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

Manual Edition: February 2006 (build label 249)

Legal Notices

Contents

I. HP StorageWorks Virtual Array

Contents

II. HP StorageWorks Enterprise Virtual Array

Contents

III. HP StorageWorks Disk Array XP

Contents

IV. EMC Symmetrix

Contents

Contents

Printing History

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

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

Table 1　　　**Edition History**

Part Number	Manual Edition	Product
B6960-90113	October 2004	Data Protector Release A.05.50
B6960-90012	April 2006	Data Protector Release A.06.00

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 online Help for information about the Data Protector graphical user interface.

Figure 1   **Data Protector Graphical User Interface**
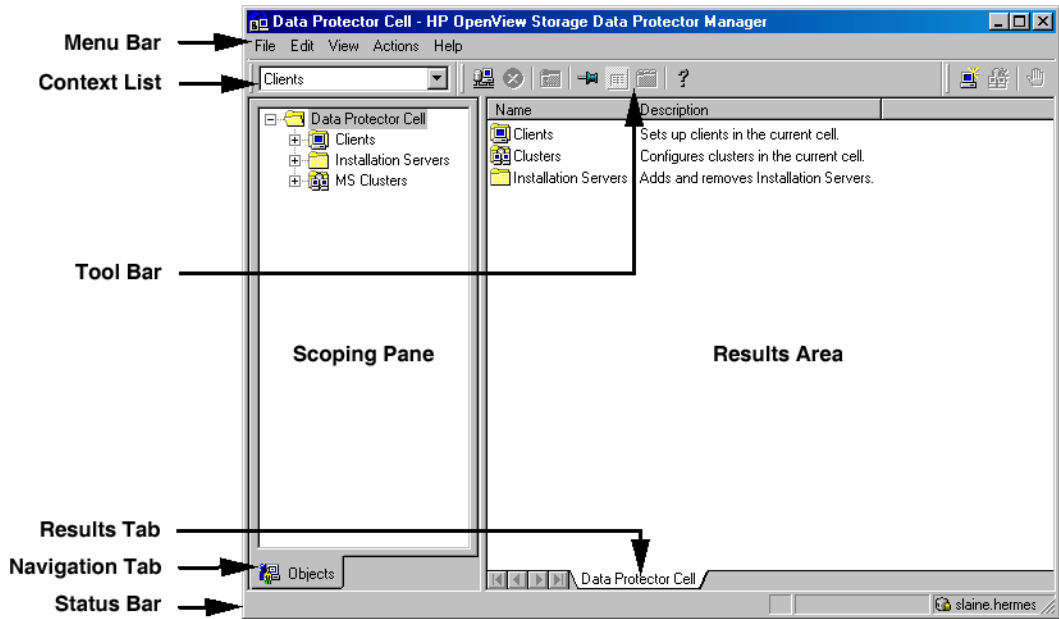
Contact Information

General Information

General information about Data Protector can be found at

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

Technical Support

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

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

http://www.hp.com/support

Information about the latest Data Protector patches can be found at

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

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

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

Documentation Feedback

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

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

Training Information

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

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

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

Data Protector Documentation

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

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

HP OpenView Storage Data Protector Concepts Guide

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

HP OpenView Storage Data Protector Installation and Licensing Guide

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

HP OpenView Storage Data Protector Troubleshooting Guide

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

HP OpenView Storage Data Protector Disaster Recovery Guide

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

HP OpenView Storage Data Protector Integration Guide

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

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

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

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

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

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

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

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

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

HP OpenView Storage Data Protector Integration Guide for HP OpenView

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

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

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

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

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

There are two versions of the manual:

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

HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide

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

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

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

HP OpenView Storage Data Protector Zero Downtime Backup Integration Guide

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

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

HP OpenView Storage Data Protector MPE/iX System User Guide

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

HP OpenView Storage Data Protector Media Operations User's Guide

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

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

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

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

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

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

Documentation Map

Abbreviations

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

Abbreviation	Manual
CLI	Command Line Interface Reference Guide
Concepts	Concepts Guide
DR	Disaster Recovery Guide
GS	Getting Started Guide
Help	Online Help
IG-IBM	Integration Guide—IBM Applications
IG-MS	Integration Guide—Microsoft Applications
IG-O/S	Integration Guide—Oracle & SAP
IG-OV	Integration Guide—HP OpenView Service Information Portal/OpenView Reporter
IG-OVOU	Integration Guide—HP OpenView Operations, UNIX
IG-OVOW	Integration Guide—HP OpenView Operations 7.1x, 7.2x, Windows
IG-OVOW	Integration Guide—HP OpenView Operations 7.5, Windows
IG-Var	Integration Guide—Sybase, Network Node Manager & NDMP
Install	Installation and Licensing Guide
MO GS	Media Operations Getting Started Guide
MO RN	Media Operations Product Announcements, Software Notes, and References
MO UG	Media Operations User Guide
MPE/iX	MPE/iX System User Guide
PA	Product Announcements, Software Notes, and References

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

Map

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

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

Integrations

Look in these manuals for details of the following integrations:

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

In This Book

The *HP OpenView Storage Data Protector Zero Downtime Backup 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:

I   HP StorageWorks Virtual Array

1   Configuration and Maintenance

Overview

This chapter describes configuration and maintenance of the Data Protector HP StorageWorks Virtual Array (VA) integration. It also provides information on the ZDB database.

Prerequisites

- Install:

  VA components:

  — Array microcode hp15 (minimum)

  — HP StorageWorks Command View SDM

  — Business Copy (BC) VA microcode and license.

  — A license for managing and controlling the VA storage system.

  Data Protector components:

  — A license for using the VA integration.

  — HP StorageWorks VA Agent.

  For installation instructions, see the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

- Make sure the same operating system (and its version) is installed on the application and backup systems.

- Connect VA to the application and backup systems through the SAN.

- 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.

- For ZDB to disk, configure a backup device (for example, a standalone file device), as you cannot configure a backup specification without selecting a device. For configuration instructions, see the online Help index: "standalone devices".

See the *HP OpenView Storage Data Protector Software Release Notes* for information on:

- General Data Protector and integration-specific limitations

- Supported platforms and integrations

- Supported backup and connectivity topologies

- Supported cluster configurations

For information on supported configurations, see the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

ZDB Database - VADB

ZDB database is referred to as **VADB**. It keeps information about:

- ZDB-to-disk and ZDB-to-disk+tape sessions. This information includes:

  — Session ID

  — Time when the session was performed

  — Name of the backup specification used in the session

  — LUNs and WWNs of disk arrays used in the session

- CRC check information calculated during the session.

- A list of snapshots not to be used by Data Protector (kept in the VA LUN exclude file).

- VA password, if LUN security is used.

- Storage Area Manager password, if HP OpenView Storage Allocater is installed.

- Filesystem and volume management system information.

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

VADB resides on the Cell Manager in:

- *UNIX:* /var/opt/omni/server/db40/vadb

- *Windows:* *<Data_Protector_home>*\db40\vadb

Configuring the Integration

To configure the integration:

- Pre-configure snapshots (if HP StorageWorks AutoPath is installed and configured on the backup system). See "Configuring VA with HP StorageWorks AutoPath Installed" on page 7.

- Activate LUN security (if HP StorageWorks Secure Manager is enabled). See "LUN Security" on page 9.

- Provide the HP OpenView Storage Area Manager (SAM) password (if VA runs in a SAN environment with HP OpenView Storage Allocater installed). See "HP OpenView Storage Area Manager Password" on page 9.

- If needed, set VA LUN exclude file. See "VA LUN Exclude File" on page 10.

Before you start configuration, make sure you met the prerequisites described in "Overview" on page 4.

Configuring VA with HP StorageWorks AutoPath Installed

The backup system with AutoPath can only detect newly created snapshots if it is restarted after the snapshot creation. Since Data Protector can create new snapshots (representing replicas), such replicas must be pre-configured and the backup system restarted before running backups.

During ZDB sessions, the VA integration reuses these pre-configured snapshots from the replicas.

IMPORTANT   Before pre-configuring snapshots, carefully consider which backup objects representing different source volumes will be used in backup specifications. Select backup objects so that source volumes representing these objects are not included in more than one backup specification.

To pre-configure snapshots that will represent a replica:

1. Create a backup specification (see "Configuring a Backup Specification" on page 19). Select the following options:

   - *ZDB to disk, ZDB to disk+tape:*

     — Leave `Track the replica for instant recovery` selected.

     — Specify `Number of replicas rotated`.

   - *ZDB to tape:*

     — Deselect `Track the replica for instant recovery`.

     — Select `Keep the replica after backup`.

     — Do not select `Use an existing replica`.

2. Run the backup (see "Running and Scheduling ZDB Sessions" on page A-2).

   - *ZDB to disk, ZDB to disk+tape:* Run as many backup sessions as specified in `Number of replicas rotated`. After that, you can reuse snapshots (new snapshots will not be created).

   - *ZDB to tape:* Run one backup session for one backup specification. For parallel backups, ensure that for each backup specification you have a replica created (one backup session performed).

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

---

3. After successfully performed backups, restart the backup system, so, AutoPath works properly and is aware of all created snapshots.

4. *ZDB to disk, ZDB to disk+tape:* Use the configured backup specification.

   ZDB-to-tape: Modify the backup specification by selecting `Use an existing replica` (leave `Keep the replica after backup` selected). As a result, snapshots are reused and AutoPath functions properly.

Repeat this procedure every time you create a new backup specification with new backup objects.

LUN Security

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

IMPORTANT   If Secure Manager is enabled, specify the VA password using the `omnidbva` command.

To activate LUN security:

1. Start Secure Manager using Command View SDM.

2. Provide the VA password:

   `omnidbva -vapasswd <VA_wwn> <password>`

   <VA_wwn>   VA node WWN.

   <password>   Password used to start Secure Manager.

