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

Manual Edition: October 2004

Manufacturing Part Number: B6960-90113

Release A.05.50

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

Legal Notices

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

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

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

UNIX® is a registered trademark of The Open Group.

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

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

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

ARM® is a registered trademark of ARM Limited.

Contents

I. HP StorageWorks Virtual Array

1. Configuration

In This Chapter 4
Prerequisites and Limitations 5
Configuring the Integration 7
Automatic Configuration of Backup System 8
Configuring VA with HP StorageWorks AutoPath Installed 9
ZDB Database—VADB 12
Querying the VADB 13
Checking the Consistency of the VADB 14
VA LUN Exclude File 14
LUN Security 16
HP OpenView Storage Area Manager Password 17
Deleting the Contents of the VADB 18

2. Backup

In This Chapter 22
Backup Process 23
ZDB Types 23
VA Backup Flows 25
Configuring a Backup Specification 29
Backup Options 36
Troubleshooting 41
Before You Begin 41
Backup Problems 41

3. Restore

In This Chapter 46
Overview 47
Restoring from Backup Media on LAN 48
Instant Recovery 49
Instant Recovery Process 50
Instant Recovery Procedure 51
Instant Recovery Options 55
Instant Recovery in a Cluster 56
Troubleshooting 57
Before You Begin 57

Instant Recovery Problems 57

II. HP StorageWorks Enterprise Virtual Array

4. Configuration

In This Chapter 62
Prerequisites and Limitations 63
Configuring the Integration 67
Automatic Configuration of Backup System 68
ZDB Database—EVADB or SMISDB 70
Setting the Login Information for CV EVA 72
Updating the Information on the EVA Hardware Configuration in the EVADB . . . 74
EVA Disk Group Pairs Configuration File 75
Querying the EVADB/SMISDB 79
Purging the SMISDB 83
Deleting the Target Volumes Created in a Specific Backup Session 84
Deleting all Target Volumes in a Specific Replica Set 84
Synchronizing the EVADB/SMISDB 85

5. Backup

In This Chapter 88
Backup Process 89
ZDB Types 89
EVA Backup Flows 92
Configuring a Backup Specification 95
Backup Options 103
Troubleshooting 109
Before You Begin 109
Backup Problems 109

6. Restore

In This Chapter 122
Overview 123
Restoring from Backup Media on LAN 124
Instant Recovery 125
Instant Recovery Process 126
Instant Recovery Procedure 127

Instant Recovery Options 131
Instant Recovery in a Cluster 131
Troubleshooting 132
Before You Begin 132
Instant Recovery Problems 132

III. HP StorageWorks Disk Array XP

7. Configuration

In This Chapter 138
Prerequisites and Limitations 139
Configuring the Integration 142
Preparing the Environment 142
Automatic Configuration of Backup System 144
ZDB Database—XPDB 145
Querying the XPDB 147
XP Command Device Handling 148
XP LDEV Exclude File 151

8. Backup

In This Chapter 156
Backup Process 157
ZDB Types 157
XP Backup Flow 158
Configuring a Backup Specification 164
Backup Options 169
Client Systems Options 169
Mirror Type Options 170
Application Options 170
Instant Recovery Option 172
Replica Management Options 173
Mirror Disk Preparation/Synchronization Options 175
Mount Options 176
Troubleshooting 178
Before You Begin 178
Creation of the Backup Specification 179
Backup Problems 179

Contents

9. Restore

In This Chapter 184
Overview 185
Restoring from Backup Media on LAN 186
Split Mirror Restore 189
Split Mirror Restore Process 189
Split Mirror Restore Procedure 190
Split Mirror Restore Options 193
Split Mirror Restore in a Cluster 195
Instant Recovery 197
Instant Recovery Process 198
Instant Recovery Procedure 199
Instant Recovery Options 203
Instant Recovery and LVM Mirroring 204
Instant Recovery in a Cluster 204
Troubleshooting 205
Before You Begin 205
Split Mirror Restore Problems 205
Instant Recovery Problems 206

IV. EMC Symmetrix

10. Configuration

In This Chapter 212
Prerequisites and Limitations 213
Configuring the Integration 214
Preparing the Environment 214
Creating the EMC Symmetrix Database File 215
Creating the Data Protector EMC Database File 216
Automatic Configuration of Backup System 216
The Data Protector EMC Log File 217

11. Backup

In This Chapter 220
Backup Process 221
ZDB Types 221
EMC Backup Flow 221

Configuring a Backup Specification 225
Backup Options 229
Backup Disk Usage 231
Testing Your Backed Up Data 232
EMC Test Options 232
Checking Your Restored Data 233
Troubleshooting 235
Before You Begin 235
Recovery Using the EMC Agent 235
Creation of the Backup Specification 237
Backup Problems 238
Error Messages 240

12. Restore

In This Chapter 244
Overview 245
Restoring from Backup Media on LAN 246
Split Mirror Restore 249
Split Mirror Restore Process 249
Split Mirror Restore Procedure 251
Split Mirror Restore Options 253
Split Mirror Restore in a Cluster 255
Troubleshooting 257
Before You Begin 257
Recovery Using the EMC Agent 257
Split Mirror Restore Problems 259
Error Messages 261

A. Appendix

In This Appendix A-2
Running and Scheduling a ZDB Session A-3
Scheduling a ZDB Session A-4
Starting an Interactive ZDB Session A-5
Alternate Paths Support A-8
HP StorageWorks AutoPath Limitations and Considerations A-9
Cluster Configurations A-10
Client on the Application System in a Cluster A-10
Cell Manager and Client on the Application System in a Cluster A-11

Contents

Client on the Application System in a Cluster, Cell Manager on the Backup System in a Cluster A-13
Cell Manager on the Backup System in a Cluster A-15
Client on the Application System in a Cluster, Cell Manager in a Cluster A-17
EMC GeoSpan for Microsoft Cluster Service Solution A-18
Instant Recovery in a Cluster A-20
MC/ServiceGuard Procedure A-20
Microsoft Cluster Server Procedure A-22
ZDB Agents Omnirc Variables A-23
Common ZDB Agents Variables A-23
VA Agent Specific Variables A-29
EVA Agent Specific Variables A-29
XP Agent Specific Variables A-33
EMC Agent Specific Variables A-34
User Scenarios—ZDB Options Exemplary Selections A-37
VA and EVA Integrations A-37
XP Integration A-39
EMC Integration A-41
EMC - Obtaining Disk Configuration Data A-42
Additional Information for Troubleshooting A-47
Listing and Unlocking Locked Backup Devices and Target Volumes A-48

Glossary

Index

Printing History

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

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

Table 1 **Edition History**

Part Number	Manual Edition	Product
B6960-90113	October 2004	Data Protector Release A.05.50

Conventions

The following typographical conventions are used in this manual.

Table 2

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

Data Protector provides a cross-platform (Windows and UNIX) graphical user interface. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for information about the Data Protector graphical user interface.

Figure 1 Data Protector Graphical User Interface

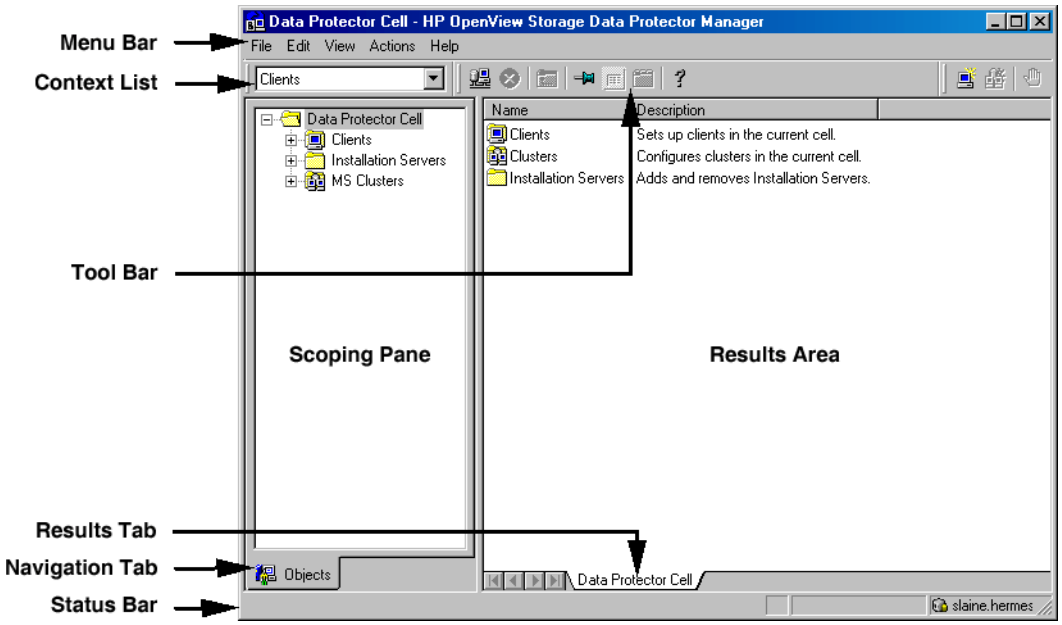

Contact Information

General Information

General information about Data Protector can be found at

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

Technical Support

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

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

http://www.hp.com/support

Information about the latest Data Protector patches can be found at

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

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

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

Documentation Feedback

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

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

Training Information

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

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

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

Data Protector Documentation

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

Manuals

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

HP OpenView Storage Data Protector Concepts Guide

This manual describes Data Protector concepts and provides background information on how Data Protector works. It is intended to be used with the task-oriented *HP OpenView Storage Data Protector Administrator's Guide*.

HP OpenView Storage Data Protector Administrator's Guide

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

HP OpenView Storage Data Protector Installation and Licensing Guide

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

HP OpenView Storage Data Protector Integration Guide

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

- *HP OpenView Storage Data Protector Integration Guide for Microsoft Applications: SQL Server 7/2000, Exchange Server 5.x, Exchange Server 2000/2003, and Volume Shadow Copy Service*

This manual describes the integrations of Data Protector with the following Microsoft applications: Microsoft Exchange Server 2000/2003, Microsoft Exchange Server 5.x, Microsoft SQL Server 7/2000, and Volume Shadow Copy Service.

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

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

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

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

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

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

HP OpenView Storage Data Protector Integration Guide for HP OpenView

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

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

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

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

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

HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide

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

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

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

HP OpenView Storage Data Protector Zero Downtime Backup Integration Guide

This manual describes how to configure and use Data Protector to perform zero downtime backup, instant recovery, and standard restore of Oracle, SAP R/3, Microsoft Exchange Server 2000/2003, and Microsoft SQL Server 2000 databases. The manual also describes how to configure and use Data Protector to perform backup and restore using the Microsoft Volume Shadow Copy Service.

HP OpenView Storage Data Protector MPE/iX System User Guide

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

HP OpenView Storage Data Protector Media Operations User's Guide

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

HP OpenView Storage Data Protector Software Release Notes

This manual gives a description of new features of HP OpenView Storage Data Protector A.05.50. It also provides information on supported configurations (devices, platforms and online database integrations, SAN, and ZDB), required patches, and limitations, as well as known problems and workarounds. An updated version of the supported configurations is available at http://www.openview.hp.com/products/datapro/spec_0001.html.

Online Help

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

In This Book

The *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide* describes how to configure and use integrations of Data Protector with the following disk arrays:

- HP StorageWorks Virtual Array
- HP StorageWorks Enterprise Virtual Array
- HP StorageWorks Disk Array XP
- EMC Symmetrix

Audience

This manual is intended for backup administrators or operators intent on configuring and using the integration of Data Protector with one of the supported disk arrays.

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:

Part I	HP StorageWorks Virtual Array
Chapter 1	“Configuration” on page 3
Chapter 2	“Backup” on page 21
Chapter 3	“Restore” on page 45
Part II	HP StorageWorks Enterprise Virtual Array
Chapter 4	“Configuration” on page 61
Chapter 5	“Backup” on page 87
Chapter 6	“Restore” on page 121
Part III	HP StorageWorks Disk Array XP
Chapter 7	“Configuration” on page 137
Chapter 8	“Backup” on page 155
Chapter 9	“Restore” on page 183
Part IV	EMC Symmetrix
Chapter 10	“Configuration” on page 211
Chapter 11	“Backup” on page 219
Chapter 12	“Restore” on page 243
Appendix A	“Appendix” on page A-1
Glossary	Definition of terms used in this manual.

I HP StorageWorks Virtual Array

1 Configuration

In This Chapter

This chapter describes the procedure for configuring Data Protector HP StorageWorks Virtual Array (VA) integration. It also provides information on the ZDB database.

For a detailed description of the installation of the Data Protector Cell Manager and clients, refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

This chapter includes the following sections:

"Prerequisites and Limitations" on page 5

"Configuring the Integration" on page 7

"Configuring VA with HP StorageWorks AutoPath Installed" on page 9

"ZDB Database—VADB" on page 12

Prerequisites and Limitations

Prerequisites

- You need a special license to use the Data Protector VA integration. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*.
- The Data Protector VA integration must be correctly installed. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*.
- The same operating system and its version must be installed on the application and the backup system.
- The following VA components are required for this integration:
  - ✓ Array microcode hp15 (minimum)
  - ✓ HP StorageWorks Command View SDM software and license for managing and controlling the VA storage system
- You should be familiar with the VA concepts and procedures. Refer to the VA related documentation.
- You should be familiar with the basic ZDB and instant recovery concepts. Refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.
- Refer to the *HP OpenView Storage Data Protector Software Release Notes* for information on:
  - ✓ general Data Protector limitations
  - ✓ supported platforms
  - ✓ supported integrations
  - ✓ supported backup topologies
  - ✓ supported connectivity topologies
  - ✓ supported cluster configurations (high availability support)

Limitations

- Limit the size of LUN 0 to 10 MB and do not use it for storing data, as LUN 0 is used as a command device and is accessed by all hosts connected to the disk array.

- Instant recovery from a ZDB-to-disk+tape session cannot be performed using the Data Protector GUI after exporting or overwriting the media used in the backup session. The backup media must not be exported or overwritten even after an object copy session. If the backup media have been exported or overwritten, perform instant recovery using the Data Protector CLI. Refer to “Instant Recovery Using the Data Protector CLI” on page 54 for more information on how to perform instant recovery using the Data Protector CLI.
- If instant recovery is performed, all target volumes (child LUNs or snapshots) for the source volume (parent LUN) involved in the instant recovery session will be deleted automatically before the restore takes place. The only target volumes (replica) that can be kept on the array are those that will be restored if the Data Protector instant recovery option `Keep the replica after the restore` is selected.

  If there are any other target volumes of the same source volume (for instance, from other backup specification or created for purposes other than Data Protector backup and restore) existing on the array, the instant recovery will abort. In this case, the target volumes concerned must be deleted before the instant recovery can be performed.
- Preview backup is not supported.
- Object copying and object mirroring is not supported for ZDB to disk.

Configuring the Integration

The following sections assume that you have chosen the desired backup configuration and connected VA to the application system(s) and the backup system. The source volumes should be presented to the application system(s) and, if necessary, mounted as filesystems on the system(s). For information on the supported configurations and their descriptions, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

Using HP StorageWorks Command View SDM, create source volumes and present them to the application system.

Using HP-UX LVM or Windows Disk Management utility, configure the filesystems on the application system, if necessary, and mount them.

IMPORTANT

For ZDB-to-disk sessions, 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 configure a backup specification for a ZDB-to-disk session. For information on configuring a standalone device, refer to the online Help index keyword “standalone devices”.

If HP StorageWorks Secure Manager is enabled for the appropriate VA, provide the password for this array using the `omnidbva` command. For instructions, refer to “LUN Security” on page 16.

If your VA is part of a SAN environment in which HP OpenView Storage Allocater is installed, perform some additional configuration steps. For instructions, refer to “HP OpenView Storage Area Manager Password” on page 17.

If your VA has the HP StorageWorks AutoPath software installed, perform some additional configuration steps. For instructions, refer to “Configuring VA with HP StorageWorks AutoPath Installed” on page 9.

To prevent Data Protector from using certain target volumes, refer to “VA LUN Exclude File” on page 14.

Automatic Configuration of Backup System

The VA integration do not require any configuration steps, such as configuring the volume groups and filesystems on the backup system. This is done by Data Protector automatically when a ZDB session is started. Based on the volume group, filesystem, and mount point configuration on the application system, Data Protector creates the same volume group and filesystem structure on the backup system. During the ZDB-to-tape and ZDB-to-disk+tape sessions, Data Protector mounts these filesystems. In case of disk images, raw device files are used. For more information on the backup system mountpoint creation, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

Configuring VA with HP StorageWorks AutoPath Installed

This section is relevant only for the Data Protector VA integration with the HP StorageWorks AutoPath software installed and configured on the backup system.

There is a limitation regarding the AutoPath software installed on VA as described below.

Problem

The backup system with AutoPath installed cannot detect the newly created snapshots if the system is not rebooted after the snapshot creation. Since Data Protector can create new snapshots (representing replica), such replicas must be pre-configured and the backup system rebooted before running backups, which enables AutoPath to function properly.

During the ZDB sessions, the Data Protector VA integration will then reuse these pre-configured snapshots from the replicas.

IMPORTANT

Before pre-configuring the snapshots, consider very carefully which backup objects that represent different source volumes will be used in backup specifications. Select the backup objects in such a way that source volumes that represent these backup objects are not included in more than one backup specification.

Action

The following is a procedure on how to pre-configure snapshots that will then represent a replica:

1. Create a backup specification as it is described in “Configuring a Backup Specification” on page 29. While creating the backup specification, consider the following options selections:

ZDB to Disk and ZDB to Disk+Tape

   a. If you intend to perform ZDB-to-disk or ZDB-to-disk+tape sessions (instant recovery enabled), use the backup options in the following way:

      - Leave the `Track the replica for instant recovery` option selected.

- Specify the `Number of replicas rotated`.

NOTE Each backup session creates one replica. Therefore, you will need to run as many backup sessions as you specified using the `Number of replicas rotated` option. Only then can snapshots representing a replica be reused, so new ones need not to be created.

ZDB to Tape

b. If you intend to perform ZDB-to-tape sessions, use the backup options in the following way:

- Deselect `Track the replica for instant recovery`.
- Select `Keep the replica after backup`.
- Do not select `Use an existing replica`.

2. Run the backup session as it is described in "Running and Scheduling a ZDB Session" on page A-3, and consider the following:

- For ZDB-to-disk and ZDB-to-disk+tape sessions, run so many backup sessions using the same backup specification that the number of replicas is reached.
- For ZDB-to-tape sessions, run the backup session for one backup specification only once. If you intend to run parallel backup sessions, make sure that for each backup specification you will have a replica created (one backup session performed).

NOTE During these backup sessions, AutoPath functionality (failover to an alternate path) will not work. Therefore, the backup sessions must complete successfully without a path failure.

3. After successfully performed backups, reboot the backup system. Now, the AutoPath software can function properly and is aware of all created snapshots.

ZDB to Disk and ZDB to Disk+Tape

4. Use the configured backup specification for your further ZDB-to-disk or ZDB-to-disk+tape sessions. The snapshots will be reused, and you will need to perform these steps again only if you create a new backup specification with new backup objects selected.

ZDB to Tape

5. Modify the backup specification so that you select also the `Use an existing replica` option (and leave the `Keep the replica after backup` option selected). Use such backup specification for your further ZDB-to-tape sessions. The snapshots will be reused and only if you create a new backup specification with new backup objects selected, you will need to perform these steps again.

IMPORTANT

If you do not select the `Use an existing replica` option for your further backup sessions, Data Protector will create a new replica with each backup session. As a consequence, AutoPath will not function properly.

ZDB Database—VADB

The **ZDB database** is in case of the VA integration referred to as **VADB**. The VADB keeps the following information:

- Information on all Data Protector VA ZDB-to-disk and ZDB-to-disk+tape sessions (instant recovery enabled). This information includes:
  - The backup session ID
  - Information on when the backup session was performed (time stamp)
  - Name of the backup specification used in the backup session
  - LUNs and World Wide Names (WWNs) of the disk arrays that were used in the backup session
- CRC check information calculated during the ZDB-to-disk or ZDB-to-disk+tape session.
- A list of snapshots that you do not want to be used by Data Protector. This information is kept in the VA LUN exclude file of the VADB. For information on the VA LUN exclude file, refer to “VA LUN Exclude File” on page 14.
- Information on the VA password, if LUN security is used.
- Information on the Storage Area Manager password, if the HP OpenView Storage Allocater software is installed (in the SAN environment only).

Information on ZDB-to-disk and ZDB-to-disk+tape sessions and the CRC check information is written in the VADB whenever a replica is created, and is deleted from the VADB whenever a replica is deleted.

Information on ZDB-to-tape sessions and some information on ZDB-to-disk+tape sessions is stored also in the Data Protector internal database (IDB).

The VADB resides on the Cell Manager in the following directory:

- On UNIX: `/var/opt/omni/server/db40/vadb`
- On Windows: `<Data_Protector_home>\db40\vadb`

The omnidbva command is used for all tasks described in this section. For more information on the omnidbva command, refer to its man page.

The omnidbva command can be run from any client within the Data Protector cell that has the User Interface component installed.

Querying the VADB

The following is the syntax of the omnidbva command when used to query the VADB for information on ZDB-to-disk and ZDB-to-disk+tape sessions (instant recovery enabled):

Syntax

```
omnidbva {-session [<session_id>] | -lun [<LUN#>]}
```

Listing all Available Backup Sessions in the VADB

To get information on all available ZDB-to-disk and ZDB-to-disk+tape sessions within the VADB, run the following command:

```
omnidbva -session [<session_id>]
```

The command displays information on the session ID, the backup specification name, and when the backup session was performed.

If *<session_id>* is specified, information on LUNs that were involved in the session, and WWN(s) of the disk array(s) involved in the session is displayed.

Listing the Information on LUNs in the VADB

In the VADB, information on LUNs that were used in the ZDB-to-disk and ZDB-to-disk+tape sessions is stored. To get a list of these LUNs, together with the information on the backup sessions (the sessions ID, time of the session, the backup specification name) that used them, run the following command:

```
omnidbva -lun [<LUN#>]
```

If *<LUN#>* is specified, only information on this specified LUN is displayed.

Checking the Consistency of the VADB

To perform a consistency check of the VADB and fix invalid entries if necessary, run the following command:

```
omnidbva -dbcheck [-force]
```

If the `-force` option is specified, the VADB will be checked/fixed without any interactive prompts.

VA LUN Exclude File

It is possible to disable Data Protector from using certain target volumes. Thus, it is possible to reserve certain target volumes, identified by their source volume (parent LUN), for purposes other than Data Protector backup and restore. A Data Protector session is aborted if the parent LUN involved in the session is listed in the VA LUN exclude file.

Manipulating the VA LUN Exclude File

The following is the syntax of the `omnidbva` command when used to manipulate the VA LUN exclude file:

Syntax

```
omnidbva -exclude {-put <filename> | -get <filename> |
-check <VA_wwn> <LUN> | -init | -delete}
```

Setting the VA LUN Exclude File

To create and set the VA LUN exclude file or to edit it, perform the following steps. To edit an existing VA LUN exclude file, skip step 1.

1. To create a template VA LUN exclude file or to overwrite an old one with the template VA LUN exclude file, use the following command:

   ```
   omnidbva -exclude -init
   ```

2. To get the file for editing, use the following command:

   ```
   omnidbva -exclude -get <filename>
   ```

   where *<filename>* is a full pathname of the file that you want to edit.

   The command reads the VA LUN exclude file from the Cell Manager and saves it as *<filename>*.

The following are the syntax and an example of the VA LUN exclude file:

VA LUN Exclude File - Template

```
#
# HP OpenView Storage Data Protector A.05.50
#
# HP StorageWorks Disk Array VA LUN Exclude File
#
# Syntax:
# [<VA wwn1>]
# <LUN>
# <LUN1>, <LUN2>, <LUN3>
# <LUN4>-<LUN5>
# [<VA wwn2>]
# ...
#
# <VA wwn> - Disk Array World Wide Name
# <LUN> - LUN number in decimal
#
# Example:
# [50060B000009295D]
# 1, 5, 10-20
# 123
# 125-220
#
#
#
# End of file
```

3. Edit the file and save it when you are done editing.
4. To copy the file to the original place, use the following command:

   ```
   omnidbva -exclude -put <filename>
   ```

The command reads the contents of the file, checks its syntax, and if the syntax is correct, copies the file to its location on the Cell Manager.

Identifying the Excluded VA LUNs

The following command checks whether a certain LUN, identified by its disk array WWN (*<VA_wwn>*) and a LUN number is specified in the VA LUN exclude file on the Cell Manager:

```
omnidbva -exclude -check <VA_wwn> <LUN#>
```

If the queried LUN is specified in the VA LUN exclude file, the command returns: `YES!`

If the queried LUN is not specified in the VA LUN exclude file, the command returns: `NO!`

Resetting the VA LUN Exclude File

The following command overwrites the current VA LUN exclude file on the Cell Manager with the template file:

```
omnidbva -exclude -init
```

Deleting the Contents of the VA LUN Exclude File

The following command empties the contents of the VA LUN exclude file on the Cell Manager:

```
omnidbva -exclude -delete
```

LUN Security

HP StorageWorks Secure Manager Virtual Array lets you set LUN permissions within VAs to protect your most critical data. It guards against LUNs being used or deleted by unauthorized servers or users.

IMPORTANT

If Secure Manager is enabled for the VA concerned, specify the password for this VA using the `omnidbva` command.

Whenever you want to create snapshots on VA, Data Protector needs to provide the correct password to Secure Manager; otherwise, the snapshot creation fails during the backup session.

Data Protector stores the password in the VADB and provides it to the Secure Manager whenever a snapshot technology on VA is used.

Activating LUN Security

To use LUN security, perform the following steps:

1. Using HP StorageWorks Command View SDM, activate Secure Manager on VA.
2. Provide the VA password to Data Protector using the following command:

   `omnidbva -vapasswd <VA_wwn> <password>`

   where `<VA_wwn>` is the VA node World Wide Name, and `<password>` is the password that you have specified during the activation of Secure Manager.

NOTE

You can find the node World Wide Name of the VA using Command View SDM. On HP-UX, you can also use a combination of the `tdlist` and `fcmsutil` commands. For information on the `fcmsutil` command, refer to the `fcmsutil` man page.

3. Select the Data Protector backup option `Integrate with VA LUN security` every time when creating a backup specification for the VA involved.

HP OpenView Storage Area Manager Password

If your VA is part of a SAN environment in which HP OpenView Storage Allocater is installed, provide the HP OpenView Storage Area Manager (SAM) password.

IMPORTANT

If you do not provide the password, the snapshot creation fails during the backup session.

To provide the password, perform the following:

- Use the `omnidbva` command as follows:

  `omnidbva -sampasswd <SAM server ID> <user> <password>`

  where

<SAM server ID> is actually the MANAGEMENT_SERVER_UID number that resides in the `PerProp` file in the *<SAM_home>*\hostagent directory.

Example of the `PerProp` File

```
UniqueID = d7cb00304c7d8347:49ba38:f29d31b081:-8000

HERMES = DOMAIN

MANAGEMENT_SERVER_UID =
d7cb00304c7d8347:7a84e4:f23ba7d999:-8000
```

Note that this number *is not* the World Wide Name of your VA.

- To run HP StorageWorks Command View SDM with Storage Allocater installed, also assign LUN 0 to the host that is running the Command View SDM software. Note that LUN 0 is used as a command device and not for storing data.

Deleting the Contents of the VADB

Deleting a Backup Session from the VADB

To delete all VADB related information on a specific ZDB-to-disk or ZDB-to-disk+tape session (identified by its session ID) and all snapshots for the replica, run the following command:

```
omnidbva -delete <session_id> [-force]
```

If the `-force` option is specified, the command will delete entries related to the backup session without any consistency checks.

IMPORTANT

As a consequence, it is not possible to perform instant recovery from the deleted ZDB-to-disk or ZDB-to-disk+tape session.

Deleting All Entries from the VADB

To reset all entries in the VADB, except the contents of the VA LUN exclude file, run the following command:

```
omnidbva -init [-force]
```

If the `-force` option is specified, the VADB will be reset without asking you for a confirmation.

IMPORTANT

The `omnidbva -init` command removes all entries from the VADB, including information on ZDB-to-disk and ZDB-to-disk+tape sessions, the LUN security, and Storage Area Manager passwords.

Although the command does not delete the snapshots for the replica, it is not possible to perform instant recovery from any ZDB-to-disk or ZDB-to-disk+tape session, because the information on these replicas was deleted.

2 Backup

In This Chapter

This chapter describes how to configure a filesystem or disk image backup of your data residing on HP StorageWorks Virtual Array (VA). The sections describe steps for configuring a ZDB using the Data Protector Graphical User Interface.

This chapter includes the following sections:

"Backup Process" on page 23

"Configuring a Backup Specification" on page 29

"Backup Options" on page 36

"Troubleshooting" on page 41

Backup Process

If you are not acquainted with the general ZDB and instant recovery concepts, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

ZDB Types

Three types of ZDB sessions are possible using the HP StorageWorks Virtual Array (VA) integration:

- **ZDB to disk**

  With this type of backup, snapshots are created, and data is kept on a disk array until reused. Data from the replica (all target volumes that are created during one backup session) is not moved to backup media (for example, tape media). A replica created in a ZDB-to-disk session can be used for instant recovery and is part of the replica set.

  Such a backup session is performed when the `Track the replica for instant recovery` backup option is selected when creating a backup specification, and the `To disk` option is selected when running or scheduling a backup.

- **ZDB to tape**

  With this type of backup, snapshots are created, and data from the replica is moved to backup media.

  If the backup option `Keep the replica after the backup` is selected, the replica remains on a disk array until reused, but cannot be used for instant recovery and is not part of the replica set. It can be reused in another ZDB-to-tape session using the same backup specification, but with the `Use an existing replica` option selected.

  If the backup option `Keep the replica after the backup` is not selected, the replica is deleted after the backup.

  Such a backup session is performed when the `Track the replica for instant recovery` backup option *is not* selected when creating a backup specification.

- **ZDB to disk+tape**

  With this type of backup, snapshots are created, and data is kept on a disk array until reused and also moved to backup media. The replica created in a ZDB-to-disk+tape session can be used for instant recovery and is part of the replica set.

  Such a backup session is performed when the `Track the replica for instant recovery` backup option is selected when creating a backup specification, and the `To disk+tape` option is selected when running or scheduling a backup.

IMPORTANT

Before creating a backup specification, consider all limitations regarding the Data Protector VA integration. For information, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

Snapshot Type

Snapshots on VA are copy-on-write snapshots with the pre-allocation of disk space. For more information on this type of snapshot, refer to *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

Replica Creation

A new replica is created in the following circumstances:

- A ZDB-to-tape session is performed in which the `Use an existing replica` option *is not* selected. Such a replica is not part of the replica set.
- A ZDB-to-disk or ZDB-to-disk+tape session is performed (the `Track the replica for instant recovery` option is selected), but the specified `Number of replicas rotated` is not reached yet. Such a replica becomes part of the replica set.

Replica Reuse

The oldest replica in the replica set is reused (a new one *is not* created) in the following circumstances:

- A ZDB-to-tape session is performed in which the `Use an existing replica` option *is* selected.

NOTE This is only possible if a disk array already contains a replica that is not part of the replica set. The replica has either been pre-configured, or has been left on a disk array from one of the previous ZDB-to-tape sessions using the same backup specification with the backup option `Keep the replica after the backup` selected.

- A ZDB-to-disk or ZDB-to-disk+tape session is performed and the specified `Number of replicas rotated` is reached.

After a Replica Creation/Reuse

If the backup option `Keep the replica after the backup` was not selected, the replica and, therefore, all snapshots created during the backup session are deleted. Otherwise, the replica is left on a storage system.

VA Backup Flows

The following are the snapshot backup flows on VA depending on the selection of backup options. For detailed information on these options, refer to “Backup Options” on page 36.

Figure 2-1 ZDB-to-Disk Session Flow

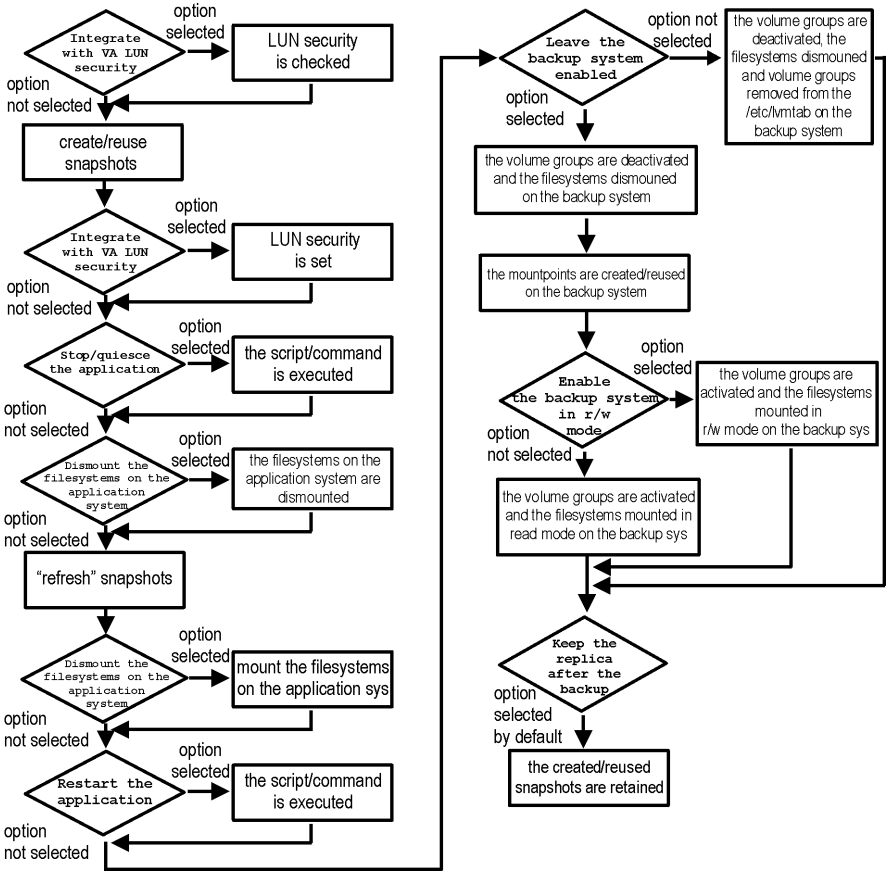

— The *create/reuse snapshots* phase represents an operation in which CommandView SDM allocates the space for snapshots. The actual writing operation is performed during the *"refresh" snapshots* phase. Since the subsequent write operations are fast because they do not need to allocate space, the time frame in which the application is in backup mode (online backup) or down (offline backup) is minimal.

— The `Enable the backup system in read/write mode` option is related to HP-UX systems only, since on HP-UX systems, filesystems are normally mounted in read-only mode; while on Windows systems, filesystems are always mounted in read/write mode.

Figure 2-2 ZDB-to-Tape and ZDB-to-Disk+Tape Session Flow

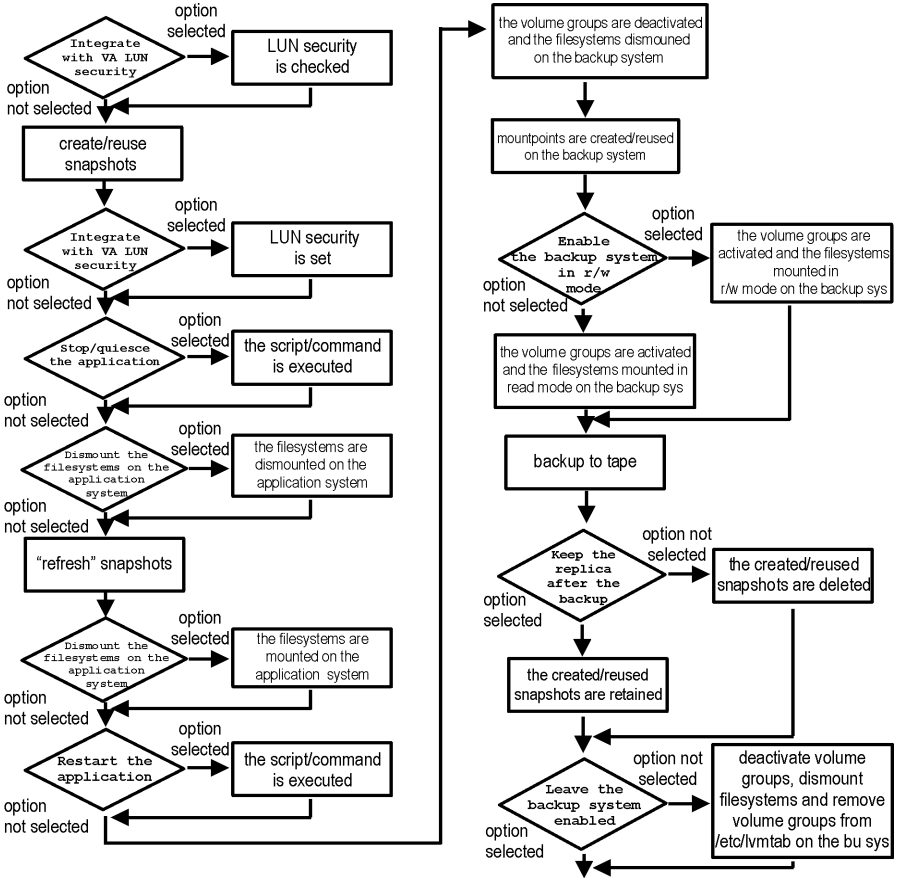

— The *create/reuse snapshots* phase represents an operation in which CommandView SDM allocates the space for snapshot creation. The actual writing operation is performed during the *"refresh" snapshots* phase. Since the subsequent write operations are fast because they do not need to allocate space, the time frame in which the application is in backup mode (online backup) or down (offline backup) is minimal.

— The `Enable the backup system in read/write mode` option is related to HP-UX systems only, since on HP-UX systems, filesystems are normally mounted in read-only mode, while on Windows systems, filesystems are always mounted in read/write mode.

— In the case of a ZDB-to-tape session, you can select the option `Keep the replica after the backup`. In the case of a ZDB-to-disk+tape session, this option is selected by default and cannot be deselected.

Configuring a Backup Specification

IMPORTANT If the backup system in your VA environment has the HP StorageWorks AutoPath software installed, you need to pre-configure the snapshots representing a replica to enable AutoPath to function properly. For the steps, refer to “Configuring VA with HP StorageWorks AutoPath Installed” on page 9. After the steps are performed, refer to “Running and Scheduling a ZDB Session” on page A-3.

Use the Data Protector GUI to create a filesystem or disk image snapshot backup specification for use with the VA integration.

1. In the Context List, select `Backup`.
2. In the Scoping Pane, expand `Backup` and then `Backup Specifications`. Right-click `Filesystem` (for both filesystem or disk image backup) and click `Add Backup`.

   The `Create New Backup` dialog box is displayed. See Figure 2-3 on page 30.

   In the `Filesystem` box, select the `Blank Filesystem Backup` template. For more information on the templates, refer to the *HP OpenView Storage Data Protector Administrator’s Guide*.

   Select the `Snapshot backup` option in the `Backup type` drop-down list and the `HP StorageWorks VA` in the `Sub type` drop-down list. For a description of other options, press **F1**.

   Click `OK`.

Figure 2-3 Create New Backup Dialog Box

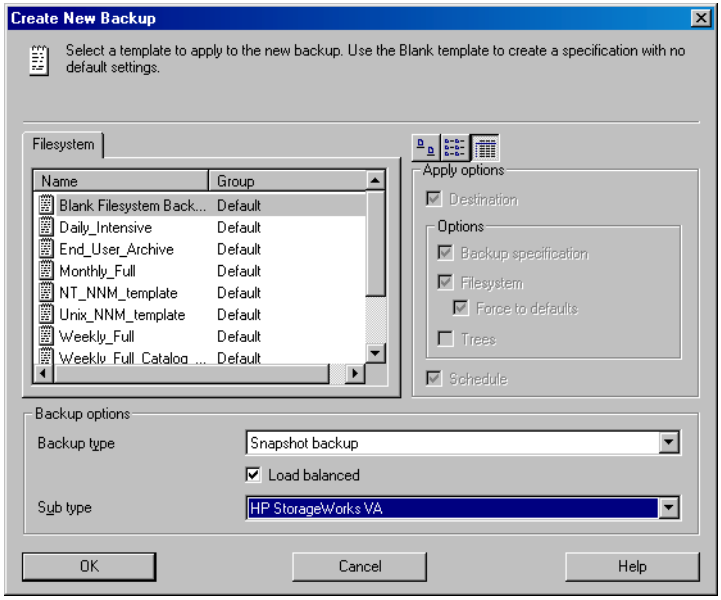

3. Under `Client systems` select the application and backup systems in the `Application system` and the `Backup system` drop-down lists.

   If Secure Manager is activated on your disk array, select the `Integrate with VA LUN security` option under `Replica management options`. Note that you need to have a password configured correctly, otherwise the backup session will fail. For information on LUN security and setting a password, see “LUN Security” on page 16.

   Specify other options in the following way:

For ZDB to Disk and ZDB to Disk+Tape

   Under `Instant recovery options`, leave the `Track the replica for instant recovery` option selected, and specify the `Number of replicas rotated`. The maximum number is 1024. See Figure 2-4 on page 31.

   For detailed information on these and other options, refer to “Backup Options” on page 36.

   Click `Next`.

IMPORTANT

You specify whether you want to perform a ZDB-to-disk or a ZDB-to-disk+tape session using the `Split mirror/snapshot backup` option when you run a backup or when you schedule a backup specification. Refer to “Running and Scheduling a ZDB Session” on page A-3 for more information.

Figure 2-4 VA Backup Options Required For ZDB-to-Disk or ZDB-to-Disk+Tape Session

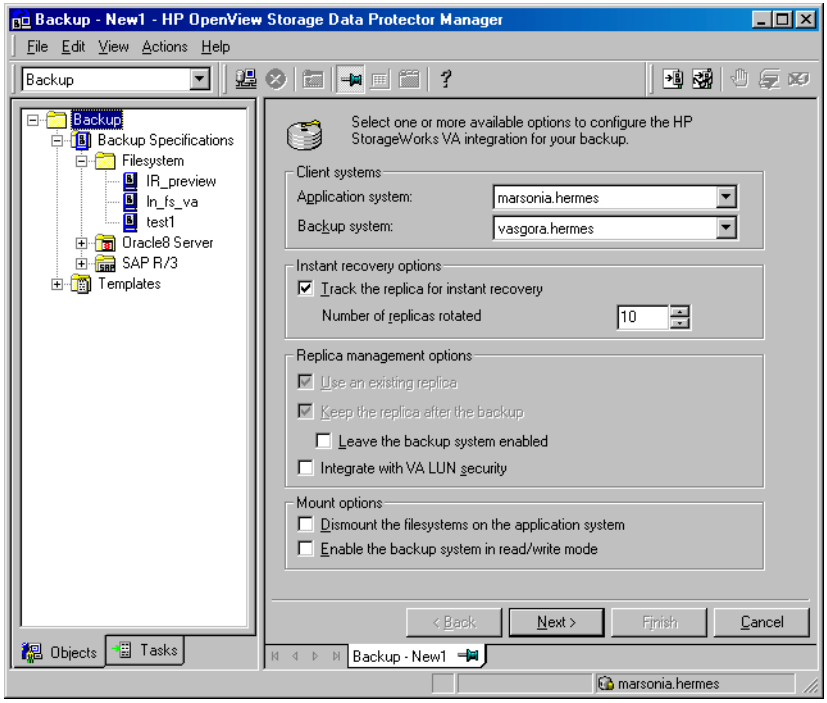

For ZDB to Tape

Deselect the `Track the replica for instant recovery` option. If you do not want the replica to be deleted after the backup session, select the `Keep the replica after the backup` option. See Figure 2-5 on page 32.

IMPORTANT

If you select the `Use an existing replica` option, ensure that there is already a replica for the same source volumes left on a disk array; otherwise, the backup session will fail.

For detailed information on these and other options, refer to “Backup Options” on page 36.

Click `Next`.

Figure 2-5 VA Backup Options Required For ZDB-to-Tape Session

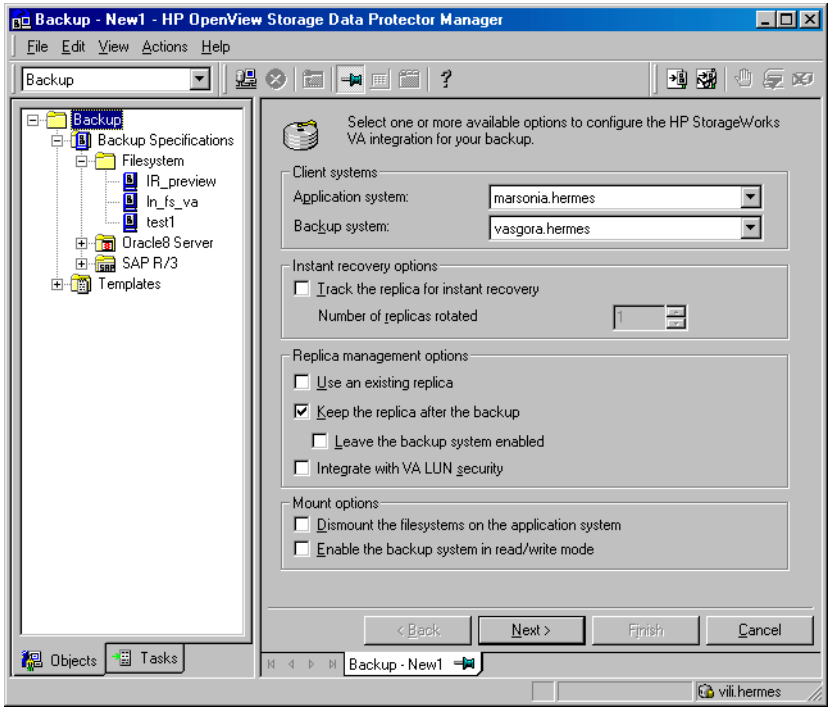

4. Depending on whether you perform a filesystem or disk image backup, proceed as follows:

Filesystem Backup

If you are configuring a filesystem backup, expand the application systems that contain the objects that you want to back up and then select what you want to back up.

IMPORTANT

On UNIX, if you intend to perform instant recovery, *select all filesystems inside the volume group* to be backed up. Otherwise, instant recovery will not be possible using the Data Protector GUI or (if you perform instant recovery using the Data Protector CLI) data can be corrupted.

Click `Next`.

Disk Image Backup

If you are configuring a disk image backup, click `Next`.

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

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

   For detailed information on the object mirror functionality, see *HP OpenView Storage Data Protector Administrator's Guide*.

NOTE

Object mirroring is not supported for ZDB to disk.

Click `Next`.

6. In the `Backup Specification Options` group box, click the `Advanced` tab and then `HP StorageWorks VA` to open the VA backup options.

   Here, you can specify the `Application options`, and modify all other options, except the `Application system` and `Backup system` options. For information on the VA backup options, refer to "Backup Options" on page 36.

   Click `OK`.

   For information on `Filesystem Options`, press **F1**.

Figure 2-6 **VA Backup Options**

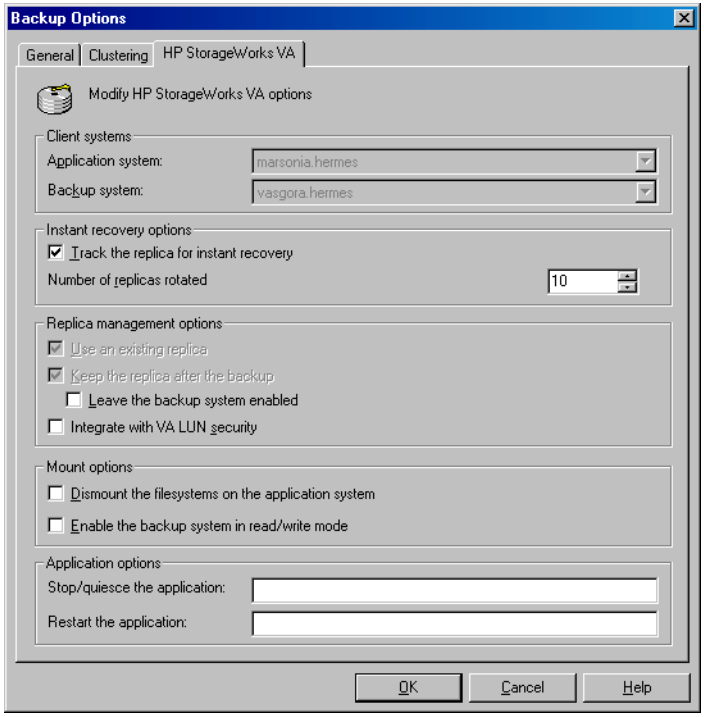

7. Follow the backup wizard to open the Data Protector scheduler (for information on scheduler, press **F1** or see "Running and Scheduling a ZDB Session" on page A-3) and then the backup summary page, where a summary of the backup specification is given.
8. Depending on whether you perform a filesystem or disk image backup, proceed as follows:

Filesystem Backup

If you are configuring a filesystem backup, click `Next`.

Disk Image Backup

For a disk image backup, proceed as follows:

a. Click `Manual add` to add the disk image objects you want to back up.

b. Select `Disk image object` and click `Next`.

c. Select the client to be backed up and click `Next`.

d. Follow the wizard to specify the `General Object Options` and the `Advanced Object Options`. For more information on these options, press **F1**.

e. In the `Disk Image Object Options` window, specify the disk image sections you want to back up.

On UNIX Systems

To specify a rawdisk section, use the following format:

`/dev/rdsk/<filename>`, for example: `/dev/rdsk/c2t0d0`

To specify a raw logical volume section, use the following format:

`/dev/vg<number>/rlvol<number>`, for example: `/dev/vg01/rlvol1`

IMPORTANT

If you intend to perform instant recovery, *specify all raw logical volumes inside the volume group* to be backed up. Otherwise, instant recovery will not be possible using the Data Protector GUI or (if you perform instant recovery using the Data Protector CLI) data can be corrupted.

On Windows Systems

Use the following format:

`\\.\PHYSICALDRIVE#`,

where # is the current number of the disk you want to back up.

For example: `\\.\PHYSICALDRIVE3`

For information on how to find the current numbers of the disks (physical drive numbers) you want to back up, refer to the online Help index keyword “disk image backups”.

f. Click `Finish` and then `Next`.

9. Save your backup specification. For more information on starting and scheduling a ZDB session, refer to “Running and Scheduling a ZDB Session” on page A-3.

Backup Options

The following tables describe the VA backup options. Refer to “VA and EVA Integrations” on page A-37 to help you understand these options.

Table 2-1 VA Client Systems Options

`Application system`	Specify the application system on which the application runs (for example, an Oracle database). In a cluster environment, specify the virtual server hostname (rather than the physical node hostname).
`Backup system`	Specify the backup system on which your data is to be backed up. In a cluster environment, specify the virtual server hostname (rather than the physical node hostname).

Table 2-2 VA Instant Recovery Options

`Track the replica for instant recovery`	Select 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. If this option is selected, you should also set the `Number of replicas rotated` parameter. Note that when this option is selected, the options `Use an existing replica` and `Keep the replica after the backup` are automatically selected. By default, this option is selected.

Table 2-2 VA Instant Recovery Options

`Number of replicas rotated`	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 reuses the oldest replica. Note that this option sets the number of replicas in the replica set for a backup specification. You need to specify this number if you have selected the `Track the replica for instant recovery` option. By default, this number is set to 1. The maximum is 1024.

Table 2-3 VA Replica Management Options

`Use an existing replica`	By default, this option is automatically selected (and cannot be deselected) if the `Track the replica for instant recovery` option is selected. If configuring a ZDB to tape, select this option to reuse an existing replica. Data Protector can reuse a replica only if the following condition is met; otherwise, the backup session will fail: 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.
`Keep the replica after the backup`	By default, this option is automatically selected (and cannot be deselected) if the `Track the replica for instant recovery` option is selected. If configuring a ZDB to tape, select 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 `Use an existing replica` selected. If this option is not selected, the replica is deleted after the backup session.

Table 2-3 VA Replica Management Options

`Leave the backup system enabled`	By default, Data Protector dismounts the filesystems (all platforms), deactives the volume groups (HP-UX systems) and removes the volume groups from `/etc/lvmtab` (HP-UX systems) on the backup system after each backup. If this option is selected, the filesystems remain mounted (all platforms), volume groups remain activated (HP-UX systems) after the backup and volume groups are not removed from `/etc/lvmtab` (HP-UX systems) after the backup. Thus, you can use the backup system for some data warehouse activity afterwards, *but not for instant recovery*. This option is available only if the `Keep the replica after the backup` option is selected. By default, this option is not selected.
`Integrate with VA LUN security`	Specify this option to apply the LUN security to the child LUNs (target volumes or snapshots) that the integration creates. *If Secure Manager is activated on VA, specify this option and configure passwords correctly; otherwise, the backup sessions will fail.* For more information, refer to “LUN Security” on page 16. By default, this option is not selected.

Table 2-4 VA Mount Options

`Dismount the filesystems on the application system`	Specify this option if you want the filesystem on the application system to be dismounted before a snapshot is created and remounted after a snapshot is created. A filesystem does not have the stop I/O functionality to flush the data from the filesystem cache to the disk and stop the I/O for the time of the snapshot. This option can be used to ensure that the data on the filesystem is consistent. If an integrated application (for example, Oracle or SAP R/3) runs on the filesystem, it controls the I/O to the disk. In this case, it is not necessary to dismount the filesystem before the snapshot creation. By default, this option is not selected.

Table 2-4 VA Mount Options

`Enable the backup system in read/write mode`	This option is related to HP-UX systems only, since on Windows systems filesystems are always mounted in read/write mode. Specify this option to have read/write access to the volume groups and filesystems on the backup system. For backup, it is sufficient to activate the backup system volume groups and filesystem in read-only mode. However, to perform other tasks after the backup has been performed, this might not be sufficient. By default, this option is not selected.

Table 2-5 VA Application Options

`Stop/quiesce the application`	Specify the optional `Stop/quiesce the application` command/script. Create this command in the `/opt/omni/lbin` (HP-UX systems) or in the *<Data_Protector_home>*`\bin` (Windows systems) directory on the application system and specify only the filename in the backup specification. This command is executed on the application system immediately before the snapshot creation is performed. It is mainly used to stop the application.
`Restart the application`	Specify the optional `Restart the application` command/script. Create this command in the `/opt/omni/lbin` (HP-UX systems) or in the *<Data_Protector_home>*`\bin` (Windows systems) directory on the application system and specify only the filename in the backup specification. This command is executed on the application system immediately after the snapshot creation is performed. It is mainly used to restart the application.

Application Options Specifics

Using the `Application options`, you can stop any application that is not integrated with the Data Protector VA integration on the application system before the snapshots are created, and restart this application after the snapshots are created.

If the execution of the `Stop/quiesce the application` command fails, the `Restart the application` command is also not executed, so a cleanup procedure should be implemented in the `Stop/quiesce the application` command. However, if the `ZDB_ALWAYS_POST_SCRIPT`

omnirc variable is set to 1, the Restart the application command is always executed if set. By default, this variable is set to 0. Refer to "ZDB Agents Omnirc Variables" on page A-23 for more information on the Data Protector VA omnirc variables.

TIP This set of options can also be used for any other operation on the application system before and after the snapshot creation.

NOTE The VA Application options are accessible from the Backup Specification Options group box (refer to step 6 on page 33), by clicking the Advanced tab and then the HP StorageWorks VA tab.

Troubleshooting

This section describes the most common problems you may encounter when using the Data Protector VA integration.

Before You Begin

- Ensure that the latest official Data Protector patches are installed. Refer to the *HP OpenView Storage Data Protector Software Release Notes* or http://www.openview.hp.com/products/datapro/spec_0001.html for an up-to-date list of supported versions, platforms, and other information.
- Examine system errors reported in `/var/opt/omni/log/debug.log` (HP-UX systems) or in `<Data_Protector_home>\log\debug.log` (Windows systems) on the application system and the backup system.

The following is a description of problems, and actions to be taken to resolve them.

Backup Problems

- **You cannot select the StorageWorks mode in the Data Protector user interface when attempting to create a backup specification.**

  Check whether the `HP StorageWorks VA Agent` integration module is installed on both the application and backup systems.

  Check the Data Protector `cell_info` file on the Data Protector Cell Manager, in `<Data_Protector_Home>\Config\server\cell\cell_info` (Windows Cell Manager) or in `/etc/opt/omni/server/cell/cell_info` (UNIX Cell Manager) to see if the needed component is installed.

  The file contents should look similar to the following example:

  ```
  -host "hpsap002.bbn.hp.com" -os "hp s800 hp-ux-11.00" -cc
  A.05.10 -da A.05.10 -ma A.05.10 -SNAPA A.05.10
  ```

- **The VA agent on the application system failed to dismount a filesystem.**

  Ensure that there is no file access. If a `Stop/quiesce the application` script has been specified, check that it stops all processes using the filesystem.

- **If you get the following message:**

Message

```
[Critical]

A system error occurred when starting the target script. The
system error code reported is 2 and the message resolves to '[2]
The system cannot find the file specified. '.
```

Description

The VA agent is not installed on the backup system.

Action

Install the `HP StorageWorks VA Agent` component on the backup system, too.

- **On Windows, snapshots cannot be mounted to the target location on the backup system.**

Message

```
[Major]

Filesystem \\.\Volume{9640da9a-6f36-11d7-bd7a-000347add7ba}
could not be mounted to C:\Program
Files\OmniBack\tmp\ranj.ipr.hermes\N\.

([145] The directory is not empty. ).
```

Description

When running a backup session with nested mountpoint objects, snapshots sometimes cannot be mounted to the target mountpoint location on the backup system. This happens when cleaning of the target mountpoint location has failed.

