibility with Microsoft Exchange Server 5.5 by emulating the Exchange Server 5.5 directory service.

See also **Information Store** and **Key Management Service**.

slot

A mechanical position in a library. Each slot can hold a medium, such as a DLT tape. Data Protector references each slot with a number. To read a medium, a robotic mechanism moves the medium from a slot into the drive.

SMB

See **split mirror backup**.

SMBF

The Session Messages Binary Files (SMBF) part of the IDB stores session messages generated during backup, object copy, restore, and media

Glossary

management sessions. One binary file is created per session. The files are grouped by year and month.

snapshot (*HP StorageWorks VA and HP StorageWorks EVA specific term*)

A form of replica produced using snapshot creation techniques. A range of snapshot types is available, with different characteristics, depending on the arrays/techniques used. Such replicas are dynamic and may be either virtual copies, still reliant upon the contents of the source volumes, or independent exact duplicates (clones), depending on the snapshot type and the time since creation.

See also **replica** and **snapshot creation**.

snapshot backup (*HP StorageWorks VA and HP StorageWorks EVA specific term*)

See **ZDB to tape**, **ZDB to disk**, and **ZDB to disk+tape**.

snapshot creation (*HP StorageWorks VA and HP StorageWorks EVA specific term*)

A replica creation technique, in which copies of source volumes are created using storage virtualization techniques. The replicas are considered to be created at one particular point-in-time, without pre-configuration, and are immediately available for use. However background

copying processes normally continue after creation.

See also **snapshot**.

source (R1) device (*EMC Symmetrix specific term*)

An EMC Symmetrix device that participates in SRDF operations with a target (R2) device. All writes to this device are mirrored to a target (R2) device in a remote EMC Symmetrix unit. An R1 device must be assigned to an RDF1 group type.

See also **target (R2) device**.

source volume (*ZDB specific term*)

A storage volume containing data to be replicated.

sparse file A file that contains data with portions of empty blocks. Examples are: -A matrix in which some or much of the data contains zeros -files from image applications -high-speed databases If sparse file processing is not enabled during restore, it might be impossible to restore this file.

split mirror (*EMC Symmetrix and HP StorageWorks Disk Array XP specific term*)

A replica created using split mirror techniques. Such a replica provides an independent, exact duplicate, or clone, of the contents of the source volumes.

See also **replica** and **split mirror creation**.

Glossary

split mirror backup (*EMC Symmetrix specific term*)

See **ZDB to tape**.

split mirror backup (*HP StorageWorks Disk Array XP specific term*)

See **ZDB to tape**, **ZDB to disk**, and **ZDB to disk+tape**.

split mirror creation (*EMC Symmetrix and HP StorageWorks Disk Array XP specific term*)

A replica creation technique, in which a pre-configured set of target volumes (a mirror) is kept synchronized with a set of source volumes until the time at which a replica of the contents of the source volumes is required. Then, the synchronization is stopped (the mirror is split) and a split mirror replica of the source volumes at the time of the split remains in the target volumes.

See also **split mirror**.

split mirror restore (*EMC Symmetrix and HP StorageWorks Disk Array XP specific term*)

A process in which data backed up in a ZDB-to-tape or a ZDB-to-disk+tape session is restored from tape media to a split mirror replica, which is then synchronized to the source volumes. Individual backup objects or complete sessions can be restored using this method.

See also **ZDB to tape**, **ZDB to disk+tape**, and **replica**.

sqlhosts file (*Informix specific term*)

An Informix connectivity-information file that contains the names of each of the database servers and any aliases to which the clients on a host computer can connect.

SRD file

The Data Protector System Recovery Data (SRD) file contains system information required for installing and configuring the operating system in case of a disaster. The SRD file is an ASCII file, generated when a CONFIGURATION backup is performed on a Windows client and stored on the Cell Manager.

SRDF (*EMC Symmetrix specific term*)

The EMC Symmetrix Remote Data Facility is a business continuation process that enables effective, real-time data replication of SLDs between dislocated processing environments. These environments could be situated within the same root computer environment or separated by long distances.