NOTE   To find the VA node WWN, use Command View SDM. On HP-UX, you can also use `tdlist` and `fcmsutil` commands. For information on the commands, see their man pages.

3. Select the option `Integrate with VA LUN security` every time you create a backup specification.

HP OpenView Storage Area Manager Password

If your VA runs in a SAN environment with HP OpenView Storage Allocater installed, provide the HP OpenView Storage Area Manager (SAM) password.

IMPORTANT   If you do not provide the password, the snapshot creation fails.

Provide the password as follows:

1. Run:

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

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

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 WWN of your VA.

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

VA LUN Exclude File

You can reserve certain target volumes, identified by their source volume
(parent LUN), for purposes other than Data Protector backup.

To set/edit the VA LUN exclude file, use the omnidbva command. See the
omnidbva man page for command syntax and examples of creating and
editing the VA LUN exclude file. The template of the file is as follows.

```
#
# 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
```

Automatic Configuration of Backup System

Necessary configuration steps, such as configuring volume groups and filesystems on the backup system, are performed 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 and mounts these filesystems during ZDB-to-tape or ZDB-to-disk+tape sessions. For more information on the backup system mountpoint creation, see the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

Maintaining the Integration

Maintenance tasks are divided into the following categories:

- Maintenance query tasks. See "Querying VADB" on page 12
- Maintenance fix tasks. See "Checking VADB Consistency" on page 12.
- Maintenance deletion tasks. See "Deleting VADB Entries" on page 12.

Querying VADB

Using the omnidbva command, you can list:

- All backup sessions stored in VADB
- LUNs used in ZDB-to-disk/ZDB-to-disk+tape sessions

See the omnidbva man page for command syntax and examples.

Checking VADB Consistency

To perform a VADB consistency check and fix invalid entries, use the omnidbva command. See the omnidbva man page for command syntax and examples.

Deleting VADB Entries

Using the omnidbva command, you can:

- Delete information on a specific session (replica), identified by the session ID
- Reset all entries in VADB (except the contents of the VA LUN exclude file and the password information)

See the omnidbva man page for command syntax and examples.

IMPORTANT   The `omnidbva` command removes all entries from VADB, including
session information, LUN security, and Storage Area Manager
passwords. Although it does not remove the target volumes that
constitute a replica, you cannot perform instant recovery from such a
replica because all session information was deleted.

2 Backup

Overview

This chapter describes configuring a filesystem and disk image ZDB using the Data Protector GUI.

You should be familiar with VA concepts and procedures and basic Data Protector ZDB and instant recovery functionality. See the VA-related documentation and the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

Snapshot Types

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

ZDB Types

You can classify ZDB sessions in two ways:

- According to the media to which data is backed up:

  — **ZDB to disk**

    The replica produced is kept on a disk array until reused. This replica becomes part of the replica set and can be used for instant recovery.

    ZDB to disk is performed if the option Track the replica for instant recovery is selected in a backup specification, and To disk is selected when running/scheduling a backup.

  — **ZDB to tape**

    The replica produced is streamed to backup media, typically tape.

    This replica is deleted after backup if the option Keep the replica after the backup *is not* selected in a backup specification. If this option *is* selected, the replica remains on a disk array until reused and becomes part of the replica set. However, it cannot be used for instant recovery.

— **ZDB to disk+tape**

The replica produced is kept on a disk array until reused and is also streamed to backup media. This replica becomes part of the replica set and can be used for instant recovery.

ZDB to disk+tape is performed if the option `Track the replica for instant recovery` is selected in a backup specification, and `To disk+tape` is selected when running/scheduling a backup.

- According to the amount of data that is backed up to tape:

Table 2-1   **Backup Types**

Full	Backs up all selected objects, even if there are no changes from the previous backup.
Incremental	Backs up changes from the last protected backup.

NOTE   The replica is created in the same way for both full and incremental ZDB. The difference between the two ZDB types is only in the amount of files that are streamed from the replica to tape.

Limitations   — Incremental ZDB is not supported for ZDB to disk.

— Incremental ZDB is supported only for filesystem backups.

Replica Creation and Reuse

A new replica is created when:

- ZDB to tape is performed, in which `Use an existing replica backup` option *is not* selected. This replica is not part of the replica set.

- ZDB to disk/ZDB to disk+tape session is performed (`Track the replica for instant recovery` selected), and the specified `Number of replicas rotated` is not reached. This replica becomes part of the replica set.

The oldest replica in the set is deleted and the new one is created when:

- ZDB to tape is performed in which `Use an existing replica` *is* selected. This is only possible if a disk array already contains a replica that is not part of the replica set (either pre-configured, or left on a disk array from a previous ZDB-to-tape session using the same backup specification, with the option `Keep the replica after the backup` selected).

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

If the option `Keep the replica after the backup` is not selected, the replica and therefore all snapshots created during the backup session are deleted.

Configuring a Backup Specification

IMPORTANT   Before you begin, consider all limitations regarding the VA integration.
For more information, see the *HP OpenView Storage Data Protector
Software Release Notes* and the *HP OpenView Storage Data Protector
Zero Downtime Backup Concepts Guide*.

Prerequisites   • If HP StorageWorks AutoPath is installed on the backup system,
pre-configure snapshots as described in "Configuring VA with HP
StorageWorks AutoPath Installed" on page 7.

To configure a backup specification:

1. In the Context List, select `Backup`.

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

   The `Create New Backup` dialog box appears.

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

   Select `Snapshot backup` as `Backup type` and `HP StorageWorks VA`
as `Sub type`. For descriptions of options, press **F1**.

   Click `OK`.

3. Under `Client systems`, select `Application system` and `Backup
system`.

   If Secure Manager is activated, select `Integrate with VA LUN
security` under `Replica management options`. The password must
be configured correctly; otherwise, the session fails. See "LUN
Security" on page 9 for instructions.

   Specify options as follows:

ZDB to disk, ZDB to disk+tape: Under `Instant recovery options`, leave `Track the replica for instant recovery` selected, and specify `Number of replicas rotated`. The maximum number is 1024. See Figure 2-1 on page 20 and "Backup Options" on page 25.

Figure 2-1   **VA Backup Options (ZDB to Disk, ZDB to Disk+Tape)**
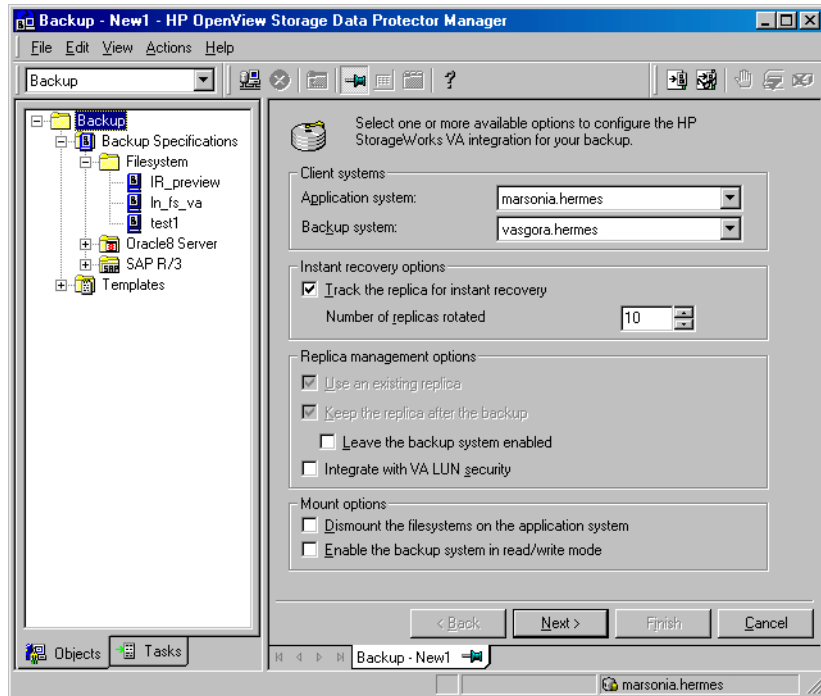
NOTE   You specify a ZDB-to-disk or ZDB-to-disk+tape session using the `Split mirror/snapshot backup` option when running/scheduling a backup. See "Running and Scheduling ZDB Sessions" on page A-2.

Click `Next`.

ZDB to tape: Deselect `Track the replica for instant recovery`. To keep the replica after backup, select `Keep the replica after the backup` as shown in Figure 2-2 on page 21. For information about options, see "Backup Options" on page 25.

Figure 2-2   **VA Backup Options (ZDB to Tape)**
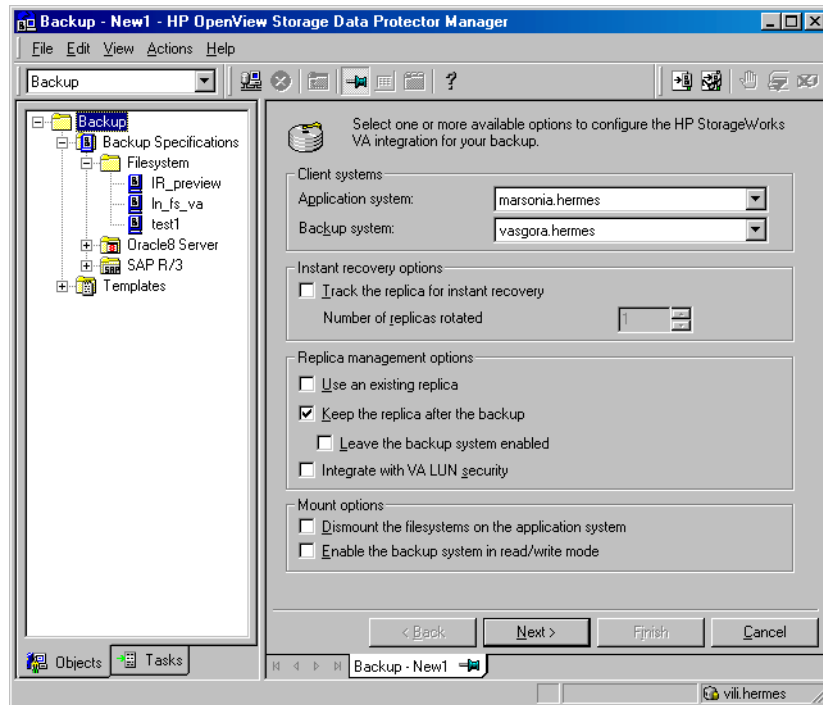
IMPORTANT   When selecting `Use an existing replica`, ensure a replica for the same source volumes exists on a disk array, otherwise the backup will fail.