Action

On the backup system, manually clean the directory where filesystems are to be mounted. For example, if your target location is *<Data_Protector_home>*\tmp (the default location), make sure that the `tmp` directory is empty.

- **The backup session fails when attempting to use an existing replica.**

Action

If using a replica that has been created in the previous backup sessions, or has been pre-configured, check that the parent LUNs used in the current session are the same as they were in the previous

backup session. The backup specification should not change in the backup objects; otherwise, there will be no existing snapshots connected with the selected backup objects.

- **A ZDB-to-disk or ZDB-to-disk+tape session fails with the following message:**

Message

```
[Critical]

IR for VA database query has returned an unexpected number of
entries.
```

Descriptions

The following reasons can cause this problem:

— A backup specification has been modified in the `Source` property page, in which you select the data to be backed up.

Action

Check that the backup specification being used has not been modified (for example, that the backup objects have not been changed) since it started to be used for a replica set rotation. If a change in the backup objects is required, but the backup specification is already in use, it must not be modified; a new backup specification should be created.

— The VADB entries are inconsistent with the state on the VA storage system. A snapshot (child LUN) that is logged in the VADB and used for the replica set rotation is missing on the VA storage system.

Action

Check if a snapshot used for the replica set rotation exists on the VA storage system. If the snapshot does not exist on the VA storage system, delete all information related to the backup session in which this snapshot was created.

To delete all information about a backup session from the VADB, run the following command:

```
omnidbva -delete <session_id> [-force]
```

3 Restore

In This Chapter

This chapter describes how to configure and run a filesystem or disk image restore of your data backed up using the Data Protector Data Protector HP StorageWorks Virtual Array (VA) integration. The sections describe steps for performing a restore using the Data Protector Graphical User Interface and the Data Protector Command Line Interface.

This chapter includes the following sections:

"Overview" on page 47

"Restoring from Backup Media on LAN" on page 48

"Instant Recovery" on page 49

"Troubleshooting" on page 57

Overview

The data backed up using the Data Protector VA integration is stored:

- On a disk array; after a ZDB-to-disk or ZDB-to-disk+tape session.
- On the backup media; after a ZDB-to-tape or ZDB-to-disk+tape session.

The data stored on a disk array is restored using instant recovery; whereas, the data stored on the backup media is restored using the standard Data Protector restore procedure. This procedure restores the data from the backup media to the application system through a LAN.

Table 3-1 Restore Types Available After a Specific ZDB Session

	Instant Recovery	**Standard Restore**
ZDB to disk	Yes	N/A
ZDB to disk+tape	Yes	Yes
ZDB to tape	N/A	Yes

From the point of view of what is restored, the difference between the two types of restore is that with instant recovery, all of the data in a replica is restored (regardless of the backup objects selection during the backup); whereas, with the standard Data Protector restore, only the backup objects selected for the restore are restored.

As far as the speed of the two types of restore is concerned, instant recovery is faster than standard Data Protector restore, because data moves faster within a disk array than from backup media through a LAN.

Restoring from Backup Media on LAN

The data backed up in ZDB-to-tape and ZDB-to-disk+tape sessions can be restored from backup media on a LAN. The restore is performed in the same way as the standard Data Protector restore. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for more information on standard Data Protector restore and for instructions on how to perform it.

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. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for more information on configuring backup devices. Refer to the "Restoring Under Another Device" section of the same guide for more information on how to perform a restore using another device.

Instant Recovery

Using the Data Protector VA integration, only the data backed up in ZDB-to-disk and ZDB-to-disk+tape sessions can be restored in instant recovery sessions.

Instant recovery can be performed:

- Using the Data Protector GUI. For information, see “Instant Recovery Using the Data Protector GUI” on page 52.
- Using the Data Protector CLI. For information, see “Instant Recovery Using the Data Protector CLI” on page 54.

The number of replicas available for instant recovery for a ZDB to disk or ZDB to disk+tape backup specification is limited by the number specified with the `Number of replicas rotated` backup option. This option sets the size of the replica set. The replicas that are available for instant recovery can be listed by selecting `Instant Recovery` in the Context List of the Data Protector GUI, and then expanding the `Restore Sessions` folder. Alternatively, the Data Protector CLI can be used by running the `omnidbva -session` command.

A replica is accessible by the backup specification name and ID of the ZDB-to-disk and ZDB-to-disk+tape sessions that produced it; for example, `2003/07/23-10[IR]`. In the Results Area, also other data is provided; for example, time when the replica was created. In this way, you can identify the correct replica containing the data for restore.

Data Protector instant recovery moves data from a replica to the source volumes directly, without involving a backup device. At the beginning of the instant recovery session, the application system needs to be disabled. This includes dismounting the filesystems and deactivation of volume groups (UNIX only) by Data Protector. Before this can be done, the status of participating filesystems and volume groups (UNIX only) is checked, and only the mounted filesystems and activated volume groups (UNIX only) are dismounted and deactivated. At the end of the session, only those filesystems that were previously dismounted will be mounted, and only those volume groups that were previously deactivated will be activated (UNIX only).

IMPORTANT

After an instant recovery session, the configuration of the restored filesystems is the same as it was at the backup time. This means that the restored filesystems are mounted to the same mount points or drive letters as they were at the backup time. In case these mount points or drive letters have some other filesystems mounted, these filesystems are automatically dismounted before instant recovery and the restored filesystems are mounted after instant recovery.

Instant Recovery Process

When an instant recovery session is started, the following happens:

- The Restore Session Manager (RSM) on the Data Protector Cell Manager starts the Data Protector VA agent (SNAPA) on the application system and sends the restore session information. This sets up a communications link with the allocated backup system and the SNAPA is started on the backup system. The two agents then have a direct communications link.
- The SNAPA queries the VADB to find out the restore objects associated with the specified ZDB-to-disk or ZDB-to-disk+tape session.
- If the `Check the data configuration consistency` instant recovery option is selected, the following takes place:

  ✓ On HP-UX systems, the volume group configurations for the source volumes stored in the VADB is compared to the selected replica volume group configurations. If the items compared do not match, the session is aborted.

  When instant recovery is performed in an MC/ServiceGuard cluster to some other node than the one that was backed up, the current volume group configuration on the node to which instant recovery is to be performed is different from the volume group configuration stored in the VADB. In such a case the VADB volume group configuration data is replaced by the current volume group configuration data on the node to which instant recovery is to be performed and the session is not aborted.

✓ The CRC check information for the data in the source volumes stored in the VADB is compared to the CRC check information for the data in the selected replica. The CRC check ensures that data in the replica was not corrupted since the ZDB-to-disk or ZDB-to-disk+tape session that produced the replica. If the items compared do not match, the session is aborted.

- The application and backup systems are disabled by deactivating the participating volume groups (HP-UX systems) and dismounting the participating filesystems.
- A check is made for any snapshots of the source volume that do not belong to the replica set with the replica to be restored. If any such snapshots are found, the restore session is aborted.
- A check is then made to verify that no snapshots in the replica set with the replica to be restored are in use. This is done by checking that all snapshots in the replica set can be locked. If this check fails, the restore session is aborted.
- If both of the previous checks have been successful, the snapshots in the replica set with the replica to be restored are deleted one by one, together with their VADB entries, except for the replica that is to be restored. This is necessary, due to hardware limitations.
- When only the replica selected for instant recovery remains, the data is copied from the replica to the source volumes.
- If the `Keep the replica after the restore` instant recovery option was not selected, the snapshots in the replica that was used for restore are deleted together with its VADB entries.
- The application system is enabled by activating the participating volume groups (HP-UX systems) and mounting the participating filesystems.

Instant Recovery Procedure

To perform a filesystem or disk image instant recovery using the Data Protector GUI, follow the procedure described in “Instant Recovery Using the Data Protector GUI” on page 52. To perform a filesystem or disk image instant recovery using the Data Protector CLI, follow the procedure described in “Instant Recovery Using the Data Protector CLI” on page 54.

Prerequisite

If you have performed a disk image backup, manually dismount the disks to be restored before instant recovery, and re-mount them afterwards.

IMPORTANT

Before starting instant recovery, carefully consider all instant recovery related limitations and considerations as stated in the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

IMPORTANT

While an instant recovery session, do not perform a ZDB session that uses the same source volumes as those to which the data is restored.

Instant Recovery Using the Data Protector GUI

To perform instant recovery using the Data Protector GUI, proceed as follows:

1. In the Context List, select `Instant Recovery`.
2. Select the backup session (replica) from which you want to perform the restore. This can be done in two ways:
   - By the backup session ID and name:

     In the Scoping Pane, expand `Restore Sessions` and select the session from the list of *all* ZDB-to-disk or ZDB-to-disk+tape sessions.
   - By type of backup (filesystem, Oracle, SAP R/3,...) and backup session name and ID:

     a. In the Scoping Pane, expand `Restore Objects`.

     A list of backed up object types (Filesystem, Disk Image, SAP R/3, Microsoft SQL Server,...) is displayed.

     b. Expand the type of object you want to restore.

     A list of all available backup specification that were used in ZDB-to-disk or ZDB-to-disk+tape sessions for the selected type of object is displayed.

c. Expand the backup specification containing the objects that you want to restore.

   A list of all available ZDB-to-disk or ZDB-to-disk+tape sessions (identified by the session ID) for the selected backup specification is displayed.

Figure 3-1 Instant Recovery Session Selection

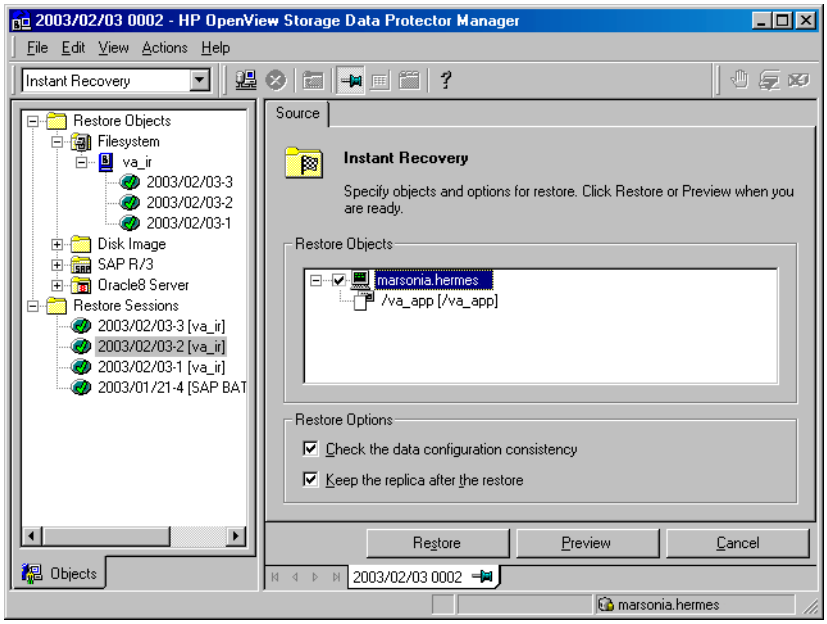

3. In the Scoping Pane, click the backup session you want to restore.

   The application system concerned and its mount points or drive letters (on Windows) representing the source volumes that were backed up during the selected ZDB-to-disk or ZDB-to-disk+tape session are displayed.

4. Check the selection box next to the application system to select the session for restore. Note that you cannot select sub-components. With instant recovery, only the complete session can be restored.

5. Select the instant recovery options that you want to use during the restore. See “Instant Recovery Options” on page 55 or press **F1** for more information on these options.

6. Click `Restore` to open the `Start instant recovery` dialog box.

7. Select `Start Restore Session` to start the instant recovery. It is recommended to test instant recovery to ensure it work properly. Click `OK`.

Instant Recovery Using the Data Protector CLI

The replica to be restored using the Data Protector CLI is identified by the ZDB-to-disk or ZDB-to-disk+tape session ID. To perform instant recovery using the Data Protector CLI, proceed as follows:

1. Get a list of all available ZDB-to-disk or ZDB-to-disk+tape sessions, identified by the session ID, using the following command:

   ```
   omnidbva -session
   ```

   From the output of the command, select the backup session you want to restore.

2. Execute the following command:

   ```
   omnir -host <application_system_name> -session
   <SessionID> -instant_restore [<INSTANT RECOVERY OPTIONS>]
   ```

   Where:

   - *<application_system_name>* is the hostname of the application system
   - *<SessionID>* is the backup session ID selected in the step 1 of this procedure.

   For *<INSTANT RECOVERY OPTIONS>* see Table 3-2 on page 55.

   This will start the instant recovery session.

Refer to the `omnidbva` and `omnir` man pages for more information on these commands.

Instant Recovery Options

Instant recovery options are set during the configuration of an instant recovery session.

Table 3-2 Instant Recovery Options

Data Protector GUI/CLI	Function
`Check the data configuration consistency/-check_config`	If this option is selected, the current volume group configuration of the volume groups involved in the instant recovery session is compared with the volume group configuration during the ZDB-to-disk or ZDB-to-disk+tape session kept in the VADB. If the volume group configuration has changed since the ZDB-to-disk or ZDB-to-disk+tape session, the restore is aborted. When instant recovery is performed in an MC/ServiceGuard cluster to some other node than the one that was backed up, the current volume group configuration on the node to which instant recovery is to be performed is different from the volume group configuration kept in the VADB. In such a case the VADB volume group configuration data is replaced by the current volume group configuration data on the node to which instant recovery is to be performed and the session is not aborted. When performing instant recovery in an MC/ServiceGuard cluster to some other node than the one that was backed up, select this option. The CRC check information for the data in the source volumes in the VADB is compared to the CRC check information for the data in the selected replica. If the items compared do not match, the session is aborted. Note that on Windows, there are fairly constant writes to drives, so it is not unusual for a CRC check to fail. By default, this option is selected.

Table 3-2 Instant Recovery Options

Data Protector GUI/CLI	Function
`Keep the replica after the restore/-keep_version`	If this option is selected, the replica from which the data was restored is left on the disk array after the restore. It is advisable to leave this option selected, in case of any problems during the restore. Even if the restore is successful, it is recommended to leave the replica in place for security until the next backup session is run. By default, this option is selected.

Instant Recovery in a Cluster

To perform instant recovery when an application or a filesystem is running in an MC/ServiceGuard or Microsoft Cluster Server on the application system, it is necessary to perform some additional steps. See “Instant Recovery in a Cluster” on page A-20 for the detailed procedure.

Troubleshooting

This section describes the most common problems you may encounter during the instant recovery when using the Data Protector VA integration.

Before You Begin

- Ensure that the latest official Data Protector patches are installed. Refer to the *HP OpenView Storage Data Protector Software Release Notes* or http://www.openview.hp.com/products/datapro/spec_0001.html for an up-to-date list of supported versions, platforms, and other information.
- Examine system errors reported in `/var/opt/omni/log/debug.log` (UNIX systems) or in `<Data_Protector_home>\log\debug.log` (Windows systems) on the application system and the backup system.

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

Instant Recovery Problems

- **Instant recovery session fails if the application system is in cluster.**

Message

```
[Critical]
Data consistency check failed!
Configuration of volume group <vg_name> has changed since the
last backup session!
```

Description

The problem occurs when the `Check data configuration consistency` instant recovery option is selected. It is probably caused by a failover to a secondary node, or the volume group configuration on the application system has changed.

Action

Make sure that the volume group configuration on the application system has not changed and/or deselect the `Check the data configuration consistency` option, and then restart the instant recovery session.

- **On Windows, instant recovery to a different cluster node fails.**

Messages

```
[Major]

Filesystem <volume_name> could not be dismounted from
<drive_letter> ([2] The system cannot find the file specified.).

[Critical]

Failed to disable the application system.

[Critical]

Failed to resolve objects for Instant Recovery.
```

Description

On Windows, the Data Protector automatic preparation of the application system cannot match the clustered volumes from one cluster node to the volumes on the other clustered node.

Action

Disable the automatic preparation of the application system as follows:

1. On the application system, enable the `omnirc` variable `ZDB_IR_MANUAL_AS_PREPARATION`. For information on the variable, see “ZDB Agents Omnirc Variables” on page A-23.
2. Manually dismount all the volumes on the application system that are to be restored.
3. Start the instant recovery session.
4. After instant recovery, manually mount the restored volumes.

II HP StorageWorks Enterprise Virtual Array

4 Configuration

In This Chapter

This chapter describes the procedure for configuring Data Protector HP StorageWorks Enterprise Virtual Array (EVA) integration. It also provides information on the ZDB database.

For a detailed description of the installation of the Data Protector Cell Manager and clients, refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

This chapter includes the following sections:

"Prerequisites and Limitations" on page 63

"Configuring the Integration" on page 67

"ZDB Database—EVADB or SMISDB" on page 70

Prerequisites and Limitations

Prerequisites

- You need a special license to use the Data Protector EVA integration. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*.
- The Data Protector EVA integration must be correctly installed. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*.
- The same operating system and its version must be installed on the application and the backup system.
- The following EVA components are required for this integration:
  - ✓ Command View EVA (CV EVA) v3.2 or above; in this case, the EVA SMI-S Agent will be used for controlling the EVA integration. If a lower version of CV EVA is installed, you'll have to use the EVA Agent (legacy).

    If you performed an upgrade to CV EVA v3.2, you need to upgrade from the EVA Agent (legacy) to the EVA SMI-S Agent. For the instructions on how to do it, refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*.
  - ✓ HP OpenView Storage Operations Manager (the HP StorageWorks SMI-S EVA provider part) starting with v1.1 to be able to use the EVA SMI-S Agent. If you plan to use the EVA Agent (legacy), you do not need to install this application.
  - ✓ HP StorageWorks Virtual Controller Software (VCS) installed on HSV Controllers. Refer to the VCS documentation.
  - ✓ A license for managing and controlling the EVA storage system.
  - ✓ SANworks Snapshot licenses installed.
- You should be familiar with the EVA concepts and procedures. Refer to the EVA-related documentation.
- You should be familiar with the basic ZDB and instant recovery concepts. Refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

- Refer to the *HP OpenView Storage Data Protector Software Release Notes* for information on:
  - ✓ general Data Protector limitations
  - ✓ supported platforms
  - ✓ supported integrations
  - ✓ supported backup topologies
  - ✓ supported connectivity topologies
  - ✓ supported cluster configurations (high availability support)

Management Appliance Limitation

- If you have multiple storage Management Appliances, only the Management Appliances that manage your EVA storage systems are supported, but not those that monitor them. Using the `omnidbeva` command, provide the login information only for these Management Appliances. Other Management Appliances should not be configured within the Data Protector cell. Note that this limitation applies only if you have the EVA Agent (legacy) installed.

VCS Limitations

- With VCS v2.x, the following is not supported:
  - — ZDB to disk and ZDB to disk+tape
  - — Instant recovery
  - — Snapclones
  - — Replica set rotation
- Only one type of target volume per source volume can exist on a disk array at the same time. For example, a snapclone of a source volume cannot be created if a vsnap or a standard snapshot of the same source volume already exists.
- The maximum of 7 snapshots (either vsnaps or standard snapshots) per source volume can exist.

NOTE

The number of snapclones per source volume is not limited.

- When a cloning process of a source volume is in progress, another snapclone of that source volume cannot be created.

- A replica cannot be reused if any snapclone from this replica has a snapshot attached.
- A replica cannot be reused for the purpose of the replica set rotation if any target volume from this replica is presented to some system.

Backup Limitations

- Only one type of target volume can be created during one ZDB session.

  Therefore, a backup session with the `Loose` (snapshot policy) backup option selected in the backup specification can be successful only when all existing target volumes of the source volumes involved in the session are of the same type (either standard snapshots, or vsnaps, or snapclones).
- For ZDB-to-disk and ZDB-to-disk+tape session, only snapclones can be used.
- Data Protector does not support creating a target volume from a snapclone. If a source volume to be backed up is found in the replica set, the backup session fails.
- Preview backup is not supported.
- Object copying and object mirroring is not supported for ZDB to disk.

Instant Recovery Limitations

- Instant recovery is possible only when snapclones were used during the backup.
- After an instant recovery session, the replica used in the session becomes the source volumes, therefore the physical location of the original application data changes.

  Additionally, if a replica to be recovered resides in other disk group than the source volumes, then this replica disk group becomes the disk group of the source volumes after instant recovery.
- After an instant recovery session, the replica restored is deleted from the replica set. Consequently, you cannot repeat the instant recovery session using the same session ID.
- Instant recovery from a ZDB-to-disk+tape session cannot be performed using the Data Protector GUI after exporting or overwriting the media used in the backup session. The backup media must not be exported or overwritten even after an object copy session. If the backup media have been exported or overwritten, perform

instant recovery using the Data Protector CLI. See “Instant Recovery Using the Data Protector CLI” on page 130 for more information on how to perform instant recovery using the Data Protector CLI.

- CRC check is not performed.

Other Limitations

- On all supported platforms (refer to the *HP OpenView Storage Data Protector Software Release Notes* for information on supported platforms), EVA supports systems with multiple FC HBA cards only when they are running HP StorageWorks Secure Path.
- Continuous Access software appliance is not supported.
- Virtual Controller Software v1.0 is not supported.
- If you upgraded the Command View EVA software to version 3.2 and, subsequently, the EVA VCS firmware to version 3.02x, you *must* upgrade the EVA Agent (legacy), used with previous versions of CV and VCS, to the EVA SMI-S Agent. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions.

Configuring the Integration

The following sections assume that you have chosen the desired backup configuration and connected EVA to the application system(s) and the backup system. The source volumes should be presented to the application system(s) and, if necessary, mounted as filesystems on the system(s). For information on the supported configurations and their descriptions, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

Using Command View EVA (CV EVA), create source volumes and present them to the application system.

Using HP-UX LVM or Windows Disk Management utility, configure the filesystems on the application system, if necessary, and mount them.

IMPORTANT

For ZDB-to-disk sessions, 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 configure a backup specification for a ZDB-to-disk session. For information on configuring a standalone device, refer to the online Help index keyword “standalone devices”.

CV EVA Login Information

Provide the login information for the CV EVA service running on the Management Appliance system. For instructions, see “Setting the Login Information for CV EVA” on page 72.

If you have multiple storage Management Appliances, provide the login information only for those Management Appliances that manage your EVA storage systems, and *not* for those that monitor them. This limitation applies only if you have the EVA Agent (legacy) installed.

If a failover from the active to the standby Management Appliance happens, proceed as follows:

- If the standby Management Appliance has the same hostname as the Management Appliance that failed, no action is needed.
- If the standby Management Appliance has a different hostname than the Management Appliance that failed, remove the Management Appliance that failed from the Data Protector configuration, and then

add the new Management Appliance to the Data Protector configuration. For information, see “Setting the Login Information for CV EVA” on page 72.

If your EVA storage system hardware configuration changes after providing the login information, perform a rescan operation. For instructions, see “Updating the Information on the EVA Hardware Configuration in the EVADB” on page 74. Note that you do not need to do the rescan if you use the EVA SMI-S Agent.

To allocate snapclones to a different disk group than the one used for the source volumes (original virtual disks), see “EVA Disk Group Pairs Configuration File” on page 75.

Automatic Configuration of Backup System

The EVA integration do not require any configuration steps, such as configuring the volume groups and filesystems on the backup system. This is done by Data Protector automatically when a ZDB session is started. Based on the volume group, filesystem, and mount point configuration on the application system, Data Protector creates the same volume group and filesystem structure on the backup system. During the ZDB-to-tape and ZDB-to-disk+tape sessions, Data Protector mounts these filesystems. In case of disk images, raw device files are used. For more information on the backup system mountpoint creation, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

Before the Automatic Preparation

Before the automatic preparation of the backup system is done, the EVA Agent (legacy) ensures that the host object representing the backup system is configured within the EVA storage system. The agent automatically detects the already configured backup system on the basis of the backup system FC HBAs. If it detects that the backup system is not configured yet, it automatically configures it.

Since EVA Agent (legacy) performs an extensive search over the EVA storage system to find the pre-configured host objects, the search can take a long time in large environments. Therefore, you can change the default Data Protector behavior by setting the `omnirc` variable `EVA_DISABLE_HBA_AUTODETECTION`. For more information on the variable, see “ZDB Agents Omnirc Variables” on page A-23.

If you set this variable, the search for the host object representing the backup system is based on the IP addresses. This type of search is faster, but it may require some pre-configuration steps in case the backup system cannot be found. In this case, manually configure the host object for the backup system through CV EVA and provide its name in the `EVA_HOSTNAMEALIASES omnirc` variable. For more information on the variable, see “ZDB Agents Omnirc Variables” on page A-23. Use this procedure also if the automatic autodetection of the backup system fails.

NOTE

The EVA SMI-S Agent does do the automatic configuration of the backup system. Therefore, before a ZDB session is started, make sure the backup system that will be involved in the session is already configured within the EVA storage system.

ZDB Database—EVADB or SMISDB

The **ZDB database** is in case of the EVA integration referred to as **EVADB** if you have the HP StorageWorks EVA Agent (legacy) installed, or as **SMISDB** if you have the HP StorageWorks EVA SMI-S Agent. The following information is kept in the EVADB/SMISDB:

- Information on the Management Appliance systems that are configured within the Data Protector cell. For each Management Appliance, the following data is stored:
  - Hostname of a Management Appliance system as recognized in the IP network
  - Port number on which CV EVA listens to requests.

    Note that the default number for CV EVA v3.0 that is used by the EVA agent (legacy) is 12301 and for CV EVA v3.2 that is used by the EVA SMI-S Agent is 5598.
  - User name and password (encoded) for logging in to the CV EVA
- Some information on the EVA hardware configuration
- Information on disk groups to be used for snapclone destination
- On UNIX, information on the volume group configuration of the source volumes
- Information on Data Protector EVA snapshot backup sessions (replicas) that are kept on disk array(s). This information includes:
  - The backup session ID
  - Information on when the backup session was performed
  - Name of the backup specification used in the backup session
  - Name, ID, and WWN of the target volume created in the backup session
  - Name and ID of the EVA storage system on which the target volume resides
  - Information on the target volume type (vsnap, standard snapshot, or snapclone)
  - ID of a source volume used in the backup session

— Information on whether the target volume can be used for instant recovery (IR flag)

— Information on whether the target volume should be deleted (purge flag)

— Names of the application and backup systems involved in the backup session

This information is written to the EVADB/SMISDB whenever a replica is created, and is deleted from the database whenever a replica is deleted.

The EVADB/SMISDB stores information on the backup sessions that have the backup option `Keep the replica after the backup` selected in the backup specification. ZDB-to-tape sessions without this backup option selected are deleted from the database after each such backup.

Information on ZDB-to-tape sessions and some information on ZDB-to-disk+tape sessions is stored also in the Data Protector internal database (IDB).

The EVADB Location

The EVADB resides on the Cell Manager in the following directory:

- On Windows:

  `<Data_Protector_home>\db40\evadb`

- On UNIX:

  `/var/opt/omni/server/db40/evadb`

The SMISDB Location

- On Windows:

  `<Data_Protector_home>\db40\smisdb`

- On UNIX:

  `/var/opt/omni/server/db40/smisdb`

The following commands are used for the tasks described in this section:

- `omnidbeva` for the EVA Agent (legacy)
- `omnidbsmis` for the EVA SMI-S Agent

For more information on these commands, refer to the corresponding man pages.

The omnidbeva and omnidbsmis commands can be run from any client within the Data Protector cell that has the User Interface component installed.

Setting the Login Information for CV EVA

Before you start ZDB sessions using the Data Protector EVA integration, you need to provide the login information for the CV EVA service running on the Management Appliance system.

EVA Agent (legacy) Limitation

If you have multiple storage Management Appliances, provide the login information only for those Management Appliances that manage your EVA storage systems. The Management Appliances that monitor your EVA storage systems are not supported and should not be configured within the Data Protector cell. Note that the EVA SMI-S Agent does not have this limitation.

Syntax

The following is the syntax of the omnidbeva and the omnidbsmis commands when used to set the login information for CV EVA:

omnidbeva

```
omnidbeva -empasswd {-list | -add <EM_hostname> [-port
<port_number>] | -remove <EM_hostname> | -update
<EM_hostname> [-port <port_number>]}
```

omnidbsmis

```
omnidbsmis -ompasswd -add <hostname> [-port <port_number>]
[-namespace <namespace>] [-user <username>] [-passwd
<password>] | -remove <hostname> [-port <port_number>]
[-namespace <namespace>] [-user <username>] | -list
[<hostname>]}
```

Adding a Management Appliance

To add a Management Appliance system to the EVADB and to define the login information for a relevant CV EVA v3.0 or lower, perform the following steps:

For CV EVA v3.0

1. omnidbeva -empasswd -add <*EM_hostname*> [-port <*port_number*>]

   where <*EM_hostname*> is the hostname of a Management Appliance system on which the relevant CV EVA is running.

   Use the option -port <*port_number*> to specify a listen port number if it is different from 12301 on CV EVA v3.0.

IMPORTANT For CV EVA v3.0, it is strongly recommended to use the default port number, which is 12301.

2. You are prompted to enter the user name and password, which Data Protector will use for logging in to CV EVA.

For CV EVA v3.2

To add a Management Appliance system to the SMISDB and to define the login information for a relevant CV EVA v3.2, run the following command:

```
omnidbsmis -ompasswd -add <hostname> [-port <port_number>]
[-namespace <namespace>] [-user <username>] [-passwd
<password>]
```

where *<hostname>* is the hostname of a Management Appliance system on which the relevant CV EVA is running.

The option `-port` *<port_number>* is used to specify a listen port number if it is different from the default one, which is 5988. The `-namespace` *<namespace>* option sets the namespace that contains the SMI-S EVA provider configuration for CV EVA. The default namespace is `root/eva`.

With the `-user` *<username>* option, you set the user of CV EVA v3.2. The default user is `administrator`. The `-passwd` *<password>* parameter is the password that will be used for logging in to CV EVA. If you did not set a password when running `omnidbsmis`, the command will ask for it interactively.

Removing a Management Appliance

To remove a Management Appliance system from a list of systems on which CV EVAs are running, run one of the following commands:

```
omnidbeva -empasswd -remove <EM_hostname>
```

if you have the EVA Agent (legacy), or:

```
omnidbsmis -ompasswd -remove <hostname> [-port
<port_number>] [-namespace <namespace>] [-user <username>]
```

if you have the EVA SMI-S Agent.

You can specify `[-port` *<port_number>*`]`, `[-namespace` *<namespace>*`]`, and/or `[-user` *<username>*`]` parameter(s) if you have more than one port/namespace/user configured on the same system. In this case, the command will remove only a specified value from the configuration.

Updating the Login Information

To update existing login information or to specify a different port number if you have the EVA Agent (legacy) installed, run the following command:

```
omnidbeva -empasswd -update <EM_hostname> [-port <port_number>]
```

You are prompted to enter the user name and password.

TIP

Use this command to specify a different port number after upgrading from HSV Element Manager v2.x to CV EVA v3.0.

NOTE

The `-update <EM_hostname>` parameter is not supported for the `omnidbsmis` command.

Getting a List of CV EVAs

To get a list of all Management Appliance systems that are configured within the Data Protector cell, run one of the following commands:

```
omnidbeva -empasswd -list
```

if you have the EVA Agent (legacy), or:

```
omnidbsmis -ompasswd -list [<hostname>]
```

if you have the EVA SMI-S Agent. The `<hostname>` value is optional: if you enter a name of the host, only the SMI-S EVA providers, configured for a specified host, will be displayed.

Updating the Information on the EVA Hardware Configuration in the EVADB

To successfully perform EVA snapshot backup sessions, Data Protector needs to obtain some information on the EVA hardware configuration. Normally, this information is automatically scanned after providing the login information for a specific Management Appliance, and is stored in the EVADB/SMISDB. However, you need to manually update the EVADB in the following cases:

- When you change Management Appliance for an EVA storage system
- When you change at least one of the HSV EVA controllers

Rescanning EVA Hardware Configuration

To rescan the EVA storage system hardware configuration and thus update the EVADB with the necessary information, run the following command:

```
omnidbeva -hwrescan [<EVA_name>]
```

If *<EVA_name>* is specified, the hardware rescan operation is performed only on the specified EVA storage system; otherwise, the rescanning is performed on all EVA storage systems that are managed by configured Management Appliance systems (systems that are logged in the EVADB).

NOTE

The `-hwrescan` option does not exist for the `omnidbsmis` command. The SMI-S Agent gets the hardware information during the execution time, so you do not need to manually rescan the EVA storage system hardware configuration.

EVA Disk Group Pairs Configuration File

It is possible to allocate snapclones to a different disk group than the one used for the source volumes (original virtual disks). Thus, you can influence the application performance, since different physical disks are used for the read and write operations on the source volumes and the replica. By defining the disk group pairs configuration file, you can also allocate a replica to low-performance disks.

Commands' Syntax

The following is the syntax of the `omnidbeva` and the `omnidbsmis` commands when used to set disk group pairs in the disk group pairs configuration file:

```
omnidbeva -dgrules {-put <filename> | -get <filename> |
-check <EVA_name> <DG_name> | -init}
```

```
omnidbsmis -dgrules {-put <filename> | -get <filename> |
-check <EVA_name> <DG_name> | -init}
```

Setting the Disk Group Pair Configuration File

To create and set the disk group pairs configuration file or edit it, perform one of the two procedures described below. Perform the first procedure that uses the `omnidbeva` command if you have the EVA Agent (legacy) installed. Carry out the second procedure that uses the `omnidbsmis` command if you have the EVA SMI-S Agent.

NOTE If you are editing an existing disk group pairs configuration file, skip step 1 in the procedure.

Using `omnidbeva`

1. To create a template disk group pairs configuration file or to overwrite an old one with the template, run the following command:

   ```
   omnidbeva -dgrules -init
   ```

2. To get the file for editing, run the following command:

   ```
   omnidbeva -dgrules -get <filename>
   ```

   where *<filename>* is a full path name of the file that you want to edit (for example, `c:\tmp\dgrules`).

   The command reads the disk group pairs configuration file from the EVADB and saves it as *<filename>*.

3. Edit the file and save it. Note that the order of defining disk group names is ignored. This means that if a source volume is found in "disk group 1", its snapclone will be created in "disk group 2", and vice versa. Note that a certain disk group can be a member of only one disk group pair.

4. To copy the file to the original place, run the following command:

   ```
   omnidbeva -dgrules -put <filename>
   ```

   The command reads the contents of the input file, checks its syntax, and uploads the file to the EVADB.

Using `omnidbsmis`

1. To create a template disk group pairs configuration file or to overwrite an old one with the template, run the following command:

   ```
   omnidbsmis -dgrules -init
   ```

2. To get the file for editing, run the following command:

   ```
   omnidbsmis -dgrules -get <filename>
   ```

   where *<filename>* is a full path name of the file that you want to edit (for example, `c:\tmp\dgrules`).

   The command reads the disk group pairs configuration file from the SMISDB and saves it as *<filename>*.

3. Edit the file and save it. Note that the order of defining disk group names is ignored. This means that if a source volume is found in “disk group 1”, its snapclone will be created in “disk group 2”, and vice versa. Note that a certain disk group can be a member of only one disk group pair.
4. To copy the file to the original place, run the following command:

   `omnidbsmis -dgrules -put <filename>`

   The command reads the contents of the input file, checks its syntax, and uploads the file to the SMISDB.

If You Upgraded the EVA Agent (legacy)

If you upgraded from the EVA Agent (legacy) to the EVA SMI-S Agent, you need to carry out additional steps described below in order to use the disk group pairs configuration file with the upgraded agent:

1. To get the disk group pairs configuration file, run the `omnidbeva -dgrules -get <filename>` command (where *<filename>* is a full path name of the file that you want to get). The command reads the disk group pairs configuration file from the EVADB and saves it as *<filename>*.
2. If necessary, edit the file and save it.
3. Run the `omnidbsmis -dgrules -put <filename>` command to upload the file to the SMISDB.

Note that instructions on how to perform the agent upgrade are described in the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

Disk Group Pairs Configuration File Template

The following is the disk group pairs configuration file template:

```
#
# HP OpenView Storage Data Protector A.05.50
#
# HP StorageWorks EVA disk group pairs configuration file
#
# Syntax:
# "<EVA box name 1>": "<disk group 1 name>", "<disk group 2 name>"
# "<EVA box name 2>": "<disk group 3 name>", "<disk group 4 name>"
#
```

```
# Example:
# "MyEVA1": "/Disk Groups/Working DG1", "/Disk Groups/Backup DG1"
# "MyEVA1": "/Disk Groups/Working DG2", "/Disk Groups/Backup DG2"
# "MyEVA2": "/Disk Groups/Working DG1", "/Disk Groups/Backup DG1"
#
#
#
# End of file
```

Identifying Disk Group Pairs

To get the information on the disk group that is in pair with a specified disk group, run one of the following commands:

`omnidbeva -dgrules -check <EVA_name> <DG_name>`

if you have the EVA Agent (legacy), or:

`omnidbsmis -dgrules -check <EVA_name> <DG_name>`

if you have the EVA SMI-S Agent.

The output of the command is the following:

```
Configured disk group pair:
        1st disk group name  :"<disk_group_name>"
        2nd disk group name  :"<disk_group_name>"
        EVA name             :"<EVA_name>"
```

If there is no rule for the specified disk group, the first and the second disk groups are the same.

Resetting the Disk Group Pairs Configuration File

The following commands overwrite the current disk group pairs configuration file in the EVADB/SMISDB with the template file:

`omnidbeva -dgrules -init`

if you have the EVA Agent (legacy), or:

`omnidbsmis -dgrules -init`

if you have the EVA SMI-S Agent.

Querying the EVADB/SMISDB

The following is the syntax of the `omnidbeva` and the `omnidbsmis` commands when used to query the EVADB/SMISDB for information on snapshot backup sessions and created target volumes:

omnidbeva

```
omnidbeva -list {-session [-ir] | -datalist | -snapshot [-ir] | -purge}
```

```
omnidbeva -show {-session <session_ID> | -datalist <datalist_name> | -snapshot <virtual disk_ID> <EVA_ID>}
```

omnidbsmis

```
omnidbsmis [-list] {-session [-ir] | -datalist}
```

```
omnidbsmis -list -purge
```

```
omnidbsmis [-show] {-session <session_ID> | -datalist <datalist_name>}
```

NOTE

The `-list` or `-show` parameters are optional and are kept because of the backward compatibility with `omnidbeva`. You will get the same output if you run the `omnidbsmis` command without these parameters, for example:

`omnidbsmis -session` to get a list of available backup sessions, or:

`omnidbsmis -session <session_ID>` to get the information on a specified backup session.

Listing all Available Backup Sessions in the EVADB/SMISDB

To get a list of all available backup sessions in the EVADB/SMISDB, identified by the session ID and the name of the backup specification, run:

```
omnidbeva -list -session [-ir]
```

if you have the EVA Agent (legacy), or:

```
omnidbsmis [-list] -session [-ir]
```

If you have the EVA SMI-S Agent installed on your system.

If the `-ir` option is used, the command lists only those sessions that are marked for instant recovery (ZDB-to-disk and ZDB-to-disk+tape sessions).

Getting Detailed Information on a Specific Backup Session

To get expanded details of a particular backup session, identified by its session ID, run one of the following commands:

```
omnidbeva -show -session <session_ID>
```

if you have the EVA Agent (legacy), or:

```
omnidbsmis [-show] -session <session_ID>
```

if you have the EVA SMI-S Agent.

The output of the command is information on all target volumes that were created in the specified backup session. The following is displayed:

- Name, ID, and WWN of the target volume that was created in the backup session
- Name and ID of the EVA storage system on which the target volume was created
- Type of the target volume created
- ID of the source volume used in the backup session
- The backup session ID
- Time stamp of the target volume created
- IR flag
- Name of the backup specification used in the backup session
- Names of the application and backup systems involved in the backup session

Listing all Available Backup Specifications in the EVADB/SMISDB

To get a list of all available backup specifications in the EVADB/SMISDB, identified by the backup specification name, run one of the following commands:

```
omnidbeva -list -datalist
```

if you have the EVA Agent (legacy), or:

```
omnidbsmis [-list] -datalist
```

if you have the EVA SMI-S Agent.

The command lists all backup specifications that have the backup sessions logged in the EVADB/SMISDB.

Listing all Backup Sessions Based on a Specific Backup Specification

To get a list of all available backup sessions, identified by the session ID, that were produced within a specific backup specification, run one of the following commands:

```
omnidbeva -show -datalist <datalist_name>
```

if you have the EVA Agent (legacy), or:

```
omnidbsmis [-show] -datalist <datalist_name>
```

if you have the EVA SMI-S Agent,

where *<datalist_name>* is the name of the backup specification.

Listing all Available Target Volumes in the EVADB

To get a list of all target volumes created by Data Protector and stored in the EVADB, run the following command:

```
omnidbeva -list -snapshot [-ir]
```

The target volumes displayed are identified by the virtual disk ID and the EVA storage system ID.

If the `-ir` option is used, the command lists only those target volumes that are marked for instant recovery.

NOTE The `-snapshot` option is not supported for the `omnidbsmis` command.

Getting Detailed Information on a Specific Target Volume

To get detailed information on a specific target volume, run the following command:

`omnidbeva -show -snapshot <virtual_disk_ID> <EVA_ID>`

The command displays the following information about the specified target volume:

- Name, ID, and WWN of the target volume
- Name and ID of the EVA storage system on which the target volume resides
- Type of the target volume
- ID of its source volume
- ID of the backup session that created the target volume
- Time stamp of the target volume
- IR flag
- Name of the backup specification that created the target volume
- Names of the application and backup systems involved in the target volume creation

NOTE The `-snapshot <virtual_disk_ID> <EVA_ID>` option is not supported for the `omnidbsmis` command.

Listing Source and Target Volumes Marked for Purging

In case some target volumes (after a ZDB-to-tape session or when reusing a target volume) or source volumes (after instant recovery) cannot be deleted although they should be at the time of backup/restore, the relevant EVA agent stores information on such virtual disks in the EVADB/SMISDB and marks them for removal at later point in time

(purging). At the beginning of each backup session, the agent checks the EVADB/SMISDB for virtual disks with the purge flag and tries to delete them.

To list all source and target volumes (virtual disks) that are marked for purging in the EVADB/SMISDB, run one of the following commands:

```
omnidbeva -list -purge
```

if you have the EVA Agent (legacy), or:

```
omnidbsmis -list -purge
```

if you have the EVA SMI-S Agent.

The source and target volumes displayed are identified by the virtual disk (either source or target) ID and the EVA storage system ID.

TIP

You can also delete such virtual disks manually using CV EVA. After the deletion, it is recommended to synchronize the EVADB/SMISDB by running `omnidbeva -sync` or `omnidbsmis -sync`.

Purging the SMISDB

The purge operation is normally run at the beginning of each backup session. The EVA agent checks the EVADB/SMISDB for the virtual disks with the purge flag and, in case of finding them, attempts to delete these objects. However, you can also run the SMISDB purge operation at any other time by executing the following command:

```
omnidbsmis -purge [-force] [-host <hostname>]
```

The virtual disks that are presented to hosts do not get deleted when running `omnidbsmis -purge`; however, if you specify the `-force` option, the removal of the elements marked for deletion will be forced, and they will be removed.

By default, the purge operation is executed from the backup system that was involved in a ZDB session. With the `-host <hostname>` option, you can choose another system that has the EVA SMI-S Agent installed to start the purge operation from. Use this option if the original backup system is no longer available.

Deleting the Target Volumes Created in a Specific Backup Session

To delete the target volumes (replicas) created in a specific backup session and consequently all information on this session, run one of the following commands:

```
omnidbeva -delete -session <session_ID> [-preview]
```

if you have the EVA Agent (legacy), or:

```
omnidbsmis -delete -session <session_ID> [-preview] [-force]
[-host <hostname>]
```

if you have the EVA SMI-S Agent.

If the `-preview` option is used, the command lists the target volumes that will be deleted if the command is run; it does not delete anything.

If the `-force` option is set, the deletion of the target volumes will be forced even if they are presented to hosts.

The `-host <hostname>` option allows you to execute the deletion from another system.

IMPORTANT It is not possible to perform instant recovery from a deleted replica.

Deleting all Target Volumes in a Specific Replica Set

To delete all target volumes (replicas) created with the backup sessions based on a specific backup specification (replica set) and consequently all information on these sessions, run one of the following commands:

```
omnidbeva -delete -datalist <datalist_name> [-preview]
```

if you have the EVA Agent (legacy), or:

```
omnidbsmis -delete -datalist <datalist_name> [-preview]
[-force] [-host <hostname>]
```

if you have the EVA SMI-S Agent.

where `<datalist_name>` is the name of the backup specification that created the replica set you want to delete.

If the `-preview` option is used, the command displays IDs of the backup sessions that created the target volumes that will be deleted if the command is run; it does not delete anything.

If the `-force` option is used, the deletion of the target volumes will be forced even if they are presented to hosts.

The `-host <hostname>` option allows you to execute the deletion from another system.

IMPORTANT It is not possible to perform instant recovery from a deleted replica.

Synchronizing the EVADB/SMISDB

The sync operation runs the relevant EVA agent (on the backup system) that synchronizes the persistent data in the EVADB/SMISDB with the most current state of the EVA storage system. If a target volume is physically missing from the EVA storage system (for example, was deleted using the EVA native GUI/CLI), the whole backup session that created this target volume is deleted from the database. The check is performed for all replica sets.

To run the synchronizing operation, use one of the following commands:

```
omnidbeva -sync [-preview]
```

if you have the EVA Agent (legacy), or:

```
omnidbsmis -sync [-preview] [-force] [-host <hostname>]
```

if you have the EVA SMI-S Agent.

If the `-preview` option is used, the command lists the replicas, identified by the backup session ID, that will be deleted if the command is run; it does not delete anything.

If the `-force` option is set, the elements, marked for deletion, will be removed even if they are presented to hosts.

If the `-host <hostname>` option is specified, you can choose another location to start the SMISDB synchronizing operation.

5 Backup

In This Chapter

This chapter describes how to configure a filesystem or disk image backup of your data residing on HP StorageWorks Enterprise Virtual Array (EVA). The sections describe steps for configuring a ZDB using the Data Protector Graphical User Interface.

This chapter includes the following sections:

“Backup Process” on page 89

“Configuring a Backup Specification” on page 95

“Backup Options” on page 103

“Troubleshooting” on page 109

Backup Process

If you are not acquainted with the general ZDB and instant recovery concepts, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

ZDB Types

Three types of ZDB sessions are possible using the HP StorageWorks Enterprise Virtual Array (EVA) integration:

- **ZDB to disk**

  With this type of backup, target volumes (snapclones) are created, and data is kept on a disk array until reused. Data from the replica (all target volumes that are created during one backup session) is not moved to backup media (for example, tape media). A replica created in a ZDB-to-disk session is used for instant recovery and is part of the replica set.

  Such a backup session is performed when the `Track the replica for instant recovery` backup option is selected when creating a backup specification, and the `To disk` option is selected when running or scheduling a backup.

- **ZDB to tape**

  With this type of backup, target volumes (any type) are created, and data from the replica is moved to backup media.

  Such a backup session is performed when the `Track the replica for instant recovery` backup option *is not* selected when creating a backup specification.

  If the backup option `Keep the replica after the backup` is selected, the replica remains on a disk array until reused, but cannot be used for instant recovery. It is part of the replica set.

  If the backup option `Keep the replica after the backup` is not selected, the replica is deleted after the backup.

- **ZDB to disk+tape**

  With this type of backup, target volumes (snapclones) are created, and data is kept on a disk array until reused and also moved to backup media. The replica created in a ZDB-to-disk+tape session can be used for instant recovery and is part of the replica set.

  Such a backup session is performed when the `Track the replica for instant recovery` backup option is selected when creating a backup specification, and the `To disk+tape` option is selected when running or scheduling a backup.

IMPORTANT

Before creating a backup specification, consider all limitations regarding the Data Protector EVA integration. For information, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

Snapshot Types

Data Protector supports all EVA snapshot types:

- Snapshots *without* the pre-allocation of disk space. On EVA, it is referred to as **vsnap** or Virtually Capacity-Free Snapshot.
- Snapshots *with* the pre-allocation of disk space. On EVA, it is referred to as **standard snapshot**.
- A full copy of a source volume (original virtual disk), independent of the original virtual disk. On EVA, it is referred to as **Snapclone**.

You select the type of snapshots in the Data Protector GUI when creating a backup specification.

For more information on the snapshot types, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

Replica Creation

A new replica is created and added to the replica set in the following circumstances:

- A ZDB-to-tape session is performed in which the `Keep the replica after the backup` option *is not* selected. Such a replica is only temporarily part of the replica set; it is deleted after the backup to tape is finished.
- A ZDB-to-tape session is performed in which the `Keep the replica after the backup` option *is* selected, but the specified `Number of`

`replicas rotated` is not reached yet. Such a replica becomes part of the replica set.

- A ZDB-to-disk or ZDB-to-disk+tape session is performed (the `Track the replica for instant recovery` option is selected), but the specified `Number of replicas rotated` is not reached yet. Such a replica becomes part of the replica set.

Replica Reuse

The oldest replica in the replica set is deleted and a new one is created in the following circumstances:

- A ZDB-to-tape session is performed in which the `Keep the replica after the backup` option *is* selected and the specified `Number of replicas rotated` is reached.
- A ZDB-to-disk or ZDB-to-disk+tape session is performed and the specified `Number of replicas rotated` is reached.

After a Replica Creation/Reuse

If the backup option `Keep the replica after the backup` was not selected, the replica and, therefore, all target volumes created during the backup session are deleted. Otherwise, the replica is left on a storage system.

EVA Backup Flows

The following are the snapshot backup flows on EVA depending on the selection of backup options. For detailed information on these options, see “Backup Options” on page 103.

Figure 5-1 ZDB-to-Disk Session Flow

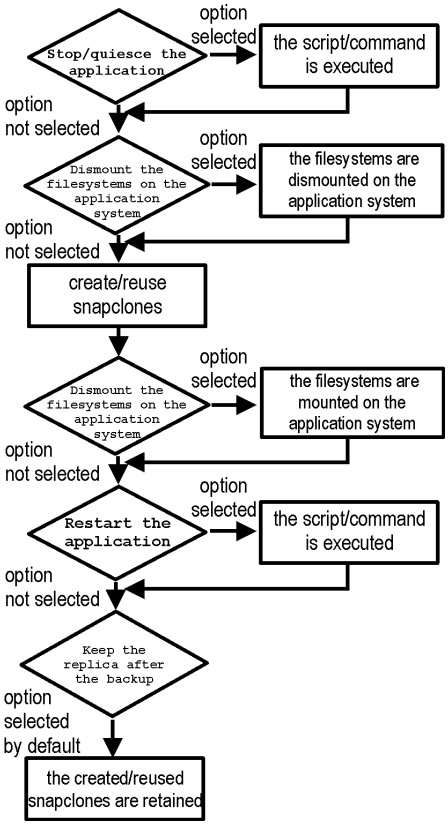

— Only snapclones are supported.

— *“Reuse”* actually means that Data Protector deletes the snapclones from the oldest replica and creates a new replica.

— Due to the EVA limitation, the creation of snapclones may fail if a target volume of another type exists on the array. Such target volumes should be deleted first.

— The `Enable the backup system in read/write mode` option is missing in the figure because it is ignored in the case of a ZDB-to-disk session.

Figure 5-2 ZDB-to-Tape and ZDB-to-Disk+Tape Session Flow

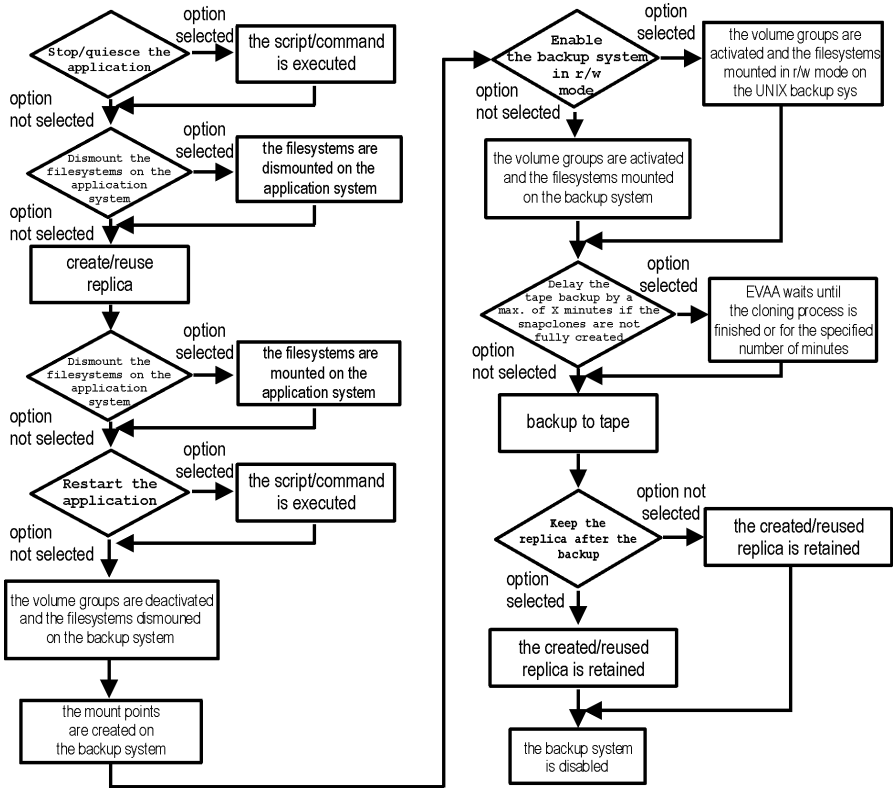

— In the case of a ZDB-to-disk+tape session, only snapclones are supported. In the case of a ZDB-to-tape session, the creation of target volumes depends on your snapshot type and policy options selected. For information on these two options, see “Backup Options” on page 103.

— *“Reuse”* actually means that Data Protector deletes the target volumes from the oldest replica and creates a new replica.

— Due to the EVA limitation, in the case of a ZDB-to-disk+tape session, the creation of snapclones may fail if a target volume of another type exists on the array. Such target volumes should be deleted first.

— The `Enable the backup system in read/write mode` option is related to UNIX systems only, since on UNIX systems, filesystems are normally mounted in read-only mode, while on Windows systems, filesystems are always mounted in read/write mode.

— In the case of a ZDB-to-tape backup, you can select the option `Keep the replica after the backup`. In the case of a ZDB-to-disk+tape backup, this option is selected by default and cannot be deselected.

  In the case of a ZDB-to-tape backup, you can select the option `Keep the replica after the backup`.

  For ZDB-to-disk+tape session, you select this option is selected by default and cannot be deselected.

Configuring a Backup Specification

Use the Data Protector GUI to create a filesystem or disk image snapshot backup specification for use with the EVA integration.

1. In the Context List, select `Backup`.
2. In the Scoping Pane, expand `Backup` and then `Backup Specifications`. Right-click `Filesystem` (for both filesystem backup and disk image backup) and click `Add Backup`.

   The `Create New Backup` dialog box is displayed. See Figure 5-3 on page 96.

   In the `Filesystem` box, select the `Blank Filesystem Backup` template. For more information on the templates, refer to the *HP OpenView Storage Data Protector Administrator's Guide*.

   In the `Backup type` drop-down list, select the `Snapshot backup` option and in the `Sub type` drop-down list, select the EVA agent that is installed on the application and the backup system (`HP StorageWorks EVA (legacy)` or `HP StorageWorks EVA SMIS-S`). See Figure 5-3 on page 96.

   Click `OK`.

Figure 5-3 Create New Backup Dialog Box

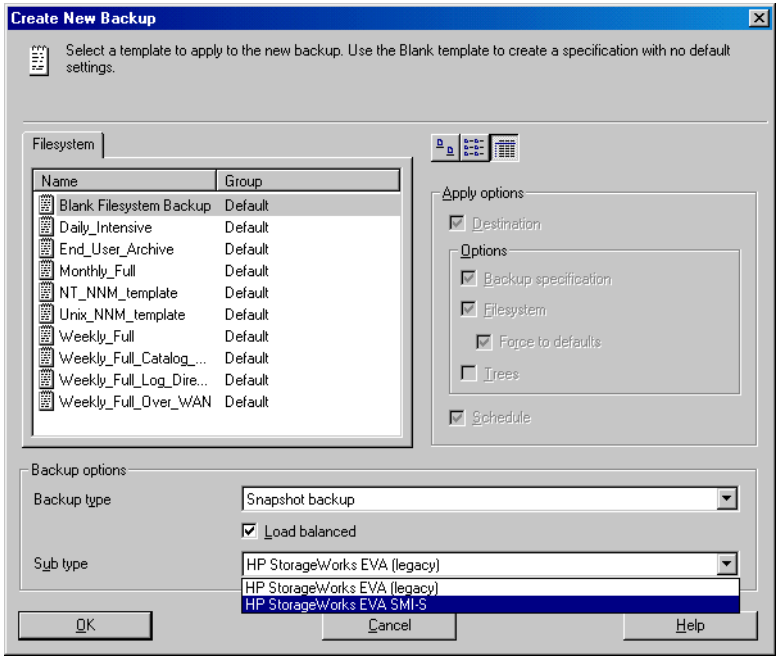

3. Under `Client systems`, select the application and backup systems in the `Application system` and the `Backup system` drop-down lists.

   Specify other options in the following way:

For ZDB to Disk and ZDB to Disk+Tape

Under `Instant recovery option`, select the `Track the replica for instant recovery` option. Under `Replica management options`, specify the `Number of replicas rotated`. Note that the maximum number for vsnaps and standard snapshots is 7. The GUI does not limit the number of replicas rotated, but the session will fail if the limit is exceeded.

Note that `Snapclone` as `Snapshot type` option and `Strict` as `Snapshot policy` option are automatically selected. See Figure 5-4 on page 98.