SSE Agent (*HP StorageWorks Disk Array XP specific term*)

A Data Protector software module that executes all tasks required for a split mirror backup integration. It communicates with the HP StorageWorks Disk Array XP storing system using the RAID Manager XP

Glossary

utility (HP-UX and Windows systems) or RAID Manager Library (Solaris systems).

sst.conf file

The file `/usr/kernel/drv/sst.conf` is required on each Data Protector Sun Solaris client to which a multi-drive library device is connected. It must contain an entry for the SCSI address of the robotic mechanism of each library device connected to the client.

st.conf file

The file `/kernel/drv/st.conf` is required on each Data Protector Solaris client with a backup device connected. It must contain device information and a SCSI address for each backup drive connected to that client. A single SCSI entry is required for a single-drive device and multiple SCSI entries are required for a multi-drive library device.

stackers

Devices with multiple slots for media storage usually with only one drive. A stacker selects media from the stack sequentially. In contrast, a library can randomly select media from its repository.

standalone file device

A file device is a file in a specified directory to which you back up data.

standard security (*MS SQL specific term*)

Standard security uses the login validation process of the Microsoft SQL Server for all connections. Standard security is useful in network environments with a variety of clients, some of which may not support trusted connections. It also provides backward compatibility for older versions of the Microsoft SQL Server.

See also **integrated security**.

Storage Group

(*Microsoft Exchange Server specific term*)

A collection of databases (stores) that share a common set of transaction log files. Exchange manages each storage group with a separate server process.

StorageTek ACS library

(*StorageTek specific term*)

Automated Cartridge System is a library system (also known as Silo) consisting of one Library Management Unit (LMU) and one to 24 Library Storage Modules (LSM) connected to the unit.

storage volume (*ZDB specific term*)

A storage volume represents an object that may be presented to an operating system or some other entity (for example, a virtualization mechanism) upon which volume management systems, file systems, or other objects may exist. The volume management

Glossary

systems, file systems are built on this storage. Typically, these can be created or exist within a storage system such as a disk array.

switchover

See failover

Sybase Backup Server API (*Sybase specific term*)

An industry-standard interface developed for the exchange of backup and recovery information between a Sybase SQL Server and a backup solution like Data Protector.

Sybase SQL Server (*Sybase specific term*)

The server in the Sybase “client-server” architecture. Sybase SQL Server manages multiple databases and multiple users, keeps track of the actual location of data on disks, maintains mapping of logical data description to physical data storage, and maintains data and procedure caches in memory.

Symmetrix Agent (SYMA) (*EMC Symmetrix specific term*)

The Data Protector software module that prepares the EMC Symmetrix environment for backup and restore operations.

System Backup to Tape (*Oracle specific term*)

An Oracle interface that handles the

actions required to load, label, and unload correct backup devices when Oracle issues a backup or restore request.

system databases (*Sybase specific term*)

The four system databases on a newly installed Sybase SQL Server are the:

- master database (master)
- temporary database (tempdb)
- system procedure database (sybsystemprocs)
- model database (model).

system disk

A system disk is a disk containing operating system files. Microsoft terminology defines the system disk as a disk containing the files required for initial step of boot process.

system partition

A system partition is a partition containing operating system files. Microsoft terminology defines a system partition as a partition containing the files required for initial step of boot process.

System State (*Windows specific term*)

The System State data comprises the Registry, COM+ Class Registration

Glossary

database, system startup files, and the Certificate Services database (if the server is a certificate server). If the server is a domain controller, Active Directory directory services and the Sysvol directory are also contained in the System State data. If the server is running the Cluster service, the System State data also includes resource registry checkpoints and the quorum resource recovery log, which contains the most recent cluster database information.

system volume/disk/partition

A volume/disk/partition containing operating system files. Microsoft terminology defines the system volume/disk/partition as the volume/disk/partition containing files required for the initial step of the boot process.

SysVol (*Windows specific term*)

A shared directory that stores the server copy of the domain's public files, which are replicated among all domain controllers in the domain.

tablespace

A part of a database structure. Each database is logically divided into one or more tablespaces. Each tablespace has data files or raw volumes exclusively associated with it.

tapeless backup (*ZDB specific term*)

See ZDB to disk.

target database (*Oracle specific term*)

In RMAN, the target database is the database that you are backing up or restoring.

target (R2) device (*EMC Symmetrix specific term*)

An EMC Symmetrix device that participates in SRDF operations with a source (R1) device. It resides in the remote EMC Symmetrix unit. It is paired with a source (R1) device in the local EMC Symmetrix unit and receives all write data from its mirrored pair. This device is not accessed by user applications during normal I/O operations. An R2 device must be assigned to an RDF2 group type. *See also source (R1) device*

target system (*Disaster Recovery specific term*)

A system after a computer disaster has occurred. The target system is typically in a non-bootable state and the goal of disaster recovery is to restore this system to the original system configuration. The difference between a crashed system and a target system is that a target system has all faulty hardware replaced.

target volume (*ZDB specific term*)

A storage volume to which data is replicated.