Click `Next`.

4. *Filesystem backup:* Expand the application systems and select the objects to be backed up.

IMPORTANT   To perform instant recovery on UNIX, select all filesystems inside the volume group. Otherwise, instant recovery may fail or require workarounds and there is a risk of data corruption.

Click Next.

Disk image backup: Click Next.

5. Select devices. Click Properties to set device concurrency, media pool, and preallocation policy. For descriptions of these options, click Help.

To create additional copies (mirrors) of backup, specify the desired number of mirrors by clicking Add mirror/Remove mirror. Select separate devices for each mirror backup.

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

NOTE   Object mirroring and object copying are 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 VA backup options.

You can specify Application options and modify all other options, except Application system and Backup system. See "Backup Options" on page 25.

In the Filesystem Options group box, click Advanced and specify filesystem options as desired. For information, press **F1**.

Windows only: If you plan to do incremental ZDB, select the Do not use archive attribute filesystem option in the WinFSOptions page.

7. Following the wizard, open the scheduler (for information, press **F1** or see "Running and Scheduling ZDB Sessions" on page A-2), and then the backup summary page.

8. *Filesystem backup:* Click Next.

Disk image backup:

a. Click `Manual add` to add disk image objects.

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

c. Select the client and click `Next`.

d. Specify `General Object Options` and `Advanced Object Options`. For information on these options, press **F1**.

e. In the `Disk Image Object Options` window, specify disk image sections.

UNIX:

Specify a rawdisk section:

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

Specify a raw logical volume section:

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

IMPORTANT   To perform instant recovery, specify all raw logical volumes inside the volume group. Otherwise, instant recovery may fail or require workarounds, and there is a risk of data corruption.

Windows:

Use the following format:

\\.\PHYSICALDRIVE#

where # is the current number of the disk to be backed up.

For information on finding current numbers of disks (physical drive numbers), see the online Help index: "disk image backups".

f. Click `Finish` and `Next`.

9. Save your backup specification. For information on starting and scheduling ZDB sessions, see "Running and Scheduling ZDB Sessions" on page A-2.

NOTE Preview backup is not supported.

Backup Options

The following tables describe VA backup options. See also "VA and EVA Integrations" on page A-29.

Table 2-2 **VA Client Systems Options**

`Application system`	System on which the application runs. In cluster environments, specify the virtual server hostname (rather than the physical node hostname).
`Backup system`	System to which the data will be backed up. In cluster environments, specify the virtual server hostname (rather than the physical node hostname).

Table 2-3 **VA Instant Recovery Options**

`Track the replica for instant recovery`	Select this option to perform a ZDB to disk/ZDB to disk+tape and leave the replica on a disk array for instant recovery. Also, specify `Number of replicas rotated`. `Use an existing replica` and `Keep the replica after the backup` are automatically selected.
	If this option is not set, you cannot perform instant recovery from the replica created or reused in this session.
	Default: selected.
`Number of replicas rotated`	During ZDB sessions, Data Protector creates a new replica and leaves it on the array until the specified `Number of replicas rotated` is reached (specify if you selected `Track the replica for instant recovery`). After that, the oldest replica is reused.
	Default: 1. Maximum: 1024.

Table 2-4　　　　　　**VA Replica Management Options**

Use an existing replica	By default, this option is automatically selected if `Track the replica for instant recovery` is set, and cannot be deselected.
	If configuring a ZDB to tape, select this option to reuse an existing replica.
	The replica can be reused only if it already exists on a disk array. Only replicas not marked for instant recovery, or replicas with snapshots not listed in the VA LUN exclude file can be reused.
Keep the replica after the backup	By default, this option is automatically selected if `Track the replica for instant recovery` is set, and cannot be deselected.
	For ZDB to tape, select this option to keep the replica on a disk array after backup. This replica is not available for instant recovery, but can be reused in future backup sessions using the same backup specification (`Use an existing replica` selected).
	If this option is not selected, the replica is deleted after backup.
Leave the backup system enabled	Available if `Keep the replica after the backup` is selected.
	By default, Data Protector dismounts filesystems (all platforms), deactivates volume groups (HP-UX) and removes volume groups from /etc/lvmtab (HP-UX) on the backup system after each backup. If this option is selected, filesystems remain mounted (all platforms), volume groups remain activated (HP-UX) and are not removed from /etc/lvmtab after backup.
	Thus, you can use the backup system for data warehouse activities, *but not for instant recovery*.
	Default: not selected.
Integrate with VA LUN security	Specify this option to apply LUN security to child LUNs (target volumes or snapshots) created by the integration.
	If Secure Manager is activated, specify this option and configure passwords correctly; otherwise, the backup will fail.
	For more information, see "LUN Security" on page 9.
	Default: not selected.

Table 2-5   **VA Mount Options**

`Dismount the filesystems on the application system`	Select this option to dismount a filesystem on the application system before snapshot creation and remount it afterwards. Filesystems do not have the stop I/O functionality to flush the data from the filesystem cache to disk and stop I/O for the snapshot creation time.
	If integrated applications (for example, Oracle) run on the filesystem, they control I/O to disk, so it is not necessary to dismount filesystems before snapshot creation.
	Default: not selected.
`Enable the backup system in read/write mode`	HP-UX only (on Windows, filesystems are always mounted in read/write mode).
	Select this option to have read/write access to volume groups and filesystems on the backup system. For backup, it is sufficient to activate backup system volume groups and filesystems in read-only mode. For other tasks, read/write mode may be needed.
	Default: not selected.

Table 2-6   **VA Application Options**

`Stop/quiesce the application`	Create the optional `Stop/quiesce the application` command in /opt/omni/lbin (HP-UX) or `<Data_Protector_home>`\bin (Windows) on the application system, and specify the filename in the backup specification.
	This command is executed on the application system before snapshot creation. It is mainly used to stop applications not integrated with Data Protector.
	If this command fails, `Restart the application` (see below) is not executed. Therefore, you need to implement a cleanup procedure in `Stop/quiesce the application`. Note that if the `ZDB_ALWAYS_POST_SCRIPT` omnirc variable is set to `1`, `Restart the application` is always executed if set (default is `0`). See "ZDB Agents Omnirc Variables" on page A-18 for details.

Table 2-6   **VA Application Options**

Restart the application	Create the optional `Restart the application` command in `/opt/omni/lbin` (HP-UX) or `<Data_Protector_home>\bin` (Windows) on the application system, and specify the filename in the backup specification.
	This command is executed on the application system after snapshot creation. It is mainly used to restart applications not integrated with Data Protector.

Charts below provide detailed backup flows according to the backup options selected.