TIP

For a ZDB-to-disk+tape backup, it is recommended to select the option `Delay the tape backup by a maximum of X minutes if the snapclones are not fully created`. In this case, the backup to tape starts only after the cloning process is finished, but not later than the specified number of minutes. Thus, you prevent degradation of the application data access times during the phase of backup to tape.

For detailed information on these and other options, see “Backup Options” on page 103.

IMPORTANT

You specify whether you want to perform a ZDB-to-disk or a ZDB-to-disk+tape session using the `Split mirror/snapshot backup` option when you run a backup or when you schedule a backup specification. See “Running and Scheduling a ZDB Session” on page A-3 for more information.

Figure 5-4 **EVA Backup Options Required For ZDB-to-Disk or ZDB-to-Disk+Tape Session**

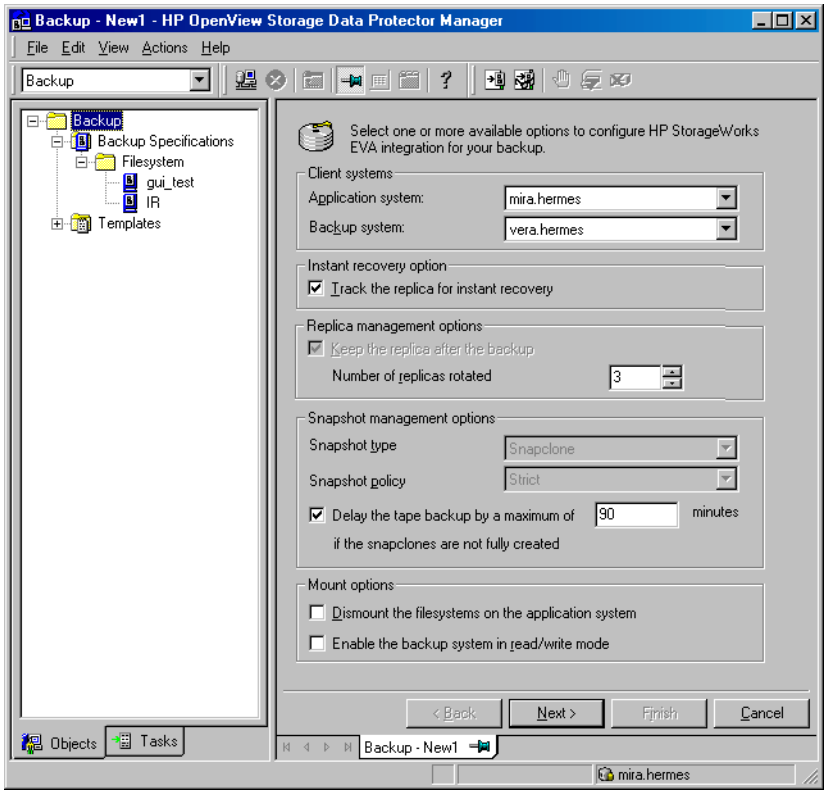

For ZDB to Tape

If you do not want the replica to be deleted after the backup session, select the `Keep the replica after the backup` option and specify the `Number of replicas rotated`. Note that the maximum number for vsnaps and standard snapshots is limited by the target EVA storage system maximum. The GUI does not limit the number of replicas rotated, but the session will fail if EVA's limit is exceeded.

Under `Snapshot management options`, select the `Snapshot type` and `Snapshot policy`.

TIP

If you select Snapclones to be created, it is recommended to select the option Delay the tape backup by a maximum of X minutes if the snapclones are not fully created. In this case, the backup to tape starts only after the cloning process is finished, but not later than the specified number of minutes. Thus, you prevent degradation of the application data access times during the phase of backup to tape.

For detailed information on these and the Mount options, see "Backup Options" on page 103.

Click Next.

4. Depending on whether you perform a filesystem or disk image backup, proceed as follows:

Filesystem Backup

If you are configuring a filesystem backup, expand the application systems that contain the objects that you want to back up and then select what you want to back up.

IMPORTANT

On UNIX, if you intend to perform instant recovery, *select all filesystems inside the volume group* to be backed up. Otherwise, instant recovery will not be possible using the Data Protector GUI or (if you perform instant recovery using the Data Protector CLI) data can be corrupted.

Click Next.

Disk Image Backup

If you are configuring a disk image backup, click Next.

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

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

   For detailed information on the object mirror functionality, refer to the *HP OpenView Storage Data Protector Administrator's Guide*.

NOTE Object mirroring is not supported for ZDB to disk.

Click `Next`.

6. In the `Backup Specification Options` group box, click the `Advanced` tab and then `HP StorageWorks EVA (legacy)` or `HP StorageWorks EVA SMI-S` to open the EVA backup options.

   You can specify the `Application options` and modify all other options, except the `Application system` and `Backup system` options. For information on the EVA backup options, see "Backup Options" on page 103.

   Click `OK`.

   For information on `Filesystem Options`, press **F1**.

Figure 5-5 All EVA Backup Options

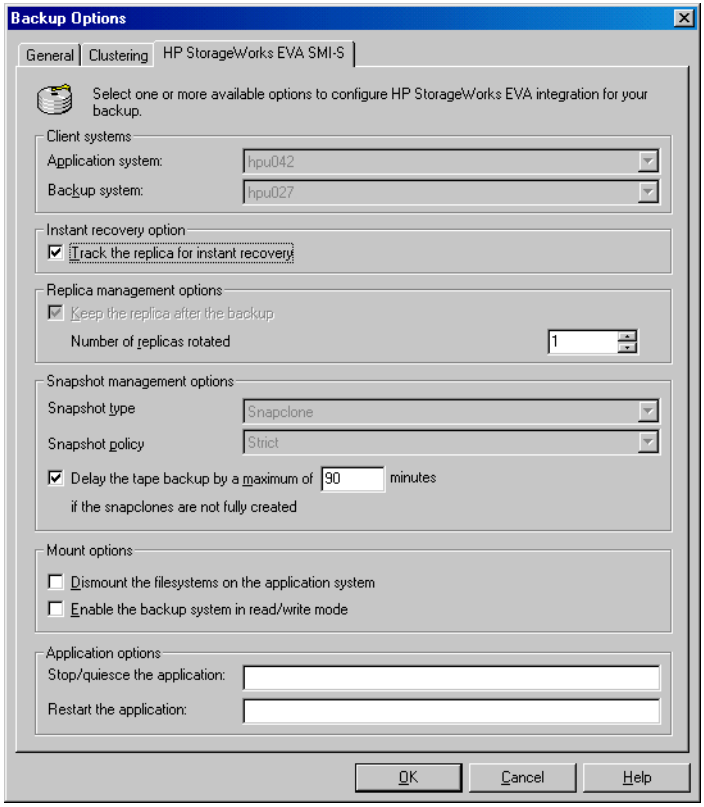

7. Follow the backup wizard to open the Data Protector scheduler (for information on scheduler, press **F1** or see "Running and Scheduling a ZDB Session" on page A-3) and then the backup summary page, where a summary of the backup specification is given.
8. Depending on whether you perform a filesystem or disk image backup, proceed as follows:

Filesystem Backup

If you are configuring a filesystem backup, click `Next`.

Disk Image Backup

For a disk image backup, proceed as follows:

a. Click `Manual add` to add the disk image objects you want to back up.

b. Select `Disk image object` and click `Next`.

c. Select the client to be backed up and click `Next`.

d. Follow the wizard to specify the `General Object Options` and the `Advanced Object Options`. For more information on these options, press **F1**.

e. In the `Disk Image Object Options` window, specify the disk image sections you want to back up.

On UNIX Systems

To specify a rawdisk section, use the following format:

`/dev/rdsk/<filename>`, for example: `/dev/rdsk/c2t0d0`

On HP-UX, to specify a raw logical volume section, use the following format:

`/dev/vg<number>/rlvol<number>`, for example: `/dev/vg01/rlvol1`

IMPORTANT

If you intend to perform instant recovery, *specify all raw logical volumes inside the volume group* to be backed up. Otherwise, instant recovery will not be possible using the Data Protector GUI or (if you perform instant recovery using the Data Protector CLI) data can be corrupted.

On Windows Systems

Use the following format:

`\\.\PHYSICALDRIVE#`,

where `#` is the current number of the disk you want to back up.

For example: `\\.\PHYSICALDRIVE3`

For information on how to find the current numbers of the disks (physical drive numbers) you want to back up, refer to the online Help index keyword “disk image backups”.

f. Click `Finish` and then `Next`.

9. Save your backup specification. For more information on starting and scheduling a ZDB session, see “Running and Scheduling a ZDB Session” on page A-3.

Backup Options

The following tables describe the EVA backup options. See “VA and EVA Integrations” on page A-37 to help you understand these options.

Table 5-1 EVA Client Systems Options

`Application system`	Specify the application system on which the application runs (for example, an Oracle database). In a cluster environment, specify the virtual server hostname (rather than the physical node hostname).
`Backup system`	Specify the backup system on which your data is to be backed up. In a cluster environment, specify the virtual server hostname (rather than the physical node hostname).

Table 5-2 EVA Instant Recovery Option

`Track the replica for instant recovery`	Select 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. If this option is selected, you should also set the `Number of replicas rotated` parameter. Note that when this option is selected, the options `Keep the replica after the backup`, `Snapclone`, and `Strict` are automatically selected. By default, this option is not selected.

Table 5-3 EVA Replica Management Options

`Keep the replica after the backup`	By default, this option is automatically selected (and cannot be deselected) if the `Track the replica for instant recovery` option is selected. If configuring a ZDB to tape, select this option to keep the replica on a disk array after the ZDB-to-tape session. By selecting this option, the replica becomes part of a replica set; therefore, you should also set the `Number of replicas rotated` parameter (see below). With ZDB-to-tape sessions, the replica *is not available for instant recovery*. If this option is not selected, the replica is deleted after the backup session.
`Number of replicas rotated`	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. Note that this option sets the number of replicas in the replica set for a backup specification. You need to specify this number if the `Keep the replica after the backup` option is selected. By default, this number is set to 1. 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.

Table 5-4 EVA Snapshot Management Options

Snapshot type	`Vsnap`	Select this option to create snapshots without the pre-allocation of disk space. For more information on this type of snapshot, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*. This is the default selection.
	`Standard snapshot`	Select this option to create snapshots with the pre-allocation of disk space. For more information on this type of snapshot, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.
	`Snapclone`	Select this option to create a clone of a source volume (original virtual disk). For more information on this type of snapshot, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.
Snapshot policy	`Strict`	With this option selected, Data Protector attempts to create snapshots of the type selected by the `Snapshot type` option. If some of the source volumes (original virtual disks) used in the backup session already have existing snapshots of a different type, the selected type of snapshots cannot be used. Such a backup session will be aborted. This is the default selection.
	`Loose`	With this option selected, Data Protector creates snapshots of a different type than specified by the `Snapshot type` option when this would help to make a successful session. 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 `Loose` option selected, Data Protector creates either vsnaps (if vsnaps already exist) or snapclones (if snapclones already exist) instead of standard snapshots. Note that Data Protector can use only one type of snapshots in a backup session. For example, if some of the source volumes (original virtual disks) used in the backup session have existing standard snapshots and some of them existing vsnaps, the backup session will be aborted.

Table 5-4 EVA Snapshot Management Options

	Delay the tape backup by a maximum of X minutes if the snapclones are not fully created	You can select this option if the selected snapshot type is Snapclone. 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. Specify also 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. With this option, you prevent degradation of the application data access times during the phase of backup to tape. The default value is 90 minutes.

Table 5-5 EVA Mount Options

Dismount the filesystems on the application system	Specify this option if you want the filesystem on the application system to be dismounted before a snapshot is created and remounted after a snapshot is created. A filesystem does not have the stop I/O functionality to flush the data from the filesystem cache to the disk and stop the I/O for the time of the snapshot. This option can be used to ensure that the data on the filesystem is consistent. If an integrated application (for example, Oracle or SAP R/3) runs on the filesystem, it controls the I/O to the disk. In this case it is not necessary to dismount the filesystem before the snapshot creation. By default, this option is not selected.
Enable the backup system in read/write mode	This option is related to UNIX systems only, since on Windows systems filesystems are always mounted in read/write mode. Specify this option to have read/write access to the volume groups and filesystems on the backup system. For backup, it is sufficient to activate the backup system volume groups and filesystem in read-only mode. However, to perform other tasks on the backup system, this might not be sufficient. By default, this option is not selected.

Table 5-6 EVA Application Options

`Stop/quiesce the application`	Specify the optional `Stop/quiesce the application` command/script. Create this command in the `/opt/omni/lbin` (UNIX systems) or in the *<Data_Protector_home>*`\bin` (Windows systems) directory on the application system and specify only the filename in the backup specification. This command is executed on the application system immediately before the snapshot creation is performed. It is mainly used to stop the application.
`Restart the application`	Specify the optional `Restart the application` command/script. Create this command in the `/opt/omni/lbin` (UNIX systems) or in the *<Data_Protector_home>*`\bin` (Windows systems) directory on the application system and specify only the filename in the backup specification. This command is executed on the application system immediately after the snapshot creation is performed. It is mainly used to restart the application.

Application Options Specifics

Using the `Application options`, you can stop any application that is not integrated with the Data Protector EVA integration on the application system before the snapshots are created, and restart this application after the snapshots are created.

If the execution of the `Stop/quiesce the application` command fails, the `Restart the application` command is also not executed, so a cleanup procedure should be implemented in the `Stop/quiesce the application` command. However, if the `ZDB_ALWAYS_POST_SCRIPT` `omnirc` variable is set to `1`, the `Restart the application` command is always executed if set. By default, this variable is set to `0`. See “ZDB Agents Omnirc Variables” on page A-23 for more information on the Data Protector ZDB Agents `omnirc` variables.

TIP

The `Application options` can also be used for any other operation on the application system before and after the snapshot creation.

NOTE

The EVA `Application options` are accessible from the `Backup Specification Options` group box (see step 6 on page 100), by clicking the `Advanced` tab and then the `HP StorageWorks EVA (legacy)` or the `HP StorageWorks EVA SMI-S` tab.

Troubleshooting

This section describes the most common problems you may encounter when using the Data Protector EVA integration.

Before You Begin

- Ensure that the latest official Data Protector patches are installed. Refer to the *HP OpenView Storage Data Protector Software Release Notes* or http://www.openview.hp.com/products/datapro/spec_0001.html for an up-to-date list of supported versions, platforms, and other information.
- Examine system errors reported in `/var/opt/omni/log/debug.log` (HP-UX or Solaris systems) or in *<Data_Protector_home>*`\log\debug.log` (Windows systems) on the application system and the backup system.

The following is a description of problems, and actions to be taken to resolve them.

Backup Problems

- **You cannot select the StorageWorks mode in the Data Protector user interface when attempting to create a backup specification.**

  Check whether the `HP StorageWorks EVA Agent (legacy)` or the `HP StorageWorks EVA SMI-S Agent` integration module is installed on both the application and backup systems.

  Check the Data Protector `cell_info` file on the Data Protector Cell Manager, in *<Data_Protector_Home>*`\Config\server\cell\cell_info` (Windows Cell Manager) or in `/etc/opt/omni/server/cell/cell_info` (UNIX Cell Manager) to see if the needed component is installed.

  The file contents should look similar to the following examples:

  ```
  -host "hpsap002.bbn.hp.com" -os "hp s800 hp-ux-11.00" -cc
  A.05.50 -da A.05.50 -ma A.05.50 -EVAA A.05.50
  ```

or:

```
-host "hpsap002.bbn.hp.com" -os "hp s800 hp-ux-11.00" -cc
A.05.50 -da A.05.50 -ma A.05.50 -SMISA A.05.50
```

- **The EVA SMI-S Agent fails to connect to the Cell Manager and retrieve configuration data.**

Message

```
[Major]

Cannot connect to the Cell Server. (Insufficient permissions.
Access denied.)
```

Description

Since the EVA SMI-S Agent is always started as an administrator's process on the application and the backup systems, the user who starts this process must be the member of either the Data Protector `admin` or `operator` user groups.

Action

Using the Data Protector GUI, check if the user is a member of the Data Protector `admin` or `operator` user group. If not, add the user to one of the above-mentioned user groups. In addition, make sure that the administrators from both the application and the backup system belong to the Data Protector `admin` or `operator`.

- **There are no HP StorageWorks SMI-S EVA provider entries configured within the SMISDB.**

  Add an HP StorageWorks SMI-S EVA provider entry to the SMISDB by running the following command:

  ```
  omnidbsmis -ompasswd -add <hostname> [-port
  <port_number>] [-namespace <namespace>] [-user
  <username>] [-passwd <password>]
  ```

- **The EVA agent (EVA Agent (legacy) or EVA SMI-S Agent) on the application system failed to dismount a filesystem.**

  Ensure that there run no processes that use the filesystem that Data Protector tries to dismount. If a `Stop/quiesce the application` script has been specified, check that it stops all processes using the filesystem.

- **On Windows, snapshots cannot be mounted to the target location on the backup system.**

Message

```
[Major]

Filesystem \\.\Volume{9640da9a-6f36-11d7-bd7a-000347add7ba}
```

```
could not be mounted to C:\Program
Files\OmniBack\tmp\ranj.ipr.hermes\N\.

([145] The directory is not empty. ).
```

Description When running a backup session with nested mountpoint objects, snapshots sometimes cannot be mounted to the target mountpoint location on the backup system. This happens when cleaning of the target mountpoint location has failed.

Action On the backup system, manually clean the directory where filesystems are to be mounted. For example, if your target location is *<Data_Protector_home>*\tmp (the default location), make sure that the tmp directory is empty.

- **Adding the login information for Command View EVA using the omnidbeva command fails.**

Message

```
Failed to add login information for Command View EVA on host
<managment_appliance_hostname>.
```

Description User name or password entered in the command prompt may be wrong.

In case you use the Command View EVA v3.0 or v3.1, note that it has two different login schemes:

— new IIS scheme

— legacy Elm scheme

Action When running the omnidbeva -empasswd -add *<management_appliance_hostname>* command, use the legacy Elm scheme for the login information. The login information may differ from what you have entered in the Command View EVA v3.0 user interface.

For setting the legacy Elm password, browse http://*<management_appliance_hostname_or_management_appliance_IP_address>*:2301/cpqlogin.htm?ChangePassword=yes".

- **The EVA Agent (legacy) fails to obtain an EVA storage system port information**.

Message

```
[MAJOR]

Failed to obtain StorageWorks EVA port information from the Cell
```

```
Manager. Please use the \"omnidbeva -hwrescan\" command to
update the port information in the internal database.
```

Description

There are two possible reasons for this error:

— Command View EVA that controls an EVA storage system is not configured within the Data Protector cell.

Action

Check whether the relevant Command View EVA is configured within the Data Protector cell by running the `omnidbeva -empasswd -list` command. The command provides a list of all Management Appliance systems that are configured within the Data Protector cell.

Proceed as follows:

— If the corresponding Management Appliance is missing, provide the login information for it. Run the following command:

```
omnidbeva -empasswd -add
<management_appliance_hostname>
```

— If the corresponding Management Appliance is configured but its password has changed, update the configuration. Run the following command:

```
omnidbeva -empasswd -update
<management_appliance_hostname>
```

— A source volume exists on an EVA storage system that is new to Data Protector or the configuration of the EVA storage system has changed significantly.

Action

Manually rescan the EVA storage system configuration by running the following command:

```
omnidbeva -hwrescan
```

- **An error occurs while parsing a Command View EVA response (the EVA Agent (legacy) problem).**

Message

```
An unknown error occurred while parsing a Command View EVA
```

response.

Description This message indicates that the Command View EVA server did not function properly. This may happen when the port number 2301 instead of 12301 is used for the Command View EVA v3.0.

Action If this is the case, use the port number 12301. To switch the port, run the following command on the application or backup system:

```
omnidbeva -empasswd -update <management_appliance_hostname> -port 12301
```

- **A Command View EVA command fails.**

Message

```
A Command View EVA command failed on the server side.
```

This message may be followed by one of the following error descriptions:

— `Receive timeout reach, no data to receive`

— `Logical Disk Sharing`

— `SCMI Target Object Does Not Exist`

— `Invalid host adapter ID`

— `Name already exists`

— `Operation rejected - The Vdisk has a sharing relationship with another object`

— `The requested allocation policy cannot be satisfied. A different allocation policy is already in use by another snapshot in the same family.`

— `SCMI Invalid Target Handle`

— `Invalid Target Handle`

Description Command View EVA that executes Data Protector requests occasionally returns an error. Data Protector retries such requests, but the number of retries or the time interval between the retries may be insufficient on the heavily loaded EVA storage systems.

Action On the application and the backup systems, adjust the `omnirc` variables `EVA_EMAPI_MAX_RETRY` and `EVA_EMAPI_RETRY_DELAY` to the needs of the EVA storage system, and restart the backup session.

- **Data Protector fails to find FC HBAs on the backup system (the EVA Agent (legacy) problem).**

Message

```
[Warning]
Failed to find any FC HBA that would connect the backup system
to the target StorageWorks EVA storage system.
StorageWorks EVA name: <EVA_name>
```

Description

Data Protector tries to automatically detect the local FC HBAs that connect the backup system to the target EVA storage system. This is necessary for the automatic backup system preparation. The detection may fail because the FC HBA drivers on the backup system do not support SNIA's Common HBA API or there is a problem with the connectivity between the backup system and the target EVA storage system.

Action

Check the connectivity and the SAN zoning between the backup system and the target storage system. Afterwards, try to update the FC HBA drivers on the backup system.

You may also decide to disable the automatic preparation of the backup system. In this case, set to `1` the `EVA_DISABLE_AUTODETECTION` `omnirc` on the backup system and ensure that the host object representing the backup system is configured within Command View EVA. If the name of this host object is different from the actual backup system hostname, specify the host object name in the `omnirc` variable `EVA_HOSTNAMEALIASES` on the backup system.

For more information on the variables, see "ZDB Agents Omnirc Variables" on page A-23.

- **Preparation of the backup system takes longer than usual.**

Message

```
[Normal]
Preparing the backup system for backup to tape.
```

Description

After the above message is displayed, it takes a long time for the backup session to continue. During this time, Data Protector performs an extensive search over the EVA storage system to check if the host object representing the backup system is already configured. This search can take a long time in large environments.

Action

You can disable the automatic preparation of the backup system by setting the `omnirc` variable `EVA_DISABLE_HBA_AUTODETECTION` on the backup system. Additionally, you need to ensure that the host object representing the backup system is configured within Command View EVA. If the name of this host object is different from the actual backup system hostname, specify the host object name in the `omnirc` variable `EVA_HOSTNAMEALIASES` on the backup system.

For more information on the variables, see "ZDB Agents Omnirc Variables" on page A-23.

- **Preparation of the backup system fails.**

Messages

```
[Major]

Failed to find any backup host object configured within the
StorageWorks EVA <EVA_name> environment that would match the
selected backup host.

[Critical]

Failed to present replica to the backup host. Aborting!
```

Description

This usually happens when you do not or cannot use the automatic preparation of the backup system and the backup system hostname cannot be found within Command View EVA. In this case, either the host object representing the backup system is not configured within Command View EVA, or the name of the host object representing the backup system is different from the actual backup system hostname.

Action

In case you want Data Protector to automatically prepare the backup system, verify that the `omnirc` variable `EVA_DISABLE_HBA_AUTODETECTION` is set to `0` (not enabled). If the variable is not enabled, the problem is probably caused by the FC HBA drivers on the backup system. The automatic preparation of the backup system is possible only if the FC HBA drivers support SNIA's Common HBA API. Therefore, update the drivers to enable the automatic preparation of the backup system.

In case you choose to disable the automatic backup system preparation, ensure that the `EVA_DISABLE_AUTODETECTION` `omnirc` variable is set to `1` on the backup system and that the host object representing the backup system is configured within Command View EVA. If the name of this host object is different from the actual backup system hostname, specify the host object name in the `omnirc` variable `EVA_HOSTNAMEALIASES` on the backup system.

For more information on the variables, see “ZDB Agents Omnirc Variables” on page A-23.

- **Data Protector fails to start the device scanning process on the backup system.**

Message

```
[Major]
Failed to lock drive scan phase for use by this session only.
Maximum number of retries for locking has been reached.
```

Description

After target volumes are created and presented to the backup system, a device scanning process is run on the backup system in order to detect the target volumes. Only one device scanning process is run at the same time. In case of parallel backup sessions, one of the sessions runs the scanning process, while the other ones wait for this process to finish before starting another scanning process. On some systems, the device scanning process may take longer, and when the timeout is reached, the backup session aborts.

Action

Increase the number of retries for running the device scanning process before aborting the backup session (the EVA agent waits for 10 seconds between the retries). To do this, set the `omnirc` variable `EVA_SCANDEVICES_LOCK_MAX_RETRY` on the backup system.

If the problem still persists, check if there is an active session hanging during the preparation or resumption of the backup system. Abort such session. Refer also to the following problem and its troubleshooting.

- **On HP-UX, the backup session freezes either during the preparation or during the resuming of the backup system. It freezes after one the following messages is displayed:**

Messages

```
[Normal]
Starting drive discovery routine.
[Normal]
Resuming the backup system.
```

Description

During the preparation of the backup system, Data Protector adds new devices to the Secure Path control and runs the device scanning process. During the resuming of the backup system, Data Protector removes the devices from the Secure Path control and runs the device scanning process.

If some other process runs the Secure Path commands or the device scanning procedure at the same time (either during the preparation or during the resuming), the backup session may freeze. To identify this problem, run the `ps -ef` command several times on the backup system and check if there are any `ioscan` or `spmgr` processes that persist in the output.

Action

Abort the backup session and stop the hanging `ioscan` and `spmgr` processes.

If the processes cannot be stopped, reboot the backup system. After the reboot, it is recommended to manually clean up the backup system. Proceed as follows:

1. On the backup system, run the `spmgr display` command to find out if any target volume created in the failed backup session is left under the Secure Path control.
2. Remove such target volumes from the Secure Path control using the `spmgr delete` command.
3. Run the `spmgr update` command, and then follow the reported instructions to make the changes persistent across reboots.
4. Using the Command View EVA user interface, delete all the presentations attached to the removed target volumes.

- **Data Protector fails to map a target volume to a valid character device file.**

Message

```
[Major]
Failed to map a target volume to a valid character device file:
Virtual disk name: <Virtual_disk_name>
StorageWorks EVA name: <EVA_name>
```

Description

Data Protector cannot find the target volumes on the backup system. The following reasons can cause this problem:

— A problem with the connectivity between the backup system and the EVA storage system.

Action

Try to manually present an EVA virtual disk from the target EVA storage system to the backup host. If you cannot get the virtual disk visible on the backup system, it is most likely that you have a

connectivity or host object configuration problem. Refer to the HP StorageWorks Enterprise Virtual Array and the HP StorageWorks Secure Path manuals to troubleshoot the problem.

— On Windows, a problem with device drivers on the backup system.

Action

If manual presentation works, it is most likely that you have a device drivers problem. This problem may occur as it takes a long time for the Windows 2000 systems to detect new devices attached to the system. By default, Data Protector waits for 300 seconds for a new physical drive to appear before aborting the session. You can increase this time period by setting the `omnirc` variable `SMB_SCAN_RDSK_TIMEOUT`. For more information on the variable, see “ZDB Agents Omnirc Variables” on page A-23.

If the problem persists, it indicates an inconsistent state of the backup system. Reboot the backup system in order to restart the device drivers and the related services.

— On HP-UX or Solaris, a problem with the configuration entry for the backup system within the EVA storage system.

Action

The problem may occur when you unpresent or delete the EVA virtual disks without removing them from the Secure Path control first. This may leave device instances in the driver configuration and thus disable detection of new devices. To solve the problem, run the `spmgr display` command on the backup system and check if any such instances exist. Typically, they will be reported as failed or inactive. Delete such instances using the `spmgr delete` command. Afterwards, run the `spmgr update` command and follow the reported instructions to make the changes persistent across reboots.

- **On Windows, Data Protector fails to find the volume on a target volume.**

Message

```
[Major]
Volume <volume_name> on target volume has not been found.
```

Descriptions

Data Protector found target volumes on the backup system, but the volumes that should have existed on the target volumes were not found in the system's device list. The following reasons can cause this problem:

— Failed installation of the "Platform Kit for EVA" on the backup system.

Action

Verify that the file `%WINDOWS%\inf\HsgCCL.inf` exists on the backup system. If the file does not exist, install the "Platform Kit for EVA" again.

— It takes unusually long time to detect new volumes on the backup system.

Action

Increase the Data Protector timeout value for detecting new volumes by setting the `omnirc` variable `SMB_SCAN_FOR_VOLUME_TIMEOUT` on the backup system. For more information on the variable, see “ZDB Agents Omnirc Variables” on page A-23.

— The backup system is in an inconsistent state.

Action

Reboot the backup system in order to restart the device drivers and the corresponding services.

6 Restore

In This Chapter

This chapter describes how to configure and run a filesystem or disk image restore of your data backed up using the Data Protector HP StorageWorks Enterprise Virtual Array (EVA) integration. The sections describe steps for performing a restore using the Data Protector Graphical User Interface and Data Protector Command Line Interface.

This chapter includes the following sections:

"Overview" on page 123

"Restoring from Backup Media on LAN" on page 124

"Instant Recovery" on page 125

"Troubleshooting" on page 132

Overview

The data backed up using the Data Protector EVA integration is stored:

- On a disk array; after z ZDB-to-disk or ZDB-to-disk+tape session.
- On the backup media; after a ZDB-to-tape or ZDB-to-disk+tape session.

The data stored on a disk array is restored using instant recovery; whereas, the data stored on the backup media is restored using the standard Data Protector restore procedure. This procedure restores the data from the backup media to the application system through a LAN.

Table 6-1 Restore Types Available After a Specific ZDB Session

	Instant Recovery	**Standard Restore**
ZDB to disk	Yes	N/A
ZDB to disk+tape	Yes	Yes
ZDB to tape	N/A	Yes

From the point of view of what is restored, the difference between the two types of restore is that with instant recovery, all of the data in the replica is restored (regardless of the backup objects selection during the backup); whereas, with the standard Data Protector restore, only the backup objects selected for the restore are restored.

As far as the speed of the two types of restore is concerned, instant recovery is faster than standard Data Protector restore, because data moves faster within a disk array than from backup media through a LAN.

Restoring from Backup Media on LAN

The data backed up in ZDB-to-tape and ZDB-to-disk+tape sessions can be restored from backup media on a LAN. The restore is performed in the same way as the standard Data Protector restore. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for more information on standard Data Protector restore and for instructions on how to perform it.

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. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for more information on configuring backup devices. Refer to the "Restoring Under Another Device" section of the same guide for more information on how to perform a restore using another device.

Instant Recovery

Using the Data Protector EVA integration, only the data backed up in ZDB-to-disk and ZDB-to-disk+tape sessions can be restored in instant recovery sessions.

Instant recovery can be performed:

- Using the Data Protector GUI. For information, see “Instant Recovery Using the Data Protector GUI” on page 128.
- Using the Data Protector CLI. For information, see “Instant Recovery Using the Data Protector CLI” on page 130.

The number of replicas available for instant recovery for a ZDB-to-disk or ZDB-to-disk+tape backup specification is limited by the number specified with the `Number of replicas rotated` backup option. This option sets the size of the replica set. The replicas that are available for instant recovery can be listed by selecting `Instant Recovery` in the Context List of the Data Protector GUI, and then expanding the `Restore Sessions` folder. Alternatively, the Data Protector CLI can be used by running the `omnidbeva -list -session -ir` command.

The replicas are accessible by the backup specification name and ID of the ZDB-to-disk and ZDB-to-disk+tape sessions that produced them; for example, `2003/07/23-10[IR]`. In the Results Area, also other data is provided; for example, time when the replica was created. In this way, you can identify the correct replica containing the data for restore.

Data Protector instant recovery moves data from a replica to the source volumes directly, without involving a backup device. At the beginning of the instant recovery session, the application system needs to be disabled. This includes dismounting the filesystems and deactivation of volume groups (UNIX only) by Data Protector. Before this can be done, the status of participating filesystems and volume groups (UNIX only) is checked, and only the mounted filesystems and activated volume groups (UNIX only) are dismounted and deactivated. At the end of the session, only those filesystems that were previously dismounted will be mounted, and only those volume groups that were previously deactivated will be activated (UNIX only).

IMPORTANT

After an instant recovery session, the configuration of the restored filesystems is the same as it was at the backup time. This means that the restored filesystems are mounted to the same mount points or drive letters as they were at the backup time. In case these mount points or drive letters have some other filesystems mounted, these filesystems are automatically dismounted before instant recovery and the restored filesystems are mounted after instant recovery.

Instant Recovery Process

When an instant recovery session is started, the following happens:

- The Restore Session Manager (RSM) on the Data Protector Cell Manager starts the relevant Data Protector EVA agent (the EVA Agent (legacy) or the EVA SMI-S Agent) on the application system and sends the restore session information. This sets up a communications link with the allocated backup system, and the agent is started on the backup system. The two agents then have a direct communications link.
- The EVA agent queries the EVADB/SMISDB to find out the restore objects associated with the specified ZDB-to-disk or ZDB-to-disk+tape session.
- If the `Check the data configuration consistency` instant recovery option is selected, the following takes place:

  ✓ On UNIX systems, the volume group configurations for the source volumes stored in the EVADB/SMISDB is compared to the selected replica volume group configurations. If the items compared do not match, the session is aborted.

  When instant recovery is performed in an MC/ServiceGuard cluster to some other node than the one that was backed up, the current volume group configuration on the node to which instant recovery is to be performed is different from the volume group configuration stored in the EVADB/SMISDB. In such a case the EVADB/SMISDB volume group configuration data is replaced by the current volume group configuration data on the node to which instant recovery is to be performed and the session is not aborted.

- The application and backup systems are disabled by deactivating the participating volume groups (UNIX systems) and dismounting the participating filesystems.
- A check is then made to verify that no snapclone in the replica set with the replica to be restored is in use by Data Protector. This is done by checking that all snapclones in the replica set can be locked. If this check fails, the restore session is aborted.
- The EVA agent checks that no snapclone from the selected replica is presented to any host. In case that a snapclone is presented to a host, it also checks if the `Force removal of all replica presentations` option is selected. If this option is selected, the agent removes such presentations; otherwise, the instant recovery session fails.
- The EVA agent stores information on the WWNs and presentations of the source volumes that are to be replaced during the instant recovery. It then removes all presentations from the source volumes.
- Source volumes (original virtual disks) are replaced with the snapclones from the selected replica. The WWNs of the replaced source volumes are assigned to the snapclones.
- For every removed presentation on the application system, the EVA agent creates a new presentation on the snapclones (new source volumes).
- The replaced source volumes are deleted.
- All information related to the replica that was used in the instant recovery session is deleted from the EVADB/SMISDB.
- The application system is enabled by activating the participating volume groups (UNIX systems) and mounting the participating filesystems.

Instant Recovery Procedure

To perform a filesystem or disk image instant recovery using the Data Protector GUI, follow the procedure described in “Instant Recovery Using the Data Protector GUI” on page 128. To perform a filesystem or

disk image instant recovery using the Data Protector CLI, follow the procedure described in “Instant Recovery Using the Data Protector CLI” on page 130.

Prerequisite

If you have performed a disk image backup, manually dismount the disks to be restored before instant recovery, and re-mount them afterwards.

IMPORTANT

Before starting instant recovery, carefully consider all instant recovery related limitations and considerations as stated in the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

IMPORTANT

While an instant recovery session, do not perform a ZDB session that uses the same source volumes as those to which the data is restored.

Instant Recovery Using the Data Protector GUI

To perform instant recovery using the Data Protector GUI, proceed as follows:

1. In the Context List, select `Instant Recovery`.
2. Select the backup session (replica) from which you want to perform the restore. This can be done in two ways:
   - By the backup session ID and name:

     In the Scoping Pane, expand `Restore Sessions` and select the session from the list of *all* ZDB-to-disk or ZDB-to-disk+tape sessions.
   - By type of backup (filesystem, Oracle, SAP R/3,...) and backup session name and ID:

     a. In the Scoping Pane, expand `Restore Objects`.

     A list of backed up object types (Filesystem, Disk Image, SAP R/3, Microsoft SQL Server, ...) is displayed.

     b. Expand the type of object you want to restore.

A list of all available backup specification that were used in ZDB-to-disk or ZDB-to-disk+tape sessions for the selected type of object is displayed.

c. Expand the backup specification containing the objects that you want to restore.

A list of all available ZDB-to-disk or ZDB-to-disk+tape sessions (identified by the session ID) for the selected backup specification is displayed.

Figure 6-1 Instant Recovery Session Selection

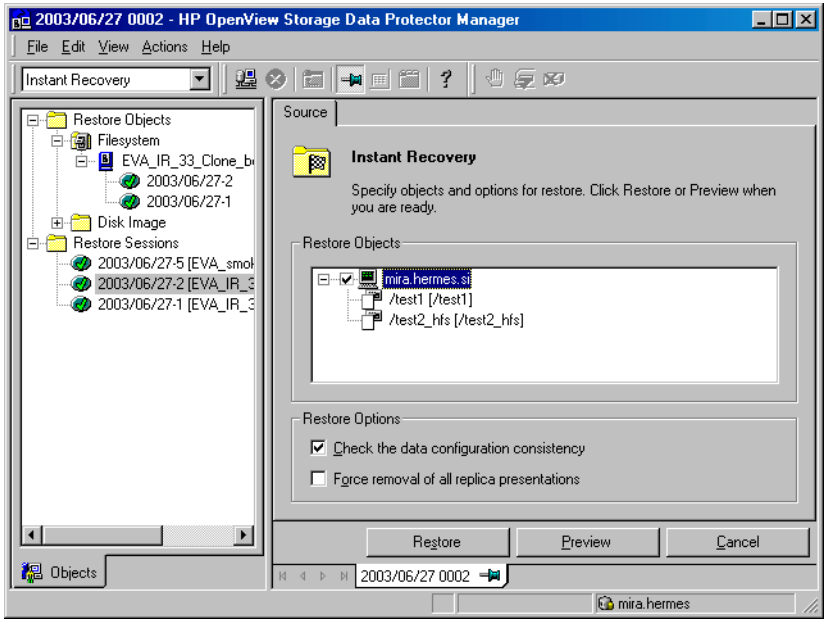

3. In the Scoping Pane, click the backup session you want to restore.

   The application system concerned and its mount points or drive letters (on Windows) representing the source volumes that were backed up during the selected ZDB-to-disk or ZDB-to-disk+tape session are displayed.

4. Check the selection box next to the application system to select the session for restore. Note that you cannot select sub-components. With instant recovery, only the complete session can be restored.

5. Select the instant recovery options that you want to use during the restore. See “Instant Recovery Options” on page 131 or press **F1** for more information on these options.
6. Click `Restore` to open the `Start instant recovery` dialog box.
7. Select `Start Restore Session` to start the instant recovery. It is recommended to test instant recovery to ensure it work properly. Click `OK`.

Instant Recovery Using the Data Protector CLI

The replica to be restored using the Data Protector CLI is identified by the ZDB-to-disk or ZDB-to-disk+tape session ID. To perform instant recovery using the Data Protector CLI, proceed as follows:

1. Get a list of all available ZDB-to-disk or ZDB-to-disk+tape sessions, identified by the session ID, using one of the following commands:

   ```
   omnidbeva -list -session -ir
   ```

   if you have the EVA Agent (legacy), or:

   ```
   omnidbsmis -list -session -ir
   ```

   if you have the EVA SMI-S Agent.

   From the output of the command, select the backup session you want to restore.

2. Execute the following command:

   ```
   omnir -host <application_system_name> -session
   <SessionID> -instant_restore [<INSTANT RECOVERY OPTIONS>]
   ```

   Where:

   - *<application_system_name>* is the hostname of the application system
   - *<SessionID>* is the backup session ID selected in the step 1 of this procedure.

   For *<INSTANT RECOVERY OPTIONS>* see Table 6-2 on page 131.

   This will start the instant recovery session.

Refer to the `omnidbeva`, `omnidbsmis`, and `omnir` man pages for more information on these commands.

Instant Recovery Options

Instant recovery options are set during the configuration of an instant recovery session.

Table 6-2 Instant Recovery Options

Data Protector GUI/CLI	Function
`Check the data configuration consistency/-check_config`	If this option is selected, the current volume group configuration of the volume groups involved in the instant recovery session is compared with the volume group configuration during the ZDB-to-disk or ZDB-to-disk+tape session kept in the EVADB/SMISDB. If the volume group configuration has changed since the ZDB-to-disk or ZDB-to-disk+tape session, the restore is aborted. When instant recovery is performed in an MC/ServiceGuard cluster to some other node than the one that was backed up, the current volume group configuration on the node to which instant recovery is to be performed is different from the volume group configuration kept in the EVADB/SMISDB. In such a case the EVADB/SMISDB volume group configuration data is replaced by the current volume group configuration data on the node to which instant recovery is to be performed and the session is not aborted. When performing instant recovery in an MC/ServiceGuard cluster to some other node than the one that was backed up, select this option. By default, this option is selected.
`Force removal of all replica presentations/-force_prp_replica`	If this option is selected and if any target volume to be restored is presented to any system, the EVA agent removes these presentations. If the option is not selected and a presentation exists, the instant recovery session is aborted. By default, this option is not selected.

Instant Recovery in a Cluster

To perform instant recovery when an application or a filesystem is running in an MC/ServiceGuard or Microsoft Cluster Server on the application system, it is necessary to perform some additional steps. See “Instant Recovery in a Cluster” on page A-20 for the detailed procedure.

Troubleshooting

This section describes the most common problems you may encounter during the instant recovery when using the Data Protector EVA integration.

Before You Begin

- Ensure that the latest official Data Protector patches are installed. Refer to the *HP OpenView Storage Data Protector Software Release Notes* or http://www.openview.hp.com/products/datapro/spec_0001.html for an up-to-date list of supported versions, platforms, and other information.
- Examine system errors reported in `/var/opt/omni/log/debug.log` (UNIX systems) or in `<Data_Protector_home>\log\debug.log` (Windows systems) on the application system and the backup system.

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

Instant Recovery Problems

- **Instant recovery fails if the application system is in cluster.**

Message

```
[Critical]
Data consistency check failed!
Configuration of volume group <vg_name> has changed since the
last backup session!
```

Description

The problem occurs when the `Check data configuration consistency` instant recovery option is selected. It is probably caused by a failover to a secondary node, or the volume group configuration on the application system has changed.

Action

Make sure that the volume group configuration on the application system has not changed and/or deselect the `Check the data configuration consistency` option, and then restart the instant recovery session.

- **On Windows, instant recovery to a different cluster node fails.**

Messages

```
[Major]
Filesystem <volume_name> could not be dismounted from
<drive_letter> ([2] The system cannot find the file specified.).
[Critical]
Failed to disable the application system.
[Critical]
Failed to resolve objects for Instant Recovery.
```

Description

On Windows, the Data Protector automatic preparation of the application system cannot match the clustered volumes from one cluster node to the volumes on the other clustered node.

Action

Disable the automatic preparation of the application system as follows:

1. On the application system, enable the `omnirc` variable `ZDB_IR_MANUAL_AS_PREPARATION`. For information on the variable, see “ZDB Agents Omnirc Variables” on page A-23.
2. Manually dismount all the volumes on the application system that are to be restored.
3. Start the instant recovery session.
4. After instant recovery, manually mount the restored volumes.

III HP StorageWorks Disk Array XP

7 Configuration

In This Chapter

This chapter describes the procedure for configuring Data Protector HP StorageWorks Disk Array XP (XP) integration. It also provides information on the ZDB database.

For a detailed description of the installation of the Data Protector Cell Manager and clients, refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

This chapter includes the following sections:

"Prerequisites and Limitations" on page 139

"Configuring the Integration" on page 142

"ZDB Database—XPDB" on page 145

Prerequisites and Limitations

Prerequisites

- You need a special license to use the Data Protector XP integration. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*.
- The Data Protector XP integration must be correctly installed. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*.
- The same operating system and its version must be installed on the application and the backup system.
- The following XP components are required for this integration:
  - ✓ To configure the integration, RAID Manager XP should be installed on both the application system and the backup system according to the installation instructions provided in the RAID Manager XP documentation.
  - ✓ RAID Manager/LIB XP is required on both the application system and the backup system and should be installed according to the installation instructions provided in the RAID Manager/LIB XP documentation.

    RAID Manager XP and RAID Manager Library are firmware dependant. You should consult your Hewlett-Packard Storage pre-sales consultant or Hewlett-Packard Sales representative for help with determining which version of RAID Manager and RAID Manager Library you should use for your version of firmware.
  - ✓ Business Copy XP or Continuous Access XP microcode and license should be installed.
- To configure the integration, the `xpinfo` utility should be installed on both the application system and the backup system. Contact a Hewlett-Packard support center to obtain this utility.
- You should be familiar with the XP concepts and procedures. Refer to the XP related documentation.
- You should be familiar with the basic ZDB and instant recovery concepts. Refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

- Refer to the *HP OpenView Storage Data Protector Software Release Notes* for information on:
  - ✓ general Data Protector limitations
  - ✓ supported platforms
  - ✓ supported integrations
  - ✓ supported backup topologies
  - ✓ supported connectivity topologies
  - ✓ supported cluster configurations (high availability support)

Limitations

- The BC1 configuration, where the backup and application systems are identical and connected to the same XP, is not recommended. To see on which operating systems the BC1 configuration is supported, check the support matrices in the *HP OpenView Storage Data Protector Software Release Notes*. Only disk image and filesystem backups are possible using the BC1 configuration.
- Asynchronous CA configuration is not supported.
- On HP-UX, the HFS and VxFS filesystems are supported.
- On Windows systems, the NTFS filesystem is supported. Windows Dynamic Disks are not supported.
- On Solaris, the UFS and VxFS filesystems are supported.
- For a split mirror restore, only the BC configuration is supported.
- Database/application restore is not supported for split mirror restore. Only filesystem and disk images can be restored using the Data Protector split mirror restore functionality.
- For an instant recovery, only the BC configuration is supported. To see on which platforms it is supported, check *HP OpenView Storage Data Protector Software Release Notes*.
- Instant recovery from a ZDB-to-disk+tape session cannot be performed using the Data Protector GUI after exporting or overwriting the media used in the backup session. The backup media must not be exported or overwritten even after an object copy session. If the backup media have been exported or overwritten, perform instant recovery using the Data Protector CLI. Refer to “Instant

Recovery Using the Data Protector CLI” on page 202 for more information on how to perform instant recovery using the Data Protector CLI.

- For an instant recovery, if a single or multiple filesystems are being restored, each filesystem must be located on a single disk if LVM is not used, or its logical volumes must be located on an integer multiple of rawdisks. In other words, no object other than those selected for an instant recovery session should share the disks that are used by objects selected for the session. Otherwise, the session will be aborted with the “Instant recovery not possible” message.
- Preview backup is not supported.
- Object copying and object mirroring is not supported for ZDB to disk.

Configuring the Integration

Preparing the Environment

The following sections explain how to prepare the application system and the backup system for use with the XP integration, regardless of whether you plan to use the CA, the BC or the combined CA and BC configuration.

Before You Begin

By now you should have chosen the desired split-mirror backup configuration and connected the XP storage devices to the application and backup systems, and assigned LUNs to the respective ports.

Solaris Systems

On Solaris systems (both on application and on backup systems), the mirrored LDEVs should be labeled and formatted by the Sun format utility. Please refer to the *HP StorageWorks Disk Array XP Operating System Configuration Guide: Sun Solaris* for more information on how to do this.

Depending on the selected backup configuration, the following prerequisites must be fulfilled:

Continuous Access (CA)

The application system should be connected to the Main Control Unit (MCU). The backup system should be connected to the Remote Control Unit (RCU). ESCON links provide communication links between the XP Main and Remote Control Units.

The main LDEVs (P-VOLs) should be connected to the application system and should have paired disks (S-VOLs) assigned. The paired LDEVs (S-VOLs) in the remote disk array should be connected to the backup system.

Business Copy (BC)

The application system and the backup system should be connected to the same XP.

The primary LDEVs (P-VOLs) should be connected to the application system and should have paired disks (S-VOLs) assigned.

First Level Mirrors

If first level mirrors are to be used, the primary LDEVs (P-VOLs) should be connected to the application system and should each have 2-3 paired disks (S-VOLs) assigned.

The mirrored LDEVs (S-VOLs) should be connected to the backup system.

Combined CA and BC

The application system should be connected to the MCU. The backup system should be connected to the RCU.

The main LDEVs (P-VOLs) should be paired to remote volumes in the Remote Control Unit (S-VOLs). The S-VOLs also function as Business Copy primary volumes (P-VOLs). they should be paired to local copies (the BC S-VOLs).

Windows Systems

On Windows systems, only the BC S-VOLs should be connected to the backup system.

HP-UX Systems

On HP-UX systems, it is recommended that only the BC S-VOL be connected to the backup system. If for any reason the CA S-VOL is connected as well, special care must be taken in case the `/etc/lvmtab` is lost in this configuration: `vgscan` can be used to recreate the volume groups, but potentially added `pvlinks` to the S-VOL must be deleted using `vgreduce`. To ensure a correct volume group configuration, the volume group should be re-imported or re-created.

HP-UX LVM Mirroring

It is recommended that the physical volume groups mirroring of LDEVs be used to ensure that each logical volume is mirrored to an LDEV on a different I/O bus. This kind of arrangement is known as **PVG-strict mirroring**. It is assumed that the disk hardware is already configured in such a way that the disk to be used as a mirror copy is connected to the system on a different bus than the bus that is used for the other (primary) copy.

When creating a volume group with LVM mirroring, the volume group must first be created (using the `vgcreate` command) with the LDEV that has its S-VOL assigned. In other words, the LVM mirror primary volume must be the LDEV that has its S-VOL assigned. Then the volume group

must be extended (using the `vgextend` command) with an LDEV that has no S-VOL assigned. In other words, the LVM mirror secondary volume must be the LDEV that has no S-VOL assigned.

For more information on using LVM, refer to the *HP-UX Managing Systems and Workgroups* manual.

Automatic Configuration of Backup System

The XP integration do not require any configuration steps, such as configuring the volume groups and filesystems on the backup system. This is done by Data Protector automatically when a ZDB session is started. Based on the volume group, filesystem, and mount point configuration on the application system, Data Protector creates the same volume group and filesystem structure on the backup system. During the ZDB-to-tape and ZDB-to-disk+tape sessions, Data Protector mounts these filesystems. In case of disk images, raw device files are used. For more information on the backup system mountpoint creation, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

IMPORTANT

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 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”.

ZDB Database—XPDB

The **ZDB database** is in case of the XP integration referred to as **XPDB**. The XPDB keeps information about the backup objects and their mirror configurations and about the XP command devices.

Backup Objects

The objects and their mirror configurations during backup and restore sessions are kept in the XPDB for the purposes of replica set rotation and instant recovery. The XPDB is a set of ASCII files stored on the Cell Manager in the following directory:

- On UNIX: `/var/opt/omni/server/db40/xpdb`
- On Windows: `<Data_Protector_home>\db40\xpdb`

The XPDB stores data about every *split* P-VOL - S-VOL pair used with this integration. Data Protector writes to the XPDB whenever a pair is split using this integration, since the S-VOL in such a pair contains a backup of data. A pair is deleted from the XPDB whenever it is resynchronized using this integration, since when a pair is resynchronized, the prior version of data is overwritten and not accessible.

IMPORTANT Only the LDEV pairs recorded in the XPDB can be used for instant recovery.

The XPDB stores the following information about every split LDEV pair:

- SessionID of the ZDB session that involved the LDEV pair.
- LDEV configuration, volume group configuration and filesystem configuration during the backup session.
- CRC check information calculated during the ZDB session.
- A flag indicating whether the pair can be used for instant recovery purposes or not. The flag is set by the `Track the replica for instant recovery` XP backup option.

IMPORTANT If the flag is not set, a ZDB-to-disk or a ZDB-to-disk+tape session cannot be started.

The volume group configuration and the CRC check information stored in the XPDB is compared to the volume group configuration during the instant recovery session and CRC check calculated during the instant recovery session. If the compared items are not the same, the instant recovery session is aborted.

Command Devices The XP command devices are needed by any process that requires access to the XP.

The information about XP command devices is kept in the XPDB for the purpose of eliminating duplicate instance usage and over allocation. Data Protector provides the following mechanism to prevent duplicate instance usage and over allocation:

- Whenever a session is started, Data Protector queries the XPDB for a list of command devices. If there is none in the XPDB (default behavior when the first session is started), Data Protector identifies command devices and generates a list of command devices in the XPDB that are connected to every application and backup system in the cell.
- Every command device is assigned an instance number (starting from 301) and the system (hostname) having access to it. If a command device can be accessed from more than one system, the hostname identifier enables Data Protector to recognize that the command device is already assigned to be used by another system; the next available instance number is assigned to such a command device-hostname combination.
- When the list is created, every XP attached to the application and backup systems has a list of its command devices and systems having access to them (together with an instance number) assigned.
- Whenever during a session an application or backup system needs access to an XP, it uses the first assigned command device with the instance number from the list. If the command device fails, the next command device from the list assigned to a particular system is used.

If all of them fail, the session fails. If successful, a command device is used by a particular system until the end of the session, and the list of command devices is used for all consecutive sessions.

Additionally, it is possible to specify a particular command device (identified by the XP serial number and LDEV number) to be used by a particular system. Optionally, an instance number can be assigned too. If the instance number is not specified, Data Protector assigns the lowest unassigned instance number.

Below is an example of the command device entries in the XPDB:

```
Serial#  CU:Ldev  (LDEV)   Inst     System
=====================================================
35371    00:67    (103)    301      application.system1.com
35371    00:67    (103)    302      backup.system.com
35372    00:68    (104)    301      application.system2.com
35373    00:69    (105)    301      application.system3.com
```

Refer to “XP Command Device Handling” on page 148 for more information on setting command devices.

Querying the XPDB

The `omnidbxp` command is used for all tasks described in this section. For more information on this command, refer to its man page.

The `omnidbxp` command can be run from any client within the Data Protector cell that has the `User Interface` component installed.

The following is the syntax of the `omnidbxp` command when used to query the XPDB for information on backup objects:

```
omnidbxp [-ir] -session (-list | -show <session_id>)
```

```
omnidbxp [-ir] -ldev (-list | -show <SEQ> <LDEV>)
```

Listing all Available Backup Sessions in the XPDB

To list all available backup sessions in the XPDB, run the following command:

```
omnidbxp -session -list
```

If the `-ir` option is used in the above command, the command lists only those sessions which are marked for instant recovery in the XPDB.

Listing all Backup System LDEVs Involved in a Certain Backup Session

To list all available backup system LDEVs involved in a backup session with a particular sessionID, run the following command:

```
omnidbxp -session -show <session_id>
```

If the `-ir` option is used in the above command, the command lists only those backup system LDEVs involved in a backup session with a particular sessionID that are marked for instant recovery in the XPDB.

Listing all Backup System LDEVs in the XPDB

To list all available backup system LDEVs in the XPDB (to list all backup system LDEVs that are still synchronized), run the following command:

```
omnidbxp -ldev -list
```

If the `-ir` option is used in the above command, the command lists only those backup system LDEVs in the XPDB that are marked for instant recovery in the XPDB.

Listing the XPDB Information about a Certain Pair

To list the complete information stored in the XPDB about a certain pair (sessionID, CRC, IR flag, primary XP #, primary LDEV #, primary port #, mirror type, MU#, date and time, application system and backup system hostnames), run the following command:

```
omnidbxp -ldev -show <SEQ> <LDEV>
```

where *<SEQ>* is the XP frame array number and *<LDEV>* is the backup system LDEV#.

If the `-ir` option is used in the above command, the command lists only the information for the pairs marked for instant recovery that are stored in the XPDB.

XP Command Device Handling

The `omnidbxp` command is used for all tasks described in this section. For more information on this command, refer to its man page.

The `omnidbxp` command can be run from any client within the Data Protector cell that has the `User Interface` component installed.

The following is the syntax of the `omnidbxp` command when used to query the XPDB for information on command devices and to configure the command device usage for this integration:

```
omnidbxp -cm (-add <serial> (<CU:LDEV> | <LDEV>) <hostname>
[<instance>] | -update <serial> (<CU:LDEV> | <LDEV>))
<hostname> [<instance>])
```

```
omnidbxp -cm -remove (all | <serial> [(<CU:LDEV> | <LDEV>)
[<hostname>]])
```

```
omnidbxp -cm -list
```

Listing all Command Devices in the XPDB

To list all command devices in the XPDB, run the following command:

```
omnidbxp -cm -list
```

Adding, Updating and Removing Command Devices Information

To add a command device to the list of command devices in the XPDB, run the following command:

Adding

```
omnidbxp -cm -add <serial> (<CU:LDEV> | <LDEV>) <hostname>
[<instance>]
```

The command will add the command device identified by the serial number of the XP (*<serial>*) and serial number of the command device in the hexadecimal or decimal format (*<CU:LDEV>* | *<LDEV>*) to the XPDB, and assign it the hostname of the system accessing it (*<hostname>*) and optionally the instance number (*<instance>*). If the instance number is not specified, Data Protector assigns the lowest unassigned instance number.

The instance number must be any number in the range between 301 and 399.

NOTE

The command does not check whether the specified command device or system exist, it only checks if the optional instance number specified is within the correct range and if the command device together with the instance number is not already assigned to be used by some other system. If checks fail, the command fails with an error message.