Glossary

Terminal Services (*Windows specific term*)

Windows Terminal Services provide a multi-session environment that allows clients to access a virtual Windows desktop session and Windows-based programs running on the server.

thread (*MS SQL Server specific term*)

An executable entity that belongs to only one process. It comprises a program counter, a user-mode stack, a kernel-mode stack, and a set of register values. Several threads can run at a time within one process.

TimeFinder (*EMC Symmetrix specific term*)

A business continuation process that creates an instant copy of single or multiple Symmetrix Logical Devices (SLDs). The instant copy is created on specially preconfigured SLDs called BCVs and is accessible via a separate device address to the system(s).

TLU

Tape Library Unit.

TNSNAMES.ORA (*Oracle and SAP R/3 specific term*)

A network configuration file that contains connect descriptors mapped to service names. The file may be maintained centrally or locally, for use by all or individual clients.

transaction

A mechanism for ensuring that a set of actions is treated as a single unit of work. Databases use transactions to keep track of database changes.

transaction backup

Transaction backups generally use fewer resources than database backups, so they can be created more frequently than database backups. By applying transaction backups, you can recover the database to a specific point in time prior to when a problem occurred.

transaction backup (*Sybase and SQL specific term*)

A backup of the transaction log providing a record of changes made since the last full or transaction backup.

transaction log backup

Transaction log backups generally use fewer resources than database backups so they can be created more frequently than database backups. By applying transaction log backups, you can recover the database to a specific point in time.

transaction log files

Files that record transactions of the database modifications, and provide fault tolerance in case of a database disaster.

Glossary

transaction logs (*Data Protector specific term*)

Keeps track of IDB changes. The archiving of transaction logs should be enabled to prevent you from losing the transaction log files that are created after the last IDB backup and are necessary for IDB recovery.

transaction log table (*Sybase specific term*)

A system table in which all changes to the database are automatically recorded.

transportable snapshot (*MS VSS specific term*)

A shadow copy that is created on the application system and can be presented to the backup system which performs the backup.

See also **Microsoft Volume Shadow Copy service (VSS)**.

TSANDS.CFG file (*Novell NetWare specific term*)

A file that allows you to specify the names of containers where you want backups to begin. It is text file located in the SYS:SYSTEM\TSA directory on the server where TSANDS.NLM is loaded.

unattended operation

See **lights-out operation**.

user account

You can use Data Protector only if you have a Data Protector user account,

which restricts unauthorized access to Data Protector and to backed up data. Data Protector administrators create this account specifying a user logon name, the systems from which the user can log on, and a Data Protector user group membership. This is checked whenever the user starts the Data Protector user interface or performs specific tasks.

user disk quotas

NTFS quota-management support enables an enhanced tracking mechanism and control over disk space usage on shared storage volumes. Data Protector backs up user disk quotas on the whole system and for all configured users at a time.

user group

Each Data Protector user is member of a User Group. Each User Group has a set of user rights that are given to every user in that User Group. The number of User Groups with their associated user rights can be defined as desired. Data Protector provides three default user groups: admin, operator, and user.

user profile (*Windows specific term*)

Configuration information retained on a user basis. This information includes desktop settings, screen colors, network connections, and so on. When the user logs on, the user profile is loaded and the Windows environment is set accordingly.

Glossary

user rights

User rights or access rights are the permissions needed to perform specific Data Protector tasks. Configuring a backup, starting a backup session, or starting a restore session are typical user rights. Users have the access rights of the user group to which they belong.

vaulting media

The process of storing media to a safe and remote place. The media are brought back to the data center when they are needed for restore or are ready for reuse in subsequent backups. The vaulting procedure depends on your company's backup strategy and policies for data protection/reliability.

VBFS (*OmniStorage specific term*)

A Very Big File System is an extension of the standard HP-UX file system on HP-UX 9.x. It is mounted to a directory the same way as any HP-UX file system. In a VBFS, only the superblock, the inode and the 'extended attribute' information remain permanently on the hard disk and are never migrated. *See also MFS.*

verify

A function that lets you check whether the Data Protector data on a specified medium is readable. Additionally, consistency within each block can be

checked if the backup was performed with the cyclic redundancy check (CRC) option ON.

Virtual Controller Software (VCS)

(HP StorageWorks EVA specific term)

The firmware that manages all aspects of storage system operation, including communication with Command View EVA through the HSV controllers.

See also Command View (CV) EVA.