Figure 2-3   **ZDB-to-Disk Session**
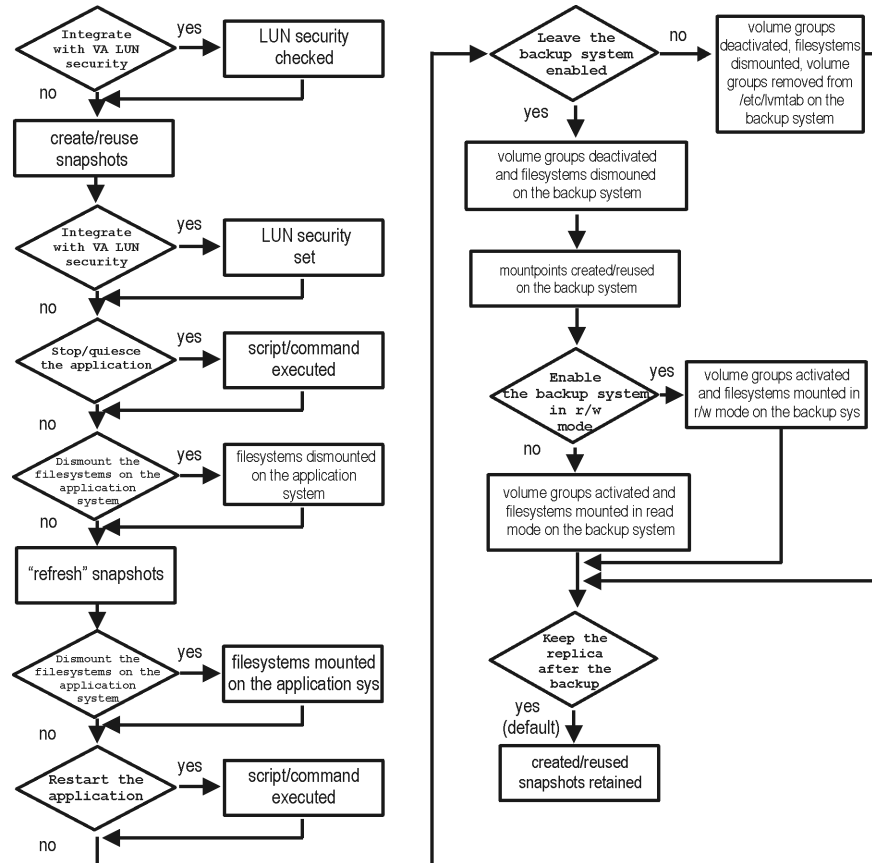
Figure 2-4   **ZDB-to-Tape, ZDB-to-Disk+Tape Session**
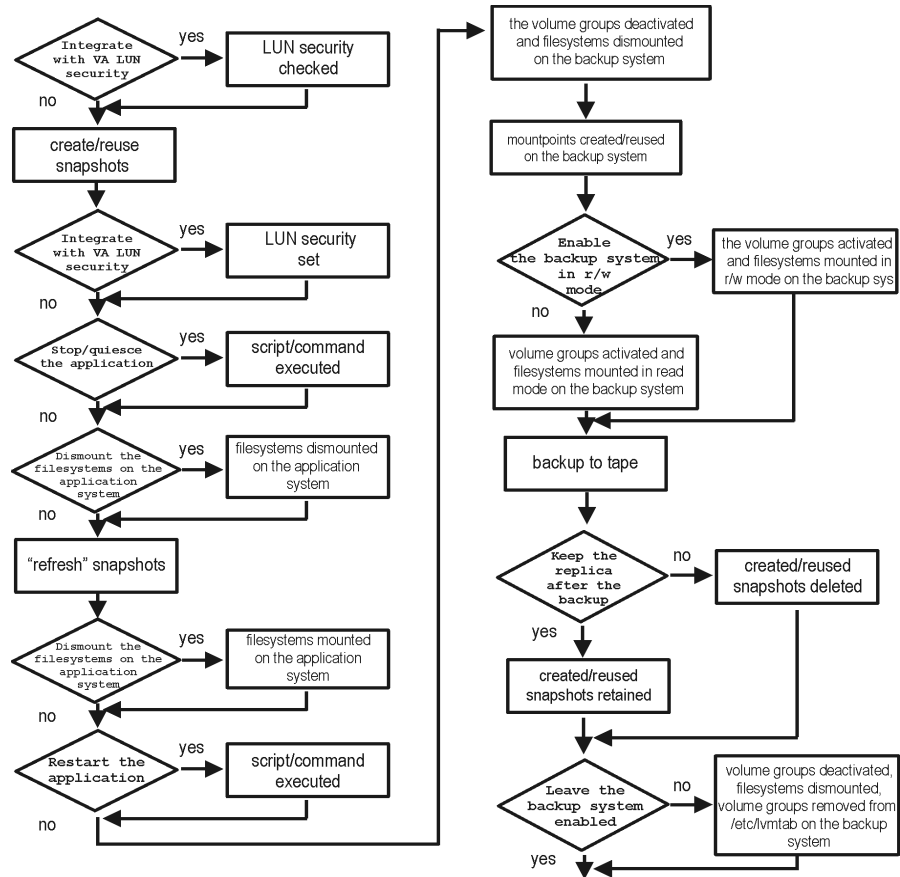
— The *create/reuse snapshots* phase represents an operation in which CommandView SDM allocates space for snapshot creation. The actual writing operation is performed during the *"refresh" snapshots* phase. Since the subsequent write operations are fast because they do not need to allocate space, the time frame in which the application is in backup mode (online backup) or down (offline backup) is minimal.

— For ZDB to tape, you can select the option Keep the replica after the backup. For ZDB to disk+tape, this option is selected by default and cannot be deselected.

3 Restore

Overview

This chapter describes configuring and running a filesystem or disk image restore of the data backed up using the VA integration. The sections describe restore procedures using the Data Protector GUI and CLI.

The data backed up in a ZDB session can be stored on a disk array (ZDB to disk, ZDB to disk+tape), or on backup media (ZDB to tape, ZDB to disk+tape).

Available restore types are:

- Restore from backup media on LAN (standard restore). See "Standard Restore" on page 33.

- Instant recovery. See "Instant Recovery" on page 34.

Table 3-1   **Restore Types**

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

Standard Restore

Data backed up in ZDB-to-tape and ZDB-to-disk+tape sessions can be restored from the backup media to the application system through a LAN. For information on this restore type, see the *HP OpenView Storage Data Protector Administrator's Guide*.

TIP You can improve the data transfer rate by connecting a backup device to the application system. See the *HP OpenView Storage Data Protector Administrator's Guide* for more information on configuring backup devices. See "Restoring Under Another Device" in the same guide for more information on restoring using another device.

Instant Recovery

Instant recovery restores data directly from a replica to source volumes, without involving a backup device. All data in the replica is restored (regardless of selections during backup). For instant recovery concepts, see the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

You can perform instant recovery using Data Protector GUI (see "Instant Recovery Using the GUI" on page 35) or CLI (see "Instant Recovery Using the CLI" on page 37).

The number of replicas available for instant recovery is limited by Number of replicas rotated, which sets the size of the replica set. You can view these replicas in the Instant Recovery GUI context by expanding Restore Sessions. Replicas are identified by the backup specification name and the session ID. Other information, such as time when the replica was created, is also provided. Alternately, you can run the omnidbva -session CLI command.

Considerations
- When instant recovery starts, the application system needs to be disabled. This includes dismounting filesystems and deactivating volume groups (UNIX). Before this is done, filesystems' and volume groups' status is checked, and only mounted filesystems and activated volume groups are dismounted and deactivated. At the end of the session, only previously dismounted filesystems are mounted, and only previously deactivated volume groups are activated.

- All target volumes (child LUNs or snapshots) for the source volume (parent LUN) involved in instant recovery are deleted automatically before the recovery. The only target volumes (replica) kept on the array are those restored if the Keep the replica after the restore option is selected.

  If there are other target volumes of the same source volume on the array (for instance, from another backup specification or created for purposes other than Data Protector backup and restore), instant recovery fails. Therefore, these target volumes must be deleted before instant recovery.

- In a ZDB to disk+tape session, the streaming of files from the replica to tape impacts the files access time, change time (UNIX), and archive (Windows) attributes, depending on how you select the Do not

preserve access time attributes and Do not use archive
attribute filesystem options. For more information on the options,
see online Help.

IMPORTANT   After instant recovery, restored filesystems are mounted to the same
mount points/drive letters as they were at the backup time. If these
mount points/drive letters have other filesystems mounted, these
filesystems are automatically dismounted before instant recovery, and
restored filesystems are mounted afterwards.

For more information about VA instant recovery considerations and
limitations, see the *HP OpenView Storage Data Protector Software
Release Notes* and the *HP OpenView Storage Data Protector Zero
Downtime Backup Concepts Guide*.

Instant Recovery Procedure

Prerequisites   • For disk image instant recovery, manually dismount the disks to be
restored before instant recovery, and re-mount them afterwards.

IMPORTANT   During an instant recovery session, do not perform ZDB using the source
volumes to which the data is being restored.

Instant Recovery Using the GUI

1. In the Context List, select Instant Recovery.

2. Select the backup session (replica) from which you want to perform
   the recovery. This can be done by selecting:

   • Backup session ID and name (in the Scoping Pane, expand
     Restore Sessions and select the session from a list of
     ZDB-to-disk/ZDB-to-disk+tape sessions)

   • Backup type (filesystem, Oracle, SAP R/3,...) and backup session
     name and ID:

     a. In the Scoping Pane, expand Restore Objects.

Backed up object types (Filesystem, Disk Image, SAP R/3, Microsoft SQL Server, ...) are displayed.

b. Expand the object type you want to restore.

All available backup specifications used in ZDB-to-disk or ZDB-to-disk+tape sessions for the selected object type are displayed.

c. Expand the backup specification containing the required objects. Available sessions are displayed:

Figure 3-1   Selecting a Session
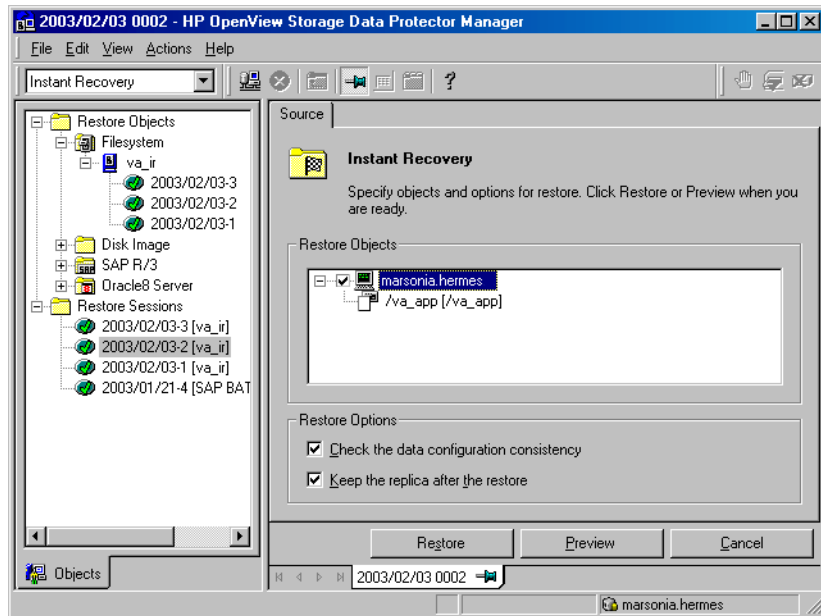
3. In the Scoping Pane, click the backup session you want to restore.

The application system and its mount points/drive letters representing source volumes backed up during the selected session are displayed.

4. Check the selection box next to the application system to select the session for restore. You cannot select sub-components because instant recovery restores the complete session.

5. Specify instant recovery options (see "Instant Recovery Options" on page 38 or press **F1** for information).

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

7. Select Start Restore Session. It is recommended to test instant recovery first to ensure it works properly.

IMPORTANT　　　You cannot use the GUI to perform instant recovery from ZDB to disk+tape after exporting or overwriting the media used in the session. Use the CLI instead. Note that backup media must not be exported or overwritten even after an object copy session.

Instant Recovery Using the CLI

1. List all available ZDB-to-disk/ZDB-to-disk+tape sessions (identified by the session ID):

   omnidbva -session

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

2. Execute:

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

   Where:

   <application_system_name> Application system hostname.