Updating

To update the information about a command device in the XPDB, run the following command.

```
omnidbxp -cm -update <serial> (<CU:LDEV> | <LDEV>)
<hostname> [<instance>]
```

The command will update the XPDB information about the command device identified by the serial number of the XP (*<serial>*), the serial number of the command device in hexadecimal or decimal format (*<CU:LDEV>* | *<LDEV>*) and the specified hostname of the system accessing it (*<hostname>*), by assigning the newly specified instance number (*<instance>*) to the XP serial number, serial number of command device and hostname combination. If the instance number is not specified, Data Protector assigns the lowest unassigned instance number.

The instance number must be any number in the range between 301 and 399.

NOTE The command does not check whether the specified command device or system exist, it only checks if the optional instance number specified is within the correct range and if the command device together with the instance number is not already assigned to be used by some other system. If checks fail, the command fails with an error message.

Removing

To remove the information about all command devices, about command devices within a specific XP, or about a specific command device from the XPDB, run one of the following commands:

```
omnidbxp -cm -remove all
```

```
omnidbxp -cm -remove <serial>
```

```
omnidbxp -cm -remove <serial> (<CU:LDEV> | <LDEV>)
[<hostname>]
```

The first command will remove the information about all command devices from the XPDB.

The second command will remove the information about command devices within a specific XP identified by the serial number of this XP (*<serial>*).

The third command will remove the information about command devices identified by the serial number of the XP (*<serial>*), the serial number of command device in hexadecimal or decimal format (*<CU*:*LDEV>* | *<LDEV>*) and optionally by the (*<hostname>*).

IMPORTANT

When removing the information about the command device without specifying the system (*<hostname>*), the command deletes all entries for the specified command device, regardless of the system(s) assigned to it.

XP LDEV Exclude File

It is possible to disable Data Protector from using certain LDEVs on the backup system. Thus, it is possible to reserve certain LDEVs for other purposes than Data Protector backup and restore. A Data Protector session is aborted if the replica set involved in the session contains an excluded LDEV.

The list of disabled mirrors is kept in the Data Protector HP StorageWorks Disk Array XP LDEV exclude file (XP LDEV exclude file) on the Cell Manager in the following directory:

- On UNIX: `/var/opt/omni/server/db40/xpdb/exclude/XPexclude`
- On Windows: *<Data_Protector_home>*`\db40\xpdb\exclude\XPexclude`

The mirrors set in this file must be the backup system LDEVs identified by the backup system LDEV#.

Syntax

The following is the syntax of the XP LDEV exclude file:

#HP OpenView Storage Data Protector A.05.50

#HP StorageWorks Disk Array XP LDEV Exclude File

#

#[*<XP1 >*]

#*<LDEV>*

#*<LDEV1>*, *<LDEV2>*, *<LDEV3>*

#*<LDEV1>*-*<LDEV2>*

```
#[<XP2 >]
#…
#
#<XP> - disk array serial/sequence number
#<LDEV> - CU#:LDEV number in decimal format
#
#End of file
```

Example

The following is an example of the XP LDEV exclude file:

```
#HP OpenView Storage Data Protector A.05.50
#HP StorageWorks Disk Array XP LDEV Exclude File
#
[35241]
3603, 3610, 3620-3625 # SAP R/3 data
2577 # Oracle archive logs
2864-3527 # Oracle database
#End of file
```

Manipulating the XP LDEV Exclude File

The `omnidbxp` command is used for all tasks described in this section. For more information on this command, refer to its man page.

The `omnidbxp` command can be run from any client within the Data Protector cell that has the `User Interface` component installed.

The following is the syntax of the `omnidbxp` command when used to manipulate the XP LDEV exclude file:

```
omnidbxp -exclude (-put <filename> | -get <filename> |
-check <SEQ> <LDEV> | -init | -delete)
```

Setting and Changing the XP LDEV Exclude File

1. Run the following command:

   ```
   omnidbxp -exclude -get <filename>
   ```

   This command reads the XP LDEV exclude file from the Cell Manager and saves it as *<filename>*.

2. Edit the copied XP LDEV exclude file from the previous step and save it when you are done editing.
3. Run the following command:

   `omnidbxp -exclude -put <filename>`

   This command reads the contents of *<filename>*, checks its syntax and if the syntax is correct copies the file to its position on the Cell Manager.

Identifying Excluded LDEVs

The following command checks whether a certain LDEV, identified by its backup system disk array serial/sequence number, and LDEV number is specified in the XP LDEV exclude file on the Cell Manager:

`omnidbxp -exclude -check <SEQ> <LDEV>`, where:

<SEQ> is the backup system disk array serial/sequence number and

<LDEV> is the S-VOL CU#:LDEV in decimal format.

If the queried LDEV is specified in the XP LDEV exclude file, the command returns: `YES!`

If the queried LDEV is not specified in the XP LDEV exclude file, the command returns: `NO!`

Resetting the XP LDEV Exclude File

The following command overwrites the current XP LDEV exclude file on the Cell Manager with the template XP LDEV exclude file:

```
omnidbxp -exclude -init
```

Deleting the Contents of the XP LDEV Exclude File

The following command empties the contents of the XP LDEV exclude file on the Cell Manager:

```
omnidbxp -exclude -delete
```

8 Backup

In This Chapter

This chapter describes how to configure a filesystem or disk image backup of your data residing on HP StorageWorks Disk Array XP (XP). The sections describe steps for configuring a ZDB using the Data Protector Graphical User Interface.

This chapter includes the following sections:

"Backup Process" on page 157

"Configuring a Backup Specification" on page 164

"Backup Options" on page 169

"Troubleshooting" on page 178

Backup Process

If you are not acquainted with the general ZDB and instant recovery concepts, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

ZDB Types

Three types of ZDB sessions are possible using the XP integration:

- **ZDB to disk**

  ZDB to disk is possible only using the BC configuration.

  With this type of backup, mirrors are created, and data is kept on a disk array until reused according to the replica set rotation. Data from the replica (all S-VOLs that are created during one backup session) is not moved to backup media (for example, tape media). A replica created in a ZDB-to-disk session can be used for instant recovery and is part of the replica set.

  Such a backup session is performed when the `Track the replica for instant recovery` backup option is selected when creating a backup specification, and the `To disk` option is selected when running or scheduling a backup.

- **ZDB to tape**

  With this type of backup, mirrors are created, and data from the replica is moved to backup media.

  If the backup option `Keep the replica after the backup` is selected, the replica remains on a disk array until reused in the next split mirror backup session using the same LDEV pairs, but cannot be used for instant recovery and is not part of the replica set.

  If the backup option `Keep the replica after the backup` is not selected, the replica is synchronized with the original after the backup.

  Such a backup session is performed when the `Track the replica for instant recovery` backup option *is not* selected when creating a backup specification.

- **ZDB to disk+tape**

  ZDB to disk+tape is possible only using the BC configuration.

  With this type of backup, mirrors are created, and data is kept on a disk array until reused according to the replica set rotation. Data is also moved to backup media. The replica created in a ZDB-to-disk+tape session can be used for instant recovery and is part of the replica set.

  Such a backup session is performed when the `Track the replica for instant recovery` backup option is selected when creating a backup specification, and the `To disk+tape` option is selected when running or scheduling a backup.

IMPORTANT Before creating a backup specification, consider all limitations regarding the Data Protector XP integration. For information, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

XP Backup Flow

General Concept The XP split mirror backup can be described as follows:

- In the first phase, if the application system data is not yet synchronized to the backup system, it is synchronized to the backup system.

  During this phase, the synchronization is performed on the level of the participating volume groups (UNIX systems) or disks (Windows systems). Therefore, if multiple filesystems or disk images are configured in the same volume group (UNIX systems) or on the same disk (Windows systems), the *whole* volume group or disk (all filesystems or disk images in this volume group or on disk) is synchronized to the backup system regardless of the filesystems, disk images or application objects selected for backup.

- In the second phase, the synchronized backup system data is backed up to a backup device.

  During this phase, only the filesystems, disk images or application objects selected for backup are backed up to a backup device.

NOTE

In the case of ZDB to disk, the second phase does not occur. Data backed up using the ZDB-to-disk functionality can be restored only by using the instant recovery functionality.

IMPORTANT

Such a concept enables a restore of selected objects (filesystems or disk images) for a split mirror restore and for a restore from backup media on LAN (filesystems, disk images or application objects), but not for an instant recovery.

With an instant recovery, the links from the application to the backup system are *not* synchronized before the restore, whereas with a split mirror restore they *are*, thus enabling the restore of selected objects by establishing the current state of the application system data on the backup system, and then restoring the selected objects to the backup system and finally resynchronizing the backup system to the application system.

Detailed Flow

When a backup is started, the following happens:

- The Data Protector Backup Session Manager (BSM) starts the Data Protector HP StorageWorks XP Agent (XP Agent or SSEA) on both the application and backup systems.
- The BSM sends information about what needs to be backed up to the XP Agent. The XP Agent identifies the LDEVs for the data to be backed up (resolve process).
- If a replica set is defined, the replica set rotation is applied to the specified replica set to select the LDEVs for use in the backup session. Refer to *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide* for more information on the replica set and on the replica set rotation.
- If the `Abort the session if the mirror disks are not synchronized` option is selected and LDEVs are not synchronized, the session is aborted.

If one of the following options is selected:

- ✓ Prepare/resync the mirror disks at the start of the backup
- ✓ Force resync at the start of the backup session,

the pairs are incrementally synchronized. Before the links are synchronized, filesystems on the backup system are dismounted. On HP-UX and Solaris systems, volume groups or disk groups are deactivated.

- If a script/command is set by the Stop/quiesce the application option, this script/command is executed on the application system. The script can perform various tasks, such as stopping an application.
- If the Dismount Filesystems on the Application System option is selected, the filesystems on the application system are dismounted before the split and remounted after the split.
- The pairs are split. After the split operation, the pairs are in PSUS/SSUS mode.
- If a script/command is set by the Restart the application option, this script/command is executed on the application system. The script can perform various tasks, such as restarting an application.

  By default, if the Stop/quiesce the application script/command fails, the Restart the application script/command is not executed, so a cleanup procedure should be implemented in the Stop/quiesce the application script/command.

  If the ZDB_ALWAYS_POST_SCRIPT omnirc variable is set to 1, the restart application script/command is always executed if set. By default the variable is set to 0. Refer to “ZDB Agents Omnirc Variables” on page A-23 for more information on XP omnirc variables.
- The backup system is prepared for backup. The mountpoint for the backed up filesystem is created on the backup system as described in *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*. On HP-UX and Solaris systems, the backup volume groups (HP-UX systems) or backup disk groups (Solaris systems) are activated. The filesystems are then mounted.
- If the option Enable the backup system in read/write mode is selected on HP-UX and Solaris systems, the volume group is activated in read/write mode (HP-UX systems) and the filesystem is

mounted in read/write mode (HP-UX and Solaris systems). In this case, the filesystem check is executed before the filesystems are mounted.

- The backup system is backed up.

NOTE With ZDB to disk, the backup system is not backed up.

- If the `Leave the backup system enabled` option is not selected, the filesystems on the backup system are dismounted; on HP-UX and Solaris systems, volume groups (HP-UX systems) or disk groups (Solaris systems) are deactivated. On HP-UX systems, the volume groups are removed from `/etc/lvmtab`. Note that on HP-UX systems, the volume groups are not removed from `/etc/lvmtab` if the `SSEA_BACKUP_VG_EXIST` `omnirc` variable is set to `1`. Refer to "ZDB Agents Omnirc Variables" on page A-23 for more information on XP `omnirc` variables.

  If this option is selected, the filesystems are left mounted; on HP-UX and Solaris systems, volume groups/disk groups remain active and on HP-UX systems, the volume groups are not removed from `/etc/lvmtab`.

- If the `At the end of the backup, prepare/resync the mirror disks for the next backup` option is selected the following is possible:

  ✓ If one or no replica is set by the `MU Number(s)` option, the pairs are left split if the `Keep the replica after the backup` option is selected, or resynchronized if the `Keep the replica after the backup` option is not selected.

  ✓ If more than one replica is set by the `MU Number(s)` option (a replica set is defined), the pairs used in the current session are left split regardless of the `Keep the replica after the backup` option selection and the replica set rotation scheme is applied to the specified replica set to select the replica for use in the next backup session; the selected pairs are resynchronized.

- If any of the described operations fail, the BSM aborts the backup session and stops the session according to the specified XP options. For information on XP options, refer to "Backup Options" on page 169.

Figure 8-1 **Filesystem Split Mirror Backup Flow**

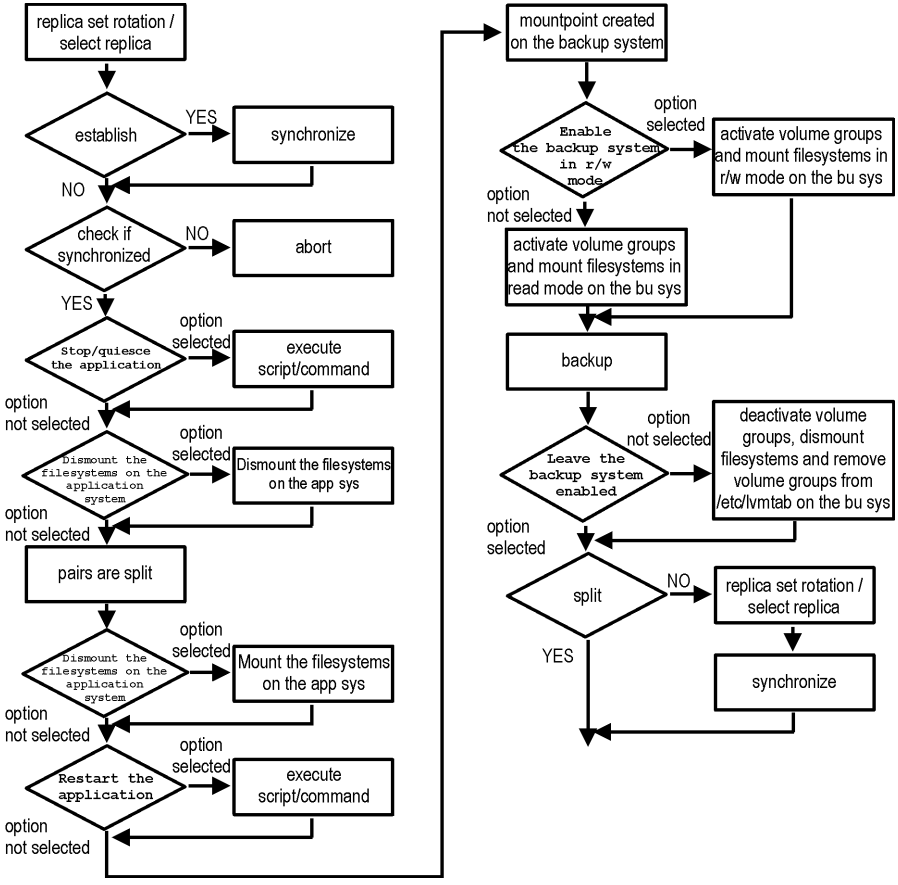

NOTE During a ZDB-to-disk session, the data is not moved to backup media, as presented in Figure 8-1 on page 162.

In Figure 8-1 on page 162, the “establish” and “split” checks are dependent on the following XP backup option selections:

Table 8-1

`At the end of the backup, prepare/resync the mirror disks for the next backup`	split = NO

Table 8-1

`Force resync at the start of the backup session`	establish = YES
`Abort the session if the mirror disks are not synchronized`	establish = NO
`Prepare/resync the mirror disks at the start of the backup`	split = YES establish = YES
`MU Number(s)` is set to `1`, `2`, `0` respectively or left empty; or `Keep the replica after the backup` option is selected	split = YES

NOTE The simultaneous selection of options in the first and the last row of Table 8-1 is conflicting. In such a situation, the “split” check is set to YES.

Configuring a Backup Specification

Use the Data Protector GUI to create a filesystem or disk image backup specification for use with the XP integration.

A ZDB-to-disk backup specification is configured in the same way as a ZDB-to-tape or a ZDB-to-disk+tape backup specification.

1. In the Context List, select `Backup`.
2. In the Scoping Pane, expand `Backup` and then `Backup Specifications`. Right-click `Filesystem` (for both filesystem backup and disk image backup) and click `Add Backup`.

   The `Create New Backup` dialog box is displayed.

   In the `Filesystem` box, select the `Blank Filesystem Backup` template. For more information on templates, refer to the *HP OpenView Storage Data Protector Administrator's Guide*.

   Select the backup type `Split mirror backup` and the sub type `HP StorageWorks XP`. See online Help for a description of other options.

   Click `OK`.

Figure 8-2 Create New Backup Dialog Box

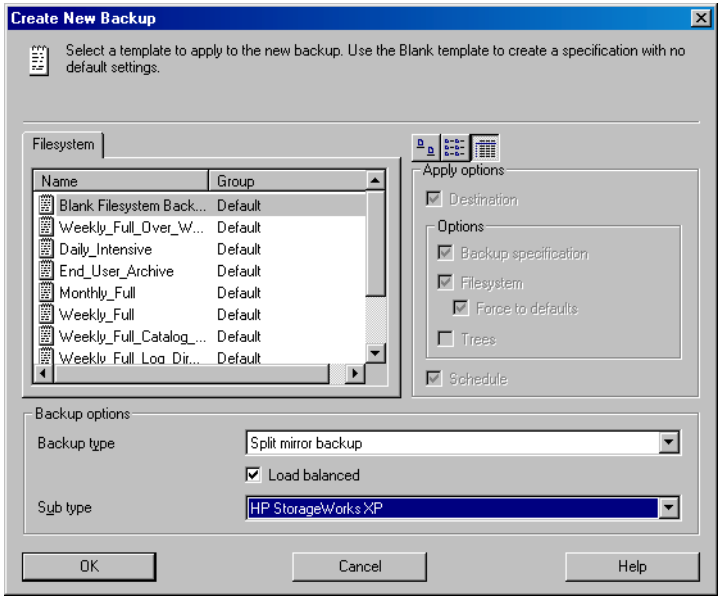

3. Under `Client systems`, select the application system and the backup system.

   Under `Mirror type`, specify the XP mirror configuration you will use for the backup.

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

   For detailed information on all options, refer to “Backup Options” on page 169.

   Click `Next`.

4. Depending on whether you perform a filesystem or disk image backup, proceed as follows:

Filesystem Backup

If you are configuring a filesystem backup, expand the application systems that contain the objects that you want to back up and then select what you want to back up.

IMPORTANT

On UNIX, if you intend to perform instant recovery, *select all filesystems inside the volume group* to be backed up. Otherwise, instant recovery will not be possible using the Data Protector GUI or (if you perform instant recovery using the Data Protector CLI) data can be corrupted.

Click `Next`.

Disk Image Backup

If you are configuring a disk image backup, click `Next`.

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

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

   For detailed information on the object mirror functionality, see *HP OpenView Storage Data Protector Administrator's Guide*.

NOTE

Object mirroring is not supported for ZDB to disk.

Click `Next`.

6. Under `Backup Specification Options`, click `Advanced` and then the `HP StorageWorks XP` tab to open the XP backup options.

   Here, you can specify the XP `Application options`, and modify all the other options, except the `Application system` and the `Backup system` options. Click `OK` when you are finished modifying the options.

   For information on XP backup options refer to "Backup Options" on page 169.

   For information on `Filesystem Options`, see online Help.

Figure 8-3 XP Backup Options

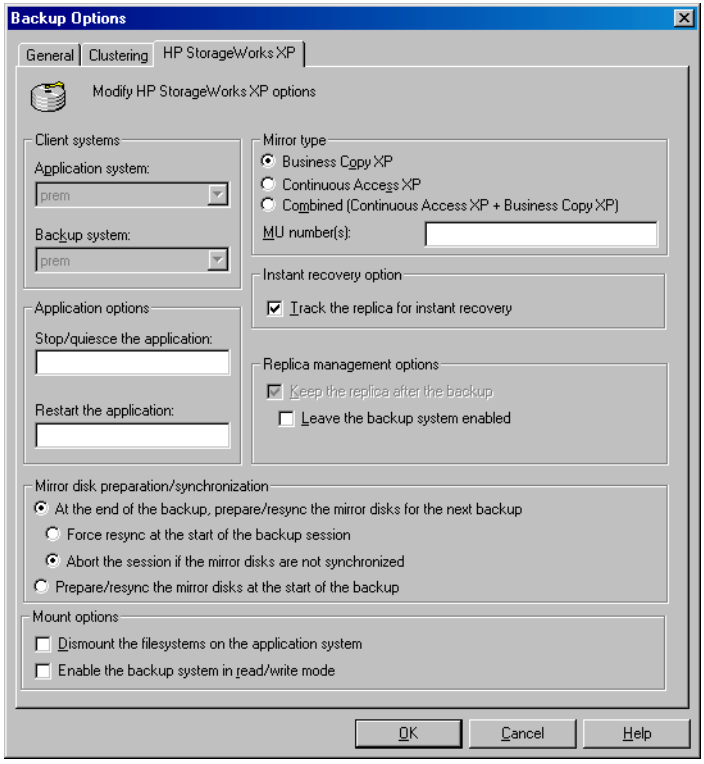

7. Follow the backup wizard to open the `Schedule` page and `Backup Object Summary` page, where a summary of the backup specification is given.
8. Depending on whether you perform a filesystem or disk image backup, proceed as follows:

Filesystem Backup

If you are configuring a filesystem backup, click `Next`.

Disk Image Backup

For a disk image backup, proceed as follows:

a. Click `Manual add` to add the disk image objects you want to back up.

b. Select `Disk image object` and click `Next`.

c. Select the client to be backed up and click `Next`.

d. Follow the wizard to specify the `General Object Options` and the `Advanced Object Options`. For more information on these options, press **F1**.

e. In the `Disk Image Object Options` window, specify the disk image sections you want to back up.

On UNIX Systems

To specify a rawdisk section, use the following format:

`/dev/rdsk/<filename>`, for example: `/dev/rdsk/c2t0d0`

To specify a raw logical volume section, use the following format:

`/dev/vg<number>/rlvol<number>`, for example: `/dev/vg01/rlvol1`

IMPORTANT

If you intend to perform instant recovery, *specify all raw logical volumes inside the volume group* to be backed up. Otherwise, instant recovery will not be possible using the Data Protector GUI or (if you perform instant recovery using the Data Protector CLI) data can be corrupted.

On Windows Systems

Use the following format:

`\\.\PHYSICALDRIVE#`,

where `#` is the current number of the disk you want to back up.

For example: `\\.\PHYSICALDRIVE3`

For information on how to find the current numbers of the disks (physical drive numbers) you want to back up, refer to the online Help index keyword “disk image backups”.

f. Click `Finish` and then `Next`.

9. Save your backup specification. For information on starting and scheduling the backup session, refer to “Running and Scheduling a ZDB Session” on page A-3.

Backup Options

The XP backup options are grouped in six sets. These sets are described on the following pages. Refer to “XP Integration” on page A-39 to help you understand these options.

TIP The XP `Mirror disk preparation/synchronization` options set defines whether the replica set rotation is applied or not.

Client Systems Options

Using this set of options, you can specify the application and the backup systems to be involved in the backup session.

Table 8-2 Client Systems Options

Data Protector GUI	Function
`Application system`	Specify the application system on which the application runs (for example an Oracle database). In a cluster environment, specify the virtual server hostname (rather than the physical node hostname).
`Backup system`	Specify the backup system on which your data is to be backed up. In a cluster environment, specify the virtual server hostname (rather than the physical node hostname).

Mirror Type Options

Using this set of options, you can specify the mirror type of the XP configuration and the replica set or the replica to be used in the backup session.

Table 8-3 Mirror Type Options

`Mirror Type`	Specify the XP mirror configuration you will use for the backup: Business copy XP, Continuous Access XP, or Combined (Continuous Access XP + Business Copy XP).

`MU number(s)`	This option is enabled only if the Business copy XP `Mirror Type` option is selected.
	
	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.
	
	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:
	
	`1`
	
	`1-2`
	
	`2,0,1`
	
	If the sequence is specified, it does not set the order in which the replicas are used. They are used according to the algorithm described in *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.
	
	If a range is entered, it must be specified in ascending order.
	
	If this option is not specified, the MU number 0 is set.

Application Options

Using this set of options, you can stop any application that is not integrated with the XP integration on the application system mirrored disks before the split operation, and restart this application after the split operation.

TIP This set of options can also be used for any other operation on the application system before and after the split operation.

NOTE The XP `Application options` are accessible from the `Backup Specification Options` group box (refer to step 6 on page 166), by clicking the `Advanced` tab and then the `HP StorageWorks XP` tab.

Table 8-4 Application Options

`Stop/quiesce the application`	Specify the optional stop/quiesce the application command/script. Create this command/script in the `/opt/omni/lbin` (HP-UX or Solaris systems) or in the *<Data_Protector_home>*`\bin` (Windows systems) directory on the application system and specify only the filename in the backup specification. The `Stop/quiesce the application` command/script is executed on the application system before the links are split. It can be used, for example, to stop the application. If the `Stop/quiesce the application` command/script fails, the `Restart the application` command/script is not executed, so a cleanup procedure should be implemented in the `Stop/quiesce the application` command/script. If the `ZDB_ALWAYS_POST_SCRIPT omnirc` variable is set to `1`, the `Restart the application` script/command is always executed if set. By default the variable is set to `0`. Refer to "ZDB Agents Omnirc Variables" on page A-23 for more information on XP `omnirc` variables.
`Restart the application`	Specify the optional restart the application command/script. Create this command in the `/opt/omni/lbin` (HP-UX or Solaris systems) or in the *<Data_Protector_home>*`\bin` (Windows systems) directory on the application system and specify only the filename in the backup specification. The `Restart the application` command/script is executed on the application system immediately after the links are split. It can be used, for example, to restart the application.

Instant Recovery Option

Using this option, you can mark a replica for instant recovery.

Table 8-5 Instant Recovery Option

`Track the replica for instant recovery`	This option is enabled only if the Business copy XP `Mirror Type` option is selected. Select this option to enable ZDB to disk or ZDB to disk+tape and instant recovery from the replica. If this option is not set, it is not possible to perform ZDB to disk or ZDB to disk+tape and instant recovery from the replica. However, this option does not influence the replica set rotation. By default, this option is selected.

CAUTION If you selected the `Track the replica for instant recovery` option, do not manually resynchronize the affected mirrors afterwards, otherwise instant recovery will not be possible.

Replica Management Options

Using this set of options, you can control the split/resync behavior of the mirror disks, the mounting of backup system filesystems and activation of backup system volume/disk groups, all after the backup session has finished.

Table 8-6 Replica Management Options

`Keep the replica after the backup`	This option is selected and disabled if the `Track the replica for instant recovery` option is selected. If this option is selected, 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. If this option is not selected, the disks involved in the backup session are resynchronized after the backup session, only if one or no replica is set by the `MU Number(s)` option. If more than one replica is set by the `MU Number(s)` option, the disks involved in the backup session will remain split after the backup session. By default, this option is selected.

Table 8-6 Replica Management Options

`Leave the backup system enabled`	By default, Data Protector dismounts the filesystems (all platforms), deactivates the volume/disk groups (HP-UX and Solaris systems) and removes the volume groups from `/etc/lvmtab` (HP-UX systems) on the backup system after each backup. Note that on HP-UX systems, the volume groups are not removed from `/etc/lvmtab` if the `SSEA_BACKUP_VG_EXIST` `omnirc` variable is set to `1`. Refer to "ZDB Agents Omnirc Variables" on page A-23 for more information on XP `omnirc` variables. If this option is selected, the filesystems remain mounted (all platforms), volume/disk groups remain activated (HP-UX and Solaris systems) and volume groups are not removed from `/etc/lvmtab` (HP-UX systems) after the backup. Thus, you can use the backup system for some data warehouse activity afterwards, *but not for instant recovery*. This option is available only if the `Keep the replica after the backup` option is selected. By default, this option is not selected.

CAUTION

By selecting the `Leave the backup system enabled` option you will not be able to use the affected replica for the instant recovery purposes unless you also select the `Track the replica for instant recovery` option.

Mirror Disk Preparation/Synchronization Options

This set of options controls whether the synchronization is done immediately after the backup or before the next backup.

Table 8-7 Mirror Disk Preparation/Synchronization Options

`At the end of the backup, prepare/resync the mirror disks for the next backup`	If this option is selected, the next replica is prepared according to the replica set rotation (resynchronized with the P-VOLs) for the next backup session at the end of the current backup session. By default, this option is selected. If this option is not selected, the next two options are disabled.
`Force resync at the start of the backup session`	If this option is selected, a resync will be initiated before the backup. By default, this option is not selected.
`Abort the session if the mirror disks are not synchronized`	If the mirror disks are not synchronized at the start of the backup, the backup is aborted. By default, this option is selected.
`Prepare/resync the mirror disks at the start of the backup`	If this option is selected, the mirror disks in the replica selected for the current backup session are resynchronized with the P-VOLs at the start of the current backup session. By default, this option is not selected.

Mount Options

Using this set of options, you can force the application system's mirrored filesystems to flush the filesystem cache to the application system mirrored disks, and/or you can force the backup system volume/disk groups and filesystems to be activated in read/write mode.

Table 8-8 Mount Options

`Dismount the filesystems on the application system`	If this option is selected, the filesystem on the application system is dismounted before the split and mounted after the split. Sometimes, this is the only way to ensure that the data on the filesystem is consistent. A filesystem does not have the stop I/O functionality to flush the data from the filesystem cache to the disk and stop the I/O for the time of the split. The only way to back up a filesystem in split mirror mode is to dismount the mount point on the application system. If an integrated application (for example, Oracle or SAP) runs on the filesystem, it controls the I/O to the disk. In this case it is not necessary to dismount the filesystem before the split. By default, this option is not selected.
`Enable the backup system in read/write mode`	This option is related to HP-UX and Solaris systems only, since on Windows systems filesystems are always mounted in read/write mode. By default, the volume/disk groups and filesystem on the backup system are activated and mounted in read-only mode. If this option is selected, the volume/disk groups and filesystem on the backup system will be activated in read/write mode. Specify this option if you want to have read/write access to the volume/disk groups and filesystems on the backup system. For backup, it is sufficient to activate the backup system volume/disk groups and filesystem in read-only mode. However, if you want to perform other tasks after the backup has been performed, this might not be sufficient. By default, this option is not selected.

IMPORTANT

If you selected the `Enable the backup system in read/write mode` option and you intend to use the replica involved for instant recovery, the replica will include all modifications that you may have made on the backup system while offline.

Business Copy/1

The BC1 configuration, where the backup and application systems are identical and connected to the same XP, is not recommended. To see on which operating systems the BC1 configuration is supported, check the support matrices in the *HP OpenView Storage Data Protector Software Release Notes*. Only disk image and filesystem backups are possible using the BC1 configuration.

Use the `Stop/quiesce the application` command/script on the application system to create a consistent replica of your application data on the disk (for example, either put your application in halt mode or shut down the application).

Troubleshooting

This section describes the most common problems you may encounter when using the Data Protector XP integration.

Before You Begin

- Ensure that the latest official Data Protector patches are installed. Refer to the *HP OpenView Storage Data Protector Software Release Notes* or http://www.openview.hp.com/products/datapro/spec_0001.html for an up-to-date list of supported versions, platforms, and other information.
- Examine system errors reported in `/var/opt/omni/log/debug.log` (HP-UX or Solaris systems) or in `<Data_Protector_home>\log\debug.log` (Windows systems) on the application system and the backup system.
- Ensure that the RAID Manager/LIB XP is correctly installed on both the application system and the backup system. If the RAID Manager/LIB XP is correctly installed, the command `/usr/omni/bin/ssea` should return the following output:

  HP OpenView Storage Data Protector A.05.50: `SSEA, internal build <build_ID>, built on <date>`

  and the command `ls <RMLIB_home>/libsvrrm.sl` (HP-UX systems) or `ls <RMLIB_home>/libsvrrm.so` (Solaris systems) should return the following output:

  `<RMLIB_home>/libsvrrm.sl` (HP-UX systems)

  `<RMLIB_home>/libsvrrm.so` (Solaris systems)

  To check whether the RAID Manager/LIB XP is installed on Windows systems, check the contents of the `<RMLIB_home>` directory. It should contain the `libsvrrm.dll` file.

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

Creation of the Backup Specification

- **You cannot select the StorageWorks mode in the Data Protector user interface when attempting to create a backup specification.**

Action

Check whether the `HP StorageWorks XP Agent` software component (XP component) is installed on both the application and backup systems.

Check the Data Protector `cell_info` file on the Data Protector Cell Manager, in *<Data_Protector_home>*`\Config\server\Cell\cell_info` (Windows Cell Manager) or in `/etc/opt/omni/server/cell/cell_info` (UNIX Cell Manager) to see if the XP component is installed.

An exemplary file should look similar to the following:

```
-host "hpsap001.bbn.hp.com" -os "hp s800 hp-ux-11.00" -cc
A.05.50 -da A.05.50 -ssea A.05.50

-host "hpsap002.bbn.hp.com" -os "hp s800 hp-ux-11.00" -cc
A.05.50 -da A.05.50 -ma A.05.50 -ssea A.05.50
```

Backup Problems

- **The XP Agent on the application system failed to dismount a filesystem.**

Action

Check the `Stop/quiesce the application` script. In this script, stop all processes using the filesystem.

- **The XP Agent on the backup system failed to mount a filesystem.**

Action

Check that the mountpoint directory has been created on the backup system.

- **The XP Agent failed to synchronize the pairs, so the split failed.**

Action

To successfully split the pair, XP Agent first checks the status of the pairs. The pairs can only be split (in PSUS/SSUS status) after all pairs have been synchronized (in PAIR status). XP Agent checks the

status of links after every 2 seconds and retries 10 times. To change these values, change the variables SSEA_SYNC_RETRY=*<number of retries>* and SSEA_SYNC_SLEEP_TIME=*<sleep time in seconds>*.

Increase the time frame for synchronization by setting the variables SSEA_SYNC_RETRY and SSEA_SYNC_SLEEP_TIME in the file /opt/omni/.omnirc (HP-UX or Solaris systems) or in the *<Data_Protector_home>*\omnirc (Windows systems) file on the backup system.

Refer to "ZDB Agents Omnirc Variables" on page A-23 for more information on these variables.

- **The XP Agent reports that a P-VOL has no paired S-VOL.**

Action

Check the XP configuration.

If a Business Copy configuration is used, then all P-VOLs on the application system must have associated BC S-VOLs on the backup system.

If a Continuous Access configuration is used, then all P-VOLs on the application system must have associated CA S-VOLs on the backup system.

If a Combined (CA+BC) configuration is used, then all P-VOLs on the application system must have associated CA S-VOLs on the backup system and all S/P-VOLs must have BC S-VOLs.

- **The XP Agent reports an invalid pair state on LDEVs.**

Action

Check the state of the link. If the link is split, use the Prepare/resync the mirror disks at the start of the backup option.

Try to configure and start RAID Manager XP instances manually. You can get a list of LDEVs from the backup session report.

- **The XP Agent reports missing details for a specific LDEV/MU# and then fails the session:**

Message

```
[Warning] From: SSEA@machine_app.company.com ""  Time:
17.10.2001. 10:41:27

Failed to get a BC pair for LDEV 55, MU# 1 in RAID 35371.

(Details unknown.)
```

```
[Normal] From: SSEA@machine_app.company.com ""  Time:
17.10.2001. 10:41:27

Resolving of backup objects on the application system
completed.

[Normal] From: SSEA@machine_bu.company.com ""  Time:
17.10.2001. 10:41:27

Resolving backup objects on the backup system.

[Critical] From: SSEA@machine_bu.company.com ""  Time:
17.10.2001. 10:41:29

Resolving of backup objects on the backup system failed.
```

Action

In the backup specification, specify an existing and configured LDEV/MU# on the backup system, or make sure that the LDEV/MU# referred to in the output is not set in the XP LDEV exclude file.

Restart the session.

9 Restore

In This Chapter

This chapter describes how to configure and run a filesystem or disk image restore of your data backed up using the Data Protector Data Protector HP StorageWorks Disk Array (XP) integration. The sections describe steps for performing a restore using the Data Protector Graphical User Interface and the Data Protector Command Line Interface.

This chapter includes the following sections:

"Overview" on page 185

"Restoring from Backup Media on LAN" on page 186

"Split Mirror Restore" on page 189

"Instant Recovery" on page 197

"Troubleshooting" on page 205

Overview

There are three methods of restoring data:

- Restore from backup media on LAN. Refer to “Restoring from Backup Media on LAN” on page 186.
- Split mirror restore. Only BC and BC1 split mirror restore is supported in this release. Refer to “Split Mirror Restore” on page 189.
- Instant recovery. This is possible only with the BC and BC1 configuration. Refer to “Instant Recovery” on page 197.

Restoring from Backup Media on LAN

Data is restored directly to the application system. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for more information on various restore options available. The restore procedure described here provides only a general description of how to restore an object that was backed up using the split mirror backup functionality from backup media on LAN.

TIP

You can improve the data transfer rate when restoring by connecting the backup device to the application system and configuring it using the Data Protector GUI. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for more information on configuring backup devices. Refer to the "Restoring Under Another Device" section of the same guide for more information on how to perform a restore using another device.

1. In the `Context List`, select `Restore`.
2. Select the objects for restore and click them to display their properties in the Scoping Pane.

   Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for more information on various restore options available. Be sure to make the following selections in the Scoping Pane in order to restore an object that was backed up using the split mirror backup functionality from backup media on LAN:

   - Select the application system as the `Target client` under the `Destination` tag.

Figure 9-1 Selecting the Application System

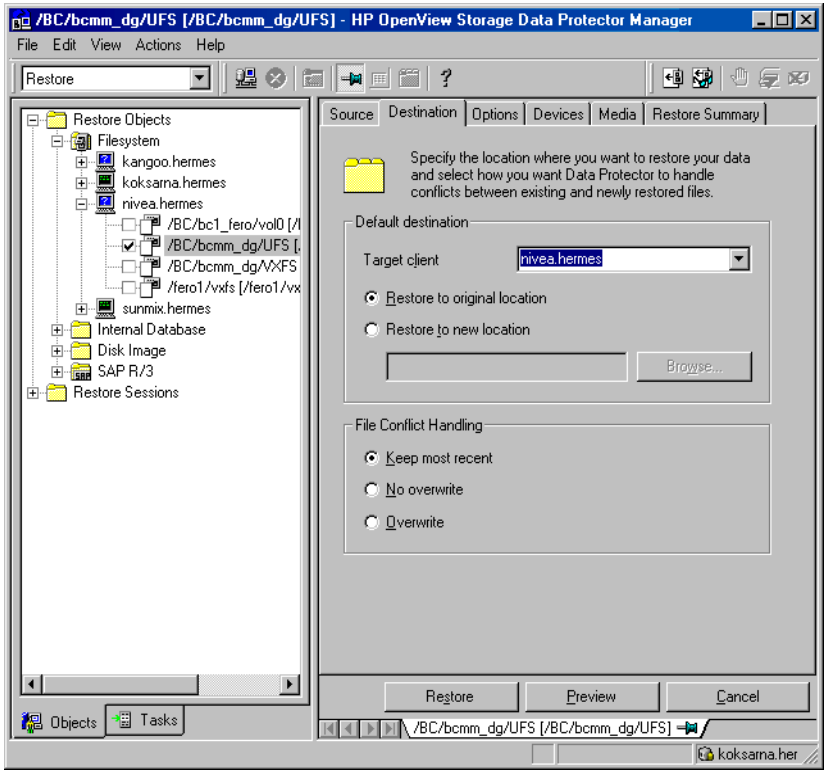

3. After you have set the restore options, click the `Restore` button. The `Start Restore Session` dialog box is displayed.
4. Click `Next` to specify the report level and network load. Click `Next`.
5. This step and figure Figure 9-2 on page 188 are relevant only if you have the EMC Symmetrix software component installed on the application system. Select `StorageWorks XP restore`. Click `Next` to display the `Start Restore Session` dialog window.

Figure 9-2 Selecting the XP Restore

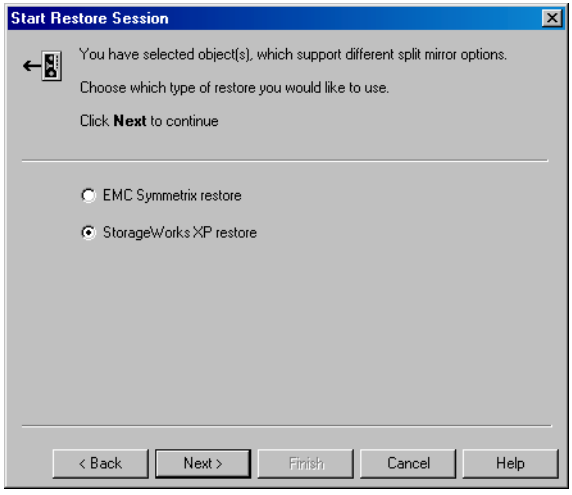

6. In the `Start Restore Session` dialog window, select `Disabled` in the `Mirror mode` drop-down list. This sets a direct restore to the application system.

Figure 9-3 Selecting a Restore from Backup Media on LAN

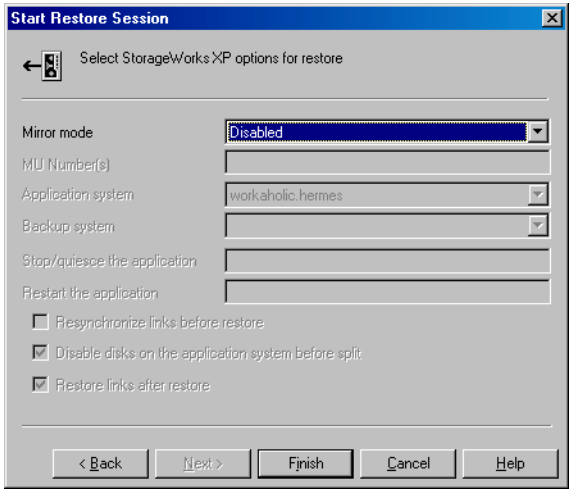

7. Click `Finish` to start the restore session.

Split Mirror Restore

It is possible to restore filesystems and disk images using the split mirror restore only in a BC configuration.

Data is restored from backup media on LAN to the mirror LDEVs (S-VOLs) and then moved to the original LDEVs (P-VOLs) using a BC configuration. The procedure consists of the following automated steps:

1. If a replica set is defined, the replica set rotation is applied to the specified replica set to select the replica to use in the restore session. Refer to *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide* for more information on the replica set and the replica set rotation.
2. Preparing the backup and application systems for restore.
3. Restoring data from backup media on LAN to the backup system and synchronizing this data to the application system.

Split Mirror Restore Process

Refer to Figure 9-4 on page 190 for an overview of the filesystem split mirror restore flow.

Figure 9-4 Filesystem Split Mirror Restore Flow

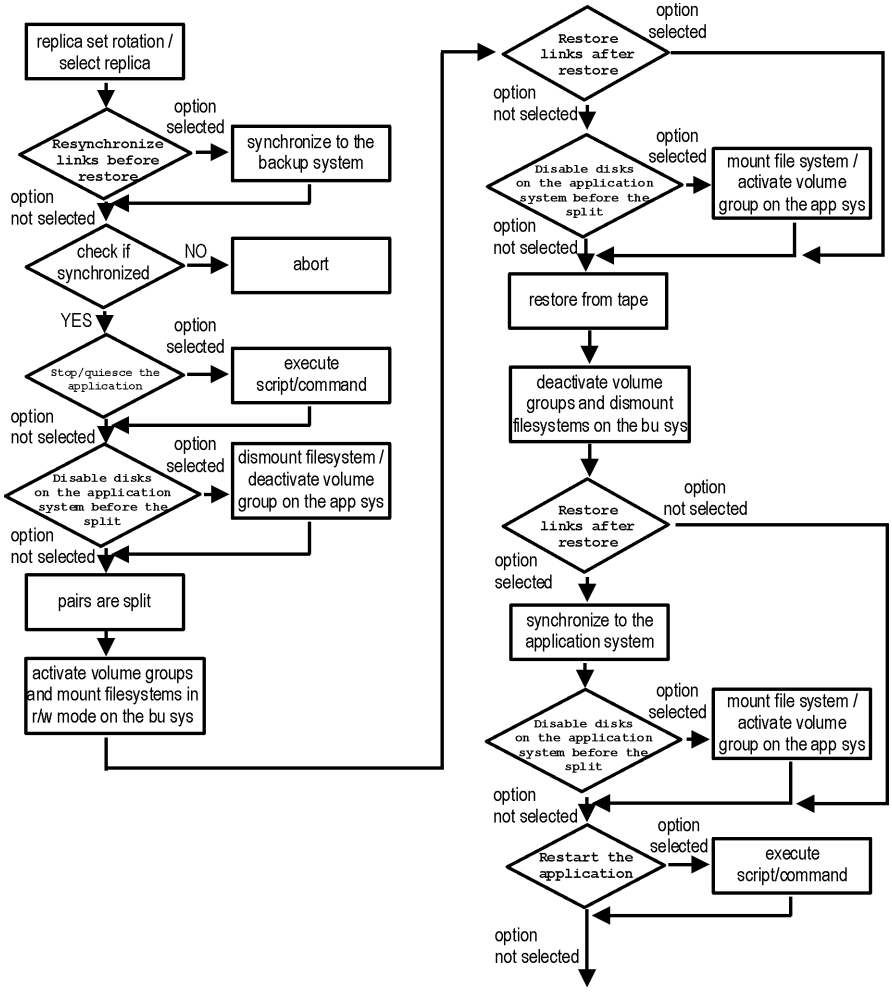

Split Mirror Restore Procedure

Follow the procedure below to perform a BC split mirror restore of a filesystem or a disk image. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for more information on the general restore process.

1. In the `Context List`, select `Restore`.

2. Select the objects for restore and click them to display their properties in the Scoping Pane.

NOTE To perform a split mirror BC restore, you need to select the application system as the `Target client` under the `Destination` tag. If the backup system is selected, a normal restore to the backup system is performed.

3. 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 Figure 9-5 on page 191 are relevant only if you have the EMC Symmetrix software component installed on the *target* client. Select `StorageWorks XP restore`. Click `Next`.

Figure 9-5 Selecting the XP Restore

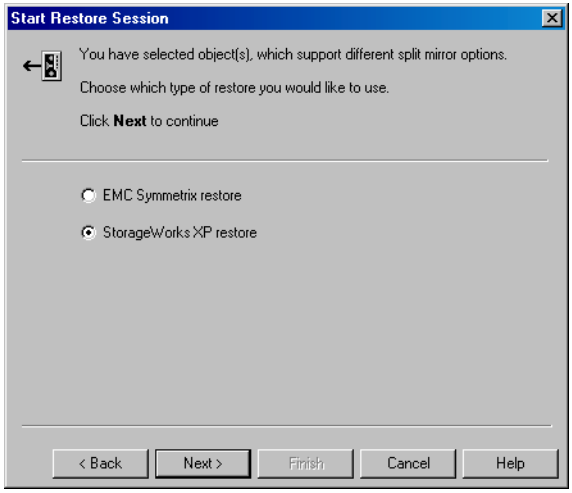

6. Specify the XP split mirror restore options. Refer to "XP Split Mirror Restore Options" on page 193 for more information on these options.

Figure 9-6 XP Split Mirror Restore Options

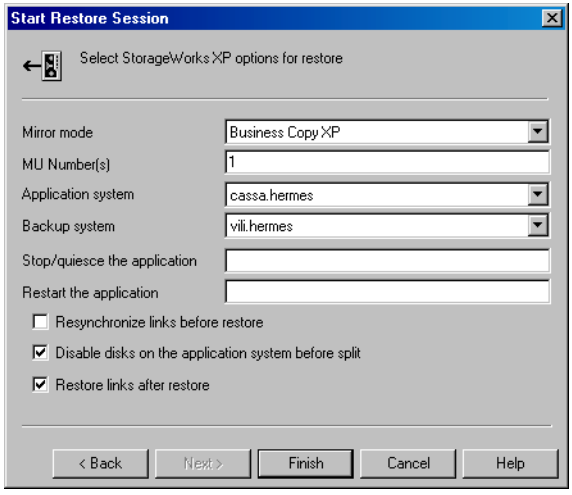

7. Click `Finish` to start the restore session.

IMPORTANT It is not possible to start split mirror backup, split mirror restore or instant recovery sessions using the same disk on the application system at the same time. A split mirror backup, split mirror restore or instant recovery session must be started only after the preceding session using the same disk on the application system has finished the synchronization operation, otherwise the session will fail.

NOTE If the LVM Mirroring configuration is used, a warning message is issued in the Data Protector monitor during the restore since the volume group LDEVs in the physical volume group on the application system do not have their BC pairs assigned. This warning message should be ignored.

Split Mirror Restore Options

Table 9-1 on page 193 shows the XP split mirror restore options, along with their descriptions:

Table 9-1 XP Split Mirror Restore Options

Data Protector GUI	Function
Mirror mode	Specify the type of XP configuration used for the restore. Only the BC configuration is supported for split mirror restore.
MU Number(s)	Specify a specific replica to be used in the restore session, or a replica set from which the integration, according to the replica set rotation scheme, selects one replica to be used in the restore session. If this option is not specified, the MU# 0 is set.
Application system	Specify the application system on which your data is to be restored. In a cluster environment, specify the virtual server hostname (rather than the physical node hostname).
Backup system	Specify the backup system. In a cluster environment, specify the virtual server hostname (rather than the physical node hostname).

Table 9-1 XP Split Mirror Restore Options

Data Protector GUI	Function
`Stop/quiesce the application`	Specify the optional stop/quiesce the application command/script. Create this command in the `/opt/omni/lbin` (HP-UX or Solaris systems) or in the *<Data_Protector_home>*`\bin` (Windows systems) directory of the application system. The `Stop/quiesce the application` command/script is executed on the application system before the links are split. It can be used, for example, to stop the application or to dismount the filesystem that is not to be restored in the active session and is mounted to the same volume or disk group in which the filesystem that is to be restored is mounted. If the `Stop/quiesce the application` command/script fails, then the `Restart the application` command/script is not executed, so a cleanup procedure should be implemented in the `Stop/quiesce the application` command/script. If the `ZDB_ALWAYS_POST_SCRIPT` `omnirc` variable is set to `1`, the `Restart the application` script/command is always executed if set. By default the variable is set to `0`. Refer to "ZDB Agents Omnirc Variables" on page A-23 for more information on XP `omnirc` variables.
`Restart the application`	Specify the optional restart the application command/script. Create this command/script in the `/opt/omni/lbin` (HP-UX or Solaris systems) or in the *<Data_Protector_home>*`\bin` (Windows system) directory of the application system. The `Restart the application` command/script is executed on the application system immediately after the links are resynchronized. It can be used, for example, to restart the application or to mount a filesystem.
`Resynchronize links before restore`	This option is used to synchronize pairs, that is, to move data to backup disks. This is necessary to prepare the disks for restore and to enable accurate restores. If the pairs were split before the restore, and only some files need to be restored, then this option can be used to update the backup system. This will ensure that the correct data is resynchronized to the application system. This option is not selected by default.

Table 9-1 XP Split Mirror Restore Options

Data Protector GUI	Function
`Disable disks on the application system before split`	Disks on the application system are disabled, that is, the filesystems are dismounted (HP-UX and Solaris systems) and volume groups are deactivated (HP-UX systems). This is performed before the split. The disks are enabled after the links are restored. Note that only filesystems selected for restore are dismounted. If other filesystems exist in the volume or disk group, a `Stop/quiesce the application` command/script and `Restart the application` command/script must be created to dismount these filesystems. You must always select this option for restore when you want to move data from the backup to the application system, that is, to incrementally restore links. The application system disks have to be disabled to provide data integrity after the links are restored, that is, data is moved.
`Restore links after restore`	With this option enabled, the XP Agent incrementally restores the links for LDEVs that were successfully restored by Data Protector from the backup media on LAN. The XP Agent also incrementally re-establishes the links for LDEVs that were not successfully restored by Data Protector from the backup media on LAN.

Split Mirror Restore in a Cluster

To perform a Data Protector split mirror restore when a filesystem is running in an MC/ServiceGuard or a Microsoft Cluster Server on the application system, it is necessary to perform some *additional* steps.

MC/ServiceGuard Procedure

To perform a Data Protector split mirror restore when a filesystem is running in an MC/ServiceGuard on the application system, it is necessary to stop the filesystem cluster package in order to be able to activate the mirrored volume groups on the application system in normal mode before the split mirror restore session is started.

Follow the procedure below to perform a Data Protector split mirror restore to the application system in an MC/ServiceGuard cluster:

1. Stop the filesystem cluster package:

   ```
   cmhaltpkg <app_pkg_name>
   ```

   This will stop the filesystem services and dismount the mirrored volume group filesystem.

2. Deactivate the mirrored volume group from the cluster mode and activate it in the normal mode:

   ```
   vgchange -c n /dev/<mirror_vg_name>
   ```

   ```
   vgchange -q n -a y /dev/<mirror_vg_name>
   ```

3. Mount the mirrored volume group filesystem:

   ```
   mount /dev/<mirror_vg_name>/<lv_name> /<mountpoint>
   ```

4. Using Data Protector, start the Data Protector split mirror restore session. Refer to “Split Mirror Restore Procedure” on page 190 for a detailed procedure.

IMPORTANT When specifying the application system during the split mirror restore procedure, make sure to specify the hostname of the application system *node* on which the mirrored volume group was activated in the normal mode when performing step 2 of this procedure.

5. When the split mirror restore session has finished, dismount the mirrored volume group filesystem:

   ```
   umount /<mountpoint>
   ```

6. Deactivate the mirrored volume group in the normal mode and activate it in the cluster mode:

   ```
   vgchange -a n /dev/<mirror_vg_name>
   ```

   ```
   vgchange -c y /dev/<mirror_vg_name>
   ```

7. Start the filesystem cluster package:

   ```
   cmrunpkg <app_pkg_name>
   ```

Instant Recovery

Data Protector instant recovery takes advantage of split mirror technology to provide instant data restore.

Using the Data Protector XP integration, only the data backed up in ZDB-to-disk and ZDB-to-disk+tape sessions can be restored in instant recovery sessions.

Instant recovery can be performed using the Data Protector GUI or using the Data Protector CLI.

Data Protector instant recovery moves data from the backup to the application system directly, without involving a backup device. At the beginning of the instant recovery session, the application system needs to be disabled. This includes dismounting the filesystems and deactivation of volume/disk groups (HP-UX and Solaris only). Before this can be done, the status of participating filesystems and volume/disk groups (HP-UX and Solaris only) is checked and only the mounted filesystems and activated volume/disk groups are dismounted and deactivated. At the end of the session, only those filesystems that were previously dismounted will be mounted, and only those volume groups that were previously deactivated will be activated.

IMPORTANT

After the instant recovery session, the configuration of the restored filesystems is the same as it was at the backup time. This means that the restored filesystems are mounted to the same mount points or drive letters as they were at the backup time. In case these mount points or drive letters have some other filesystems mounted, these filesystems are automatically dismounted before instant recovery and the restored filesystems are mounted after instant recovery.

Up to three replicas can exist at the same time. XP allows up to three replicas, and each can have additional two copies if cascading is used. Refer to *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

Limitations

Data Protector can only use first level mirrors for backup and instant recovery purposes. The additional six (cascading) copies are not supported. Instant recovery is only possible using the HP StorageWorks BusinessCopy XP configuration (BC and BC1 configurations). Refer to “Prerequisites and Limitations” on page 139 for platform and backup limitations with instant recovery.

IMPORTANT

The Data Protector instant recovery functionality does not provide recovery of a database or an application. It only synchronizes the application system LDEVs from their mirrors on the backup system. To recover a database or an application, additional database or application recovery-related steps must be performed.

XPDB

The objects and their mirror configurations during backup and restore sessions are kept in the XPDB for the purpose of replica set rotation and instant recovery. Refer to “ZDB Database—XPDB” on page 145 for more information on the XPDB.

During an instant recovery, the data in the specified replica (left unchanged for the purpose of instant recovery) is synchronized to the application system disks without restoring from backup media. Prior to the instant recovery, Data Protector *can* check the following:

- volume group configuration and
- CRC of the mirror copy

These checks assure that data in the replica was left intact since the backup session. If any of the checks fail, the session fails.

Once the replica is restored, it can be left unchanged or it can be resynchronized, depending on the instant recovery options. Refer to “Instant Recovery Options” on page 203 for more information on instant recovery options.

Instant Recovery Process

When an instant recovery is started, the following happens:

- The Restore Session Manager (RSM) starts the XP Agent. The XP Agent reads the information about the selected LDEV pairs from the XPDB.

If the `Check the data configuration consistency` instant recovery option was selected, the following takes place:

✓ The volume group configuration for the selected LDEV pairs stored in the XPDB is compared to the current volume group configuration. If the items compared do not match, the session is aborted.

When instant recovery is performed in an MC/ServiceGuard cluster to some other node than the one that was backed up, the current volume group configuration on the node to which instant recovery is to be performed is different from the volume group configuration stored in the XPDB. In such a case the XPDB volume group configuration data is replaced by the current volume group configuration data on the node to which instant recovery is to be performed and the session is not aborted.

✓ The CRC check information for the selected LDEV pairs stored in the XPDB is compared to the current CRC check information. If the items compared do not match, the session is aborted.

- The selected LDEV pairs are checked if they are synchronized. If the pairs are synchronized, the session is aborted.
- The application and backup systems are disabled by deactivating the participating volume/disk groups (HP-UX and Solaris systems) and dismounting the participating filesystems.
- The P-VOLs belonging to the selected LDEV pairs are synchronized from their paired SVOLs.
- If the `Keep the replica after the restore` instant recovery option was selected, the pairs involved are split.
- The application system is enabled by activating the participating volume/disk groups (HP-UX or Solaris systems) and mounting the participating filesystems.

Instant Recovery Procedure

To perform a filesystem or disk image instant recovery using the Data Protector GUI, follow the procedure described in “Instant Recovery Using the Data Protector GUI” on page 200. To perform a filesystem or disk image instant recovery using the Data Protector CLI, follow the procedure described in “Instant Recovery Using the Data Protector CLI” on page 202.