Virtual Device Interface (*MS SQL Server 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 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

Glossary

resources. This way all requests for a particular virtual server are cached by a specific cluster node.

volser (*ADIC and STK specific term*)

A VOLume SERial number is a label on the medium to identify the physical tape used in very large libraries. A volser is a naming convention specific to ADIC/ GRAU and StorageTek devices.

volume group

A unit of data storage in an LVM system. A volume group can consist of one or more physical volumes. There can be more than one volume group on the system.

volume mountpoint (*Windows specific term*)

An empty directory on a volume that can be used to mount another volume. The volume mount point acts as a gateway to the target volume. Provided that the volume is mounted, users and applications can refer to the data on the mounted volume by the full (merged) filesystem path as if both volumes are one integral part.

Volume Shadow Copy service

See **Microsoft Volume Shadow Copy service**.

VPO

See **OVO**.

VSS

See **Microsoft Volume Shadow Copy service**.

VxFS

Veritas Journal Filesystem.

VxVM (Veritas Volume Manager)

A Veritas Volume Manager is a system for managing disk space on Solaris platforms. A VxVM system consists of an arbitrary group of one or more physical volumes organized into logical disk groups.

Wake ONLAN

Remote power-up support for systems running in power-save mode from some other system on the same LAN.

Web reporting

The Data Protector functionality that allows you to view reports on backup status and Data Protector configuration using the Web interface.

wildcard character

A keyboard character that can be used to represent one or many characters. The asterisk (*), for example, typically represents one or more characters, and the question mark (?) typically represents a single character. Wildcard characters are often used in operating systems as a means of specifying more than one file by name.

Glossary

Windows CONFIGURATION backup

Data Protector allows you to back up Windows CONFIGURATION, including Windows Registry, user profiles, Event Logs, and WINS and DHCP server data (if configured on a system) in one step.

Windows Registry

A centralized database used by Windows to store configuration information for the operating system and the installed applications.

WINS server A system running Windows Internet Name Service software that resolves Windows networking computer names to IP addresses. Data Protector can back up WINS server data as part of the Windows configuration.

writer

(MS VSS specific term)

A process that initiates change of data on the original volume. Writers are typically applications or system services that write persistent information on a volume. Writers also participate in the shadow copy synchronization process by assuring data consistency.

XBSA interface *(Informix specific term)*

The onbar utility and Data Protector communicate with each other through

the X/Open Backup Specification Services Programmer's Interface (XBSA).

XCOPY engine *(direct backup specific term)*

A SCSI-3 copy command that allows you to copy data from a storage device having a SCSI source address to a backup device having a SCSI destination address, thus enabling direct backup. The data flows from a source device (either block or streaming, that is, disk or tape) to the destination device (either block or streaming) through XCOPY. This releases the controlling server of reading the data from the storage device into memory and then writing the information to the destination device.

See also **direct backup**.

ZDB

See **zero downtime backup (ZDB)**.

ZDB database *(ZDB specific term)*

A part of the IDB, storing ZDB related information such as source volumes, replicas and security information. The ZDB database is used for ZDB, instant recovery, and split mirror restore.

See also **zero downtime backup (ZDB)**.

ZDB to disk *(ZDB specific term)*

A form of zero downtime backup where the replica produced is kept on the disk

Glossary

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.**

Glossary

A

- advantages
 - MS Exchange Server integration, 303
 - MS SQL Server integration, 253
 - Oracle integration, 3
 - SAP R/3 integration, 145
 - VSS integration, 361
- application systems
 - SAP R/3 Database Servers, configuring, 182
 - SAP R/3 Oracle databases, configuring, 174
- architecture
 - SAP R/3 integration, 156, 157
 - VSS integration, 356, 364