   <SessionID>　　Backup session ID (step 1 of this procedure).

   For *<INSTANT RECOVERY OPTIONS>*, see Table 3-2 on page 38.

See omnidbva and omnir man pages for details.

Instant Recovery Options

Table 3-2   **Instant Recovery Options**

Data Protector GUI/CLI	Function
`Check the data configuration consistency/-check_config`	If this option is selected (default), the current configuration of participating volume groups is compared with the volume group configuration during the backup session kept in VADB. If the configuration changed since then, the restore fails.
	When instant recovery is performed in an MC/ServiceGuard cluster to another node (not the one that was backed up), leave this option selected. The current volume group configuration on the node to which instant recovery is being performed is different from the configuration kept in VADB. With this option selected, VADB volume group configuration is replaced by the current configuration on the node to which instant recovery is being performed, and the session is not aborted.
	CRC check information for the data in the source volumes is compared to the CRC check information for the data in the selected replica. If this does not match, the session fails.
`Keep the replica after the restore/-keep_version`	If this option is selected (default), the replica from which the data was restored is left on the disk array after restore.
	Even if the restore is successful, it is recommended to keep the replica until next backup.

Instant Recovery in a Cluster

Instant recovery with an application/filesystem running in an MC/ServiceGuard or Microsoft Cluster Server on the application system requires some additional steps. See "Instant Recovery in a Cluster" on page A-15 for instructions.

4 Troubleshooting

Before You Begin

This chapter contains a list of general checks and verifications and a list of problems you may encounter when using the VA integration. For general Data Protector troubleshooting information, see the *HP OpenView Storage Data Protector Troubleshooting Guide*.

✓ Ensure that the latest official Data Protector patches are installed. See the online Help index: "patches" on how to verify this.

✓ See the *HP OpenView Storage Data Protector Software Release Notes* for general Data Protector and integration-specific limitations, as well as known problem and workarounds.

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

Checks and Verifications

✓ On the application and backup systems, examine system errors reported in:

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

Windows: `<Data_Protector_home>\log\debug.log`

Backup Problems

Problem   **You cannot select StorageWorks mode in the Data Protector user interface when creating a backup specification**

Action   Check that the HP StorageWorks VA Agent integration module is installed on the application and backup systems. To do that, open the cell info file located in:

Windows Cell Manager:
<Data_Protector_Home>\Config\server\cell\cell_info

UNIX Cell Manager: /etc/opt/omni/server/cell/cell_info

File contents should look similar to the following:

-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

Problem   **On the application system, dismounting a filesystem fails**

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

Problem   **System error is reported**

[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. '.

VA agent is not installed on the backup system.

Action   Install the HP StorageWorks VA Agent component on the backup system.

Problem   **Snapshots cannot be mounted to the target location on the backup system (Windows)**

[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.).

When backup with nested mountpoint objects is run, snapshots cannot be mounted to the target mountpoint location on the backup system if cleaning of the target mountpoint location fails.

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 (default), empty the tmp directory.

Problem   **Backup session fails when using an existing replica**

Action   When using a pre-configured replica or the replica created in previous backup sessions, check that parent LUNs used in the current session are the same as in the previous backup session. The backup specification should not change in the backup objects; otherwise, there will be no existing snapshots connected with selected backup objects.

Problem   **ZDB to disk or ZDB to disk+tape fails with the following message:**

[Critical]

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

This may be caused by:

- The backup specification was modified in the Source property page (where you select the backup data).

Action   Check that the backup specification was not modified (backup objects remained unchanged) since you started to use it for replica set rotation. If a change is required, but the backup specification is already in use, create a new backup specification.

- The snapshot (child LUN) logged in VADB and used for replica set rotation is missing on the VA storage system.

Action   Check if a snapshot used for replica set rotation exists on the VA storage system. If it is missing, delete all information about the backup session that created this snapshot:

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

Instant Recovery Problems

Problem   **Instant recovery session fails if the application system is in a cluster**

[Critical]

Data consistency check failed!

Configuration of volume group *<vg_name>* has changed since the last backup session!

The problem occurs if the option Check data configuration consistency is selected and may be caused by:

- Failover to a secondary node

- Changed volume group configuration on the application system

Action   Ensure that the volume group configuration on the application system has not changed and/or deselect Check the data configuration consistency, and then restart instant recovery.

Problem   **Instant recovery to a different cluster node fails (Windows)**

[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.

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

Action   Disable the automatic preparation of the application system:

1. On the application system, enable the ZDB_IR_MANUAL_AS_PREPARATION variable (see "ZDB Agents Omnirc Variables" on page A-18) and manually dismount the volumes to be restored.

2. Start instant recovery.

3. After instant recovery, manually mount restored volumes.

- If the Check the data configuration consistency instant recovery option is selected, the following happens:

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

  ✓ The CRC check information for the data in the source volumes stored in 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.

- 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 session is aborted.

- A check is 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 session is aborted.

II HP StorageWorks Enterprise Virtual Array

5   Configuration and Maintenance

Overview

This chapter describes configuration and maintenance of the Data
Protector HP StorageWorks Enterprise Virtual Array (EVA) integration.
It also provides information on the ZDB database.

Prerequisites

- Install:

    EVA components:

    — HP StorageWorks Virtual Controller Software (VCS) and
      Command View (CV) EVA. See the VCS and CV EVA
      documentation for installation instructions. See the *HP OpenView
      Storage Data Protector Software Release Notes* for information on
      supported product versions.

    — HP OpenView Storage Operations Manager (HP StorageWorks
      SMI-S EVA provider part) starting with v1.1 to use SMI-S Agent.
      If you use EVA Agent (legacy), do not install this application.

    — Business Copy (BC) EVA microcode and license. TBD

    — A license for controlling the EVA storage system.

    — SANworks Snapshot licenses.

    Data Protector components:

    — A license for using the EVA integration.

    — HP StorageWorks EVA Agent (legacy) or HP StorageWorks EVA
      SMI-S Agent. Note that you may need to upgrade EVA Agent
      (legacy) to SMI-S Agent.

    For installation/upgrade instructions, see the *HP OpenView Storage
    Data Protector Installation and Licensing Guide*.

- Make sure the same operating system (and its version) is installed on
  the application and backup systems.

- Connect EVA to the application and backup systems through the
  SAN.

- Using 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.

- For ZDB to disk, configure a backup device (for example, a standalone file device), as you cannot configure a backup specification without selecting a device. For instructions, see the online Help index: "standalone devices".

See the *HP OpenView Storage Data Protector Software Release Notes* for information on:

- General Data Protector and integration-specific limitations

- Supported platforms and integrations

- Supported integrations

- Supported backup and connectivity topologies

For information on supported configurations, see the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

ZDB Database

ZDB database is referred to as:

- **EVADB** if you use HP StorageWorks EVA Agent (legacy)

- **SMISDB** if you use HP StorageWorks EVA SMI-S Agent

EVADB/SMISDB keeps information about:

- Management Appliance systems TBD configured in the Data Protector cell. For each system, the following is stored:

  — Hostname as recognized in the IP network.

  — Port numbers through which the agents communicate. The default numbers are:
  EVA Agent (legacy) - 12301, SMI-S Agent - 5988.

  — User name and encoded password for CV EVA (EVA Agent (legacy)) or SMI-S EVA provider (SMI-S Agent) login

- EVA hardware configuration (EVA Agent (legacy)).

- Disk groups for snapclone destination.

- Volume group configuration of source volumes (UNIX only).

- Replicas kept on the array. This information includes:

  — ID of the session that produced the replica

  — Time when the session was performed

  — Name of the backup specification used in the session

  — Name, ID, and WWN of the target volume created in the session

  — Name and ID of the EVA storage system on which the target volume resides

  — Target volume type (vsnap, standard snapshot, snapclone)

  — ID of the source volume used in the session

  — IR flag (indicating that the target volume can be used for instant recovery)

  — Purge flag (indicating that the target volume must be deleted)

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

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

- Filesystem and volume management system information.

EVADB/SMISDB resides on the Cell Manager in:

- *Windows:* `<Data_Protector_home>`\db40\evadb or `<Data_Protector_home>`\db40\smisdb

- *UNIX:* /var/opt/omni/server/db40/evadb or

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

Configuring the Integration

To configure the integration:

- Provide the login information for CV EVA running on the Management Appliance system. See "Setting Login Information for CV EVA" on page 54.

- If needed, set disk group pairs. See "EVA Disk Group Pairs Configuration File" on page 55.

Before you start configuration, make sure you met the prerequisites described in "Overview" on page 50.

Setting Login Information for CV EVA

Before starting ZDB sessions, provide the login information for CV EVA running on the Management Appliance system.

To set, delete, and list the login information, use omnidbeva or omnidbsmis commands. See omnidbeva and omnidbsmis man pages for command syntax and examples.

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

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

- If standby and failed Management Appliances have the same hostname, no action is needed.

- If standby and failed Management Appliance have different hostnames, remove the failed system from the Data Protector configuration, and then add the new Management Appliance.

NOTE   If the EVA hardware configuration changes after providing the login information, perform a rescan operation (see "Updating the Information on EVA Hardware Configuration in EVADB" on page 57). This applies only if you use EVA Agent (legacy).

IMPORTANT   It is recommended to use default port numbers (EVA Agent (legacy) - 12301, SMI-S Agent - 5988).

EVA Disk Group Pairs Configuration File

You can allocate snapclones to a different disk group from that used for source volumes (original virtual disks). Thus, you influence the application performance, since different physical disks are used for read and write operations on source volumes and the replica.