Prerequisite

If you have performed a disk image backup, you must manually dismount the disks to be restored before instant recovery, and re-mount them afterwards.

NOTE

When performing an instant recovery, a specific ZDB-to-disk or ZDB-to-disk+tape session rather than a specific replica is selected in the Data Protector User Interface.

Instant Recovery Using the Data Protector GUI

To perform instant recovery using the Data Protector GUI, proceed as follows:

1. In the `Context List`, select `Instant Recovery`.

   You can select a specific ZDB-to-disk or ZDB-to-disk+tape session by expanding the `Restore Sessions` item in the Scoping Pane—a list of all ZDB-to-disk and ZDB-to-disk+tape sessions sorted by date is displayed, or you can select a specific ZDB-to-disk or ZDB-to-disk+tape session as follows:

   a. Expand the `Restore Objects` item in the Scoping Pane - a list of backed up objects types (Filesystem, Disk Image, SAP R/3, MS SQL Server,...) is displayed in the Scoping Pane.

   b. Expand the type of object you want to restore—a list of all split mirror backup specifications for this backup type is displayed in the Scoping Pane.

   c. Expand the backup specification item containing the object you want to restore—a list of all ZDB-to-disk and ZDB-to-disk+tape sessions for this backup specification (sorted by date) is displayed in the Scoping Pane.

Figure 9-7 ZDB-to-Disk or ZDB-to-Disk+Tape Session Selection

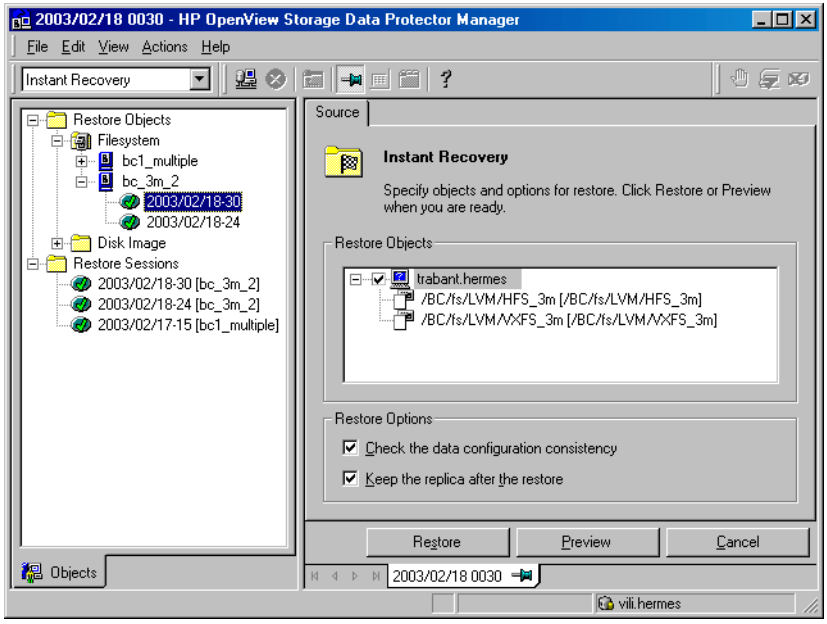

2. In the Scoping Pane, click the ZDB-to-disk or ZDB-to-disk+tape session you want to restore—the application system and its mountpoints (UNIX systems) or drive letters (Windows systems) that were backed up on the application system during the selected ZDB-to-disk or ZDB-to-disk+tape session are displayed in the Results Area.
3. Select the application system to be restored and specify the XP instant recovery options. Refer to “Instant Recovery Options” on page 203 for more information on these options.
4. Click the `Restore` button to start the instant recovery or the `Preview` button to preview (filesystem backup only) the instant recovery.
5. Select `Start Restore Session` to start the instant recovery session or select `Start Preview Session` to start the preview (filesystem restore only) of the instant recovery session. Click `OK`.

IMPORTANT

It is not possible to start split mirror backup, split mirror restore or instant recovery sessions using the same disk on the application system at the same time. A split mirror backup, split mirror restore or instant recovery session must be started only after the preceding session using the same disk on the application system has finished the synchronization operation, otherwise the session will fail.

Instant Recovery Using the Data Protector CLI

The replica to be restored using the Data Protector CLI is identified by the ZDB-to-disk or ZDB-to-disk+tape session ID. To perform instant recovery using the Data Protector CLI, proceed as follows:

1. Get a list of all available ZDB-to-disk or ZDB-to-disk+tape sessions, identified by the session ID, using the following command:

   ```
   omnidbxp -ir -session -list
   ```

   From the output of the command, select the backup session you want to restore.

2. Execute the following command:

   ```
   omnir -host <application_system_name> -session
   <SessionID> -instant_restore [<INSTANT RECOVERY OPTIONS>]
   ```

   Where:

   - *<application_system_name>* is the hostname of the application system
   - *<SessionID>* is the backup session ID selected in the step 1 of this procedure.

   For *<INSTANT RECOVERY OPTIONS>* see Table 9-2 on page 203.

   This will start the instant recovery session.

Refer to the `omnidbxp` and `omnir` man pages for more information on these commands.

Instant Recovery Options

XP instant recovery options are set during the configuration of an instant recovery session.

Table 9-2 XP Instant Recovery Options

Data Protector GUI	Function
`Check the data configuration consistency/-check_config`	If this option is selected, the current volume group configuration of the participating volume groups is compared with the volume group configuration during the ZDB-to-disk or ZDB-to-disk+tape session kept in the XPDB. If the volume group configuration has changed since the ZDB-to-disk or ZDB-to-disk+tape session, the session is aborted. When instant recovery is performed in an MC/ServiceGuard cluster to some other node than the one that was backed up, the current volume group configuration on the node to which instant recovery is to be performed is different from the volume group configuration kept in the XPDB. In such a case the XPDB volume group configuration data is replaced by the current volume group configuration data on the node to which instant recovery is to be performed and the session is not aborted. When performing instant recovery in an MC/ServiceGuard cluster to some other node than the one that was backed up, select this option. The CRC check information for the selected LDEV pairs stored in the XPDB is compared to the current CRC check information. If the items compared do not match, the session is aborted. By default, this option is selected.
`Keep the replica after the restore/-keep_version`	If this option is selected, the LDEV pairs involved in the current instant recovery session are split after the session is finished. By default, this option is selected.

Instant Recovery and LVM Mirroring

If one of the supported LVM Mirroring configurations is used, the logical volume containing LVM mirroring disks must first be reduced (using the `lvreduce -m 0` command) by excluding the LVM mirror disk (the one without the S-VOL assigned) from the logical volume. Only after this step can an instant recovery be performed.

This assures that data recovered using the instant recovery functionality does not get overwritten by the version of data on the LVM mirror disk (the one without the S-VOL assigned).

After instant recovery is performed, the logical volume containing LVM mirroring disks must be extended (using the `lvextend -m` command) with the LVM mirror disk (the one without the S-VOL assigned) that was previously excluded from the logical volume.

Instant Recovery in a Cluster

To perform instant recovery when an application or a filesystem is running in an MC/ServiceGuard or Microsoft Cluster Server on the application system, it is necessary to perform some additional steps. See “Instant Recovery in a Cluster” on page A-20 for the detailed procedure.

Troubleshooting

This section describes the most common problems you may encounter during the split mirror restore or instant recovery when using the Data Protector XP integration.

Before You Begin

- Ensure that the latest official Data Protector patches are installed. Refer to the *HP OpenView Storage Data Protector Software Release Notes* or http://www.openview.hp.com/products/datapro/spec_0001.html for an up-to-date list of supported versions, platforms, and other information.
- Examine system errors reported in `/var/opt/omni/log/debug.log` (UNIX systems) or in `<Data_Protector_home>\log\debug.log` (Windows systems) on the application system and the backup system.

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

Split Mirror Restore Problems

- **The XP Agent reports that a filesystem cannot be dismounted and that the corresponding volume group cannot be deactivated, and then fails the session:**

Message

```
[Major] From: SSEA@machine.company.com ""  Time: 17.10.2001.
11:06:46

Filesystem /dev/bc_nested/hfs could not be dismounted from

/BC/fs/HFS

/usr/sbin/vgchange -a n /dev/bc_nested

[Major] From: SSEA@machine.company.com ""  Time: 17.10.2001.
11:06:47

[224:8] Volume group /dev/bc_nested could not be
deactivated.
```

Action

Make sure that the filesystem/volume group is not in use (for example, you are positioned in the filesystem mountpoint directory) and then restart the session.

- **The XP Agent reports that a LDEV pair is in "STAT_COPY" state at the beginning of the split mirror restore session and then fails the split mirror restore session:**

Message

```
[Critical] From: SSEA@machine.company.com ""  Time:
16.10.2001. 17:25:00

The following BC pairs have an invalid status for the
requested operation:

SEQ#    LDEV         Port   TID  LUN  MU#  Status     SEQ#
LDEV

-----------------------------------------------------

35371   00A8h ( 168)  CL1-D    1    3    0  STAT_COPY  35371
01A5h ( 421)

35371   00A8h ( 168)  CL1-D    1    3    0  STAT_COPY  35371
01A6h ( 422)

   -----------------------------------------------------

[Critical] From: SSEA@machine.company.com ""  Time:
16.10.2001. 17:25:00

Failed to resolve objects for Instant Recovery.
```

Action

Wait until the LDEV pair is in "PAIR" or "PSUS/SSUS" status and then restart the split mirror restore session.

Instant Recovery Problems

- **The XP Agent reports that a LDEV pair is in "STAT_COPY" or "STAT_PAIR" state at the beginning of the instant recovery session and then fails the instant recovery session:**

Message

```
[Critical] From: SSEA@machine.company.com ""  Time:
16.10.2001. 17:25:00

The following BC pairs have an invalid status for the
requested operation:

SEQ#    LDEV         Port   TID  LUN  MU#  Status     SEQ#
LDEV
```

```
-----------------------------------------------------

35371   00A8h ( 168)  CL1-D    1    3    0  STAT_COPY  35371
01A5h ( 421)

35371   00A8h ( 168)  CL1-D    1    3    0  STAT_COPY  35371
01A6h ( 422)

-----------------------------------------------------

[Critical] From: SSEA@machine.company.com ""  Time:
16.10.2001. 17:25:00

Failed to resolve objects for Instant Recovery.
```

Action Wait until the LDEV pair is in “PAIR” or “PSUS/SSUS” status and then restart the instant recovery session.

- **When performing an instant recovery to the application systems in a cluster with the `Check the data configuration consistency` option selected, the XP Agent reports that the configuration of a volume group involved has changed since the last backup session and then fails the instant recovery session:**

```
[Critical] From: SSEA@machine.company.com ""  Time:
11/8/2001 11:43:09 AM

Data consistency check failed!

Configuration of volume group /dev/vg_sap has changed since
the last backup session!
```

Description The problem was probably caused by a failover of a secondary node, or something has changed in the volume group on the application system.

Action Make sure that the volume group configuration on the application system has not changed and/or deselect the `Check the data configuration consistency` option and restart the instant recovery session.

IV EMC Symmetrix

10 Configuration

In This Chapter

This chapter describes the procedure for configuring Data Protector EMC Symmetrix (EMC) integration.

For a detailed description of the installation of the Data Protector Cell Manager and clients, refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

This chapter includes the following sections:

"Prerequisites and Limitations" on page 213

"Configuring the Integration" on page 214

NOTE

The steps in this chapter require specific information regarding disk identification. You will also need to check the disk configuration (SRDF or TimeFinder, for example). Instructions on how to obtain the necessary information, are provided in "EMC - Obtaining Disk Configuration Data" on page A-42.

Prerequisites and Limitations

Prerequisites

- You need a special license to use the Data Protector EMC integration. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*.
- The same operating system and its version must be installed on the application and the backup system.
- The Data Protector EMC integration must be correctly installed. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*.
- EMC Solution Enabler must be installed and enabled.
- You should also be familiar with the following:
  - ✓ Basic ZDB concepts. Refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.
  - ✓ The EMC Command-Line Interface
  - ✓ Logical Volume Manager concepts
- Refer to the *HP OpenView Storage Data Protector Software Release Notes* for information on:
  - ✓ general Data Protector limitations
  - ✓ supported platforms
  - ✓ supported integrations
  - ✓ supported backup topologies
  - ✓ supported connectivity topologies
  - ✓ supported cluster configurations (high availability support)

Limitations

- The EMC Multi-BCV environments are not supported.
- ZDB to disk, ZDB to disk+tape, and instant recovery are not supported. Only ZDB to tape is supported.
- Preview backup is not supported.

Configuring the Integration

Preparing the Environment

The following sections explain how to prepare the application system and the backup system for use with the Data Protector EMC integration, regardless of whether you plan to use the TF, the SRDF or the Combined configuration.

Before You Begin

By now you should have chosen the desired split-mirror backup configuration and connected the EMC storage devices to the application and backup systems.

Depending on the selected backup configuration, the following prerequisites must be fulfilled:

Symmetrix Remote Data Facility (SRDF) The application system should be connected to the Application (R1) Symmetrix. The backup system should be connected to the Backup (R2) Symmetrix.

The main Source (R1) Devices should be connected to the application system and should have paired disks assigned. The paired Target (R2) Devices in the remote disk array should be connected to the backup system.

Time Finder (TF) The application system and the backup system should be connected to the same EMC disk array.

The Standard Devices should be connected to the application system and should have paired disks assigned.

The BCV Devices should be connected to the backup system.

Combined (SRDF + TimeFinder) The application system should be connected to the Application (R1) Symmetrix. The backup system should be connected to the Backup (R2) Symmetrix.

The main Source (R1) Devices should be paired to Target (R2) Devices in the Backup (R2) Symmetrix. The Backup (R2) Symmetrix Target (R2) Devices also functions as TimeFinder Standard Devices. They should be paired to BCV (R2) Devices.

It is recommended that only the TimeFinder BCV (R2) Devices be connected to the backup system. If for any reason the SRDF Target (R2) Devices are connected as well, special care must be taken in case the `/etc/lvmtab` is lost in this configuration: `vgscan` can be used to recreate the volume groups, but potentially added `pvlinks` to the SRDF Target (R2) Devices must be deleted using `vgreduce`. To ensure a correct volume group configuration, the volume group can be re-imported or re-created.

Creating the EMC Symmetrix Database File

Before You Begin

Ensure that the EMC Solution Enabler software is installed on the application and backup systems.

The EMC Symmetrix database file contains the physical configuration information of SCSI parameters that define your entire storage complex. It is located in:

`/var/symapi/db/symapi_db.bin` (HP-UX systems)

`<symapi_home>\db\symapi_db.bin` (Windows systems)

Typically, when your configuration has changed or if this is your first command-line session, you should rebuild the EMC Symmetrix database file with the most complete and current information concerning all physical devices connected via the SCSI buses to your system.

To scan the hardware and rebuild the database, enter:

```
symcfg discover
```

This command scans all SCSI buses on the System, not only those connected to EMC units. This can take a significant amount of time to complete.

You can display what is in the EMC Symmetrix database file with the commands:

- `syminq -sym`, which displays all EMC devices.
- `symbcv list dev`, which lists all BCV devices that are configured on the EMC unit.
- `symrdf list`, which lists all RDF disk devices known to the system.

Refer to “EMC - Obtaining Disk Configuration Data” on page A-42 for more information on these commands.

Creating the Data Protector EMC Database File

The Data Protector EMC database file is essentially the same as the EMC Symmetrix database file. It is used only by the EMC Agent (SYMA) to store configuration information. Create this database file before you start running Data Protector backups as well as each time your disk configuration has changed.

Alternatively, you can also enable the `Run discovery of Symmetrix environment` backup option in the backup specification. However, this operation is time consuming because it checks disk configuration through low-level SCSI commands. We recommend that you create the Data Protector EMC database file on both the application and backup systems using the EMC Agent as follows:

HP-UX

```
/opt/omni/lbin/syma -init
```

This command creates the following Data Protector EMC database file on both the application and backup system:

```
/var/opt/omni/client/emc/symm.bin
```

Windows

```
<Data_Protector_home>\bin\syma -init
```

This command creates the following Data Protector EMC database file on both the application and backup system:

```
<Data_Protector_home>\Config\Client\EMC\symm.bin.
```

Automatic Configuration of Backup System

The EMC integration do not require any configuration steps, such as configuring the volume groups and filesystems on the backup system. This is done by Data Protector automatically when a ZDB session is started. Based on the volume group, filesystem, and mount point configuration on the application system, Data Protector creates the same volume group and filesystem structure on the backup system. During the ZDB sessions, Data Protector mounts these filesystems. In case of disk images, raw device files are used. For more information on the backup system mountpoint creation, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

The Data Protector EMC Log File

Data Protector logs all information about objects, devices, and device groups into log files in the:

`/var/opt/omni/tmp/emc` (HP-UX systems)

or

<Data_Protector_home>`\Config\client\tmp\emc` (Windows systems)

directory on the application system and backup system. The log files are named by the session as `R1_`*<session_name>*`.log` or `R2_`*<session_name>*`.log` where *<session_name>* is composed of the sessionID, the forward slashes "/" having been replaced with dashes "-." For example:

```
R1_1999-09-13-3.log
```

```
R2_1999-09-13-3.log
```

The log includes the following information:

- Information on the resolved EMC configuration (mapping to EMC devices).
- Information about created and deleted device groups, as well as devices added to device groups.
- Information about operations on device groups, like splitting links, incremental establish, incremental restore, etc.
- Status information about backup and restore objects.

We recommend that you check both log files if the backup did not finish successfully or in case of any other problems. The logs can also be useful if you leave the links split after a backup or restore session.

11 Backup

In This Chapter

This chapter describes how to configure a filesystem or disk image backup of your data residing on EMC Symmetrix (EMC). The sections describe steps for configuring a ZDB using the Data Protector Graphical User Interface.

This chapter includes the following sections:

"Backup Process" on page 221

"Configuring a Backup Specification" on page 225

"Backup Options" on page 229

"Backup Disk Usage" on page 231

"Testing Your Backed Up Data" on page 232

"Troubleshooting" on page 235

Backup Process

If you are not acquainted with the general ZDB concepts, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

ZDB Types

Only the ZDB to tape is possible using the Data Protector EMC integration.

ZDB to Tape

With ZDB to tape, mirrors are created, and data from the replica is moved to backup media.

If the backup option `Re-establish links after backup` is not selected, the replica remains on a disk array until reused in the next split mirror backup session using the same EMC device pairs.

If the backup option `Re-establish links after backup` is selected, the replica is synchronized with the original after the backup.

EMC Backup Flow

The Data Protector EMC split mirror backup can be described as follows:

- In the first phase, if the application system data is not yet synchronized to the backup system it is synchronized to the backup system.

  During this phase the synchronization is performed on the level of participating volume groups (HP-UX systems) or disks (Windows systems). Therefore, if multiple filesystems or disk images are configured in the same volume group (HP-UX systems) or on the same disk (Windows systems), then the *whole* volume group or disk (all filesystems or disk images in this volume group or on disk) is synchronized to the backup system regardless of the filesystems, disk images or application objects selected for backup.
- In the second phase, the synchronized backup system data is backed up to a backup device.

During this phase, only the filesystems, disk images or application objects selected for backup are backed up to a backup device.

IMPORTANT

Such a concept enables the restore of selected objects (filesystems or disk images) for a split mirror restore and for a restore from backup media on LAN (filesystems, disk images or application objects).

With a split mirror restore, the links from the application to the backup system are synchronized before the restore, thus enabling the restore of the selected objects by establishing the current state of the application system data on the backup system, and then restoring the selected objects to the backup system, and finally resynchronizing the backup system to the application system.

Detailed Flow

When a backup is started, the following happens:

- The Data Protector Backup Session Manager (BSM) starts the Data Protector EMC Agent (SYMA) on both the application and backup systems.
- The BSM sends information about what needs to be backed up to the EMC Agent. The EMC Agent identifies the EMC devices (SLDs) for the data to be backed up (resolve process).
- If the `Run discovery of Symmetrix environment` option is enabled, the Data Protector database file is updated.
- If the `Re-establish links before backup` option is selected, the pairs are incrementally synchronized. Before the links are synchronized, filesystems on the backup system are dismounted. On HP-UX, volume groups are deactivated.
- If a script/command is set by the `Split pre-exec` option, this script/command is executed on the application system. The script can perform various tasks, such as stopping an application.
- If the `SYMA_UMOUNT_BEFORE_SPLIT` variable is set in the `omnirc` file, the filesystems on the application system are dismounted before the split and mounted again after the split.
- EMC devices (SLDs) are split. After the split operation, the pairs are in SPLIT mode.

- If a script/command is set by the `Split post-exec` option, this script/command is executed on the application system. The script can perform various tasks, such as restarting an application.

  If the Split pre-exec script/command fails, the Split post-exec script/command is not executed, so a cleanup procedure should be implemented in the Split post-exec script/command.

  If the `ZDB_ALWAYS_POST_SCRIPT` `omnirc` file variable is set to `1`, the Split post-exec script/command is always executed if set. By default the variable is set to `0`. Refer to "ZDB Agents Omnirc Variables" on page A-23 for more information on Data Protector EMC `omnirc` file variables.

- The backup system is prepared for backup. The mountpoint for the backed up filesystem is created on the backup system as described in the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*. On HP-UX, the backup volume groups are activated. The filesystems are mounted.
- If the `SYMA_MOUNT_R2_READWRITE` variable is set in the `omnirc` file on HP-UX, the volume group is activated in the read/write mode and the filesystem is mounted in the read/write mode. In this case, the filesystem check is executed before the filesystem is mounted.
- The backup system is backed up.
- The filesystems on the backup system are dismounted; on HP-UX, volume groups are deactivated on the backup system.
- If the `Re-establish links after backup` option is selected, the pairs are incrementally synchronized. If the option is not selected, the pairs are left split.
- If any of the described operations fail, the BSM aborts the backup session and stops the session according to the specified EMC options. For information on EMC options, refer to "Backup Options" on page 229.

Figure 11-1 Filesystem Split Mirror Backup Flow

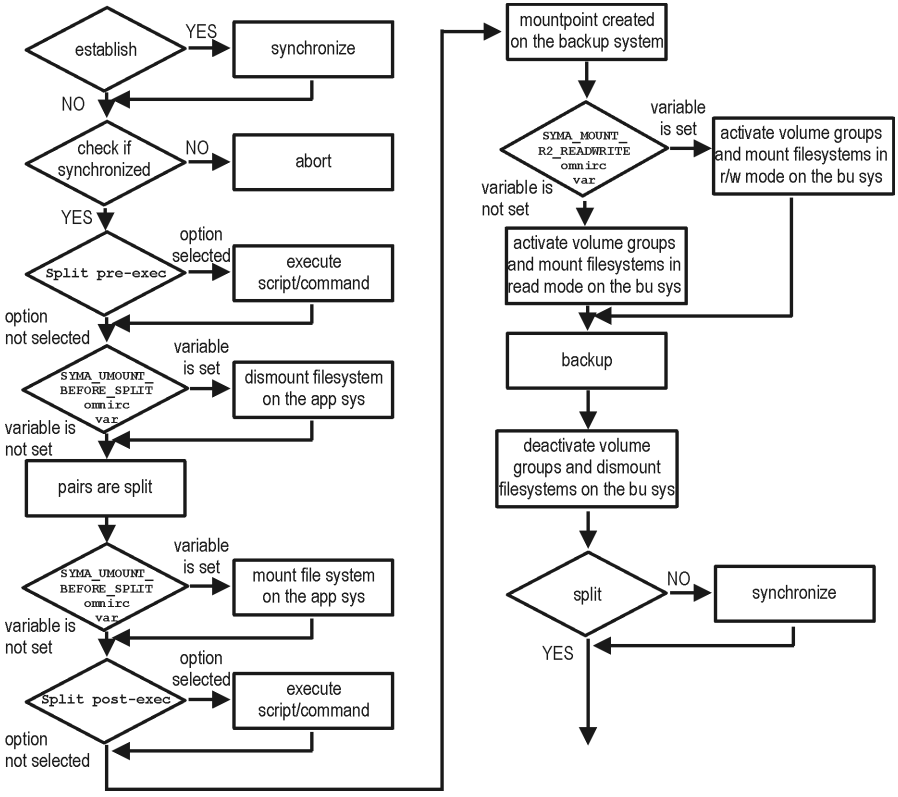

In Figure 11-1 on page 224, the “establish” and “split” checks are dependent on the following EMC backup option selections:

Table 11-1

The `Re-establish links after backup` option is selected	split = YES
The `Re-establish links before backup` option is selected	establish = YES
The `Re-establish links after backup` option is not selected	split = NO
The `Re-establish links before backup` option is not selected	establish = NO

Configuring a Backup Specification

Using the Data Protector GUI, you can create a filesystem or disk image ZDB-to-tape backup specification for the Data Protector EMC integration. Proceed as follows:

1. In the Context List, select Backup.
2. In the Scoping Pane, expand Backup and then Backup Specifications. Right-click Filesystem (for a filesystem backup) and click Add Backup.

   The Create New Backup dialog box is displayed.

   In the Backup type drop-down list, select the Split mirror backup option and select EMC Symmetrix in the Sub type drop-down list. For more information about the templates refer to the *HP OpenView Storage Data Protector Administrator's Guide*. See online Help for a description of other options. Click OK.
3. Select the application system and the backup system. Also specify the EMC configuration you will use for the backup: TimeFinder, SRDF, or Combined (SRDF + TimeFinder). See "Backup Options" on page 229 for detailed information about the other options you need to specify.

IMPORTANT

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

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

   Click Next.
4. Depending on whether you perform a filesystem or disk image backup, proceed as follows:

Filesystem Backup

   If you are configuring a filesystem backup, expand the application systems that contain the objects that you want to back up and then select what you want to back up.

   Click Next.

Disk Image Backup

If you are configuring a disk image backup, click `Next`.

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

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

   For detailed information on the object mirror functionality, see *HP OpenView Storage Data Protector Administrator's Guide*.

   Click `Next`.

6. In the `Backup Specification Options` group box, click the `Advanced` tab and then the `EMC Symmetrix` tab, to display EMC options. You can now modify the options, except the application system and the backup system.

Figure 11-2 Data Protector GUI: Backup Options Dialog Box

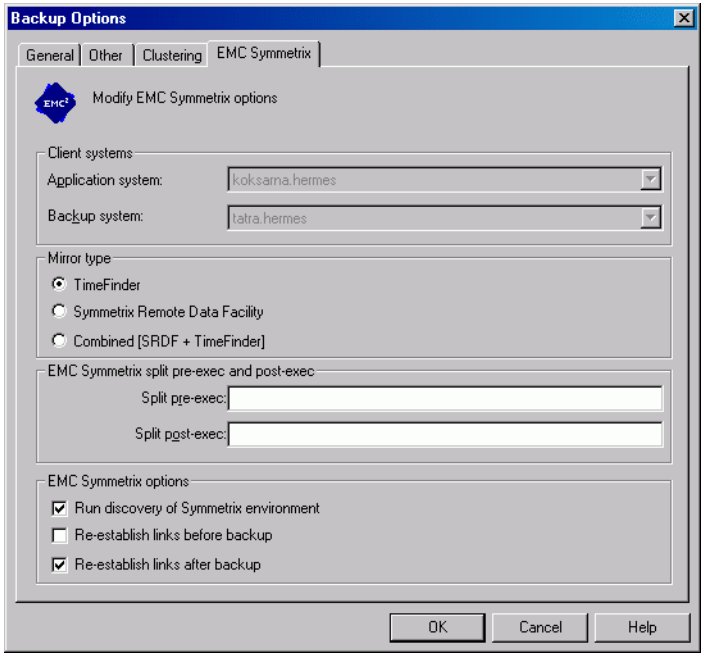

NOTE For detailed information on EMC backup specification options, see “Backup Options” on page 229.

For information on `Filesystem Options`, refer to the online Help.

7. Follow the backup wizard to open the Data Protector scheduler and then the `Backup Summary` dialog box, where a summary of the specified options is given.
8. Depending on whether you perform a filesystem or disk image backup, proceed as follows:

Filesystem Backup

If you are configuring a filesystem backup, click `Next`.

Disk Image Backup

For a disk image backup, proceed as follows:

a. Click `Manual add` to add the disk image objects you want to back up.

b. Select `Disk image object` and click `Next`.

c. Select the client to be backed up and click `Next`.

d. Follow the wizard to specify the `General Object Options` and the `Advanced Object Options`. For more information on these options, press **F1**.

e. In the `Disk Image Object Options` window, specify the disk image sections you want to back up.

To specify a rawdisk section, use the following format:

HP-UX

`/dev/rdsk/<filename>`, for example: `/dev/rdsk/c2t0d0`

To specify a raw logical volume section, use the following format:

`/dev/vg<number>/rlvol<number>`, for example: `/dev/vg01/rlvol1`

Windows

`\\.\PHYSICALDRIVE<#>`

Where `<#>` is the number of the disk you want to back up.

For example: `\\.\PHYSICALDRIVE3`.

For information on how to find the current numbers of the disks (physical drive numbers) you want to back up, refer to the *HP OpenView Storage Data Protector Administrator's Guide*.

f. Click `Finish` and then `Next`.

9. Save your backup specification. For information on starting and scheduling the backup session, refer to "Running and Scheduling a ZDB Session" on page A-3.

Backup Options

NOTE

Besides the EMC backup options, which are described in this section, there are two `omnirc` file variables that impact the backup behavior: `SYMA_UMOUNT_BEFORE_SPLIT` and `SYMA_MOUNT_R2_READWRITE`. Refer to "ZDB Agents Omnirc Variables" on page A-23 and to "Backup Process" on page 221 for more information on the two options.

EMC backup options are described as follows:

Table 11-2 EMC Backup Options in GUI and their Functions

Data Protector GUI	Function
`Application system`	Specify the application system on which the application runs. In a cluster environment, specify the virtual server hostname (rather than the physical node hostname).
`Backup system`	Specify the backup system on which your data is backed up. In a cluster environment, specify the virtual server hostname (rather than the physical node hostname). In the EMC GeoSpan for MSCS environment, select the backup system for the active node. After a failover, select the backup system for the currently active node and save the backup specification.
`Mirror type`	Specify the EMC configuration you will use for the backup: TimeFinder, Symmetrix Remote Data Facility, or Combined (SRDF + TimeFinder). In the EMC GeoSpan for MSCS environment, specify the TimeFinder configuration.
`Split pre-exec`	Specify the optional `Split pre-exec` command. Create this command in the `/opt/omni/lbin` (HP-UX systems) or *<Data_Protector_home>*`\bin` (Windows systems) directory on the application system. The `Split pre-exec` command is executed on the application system before the links are split. It is used, for example, to stop the application.

Table 11-2 EMC Backup Options in GUI and their Functions

Data Protector GUI	Function
`Split post-exec`	Specify the optional `Split post-exec` command. Create this command in the `/opt/omni/lbin` (HP-UX systems) or *<Data_Protector_home>*`\bin` (Windows systems) directory on the application system. The `Split post-exec` command is executed on the application system directly after the links are split. It is used, for example, to restart the application. If the Split pre-exec script/command fails, the Split post-exec script/command is not executed, so a cleanup procedure should be implemented in the Split post-exec script/command. If the `ZDB_ALWAYS_POST_SCRIPT` `omnirc` file variable is set to `1`, the Split post-exec script/command is always executed if set. By default the variable is set to `0`. Refer to "ZDB Agents Omnirc Variables" on page A-23 for more information on EMC `omnirc` variables.
`Run discovery of Symmetrix environment`	This option builds or re-builds the Data Protector EMC database on both the application and backup systems. Its functionality is the same as that of the Data Protector command `syma -init`. See "Creating the Data Protector EMC Database File" on page 216 for more information. This option is enabled by default. It takes a few minutes to run a discovery.
`Re-establish links before backup`	To maintain data integrity, Data Protector requires that the application and backup disks are synchronized before backup. With this option enabled, Data Protector automatically synchronizes disks before backup. This may be necessary if you disable the `Re-establish links after backup` option or if you used some other EMC commands that left the links split. Do not confuse this incremental establish option with the (full) establish option. This option is disabled by default.
`Re-establish links after backup`	This option allows you to re-establish the links between the application and mirrored devices after backup. If this option is disabled, the links are left split after the backup. In this case, you can use the mirrored devices on the backup system after the backup is finished. Use the option `Re-establish links before backup` to synchronize the disks before the next backup is started. This option is enabled by default.

Backup Disk Usage

If mirrored devices have not been re-established after backup, they still contain the last version of backed up data. You can use these mirrored devices to quickly restore or view your data.

NOTE Data can be quickly restored only by using the EMC device mirroring facilities.

To view this data, enable the mirrored devices, that is, activate the volume groups (HP-UX systems) and mount the filesystems. Information about volume groups (HP-UX systems) and filesystems is found in the EMC Agent log file, `/var/opt/omni/tmp/emc/R2_<session_name>.log` (HP-UX systems) or `<Data_Protector_home>\Config\client\tmp\emc\R2_<session_name>.log` (Windows systems), where *<session_name>* is composed of the sessionID, the forward slashes "/" having been replaced with dashes "-".

Testing Your Backed Up Data

This section shows how to periodically test the validity of your backed up data. To do so, either restore your data to the backup system or use the mirrored devices that have not been re-established after a backup and then test data integrity. Meanwhile, your applications run normally on the application system.

To restore to the backup system, follow the steps described in “Split Mirror Restore Procedure” on page 251 and set the EMC split mirror restore options as listed in “EMC Test Restore Options in GUI” on page 232.

EMC Test Options

The following table shows EMC test restore options as they appear in the Data Protector GUI along with their functions. Refer to “Split Mirror Restore Options” on page 253 for a more detailed description of these options.

NOTE

The `omnirc` file variable `SYMA_UMOUNT_BEFORE_SPLIT` must be set to `0` (default) and the `omnirc` file variable `SYMA_MOUNT_R2_READWRITE` must set to `1` for the test purposes. Refer to “ZDB Agents Omnirc Variables” on page A-23 and to the “Backup Process” on page 221 for more information on these two options.

Table 11-3 EMC Test Restore Options in GUI

Data Protector GUI	Function
`EMC Symmetrix mode`	Specify the EMC configuration you will use for the test backup: TimeFinder, SRDF, or Combined (SRDF + TimeFinder). In the EMC GeoSpan for MSCS environment, specify the TimeFinder configuration.
`Application system`	Specify the application system on which the application runs. In a cluster environment, specify the virtual server hostname (rather than the physical node hostname).

Table 11-3 EMC Test Restore Options in GUI

Data Protector GUI	Function
`Backup system`	Specify the backup system on which your data is to be restored. In a cluster environment, specify the virtual server hostname (rather than the physical node hostname). In the EMC GeoSpan for MSCS environment, select the backup system for the active node. After a failover, select the backup system for the currently active node and save the backup specification.
`Run discovery of the Symmetrix environment`	You must disable this option.
`Re-establish links before restore`	You can either enable or disable this option.
`Disable disks on application client before split`	You must disable this option upon testing your backup, that is, disks on the application system should not be disabled. The `Restore links after restore` option is also disabled, so that applications on the application system run uninterrupted. The restored data must not be moved to the application system for test purposes. This would cause integrity problems.
`Restore links after restore`	You must disable this option, leaving the links in a split state. You can then check the integrity of restored data on the backup system.

Checking Your Restored Data

As the `Restore links after restore` option is disabled for test purposes, the mirrored devices contain the restored version of data.

To view this data, you first enable the mirrored devices and mount the filesystems.

Manually re-establish the links using the appropriate EMC CLI command (`symrdf` or `symmir`), or enable the option `Re-establish links before backup` (for backup), or `Re-establish links before restore` (for restore) for the next backup or restore session.

NOTE Do not restore the data to the application system for the test purposes. If you do so, you will lose all your data written to the mirrored devices on the application system after the split.

Troubleshooting

This section describes the most common problems you may encounter when using the Data Protector EMC integration.

Before You Begin

- Ensure that the latest official Data Protector patches are installed. Refer to the *HP OpenView Storage Data Protector Software Release Notes* or http://www.openview.hp.com/products/datapro/spec_0001.html for an up-to-date list of supported versions, platforms, and other information.
- Examine system errors reported in `/var/opt/omni/log/debug.log` (HP-UX or Solaris systems) or in *<Data_Protector_home>*`\log\debug.log` (Windows systems) on the application system and the backup system.

The following is a description of problems, and actions to be taken to resolve them.

Recovery Using the EMC Agent

If a backup or other operation did not finish successfully, then the EMC environment is left in an undefined state, for example, with links split, device groups not deleted in the Data Protector EMC database file, filesystems on the backup system mounted, volume groups on the backup system activated, and so on. In this case, invoke the EMC Agent (SYMA) `recovery` command to recover the environment. Information about EMC Agent objects, device groups, and volume groups is logged into the EMC Agent recovery files:

HP-UX

`/var/opt/omni/emc/symmR1.rec`

`/var/opt/omni/emc/symmR2.rec`

Windows

<Data_Protector_home>`\Config\Emc\symmR1.rec`

<Data_Protector_home>`\Config\Emc\symmR2.rec`

At the time a record is entered, it is marked as a valid record. It is marked as invalid if the session does not terminate normally. Invalid records are automatically deleted when the EMC Agent recovery file exceeds a certain value, by default, `SYMA_REC_FILE_LIMIT = 102400` bytes.

To recover the environment, invoke the following command to re-establish the links and the delete device groups. The next split mirror backup or split mirror restore session will dismount filesystems and de-activate volume groups on the backup system.

- On the application system:

HP-UX

```
/opt/omni/lbin/syma -r1 -session <sessionID> -recovery
```

Windows

```
<Data_Protector_home>\bin\syma -r1 -session <sessionID>
-recovery
```

- On the backup system:

HP-UX

```
/opt/omni/lbin/syma -no_r1 -session <sessionID> -recovery
[-split]
```

Windows

```
<Data_Protector_home>\bin\syma -no_r1 -session
<sessionID> -recovery [-split]
```

You can obtain the *<sessionID>* from the Data Protector GUI as shown in the figure Figure 11-3 on page 237.

Figure 11-3 Obtaining the Session ID

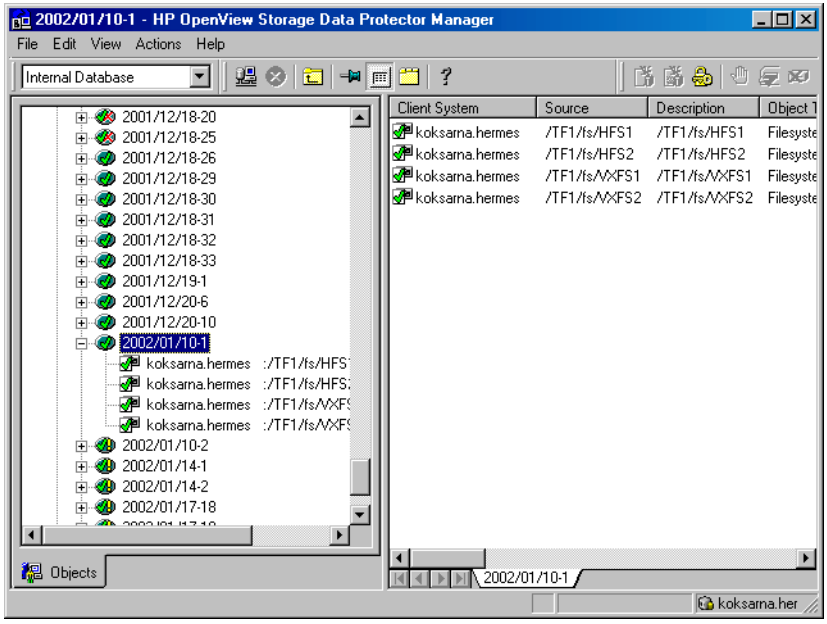

The `split` option disables the synchronization of links.

This command reads the recovery file and recovers the state of the environment before the session.

NOTE Do not edit or restore the EMC Agent recovery file.

Creation of the Backup Specification

- **You are unable to select the EMC mode in the Data Protector GUI when attempting to create a backup specification.**

Action

Check whether the EMC Agent software component is installed on both the application and backup systems.

Check the Data Protector `cell_info` file on the Data Protector Cell Manager in *<Data_Protector_home>*`\Config\Server\cell\cell_info`

(Windows Cell Manager) or in `/etc/opt/omni/server/cell/cell_info` (HP-UX or Solaris Cell Manager) to see if the EMC Agent software component is installed.

An exemplary file should look similar to the following:

```
-host "hpsap001.bbn.hp.com" -os "hp s800 hp-ux-11.00" -cc
A.05.50 -da A.05.50 -emc A.05.50

-host "hpsap002.bbn.hp.com" -os "hp s800 hp-ux-11.00" -cc
A.05.50 -da A.05.50 -ma A.05.50 -emc A.05.50
```

Backup Problems

- **The EMC Agent on the application system failed to dismount a filesystem.**

Action Check the `Split pre-exec` script. In this script, stop all processes using the filesystem.

- **The EMC Agent failed to synchronize the disks, so the split failed.**

Description To successfully split the disks, the EMC Agent first checks the status of the links. The links can be split only after all devices have been synchronized. The EMC Agent checks the status of links after every 30 seconds and retries 15 times. To change these values, use the variables `SYMA_SYNC_RETRY=<number of retries>` and `SYMA_SLEEP_FOR_SYNC=<sleep time in seconds>`.

Action You can increase the time frame for synchronization by setting the variables `SYMA_SLEEP_FOR_SYNC` in the `<Data_Protector_home>\omnirc` (Windows systems) or in the `/opt/omni/.omnirc` (HP-UX systems) file on the backup system. For additional information, refer to “ZDB Agents Omnirc Variables” on page A-23.

- **The EMC Agent reports that an EMC device is not part of a BCV pair.**

Action Check the selected EMC configuration. If the TimeFinder or Combined (SRDF + TimeFinder) configuration is used, then all backup disks on the application system must have an associated BCV device on the backup system.

- **The EMC Agent reports that a device group could not be created.**

Action Check if any of the previous sessions were improperly stopped, and run the EMC Agent recovery operation for the session on the backup system.

- **The EMC Agent reports that it failed to add a device into a device group or associate a BCV to a device group.**

Action Check if any of the previous backups were improperly stopped, and then run the EMC Agent recovery operation for the session on the backup system. See “Recovery Using the EMC Agent” on page 235 for instructions.

- **The EMC Agent reports that the volume group on the backup system cannot be de-activated.**

Action Check the backup volume group. If there are any processes running on the volume group filesystem, stop them to release the filesystem.

- **The EMC Agent reports that it failed to rebuild the Data Protector EMC database.**

Action Run the discover operation `<Data_Protector_home>\bin\syma -init` (Windows systems) or `/opt/omni/lbin/syma -init` (HP-UX systems). This must be done on both the application system and backup system. If the operation succeeds, then disable the `Run discovery of Symmetrix environment` option in the Data Protector user interface and restart the backup.

If the discovery operation fails, run the `symcfg -discover` command.

- **Resolving on some object failed**.

Action Check the EMC Agent log file on the application system and ensure that all objects that are logged into this file are created on the mirrored EMC devices.

- **EMC Agent reports an invalid link state on the EMC device.**

Action Check the state of the link. If the link is split, use the `Re-establish links before backup` option.

Error Messages

The following section provides information on error messages and how to solve them.

Message

```
[Major] From: SYMA@Backup (R2) System ""  Time: 04/03/99 09:18:34

[223:324] SYMA-R2 Could not add device 048 from Symmetrix
000282600317 to device group SYMA_REG_1999-03-04-2_0.

(SYMAPI-The device is already a member of a device group)
```

The EMC Agent failed in some of the sessions beforehand, and therefore backup was not completed. As the result you may get this message as well as the others that follow.

Actions

- Run a recovery of the failed session to create a consistent environment.
- Check that the `/var` directory is not 100% full. If `/var` is full, the EMC Agent does not have enough space to write its record into the file, therefore the session fails.

Message

```
[Major] From: SYMA@Application (R1) System ""  Time: 11/03/99
15:06:22

[223:193] SYMA-R2 Could not activate volume group /dev/tf1_fs2_b
```

The cause may either be that the backup volume group is not deactivated or that there is a problem with the configuration.

Actions

- Run the same backup with debug on, and then check the EMC Agent R2 debug file on the backup system for LVM error messages.
- Try manually splitting the links on the backup system and activating the backup volume group. If you are unable to do this, the backup may fail with an error `[223:193]`.

Message

```
[Major] From: SYMA@Application (R1) System ""  Time: 3/31/99
11:32:58 AM

[223:406] Failed to initialize the SYMAPI session

(SYMAPI-The version of the symapi library is too old; please
upgrade to a newer version of SYMAPI)
```

Action

- Check the EMC Solution Enabler version.

Message

```
[Major] From: SYMA@Application (R1) System ""  Time: 6/30/99
10:57:00 AM

[223:408] Failed to re-sync Symmetrix database. (SYMAPI-No
Symmetrix devices were found)
```

Action

- Run the same session with the option `Run discovery of Symmetrix environment` set.

Message

```
[Major] From: SYMA@Application (R1) System ""  Time: 3/31/99
2:17:43 PM

[223:407] Failed to rescan host devices and rebuild
Symmetrix database

SYMAPI-Error opening the gatekeeper device for communication
to the Symmetrix)
```

Actions

- Run `symcfg discover`. If the problem persists, check the pseudo-devices file.
- If the device you want to use as a gatekeeper or BCV device is accessed through the HP-PB (NIO) SCSI bus controller, you must create pseudo-devices for all gatekeepers and BCV devices.
- See the README file in `/var/symapi/config/README.pseudo_devices`.

Message

```
[Major] From: SYMA@Backup (R2) System ""  Time: 5/11/99 12:01:11 PM

[223:335] SYMA-R2 Failed to synchronize SRDF links in device
group SYMA_RDF2_1999-05-11-21_0 before backup. (SYMAPI-The
operation failed because another process has an exclusive
lock on a locally-attached Symmetrix)

[Major] From: SYMA@Backup (R2) System ""  Time: 5/11/99 12:01:13 PM

SYMA-R2 Invalid SRDF link state of device 000 from Symmetrix
000282600317 (links state=103)
```

Devices are not synchronized.

Action Manually establish links or use the option `Establish Links Before Backup`. If the problem persists, run:

CAUTION See the `symrdf` man page about the bypass option before running the following:

```
symrdf -g <Dg_name> establish -bypass
```

Message

```
[Major] From: SYMA@twingo ""  Time: 6/7/99 1:08:30 PM

[223:301] SYMA-R2 Device 006 from Symmetrix 000182600287 is not
part of a BCV pair
```

Actions

- Check the backup options in the backup specification.
- Check the configuration in the backup specification.

Message

```
[Major] From: SYMA@Backup (R2) System ""  Time: 8/4/99 3:26:27 PM

SYMA-R2 Invalid SRDF link state of device 001 from Symmetrix
000282600317 (links state=103)

[Major] From: SYMA@Backup (R2) System ""  Time: 8/4/99 3:26:28 PM

[223:361] SYMA-R2 Split of links(s), which belong to the object
/dev/rdsk/c1t8d0, has failed. (Unexpected state of rdf link)
```

The connection between EMC R1 and R2 devices is not established.

Action

- Run the same session with the option `Re-establish links before backup`.

Message

```
[Major] From: SYMA@Backup (R2) System ""  Time: 8/30/99 11:37:12 AM

[223:125] SYMA-R2 Resolving of object /RDF/fs/HFS has failed
(Volume group is not deactivated)
```

The volume group on the backup system is still activated.

Action

- On the backup system, split the links and deactivate the backup volume group. Re-establish the links manually or select the option `Re-establish links before backup` in the Data Protector GUI backup specification.

12 Restore

In This Chapter

This chapter describes how to configure and run a filesystem or disk image restore of your data backed up using the Data Protector Data Protector EMC Symmetrix (EMC) integration. The sections describe steps for performing a restore using the Data Protector Graphical User Interface.

This chapter includes the following sections:

"Overview" on page 245

"Restoring from Backup Media on LAN" on page 246

"Split Mirror Restore" on page 249

"Troubleshooting" on page 257

Overview

There are two methods of restoring data:

- Restore from backup media on LAN. Refer to “Restoring from Backup Media on LAN” on page 246.
- Split mirror restore. Refer to “Split Mirror Restore” on page 249.

Restoring from Backup Media on LAN

Data is restored directly to the application system. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for more information on various restore options available. The restore procedure described here provides only a general description of how to restore an object that was backed up using the split mirror backup functionality from backup media on LAN.

TIP

You can improve the data transfer rate when restoring by connecting the backup device to the application system and configuring it using the Data Protector GUI. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for more information on configuring backup devices. Refer to the "Restoring Under Another Device" section of the same guide for more information on how to perform a restore using another device.

1. In the `Context List`, select `Restore`.
2. Select the objects for restore and click them to display their properties in the Scoping Pane.

   Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for more information on various restore options available. Be sure to make the following selection in the Scoping Pane in order to restore an object that was backed up using the split mirror backup functionality from backup media on LAN:

   - Select the application system as the `Target client` under the `Destination` tag.

Figure 12-1 Selecting the Application System

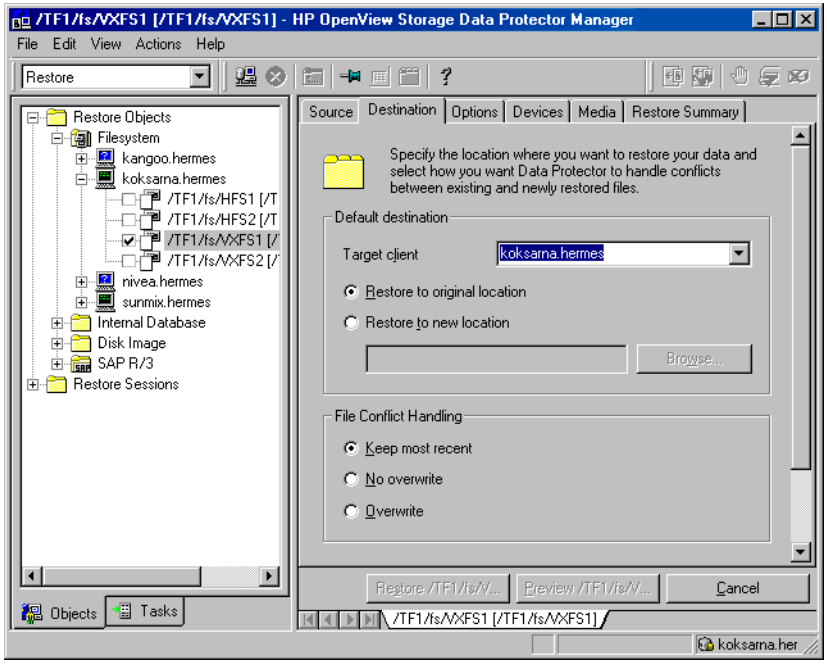

3. After you have set the restore options, click the `Restore` button. The `Start Restore Session` dialog box is displayed.
4. Click `Next` to specify the report level and network load. Click `Next`.
5. In the `Start Backup Session` dialog window, select `Disabled` in the `EMC Symmetrix mode` drop-down list. This sets a restore from backup media on LAN.

Figure 12-2 Selecting a Restore from Backup Media on LAN

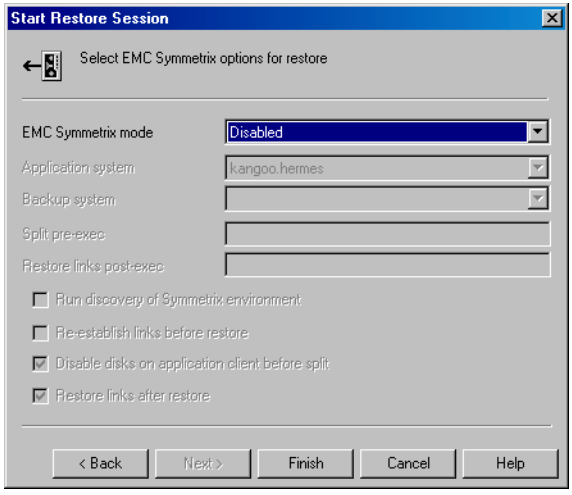

6. Click `Finish` to start the restore session.

Split Mirror Restore

Data is restored to the backup system and then moved to the application system by resynchronizing. The procedure consists of the following automated steps:

1. Preparing the backup and application systems for restore.
2. Restoring data from backup media on LAN to the backup system and synchronizing this data to the application system.

Split Mirror Restore Process

The following happens during a split mirror restore session:

1. The Data Protector Restore Session Manager (RSM) sends information about what needs to be restored to the EMC Agent.
2. Links are re-established if the `Re-establish links before restore` option is enabled.
3. `Restore Split pre-exec` is started. The following might be set in the `Restore Split pre-exec` script/command:

   On HP-UX, volume groups should be prepared for de-activation and all applications stopped. Filesystems that are not included in the restore session, but are physically located on the same disks as the filesystems that are included in the restore, should be dismounted. The filesystems that are included in the session are dismounted automatically.

   On Windows, all applications should be stopped.
4. The application system is prepared for restore:

   On HP-UX, the filesystems are dismounted and volume groups that are a part of the restore are de-activated. This prevents the system from using disks that are a part of the restore process.

   On Windows, the filesystems are automatically dismounted by the agent. This prevents the system from using disks that are a part of the restore process.
5. Links are split.

6. The backup system is prepared for the restore:

   On HP-UX, primary device groups are created and disks used for restore are added to these groups. Device groups are split, volume groups are activated, and filesystems are mounted.

   On Windows, primary device groups are created and disks used for restore are added to these groups. Device groups are split and filesystems mounted.

7. The restore is performed to the backup system.
8. When the restore is finished, the volume groups (on an HP-UX backup system) are de-activated and filesystems dismounted.
9. Restore device groups are created and successfully restored disks are moved to the restore device groups. Disks with failed restore objects remain in the primary device group.

   If the `Restore links after restore` option is enabled, disks in the restore device groups are incrementally restored to the application system. Disks in the primary device group are re-established.

   If the `Restore links after restore` option is disabled, the links are left in a split state. You can check the data and decide which disks you want to restore to the application system.

10. On HP-UX, volume groups on the application system are re-activated and filesystems are remounted.

    On Windows, filesystems are remounted on the application system.

11. `Restore Split post-exec` is started - it should remount filesystems and start the application.

NOTE If the `Restore links after restore` option is enabled, you must not restart stopped applications until the data is moved to the application system.

Refer to Figure 12-3 on page 251 for an overview of the filesystem split mirror restore flow.

Figure 12-3 Split Mirror Restore Flow

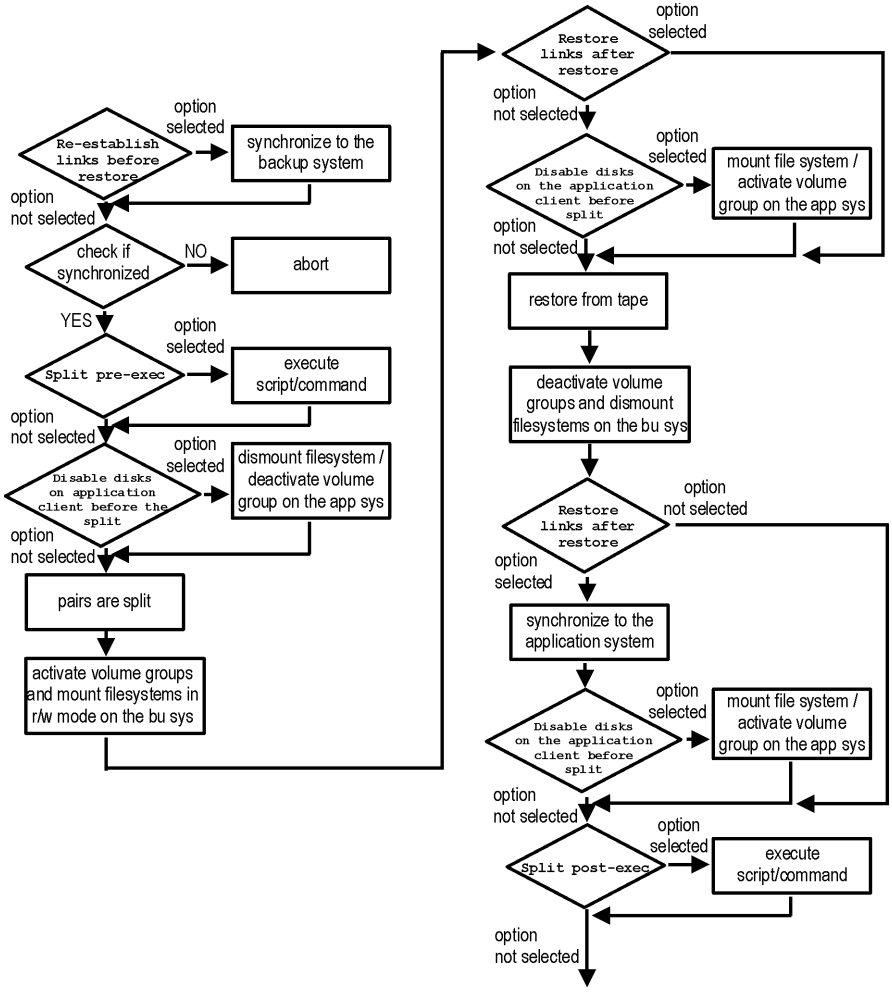

Split Mirror Restore Procedure

Follow the procedure in this section to perform a split mirror restore of a filesystem or a disk image. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for more information on the general restore procedure.

1. In the `Context List`, select `Restore`.

2. Select the objects for restore and click them to display their properties in the Scoping Pane.

NOTE

To perform a split mirror restore, you need to select the application system as the `Target client` under the `Destination` tag. If the backup system is selected, a normal restore to the backup system is performed.