B

- backing up Lotus
 - backup options, 67
- backing up MS Exchange Server, 332–335
 - backup flow, 310
 - backup options, 315
 - backup specifications, creating, 313–328
 - backup types, 303
 - scheduling backups, 332
 - starting backups, 334
- backing up MS SQL Server, 279–282
 - backup flow, 258
 - backup methods, 279
 - backup options, 277
 - backup specifications, creating, 268
 - backup types, 253
 - scheduling backups, 279
 - starting backups, 281
- backing up Oracle
 - backup options, 67
 - backup set ZDB concepts, 14–??
 - backup set ZDB session flow, 17–??
 - backup specifications, creating, 53
 - backup types, 4
 - offline ZDB, 64
 - proxy-copy ZDB concepts, 21–??
 - proxy-copy ZDB session flow, 24–??
 - scheduling backups, 76
 - starting backups, 75–79
 - starting backups, using GUI, 77
 - verifying backups, 133
- backing up SAP R/3, 213–217
 - backup concepts, scheme, 156
 - backup flow, backint mode, 157
 - backup methods, 213
 - backup options, 207
 - backup owners, configuring, 173
 - backup specifications, creating, 195
 - backup types, 146
 - backup utilities, 154
 - scheduling backups, 214
 - starting backups, 213–217
 - starting backups, using CLI, 217
 - starting backups, using GUI, 216
 - verifying backups, 233
 - ZDB concepts, 161–165
 - ZDB session flow, 163–165
- backing up VSS, 376–386
 - architecture, 356, 364
 - backup flow, 362
 - backup specifications, creating, 376
 - backup types, 357
 - example, MS Exchange Server 2003, 405
 - scheduling backups, 385
 - starting backups, 386
- backint mode, SAP R/3 integration
 - backup flow, 157
 - restore flow, 160
- backup flow
 - MS SQL Server integration, 258
 - VSS integration, 362
- backup flow, Oracle integration, 10–11
 - backup set ZDB session flow, 17–??
 - proxy-copy ZDB session flow, 24–??
- backup flow, SAP R/3 integration
 - backint mode, 157
 - ZDB session flow, 163–165
- backup methods
 - MS SQL Server integration, 279
 - SAP R/3 integration, 213
- backup options
 - Lotus integration, 67
 - MS Exchange Server integration, 315
 - MS SQL Server integration, 277
 - Oracle integration, 67
 - SAP R/3 integration, 207
- backup owners, configuring
 - SAP R/3 integration, 173
- backup specifications, creating
 - MS Exchange Server integration, 313–328
 - MS SQL Server integration, 268
 - Oracle integration, 53
 - SAP R/3 integration, 195
 - VSS integration, 376
- backup specifications, modifying

Index

- MS Exchange Server integration, 328
 - backup specifications, ownership
 - Oracle integration, 37
 - backup specifications, scheduling
 - MS Exchange Server integration, 332
 - MS SQL Server integration, 279
 - Oracle integration, 76
 - SAP R/3 integration, 214
 - VSS integration, 385
 - backup system
 - configuring Oracle client, backup set ZDB, 32
 - backup types
 - MS Exchange Server integration, 303
 - MS SQL Server integration, 253
 - Oracle integration, 4
 - SAP R/3 integration, 146
 - VSS integration, 357
 - backup utilities
 - SAP R/3 integration, 154
 - BRARCHIVE, 154, 208
 - BRBACKUP, 154, 164, 208
 - BRRESTORE, 155, 160
- ## C
- checking configuration
 - MS SQL Server integration, 266
 - Oracle integration, ZDB configuration, 50
 - SAP R/3 integration, 190
 - SAP R/3 integration, ZDB configuration, 193
 - checking database files for consistency
 - MS Exchange Server integration, 329
 - concepts
 - MS Exchange Server integration, 310–311
 - MS SQL Server integration, 258–260
 - Oracle integration, 9
 - Oracle integration, backup set ZDB concepts, 14–??
 - Oracle integration, proxy-copy ZDB concepts, 21–??
 - SAP R/3 integration, 154–160
 - SAP R/3 integration, ZDB concepts, 161–165
 - VSS integration, 361–366
 - configuration files
 - MS SQL Server integration, 261
 - SAP R/3 integration, 166–172
 - configuration files, modifying
 - SAP R/3 integration, 169
 - configuring MS SQL Server, 263–278
 - checking configuration, 266
 - configuration files, 261
 - configuring Oracle, 28–??
 - client on backup system, backup set ZDB, 32
 - offline ZDB, 64
 - Servers, 39
 - users, 37
 - using CLI, 45
 - using GUI, 40
 - verifying configuration, 132
 - ZDB methods, switching between, 51
 - configuring SAP R/3, 173–194
 - backup owners, 173
 - checking configuration, 190
 - checking ZDB configuration, 193
 - configuring Oracle databases, 174
 - Database Servers, 182
 - Inet user account, 184
 - testing ZDB configuration, 193
 - verifying configuration, 232
 - configuring VSS, 367–369
 - control files
 - Oracle integration, restore, 87
 - conventions, xi
 - creating backup specifications
 - MS Exchange Server integration, 313–328
 - MS SQL Server integration, 268
 - Oracle integration, 53
 - SAP R/3 integration, 195
 - VSS integration, 376
 - creating parameter files
 - SAP R/3 integration, 211
- ## D
- data objects
 - SAP R/3 integration, 154
 - Data Protector Database Library
 - linking with Oracle, 35
 - database recovery
 - Oracle integration, after instant recovery, 117
 - Oracle integration, options, 96
 - SAP R/3 integration, options, 226
 - SAP R/3, after instant recovery, 224
 - disaster recovery
 - Oracle integration, 83