To set disk group pairs, use omnidbeva or omnidbsmis commands. See omnidbeva and omnidbsmis man pages for command syntax and examples of manipulating the disk group pairs configuration file. The file template is as follows.

```
#
# HP OpenView Storage Data Protector A.06.00
#
# 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
```

Automatic Configuration of Backup System

When you start a ZDB session, Data Protector performs necessary configuration steps, such as configuring volume groups and filesystems on the backup system. 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 and mounts these filesystems during ZDB-to-tape or ZDB-to-disk+tape sessions. For more information on the backup system mountpoint creation, see the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

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

In large environments, the search for pre-configured hosts can be time-consuming. To skip the search and thus improve the performance, set the EVA DISABLE HBA AUTODETECTION variable to 1. After that, configure the backup system manually and use the EVA_HOSTNAMEALIASES variable to define the backup system object name. Use this procedure also if the automatic autodetection of the backup system fails.

For more information, see "ZDB Agents Omnirc Variables" on page A-18.

NOTE   SMI-S Agent does not automatically detect EVA host objects. Therefore, before a ZDB session, make sure the participating backup system is configured on the EVA storage system.

Maintaining the Integration

Maintenance tasks are divided into the following categories:

- Maintenance configuration tasks. See "Updating the Information on EVA Hardware Configuration in EVADB" on page 57

- Maintenance query tasks. See "Querying EVADB/SMISDB" on page 57

- Maintenance fix tasks. See "Synchronizing EVADB/SMISDB" on page 58.

- Maintenance deletion tasks. See "Purging EVADB/SMISDB" on page 58 and "Deleting EVADB/SMISDB Entries" on page 58.

Updating the Information on EVA Hardware Configuration in EVADB

To perform ZDB sessions, Data Protector needs to obtain information on the EVA hardware configuration (normally written to EVADB/SMISDB after you provide login details for a Management Appliance). However, you need to manually update EVADB when you change:

- Management Appliance system

- At least one of HSV EVA controllers

To rescan the EVA hardware configuration and thus update EVADB with the necessary information, use the omnidbeva command. See the omnidbeva man page for command syntax.

Querying EVADB/SMISDB

Using omnidbeva/omnidbsmis commands, you can list:

- All available backup sessions

- All available backup specifications

- All backup sessions based on a specific backup specification

- All available target volumes in EVADB

- Obsoleted volumes marked for purging

You can also get details on a specific backup session and target volume.

See `omnidbeva` and `omnidbsmis` man pages for command syntax and examples.

Synchronizing EVADB/SMISDB

During synchronization, EVA agent (on the backup system) synchronizes persistent data in EVADB/SMISDB with the current state of the EVA storage system. This ensures consistency of the ZDB database with the physical environment. If a target volume is physically missing from the storage system (for example, 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 start synchronizing, use `omnidbeva` or `omnidbsmis` commands. See `omnidbeva` and `omnidbsmis` man pages for command syntax.

Purging EVADB/SMISDB

During purge (normally started at the beginning of the backup session), EVA agent attempts to delete virtual disks marked for purging. You can also run SMISDB purge manually using the `omnidbsmis` command. See the `omnidbsmis` man page for more information.

Deleting EVADB/SMISDB Entries

Using `omnidbeva` or `omnidbsmis` commands, you can delete:

- Information about a specific backup session (replica), identified by the session ID

- Information about backup sessions based on a specific backup specification (all replicas in a replica set), identified by the backup specification name

See `omnidbeva` and `omnidbsmis` man pages for command syntax and examples.

IMPORTANT   Although the commands do not remove the target volumes that
constitute a replica, you cannot perform instant recovery from such a
replica because all session information was deleted.

6 **Backup**

Overview

This chapter describes configuring a filesystem and disk image ZDB using the Data Protector GUI.

You should be familiar with EVA concepts and procedures and basic Data Protector ZDB and instant recovery functionality. See the EVA-related documentation and the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

Considerations

- Only one type of target volume per source volume can exist on an EVA disk array at the same time.

- When a cloning process of a source volume is in progress, another snapshot (any type) of that source volume cannot be created.

- A replica cannot be reused if:

  — a snapclone from this replica has a snapshot attached

  — a target volume from this replica is presented to a system

  This applies only if EVA Agent (legacy) is used.

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

- Creating a target volume from a snapclone is not supported.

For more information on backup-related considerations, see the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

Snapshot Types

Data Protector supports the following snapshot types:

- Snapshots *with* pre-allocation of disk space (**standard snapshots**).

- Snapshots *without* pre-allocation of disk space (**vsnaps** or Virtually Capacity-Free Snapshots).

- A full copy of the source volume (original virtual disk), independent of the original virtual disk (**snapclones)**.

You select the snapshot type in the GUI when creating a backup specification.

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

ZDB Types

You can classify ZDB sessions in two ways:

- According to the media to which data is backed up:

  — **ZDB to disk**

    The replica produced (snapclone) is kept on a disk array until reused. This replica becomes part of the replica set and can be used for instant recovery.

    ZDB to disk is performed if the option Track the replica for instant recovery is selected in a backup specification, and To disk is selected when running/scheduling a backup.

  — **ZDB to tape**

    The replica produced (any type) is streamed to backup media, typically tape.

    This replica is deleted after backup if the option Keep the replica after the backup *is not* selected in a backup specification. If this option *is* selected, the replica remains on a disk array until reused and becomes part of the replica set. However, it cannot be used for instant recovery.

  — **ZDB to disk+tape**

    The replica produced (snapclone) is kept on a disk array until reused and is also streamed to backup media. This replica becomes part of the replica set and can be used for instant recovery.

    ZDB to disk+tape is performed if the option Track the replica for instant recovery is selected in a backup specification, and To disk+tape is selected when running/ scheduling a backup.

For more information on ZDB types, see the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

• According to the amount of data that is backed up to tape:

Table 6-1   **Backup Types**

Full	Backs up all selected objects, even if there are no changes from the previous backup.
Incremental	Backs up changes from the last protected backup.

NOTE   The replica is created in the same way for both full and incremental ZDB. The difference between the two ZDB types is only in the amount of files that are streamed from the replica to tape.

Limitations   — Incremental ZDB is not supported for ZDB to disk.

— Incremental ZDB is supported only for filesystem backups.

Replica Creation and Reuse

A new replica is created and added to the replica set when:

• ZDB to tape is performed, in which Keep the replica after the backup *is* selected, but the specified Number of replicas rotated is not reached.

• ZDB to disk/ZDB to disk+tape is performed (Track the replica for instant recovery selected), and the specified Number of replicas rotated is not reached.

The oldest replica in the set is deleted and the new one is created when:

• ZDB to tape is performed in which Keep the replica after the backup *is* selected and the specified Number of replicas rotated is reached.

• ZDB to disk/ZDB to disk+tape is performed and the specified Number of replicas rotated is reached.

If the option Keep the replica after the backup is not selected, the replica and therefore all target volumes created during the backup session are deleted.

Configuring a Backup Specification

IMPORTANT   Before you begin, consider all limitations regarding the EVA integration. For more information, see the *HP OpenView Storage Data Protector Software Release Notes* and the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

1. In the Context List, select `Backup`.

2. Under `Client systems`, select `Application system` and `Backup system`.

   Specify options as follows:

   ZDB to disk, ZDB to disk+tape: Under `Instant recovery option`, select `Track the replica for instant recovery`. Under `Replica management options`, specify `Number of replicas rotated`. The maximum number for vsnaps and standard snapshots is 7 per source volume. GUI does not limit the number of replicas rotated, but the session fails if the limit is exceeded.

   A `Snapshot type` of `Snapclone` and `Snapshot policy` of `Strict` are automatically selected.

   See Figure 6-1 on page 66 and "Backup Options" on page 70.

Figure 6-1   **EVA Backup Options (ZDB to Disk, ZDB to Disk+Tape)**
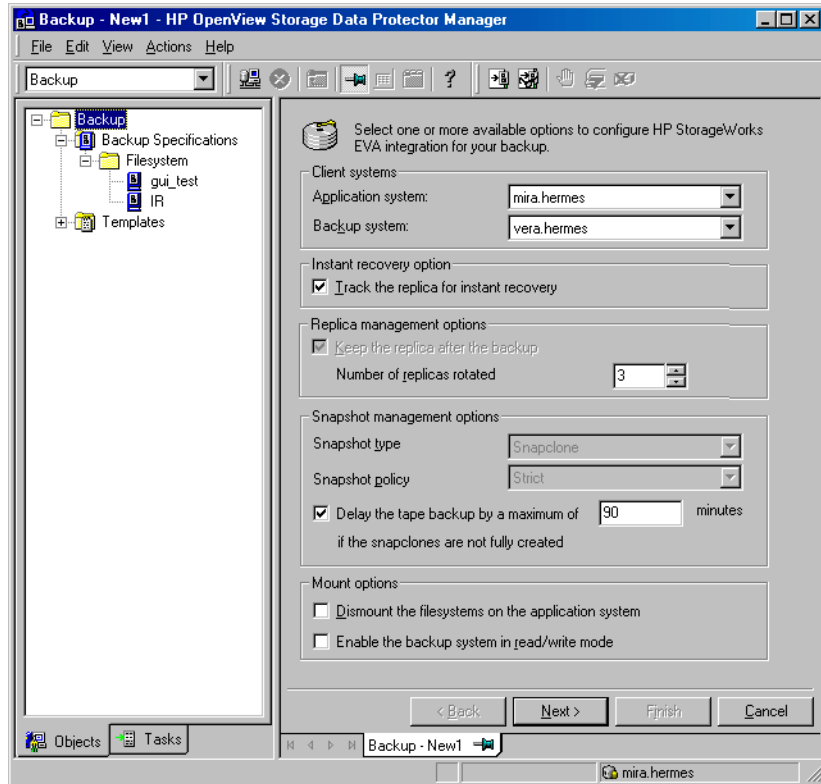
NOTE   You specify a ZDB-to-disk or ZDB-to-disk+tape session using the Split
mirror/snapshot backup option when running/scheduling a backup.
See "Running and Scheduling ZDB Sessions" on page A-2.

Click Next.