3. 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. Specify the EMC split mirror restore options. Refer to “Split Mirror Restore Options” on page 253 for more information on these options.

Figure 12-4 EMC Split Mirror Restore Options

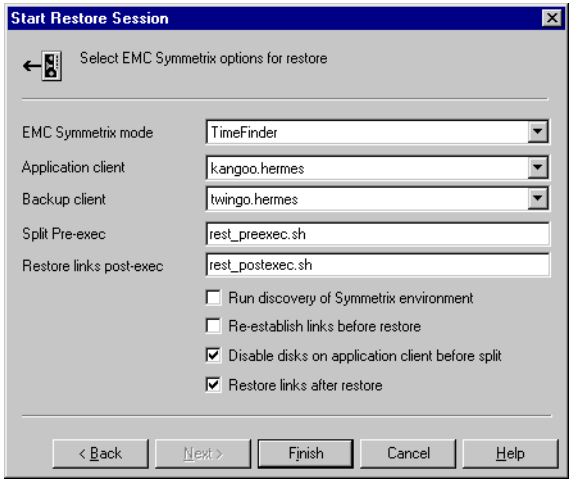

6. Click `Finish` to start the restore session.

IMPORTANT

It is not possible to start split mirror backup or split mirror restore sessions using the same disk on the application system at the same time. A split mirror backup or split mirror restore session must be started only

after the preceding session using the same disk on the application system has finished the synchronization operation, otherwise the session will fail.

Split Mirror Restore Options

The following table shows the EMC split mirror restore options as they appear in the Data Protector GUI, along with their functions:

Table 12-1 EMC Split Mirror Restore Options in the Data Protector GUI and their Functions

Data Protector GUI	Function
`EMC Symmetrix mode`	Specify the EMC Symmetrix configuration you will use for the restore: TimeFinder, SRDF, or Combined (SRDF + TimeFinder).
`Application system`	Specify the application system on which the application runs. In a cluster environment, specify the virtual server hostname (rather than the physical node hostname).
`Backup system`	Specify the backup system on which your data is first restored. In a cluster environment, specify the virtual server hostname (rather than the physical node hostname).

Table 12-1 EMC Split Mirror Restore Options in the Data Protector GUI and their Functions

Data Protector GUI	Function
`Split pre-exec`	Specify the `Split pre-exec` command, which is run before the links are split. Create the command in the `/opt/omni/lbin` (HP-UX systems) or *<Data_Protector_home>*`\bin` (Windows systems) directory on the application system. The command can be used for stopping the application and dismounting filesystems (HP-UX systems only) that are not to be restored in the active session, and are mounted to the volume groups that will be restored in the same session. This prepares the volume groups for de-activation. If the Split pre-exec script/command fails, the Restore links post-exec script/command is not executed, so a cleanup procedure should be implemented in the Restore links post-exec script/command. If the `ZDB_ALWAYS_POST_SCRIPT` `omnirc` file variable is set to `1`, the Restore links post-exec script/command is always executed if set. By default the variable is set to `0`. Refer to “ZDB Agents Omnirc Variables” on page A-23 for more information on Data Protector HP EMC `omnirc` file variables.
`Restore links post-exec`	Specify the `Restore links post-exec` command, which is run after the links are restored. Create the command in the `/opt/omni/lbin` (HP-UX systems) or *<Data_Protector_home>*`\bin` (Windows systems) directory on the application system. It is used to remount filesystems (HP-UX systems only) and to restart the application. Do not use this command to enable the application if you have disabled the option `Re-establish links after restore`. In such a case, applications using restored disks should not be restarted until the links are established manually or a decision that the data will not be moved to the application system has been made.
`Run discovery of Symmetrix environment`	This option builds or re-builds the Data Protector EMC database on both the application and backup systems. Its functionality is the same as that of the Data Protector `syma -init` command. See “Creating the Data Protector EMC Database File” on page 216 for more information. This option is disabled by default.

Table 12-1 EMC Split Mirror Restore Options in the Data Protector GUI and their Functions

Data Protector GUI	Function
`Re-establish links before restore`	This option is used to synchronize split disks, that is, to move data to backup disks. This is necessary to prepare the disks for restore and to enable accurate restores. This option is disabled by default.
`Disable disks on application client before split`	Disks on the application system are disabled, that is, the filesystems are dismounted and volume groups de-activated (HP-UX systems only). This is performed before the split. The disks are enabled after the links are restored. Note that only filesystems selected for restore are dismounted. You must always select this option for restore when you want to move data from the backup to the application system, that is, to incrementally restore links. The application system disks have to be disabled to provide data integrity after the links are restored, that is, data is moved.
`Restore links after restore`	With this option enabled, the EMC Agent incrementally restores the links of devices that were successfully restored by Data Protector to the backup system. The EMC Agent also incrementally re-establishes the links of devices that were not successfully restored by Data Protector to the backup system.

Split Mirror Restore in a Cluster

To perform a Data Protector split mirror restore when a filesystem is running in an MC/ServiceGuard or a Microsoft Cluster Server on the application system, it is necessary to perform some additional steps.

MC/ServiceGuard Procedure

To perform a Data Protector split mirror restore when a filesystem is running in an MC/ServiceGuard on the application system, it is necessary to stop the filesystem cluster package in order to activate the mirrored volume groups on the application system in normal mode before the split mirror restore session is started.

Follow the procedure below to perform a Data Protector split mirror restore to the application system in an MC/ServiceGuard cluster:

1. Stop the filesystem cluster package:

   ```
   cmhaltpkg <app_pkg_name>
   ```

   This will stop the filesystem services and dismount the mirrored volume group filesystem.

2. Deactivate the mirrored volume group from the cluster mode and activate it in the normal mode:

   ```
   vgchange -c n /dev/<mirror_vg_name>
   ```

   ```
   vgchange -q n -a y /dev/<mirror_vg_name>
   ```

3. Mount the mirrored volume group filesystem:

   ```
   mount /dev/<mirror_vg_name>/<lv_name> /<mountpoint>
   ```

4. Using Data Protector, start the Data Protector split mirror restore session. Refer to “Split Mirror Restore Procedure” on page 251 for a detailed explanation.

IMPORTANT When specifying the application system during a split mirror restore procedure, make sure to specify the hostname of the application system *node* on which the mirrored volume group was activated in the normal mode when performing step 2 of this procedure.

5. When the split mirror restore session has finished, dismount the mirrored volume group filesystem:

   ```
   umount /<mountpoint>
   ```

6. Deactivate the mirrored volume group in the normal mode and activate it in the cluster mode:

   ```
   vgchange -a n /dev/<mirror_vg_name>
   ```

   ```
   vgchange -c y /dev/<mirror_vg_name>
   ```

7. Start the filesystem cluster package:

   ```
   cmrunpkg <app_pkg_name>
   ```

Troubleshooting

This section describes the most common problems you may encounter during the restore when using the Data Protector EMC integration.

Before You Begin

- Ensure that the latest official Data Protector patches are installed. Refer to the *HP OpenView Storage Data Protector Software Release Notes* or http://www.openview.hp.com/products/datapro/spec_0001.html for an up-to-date list of supported versions, platforms, and other information.
- Examine system errors reported in `/var/opt/omni/log/debug.log` (HP-UX systems) or in `<Data_Protector_home>\log\debug.log` (Windows systems) on the application system and the backup system.

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

Recovery Using the EMC Agent

If a backup or other operation did not finish successfully, then the EMC environment is left in an undefined state, for example, with links split, device groups not deleted in the Data Protector EMC database file, filesystems on the backup system mounted, volume groups on the backup system activated, and so on. In this case, invoke the EMC Agent (SYMA) `recovery` command to recover the environment. Information about EMC Agent objects, device groups, and volume groups is logged into the EMC Agent recovery files:

HP-UX

`/var/opt/omni/emc/symmR1.rec`

`/var/opt/omni/emc/symmR2.rec`

Windows

`<Data_Protector_home>\Config\Emc\symmR1.rec`

`<Data_Protector_home>\Config\Emc\symmR2.rec`

At the time a record is entered, it is marked as a valid record. It is marked as invalid if the session does not terminate normally. Invalid records are automatically deleted when the EMC Agent recovery file exceeds a certain value, by default, `SYMA_REC_FILE_LIMIT = 102400` bytes.

To recover the environment, invoke the following command to re-establish the links and delete device groups. The next split mirror backup or split mirror restore session will dismount filesystems and de-activate volume groups (on HP-UX systems) on the backup system:

- On the application system:

HP-UX

```
/opt/omni/lbin/syma -r1 -session <sessionID> -recovery
```

Windows

```
<Data_Protector_home>\bin\syma -r1 -session <sessionID>
-recovery
```

- On the backup system:

HP-UX

```
/opt/omni/lbin/syma -no_r1 -session <sessionID> -recovery
[-split]
```

Windows

```
<Data_Protector_home>\bin\syma -no_r1 -session
<sessionID> -recovery [-split]
```

You can obtain the *<sessionID>* from the Data Protector GUI as shown in the Figure 12-5 on page 259.

Figure 12-5 Obtaining the Session ID

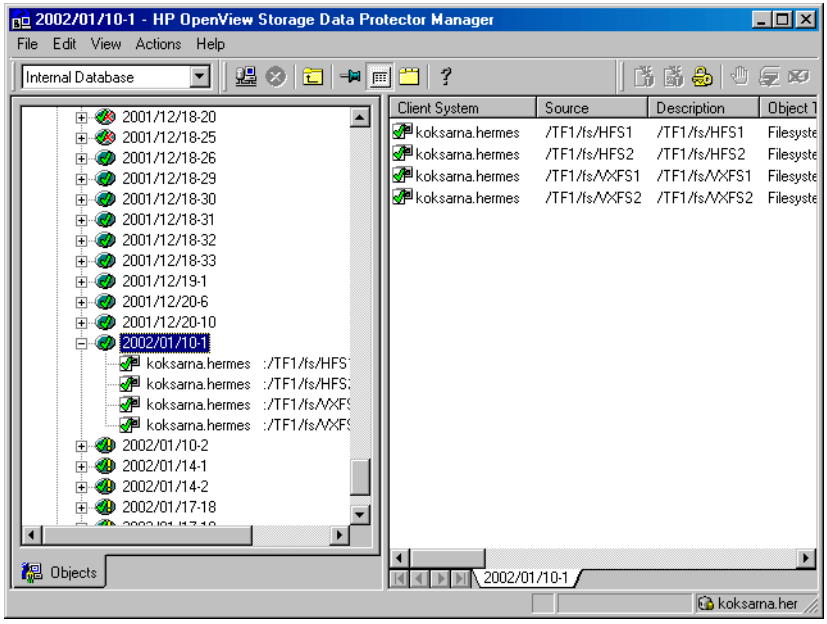

The `split` option disables the synchronization of links.

This command reads the recovery file and recovers the state of the environment before the session.

NOTE Do not edit or restore the EMC Agent recovery file.

Split Mirror Restore Problems

- **The EMC Agent on the application system failed to de-activate volume groups during the restore session (HP-UX systems only).**

Action

Check the `Split pre-exec` script. In this script, stop all processes using the affected volume groups and dismount all filesystems created on these volume groups that are not to be restored in the current session.

- **The EMC Agent failed to synchronize the disks, so the split failed.**

Description

To successfully split the disks, the EMC Agent first checks the status of the links. The links can be split only after all devices have been synchronized. The EMC Agent checks the status of the links after every 30 seconds and retries 15 times. To change these values, use the variables `SYMA_SYNC_RETRY=<number of retries>` and `SYMA_SLEEP_FOR_SYNC=<sleep time in seconds>`.

Action

You can increase the time frame for synchronization by setting the variable `SYMA_SLEEP_FOR_SYNC` in the `<Data_Protector_home>\omnirc` (Windows systems) or `/opt/omni/.omnirc` (HP-UX systems) file on the backup system. For additional information, refer to "ZDB Agents Omnirc Variables" on page A-23.

- **The EMC Agent reports that an EMC device is not part of a BCV pair.**

Action

Check the selected EMC configuration. If the TimeFinder or Combined (SRDF + TimeFinder) configuration is used, then all backup disks on the application system must have an associated BCV device on the backup system.

- **The EMC Agent reports that a device group could not be created.**

Action

Check if any of the previous sessions were improperly stopped, and run the EMC Agent recovery operation for the session on the backup system. See "Recovery Using the EMC Agent" on page 257 for instructions.

- **The EMC Agent reports that it failed to add a device into a device group or associate a BCV to a device group.**

Action

Check if any of the previous backups were improperly stopped, and then run the EMC Agent recovery operation for the session on the backup system. See "Recovery Using the EMC Agent" on page 257 for instructions.

- **The EMC Agent reports that it failed to rebuild the Data Protector EMC database.**

Action

Run the discover operation `<Data_Protector_home>\bin\syma -init` (Windows systems), or `/opt/omni/lbin/syma -init` (HP-UX systems). This has to be done on both the application system and backup system. If the operation succeeds, then disable the `Run discovery of Symmetrix environment` option in the Data Protector user interface and restart the restore.

If the discovery operation fails, run the `symcfg -discover` command.

- **Resolving of some object failed.**

Action

Check the EMC Agent log file on the application system and ensure that all objects that are logged into this file are created on the mirrored EMC devices.

- **EMC Agent reports an invalid link state on the EMC device.**

Action

Check the state of the link. If the link is split, use the `Re-establish links before restore` option.

Error Messages

The following section provides information on error messages, and how to solve them.

Message

```
[Major] From: SYMA@Backup (R2) System ""  Time: 04/03/99 09:18:34
[223:324] SYMA-R2 Could not add device 048 from Symmetrix
000282600317
to device group SYMA_REG_1999-03-04-2_0.
(SYMAPI-The device is already a member of a device group)
```

The EMC Agent failed in some of the sessions beforehand and therefore the recovery was not completed. As the result you may get this message, as well as others that follow.

Actions

- Run a recovery of the failed session to create a consistent environment.
- Check that the `/var` directory is not 100% full. If `/var` is full, the EMC Agent does not have enough space to write its record into the file, therefore the session fails.

Message

```
[Major] From: SYMA@Application (R1) System ""  Time: 11/03/99
15:06:22

[223:193] SYMA-R2 Could not activate volume group /dev/tf1_fs2_b
```

The cause may either be that the volume group is not deactivated or that there is a problem with the configuration.

Actions

- Run the same restore with debug on, and then check the EMC Agent R2 debug file on the backup system for LVM error messages.
- Try manually splitting the links on the backup system and activating the volume group. If you are unable to do this, the restore may fail with an error `[223:193]`.

Message

```
[Major] From: SYMA@Application (R1) System ""  Time: 3/31/99
11:32:58 AM

[223:406] Failed to initialize the SYMAPI session

(SYMAPI-The version of the symapi library is too old; please
upgrade to a newer version of SYMAPI)
```

Action

- Check the EMC Solution Enabler version.

Message

```
[Major] From: SYMA@Application (R1) System ""  Time: 6/30/99
10:57:00 AM

[223:408] Failed to re-sync Symmetrix database. (SYMAPI-No
Symmetrix devices were found)
```

Action

- Run the same session with the option `Run discovery of Symmetrix environment` set.

Message

```
[Major] From: SYMA@Application (R1) System ""  Time: 3/31/99
2:17:43 PM

[223:407] Failed to rescan host devices and rebuild Symmetrix
database

SYMAPI-Error opening the gatekeeper device for communication to the
Symmetrix)
```

Actions

- Try to run `symcfg discover`. If that does not help, check the pseudo-devices file.
- If the device you want to use as a gatekeeper or BCV device is accessed through the HP-PB (NIO) SCSI bus controller, you must create pseudo-devices for all gatekeepers and BCV devices.

- See the README file in /var/symapi/config/README.pseudo_devices.

Message

```
[Major] From: SYMA@Backup (R2) System ""  Time: 5/11/99 12:01:11 PM

[223:335] SYMA-R2 Failed to synchronize SRDF links in device group
SYMA_RDF2_1999-05-11-21_0 before backup. (SYMAPI-The operation
failed because another process has an exclusive lock on a
locally-attached Symmetrix)

[Major] From: SYMA@Backup (R2) System ""  Time: 5/11/99 12:01:13 PM

SYMA-R2 Invalid SRDF link state of device 000 from Symmetrix
000282600317 (links state=103)
```

Devices are not synchronized.

Action

Manually establish links or use the option `Re-establish Links Before Restore`. If the problem persists, run:

CAUTION

See the `symrdf` man page about the bypass option before running the following:

```
symrdf -g <Dg_name> establish -bypass
```

Message

```
[Major] From: SYMA@twingo ""  Time: 6/7/99 1:08:30 PM

[223:301] SYMA-R2 Device 006 from Symmetrix 000182600287 is not
part of a BCV pair
```

Actions

- Check the restore options.

Message

```
[Major] From: SYMA@Backup (R2) System ""  Time: 8/4/99 3:26:27 PM

SYMA-R2 Invalid SRDF link state of device 001 from Symmetrix
000282600317 (links state=103)

[Major] From: SYMA@Backup (R2) System ""  Time: 8/4/99 3:26:28 PM

[223:361] SYMA-R2 Split of links(s), which belong to the object
/dev/rdsk/c1t8d0, has failed. (Unexpected state of rdf link)
```

The connection between EMC R1 and R2 devices is not established.

Action

- Run the same session with the option `Re-establish links before restore`.

Message

```
[Major] From: SYMA@Backup (R2) System ""  Time: 8/30/99 11:37:12 AM
[223:360] SYMA-R2 Resolving of object /RDF/fs/HFS has failed
(Volume group is not deactivated)
```

The volume group on the backup system is still activated.

Action

- On the backup system, split the links and deactivate the backup volume group. Re-establish the links manually or select the option `Re-establish links before restore` in the Data Protector GUI backup specification.

A Appendix

In This Appendix

This appendix gives information on the following topics:

- “Running and Scheduling a ZDB Session” on page A-3
- “Alternate Paths Support” on page A-8
- “Cluster Configurations” on page A-10
- “Instant Recovery in a Cluster” on page A-20
- “ZDB Agents Omnirc Variables” on page A-23
- “User Scenarios—ZDB Options Exemplary Selections” on page A-37
- “EMC - Obtaining Disk Configuration Data” on page A-42
- “Listing and Unlocking Locked Backup Devices and Target Volumes” on page A-48

Running and Scheduling a ZDB Session

To run a ZDB session (ZDB to tape, ZDB to disk, or ZDB to disk+tape), use one of the following methods:

- Schedule a backup of an existing filesystem or disk image backup specification using the Data Protector scheduler.

  Refer to “Scheduling a ZDB Session” on page 4 and to the “Scheduling Unattended Backups” section in the *HP OpenView Storage Data Protector Administrator’s Guide* for general information on Data Protector schedule functionality.

- Start an interactive backup of an existing filesystem or disk image backup specification. You can start a backup using the Data Protector GUI or the Data Protector command-line interface.

  To start an interactive backup of an existing filesystem or disk image backup specification, refer to “Starting an Interactive ZDB Session” on page 5.

If you intend to perform concurrent ZDB sessions, refer also to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide* for information on limitations when running concurrent ZDB sessions.

If you intend to perform concurrent ZDB sessions using multiple application systems, refer also to the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide* for information on limitations when running concurrent ZDB sessions using multiple application systems.

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.

On XP, if the LVM Mirroring configuration is used, a warning message is issued on the Data Protector monitor during the backup, since the volume group LDEVs in the physical volume group on the application system do not have their BC pairs assigned. This warning message should be ignored.

Scheduling a ZDB Session

To schedule a filesystem or a disk image backup on a specific date, you can create a new backup specification, or modify an existing one. For detailed steps, refer to the online Help index keyword “scheduling backups on specific dates and times”.

Figure A-1 Scheduling a ZDB-to-Disk+Tape or a ZDB-to-Disk Session Using the Data Protector Scheduler—VA, EVA, and XP

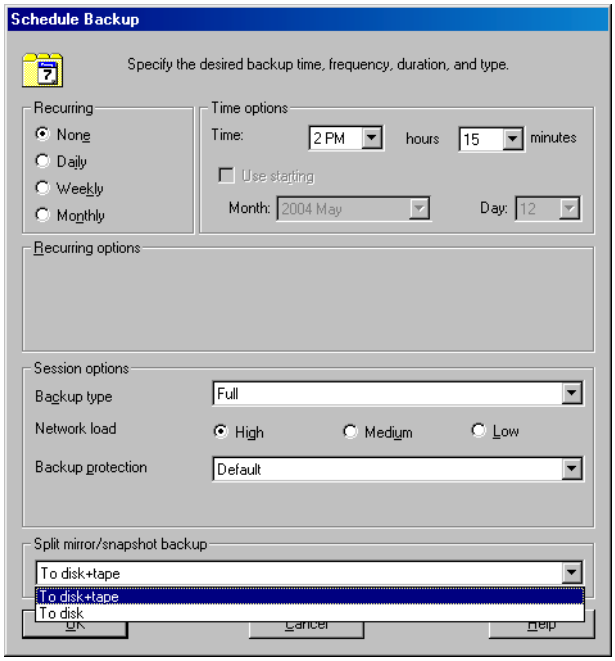

Figure A-2 **Scheduling a ZDB-to-Tape Session Using the Data Protector Scheduler—VA, EVA, XP, and EMC**

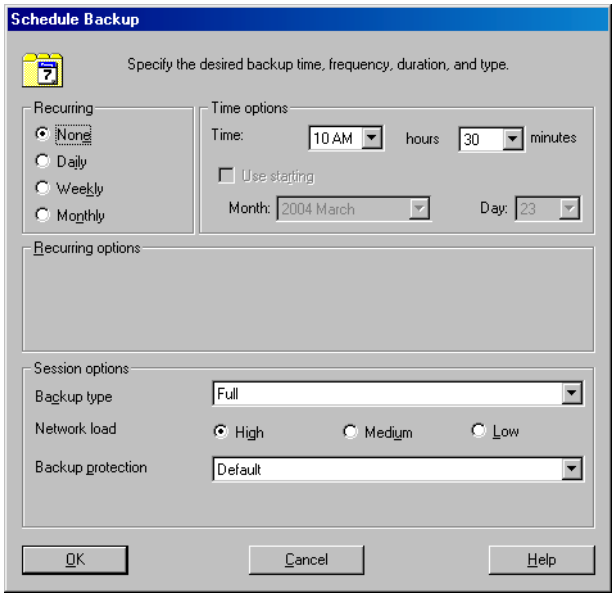

Starting an Interactive ZDB Session

Using the GUI

After you have saved your backup specification, you can start an interactive backup of an existing filesystem or disk image backup specification. Use the Data Protector GUI in the following way:

1. In the `Context List`, select `Backup`.
2. In the Scoping Pane, expand `Backup`, `Backup Specification`, and then `Filesystem`. Right-click on the backup specification for which you want to start an interactive backup and select `Start Backup`.
3. The `Start Backup` dialog box appears.

   The `Backup type` (`full` or `incr`) is ignored for ZDB sessions. It is always full.

   For information on `Network load`, press **F1**.

In the case of a ZDB-to-disk or a ZDB-to-disk+tape session on VA, EVA, or XP, specify the `Split mirror/snapshot backup` type as follows:

- To specify a **ZDB-to-disk** session, select `To disk` in the `Split mirror/snapshot backup` drop-down list.
- To specify a **ZDB-to-disk+tape** session, select `To disk+tape` in the `Split mirror/snapshot backup` drop-down list.

Figure A-3 Starting a ZDB-to-Disk or ZDB-to-Disk+Tape Session Interactively—VA, EVA, and XP

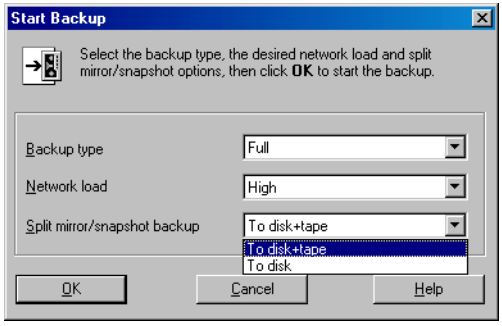

Figure A-4 Starting a ZDB-to-Tape Session Interactively—VA, EVA, XP, and EMC

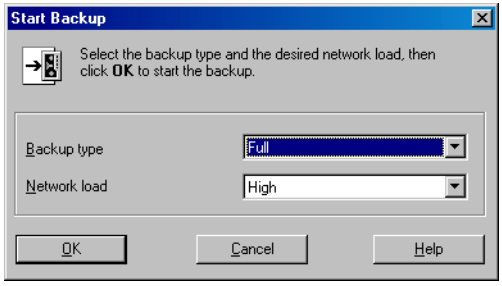

For more information on these options, press **F1**.

4. Click `OK` to start the backup session.

Using the CLI

To start a filesystem or a disk image **ZDB-to-tape** or **ZDB-to-disk+tape** session using the Data Protector CLI, use:

```
omnib -datalist <Name>
```

To start a filesystem or a disk image **ZDB-to-disk** session (VA, EVA, or XP only) using the Data Protector CLI, use:

```
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.

Alternate Paths Support

For systems configured with multiple host adapters and connections to a disk array, alternate paths solution perform dynamic load balancing while monitoring each path to ensure that the I/O is actually completing its transactions. In the event of a failure of any part of a path between the disk array and a server, alternate path software automatically switches to an alternate path, removing the failed path from the I/O rotation without any loss of data. The switchover is completely transparent to applications, so normal operation continue without downtime.

Data Protector ZDB integrations support the following alternate paths solutions:

	VA	**EVA**	**XP**	**EMC**
HP StorageWorks AutoPath (HP-UX and Windows)	Yes	No	Yes	No
HP StorageWorks Secure Path (HP-UX, Solaris, and Windows)	No	Yes	No	No
HP-UX Logical Volume Manager Alternate Links (HP-UX)	Yes	No	Yes	Yes
Sun Alternate Pathing Driver	No	No	Yes	No
Veritas Volume Manager Dynamic Multipathing	No	No	Yes	Yes
EMC Power Path	No	No	No	Yes

The Data Protector ZDB integrations alternate paths support means:

- Data Protector can handle situations where an alternate path solution automatically switches to an alternate path when a failure of any part of a path between a disk array and a server occurs.

- Data Protector, when creating target volumes in a replica, detects and uses the alternate path solution configured on the backup system for the newly created target volumes.
- The alternate paths load balancing is enabled for target volumes in a replica for the HP StorageWorks AutoPath (on VA and XP). This can improve performance during ZDB-to-tape and ZDB-to-disk+tape sessions.

  The HP StorageWorks AutoPath load balancing policy on the backup system can be controlled using the variables in the `omnirc` file. The variable used for this purpose is `OB2AUTOPATH_BALANCING_POLICY`. By default, Data Protector uses the HP StorageWorks AutoPath Round Robin load balancing policy.

  Refer to "ZDB Agents Omnirc Variables" on page A-23 for more information on the alternate paths load balancing policy control using the `omnirc` file.

HP StorageWorks AutoPath Limitations and Considerations

- On VA, due to certain HP StorageWorks AutoPath (AutoPath) limitations, additional configuration steps must be performed to ensure that the Data Protector VA integration using the AutoPath solution will work. Refer to "Configuring VA with HP StorageWorks AutoPath Installed" on page 9 for more information.
- On EVA, if the HP StorageWorks Secure Path is used, Data Protector uses the load balancing policy configured by the HP StorageWorks Secure Path; it is not possible to change the load balancing policy using Data Protector.
- During a Data Protector session, when the AutoPath `Shortest Queue Length` load balance policy is set, and if the failover to an alternate path occurs, the session will complete with errors.
- If backing up disk images without using raw logical volumes (rlvols) and if a failover to an alternate path occurs during the session, the session will complete with errors. If raw logical volumes (rlvols) are used, the failover to an alternate path is successful and the session is completed successfully.

Cluster Configurations

The Data Protector ZDB integrations provide support for:

- MC/ServiceGuard on HP-UX (for all supported disk arrays)
- Veritas Cluster on Solaris, and Microsoft Cluster Server (for VA, EVA, and XP)
- EMC GeoSpan for Microsoft Cluster Service (for EMC)

For more information on Data Protector cluster support and concepts, refer to the *HP OpenView Storage Data Protector Concepts Guide* and the *HP OpenView Storage Data Protector Administrator's Guide*.

The following sections provide an overview of supported Data Protector cluster ZDB configurations.

In figures A-5 to A-10, a Data Protector *application* backup configuration is presented. For *filesystem and disk image* backup, the Data Protector Disk Agent must be installed instead of a Data Protector integration software component, and an application database and binaries do not need to be installed as presented in the figures.

For an application in a cluster, a floating IP address can be used rather than a static one. This allows a successful start of a backup even after a local failover.

Client on the Application System in a Cluster

Configuration Behavior

In such a configuration the Data Protector Cell Manager is not installed in a cluster; only a Data Protector application client is installed in a cluster on the application system. If the application failover occurs during a Data Protector backup session, the backup session fails and the session must be restarted manually. If the application failover occurs before a Data Protector backup session is started, the session is completed successfully.

Components Distribution

The following must be installed on participating systems:

- The application binaries, Data Protector integration software component, and Data Protector disk array agent must be installed on the application system on all cluster nodes on local disks.

- The application database must be installed on the application system cluster shared disk. This cluster shared disk must be a disk array replicated disk.
- Data Protector integration software component, Data Protector disk array agent, and the Data Protector General Media Agent must be installed on the backup system on local disks.

Figure A-5 **Data Protector Client on the Application System in a Cluster**

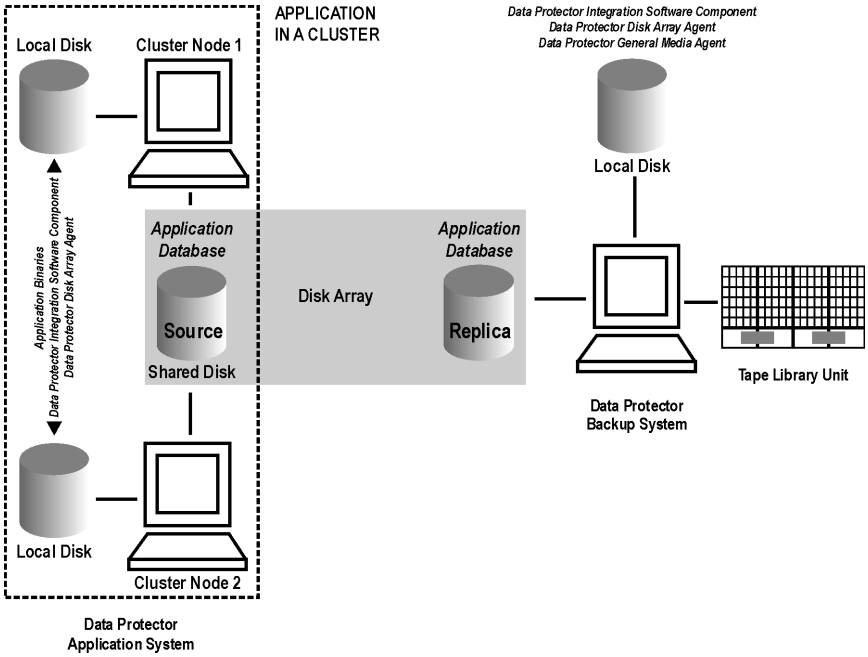

Cell Manager and Client on the Application System in a Cluster

Limitations

- This configuration is not supported on Veritas Cluster.
- On XP and EMC, split mirror restore is not possible in this configuration.

Configuration Behavior

In such a configuration, the Data Protector Cell Manager and an application are installed in a cluster on the application system.

If the application or Data Protector failover occurs during a Data Protector backup session, the failed backup session is automatically restarted.

If the application or Data Protector failover occurs before a Data Protector backup session is started, the session is completed successfully.

Components Distribution

The following must be installed on participating systems:

- The application binaries, Data Protector integration software component, and Data Protector disk array agent must be installed on the application system on all cluster nodes on local disks.
- The application database must be installed on its own application system cluster shared disk. This cluster shared disk must be a disk array replicated disk.
- The Data Protector Cell Manager must be installed on its own application system cluster shared disk.
- Data Protector integration software component, Data Protector disk array agent, and the Data Protector General Media Agent must be installed on the backup system on local disks.
- The Cell Manager cluster's critical resources must be configured in the same cluster group or package as those for the application being backed up.

Figure A-6 **Data Protector Cell Manager and Data Protector Client on the Application System in a Cluster**

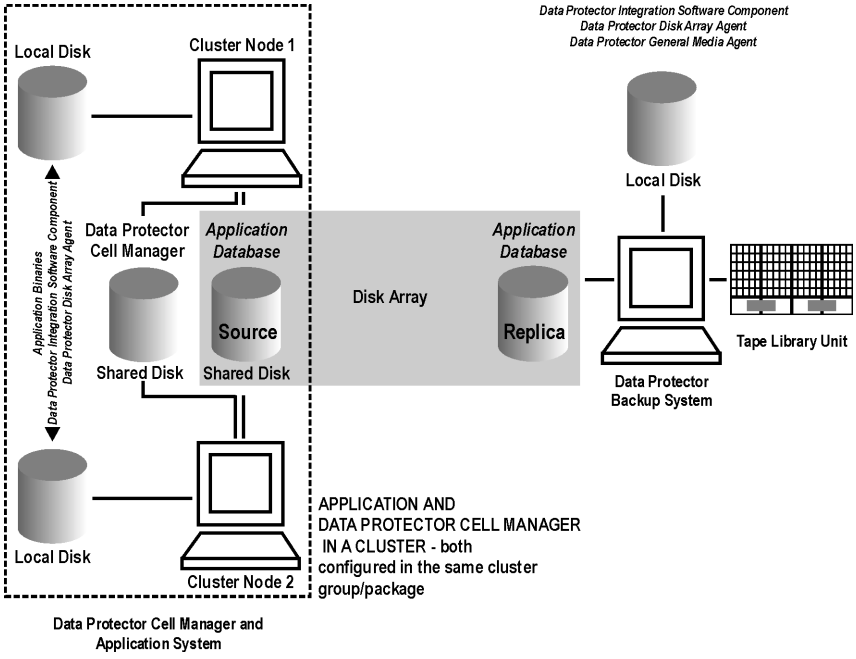

Client on the Application System in a Cluster, Cell Manager on the Backup System in a Cluster

Limitation

This configuration is not supported on Veritas Cluster.

Configuration Behavior

In such a configuration, the Data Protector Cell Manager is installed in a cluster on the backup system, and a Data Protector application client is installed in a cluster on the application system. If an application failover occurs during a Data Protector backup session, the backup session fails and the session must be restarted manually. If the application failover occurs before a Data Protector backup session, the session is completed successfully. If a Data Protector Cell Manager failover occurs during a Data Protector backup session, the failed backup session is automatically restarted, provided the appropriate Data Protector option (`Restart backup of failed objects`) is set.

Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for information on defining all possible Data Protector cluster-related options for the case of a failover of the Cell Manager.

If a Data Protector Cell Manager failover occurs before a Data Protector backup session is started, then the session is completed successfully.

Components Distribution

The following must be installed on participating systems:

- The application binaries, Data Protector integration software component, and Data Protector disk array agent must be installed on the application system on all cluster nodes on local disks.
- The application database must be installed on the application system cluster shared disk. This cluster shared disk must be a disk array replicated disk.
- The Data Protector Cell Manager must be installed on the backup system cluster shared disk.
- Data Protector integration software component, Data Protector disk array agent, and the Data Protector General Media Agent must be installed on the backup system on all cluster nodes on local disks.

Figure A-7 **Data Protector Client on the Application System in a Cluster, Data Protector Cell Manager on the Backup System in a Cluster**

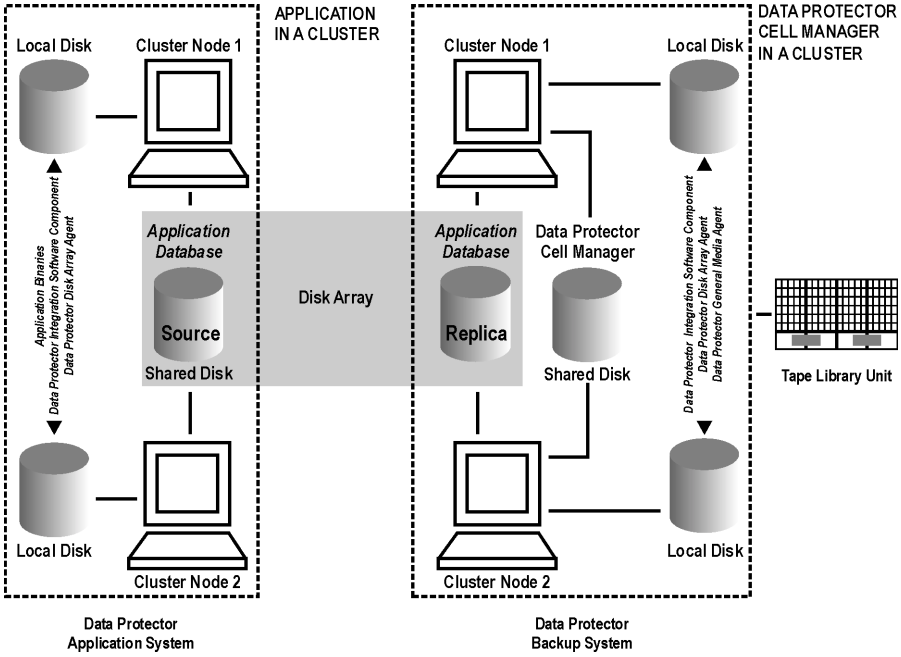

Cell Manager on the Backup System in a Cluster

Limitation

This configuration is not supported on Veritas Cluster.

Configuration Behavior

In such a configuration, the Data Protector Cell Manager is installed in a cluster on the backup system. If a Data Protector Cell Manager failover occurs during a Data Protector backup session, the failed backup session is automatically restarted, provided the appropriate Data Protector option (`Restart backup of failed objects`) is set.

Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for information on defining all possible Data Protector cluster-related options for the case of a failover of the Cell Manager.

If a Data Protector Cell Manager failover occurs before a Data Protector backup session is started, then the session is completed successfully.

Components Distribution

The following must be installed on participating systems:

- The application binaries, Data Protector integration software component, and Data Protector disk array agent must be installed on the application system.
- The application database must be installed on the application system. This disk must be a disk array replicated disk.
- The Data Protector Cell Manager must be installed on the backup system cluster shared disk.
- Data Protector integration software component, Data Protector disk array agent, and the Data Protector General Media Agent must be installed on the backup system on all cluster nodes on local disks.

Figure A-8 Data Protector Cell Manager on the Backup System in a Cluster

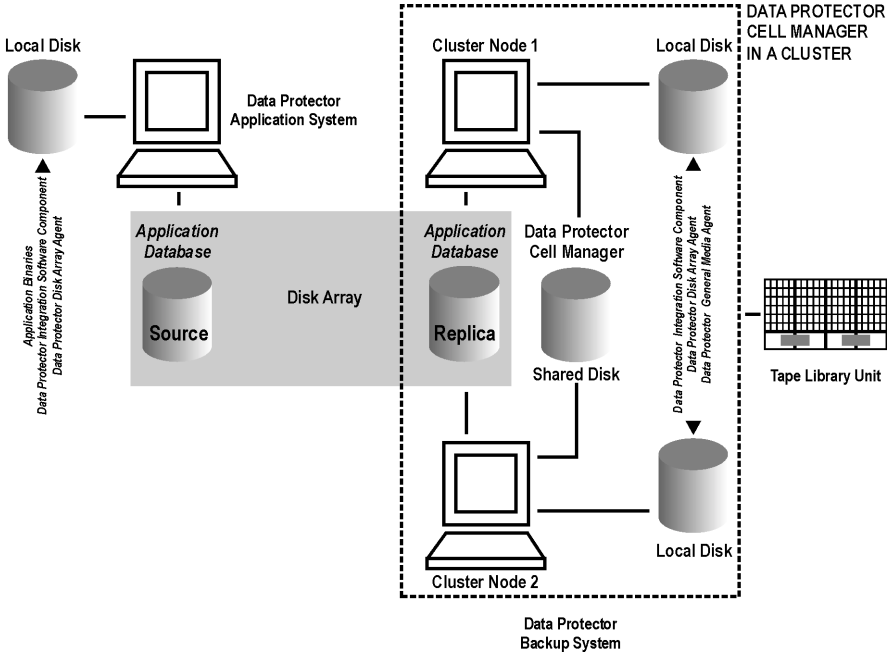

Client on the Application System in a Cluster, Cell Manager in a Cluster

Limitation

This configuration is not supported on Veritas Cluster.

Configuration Behavior

In such a configuration, the Data Protector Cell Manager is installed in a cluster on any system that is not a backup or an application system, and a Data Protector application client is installed in a cluster on the application system. If the application failover occurs during a Data Protector backup session, the backup session fails and the session must be restarted manually. If an application failover occurs before a Data Protector backup session is started, the session is completed successfully. If a Data Protector Cell Manager failover occurs during a Data Protector backup session, the failed backup session is automatically restarted, provided the appropriate Data Protector option (`Restart backup of failed objects`) is set.

Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for information on defining all possible Data Protector cluster-related options for the case of a failover of the Cell Manager.

If a Data Protector Cell Manager failover occurs before a Data Protector backup session is started, then the session is completed successfully.

Components Distribution

The following must be installed on participating systems:

- The application binaries, Data Protector integration software component, and Data Protector disk array agent must be installed on the application system on all cluster nodes on local disks.
- The application database must be installed on the application system cluster shared disk. This cluster shared disk must be a disk array replicated disk.
- The Data Protector Cell Manager must be installed on any system cluster shared disk.
- Data Protector integration software component, Data Protector disk array agent, and the Data Protector General Media Agent must be installed on the backup system on local disks.

Figure A-9 **Data Protector Client on the Application System in a Cluster, Data Protector Cell Manager in a Cluster**

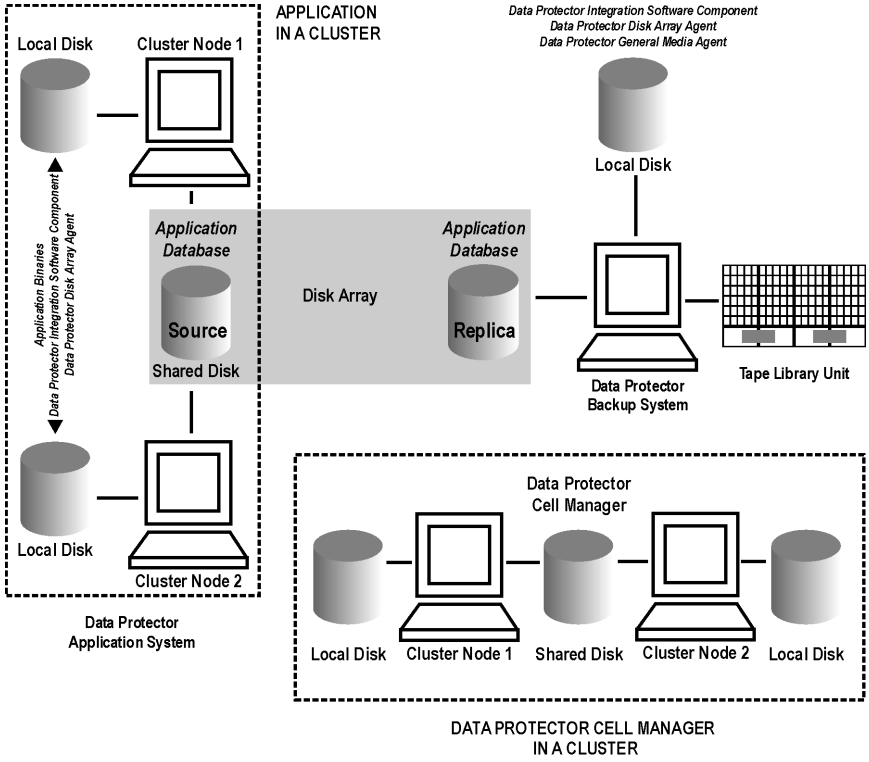

EMC GeoSpan for Microsoft Cluster Service Solution

In such a configuration, the EMC Symmetrix SRDF links are controlled by the EMC GeoSpan, whereas the EMC Symmetrix TF links are controlled by Data Protector.

Configuration Behavior

The Data Protector Cell Manager is not installed in a cluster; only a Data Protector application client is installed in a cluster on the application system. If the application or hardware failover occurs during a Data Protector backup session, the backup session fails. The failed session must be restarted manually and the backup system in the backup specification must be set as the backup system for the active node. If the application failover occurs before a Data Protector backup

session is started, the session is completed successfully if the backup system in the backup specification is set as the backup system for the active node.

Components Distribution

The following must be installed on participating systems:

- The application binaries, Data Protector integration software component, and EMC Agent must be installed on the application system on all cluster nodes on local disks.
- The application database must be installed on the application system cluster shared disk. This cluster shared disk must be a disk array replicated disk.
- Data Protector integration software component, EMC Agent, and the Data Protector General Media Agent must be installed on the backup system on local disks.

Figure A-10 EMC GeoSpan for Microsoft Cluster Service Solution

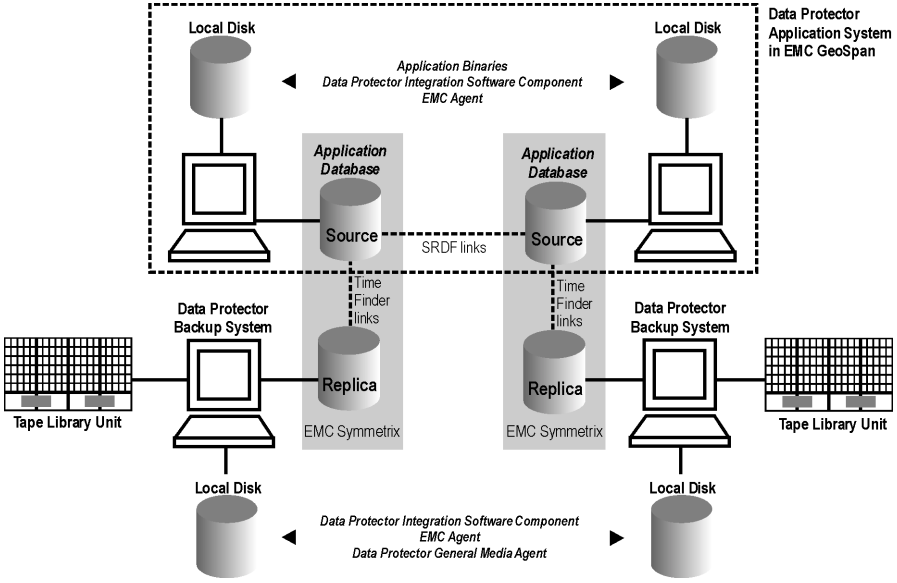

Instant Recovery in a Cluster

To perform a Data Protector instant recovery, when an application or a filesystem is running in an MC/ServiceGuard or a Microsoft Cluster Server on the application system, it is necessary to perform some *additional* steps.

IMPORTANT

If HP-UX LVM mirroring is also used, please refer also to “Instant Recovery and LVM Mirroring” on page 204.

MC/ServiceGuard Procedure

To perform a Data Protector instant recovery when an application or a filesystem is running in an MC/ServiceGuard on the application system, it is necessary to shut down the application without causing a failover (this involves stopping and the restarting the application cluster package). After the resynchronization to the application system has finished, it is also necessary to enable the replicated volume groups on the application system in the exclusive mode.

Prerequisite

Enabling the replicated volume groups on the application system in the exclusive mode is performed by Data Protector if the `ZDB_IR_VGCHANGE_A` variable on the application system is set appropriately (if it is set to `vgchange -a e`). Refer to “ZDB Agents Omnirc Variables” on page 23 for more information on how to set the variable in order to perform an instant recovery in an MC/ServiceGuard cluster.

Follow the procedure below to perform a Data Protector instant recovery to the application system in an MC/ServiceGuard cluster:

1. Stop the application cluster package:

   `cmhaltpkg <app_pkg_name>`

2. In order to be able to shutdown the application (database) running in the cluster without causing a failover, disable the execution of those lines in the *shell script for starting, shutting down and monitoring the database* that monitor the application processes, by commenting such lines (putting a # sign at the beginning of the line).

NOTE The following is an example of such a line in the Oracle8 *shell script for starting, shutting down and monitoring the database* as provided by MC/ServiceGuard (`ORACLE.sh`):