E

- examples, MS Exchange Server integration
 - creating backup specifications, 325
- examples, Oracle integration
 - restoring using RMAN, 102
- examples, SAP R/3 integration
 - restoring, 240–249
- examples, VSS integration
 - backing up, 405
 - Microsoft Exchange Server 2003, ??–408
 - MS Exchange Server 2003, 405–??
 - restoring, 407

F

- finding users
 - Oracle integration, 37

I

- Inet user account, configuring
 - SAP R/3 integration, 184
- instant recovery
 - MS Exchange Server integration, 346–349
 - MS SQL Server integration, 290–292, 397–??
 - Oracle integration, 117–125
 - Oracle integration, database recovery after, 122
 - RAC preparation steps, 118
 - reconfiguring an Oracle instance, A-3
 - SAP R/3 integration, 224–228
 - SAP R/3 integration, options for database recovery, 226
- interactive backups
 - MS Exchange Server integration, 334
 - MS SQL Server integration, 281
 - Oracle integration, 77
 - SAP R/3 integration, 215
 - VSS integration, 386
- introduction
 - MS Exchange Server integration, 303
 - MS SQL Server integration, 253
 - Oracle integration, 3
 - SAP R/3 integration, 145
 - VSS integration, 355

L

- limitations
 - MS Exchange Server integration, 309

- MS SQL Server integration, 257
- SAP R/3 integration, 152
- VSS integration, 359
- linking to the Data Protector Database Library
 - Oracle integration, 35
- Lotus backup
 - backup options, 67

M

- modifying backup specifications
 - MS Exchange Server integration, 328
- modifying configuration files
 - SAP R/3 integration, 169
- modifying parameter files
 - SAP R/3 integration, 211
- MS Exchange Server backup, 332–335
 - backup flow, 310
 - backup options, 315
 - backup specifications, creating, 313–328
 - backup types, 303
 - scheduling backups, 332
 - starting backups, 334
- MS Exchange Server integration
 - advantages, 303
 - backup, 332–335
 - concepts, 310–311
 - introduction, 303
 - limitations, 309
 - prerequisites, 306
 - restore, 336–349
 - troubleshooting, 350–352
- MS Exchange Server restore, 336–349
 - instant recovery, 346–349
 - point in time recovery, 336, 337
 - restore types, 304
 - roll forward recovery, 336, 341
 - to the application system, 336
- MS Exchange Server troubleshooting, 350–352
- MS SQL Server backup, 279–282
 - backup flow, 258
 - backup methods, 279
 - backup options, 277
 - backup specifications, creating, 268
 - backup types, 253
 - scheduling backups, 279
 - starting backups, 281
- MS SQL Server configuration, 263–278
 - checking configuration, 266

- configuration files, 261
- MS SQL Server integration
 - advantages, 253
 - backup, 279–282
 - backup types, 253
 - concepts, 258–260
 - configuration, 263–278
 - introduction, 253
 - limitations, 257
 - prerequisites, 256
 - restore, 283–292
 - restore types, 254
 - troubleshooting, 293–300
- MS SQL Server restore, 283–292
 - instant recovery, 290–292, 397–??
 - restore flow, 259
 - restore options, 287
 - restore types, 254
 - using GUI, 284
- MS SQL Server troubleshooting, 293–300

O

- omnirc variables, A-9–A-14
- Oracle backup
 - backup concepts, scheme, 12
 - backup set ZDB concepts, 14–??
 - backup set ZDB session flow, 17–??
 - backup specifications, creating, 53
 - backup types, 4
 - proxy-copy concepts, 21–??
 - proxy-copy ZDB session flow, 24–??
 - scheduling backups, 76
 - starting backups, 75–79
 - starting backups, using GUI, 77
 - verifying backups, 133
- Oracle configuration
 - checking ZDB configuration, 50
 - testing ZDB configuration, 50
 - users, 37
 - verifying configuration, 132
 - ZDB methods, switching between, 51
- Oracle integration
 - advantages, 3
 - backup set ZDB concepts, 14–??
 - backup set ZDB session flow, 17–??
 - backup types, 4
 - concepts, 9
 - configuration, 28–??
 - configuring Servers, 39