ZDB to tape: To keep the replica after backup, select Keep the
replica after the backup and specify Number of replicas
rotated. The maximum number for vsnaps and standard snapshots
is limited by the target EVA storage system. GUI does not limit the
number of replicas rotated, but the session fails if the limit is
exceeded.

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

TIP　　　　　For ZDB to disk+tape and ZDB to tape, select Delay the tape backup by a maximum of X minutes if the snapclones are not fully created. In this case, backup to tape starts when the cloning process finishes, but not later than the specified number of minutes. Thus, you prevent degradation of the application data access times during backup to tape.

For detailed information on these and Mount options, see "Backup Options" on page 70.

Click Next.

3. *Filesystem backup:* Expand application systems and select the objects to be backed up.

IMPORTANT　　To perform instant recovery on UNIX, select all filesystems inside the volume group. Otherwise, instant recovery may fail or require workarounds and there is a risk of data corruption.

Click Next.

Disk image backup: Click Next.

4. Select devices. Click Properties to set device concurrency, media pool, and preallocation policy. For descriptions of these options, click Help.

To create additional copies (mirrors) of backup, specify the desired number of mirrors by clicking Add mirror/Remove mirror. Select separate devices for each mirror backup.

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

NOTE　　　　Object mirroring and object copying are not supported for ZDB to disk.

Click Next.

5. 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 EVA backup options.

   You can specify Application options and modify all other options,
   except Application system and Backup system. See "Backup
   Options" on page 70.

   In the Filesystem Options group box, click Advanced and specify
   filesystem options as desired. For information, press **F1**.

   Windows only: If you plan to do incremental ZDB, select the Do not
   use archive attribute filesystem option in the WinFSOptions
   page.

6. Following the wizard, open the scheduler (for information, press **F1** or
   see "Running and Scheduling ZDB Sessions" on page A-2), and then
   the backup summary page.

7. *Filesystem backup:* Click Next.

   Disk image backup:

   a. Click Manual add to add disk image objects.

   b. Select Disk image object and click Next.

   c. Select the client and click Next.

   d. Specify General Object Options and Advanced Object
      Options. For information on these options, press **F1**.

   e. In the Disk Image Object Options window, specify disk image
      sections.

      UNIX:

      Specify a rawdisk section:

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

      On HP-UX, specify a raw logical volume section:

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

IMPORTANT For instant recovery, specify all raw logical volumes inside the volume group. Otherwise, instant recovery may fail or require workarounds, and there is a risk of data corruption.

Windows:

Use the following format:

`\\.\PHYSICALDRIVE#`

where # is the current number of the disk to be backed up.

For information on finding current disk numbers (physical drive numbers), see the online Help index: "disk image backups".

   f. Click `Finish` and `Next`.

8. Save your backup specification. For information on starting and scheduling ZDB sessions, see "Running and Scheduling ZDB Sessions" on page A-2.

NOTE Preview backup is not supported.

Backup Options

The following tables describe EVA backup options. See also "VA and EVA Integrations" on page A-29.

Table 6-2   **EVA Client Systems Options**

Application system	System on which the application runs. In cluster environments, specify the virtual server hostname (rather than the physical node hostname).
Backup system	System to which your data will be backed up. In cluster environments, specify the virtual server hostname (rather than the physical node hostname).

Table 6-3   **EVA Instant Recovery Option**

Track the replica for instant recovery	Select this option to perform a ZDB to disk/ZDB to disk+tape and leave the replica on a disk array for instant recovery. Also, specify Number of replicas rotated. Keep the replica after the backup, Snapclone, and Strict are automatically selected.
	If this option is not set, you cannot perform instant recovery from the replica created or reused in this session.
	Default: selected.

Table 6-4   **EVA Replica Management Options**

Keep the replica after the backup	By default, this option is automatically selected if Track the replica for instant recovery is set, and cannot be deselected.
	If configuring a ZDB to tape, select this option to keep the replica on a disk array after backup. The replica becomes part of a replica set (specify Number of replicas rotated), but it *is not available for instant recovery*.
	If this option is not selected, the replica is deleted after backup.

Table 6-4	**EVA Replica Management Options**

`Number of replicas rotated`	During ZDB sessions, Data Protector creates a new replica and leaves it on the array until the specified `Number of replicas rotated` is reached (specify if you selected `Keep the replica after the backup`). After that, the oldest replica is reused.
	Default: 1.
	The maximum number for vsnaps and standard snapshots is 7. Data Protector does not limit the number of replicas rotated, but the session fails if the limit is exceeded.

Table 6-5	**EVA Snapshot Management Options**

Snapshot type	`Vsnap`	Creates snapshots without the pre-allocation of disk space. This is the default selection.
	`Standard snapshot`	Creates snapshots with the pre-allocation of disk space.
	`Snapclone`	Creates a clone of the source volume.
For more information on snapshot types, see the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*		
Snapshot policy	`Strict`	Data Protector attempts to create snapshots as specified by the `Snapshot type` option. If the source volumes used in the session have existing snapshots of a different type, the selected type is not created (the session is aborted). Default: selected.
	`Loose`	Data Protector creates snapshots of a different type than specified in `Snapshot type` (if this helps to complete a session successfully). For example, if you select standard snapshots, but Data Protector detects that standard snapshots cannot be created because vsnaps/snapclones of the source volumes already exist in a replica set, it creates either vsnaps or snapclones instead of standard snapshots. Note that Data Protector can use only one type of snapshots in a backup session. For example, if the source volumes used in a session have existing standard snapshots/vsnaps, the backup session is aborted.

Table 6-5 **EVA Snapshot Management Options**

`Delay the tape backup by a maximum of X minutes if the snapclones are not fully created`	Available if `Snapclone` is selected as a snapshot type. Prevents degradation of the application data access times by delaying moving data to tape until the cloning process completes (ZDB to tape, ZDB to disk+tape). Set the maximum waiting time. When the specified time is reached, backup to tape starts (even if the cloning process is not finished). Default: 90 minutes.

Table 6-6 **EVA Mount Options**

`Dismount the filesystems on the application system`	Select this option to dismount a filesystem on the application system before snapshot creation and remount it afterwards. Filesystems do not have the stop I/O functionality to flush the data from the filesystem cache to disk and stop I/O for the snapshot creation time. If integrated applications (for example, Oracle) run on the filesystem, they control I/O to disk, so it is not necessary to dismount filesystems before snapshot creation. Default: not selected.
`Enable the backup system in read/write mode`	HP-UX only (on Windows, filesystems are always mounted in read/write mode). Select this option to have read/write access to 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. For other tasks, read/write mode may be needed. Default: not selected.

Table 6-7   **EVA Application Options**

Stop/quiesce the application	Create the optional `Stop/quiesce the application` command in /opt/omni/lbin (UNIX) or *<Data_Protector_home>*\bin (Windows) on the application system, and specify the filename in the backup specification.
	This command is executed on the application system before snapshot creation. It is mainly used to stop applications not integrated with Data Protector.
	If this command fails, `Restart the application` (see below) is not executed. Therefore, you need to implement a cleanup procedure in `Stop/quiesce the application`. Note that if the `ZDB_ALWAYS_POST_SCRIPT` variable is set to `1`, `Restart the application` is always executed if set (default is `0`). See "ZDB Agents Omnirc Variables" on page A-18 for details.
Restart the application	Create the optional `Restart the application` command in /opt/omni/lbin (UNIX) or *<Data_Protector_home>*\bin (Windows) on the application system, and specify the filename in the backup specification.
	This command is executed on the application system after snapshot creation. It is mainly used to restart applications not integrated with Data Protector.

Charts below provide detailed backup flows according to the backup options selected.

Figure 6-2   **ZDB-to-Disk Session**
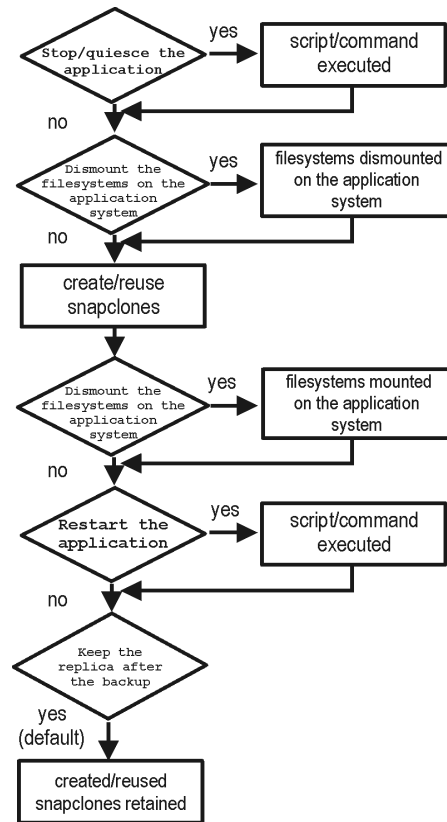
Figure 6-3   **ZDB-to-Tape, ZDB-to-Disk+Tape Session**
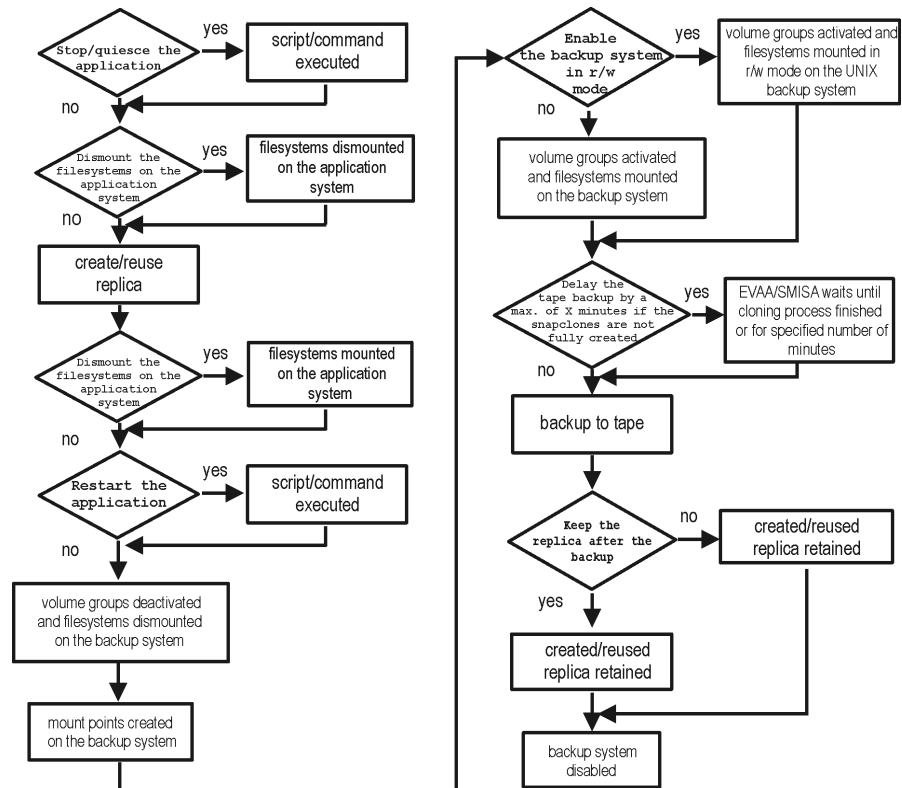
— *"Reuse"* means that target volumes from the oldest replica are deleted and a new replica is created.