```
#set -A MONITOR_PROCESSES ora_pmon_${SID_NAME}
ora_dbw0_${SID_NAME} ora_ckpt_${SID_NAME}
ora_smon_${SID_NAME} ora_lgwr_${SID_NAME}
ora_reco_${SID_NAME} ora_arc0_${SID_NAME}
```

3. After the *shell script for starting, shutting down and monitoring the database* has been changed, restart the application cluster package:

   `cmrunpkg <app_pkg_name>`

4. Shutdown the application (database).
5. Using Data Protector, start the Data Protector instant recovery session. For a detailed procedure refer to one of the following sections:
   - “Instant Recovery Procedure” on page 51 (for VA)
   - “Instant Recovery Procedure” on page 127 (for EVA)
   - “Instant Recovery Procedure” on page 199 (for XP)

IMPORTANT When following the procedure and if you are performing instant recovery in an MC/ServiceGuard cluster to some other node than the one that was backed up, make sure to select the `Check the data configuration consistency` instant recovery option.

6. When the instant recovery session has finished, stop the application cluster package:

   `cmhaltpkg <app_pkg_name>`

7. Uncomment the lines (delete a # sign at the beginning of the line) that were commented in step 2 of this procedure, in order to re-enable a failover of the application.
8. After the *shell script for starting, shutting down and monitoring the database* has been changed, restart the application cluster package:

   `cmrunpkg <app_pkg_name>`
9. Recover the database. Refer to the database manufacturer's documentation for detailed procedures.

Microsoft Cluster Server Procedure

To perform a Data Protector instant recovery when an application or a filesystem is running in an MS Cluster Server on the application system, it is necessary to take the application cluster resource offline.

Follow the procedure below to perform a Data Protector instant recovery to the application system in an MS Cluster Server:

1. Using the Cluster Administrator utility or Cluster CLI, take the application cluster resource offline. Refer to the Microsoft Cluster Server documentation for detailed instructions.
2. Shutdown the application (database).
3. Using Data Protector, start the Data Protector instant recovery session. For a detailed procedure refer to one of the following sections:
   - "Instant Recovery Procedure" on page 51 (for VA)
   - "Instant Recovery Procedure" on page 127 (for EVA)
   - "Instant Recovery Procedure" on page 199 (for XP)
4. Restart the application (database).
5. Recover the database. Refer to the database manufacturer's documentation for detailed procedures.
6. Using the Cluster Administrator utility or CLI, put the application cluster resource online. Refer to the Microsoft Cluster Server documentation for detailed instructions.

ZDB Agents Omnirc Variables

The Data Protector ZDB agents 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 agents customizing. Refer to the *HP OpenView Storage Data Protector Administrator's Guide* for information on how to use the `omnirc` file.

For more information on `omnirc` file variables, refer to the following sections:

- "Common ZDB Agents Variables" on page A-23 (for common ZDB agents variables)
- "VA Agent Specific Variables" on page A-29 (for VA Agent specific variables)
- "EVA Agent Specific Variables" on page A-29 (for EVA Agent specific variables)
- "XP Agent Specific Variables" on page A-33 (for XP Agent specific variables)
- "EMC Agent Specific Variables" on page A-34 (for EMC Agent specific variables)

Common ZDB Agents Variables

This section explains the `omnirc` file variables that can be set for any ZDB agent.

ZDB_PRESERVE_MOUNTPOINTS: Data Protector creates unique mount points and mounts them automatically on the backup system.

The `ZDB_PRESERVE_MOUNTPOINTS` variable controls how the mount points on the backup system are created. Possible values are `0` and `1`. The default value is `0`.

NOTE The creation of mount points is also influenced by the ZDB_MULTI_MOUNT and the ZDB_MOUNT_PATH variables.

If the ZDB_PRESERVE_MOUNTPOINTS variable is set to 0, the mount point for the backed up filesystem is created in the following directories on the backup system:

- If the ZDB_MULTI_MOUNT is set to 1:
  - — On VA and EVA:

    <BU_MOUNT_PATH>/<application_system_name>/<mountpoint_name_on_application_system>_<SessionID>
  - — On XP:

    <BU_MOUNT_PATH>/<application_system_name>/<mountpoint_name_on_application_system>_<LDEV_MU#>
- If the ZDB_MULTI_MOUNT is set to 0 or not used:

  <BU_MOUNT_PATH>/<application_system_name>/<mountpoint_name_on_application_system>

where *<BU_MOUNT_PATH>* is:

- In the case of a UNIX Data Protector client:

  /var/opt/omni/tmp, if the ZDB_MOUNT_PATH is not specified or

  <ZDB_MOUNT_PATH>, if the ZDB_MOUNT_PATH is set to *<ZDB_MOUNT_PATH>*.
- In the case of a Windows Data Protector client:

  <Data_Protector_home>\tmp, if the ZDB_MOUNT_PATH is not specified or

  <ZDB_MOUNT_PATH>, if the ZDB_MOUNT_PATH is set to *<ZDB_MOUNT_PATH>*.

If the ZDB_PRESERVE_MOUNTPOINTS variable is set to 1,the mount point for the backed up filesystem is created in the:

/*<mountpoint_name_on_application_system>* (UNIX Data Protector client)

\<*mountpoint_name_on_application_system*> or <*Drive_letter_on_the_app_system*> (Windows Data Protector client)

directory or drive letter on the backup system.

IMPORTANT

In the cases stated below, the ZDB_PRESERVE_MOUNTPOINTS variable is set to 1 and its override is ignored; the ZDB_MULTI_MOUNT and the ZDB_MOUNT_PATH variables are ignored if this variable is set to 1:

disk images

Oracle8/9

SAP R/3

MS SQL Server 2000

ZDB_MULTI_MOUNT: This variable specifies, together with the ZDB_PRESERVE_MOUNTPOINTS and with the ZDB_MOUNT_PATH variables, how the mount points are created on the backup system.

NOTE

On VA, this variable is ignored and set to 1.

On EMC, this variable is ignored and set to 0.

If the ZDB_MULTI_MOUNT variable is set to 1, the Data Protector SessionID (VA and EVA) or LDEV MU# (XP) is appended at the end of mount point path on the backup system, thus enabling every group of mount points for one replica in the replica set to be mounted to their own mount points.

If the ZDB_MULTI_MOUNT variable is set to 0, the integration always mounts the selected group of mount points for one replica in the replica set to the same mount points on the backup system.

The default value is 1. The ZDB_MULTI_MOUNT variable is ignored if the ZDB_PRESERVE_MOUNTPOINTS is set to 1.

ZDB_MOUNT_PATH: This variable specifies, together with the ZDB_PRESERVE_MOUNTPOINTS and ZDB_MULTI_MOUNT variables, how the mount points are created on the backup system.

By default, the ZDB_MOUNT_PATH variable is not set.

If the ZDB_MOUNT_PATH variable is not set, the first part of the mount point path is set as:

- /var/opt/omni/tmp (UNIX Data Protector client), or
- <*Data_Protector_home*>\tmp (Windows Data Protector client).

Specify the first part of the mount point path to set this variable.

This variable is ignored if the ZDB_PRESERVE_MOUNTPOINTS is set to 1.

ZDB_ALWAYS_POST_SCRIPT: By default, the command/script set by the Restart the application option is not executed if the command/script set by the Stop/quiesce the application option fails.

The default value is 0. Possible values are 0 and 1.

If this variable is set to 1, the command/script set by the Restart the application option is always executed if set.

ZDB_BACKUP_VG_EXIST: For systems configured with multiple host adapters (HBA) and connections to a disk array, alternate paths solution perform dynamic load balancing. By default, during the preparation for backup and restore, Data Protector always creates a volume group with the disk on the first HBA as the primary path.

If you would like to disable autoconfiguration of volume groups on backup host and load balance the data read from the disk across multiple paths manually, set this variable to 1. Backup volume group that already exists will be used in the next backup or restore session.

NOTE If this variable is set on HP-UX systems, the volume groups are not removed from /etc/lvmtab on the backup system after each backup. Refer also to “Replica Management Options” on page 173.

The default value is 0.

ZDB_IR_VGCHANGE: This is an HP-UX instant recovery related variable. It specifies the mode in which the replicated volume groups on the application system are activated *after* the data is restored to the source volumes.

NOTE This variable is not supported on EMC.

NOTE The variable can be set on the application system only.

By setting this variable, the volume groups can be activated in the following modes:

- **exclusive**; the variable must be set as follows:
  `ZDB_IR_VGCHANGE_A=vgchange -a e`
- **shared**; the variable must be set as follows:
  `ZDB_IR_VGCHANGE_A=vgchange -a s`
- **normal** (default); the variable must be set as follows:
  `ZDB_IR_VGCHANGE_A=vgchange -q n -a y`

IMPORTANT Use the exclusive mode to enable instant recovery if an application/filesystem is running on the MC/ServiceGuard cluster on the application system.

SMB_SCAN_RDSK_TIMEOUT: During the backup system preparation, the system is scanned for new devices. When such device is detected, it appears on the backup system as a new physical drive. This variable sets the maximum time period (in seconds) during which a ZDB Agent on the backup system waits for a new physical drive to appear.

The variable is applicable only on Windows. The default value (30 seconds) is usually sufficient, unless there are some configuration problems on the backup system.

SMB_SCAN_FOR_VOLUME_TIMEOUT: This variable sets the maximum time period (in seconds) during which a ZDB Agent on the backup system waits for a new volume to appear on the backup system. This takes place after the corresponding physical drive was detected during the backup system preparation.

The variable is applicable only on Windows. The default value (300 seconds) is usually sufficient, unless there are some configuration problems on the backup system.

OB2AUTOPATH_BALANCING_POLICY: The HP StorageWorks AutoPath alternate paths load balancing is enabled for configured AutoPath alternate paths on the backup system. This can improve performance during ZDB-to-tape and ZDB-to-disk+tape sessions. This variable specifies the AutoPath load balancing policy used.

AutoPath provides enhanced data availability for systems configured with multiple host adapters and connections to a disk array. When several alternate paths are available, AutoPath will dynamically balance data load between the alternate paths to achieve optimum performance.

The default value is `1`.

Possible values are:

- `0` [none]–No Load Balance policy
- `1` [RR]–Round Robin policy
- `2` [SQL]–Shortest Queue Length policy (Default)

IMPORTANT

During a Data Protector session, when the AutoPath `Shortest Queue Length` load balance policy is set, and if the failover to an alternate path occurs, Data Protector will fail the session.

- `3` [SST]–Shortest Service Time policy

Refer to the HP StorageWorks AutoPath documentation for more information.

NOTE

This variable is not supported on EMC and EVA.

ZDB_IR_MANUAL_AS_PREPARATION: If you want to manually prepare the application system for instant recovery, set this variable to `1`. Manual preparation includes dismounting the filesystems that are to be restored and disabling the volume groups if they are configured. After an instant recovery session, you need to manually enable the volume groups and mount the filesystems. You can use this variable also when the automatic preparation of the application system for instant recovery has failed because the application data configuration has changed after the backup. For example, if a fail-over to a secondary cluster node occurred

between the backup session and the instant recovery session time, Data Protector may have problems with matching the secondary node resources to the resources that existed on the primary node during the backup time.

By default, this variable is not enabled (the value is `0`).

VA Agent Specific Variables

There are no specific VA `omnirc` file variables.

Refer to “Common ZDB Agents Variables” on page A-23 for explanations of `omnirc` file variables that can be set for the Data Protector VA Agent (SNAPA).

EVA Agent Specific Variables

This section explains the `omnirc` file variables that can be set for the Data Protector EVA Agent (EVAA).

Refer also to “Common ZDB Agents Variables” on page A-23 for explanations of other `omnirc` file variables that can be set for the Data Protector EVA Agent.

EVA_HOSTNAMEALIASES: This variable is used only when the `EVA_DISABLE_HBA_AUTODETECTION` variable is enabled or if the autodetection of the FC HBAs on the backup system fails.

As a part of a ZDB session, EVA Agent searches Command View EVA (CV EVA) for hostname objects that match the backup system hostname. The search is done by the hostname object name. By default, EVA Agent will only search for two names:

- full backup system hostname (as seen on the IP network)
- short backup system hostname (as seen on the IP network)

If you wish to add more hostnames to the search, you can do this by specifying hostname object names for this variable.

By default, no hostname objects are specified.

Example:

Suppose your backup host is represented within the CV EVA by the following two host objects:

- `/Hosts/Backup hosts/MyHost_Port1`
- `/Hosts/Backup hosts/MyHost_Port2`

You can force the Data Protector EVA client to find these two host objects by setting:

`EVA_HOSTNAMEALIASES=MyHost_Port1,MyHost_Port2`

EVA_DISABLE_HBA_AUTODETECTION: By default, EVA Agent attempts to perform a fully automatic backup system preparation during a ZDB-to-tape or ZDB-to-disk+tape session. This procedure is based on detection of the FC HBAs on the backup system, which is followed by an extensive search through the EVA storage system configuration. By setting this variable to 1, you can skip the automatic configuration of the backup system within the EVA storage system, which will improve performance. In this case, the backup system must be configured manually within the EVA storage system and the `EVA_HOSTNAMEALIASES` variable should be used to define the backup system object name.

By default, this variable is not enabled (the value is 0).

EVA_EMAPI_MAX_RETRY: Some Command View EVA commands that the Data Protector EVA client executes tend to fail when the target EVA system is under heavy load. In this case, Data Protector will wait for the period set by the `EVA_EMAPI_RETRY_DELAY` variable and try to execute the command again. This variable lets you define how many times Data Protector should retry before concluding that a command has failed. You might need to increase the default value when your EVA system is under extremely heavy load.

The default value is 10.

EVA_EMAPI_RETRY_DELAY: Some Command View EVA commands that the Data Protector EVA client executes tend to fail when the target EVA system is under heavy load. In this case, Data Protector will wait for the period set by the this variable (in seconds) and then try to execute the command again. Refer also to the variable `EVA_EMAPI_MAX_RETRY`, which sets the maximum number of retires.

The default value is 20.

TIP Since each Command View EVA command increases the load on your EVA system, it is usually better to increase the delay than the number of retries.

EVA_GETOBJID_MAX_RETRY: When creating new Command View EVA objects, the Data Protector EVA client sends the creation requests to Command View EVA and then collects data about the created objects. When the target EVA system is under heavy load, it may happen that object creation takes some time. Therefore, Data Protector has to retry to collect the created object until all of them are found. With this variable, you can define how many times Data Protector should retry before concluding that a command has failed. You might need to increase the default value when your EVA system is under extremely heavy load.

The default value is 10.

Refer also to the `EVA_GETOBJID_RETRY_DELAY` variable, which sets the delay between the retries.

EVA_GETOBJID_RETRY_DELAY: When creating new Command View EVA objects, the Data Protector EVA client sends the creation requests to Command View EVA and then collects data about the created objects. When the target EVA system is under heavy load, it may happen that object creation takes some time. Therefore, Data Protector has to retry to collect the data about the created objects until all of them are found. This variable lets you define how long (in seconds) Data Protector should wait between the retries.

The default value is 10.

Refer also to the `EVA_GETOBJID_MAX_RETRY` variable, which sets the number of retries before concluding that a command has failed.

TIP Since each Command View EVA command increases the load on your EVA system, it is usually better to increase the delay than the number of retries.

EVA_MSGWAITING_INTERVAL: This variable specifies the time interval (in minutes) between messages that report the progress of the snapclone creation process. Snapclone creation progress is monitored

only during the ZDB-to-tape and ZDB-to-disk+tape sessions immediately after preparation of the backup system. It will happen only if snapclones are selected as target replica type and the backup option `Delay the tape backup by a maximum of [XX] minutes if the snapclones are not fully created` is selected in the GUI.

By default, the message is displayed every 10 minutes.

EVA_CLONECREATION_QUERY_INTERVAL: This variable specifies the time interval (in minutes) between queries of the EVA storage system for checking the progress of the snapclone creation process. Such querying takes place during the ZDB-to-tape and ZDB-to-disk+tape sessions immediately after preparation of the backup system. It will happen only if snapclones are selected as target replica type and the backup option `Delay the tape backup by a maximum of [XX] minutes if the snapclones are not fully created` is selected in the GUI. A shorter time interval ensures that snapclone completion is detected more promptly, but it also increases load on the EVA storage system.

The default value is 5 minutes.

EVA_SCANDEVICES_LOCK_MAX_RETRY: After target volumes are created and presented to the backup system, a device scanning process is run on the backup system in order to detect the target volumes. Only one device scanning process is run at the same time. In case of parallel backup sessions, one of the sessions runs the scanning process, while the other ones wait for this process to finish before starting another scanning process.

The variable defines how many times EVA Agent should retry to run the device scanning process before aborting the backup session. EVA Agent waits for 10 seconds between the retries. On some systems, the device scanning process may take longer; therefore, you may want to increase the default setting.

The default value is 100.

XP Agent Specific Variables

This section explains the `omnirc` file variables that can be set for the XP Agent (SSEA).

Refer also to “Common ZDB Agents Variables” on page A-23 for explanations of other `omnirc` file variables that can be set for the XP Agent.

SSEA_SPLIT_REPORT_RATE: During the split of mirrored disks, the XP Agent checks the status of mirrored disks within a check interval specified by the `SSEA_SPLIT_SLEEP_TIME` variable for so many times as specified by the `SSEA_SPLIT_RETRY` variable. The `SSEA_SPLIT_REPORT RATE` variable specifies the rate of displaying the status of the mirrored disks to the Data Protector Monitor. For example, if the `SSEA_SPLIT_SLEEP_TIME` variable is set to 2 seconds and the `SSEA_SPLIT_REPORT_RATE` is set to 5, the status will be displayed to monitor every 5th check, that is every 10 seconds. The default value is 5.

SSEA_SPLIT_RETRY: During the split of mirrored disks, the XP Agent checks the status of mirrored disks within a check interval specified by the `SSEA_SPLIT_SLEEP_TIME` variable. The `SSEA_SPLIT_RETRY` variable sets the number of retries for these checks. The default value is 120. If there is no progress after the specified number of status checks, the split is aborted.

SSEA_SPLIT_SLEEP_TIME: During the split of mirrored disks, the XP Agent checks the status of mirrored disks for so many times as specified by the `SSEA_SPLIT_RETRY` variable. The `SSEA_SPLIT_SLEEP_TIME` variable sets the time interval (in seconds) between these status checks. The default value is 2.

SSEA_SYNC_REPORT_RATE: During the resynchronization of mirrored disks, the XP Agent checks the status of mirrored disks within a check interval specified by the `SSEA_SYNC_SLEEP_TIME` variable for so many times as specified by the `SSEA_SYNC_RETRY` variable. The `SSEA_SYNC_REPORT_RATE` variable specifies the rate of displaying the status of the mirrored disks to the `Data Protector` Monitor. For example, if the `SSEA_SYNC_SLEEP_TIME` variable is set to 5 seconds and the `SSEA_SPLIT_REPORT_RATE` is set to 2, the status will be displayed to monitor every 2nd check, that is every 10 seconds. The default value is 2.

SSEA_SYNC_RETRY: During the resynchronization of mirrored disks, the XP Agent checks the status of mirrored disks within a check interval specified by the `SSEA_SYNC_SLEEP_TIME` variable. The `SSEA_SYNC_RETRY`

variable sets the number of retries for these checks. The default value is 10. If there is no progress after the specified number of status checks, the resynchronization is aborted.

SSEA_SYNC_SLEEP_TIME: During the resynchronization of mirrored disks, the XP Agent checks the status of mirrored disks for so many times as specified by the `SSEA_SYNC_RETRY` variable. The `SSEA_SYNC_SLEEP_TIME` variable sets the time interval (in seconds) between these status checks. The default value is 5.

SSEA_WAIT_PAIRS_PROPER_STATUS: All disk pairs must be in proper status which is either `STAT_PSUS/SSUS` or `STAT_PAIR` before a process continues. This variable defines the maximum time period in minutes that application agent waits for disk pairs to change to proper status. The default value is 120 minutes.

EMC Agent Specific Variables

This section explains the `omnirc` file variables that can be set for the Data Protector EMC Agent (SYMA).

Refer also to “Common ZDB Agents Variables” on page A-23 for explanations of other `omnirc` file variables that can be set for the Data Protector EMC Agent.

SYMA_LOCK_RETRY, **SYMA_SLEEP_FOR_LOCK**, and **SYMA_REC_FILE_LIMIT**: Data Protector EMC integration related options. Each time the EMC Agent calls the WideSky library, it has to initiate the WideSky session, which locks the EMC Symmetrix database file. Other sessions have to wait to get the lock. The default value for the number of retries is 15 and the default sleep time is 30s. You can change this by setting the variable `SYMA_LOCK_RETRY=<number of retries>` and `SYMA_SLEEP_FOR_LOCK=<sleep time in seconds>`.

To successfully split the disks, the EMC Agent first checks the status of links. The links can be split only after all devices have been synchronized. The EMC Agent checks the status of links every 30 seconds and retries 15 times. To change this, set the variables `SYMA_SYNC_RETRY=<number of retries>` and `SYMA_SLEEP_FOR_SYNC=<sleep time in seconds>`.

These two variables are also used during incremental restore of device groups. The EMC Agent starts the incremental restore only when there are no longer any write pending tracks to devices in the restore device group. The EMC Agent checks the number of write pending track every

30 seconds and retries 15 times. To change this set the variables SYMA_SYNC_RETRY=*<number of retries>* and SYMA_SLEEP_FOR_SYNC=*<sleep time in seconds>*.

Invalid records are automatically deleted when the EMC Agent recovery file exceeds a certain value, by default, SYMA_REC_FILE_LIMIT = 102400 bytes.

SYMA_MOUNT_R2_READWRITE:

Possible values are 0 and 1. The default value is 0.

If this variable is set to 0, the volume groups and filesystems on the backup system are activated and mounted in read-only mode.

If this variable is set to 1, the volume groups and filesystem on the backup system will be activated in read/write mode.

For backup, it is sufficient to activate the backup system volume groups and filesystems in read-only mode. Should the mirror be used for DSS or other tasks after the backup, this might not be sufficient.

SYMA_SLEEP_FOR_SYNC:

Default: 30 seconds

Number of seconds, between last unsuccessful link status check (see SYMA_SYNC_RETRY) and next link status test.

SYMA_SYNC_RETRY:

Default: 15

A successful split and restore operation of the links (SRDF, TimeFinder) requires a valid status of the links (for example, devices not fully synchronized, existing write pending tracks,...). EMC Agent agent is sequentially checking for the invalid link statuses (before split or restore operation) until they are in the valid condition but not for more then SYMA_SYNC_RETRY times.

SYMA_UMOUNT_BEFORE_SPLIT:

Possible values are 0 and 1. The default value is 0.

If this variable is set to 1, the filesystem on the application system is dismounted before the split and mounted after the split. Sometimes, this is the only way to ensure that the data on the filesystem is consistent. A filesystem does not have a stop I/O functionality to flush the data from the filesystem cache to the disk and stop the I/O for the time of the split. The only way to back up a filesystem in split mirror mode is to dismount

the mount point on the application system. If an application (Oracle or SAP R/3) runs on the filesystem, it controls the I/O to the disk. In such a case it is not required to dismount the filesystem before the split.

If this variable is set to `0`, the filesystem on the application system is *not* dismounted before the split, therefore there is no need to mount it after the split.

User Scenarios—ZDB Options Exemplary Selections

This section provides descriptions of some exemplary backup policies, together with ZDB options to be selected to implement these policies.

VA and EVA Integrations

On VA, with all the examples, if HP StorageWorks Secure Manager Virtual Array is used, make sure that you set the password information using the `omnidbva` command and select the `Integrate with VA LUN security` option. For more information on setting the password information for the Secure Manager, refer to “LUN Security” on page 16.

Example 1

Data is to be backed up to tape media once a day (during the night). During the day, three copies of data need to be available for instant recovery.

To implement such backup policy, consider the following snapshot backup options selection:

- Select the `Track the replica for instant recovery` option
- Set the `Number of replicas rotated` option to `3`.

The following options are automatically set with the above options selected:

- `Use an existing replica` (on VA only)
- `Keep the replica after the backup`
- `Snapclone` (on EVA only)
- `Strict` (on EVA only)

Using the Data Protector scheduler and its `Split mirror/snapshot backup` drop-down list, you need to schedule the backup specification to start three ZDB-to-disk sessions during the day and one ZDB-to-disk+tape session every day during the night.

Example 2

Data is to be backed up to tape media every three hours. The replica created every three hours is used for data mining but not for instant recovery and can be obsoleted after these three hours.

To implement such backup policy, consider the following snapshot backup options selection:

- Do not select the `Track the replica for instant recovery` option.
- Select the `Keep the replica after the backup` option.
- On EVA, leave the `Number of replicas rotated` option set to `1`.
- On VA, select the `Leave the backup system enabled` option.
- Optionally, select the `Enable the backup system in read/write mode` option (on UNIX systems only).
- The `omnirc` file variable `ZDB_ORA_INCLUDE_CF_OLF` must be set to `1`. Refer to "ZDB Agents Omnirc Variables" on page A-23 for more information on this variable.

Using the Data Protector scheduler, schedule the backup specification to start a backup session every three hours every day.

On VA, after the first backup session is finished, change and save the backup specification by selecting the `Use an existing replica` option. If this option is not set after the first backup session, every new session for this backup specification will create a new replica without removing it.

Example 3

Data is to be backed up to tape media every three hours. The replica created every three hours must be available for instant recovery for 12 hours.

To implement such backup policy, consider the following snapshot backup options selection:

- Select the `Track the replica for instant recovery` option
- Set the `Number of replicas rotated` option to 4.

The following options are automatically set with the above options selected:

- `Use an existing replica` (on VA only)
- `Keep the replica after the backup`
- `Snapclone` (on EVA only)
- `Strict` (on EVA only)

Using the Data Protector scheduler and its `Split mirror/snapshot backup` drop-down list, schedule the backup specification to start eight ZDB-to-disk+tape sessions every three hours every day.

XP Integration

Example 1

A replica set is configured, replicas in the replica set are to be kept for instant recovery, next replica is to be prepared according to the replica set rotation after the backup and force-synchronized (if it is not already synchronized) before the next backup.

To implement such backup policy, the following XP backup options should be selected:

- `Keep the replica after the backup`
- `Track the replica for instant recovery`
- `At the end of the backup, prepare/resync the mirror disks for the next backup`
- `Force resync at the start of the backup session`

Example 2

A replica set is configured, *all* replicas in the replica set are to be used for offline data processing after the backup, next replica is to be prepared according to the replica set rotation after the backup, the next backup session is to be aborted if the data processing has not been finished at the start of the next backup session.

NOTE

This example is based on the assumption that offline data processing involves tasks such as splitting the links before data processing and resynchronizing the links when data processing is finished.

To implement such backup policy, the following XP backup options should be selected:

- `Keep the replica after the backup`
- `Leave the backup system enabled`
- `At the end of the backup, prepare/resync the mirror disks for the next backup`
- `Abort the session if the mirror disks are not synchronized`

Example 3

A replica set is configured, versions on replicas in the replica set are to be used for on-demand offline data processing (meaning that the links are split on demand and the backup system is prepared for offline data processing manually) after the backup and are not to be used for the instant recovery, a replica to be used in a backup session is to be prepared at the start of a backup session.

To implement such backup policy, the following XP backup options should be set:

- The `Keep the replica after the backup` option should not be selected
- The `Prepare/resync the mirror disks at the start of the backup` option should be selected.

Example 4

A single BC, CA or combined BC+CA replica is configured, version on the replica is to be used for offline data processing after the backup, the replica is to be prepared at the start of a backup session.

To implement such backup policy, the following XP backup options should be selected:

- `Keep the replica after the backup`
- `Leave the backup system enabled`
- `Prepare/resync the mirror disks at the start of the backup`

Conflicting Options

In case *a single* BC, CA or combined CA+BC replica is configured and the following two options are set:

- `Keep the replica after the backup`
- `At the end of the backup, prepare/resync the mirror disks for the next backup`

the second option is ignored, since the replica to be kept is at the same time the replica to be prepared for the next backup session.

NOTE

Such a conflicting situation can also arise even though a replica set is configured, depending on the replica set selection and on the XP LDEV exclude file.

EMC Integration

Example 1

After a split mirror backup session, the replica is to be discarded and prepared for the next backup at the end of the backup session.

To implement such backup policy, the following EMC backup options should be selected:

- `Re-establish links after backup`

and the following EMC backup option should *not* be selected

- `Re-establish links before backup`

Example 2

After a split mirror backup session, the replica is to be used for offline data processing and prepared for backup at the start of the next backup session.

To implement such backup policy, the following EMC backup options should be selected:

- `Re-establish links before backup`

and the following EMC backup option should *not* be selected

- `Re-establish links after backup`

EMC - Obtaining Disk Configuration Data

This section contains information on how to obtain disk information that is necessary during installation and configuration.

The following examples show how to choose and check EMC devices (disks) for the right type of connection (TimeFinder, SRDF, SRDF+TimeFinder):

To verify that the disks you intend to use are in the correct EMC configuration, run:

- `syminq` to display disk type, which can be blank, R1, R2, or BCV.
- `symbcv list` to display SLD-BCV pairs.
- `symrdf list` to display RDF1 - RDF2 pairs.

Example A-1 **Example 1**

The application system is connected to the Primary (R1) Symmetrix and the backup system to the Secondary (R2) Symmetrix. Disks 008 and 009 on the application system can be used for SRDF or combined SRDF+TimeFinder configuration. To verify that the disks are in the correct EMC configuration, proceed as follows:

1. Run `syminq` on the application system and search for disk numbers in the “Ser Num” column of the output.

Table A-1 **Finding the Disk Serial Number by Using `syminq`**

Device			Product		Device	
Name	**Type**	**Vendor**	**ID**	**Rev**	**Ser Num**	**Cap (KB)**
HP-UX: /dev/rdsk/c1t9d1 Windows: \\.\PHYSICALDRIVE1	R1	EMC	SYMMETRIX	5264	87008150	2817120
HP-UX: /dev/rdsk/c1t9d2 Windows: \\.\PHYSICALDRIVE2	R1	EMC	SYMMETRIX	5264	87009150	2817120

From the column “Type,” you can see that the disks are of type R1, which is required for SRDF and combined SRDF+TimeFinder configurations.

2. To check if the disks have the same serial number on the backup system or if the number differs, run `symrdf list` on the backup system.

Table A-2 Finding the Disk Serial Number by Using `symrdf list`

Local Device View									
Status Modes					RDF States				
Sym Dev	Rdev	RDF Typ:D	SA RA LNK	Mode Dom ACp	R1 Inv Tracks	R2 Inv Tracks	Dev	Rdev	Pair
008	008	R2:1	RW WD RW	SYN DIS OFF	0	0	WD	RW	Synch
009	009	R2:1	RW WD RW	SYN DIS OFF	0	0	WD	RW	Synch

You can see from the first two columns that the disks have the same numbers on both hosts.

3. When you find out the disk number on the backup system, you can query additional information by running `syminq` and look for disks 008 and 009.

4. If you have combined SRDF+Time Finder configuration, proceed as follows:

   a. Run `symbcv list` on the backup system to find associated BCVs.

Table A-3 Finding BCVs by Using `symbcv list`

BCV Device				Standard Device		Status	
Physical	Sym	RDF Att.	Inv. Tracks	Physical	Sym	Inv. Tracks	BCV<=>STD
HP-UX: c1t8d0 Windows: DRIVE5	038	+	0	HP-UX: c1t1d0 Windows: Not Visible	008	0	Synch

Table A-3 Finding BCVs by Using `symbcv list`

BCV Device				Standard Device		Status	
Physical	**Sym**	**RDF Att.**	**Inv. Tracks**	**Physical**	**Sym**	**Inv. Tracks**	**BCV<=>STD**
HP-UX: c1t8d1 Windows: DRIVE6	039	+	0	HP-UX: c1t1d1 Windows: Not Visible	009	0	Synch

From the output, you can see which BCV belongs to which SLD. The first few columns contain information about BCVs, the rest of the columns contain information about corresponding SLDs.

b. To be certain that the disks are correct, again run `syminq` on the backup system and search for BCVs under disk numbers 038 and 039. The disk you find should be BCV. The output of `syminq` should resemble the following:

Table A-4 Additional Check With `syminq` on the Backup System

Device			Product		Device	
Name	**Type**	**Vendor**	**ID**	**Rev**	**Ser Num**	**Cap (KB)**
HP-UX: /dev/rdsl/c1t8d0 Windows: \\PHYSICALDRIVE5	BCV	EMC	Symmetrix	5264	87038150	N/A
HP-UX: /dev/rdsl/c1t8d1 Windows: \\PHYSICALDRIVE6	BCV	EMC	Symmetrix	5264	87039150	N/A

Example A-2 **Example 2**

Both the application system and the backup system are connected to the same EMC box. Disks 048 and 049 on application system can be used for the TimeFinder configuration. To check if disks are in the expected EMC configuration, proceed as follows:

1. Run `syminq` on the application system and search for disk numbers in the “Ser Num” column of the output.

Table A-5 **Using `syminq` on the Application System to Find Disk Numbers**

Device			Product		Device	
Name	**Type**	**Vendor**	**ID**	**Rev**	**Ser Num**	**Cap (KB)**
HP-UX: /dev/rdsk/c0t0d0 Windows: \\.\PHYSICALDRIVE1		EMC	Symmetrix	5264	87048150	2817120
HP-UX: /dev/rdsk/c0t0d1 Windows: \\.\PHYSICALDRIVE2		EMC	Symmetrix	5264	87049150	2817120

From the column “Type,” you can see that the type of those disks is blank. However, it may also be R1 or R2, and the disks must have associated BCVs. These are all requirements for TimeFinder configurations.

2. Run `symbcv list` on the backup system and find your disk there.

Table A-6 Using `symbcv list` on the Backup System

BCV Device				Standard Device		Status	
Physical	**Sym**	**RDF att.**	**Inv. Tracks**	**Physical**	**Sym**	**Inv. Tracks**	**BCV<=>STD**
HP-UX: c0t5d0 Windows: DRIVE13	028	+	0	HP-UX: c0t10d0 Windows: Not Visible	048	0	Synch
HP-UX: c0t5d1 Windows: DRIVE14	029	+	0	HP-UX: c0t10d1 Windows: Not Visible	049	0	Synch

In the output, you can find which BCV belongs to SLD. First few columns hold information about BCVs and the rest about corresponding SLDs.

You can double-check the BCV by running `syminq` on the backup system. The disk you find should be BCV.

Table A-7 Additional Check With `syminq` on the Backup System

Device			Product		Device	
Name	**Type**	**Vendor**	**ID**	**Rev**	**Ser Num**	**Cap (KB)**
HP-UX: /dev/rdsk/c0t5d0 Windows: \\.\PHYSICALDRIVE5	BCV	EMC	Symmetrix	5264	17028150	2817120
HP-UX: /dev/rdsk/c0t5d1 Windows: \\.\PHYSICALDRIVE6	BCV	EMC	Symmetrix	5264	17029150	2817120

Additional Information for Troubleshooting

You can use the information in the previous section not only for installation and configuration, but also for troubleshooting the configuration. The following commands can also be used for obtaining information necessary for troubleshooting.

HP-UX

Check the created volume groups and find out which physical devices belong to which volume groups, by running the following commands:

On the application system:

- `strings /etc/lvmtab`

  The command displays all volume groups and devices that are in the volume groups on the host.

- `vgdisplay -v /dev/VG_name`

  The command displays which logical volumes and devices are in the volume group.

On the backup system:

- `/usr/symcli/bin/symdg list`

  The command displays device group names and additional information about them.

- `/usr/symcli/bin/symdg show DgName`

  The command displays more detailed information about devices and associated BCVs.

Windows

Run `symntctl` with additional parameters to get more information about disks, signatures, and drives. Refer to EMC documentation for more information.

On the backup system, run:

- `symdg list` - the command displays device group names and additional information about them
- `symdg show DgName` - the command displays detailed information about devices and associated BCVs

Listing and Unlocking Locked Backup Devices and Target Volumes

In certain situations (backup or restore session crashes), devices and target volumes remain locked, even though the General Media Agent or ZDB agents are not running. By default, such devices and target volumes are unlocked after 60 min.

You can list all locked devices and target volumes and unlock them.

The `omnidbutil` command options used for listing and unlocking backup devices and target volumes are: `-show_locked_devs` and `-free_locked_devs`.

To list locked backup devices and target volumes, run the following command:

```
omnidbutil -show_locked_devs
```

To unlock all backup devices and target volumes, run the following command:

```
omnidbutil -free_locked_devs
```

To unlock a specific backup device, run the following command:

```
omnidbutil -free_locked_devs <devname>
```

VA and EVA

To unlock a specific target volume, run the following command:

```
omnidbutil -free_locked_devs <wwn_lun>
```

where *<wwn>* is the VA or EVA world-wide-name and *<lun>* is the logical unit number (LUN).

XP

To unlock a specific target volume, run the following command:

```
omnidbutil -free_locked_devs <serial_ldev>
```

where *<serial>* is the XP serial number and *<ldev>* is the LDEV number.

For more information on the `omnidbutil` command, refer to its man page.

Glossary

access rights
See **user rights**.

ACSLS *(StorageTek specific term)*
The Automated Cartridge System Library Server (ACSLS) software that manages the Automated Cartridge System (ACS).

Active Directory *(Windows specific term)*
The directory service in a Windows network. It contains information about resources on the network and makes them accessible to users and applications. The directory services provide a consistent way to name, describe, locate, access and manage resources regardless of the physical system they reside on.

AML *(EMASS/GRAU specific term)*
Automated Mixed-Media library.

application agent
A component needed on a client to back up or restore online database integrations.
See also **Disk Agent**.

application system *(ZDB specific term)*
A system the application or database runs on. The application or database data is located on source volumes.
See also **backup system** and **source volume**.

archived redo log *(Oracle specific term)*
Also called offline redo log. If the Oracle database operates in the ARCHIVELOG mode, as each online redo log is filled, it is copied to one (or more) archived log destination(s). This copy is the Archived Redo Log. The presence or absence of an Archived Redo Log is determined by the mode that the database is using:

- ARCHIVELOG - The filled online redo log files are archived before they are reused. The database can be recovered from an instance and disk failure. The “hot” backup can be performed only when the database is running in this mode.

- NOARCHIVELOG - The filled online redo log files are not archived.

See also **online redo log.**

archive logging *(Lotus Domino Server specific term)*
Lotus Domino Server database mode where transaction log files are overwritten only after they have been backed up.

ASR Set
A collection of files stored on several diskettes, required for proper reconfiguration of the replacement disk

(disk partitioning and logical volume configuration) and automatic recovery of the original system configuration and user data that was backed up during the full client backup.
These files are stored as an ASR archive file on the Cell Manager (in *<Data_Protector_home>*\Config\Server\dr\asr on a Windows Cell Manager or in /etc/opt/omni/server/dr/asr/ on a UNIX Cell Manager) as well as on the backup medium. The ASR archive file is extracted to three diskettes for 32-bit Windows systems or four diskettes for 64-bit Windows systems after a disaster occurs. You need these diskettes to perform ASR.

autochanger
See **library**

autoloader
See **library**

BACKINT *(SAP R/3 specific term)*
SAP R/3 backup programs can call the Data Protector backint interface program via an open interface, which enables them to communicate with Data Protector software. For backup and restore, SAP R/3 programs issue orders for the Data Protector backint interface.

backup API
The Oracle interface between the Oracle backup/restore utility and the backup/restore media management layer. The interface defines a set of routines to allow the reading and writing of data to the backup media, the creation, searching and removing the backup files.

backup chain
This relates to a situation where full and incremental backups are performed. Based on the level of the incremental backups used (Incr, Incr 1, Incr 2, and so on), simple or rather complex dependencies of incrementals to previous incrementals can exist. The backup chain are all backups, starting from the full backup plus all the dependent incrementals up to the desired point in time.

backup device
A device configured for use with Data Protector, which can write data to and read data from storage media. This can be, for example, a standalone DDS/DAT drive or a library.

backup generation
One backup generation includes one full backup and all incremental backups until the next full backup.

backup ID
An identifier of an integration object that equals the session ID of the backup of this object. The backup ID is preserved when an object is copied, exported, or imported.

backup object
A backup unit that contains all items backed up from one disk volume (logical disk or mount point). The backed up items can be any number of files, directories, or the entire disk or mount point. Additionally, a backup object can be a database entity or a disk image (rawdisk).

A backup object is defined by:

- Client name: hostname of the Data Protector client where the backup object resides.
- Mount point: the access point in a directory structure (drive on Windows and mount point on UNIX) on the client where the backup object is located.
- Description: uniquely defines backup objects with identical client name and mount point.
- Type: backup object type (for example filesystem or Oracle).

backup owner
Each backup object in the IDB has an owner. The default owner of a backup is the user who starts the backup session.

backup session
A process that creates a copy of data on storage media. The activities are specified in a backup specification or an interactive session. All clients configured in one backup specification are backed up together in one backup session using the same backup type (full or incremental). The result of a backup session is a set of media, which was written to, also called the backup or media set.
See also **incremental backup** and **full backup.**

backup set
A complete set of integration objects associated with a backup.

backup set *(Oracle specific term)*
A logical grouping of backed up files that are created using the RMAN backup command. A backup set is a complete set of files associated with a backup. The files can be multiplexed to increase performance. A backup set contains either datafiles or archived logs, but not both together.

backup specification
A list of objects to be backed up, together with a set of devices or drives to be used, backup options for all objects in the specification, days and time that you want backups to be performed. The objects are entire disks/volumes or parts of them such as files, directories, or even the Windows

Registry for example. File selection lists such as include-lists and exclude-lists can be specified.

backup system *(ZDB specific term)*
A system connected to target volumes of one or multiple application systems. The backup system is typically connected to a backup device to perform the backup of the data in a replica.
See also **application system**, **target volume**, and **replica**.

backup types
See **incremental backup**, **differential backup**, **transaction backup**, **full backup** and **delta backup**.

backup view
Data Protector provides different views for backup specifications: By Type - according to the type of data available for backups/templates. Default view. By Group - according to the group to which backup specifications/templates belong. By Name - according to the name of backup specifications/templates. By Manager - if you are running MoM, you can also set the Backup view according to the Cell Manager to which backup specifications/templates belong.

BC *(EMC Symmetrix specific term)*
Business 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

maintain internal copies of HP StorageWorks Virtual Array LUNs for data backup or data duplication within the same virtual array. The copies (child or Business Copy LUNs) can be used for various purposes, such as backup, data analysis or development. When used for backup purposes, the original (parent) LUNs are connected to the application system and the Business Copy (child) LUNs are connected to the backup system.
See also **HP StorageWorks Virtual Array LUN**, **application system**, and **backup system**.

BCV *(EMC Symmetrix specific term)*
Business Continuance Volumes, or BCV devices, are dedicated SLDs that are pre-configured in the ICDA on which the business continuation operation runs. BCV devices are assigned separate SCSI addresses, differing from the addresses used by the SLDs they mirror. The BCV devices are used as splittable mirrors of the primary EMC Symmetrix SLDs that need to be protected.
See also **BC** and **BC Process**.

Boolean operators
The Boolean operators for the full text search functionality of the online Help system are AND, OR, NOT, and NEAR. Used when searching, they enable you to define your query precisely by creating a relationship between search terms. If no operator is specified in a multi-word search, AND is used by default. For example, the query manual disaster recovery is equivalent to manual AND disaster AND recovery.

boot volume/disk/partition
A volume/disk/partition with files required for the initial step of the boot process. Microsoft terminology defines the boot volume/disk/partition as a volume/disk/partition containing the operating system files.

BRARCHIVE *(SAP R/3 specific term)*
An SAP R/3 backup tool that allows you to archive redo log files. BRARCHIVE also saves all the logs and profiles of the archiving process.
See also **SAPDBA**, **BRBACKUP** and **BRRESTORE**.

BRBACKUP *(SAP R/3 specific term)*
An SAP R/3 backup tool that allows an online or offline backup of the control file, of individual data files, or of all tablespaces and, if necessary, of the online redo log files.
See also **SAPDBA**, **BRARCHIVE** and **BRRESTORE**.

BRRESTORE *(SAP R/3 specific term)*
An SAP R/3 tool that can be used to restore files of the following type:

- Database data files, control files, and online redo log files saved with BRBACKUP

- Redo log files archived with BRARCHIVE
- Non-database files saved with BRBACKUP

You can specify files, tablespaces, complete backups, log sequence numbers of redo log files, or the session ID of the backup.
See also **SAPDBA**, **BRBACKUP** and **BRARCHIVE**.

BSM
The Data Protector Backup Session Manager controls the backup session. This process always runs on the Cell Manager system.

CA *(HP StorageWorks Disk Array XP specific term)*
Continuous Access XP allows you to create and maintain remote copies of HP StorageWorks Disk Array XP LDEVs for purposes such as data duplication, backup, and disaster recovery. CA operations involve the main (primary) disk arrays and the remote (secondary) disk arrays. The main disk arrays contain the CA primary volumes (P-VOLs), which contain the original data and are connected to the application system. The remote disk arrays contain the CA secondary volumes (S-VOLs) connected to the backup system.
See also **BC** *(HP StorageWorks Disk Array XP specific term)*, **Main Control Unit** and **HP StorageWorks Disk Array XP LDEV**.

CAP *(StorageTek specific term)*
Cartridge Access Port is a port built into the door panel of a library. The purpose is to enter or eject media.

catalog protection
Defines how long information about backed up data (such as file names and file versions) is kept in the IDB.
See also **data protection**.

CDB
The Catalog Database is a part of the IDB that contains information about backups, object copies, restores, media management sessions, and backed up data. Depending on the selected logging level, it also contains file names and file versions. This part of the database is always local to the cell.
See also **MMDB**.

CDF file *(UNIX specific term)*
A Context Dependent File is a file consisting of several files grouped under the same pathname. The system ordinarily selects one of the files using the context of the process. This mechanism allows machine dependent executables, system data, and device files to work correctly from all hosts in a cluster while using the same pathname.

cell
A set of systems that are under the control of a Cell Manager. The cell typically represents the systems on a site or an organizational entity, which are connected to the same LAN. Central control is available to administer the backup and restore policies and tasks.

Cell Manager
The main system in the cell where the essential Data Protector software is installed and from which all backup and restore activities are managed. The GUI used for management tasks can be located on a different system. Each cell has one Cell Manager system.

centralized licensing
Data Protector allows you to configure centralized licensing for the whole enterprise environment consisting of several cells. All Data Protector licenses are installed and kept on the Enterprise Cell Manager system. You can then allocate licenses to specific cells to suit your needs.
See also **MoM**.

Centralized Media Management Database (CMMDB)
See **CMMDB**.

channel *(Oracle specific term)*
An Oracle Recovery Manager resource allocation. Every allocated channel starts a new Oracle process, which performs backup, restore, and recovery actions. The type of channel allocated determines the type of media used:

- type “disk”
- type ‘SBT_TAPE’

If the specified channel is type ‘SBT_TAPE’ and Oracle is integrated with Data Protector, the server process will attempt to read backups from or write data files to Data Protector.

circular logging *(Microsoft Exchange Server and Lotus Domino Server specific term)*
Circular logging is a Microsoft Exchange Server database and Lotus Domino Server database mode, in which the transaction log file contents are periodically overwritten after the corresponding data is committed to the database. Circular logging reduces disk storage space requirements.

client backup
A backup of all filesystems mounted on a client. Filesystems mounted on the client after the backup specification was created are not automatically detected.

client backup with disk discovery
A backup of all filesystems mounted on a client. When the backup starts, Data Protector discovers the disks on the clients. Client backup with disk

discovery simplifies backup configuration and improves backup coverage of systems that often mount or dismount disks.

client or **client system**
Any system configured with any Data Protector functionality and configured in a cell.

cluster-aware application
It is an application that supports the cluster Application Programming Interface. Each cluster-aware application declares its own critical resources (disk volumes (on Microsoft Cluster Server), volume groups (on MC/ServiceGuard), application services, IP names and addresses ...).

CMD Script for OnLine Server
(Informix specific term)
Windows CMD script that is created in INFORMIXDIR when Informix OnLine Server is configured. The CMD script is a set of system commands that export environment variables for OnLine Server.

CMMDB
The Data Protector Centralized Media Management Database (CMMDB) is the result of merging MMDBs from several cells in the MoM environment. It allows you to share high-end devices and media across multiple cells in a MoM environment. One cell can control the robotics, serving the devices that are connected to systems in other cells. The CMMDB must reside on the MoM Manager. A reliable network connection between the MoM cell and the other Data Protector cells is highly recommended
See also **MoM**.

COM+ Registration Database
(Windows specific term)
The COM+ Registration Database and the Windows Registry store COM+ application attributes, class attributes, and computer-level attributes. This guarantees consistency among these attributes and provides common operation on top of these attributes.

command-line interface
A set of DOS and UNIX like commands that you can use in shell scripts to perform Data Protector configuration, backup, restore, and management tasks.

Command View (CV) EVA *(HP StorageWorks EVA specific term)*
The user interface that allows you to configure, manage, and monitor your HP StorageWorks EVA storage system. It is used to perform various storage management tasks, for example, creating virtual disk families, managing storage system hardware, creating snapclones and snapshots of virtual disks. The Command View EVA software runs on the HP OpenView

Storage Management Appliance, and is accessed by a Web browser.
See also **HP StorageWorks EVA Agent (legacy)** and **HP StorageWorks EVA SMI-S Agent**.

concurrency
See **Disk Agent concurrency**.

control file *(Oracle and SAP R/3 specific term)*
An Oracle data file that contains entries specifying the physical structure of the database. It provides database consistency information used for recovery.

CRS
The Cell Request Server process (service) runs on the Data Protector Cell Manager. It starts and controls the backup and restore sessions. The service is started as soon as Data Protector is installed on the Cell Manager.
CRS runs under the account root on UNIX systems, and under any Windows account. By default, it runs under the account of the user, specified at installation time.

CSM
The Data Protector Copy Session Manager process controls the object copy session and runs on the Cell Manager system.

data file *(Oracle and SAP R/3* specific term)
A physical file created by Oracle that contains data structures such as tables and indexes. A data file can only belong to one Oracle database.

data protection
Defines how long the backed up data on media remains protected, that is, Data Protector will not overwrite it. When the protection expires, Data Protector will be able to reuse the media in one of the next backup sessions.
See also **catalog protection**.

Data Protector Event Log
A central repository of all Data Protector related notifications. By default, all notifications are sent to the Event Log. The Event Log is accessible only to Data Protector users in the Admin group and to Data Protector users who are granted the Reporting and notifications user rights. You can view or delete all events in the Event Log.

Data Protector user account
You can use Data Protector only if you have a Data Protector user account, which restricts unauthorized access to Data Protector and to backed up data. Data Protector administrators create this account specifying a user logon name, the systems from which the user can log on, and a Data Protector user group

membership. This is checked whenever the user starts the Data Protector user interface or performs specific tasks.

data stream
Sequence of data transferred over the communication channel.

database library
A Data Protector set of routines that enables data transfer between Data Protector and a server of an online database integration, for example, the Oracle Server.

database parallelism
More than one database is backed up at a time if the number of available devices allows you to perform backups in parallel.

database server
A computer with a large database stored on it, such as the SAP R/3 or Microsoft SQL database. A server has a database that can be accessed by clients.

Dbobject *(Informix specific term)*
An Informix physical database object. It can be a blobspace, dbspace, or logical-log file.

DC directory
The Detail Catalog (DC) directory consists of DC binary files, which store information about file versions. It represents the DCBF part of the IDB, which occupies approximately 80% of the IDB. The default DC directory is called the dcbf directory and is located in the <*Data_Protector_home*>\db40 directory on a Windows Cell Manager and in the /var/opt/omni/server/db40 directory on a UNIX Cell Manager. You can create more DC directories and locate them as appropriate to you. Up to 10 DC directories are supported per cell. The default maximum size of a DC directory is 4 GB.

DCBF
The Detail Catalog Binary Files (DCBF) part of the IDB stores information about file versions and attributes. It occupies approximately 80% of the IDB. One DC binary file is created for each Data Protector medium used for backup.

delta backup
A delta backup is a backup containing all the changes made to the database from the last backup of any type.
See also **backup types**

device
A physical unit which contains either just a drive or a more complex unit such as a library.

device chain
A device chain consists of several standalone devices configured for sequential use. When a medium in one

device gets full, the backup automatically continues on a medium in the next device in the device chain.

device group *(EMC Symmetrix specific term*)
A logical unit representing several EMC Symmetrix devices. A device cannot belong to more than a single device group. All devices in a device group must be on the same EMC Symmetrix unit. You can use a device group to identify and work with a subset of the available EMC Symmetrix devices.

device streaming
A device is streaming if it can feed enough data to the medium to keep it moving forward continuously. Otherwise, the tape has to be stopped, the device waits for more data, reverses the tape a little and resumes to write to the tape, and so on. In other words, if the data rate written to the tape is less or equal the data rate which can be delivered to the device by the computer system, then the device is streaming. Streaming significantly improves the performance of the device and use of space.

DHCP server
A system running the Dynamic Host Configuration Protocol (DHCP) providing dynamic configuration of IP addresses and related information.

differential backup
An incremental backup (incr) based on any previous Data Protector backup (full or any incremental), which must still be protected.
See **incremental backup**.

differential backup *(MS SQL specific term)*
A database backup that records only the data changes made to the database after the last full database backup.
See also **backup types**.

differential database backup
A differential database backup records only those data changes made to the database after the last full database backup.

direct backup
A SAN-based backup solution in which data movement directly from disk to tape (or to other secondary storage) is facilitated by the SCSI Extended Copy (Xcopy) command. Direct backup lessens the backup I/O load on systems in a SAN environment. The data movement is facilitated directly from disk to tape (or to other secondary storage) by the SCSI Extended Copy (XCopy) command. The command is provided by any element of the infrastructure including bridges, switches, tape libraries, and disk subsystems.
See also **XCopy engine**.

directory junction *(Windows specific term)*
Directory junctions use the reparse point concept of Windows. An NTFS 5 directory junction allows you to redirect a directory/file request to another location.

Directory Store (DS) *(Microsoft Exchange specific term)*
A part of the Microsoft Exchange Server directory. The Microsoft Exchange Server directory contains objects used by Microsoft Exchange applications in order to find and access services, mailboxes, recipients, public folders, and other addressable objects within the messaging system.
See also **Information Store (MDB)**.

disaster recovery
A process to restore a client's main system disk to a state close to the time when a (full) backup was performed.

Disk Agent
A component needed on a client to back it up and restore it. The Disk Agent controls reading from and writing to a disk. During a backup session, the Disk Agent reads data from a disk and sends it to the Media Agent, which then moves it to the device. During a restore session the Disk Agent receives data from the Media Agent and writes it to the disk.

Disk Agent concurrency
The number of Disk Agents that are allowed to send data to one Media Agent concurrently.

disk discovery
The detection of disks during client backup with disk discovery. During this backup, Data Protector discovers (detects) the disks that are present on the client — even though they might not have been present on the system when the backup was configured — and backs them up. This is particularly useful in dynamic environments, where configurations change rapidly. After the disks are expanded, each inherits all options from its master client object. Even if pre- and post-exec commands are specified once, they are started many times, once per each object.

disk group *(Veritas Volume Manager specific term)*
The basic unit of data storage in VxVM system. A disk group can consist of one or more physical volumes. There can be more than one disk group on the system.

disk image (rawdisk) backup
A high-speed backup where Data Protector backs up files as bitmap images. A disk image (rawdisk) backup does not track the files and directory structure stored on the disk, but stores a disk image structure on byte level. You

can perform a disk image backup of either specific disk sections or a complete disk.

disk quota
A concept to manage disk space consumption for all or a subset of users on a computer system. This concept is used by several operating system platforms.

disk staging
The process of backing up data in several phases to improve the performance of backups and restores, reduce costs of storing the backed up data, and increase the data availability and accessibility for restore. The backup stages consist of backing up data to one media type first (for example disk) and later copying it to a different media type (for example tape).

Distributed File System (DFS)
A service that connects file shares into a single namespace. The file shares can reside on the same or on different computers. DFS provides client access to the resources in a location-transparent manner.

DMZ
The Demilitarized Zone (DMZ) is a network inserted as a "neutral zone" between a company's private network (intranet) and the outside public network (Internet). It prevents outside users from getting direct access to company servers in the intranet.

DNS server
In the DNS client-server model, this is the server containing information about a portion of the DNS database that makes computer names available to client resolvers querying for name resolution across the Internet.

domain controller
A server in a network that is responsible for user security and verifying passwords within a group of other servers.

DR image
Data required for temporary disaster recovery operating system (DR OS) installation and configuration.

DR OS
A disaster recovery operating system is an operating system environment in which disaster recovery runs. It provides Data Protector a basic runtime environment (disk, network, tape, and filesystem access). The OS has to be installed and configured before the Data Protector disaster recovery can be performed. DR OS not only hosts the Data Protector disaster recovery process but is also a part of the restored system because it replaces its own configuration data with the original configuration data.

drive
A physical unit that receives data from a computer system and can write it onto a magnetic medium (typically a tape drive). It can also read the data from the medium and send it to the computer system.

drive index
A number that identifies the mechanical position of a drive inside a library device. This number is used by the robotic control to access a drive.

dynamic client
See **client backup with disk discovery**.

EMC Symmetrix Agent (SYMA)

(EMC Symmetrix specific term)
See **Symmetrix Agent (SYMA)**

emergency boot file *(Informix specific term)*
An Informix configuration file that resides in the <*INFORMIXDIR*>\etc directory (on HP-UX) or <*INFORMIXDIR*>/etc directory (on Windows) and is called ixbar.<*server_id*>, where <*INFORMIXDIR*> is the OnLine Server home directory and <*server_id*> is the value of the SERVERNUM configuration parameter. Each line of the emergency boot file corresponds to one backup object.

Enterprise Backup Environment
Several cells can be grouped together and managed from a central cell. The enterprise backup environment includes all clients located in several Data Protector cells which are managed and administered from a central cell using the Manager-of-Managers concept.
See also **MoM**.

Event Logs
Files in which Windows logs all events, such as the starting or stopping of services and the logging on and off of users. Data Protector can back up Windows Event Logs as part of the Windows configuration backup.

exchanger
Also referred to as SCSI Exchanger.
See also **library**.

exporting media
A process that removes all data about backup sessions, such as systems, objects, and file names, which reside on the media from the IDB. Information about the media and their relation to a pool is also removed from the IDB. The data on the media remains unchanged.
See also **importing media.**

Extensible Storage Engine (ESE)

(Microsoft Exchange Server 2000/2003 specific term)
A database technology used as a storage system for information exchange in Microsoft Exchange Server 2000/2003.

failover
Transferring of the most important cluster data, called group (on Windows) or package (on Unix) from one cluster node to another. A failover can occur mostly because of software or hardware failures or maintenance on the primary node.

FC bridge
See **Fibre Channel bridge**

Fibre Channel
An ANSI standard for high-speed computer interconnection. Using either optical or copper cables, it allows the high speed bidirectional transmission of large data files and can be deployed between sites kilometers apart.
Fibre Channel connects nodes using three physical topologies: point-to-point, loop, and switched.

Fibre Channel bridge
A Fibre Channel bridge or multiplexer provides the ability to migrate existing parallel SCSI devices, like RAID arrays, solid state disks (SSD), and tape libraries to a Fibre Channel environment. On one side of the bridge or multiplexer there is a Fibre Channel interface while on the other side there are parallel SCSI ports. The bridge or multiplexer enables SCSI packets to be moved between the Fibre Channel and parallel SCSI devices.

file depot
A file containing the data from a backup to a file library device.

file jukebox device
A device residing on disk consisting of multiple slots used to store file media.

file library device
A device which resides on a disk emulating a library with several media, hence containing multiple files, referred to as file depots.

File Replication Service (FRS)
A Windows service that replicates the domain controller store logon scripts and group policies. FRS also enables replication of Distributed File System (DFS) shares between systems and allows any server to perform replication activity.

file version
The same file can be backed up multiple times in case of full backups and incremental backups (if the file changed). If the log level ALL is selected for backup, Data Protector

retains one entry in the IDB for the filename itself and one for each version (date/time) of the file.

filesystem
The organization of files on a hard disk. A filesystem is backed up so that the file attributes and the file contents are stored on the backup media.

first level mirror (*HP StorageWorks Disk Array XP specific term*)
HP StorageWorks Disk Array XP allows up to three mirror copies of a Primary Volume and each of these copies can have additional two copies. The three mirror copies are called first level mirrors.
See also **Primary Volume**, and **MU numbers**.

fnames.dat
The fnames.dat files of the IDB contain information on the names of the backed up files. Typically, these files occupy about 20% of the IDB, if filenames are stored.

formatting
A process that erases any data contained on a medium and prepares it for use with Data Protector. Information about media (media ID, description, and location) is saved in the IDB as well as on the respective media (media header). Data Protector media with protected data are not formatted until the protection expires or the media are unprotected/recycled.

free pool
An auxiliary source of media for use by media pools when they run out of media. The media pools must be configured to use free pools.

full backup
A backup in which all selected objects are backed up, whether or not they have been recently modified.
See also **backup types**.

full database backup
A backup of all data in a database, not only the data that has been changed after the last (full or incremental) backup of the database. A full database backup does not depend on any other backup.

full mailbox backup
A full mailbox backup is a backup of the entire mailbox content.

global options file
A file that allows you to customize Data Protector. It explains the global options, which cover various aspects of Data Protector, typically time-outs and limits, and affect the entire Data Protector cell. The file is located in the /etc/opt/omni/server/options directory on HP-UX and Solaris systems and in the

<Data_Protector_home>\Config\Server\Options directory on Windows systems.

group *(Microsoft Cluster Server specific term*)
A collection of resources (for example disk volumes, application services, IP names and addresses) that are needed to run a specific cluster-aware applications.

GUI
A cross-platform (HP-UX, Solaris, and Windows) graphical user interface, provided by Data Protector for easy access to all configuration, administration, and operation tasks.

hard recovery *(Microsoft Exchange Server specific term)*
A Microsoft Exchange Server database recovery that is performed after a restore by the database engine, using transaction log files.

heartbeat
A cluster data set with a time stamp carrying information about the operational status of a particular cluster node. This data set or packet is distributed among all cluster nodes.

Hierarchical Storage Management (HSM)
A method for optimizing the use of expensive hard disk storage by migrating less frequently used data to less expensive optical platters. When needed, the data is migrated back to hard disk storage. This balances the need for fast retrieval from hard disk with the lower cost of optical platters.

Holidays file
A file that contains information about holidays. You can set different holidays by editing the Holidays file: /etc/opt/omni/server/Holidays on the UNIX Cell Manager and *<Data_Protector_home>*\Config\Server\holidays on the Windows Cell Manager.

host backup
See **client backup with disk discovery**.

hosting system
A working Data Protector client used for Disk Delivery Disaster Recovery with a Data Protector Disk Agent installed.

HP ITO
See **OVO**.

HP OpC
See **OVO**.

HP OpenView SMART Plug-In (SPI)
A fully integrated, out-of-the-box solution which "plugs into" HP OpenView Operations, extending the managed domain. Through the Data Protector integration, which is implemented as an HP OpenView

SMART Plug-In, a user can have an arbitrary number of Data Protector Cell Managers monitored as an extension to HP OpenView Operations (OVO).

HP OVO
See **OVO**.

HP StorageWorks Disk Array XP LDEV
A logical partition of a physical disk within an HP StorageWorks Disk Array XP. LDEVs are entities that can be replicated in the Continuous Access XP (CA) and Business Copy XP (BC) configurations, or can be used as standalone entities.
See also **BC** *(HP StorageWorks Disk Array XP specific term)*, **CA** *(HP StorageWorks Disk Array XP specific term)*, and **replica**.

HP StorageWorks EVA Agent (legacy)
A Data Protector software module that executes all tasks required for the HP StorageWorks Enterprise Virtual Array integration operating on HP StorageWorks EVA with Command View (CV) EVA software v3.1 or lower, and the EVA VCS firmware v3.01x or lower.
See also **Command View (CV) EVA** and **HP StorageWorks EVA SMI-S Agent**.

HP StorageWorks EVA SMI-S Agent
A Data Protector software module that executes all tasks required for the HP StorageWorks Enterprise Virtual Array integration operating on HP StorageWorks EVA with Command View (CV) EVA software starting with v3.2. With the EVA SMI-S Agent, the control over the array is established through HP StorageWorks SMI-S EVA provider, which directs communication between incoming requests and CV EVA.
See also **Command View (CV) EVA**, **HP StorageWorks SMI-S EVA provider**, and **HP StorageWorks EVA Agent (legacy)**.

HP StorageWorks SMI-S EVA provider
An interface used for controlling HP StorageWorks Enterprise Virtual Array. SMI-S EVA provider runs as a separate service on the HP OpenView Storage Management Appliance system and acts as a gateway between incoming requests and Command View EVA. With the Data Protector HP StorageWorks EVA integration, SMI-S EVA provider accepts standardized requests from the EVA SMI-S Agent, communicates with Command View EVA for information or method invocation, and returns standardized responses.
See also **HP StorageWorks EVA SMI-**

S Agent and **Command View (CV) EVA**.

HP StorageWorks Virtual Array LUN
A logical partition of a physical disk within an HP StorageWorks Virtual Array. LUNs are entities that can be replicated in the HP StorageWorks Business Copy VA configuration, or can be used as standalone entities.
See also **BC VA** and **replica**.

HP VPO
See **OVO**.

ICDA *(EMC Symmetrix specific term)*
EMC's Symmetrix Integrated Cached Disk Arrays (ICDA) is a disk array device that combines a set of physical disks, a number of FWD SCSI channels, an internal cache memory, and control and diagnostic software commonly referred to as the microcode.

IDB
The Data Protector Internal Database is an embedded database located on the Cell Manager that keeps information regarding which data is backed up, on which media it is backed up, how backup and restore sessions are run, and which devices and libraries are configured.

importing media
A process that re-reads all data about backup sessions which are on the medium back into the IDB. This then allows for fast and convenient access to the data on the media.
See also **exporting media.**

incremental backup
A backup that selects only files that have changed since a previous backup. Several levels of incremental backup are available, allowing selective backup of only files that have changed since the last incremental backup.
See also **backup types**.

incremental backup *(Microsoft Exchange Server specific term)*
A backup of the Microsoft Exchange Server data that has changed since the last full or incremental backup. With the incremental backup, only the transaction log files are backed up.
See also **backup types**.

incremental mailbox backup
An incremental mailbox backup backs up all the changes made to the mailbox after the last backup of any type.

incremental1 mailbox backup
An incremental1 mailbox backup backs up all the changes made to the mailbox after the last full backup.

incremental (re)-establish *(EMC Symmetrix specific term)*
A BCV or SRDF control operation. In BCV control operations, an incremental establish causes the BCV device to be synchronized incrementally and to function as an EMC Symmetrix mirrored medium. The EMC Symmetrix devices must have been previously paired.
In SRDF control operations, an incremental establish causes the target (R2) device to be synchronized incrementally and to function as an EMC Symmetrix mirrored medium. The EMC Symmetrix devices must have been previously paired.

incremental restore *(EMC Symmetrix specific term)*
A BCV or SRDF control operation.
In BCV control operations, an incremental restore reassigns a BCV device as the next available mirror of the standard device in the pair. However, the standard devices are updated with only the data that was written to the BCV device during the time of the original pair split, and the data that was written to the standard device during the split is overwritten with data from the BCV mirror. In SRDF control operations, an incremental restore reassigns a target (R2) device as the next available mirror of the source (R1) device in the pair. However, the source (R1) devices are updated with only the data that was written to the target (R2) device during the time of the original pair split, and the data that was written to the source (R1) device during the split is overwritten with data from the target (R2) mirror.

Inet
A process that runs on each UNIX system or service that runs on each Windows system in the Data Protector cell. It is responsible for communication between systems in the cell and for starting other processes needed for backup and restore. The Inet service is started as soon as Data Protector is installed on a system. The Inet process is started by the inetd daemon.

Information Store *(Microsoft Exchange Server 2000/2003 specific term)*
The Microsoft Exchange Server 2000/ 2003 service that is responsible for storage management. Information Store in Microsoft Exchange Server 2000/ 2003 manages two kinds of stores: mailbox stores and public folder stores. A mailbox store consists of mailboxes that belong to individual users. A public folder store contains public folders and messages that are shared among several users.
See also **Key Management Service** and **Site Replication Service**.

Information Store *(Microsoft Exchange Server 5.5 specific term)*
This is the default message store provider for the Microsoft Exchange Server 5.5. Information Store consists of the following stores:

- public information store
- private information store
- personal folder store
- offline information store.

The public information store contains public folders and messages that can be shared among multiple users and applications. A single public store is shared by all users within an Exchange Server 5.5 organization, even if multiple Exchange Servers are used. The private information store consists of mail boxes that can belong to users or to applications. The mail boxes reside on the server running the Exchange Server 5.5.
See also **Directory Store (DS)**.

initializing
See **formatting**.

Installation Server
A computer system that holds a repository of the Data Protector software packages for a specific architecture. The Installation Server is used for remote installation of Data Protector clients. In mixed environments at least two Installation Servers are needed: one for UNIX systems and one for Windows systems.

instant recovery (*ZDB specific term*)
A process in which a replica, produced by a ZDB-to-disk or a ZDB-to-disk+tape sessions, is used to restore the contents of the source volumes to their states at the time at which the replica was created, avoiding the need to perform a restore from tape. Depending on the application/database concerned, this may be all that is required, or other steps, such as the application of transaction log files, may be required for full recovery.
See also **replica**, **zero downtime backup (ZDB)**, **ZDB to disk**, and **ZDB to disk+tape**.

integrated security *(MS SQL specific term)*
Integrated security allows the Microsoft SQL Server to use Windows authentication mechanisms to validate Microsoft SQL Server logins for all connections. Using integrated security means that users have one password for both Windows and Microsoft SQL Server. Integrated security should be used in environments where all clients support trusted connections. Connections validated by Windows Server and accepted by Microsoft SQL

Server are referred to as trusted connections. Only trusted connections are allowed.

integration object
A backup object of a Data Protector integration, such as Oracle or SAP DB.

Internet Information Server (IIS)

(Windows specific term)
Microsoft Internet Information Server is a network file and application server that supports multiple protocols. Primarily, IIS transmits information in Hypertext Markup Language (HTML) pages by using the Hypertext Transport Protocol (HTTP).

IP address
Internet Protocol address is a numeric address of a system used to uniquely identify the system on the network. The IP address consists of four groups of numbers separated by periods (full stops).

ISQL *(Sybase specific term)*
A Sybase utility used to perform system administration tasks on Sybase SQL Server.

ITO
See **OVO**.

jukebox
See **library**.

jukebox device
A device consisting of multiple slots used to store either optical or file media. When being used to store file media, the jukebox device is known as the “file jukebox device”.

Key Management Service *(Microsoft Exchange Server 2000/2003 specific term)*
The Microsoft Exchange Server 2000/ 2003 service that provides encryption functionality for enhanced security.
See also **Information Store** and **Site Replication Service**.

LBO *(EMC Symmetrix specific term)*
A Logical Backup Object (LBO) is an object of data storage/retrieval in the EMC Symmetrix environment. It is stored/retrieved by EMC Symmetrix as one entity and can only be restored as a whole.

library
Also called autochanger, jukebox, autoloader, or exchanger. A library contains media in repository slots. Each slot holds one medium (for example, DDS/DAT). Media are moved between slots and drives by a robotic mechanism, allowing random access to media. The library can contain multiple drives.

lights-out operation or **unattended operation**
A backup or restore operation that takes

place outside of normal business hours without an operator. This implies that no operator personnel is present to work with the backup application or service mount requests, for example.

LISTENER.ORA *(Oracle specific term)*
An Oracle configuration file that describes one or more Transparent Network Substrate (TNS) listeners on a server.

load balancing
By default, Data Protector automatically balances the usage of devices selected for backup, so that they are used evenly. Load balancing optimizes the device usage by balancing the number of objects written to each device. Since load balancing is done automatically during backup time, you do not need to manage how the data is actually backed up. You just specify the devices to be used. If you do not want to use load balancing, you can select which device will be used for each object in the backup specification. Data Protector will access the devices in the specified order.

local and remote recovery
Remote recovery is performed if all Media Agent hosts specified in the SRD file are accessible. If any of them fails, the disaster recovery process fails over to the local mode. This means that the target system is searched for locally attached devices. If only one device is found, it is automatically used. Otherwise, Data Protector prompts you to select the device, which will be used for restore.

lock name
You can configure the same physical device several times with different characteristics, by using different device names.

The lock name is a user specified string that is used for locking all such device configurations to prevent collision if several such devices (device names) are used concurrently. Use an identical lock name for all device definitions which use the same physical device.

log_full shell script *(Informix UNIX specific term)*
A script provided by ON-Bar that you can use to start backing up logical-log files when OnLine Server issues a log-full event alarm. The Informix ALARMPROGRAM configuration parameter defaults to the <*INFORMIXDIR*>/etc/log_full.sh, where <*INFORMIXDIR*> is the OnLine Server home directory. If you do not want logical logs to be backed up continuously, set the ALARMPROGRAM configuration parameter to <*INFORMIXDIR*>/etc/no_log.sh.

logging level
The logging level determines the amount of details on files and directories written to the IDB during backup or object copying. You can always restore your data, regardless of the logging level used during backup. Data Protector provides four logging levels: Log All, Log Directories, Log Files, and No Log. The different logging level settings influence the IDB growth, backup speed, and the convenience of browsing data for restore.

logical-log files
This applies to online database backup. Logical-log files are files in which modified data is first stored before being flushed to disk. In the event of a failure, these logical-log files are used to roll forward all transactions that have been committed as well as roll back any transactions that have not been committed.

login ID *(MS SQL Server specific term)*
The name a user uses to log on to Microsoft SQL Server. A login ID is valid if Microsoft SQL Server has an entry for that user in the system table syslogin.

login information to the Oracle Target Database *(Oracle and SAP R/3 specific term)*
The format of the login information is <user_name>/<password>@<service>, where:

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

- <password> is a string used for data security and known only to its owner. Passwords are entered to connect to an operating system or software application. The password has to be the same as the password specified in the Oracle password file (orapwd), which is used for authentication of users performing database administration.

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

login information to the Recovery Catalog Database *(Oracle specific term)*
The format of the login information to the Recovery (Oracle) Catalog Database is <user_name>/

<password>@<service>, where the description of the user name, password, and service name is the same as in the Oracle SQL*Net V2 login information to the Oracle target database. In this case, <service> is the name of the service to the Recovery Catalog Database, not the Oracle target database.

Note that the Oracle user specified here has to be the owner of the Oracle Recovery (Oracle) Catalog.

Lotus C API *(Lotus Domino Server specific term)*
An interface for the exchange of backup and recovery information between Lotus Domino Server and a backup solution, like Data Protector.

LVM
A Logical Volume Manager is a subsystem for structuring and mapping physical disk space to logical volumes on UNIX systems. An LVM system consists of several volume groups, where each volume group has several volumes.

Magic Packet
See **Wake ONLAN**.

mailbox (*Microsoft Exchange Server specific term)*
The location to which e-mail is delivered, which is set up by the administrator for each user. If a set of personal folders is designated as the e-mail delivery location, e-mail is routed from the mailbox to this location.

Mailbox Store (*Microsoft Exchange Server 2000/2003 specific term)*
A part of the Information Store that maintains information about user mailboxes. A mailbox store consists of a binary rich-text .edb file and a streaming native internet content .stm file.

Main Control Unit (MCU) *(HP StorageWorks Disk Array XP specific term)*
An HP StorageWorks XP disk array that contains the primary volumes for the Continuous Access configuration and acts as a master device.
See also **BC** *(HP StorageWorks Disk Array XP specific term)*, **CA** *(HP StorageWorks Disk Array XP specific term)*, and **HP StorageWorks Disk Array XP LDEV**.

Manager-of-Managers (MoM)
See **Enterprise Cell Manager**.

Media Agent
A process that controls reading from and writing to a device, which reads from or writes to a medium (typically a tape). During a backup session, a Media Agent receives data from the Disk Agent and sends it to the device for writing it to the medium. During a restore session, a Media Agent locates data on the backup

medium and sends it to the Disk Agent. The Disk Agent then writes the data to the disk. A Media Agent also manages the robotics control of a library.

MAPI *(Microsoft Exchange specific term)*
The MAPI (Messaging Application Programming Interface) is the programming interface that lets applications and messaging clients interact with messaging and information systems.

media allocation policy
Determines in which sequence media are used for backup. The Strict allocation policy directs Data Protector to prompt for a specific medium. The Loose policy directs Data Protector to prompt for any suitable medium. The Formatted First policy directs Data Protector to give preference to unknown media, even if unprotected media are available in the library.

media condition
The quality of a medium as derived from the media condition factors. Heavy usage and age result in an increased number of read and write errors with tape media. Media need to be replaced when they are marked as POOR.

media condition factors
The user-assigned age threshold and overwrite threshold used to determine the state of a medium.

media ID
A unique identifier assigned to a medium by Data Protector.

media label
A user-defined identifier used to describe a medium.

media location
A user-defined physical location of a medium, such as "building 4" or "off-site storage".

media management session
A session performing some action on a medium, such as initializing, scanning the content, verifying data on a medium, or copying a medium.

media pool
A set of media of the same type (such as DDS) used and tracked as a group. Media are formatted and assigned to a media pool.

media set
The result of a backup session is data backed up on a group of media called media set. Depending on the media usage policy, several sessions can share the same media.

media type
The physical type of media, such as DDS or DLT.

media usage policy
The media usage policy controls how new backups are added to the already used media. It can be Appendable, Non-Appendable, or Appendable for incrementals only.

merging
This defines one mode to resolve file conflicts during restore. If the file to be restored already exists at the destination, the one with the more recent modification date is kept. Files not present on the disk are always restored. *See also* **overwrite**.

MFS
The Migrating File System enables a standard JFS filesystem with migration capabilities (on HP-UX 11.00). The MFS is accessed via a standard filesystem interface (DMAPI), it is mounted to a directory the same way as any HP-UX filesystem. In an MFS, only the superblock, the inode and the 'extended attribute' information remain permanently on the hard disk and are never migrated.
See also **VBFS**.

Microsoft Exchange Server
A “client-server” messaging and a workgroup system that offers a transparent connection to many different communication systems. It provides users with an electronic mail system, individual and group scheduling, online forms, and workflow automation tools. It provides a developer with a platform on which to build custom information-sharing and messaging-service applications.

Microsoft Management Console (MMC) *(Windows specific term)*
An administration model for Windows-based environments. It provides a simple, consistent, and integrated administration user interface allowing management of many applications through the same GUI, provided that the applications adhere to the MMC model.

Microsoft SQL Server 7.0/2000
A database management system designed to meet the requirements of distributed "client-server" computing.

Microsoft Volume Shadow Copy service (VSS)
A software service that provides a unified communication interface to coordinate backup and restore of a VSS-aware application regardless of its specific features. This service collaborates with the backup application, writers, shadow copy providers, and the operating system kernel to implement the management of volume shadow copies and shadow copy

sets.
See also **shadow copy, shadow copy provider, writer**.

mirror *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)*
See **target volume**.

mirror rotation *(HP StorageWorks Disk Array XP specific term)*
See **replica set rotation**.

MMD
The Media Management Daemon process (service) runs on the Data Protector Cell Manager and controls media management and device operations. The process is started when Data Protector is installed on the Cell Manager.

MMDB
The Media Management Database (MMDB) is a part of the IDB that contains information about media, media pools, devices, libraries, library drives, and slots configured in the cell, as well as the Data Protector media used for backup. In an enterprise backup environment, this part of the database can be common to all cells.
See also **CMMDB**, **CDB**.

MoM
Several cells can be grouped together and managed from a central cell. The management system of the central cell is the Manager-of-Managers (MoM). The MoM allows you to configure and manage multiple cells from a central point.

mount request
A screen prompt that tells you to insert a specific medium into a device. Once you respond to the mount request by providing the required medium and confirm the mount request, the session continues.

mount point
The access point in a directory structure for a disk or logical volume, for example /opt or d:. On UNIX, the mount points are displayed using the bdf or df command.

MSM
The Data Protector Media Session Manager, which runs on the Cell Manager and controls media sessions, such as copying media.

MU number (*HP StorageWorks Disk Array XP specific term*)
A Mirror Unit number is an integer number (0, 1 or 2), used to indicate a first level mirror.
See also **first level mirror**.

multi-drive server
A license that allows you to run an unlimited number of Media Agents on a

single system. This license, which is bound to the IP address of the Cell Manager, is no longer available.

obdrindex.dat
An IDB file with information about IDB backups, media, and devices used for the backup. This information can significantly simplify IDB recovery. It is recommended to relocate the file, together with IDB transaction logs, on a separate physical disk from other IDB directories, and, additionally, to make a copy of the file and locate it where you want.

OBDR capable device
A device that can emulate a CD-ROM drive loaded with a bootable disk and can thus be used as a backup or boot device for disaster recovery purposes.

object
See **backup object**

object copy
A copy of a specific object version that is created during an object copy session or a backup session with object mirroring.

object copy session
A process that creates an additional copy of the backed up data on a different media set. During an object copy session, the selected backed up objects are copied from the source to the target media.

object copying
The process of copying selected object versions to a specific media set. You can select object versions from one or several backup sessions to be copied.

Object ID *(Windows specific term)*
The object IDs (OIDs) enable access to NTFS 5 files no matter where in the system the files reside. Data Protector treats the OIDs as alternate streams of the files.

object mirror
A copy of a backup object created using object mirroring. Object mirrors are often referred to as object copies.

object mirroring
The process of writing the same data to several media sets during a backup session. Data Protector enables you to mirror all or some backup objects to one or more media sets.

offline backup
A backup during which an application database cannot be used by the application.

- For simple backup methods (non ZDB), the database is generally put into a quiescent state that allows use

by the backup system, but not the application, for the whole backup period (~minutes/hours). For instance, for backup to tape, until streaming of data to the tape is finished.

- For ZDB methods, the database is also put into the quiescent state, but for the period of the data replication process only (~seconds). Normal database operation can then be resumed for the rest of the backup process.

See also **zero downtime backup (ZDB)** and **online backup**.

offline recovery
Offline recovery is performed if the Cell Manager is not accessible, for example, due to network problems. Only standalone and SCSI library devices can be used for offline recovery. Recovery of the Cell Manager is always offline.

offline redo log
See **archived redo log**

OmniStorage
Software providing transparent migration of less frequently used data to the optical library while keeping more frequently used data on the hard disk. HP OmniStorage runs on HP-UX systems.

On-Bar *(Informix specific term)*
A backup and restore system for OnLine Server. ON-Bar enables you to create a copy of your OnLine Server data and later restore the data. The ON-Bar backup and restore system involves the following components:

- onbar utility
- Data Protector, as the backup solution
- XBSA interface
- ON-Bar catalog tables, which are used to back up dbobjects and track instances of dbobjects through multiple backups.

onbar utility *(Informix specific term)*
The Informix utility that communicates backup and restore requests to OnLine Server. The utility uses XBSA to exchange control data and back up and restore data with Data Protector.

ONCONFIG *(Informix specific term)*
An environment variable that specifies the name of the active ONCONFIG configuration file. If the ONCONFIG environment variable is not present, OnLine uses the configuration values from the file <*INFORMIXDIR*>/etc/onconfig (on HP-UX) or <*INFORMIXDIR*>\etc\onconfig (on Windows).

online backup
A backup performed while a database application remains available for use. The database is placed into a special backup mode of operation for the time period that the backup application requires access to the original data objects. During this period, the database is fully operational, but there may be a small performance impact and log files may grow very quickly.

- For simple backup methods (non ZDB), backup mode is required for the whole backup period (~minutes/ hours). For instance, for backup to tape, until streaming of data to tape is finished.

- For ZDB methods, backup mode is required for the short period of the data replication process only (~seconds). Normal database operation can then be resumed for the rest of the backup process.

In some cases, transaction logs may also have to be backed up to allow a consistent database to be restored.
See also **zero downtime backup (ZDB)** and **offline backup**.

online redo log *(Oracle specific term)*
Redo logs that have not been archived, but are either available to the instance for recording database activity or are filled and waiting to be archived or reused.
See also **archived redo log**.

OnLine Server *(Informix specific term)*
Refers to INFORMIX-OnLine Dynamic Server.

OpC
See **OVO**.

Oracle instance *(Oracle specific term)*
Each installation of an Oracle database on one or more systems. One computer system can have several instances of a database running.

ORACLE_SID *(Oracle specific term)*
A unique name for an Oracle Server instance. To switch among Oracle Servers, specify the desired <*ORACLE_SID*>. The <*ORACLE_SID*> is included in the CONNECT DATA parts of the connect descriptor in a TNSNAMES.ORA file and in the definition of the TNS listener in the LISTENER.ORA file.

original system
The system configuration backed up by Data Protector before a computer disaster hits the system.

overwrite
An option that defines one mode to resolve file conflicts during restore. All

files are restored from a backup even if they are older than existing files.
See also **merging**.

OVO
HP OpenView Operations for Unix provides powerful capabilities for operations management of a large number of systems and applications on in a network. Data Protector provides an integration into this management product. This integration is implemented as a SMART Plug-In for OVO management servers on HP-UX and Solaris. Earlier versions of OVO were called IT/Operation, Operations Center and Vantage Point Operations.
See also **merging**.

ownership
The ownership of a backup determines who can restore from the backup. The user who starts an interactive backup is the session owner. If a user starts an existing backup specification without modifying it, the session is not considered interactive. In that case, if the backup owner has been defined in the backup specification, they remain the session owner. Otherwise, the session owner becomes the user who started the backup in question. For the scheduled backups, by default, the session owner is for the UNIX Cell Manager: root.sys@<Cell Manager>, and for the Windows Cell Manager, the user that was specified during the installation of the Cell Manager. It is possible to modify the ownership, so that the specific user becomes the session owner.

P1S file
P1S file contains information on how to format and partition all disks installed in the system during Enhanced Automated Disaster Recovery (EADR). It is created during full backup and is saved on backup medium and on Cell Manager into <*Data_Protector_home*>\Config\Server\dr\p1s directory on a Windows Cell Manager or in /etc/opt/omni/server/dr/p1s directory on a UNIX Cell Manager with the filename recovery.p1s.

package *(MC/ServiceGuard and Veritas Cluster specific term)*
A collection of resources (for example volume groups, application services, IP names and addresses) that are needed to run a specific cluster-aware application.

pair status *(HP StorageWorks Disk Array XP specific term*)
A mirrored pair of disks can have various status values depending on the action performed on it. The three most important status values are:

- COPY - The mirrored pair is currently resynchronizing. Data is transferred from one disk to the other. The disks do not contain the same data.

- PAIR - The mirrored pair is completely synchronized and both disks (the primary volume and the mirrored volume) contain identical data.

- SUSPENDED - The link between the mirrored disks is suspended. That means that both disks are accessed and updated independently. However, the mirror relationship is still maintained and the pair can be resynchronized without transferring the complete disk.

parallel restore
Restoring backed up data to multiple disks at the same time (that is, in parallel) by running multiple Disk Agents, that receive data from one Media Agent. For the parallel restore to work, select data that is located on different disks or logical volumes and during backup, the data from the different objects must have been sent to the same device using a concurrency of 2 or more. During a parallel restore, the data for multiple objects selected for restore is read from media at the same time, thereby improving performance.

parallelism
The concept of reading multiple data streams from an online database.

physical device
A physical unit that contains either a drive or a more complex unit such as a library.

post-exec
A backup option that executes a command or script after the backup of an object or after the entire session completes. Post-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows and as shell scripts on UNIX.
See also **pre-exec**.

pre- and post-exec commands
Pre- and post-exec commands are used to perform additional action before and after a backup or restore session. They are not supplied by Data Protector. You need to create your own commands. They can be written as executables or batch files on Windows and as shell scripts on UNIX.

prealloc list
A subset of media in a media pool that specifies the order in which media are used for backup.

pre-exec
A backup option that executes a command or script before the backup of an object or before the entire session is started. Pre-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows and as shell scripts on UNIX.
See also **post-exec**.

Primary Volume (P-VOL) *(HP StorageWorks Disk Array XP specific term)*
Standard HP StorageWorks Disk Array XP LDEVs that act as a primary volume for the CA and BC configurations. The P-VOL is located in the MCU.
See also **Secondary Volume (S-VOL)**.

Private Information Store *(Microsoft Exchange Server 5.5 specific term)*
A part of the Information Store that maintains information in user mailboxes. A mailbox store consists of a binary rich-text .edb file.

protection
See **data protection** and also **catalog protection**.

public folder store (*Microsoft Exchange Server 2000/2003 specific term)*
The part of the Information Store that maintains information in public folders. A public folder store consists of a binary rich-text .edb file and a streaming native internet content .stm file.

public/private backed up data
When configuring a backup, you can select whether the backed up data will be:

- public, that is visible (and accessible for restore) to all Data Protector users
- private, that is, visible (and accessible for restore) only to the owner of the backup and administrators

RAID
Redundant Array of Inexpensive Disks.

RAID Manager Library *(HP StorageWorks Disk Array XP specific term)*
The RAID Manager Library is used internally by Data Protector on Solaris systems to allow access to HP StorageWorks Disk Array XP configuration, status, and performance data and to key HP StorageWorks Disk Array XP features through the use of function calls translated into a sequence of low level SCSI commands.

RAID Manager XP *(HP StorageWorks Disk Array XP specific term)*
The RAID Manager XP application

provides an extensive list of commands to report and control the status of the CA and BC applications. The commands communicate through a RAID Manager instance with the HP StorageWorks Disk Array XP Disk Control Unit. This instance translates the commands into a sequence of low level SCSI commands.

rawdisk backup
See **disk image backup**.

RCU *(HP StorageWorks specific term)*
The Remote Control Unit acts as a slave of an MCU in a CA configuration. In bidirectional configurations, the RCU can act as an MCU.

RDBMS
Relational Database Management System.

RDF1/RDF2 *(EMC Symmetrix specific term)*
A type of SRDF device group. Only RDF devices can be assigned to an RDF group. An RDF1 group type contains source (R1) devices and an RDF2 group type contains target (R2) devices.

RDS
The Raima Database Server process (service) runs on the Data Protector Cell Manager and manages the IDB. The process is started when Data Protector is installed on the Cell Manager.

Recovery Catalog *(Oracle specific term)*
A set of Oracle tables and views that are used by Recovery Manager to store information about Oracle databases. This information is used by Recovery Manager to manage the backup, restore, and recovery of Oracle databases. The recovery catalog contains information about:

- The physical schema of the Oracle target database
- Data file and archived log backup sets
- Data file copies
- Archived Redo Logs
- Stored scripts.

Recovery Catalog Database *(Oracle specific term)*
An Oracle database that contains a recovery catalog schema. You should not store the recovery catalog in your target database.

RecoveryInfo
When backing up Windows configuration files, Data Protector collects the information about the current system configuration (information on disk layout, volume,

and network configuration). This information is needed for disaster recovery.

Recovery Manager (RMAN) *(Oracle specific term)*
An Oracle command-line interface that directs an Oracle Server process to back up, restore, or recover the database it is connected to. RMAN uses either the recovery catalog or the control file to store information about backups. This information can be used later in restore sessions.

recycle
A process that removes the data protection from all backed up data on a medium, allowing Data Protector to overwrite it during one of the next backups. Data that belongs to the same session(s) but resides on other media is also unprotected. Recycling does not actually alter the data on the medium.

redo log *(Oracle specific term)*
Every Oracle database has a set of two or more redo log files. The set of redo log files for the database is known as the database's redo log. Oracle uses the redo log to record all changes made to data.

Remote Control Unit *(HP StorageWorks Disk Array XP specific term)*
The Remote Control Unit (RCU) acts as a slave of an MCU in a CA configuration. In bidirectional configurations, the RCU can act as an MCU.

Removable Storage Management Database *(Windows specific term)*
A Windows service used for managing removable media (such as tapes and disks) and storage devices (libraries). Removable Storage allows applications to access and share the same media resources.

reparse point *(Windows specific term)*
A system-controlled attribute that can be associated with any directory or file. The value of a reparse attribute can have user-controlled data. The format of the data is understood by the application that stored the data and a filesystem filter that was installed to interpret the data and process such files. Whenever the filesystem encounters a file with a reparse point, it attempts to find the filesystem filter associated with the data format.

replica *(ZDB specific term)*
An image, at a particular point in time, of the data in source volumes that contain user-specified backup objects. Depending on the hardware/software with which it is created, the image may be an independent exact duplicate (clone) of the storage blocks at the physical disk level (for example, a split mirror), or a virtual copy (for example, a

snapshot). From a host's perspective, on a basic UNIX or Windows system, the complete physical disk containing a backup object is replicated. However, if a volume manager is used on UNIX, the whole volume/disk group containing a backup object is replicated.
See also **snapshot**, **snapshot creation**, **split mirror**, and **split mirror creation**.

replica set *(ZDB specific term)*
A group of replicas, all created using the same backup specification.
See also **replica** and **replica set rotation**.

replica set rotation *(ZDB specific term)*
The use of a replica set for regular backup production: Each time the same backup specification requiring the use of a replica set is run, a new replica is created and added to the set, until the maximum number for the set is reached. After that, the oldest replica in the set is replaced, maintaining the maximum number of replicas in the set.
See also **replica** and **replica set**.

restore session
A process that copies data from backup media to a client.

RMAN *(Oracle specific term)*
See **Recovery Manager**.

RSM
The Data Protector Restore Session Manager controls the restore session. This process always runs on the Cell Manager system.

RSM *(Windows specific term)*
Removable Storage Manager (RSM) includes a media management service that facilitates communication among applications, robotic changers, and media libraries. It enables multiple applications to share local robotic media libraries and tape or disk drives and to manage removable media.

SAPDBA *(SAP R/3 specific term)*
An SAP R/3 user interface that integrates the BRBACKUP, BRARCHIVE, and BRRESTORE tools.

scan
A function that identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library).

scanning
A function which identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library). It is useful to perform a scan and check the actual media in the

device if someone has manually manipulated media without using Data Protector to eject or enter, for example.

Scheduler
A function that controls when and how often automatic backups occur. By setting up a schedule, you automate the start of backups.

Secondary Volume (S-VOL) *(HP StorageWorks Disk Array XP specific term*)
Secondary Volumes, or S-VOLs, are XP LDEVs that act as a secondary CA or BC mirror of another LDEV (a P-VOL). In the case of CA, S-VOLs can be used as failover devices in a MetroCluster configuration. The S-VOLs are assigned separate SCSI addresses, different from the addresses used by the P-VOLs.
See also **Primary Volume (P-VOL).**

session
See **backup session**, **media management session,** and **restore session**.

session ID
An identifier of a backup, restore, object copy, or media management session, consisting of the date when the session ran and a unique number.

session key
This environment variable for the Pre- and Post-exec script is a Data Protector unique identification of any session, including preview sessions. The session key is not recorded in the database, and it is used for specifying options for the omnimnt, omnistat and omniabort CLI commands.

shadow copy *(MS VSS specific term)*
A volume that represents a duplicate of the original volume at a certain point in time. The data is then backed up from the shadow copy and not from the original volume. The original volume continues to changes as the backup process continues, but the shadow copy of the volume remains constant.
See also **Microsoft Volume Shadow Copy service**.

shadow copy provider *(MS VSS specific term)*
An entity that performs the work on creating and representing the volume shadow copies. Providers own the shadow copy data and expose the shadow copies. Providers can be software (for example, system providers) or hardware (local disks, disk arrays).
See also **shadow copy.**

shadow copy set *(MS VSS specific term)*
A collection of shadow copies created at the same point in time.
See also **shadow copy**.

shared disks
A Windows disk on another system that has been made available to other users on the network. Systems with shared disks can be backed up without a Data Protector Disk Agent installed.

SIBF
The Serverless Integrations Binary Files (SIBF) is a part of the IDB that stores raw NDMP meta data. This data is necessary to perform restore of NDMP objects.

Site Replication Service *(Microsoft Exchange Server 2000/2003 specific term)*
The Microsoft Exchange Server 2000/ 2003 service that permits compatibility with Microsoft Exchange Server 5.5 by emulating the Exchange Server 5.5 directory service.
See also **Information Store** and **Key Management Service**.

slot
A mechanical position in a library. Each slot can hold a medium, such as a DLT tape. Data Protector references each slot with a number. To read a medium, a robotic mechanism moves the medium from a slot into the drive.

SMB
See **split mirror backup**.

SMBF
The Session Messages Binary Files (SMBF) part of the IDB stores session messages generated during backup, object copy, restore, and media management sessions. One binary file is created per session. The files are grouped by year and month.

snapshot *(HP StorageWorks VA and HP StorageWorks EVA specific term)*
A form of replica produced using snapshot creation techniques. A range of snapshot types is available, with different characteristics, depending on the arrays/techniques used. Such replicas are dynamic and may be either virtual copies, still reliant upon the contents of the source volumes, or independent exact duplicates (clones), depending on the snapshot type and the time since creation.
See also **replica** and **snapshot creation**.

snapshot backup *(HP StorageWorks VA and HP StorageWorks EVA specific term)*
See **ZDB to tape**, **ZDB to disk**, and **ZDB to disk+tape**.

snapshot creation *(HP StorageWorks VA and HP StorageWorks EVA specific term)*
A replica creation technique, in which copies of source volumes are created using storage virtualization techniques. The replicas are considered to be created

at one particular point-in-time, without pre-configuration, and are immediately available for use. However background copying processes normally continue after creation.
See also **snapshot**.

source (R1) device *(EMC Symmetrix specific term)*
An EMC Symmetrix device that participates in SRDF operations with a target (R2) device. All writes to this device are mirrored to a target (R2) device in a remote EMC Symmetrix unit. An R1 device must be assigned to an RDF1 group type.
See also **target (R2) device**.

source volume *(ZDB specific term)*
A storage volume containing data to be replicated.

sparse file A file that contains data with portions of empty blocks. Examples are: -A matrix in which some or much of the data contains zeros -files from image applications -high-speed databases If sparse file processing is not enabled during restore, it might be impossible to restore this file.

split mirror *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)*
A replica created using split mirror techniques. Such a replica provides an independent, exact duplicate, or clone, of the contents of the source volumes.
See also **replica** and **split mirror creation**.

split mirror backup *(EMC Symmetrix specific term)*
See **ZDB to tape**.

split mirror backup *(HP StorageWorks Disk Array XP specific term)*
See **ZDB to tape**, **ZDB to disk**, and **ZDB to disk+tape**.

split mirror creation *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)*
A replica creation technique, in which a pre-configured set of target volumes (a mirror) is kept synchronized with a set of source volumes until the time at which a replica of the contents of the source volumes is required. Then, the synchronization is stopped (the mirror is split) and a split mirror replica of the source volumes at the time of the split remains in the target volumes.
See also **split mirror**.

split mirror restore *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)*
A process in which data backed up in a ZDB-to-tape or a ZDB-to-disk+tape session is restored from tape media to a split mirror replica, which is then synchronized to the source volumes. Individual backup objects or complete

sessions can be restored using this method.
See also **ZDB to tape**, **ZDB to disk+tape**, and **replica**.

sqlhosts file *(Informix specific term)*
An Informix connectivity-information file that contains the names of each of the database servers and any aliases to which the clients on a host computer can connect.

SRD file
The Data Protector System Recovery Data (SRD) file contains system information required for installing and configuring the operating system in case of a disaster. The SRD file is an ASCII file, generated when a CONFIGURATION backup is performed on a Windows client and stored on the Cell Manager.

SRDF *(EMC Symmetrix specific term)*
The EMC Symmetrix Remote Data Facility is a business continuation process that enables effective, real-time data replication of SLDs between dislocated processing environments. These environments could be situated within the same root computer environment or separated by long distances.

SSE Agent *(HP StorageWorks Disk Array XP specific term)*
A Data Protector software module that executes all tasks required for a split mirror backup integration. It communicates with the HP StorageWorks Disk Array XP storing system using the RAID Manager XP utility (HP-UX and Windows systems) or RAID Manager Library (Solaris systems).

sst.conf file
The file /usr/kernel/drv/sst.conf is required on each Data Protector Sun Solaris client to which a multi-drive library device is connected. It must contain an entry for the SCSI address of the robotic mechanism of each library device connected to the client.

st.conf file
The file /kernel/drv/st.conf is required on each Data Protector Solaris client with a backup device connected. It must contain device information and a SCSI address for each backup drive connected to that client. A single SCSI entry is required for a single-drive device and multiple SCSI entries are required for a multi-drive library device.

stackers
Devices with multiple slots for media storage usually with only one drive. A stacker selects media from the stack sequentially. In contrast, a library can randomly select media from its repository.

standalone file device
A file device is a file in a specified directory to which you back up data.

standard security *(MS SQL specific term)*
Standard security uses the login validation process of the Microsoft SQL Server for all connections. Standard security is useful in network environments with a variety of clients, some of which may not support trusted connections. It also provides backward compatibility for older versions of the Microsoft SQL Server.
See also **integrated security**.

Storage Group
(Microsoft Exchange Server 2000/2003 specific term)
A collection of databases (stores) that share a common set of transaction log files. Exchange manages each storage group with a separate server process.

StorageTek ACS library
(StorageTek specific term)
Automated Cartridge System is a library system (also known as Silo) consisting of one Library Management Unit (LMU) and one to 24 Library Storage Modules (LSM) connected to the unit.

storage volume *(ZDB specific term)*
A storage volume represents an object that may be presented to an operating system or some other entity (for example, a virtualization mechanism) upon which volume management systems, file systems, or other objects may exist. The volume management systems, file systems are built on this storage. Typically, these can be created or exist within a storage system such as a disk array.

switchover
See **failover**

Sybase Backup Server API *(Sybase specific term)*
An industry-standard interface developed for the exchange of backup and recovery information between a Sybase SQL Server and a backup solution like Data Protector.

Sybase SQL Server *(Sybase specific term)*
The server in the Sybase "client-server" architecture. Sybase SQL Server manages multiple databases and multiple users, keeps track of the actual location of data on disks, maintains mapping of logical data description to physical data storage, and maintains data and procedure caches in memory.

Symmetrix Agent (SYMA) *(EMC Symmetrix specific term)*
The Data Protector software module that prepares the EMC Symmetrix environment for backup and restore operations.

System Backup to Tape *(Oracle specific term)*
An Oracle interface that handles the actions required to load, label, and unload correct backup devices when Oracle issues a backup or restore request.

system databases *(Sybase specific term)*
The four system databases on a newly installed Sybase SQL Server are the:

- master database (master)
- temporary database (tempdb)
- system procedure database (sybsystemprocs)
- model database (model).

system disk
A system disk is a disk containing operating system files. Microsoft terminology defines the system disk as a disk containing the files required for initial step of boot process.

system partition
A system partition is a partition containing operating system files. Microsoft terminology defines a system partition as a partition containing the files required for initial step of boot process.

System State *(Windows specific term)*
The System State data comprises the Registry, COM+ Class Registration database, system startup files, and the Certificate Services database (if the server is a certificate server). If the server is a domain controller, Active Directory directory services and the Sysvol directory are also contained in the System State data. If the server is running the Cluster service, the System State data also includes resource registry checkpoints and the quorum resource recovery log, which contains the most recent cluster database information.

system volume/disk/partition
A volume/disk/partition containing operating system files. Microsoft terminology defines the system volume/disk/partition as the volume/disk/partition containing files required for the initial step of the boot process.

SysVol (*Windows specific term)*
A shared directory that stores the server copy of the domain's public files, which are replicated among all domain controllers in the domain.

tablespace
A part of a database structure. Each database is logically divided into one or more tablespaces. Each tablespace has data files or raw volumes exclusively associated with it.

tapeless backup *(ZDB specific term)*
See **ZDB to disk**.

target database *(Oracle specific term)*
In RMAN, the target database is the database that you are backing up or restoring.

target (R2) device *(EMC Symmetrix specific term)*
An EMC Symmetrix device that participates in SRDF operations with a source (R1) device. It resides in the remote EMC Symmetrix unit. It is paired with a source (R1) device in the local EMC Symmetrix unit and receives all write data from its mirrored pair. This device is not accessed by user applications during normal I/O operations. An R2 device must be assigned to an RDF2 group type.
See also **source (R1) device**

target system *(Disaster Recovery specific term)*
A system after a computer disaster has occurred. The target system is typically in a non-bootable state and the goal of disaster recovery is to restore this system to the original system configuration. The difference between a crashed system and a target system is that a target system has all faulty hardware replaced.

target volume *(ZDB specific term)*
A storage volume to which data is replicated.

Terminal Services *(Windows specific term)*
Windows Terminal Services provide a multi-session environment that allows clients to access a virtual Windows desktop session and Windows-based programs running on the server.

thread *(MS SQL Server 7.0/2000 specific term)*
An executable entity that belongs to only one process. It comprises a program counter, a user-mode stack, a kernel-mode stack, and a set of register values. Several threads can run at a time within one process.

TimeFinder *(EMC Symmetrix specific term)*
A business continuation process that creates an instant copy of single or multiple Symmetrix Logical Devices (SLDs). The instant copy is created on specially preconfigured SLDs called BCVs and is accessible via a separate device address to the system(s).

TLU
Tape Library Unit.

TNSNAMES.ORA *(Oracle and SAP R/3 specific term)*
A network configuration file that

contains connect descriptors mapped to service names. The file may be maintained centrally or locally, for use by all or individual clients.

transaction
A mechanism for ensuring that a set of actions is treated as a single unit of work. Databases use transactions to keep track of database changes.

transaction backup
Transaction backups generally use fewer resources than database backups, so they can be created more frequently than database backups. By applying transaction backups, you can recover the database to a specific point in time prior to when a problem occurred.

transaction backup *(Sybase and SQL specific term*)
A backup of the transaction log providing a record of changes made since the last full or transaction backup.

transaction log backup
Transaction log backups generally use fewer resources than database backups so they can be created more frequently than database backups. By applying transaction log backups, you can recover the database to a specific point in time.

transaction log files
Files that record transactions of the database modifications, and provide fault tolerance in case of a database disaster.

transaction logs (*Data Protector specific term)*
Keeps track of IDB changes. The archiving of transaction logs should be enabled to prevent you from losing the transaction log files that are created after the last IDB backup and are necessary for IDB recovery.

transaction log table *(Sybase specific term)*
A system table in which all changes to the database are automatically recorded.

transportable snapshot *(MS VSS specific term)*
A shadow copy that is created on the application system and can be presented to the backup system which performs the backup.
See also **Microsoft Volume Shadow Copy service (VSS)**.

TSANDS.CFG file *(Novell NetWare specific term)*
A file that allows you to specify the names of containers where you want backups to begin. It is text file located in the SYS:SYSTEM\TSA directory on the server where TSANDS.NLM is loaded.

unattended operation
See **lights-out operation**.

user account
You can use Data Protector only if you have a Data Protector user account, which restricts unauthorized access to Data Protector and to backed up data. Data Protector administrators create this account specifying a user logon name, the systems from which the user can log on, and a Data Protector user group membership. This is checked whenever the user starts the Data Protector user interface or performs specific tasks.

user disk quotas
NTFS quota-management support enables an enhanced tracking mechanism and control over disk space usage on shared storage volumes. Data Protector backs up user disk quotas on the whole system and for all configured users at a time.

user group
Each Data Protector user is member of a User Group. Each User Group has a set of user rights that are given to every user in that User Group. The number of User Groups with their associated user rights can be defined as desired. Data Protector provides three default user groups: admin, operator, and user.

user profile *(Windows specific term)*
Configuration information retained on a user basis. This information includes desktop settings, screen colors, network connections, and so on. When the user logs on, the user profile is loaded and the Windows environment is set accordingly.

user rights
User rights or access rights are the permissions needed to perform specific Data Protector tasks. Configuring a backup, starting a backup session, or starting a restore session are typical user rights. Users have the access rights of the user group to which they belong.

vaulting media
The process of storing media to a safe and remote place. The media are brought back to the data center when they are needed for restore or are ready for reuse in subsequent backups. The vaulting procedure depends on your company's backup strategy and policies for data protection/reliability.

VBFS *(OmniStorage specific term)*
A Very Big File System is an extension of the standard HP-UX file system on HP-UX 9.x. It is mounted to a directory the same way as any HP-UX file system. In a VBFS, only the superblock, the inode and the 'extended attribute'

information remain permanently on the hard disk and are never migrated.
See also **MFS**.

verify
A function that lets you check whether the Data Protector data on a specified medium is readable. Additionally, consistency within each block can be checked if the backup was performed with the cyclic redundancy check (CRC) option ON.

Virtual Controller Software (VCS)

(HP StorageWorks EVA specific term)
The firmware that manages all aspects of storage system operation, including communication with Command View EVA through the HSV controllers.
See also **Command View (CV) EVA**.

Virtual Device Interface *(MS SQL Server 7.0/2000 specific term)*
This is a SQL Server 7.0/2000 programming interface that allows fast backup and restore of large databases.

virtual disk *(HP StorageWorks EVA specific term)*
A unit of storage allocated from an HP StorageWorks Enterprise Virtual Array storage pool. Virtual disks are the entities that are replicated using the HP StorageWorks Enterprise Virtual Array snapshot functionality.
See also **source volume** and **target volume**.

virtual server
A virtual machine in a cluster environment defined in a domain by a network IP name and address. Its address is cached by the cluster software and mapped to the cluster node that is currently running the virtual server resources. This way all requests for a particular virtual server are cached by a specific cluster node.

volser *(ADIC and STK specific term)*
A VOLume SERial number is a label on the medium to identify the physical tape used in very large libraries. A volser is a naming convention specific to ADIC/GRAU and StorageTek devices.

volume group
A unit of data storage in an LVM system. A volume group can consist of one or more physical volumes. There can be more than one volume group on the system.

volume mountpoint (*Windows specific term*)
An empty directory on a volume that can be used to mount another volume. The volume mount point acts as a gateway to the target volume. Provided that the volume is mounted, users and applications can refer to the data on the

mounted volume by the full (merged) filesystem path as if both volumes are one integral part.

Volume Shadow Copy service
See **Microsoft Volume Shadow Copy service**.

VPO
See **OVO**.

VSS
See **Microsoft Volume Shadow Copy service**.

VxFS
Veritas Journal Filesystem.

VxVM (Veritas Volume Manager)
A Veritas Volume Manager is a system for managing disk space on Solaris platforms. A VxVM system consists of an arbitrary group of one or more physical volumes organized into logical disk groups.

Wake ONLAN
Remote power-up support for systems running in power-save mode from some other system on the same LAN.

Web reporting
The Data Protector functionality that allows you to view reports on backup status and Data Protector configuration using the Web interface.

wildcard character
A keyboard character that can be used to represent one or many characters. The asterisk (*), for example, typically represents one or more characters, and the question mark (?) typically represents a single character. Wildcard characters are often used in operating systems as a means of specifying more than one file by name.

Windows CONFIGURATION backup
Data Protector allows you to back up Windows CONFIGURATION, including Windows Registry, user profiles, Event Logs, and WINS and DHCP server data (if configured on a system) in one step.

Windows Registry
A centralized database used by Windows to store configuration information for the operating system and the installed applications.

WINS server A system running Windows Internet Name Service software that resolves Windows networking computer names to IP addresses. Data Protector can back up WINS server data as part of the Windows configuration.

writer
(MS VSS specific term)
A process that initiates change of data on the original volume. Writers are typically applications or system services that write persistent information on a volume. Writers also participate in the shadow copy synchronization process by assuring data consistency.

XBSA interface *(Informix specific term)*
The onbar utility and Data Protector communicate with each other through the X/Open Backup Specification Services Programmer's Interface (XBSA).

XCopy engine *(direct backup specific term)*
A SCSI-3 copy command that allows you to copy data from a storage device having a SCSI source address to a backup device having a SCSI destination address, thus enabling direct backup. The data flows from a source device (either block or streaming, that is, disk or tape) to the destination device (either block or streaming) through XCopy. This releases the controlling server of reading the data from the storage device into memory and then writing the information to the destination device.
See also **direct backup**.

ZDB
See **zero downtime backup (ZDB)**.

ZDB database *(ZDB specific term)*
A part of the IDB, storing ZDB related information such as source volumes, replicas and security information. The ZDB database is used for ZDB, instant recovery, and split mirror restore.
See also **zero downtime backup (ZDB)**.

ZDB to disk *(ZDB specific term)*
A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time. Multiple replicas, produced using the same backup specification at different times, can be kept in a replica set. A replica from ZDB to disk can be restored using the instant recovery process.
See also **zero downtime backup (ZDB)**, **ZDB to tape**, **ZDB to disk+tape**, **instant recovery**, and **replica set rotation**.

ZDB to disk+tape *(ZDB specific term)*
A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time, in the same way as with ZDB to disk. However, data in the replica is also streamed to a backup medium, as with ZDB to tape. If this backup method is used, data backed up in the same session can be restored

using the instant recovery process, the standard Data Protector restore from tape, or on split mirror arrays, split mirror restore.
See also **zero downtime backup (ZDB)**, **ZDB to disk**, **ZDB to tape**, **instant recovery**, **replica**, and **replica set rotation**.

ZDB to tape *(ZDB specific term)*
A form of zero downtime backup where data in the replica produced is streamed to a backup medium, typically tape. Instant recovery is not possible from such a backup, so the replica need not be retained on the disk array after backup completion. The backed-up data can be restored using standard Data Protector restore from tape. On split mirror arrays, split mirror restore can also be used.
See also **zero downtime backup (ZDB)**, **ZDB to disk**, **instant recovery**, **ZDB to disk+tape**, and **replica**.

zero downtime backup (ZDB)
A backup approach in which data replication techniques, provided by a disk array, are used to minimize the impact of backup operations on an application system. A replica of the data to be backed up is created first. All subsequent backup operations are performed on the replicated data rather than the original data, while the application system can be returned to normal operation.
See also **ZDB to disk**, **ZDB to tape**, **ZDB to disk+tape**, and **instant recovery**.

A

alternate paths, A-8

B

backup
  configuring for EMC, 225–228
  configuring for EVA, 95–102
  configuring for VA, 29–35
  configuring for XP, 164–168
  EMC flow, 221
  EMC troubleshooting, 235
  EVA flows, 92–94
  EVA troubleshooting, 109
  running, A-3–A-7
  scheduling, A-3–A-7
  VA flows, 25–28
  VA troubleshooting, 41
  XP flow, 158
  XP troubleshooting, 178
backup devices
  unlocking, A-48
backup options
  EMC, 229–230
  EMC exemplary selections, A-41
  EVA, 103–108
  EVA exemplary selections, A-37
  VA, 36–40
  VA exemplary selections, A-37
  XP, 169–177
  XP exemplary selections, A-39
backup process
  EMC, 221
  EVA, 89
  VA, 23
  XP, 157
backup system
  automatic configuration on EMC, 216
  automatic configuration on EVA, 68
  automatic configuration on VA, 8
  automatic configuration on XP, 144
  configuring for EVA, 68
backup types
  EMC, 221
  EVA, 89
  VA, 23
  XP, 157
BC1 configuration, 177
Business Copy XP configuration (BC)
  preparing environment, 142

C

cluster configurations, A-10–A-18
cluster specifics
  EMC split mirror restore, 255
  instant recovery, A-20
  XP split mirror restore, 195
Combined (SRDF + TimeFinder)
    configuration
  preparing environment, 214
Combined Continuous Access XP and
    Business Copy XP configuration
  preparing environment, 143
command device
  concepts, XP, 146
  handling, XP, 148
Command View EVA *see* CV EVA
commands
  omnidbeva, 71
  omnidbsmis, 71
  omnidbutil, A-48
  omnidbva, 13
  omnidbxp, 147
concepts
  command device, XP, 146
  EMC ZDB to tape, 221
  EVA ZDB to disk, 89
  EVA ZDB to disk+tape, 90
  EVA ZDB to tape, 89
  VA ZDB to disk, 23
  VA ZDB to disk+tape, 24
  VA ZDB to tape, 23
  XP ZDB to disk, 157
  XP ZDB to disk+tape, 158
  XP ZDB to tape, 157
configuring
  backup system for EVA, 68
  EMC, 214
  EMC backup, 225–228
  EVA, 67
  EVA backup, 95–102
  EVA disk group pairs, 75–78
  EVA login information, 72
  scanning EVA hardware configuration, 74
  VA, 7
  VA backup, 29–35
  VA LUN exclude file, 14–16
  XP, 142
  XP backup, 164–168
  XP LDEV exclude file, 151–153

Continuous Access XP configuration (CA)
  preparing environment, 142
conventions, xi
creating
  Data Protector EMC database file, 216
  EMC Symmetrix database file, 215
CV EVA, 67

D

Data Protector EMC database file
  creating, 216
  See also EMC Symmetrix database file
disk configuration data
  obtaining, EMC, A-42
disk group pairs configuration file, EVA, 77
disk group pairs on EVA
  defining, 75—78

E

EMC
  backup flow, 221
  backup options, 229—230
  backup process, 221
  backup types, 221
  configuring, 214
  configuring backup, 225—228
  licensing, 213
  limitations, 213
  obtaining disk configuration data, A-42
  preparing environment, 214
  prerequisites, 213
  restore on LAN, 246
  restore overview, 245
  split mirror restore, 249
  split mirror restore, procedure, 251
  troubleshooting backup, 235
  troubleshooting restore, 257
EMC log file, 217
EMC Power Path, A-8
EMC Symmetrix database file
  creating, 215
  See also Data Protector EMC database file
EVA
  backup flows, 92—94
  backup options, 103—108
  backup process, 89
  backup types, 89
  configuring, 67
  configuring backup, 95—102
  configuring backup system, 68
  CV EVA login information, 67
  EVADB, 70—85
  instant recovery flow, 126
  instant recovery options, 131
  instant recovery overview, 125
  instant recovery procedure, 127
  licensing, 63
  limitations, 64
  prerequisites, 63
  restore on LAN, 124
  restore, overview, 123
  SMISDB, 70—85
  troubleshooting backup, 109
  troubleshooting instant recovery, 132
  troubleshooting restore, 132
  ZDB to disk flow, 92
  ZDB to disk+tape flow, 93
  ZDB to tape flow, 93
EVADB, 70—85
  deleting target volumes, 84
  disk group pairs configuration file, 75—78
  hardware configuration, 74
  querying, 79—83
  setting CV EVA login information, 72
  synchronizing, 85
examples
  EMC backup options, A-41
  EVA backup options, A-37
  VA backup options, A-37
  XP backup options, A-39
exclude file
  see VA LUN exclude file
  see XP LDEV exclude file

H

HP OpenView Storage Area Manager
    password, VA, 17
HP StorageWorks AutoPath, A-8
  configuring VA, 9
  limitations and considerations, A-9
HP StorageWorks Secure Path, A-8
HP-UX Logical Volume Manager Alternate
    Links, A-8

I

instant recovery
  cluster specifics, A-20
  EVA flow, 126

EVA options, 131
EVA overview, 125
EVA procedure, 127
EVA troubleshooting, 132
LVM mirroring, XP, 204
VA flow, 50
VA options, 55
VA overview, 49
VA procedure, 51
VA troubleshooting, 57
XP flow, 198
XP options, 203
XP overview, 197
XP procedure, 199
XP troubleshooting, 206

L

licensing
EMC, 213
EVA, 63
VA, 5
XP, 139
limitations
EMC, 213
EVA, 64
VA, 5
XP, 140
LUN security, VA, 16
LVM mirroring
instant recovery on XP, 204
preparing environment, XP, 143
PVG-strict mirroring, XP, 143
LVM volume group
checking, A-47

O

omnidbeva, 71
omnidbsmis, 71
omnidbutil, A-48
omnidbva, 13
omnidbxp, 147
omnirc variables, A-23—A-36
common ZDB, A-23—A-29
EMC specific, A-34—A-36
EVA specific, A-29—A-32
VA specific, A-29
XP specific, A-33—A-34
options
backup examples, EMC, A-41
backup examples, EVA, A-37
backup examples, VA, A-37
backup examples, XP, A-39
backup, EMC, 229—230
backup, EVA, 103—108
backup, VA, 36—40
backup, XP, 169—177
instant recovery, EVA, 131
instant recovery, VA, 55
instant recovery, XP, 203
split mirror restore, EMC, 253—255
split mirror restore, XP, 193—195

P

password
for LUN security, VA, 16
Storage Area Manager, VA, 17
preparing environment
Business Copy XP configuration (BC), 142
Combined (SRDF + TimeFinder) configuration, 214
Combined Continuous Access XP and Business Copy XP configuration, 143
Continuous Access XP configuration (CA), 142
LVM mirroring, XP, 143
Symmetrix integration, 214
Symmetrix Remote Data Facility configuration (SRDF), 214
TimeFinder configuration (TF), 214
XP, 142
prerequisites
EMC, 213
EVA, 63
VA, 5
XP, 139
PVG-strict mirroring, XP, 143

Q

querying
EVADB, 79
SMISDB, 79
VADB, 13
XPDB, 147

R

restore
EMC on LAN, 246

EMC split mirror restore, 249
EMC troubleshooting, 257
EVA instant recovery, 125
EVA on LAN, 124
EVA troubleshooting, 132
overview, EMC, 245
overview, EVA, 123
overview, VA, 47
overview, XP, 185
VA instant recovery, 49
VA on LAN, 48
VA troubleshooting, 57
XP instant recovery, 197
XP split mirror restore, 189
XP troubleshooting, 205
XP, on LAN, 186
running
backup, A-3—A-7

S

scheduling
backup, A-3—A-7
security
LUN security, VA, 16
VA LUN exclude file, 14—16
XP LDEV exclude file, 151—153
SMISDB, 70—85
deleting target volumes, 84
disk group pairs configuration file, 75—78
purging, 83
querying, 79—83
setting CV EVA login information, 72
synchronizing, 85
split mirror restore
EMC, 249
EMC cluster specifics, 255
EMC flow, 249
EMC options, 253—255
EMC procedure, 251
EMC troubleshooting, 259
XP, 189
XP cluster specifics, 195
XP flow, 189
XP options, 193—195
XP procedure, 190
XP troubleshooting, 205
Storage Area Manager password, VA, 17
Sun Alternate Pathing Driver, A-8
Symmetrix Remote Data Facility configuration (SRDF)
preparing environment, 214

T

target volumes
unlocking, A-48
TimeFinder configuration (TF)
preparing environment, 214
troubleshooting
EMC backup, 235
EMC restore, 257
EMC split mirror restore, 259
EVA backup, 109
EVA instant recovery, 132
EVA restore, 132
VA backup, 41
VA instant recovery, 57
VA restore, 57
XP backup, 178
XP instant recovery, 206
XP restore, 205
XP split mirror restore, 205
typographical conventions, xi

U

unlocking
target volumes and backup devices, A-48

V

VA
backup flows, 25—28
backup options, 36—40
backup process, 23
backup types, 23
configuring, 7
configuring backup, 29—35
configuring with AutoPath, 9
instant recovery flow, 50
instant recovery options, 55
instant recovery overview, 49
instant recovery procedure, 51
licensing, 5
limitations, 5
prerequisites, 5
restore on LAN, 48
restore overview, 47
troubleshooting backup, 41

troubleshooting instant recovery, 57
troubleshooting restore, 57
VADB, 12–19
ZDB to disk flow, 26
ZDB to disk+tape flow, 27
ZDB to tape flow, 27
VA LUN exclude file, 14–16
VADB, 12–19
checking consistency, 14
deleting the contents, 18
providing Storage Area Manager password, 17
querying, 13
setting LUN security, 16
VA LUN exclude file, 14–16
Veritas Volume Manager Dynamic Multipathing, A-8

X

XP
backup flow, 158
backup options, 169–177
backup process, 157
backup types, 157
configuring, 142
configuring backup, 164–168
instant recovery flow, 198
instant recovery options, 203
instant recovery procedure, 199
instant recovery, overview, 197
licensing, 139
limitations, 140
preparing environment, 142
prerequisites, 139
restore on LAN, 186
restore overview, 185
split mirror restore, 189
split mirror restore, procedure, 190
troubleshooting backup, 178
troubleshooting instant recovery, 206
troubleshooting restore, 205
XPDB, 145–153
XP LDEV exclude file, 151–153
XPDB, 145–153
command device handling, 148
command devices, 146
querying, 147
XP LDEV exclude file, 151–153

Z

ZDB database
EVADB *see* EVADB
SMISDB *see* SMISDB
VADB *see* VADB
XPDB *see* XPDB
ZDB to disk
EVA concept, 89
EVA flow, 92
VA concept, 23
VA flow, 26
XP concept, 157
ZDB to disk+tape
EVA concept, 90
EVA flow, 93
VA concept, 24
VA flow, 27
XP concept, 158
ZDB to tape
EMC concept, 221
EVA concept, 89
EVA flow, 93
VA concept, 23
VA flow, 27
XP concept, 157

Index