- configuring Servers, using GUI, 40
- instant recovery, 117–125
- introduction, 3
- linking to the Database Library, 35
- proxy-copy ZDB concepts, 21–??
- proxy-copy ZDB session flow, 24–??
- removing the integration, 126
- removing the integration, from HP-UX, 126
- removing the integration, from Solaris and other UNIX systems, 127
- restore flow, 11
- restore types, 5
- standard restore procedure, 80–117
- troubleshooting, 129–142

Oracle restore

- control files, 87
- database items, 80
- database objects, 88
- database recovery after instant recovery, 117
- examples, using RMAN, 102
- instant recovery, 117–125
- preparing databases for restore, 102
- recovery catalog database, 84
- recovery catalog, using CLI, 116
- restore options, 96
- restore types, 5
- standard procedure, 80–117
- tablespaces and datafiles, 92
- to the application system, 82
- using another device, 117
- using GUI, 83
- using RMAN, 102
- verifying restores, 134

Oracle RMAN script, 69

Oracle troubleshooting, 129–142

- verifying backups, 133
- verifying configuration, 132
- verifying prerequisites on Oracle side, 129
- verifying restores, 134

ownership, backup specifications

- Oracle integration, 37

P

- parameter files, creating
 - SAP R/3 integration, 211
- parameter files, modifying
 - SAP R/3 integration, 211
- prerequisites

- MS Exchange Server integration, 306
- MS SQL Server integration, 256
- SAP R/3 integration, 150
- VSS integration, 359
- prerequisites, verifying
 - Oracle integration, on Oracle side, 129
 - SAP R/3 integration, on Oracle side, 230
 - SAP R/3 integration, on SAP R/3 side, 231

R

RAC

- preparing for instant recovery, 118
- RAC, configuring Oracle Servers
 - on HP-UX, 34
 - on other UNIX systems, 35
- reconfiguring an Oracle instance for instant recovery, A-3

recovery

- Oracle database after instant recovery, 117
- Oracle integration, options, 96
- SAP R/3 database after instant recovery, 224
- SAP R/3 integration, options, 226

recovery catalog database

- Oracle integration, restore, 84
- removing the Oracle integration, 126
 - from HP-UX, 126
 - from Solaris and other UNIX systems, 127

restore flow

- MS SQL Server integration, 259
- SAP R/3 integration, backint mode, 160
- VSS integration, 365

restore options

- MS SQL Server integration, 287
- VSS integration, 390

restore types

- MS Exchange Server integration, 304
- MS SQL Server integration, 254
- Oracle integration, 5
- SAP R/3 integration, 146

restoring MS Exchange Server, 336–349

- instant recovery, 346–349
- point in time recovery, 336, 337
- restore types, 304
- roll forward recovery, 336, 341
- to the application system, 336
- restoring MS SQL Server, 283–292
 - instant recovery, 290–292, 397–??
 - restore flow, 259
 - restore options, 287

- restore types, 254
 - using GUI, 284
- restoring Oracle
 - control files, 87
 - database objects, 88
 - database recovery after instant recovery, 117
 - instant recovery, 117–125
 - recovery catalog database, 84
 - recovery catalog, using CLI, 116
 - restore flow, 11
 - restore types, 5
 - standard restore procedure, 80–117
 - tablespaces and datafiles, 92
 - to the application system, 82
 - using another device, 117
 - using GUI, 83
 - using RMAN, 102
 - verifying restores, 134
- restoring SAP R/3, 218–228
 - archive log files, example, 248
 - database recovery after instant recovery, 224
 - examples, 240–249
 - full database, example, 242
 - instant recovery, 224–228
 - lost files, example, 247
 - partial, example, 246
 - preparing database for restore, 240
 - restore flow, backint mode, 160
 - restore types, 146
 - standard restore procedure, 218–224
 - to the application system, 218
 - verifying restores, 234
- restoring VSS, 387–399
 - example, MS Exchange Server 2003, 407
 - restore flow, 365
 - restore options, 390
 - using GUI, 388
- RMAN, Oracle integration
 - restore, 102
- running backups *See* starting backups

S

- SAP R/3 backup, 213–217
 - backup concepts, scheme, 156
 - backup flow, backint mode, 157
 - backup methods, 213
 - backup options, 207

- backup owners, configuring, 173
 - backup specifications, creating, 195
 - backup types, 146
 - backup utilities, 154
 - scheduling backups, 214
 - starting backups, 213–217
 - starting backups, using CLI, 217
 - starting backups, using GUI, 216
 - verifying backups, 233
 - ZDB concepts, 161–165
 - ZDB session flow, 163–165
 - SAP R/3 configuration, 173–194
 - backup owners, 173
 - checking configuration, 190
 - checking ZDB configuration, 193
 - configuring Oracle databases, 174
 - Database Servers, 182
 - Inet user account, 184
 - testing ZDB configuration, 193
 - verifying configuration, 232
 - SAP R/3 integration
 - advantages, 145
 - architecture, 157
 - backup, 213–217
 - backup types, 146
 - concepts, 154–160
 - configuration, 173–194
 - configuration files, 166–172
 - data objects, 154
 - instant recovery, 224–228
 - introduction, 145
 - limitations, 152
 - parameter files, creating, 211
 - parameter files, modifying, 211
 - prerequisites, 150
 - restore, 218–228
 - restore types, 146
 - standard restore procedure, 218–224
 - troubleshooting, 229–239
 - util_cmd, 169
 - ZDB concepts, 161–165
 - ZDB session flow, 163–165
 - SAP R/3 restore, 218–228
 - archive log files, example, 248
 - database recovery after instant recovery, 224
 - examples, 240–249
 - full database, example, 242
 - instant recovery, 224–228
 - lost files, example, 247
 - partial, example, 246
 - preparing database for restore, 240
 - restore flow, backint mode, 160
 - restore types, 146
 - standard procedure, 218–224
 - to the application system, 218
 - verifying restores, 234
 - SAP R/3 troubleshooting, 229–239
 - verifying backups, 233
 - verifying configuration, 232
 - verifying prerequisites on Oracle side, 230
 - verifying prerequisites on SAP R/3 side, 231
 - verifying restores, 234
 - sapdba, 155
 - scheduling backups
 - MS Exchange Server integration, 332
 - MS SQL Server integration, 279
 - Oracle integration, 76
 - SAP R/3 integration, 214
 - VSS integration, 385
 - starting backups
 - MS Exchange Server integration, 334
 - MS SQL Server integration, 281
 - VSS integration, 386
 - starting backups, Oracle integration, 75–79
 - using GUI, 77
 - starting backups, SAP R/3 integration, 213–217
 - using CLI, 217
 - using GUI, 216
 - switching between Oracle ZDB methods, 51
- ## T
- testing configuration
 - Oracle ZDB configuration, 50
 - SAP R/3 integration, ZDB configuration, 193
 - troubleshooting MS Exchange Server, 350–352
 - troubleshooting MS SQL Server, 293–300
 - troubleshooting Oracle, 129–142
 - verifying backups, 133
 - verifying configuration, 132
 - verifying prerequisites on Oracle side, 129
 - verifying restores, 134
 - troubleshooting SAP R/3, 229–239
 - verifying backups, 233
 - verifying configuration, 232
 - verifying prerequisites on Oracle side, 230

- verifying prerequisites on SAP R/3 side, 231
- verifying restores, 234
- troubleshooting VSS, 400–404
- typographical conventions, xi

U

- users, configuring
 - Oracle integration, 37
- users, finding
 - Oracle integration, 37
- util_cmd
 - SAP R/3 integration, 169
- util_oracle8.exe, 133

V

- verifying backups
 - Oracle integration, 133
 - SAP R/3 integration, 233
- verifying configuration
 - Oracle integration, 132
 - SAP R/3 integration, 232
- verifying prerequisites
 - Oracle integration, on Oracle side, 129
 - SAP R/3 integration, on Oracle side, 230
 - SAP R/3 integration, on SAP R/3 side, 231
- verifying restores
 - Oracle integration, 134
 - SAP R/3 integration, 234
- Volume Shadow Copy service
 - VSS integration, 355
- VSS backup, 376–386
 - architecture, 356, 364
 - backup flow, 362
 - backup specifications, creating, 376
 - backup types, 357
 - example, MS Exchange Server 2003, 405
 - scheduling backups, 385
 - starting backups, 386
- VSS configuration, 367–369
- VSS integration
 - advantages, 361
 - architecture, 356, 364
 - backup, 376–386
 - concepts, 361–366
 - configuration, 367–369
 - introduction, 355
 - limitations, 359
 - prerequisites, 359
 - restore, 387–399

- troubleshooting, 400–404
- VSS restore, 387–399
 - example, MS Exchange Server 2003, 407
 - restore flow, 365
 - restore options, 390
 - using GUI, 388
- VSS troubleshooting, 400–404

W

- writers specifics, VSS integration, 371
 - MS Exchange Server 2003, 382, 392
 - MSDE, 391

Z

- ZDB session flow
 - Oracle integration, backup set, 17–??
 - Oracle integration, proxy-copy, 24–??
 - SAP R/3 integration, 163–165