— Due to the EVA limitation, snapclone creation (ZDB to disk, ZDB to disk+tape) may fail if a target volume of another type exists on the array. Such target volumes should be deleted first.

— `Enable the backup system in read/write mode` is ignored for ZDB to disk.

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

Backup
Backup Options

7 Restore

Overview

This chapter describes configuring and running a filesystem or disk image restore of the data backed up using the EVA integration. The sections describe restore procedures using the Data Protector GUI and CLI.

The data backed up in a ZDB session can be stored on a disk array (ZDB to disk, ZDB to disk+tape), or on backup media (ZDB to tape, ZDB to disk+tape).

Available restore types are:

- Restore from backup media on LAN (standard restore). See "Standard Restore" on page 79.

- Instant recovery. See "Instant Recovery" on page 80.

Table 7-1   **Restore Types**

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

Standard Restore

Data backed up in ZDB-to-tape and ZDB-to-disk+tape sessions can be restored from the backup media to the application system through a LAN. For information on this restore type, see the *HP OpenView Storage Data Protector Administrator's Guide*.

TIP   You can improve the data transfer rate by connecting a backup device to the application system. See the *HP OpenView Storage Data Protector Administrator's Guide* for more information on configuring backup devices. See "Restoring Under Another Device" in the same guide for more information on restoring using another device.

Instant Recovery

Instant recovery restores data directly from a replica to source volumes, without involving a backup device. All data in the replica is restored (regardless of selections during backup). For instant recovery concepts, see the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

You can perform instant recovery using Data Protector GUI (see "Instant Recovery Using the GUI" on page 81) or CLI (see "Instant Recovery Using the CLI" on page 83).

The number of replicas available for instant recovery is limited by Number of replicas rotated, which sets the size of the replica set. You can view these replicas in the Instant Recovery GUI context by expanding Restore Sessions. Replicas are identified by the backup specification name and the session ID. Other information, such as time when the replica was created, is also provided. Alternately, you can use CLI to list seesions (see omnidbeva and omnidbsmis man pages for information).

Considerations
- When instant recovery starts, the application system needs to be disabled. This includes dismounting filesystems and deactivating volume groups (UNIX). Before this is done, filesystems' and volume groups' status is checked, and only mounted filesystems and activated volume groups are dismounted and deactivated. At the end of the session, only previously dismounted filesystems are mounted, and only previously deactivated volume groups are activated.

- After instant recovery, the replica used in the session becomes the source volumes (physical location of the original application data changes).

  Additionally, if a replica resides in another disk group than the source volumes, this replica disk group becomes the disk group of the source volumes.

- After instant recovery, the replica restored is deleted from the replica set. Therefore, you cannot run instant recovery using the same session ID.

- In a ZDB to disk+tape session, the streaming of files from the replica to tape impacts the files access time, change time (UNIX), and archive (Windows) attributes, depending on how you select the `Do not preserve access time attributes` and `Do not use archive attribute` filesystem options. For more information on the options, see online Help.

IMPORTANT   After instant recovery, restored filesystems are mounted to the same mount points/drive letters as they were at the backup time. If these mount points/drive letters have other filesystems mounted, these filesystems are automatically dismounted before instant recovery, and the restored filesystems are mounted afterwards.

For more information about EVA instant recovery considerations and limitations, see the *HP OpenView Storage Data Protector Software Release Notes* and the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

Instant Recovery Procedure

Prerequisites   • For disk image instant recovery, manually dismount the disks to be restored before instant recovery, and re-mount them afterwards.

IMPORTANT   During an instant recovery session, do not perform a ZDB using the source volumes to which the data is being restored.

Instant Recovery Using the GUI

1. In the Context List, select `Instant Recovery`.

2. Select the backup session (replica) from which you want to perform the recovery. This can be done by selecting:

   - Backup session ID and name (in the Scoping Pane, expand `Restore Sessions` and select the session from a list of ZDB-to-disk/ZDB-to-disk+tape sessions)

   - Backup type (filesystem, Oracle, SAP R/3,...) and backup session name and ID:

a. In the Scoping Pane, expand Restore Objects.

Backed up object types (Filesystem, Disk Image, SAP R/3, Microsoft SQL Server, ...) are displayed.

b. Expand the object type you want to restore.

All available backup specification used in ZDB-to-disk/ZDB-to-disk+tape sessions for the selected object type are displayed.

c. Expand the backup specification containing the required objects. Available sessions are displayed:

Figure 7-1   **Selecting a Session**
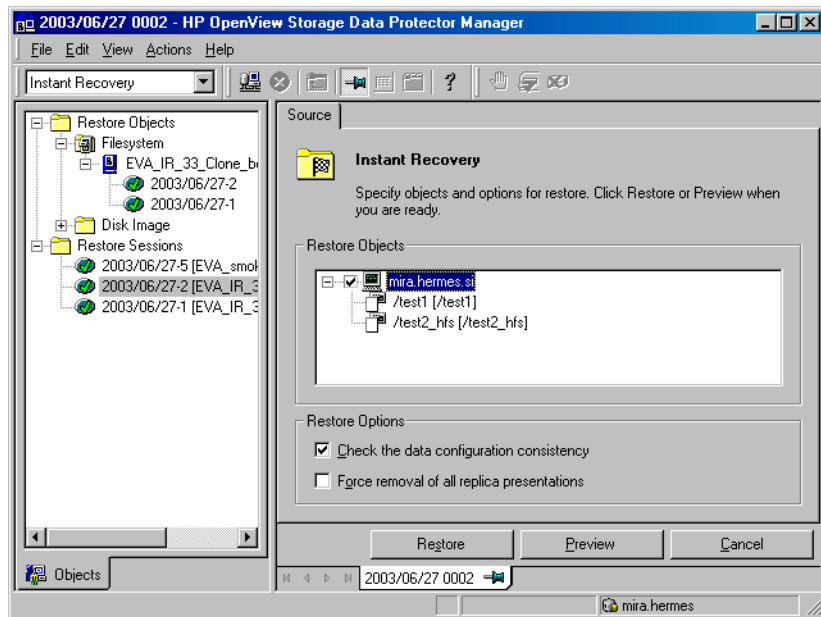
3. In the Scoping Pane, click the backup session you want to restore.

The application system and its mount points/drive letters representing source volumes backed up during the selected session are displayed.

4. Check the selection box next to the application system to select the session for restore. You cannot select sub-components because instant recovery restores the complete session.

5. Specify instant recovery options (see "Instant Recovery Options" on page 84 or press **F1** for information).

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

7. Select Start Restore Session. It is recommended to test instant recovery first to ensure it works properly.

IMPORTANT   You cannot use the GUI to perform instant recovery from ZDB to disk+tape after exporting or overwriting the media used in the session. Use the CLI instead. Note that backup media must not be exported or overwritten even after an object copy session.

Instant Recovery Using the CLI

1. List all available ZDB-to-disk/ZDB-to-disk+tape sessions (identified by the session ID):

   omnidbeva -list -session -ir

   if you have EVA Agent (legacy), or:

   omnidbsmis -list -session -ir

   if you have SMI-S Agent.

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

2. Execute:

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

   Where:

   <application_system_name> Application system hostname.

   <SessionID>   Backup session ID (step 1 of this procedure).

   For *<INSTANT RECOVERY OPTIONS>*, see Table 7-2 on page 84.

See omnidbeva, omnidbsmis, and omnir man pages for details.

Instant Recovery Options

Table 7-2   **Instant Recovery Options**

Data Protector GUI/CLI	Function
`Check the data configuration consistency/-check_config`	The current configuration of participating volume groups is compared with the volume group configuration during the backup session kept in EVADB/SMISDB. If the configuration changed since then, the restore fails. When instant recovery is performed in an MC/ServiceGuard cluster to another node (not the one that was backed up), leave this option selected. The current volume group configuration on the node to which instant recovery is being performed is different from the configuration kept in EVADB/SMISDB. With this option selected, EVADB/SMISDB volume group configuration is replaced by the current configuration on the node to which instant recovery is being performed, and the session is not aborted. Default: selected.
`Force removal of all replica presentations/-force_prp_ replica`	If this option is selected and any target volume to be restored is presented to a system, EVA agent removes these presentations. If this option is not selected but presentations exist, instant recovery fails. Default: not selected.

Instant Recovery in a Cluster

Instant recovery with an application/filesystem running in MC/ServiceGuard or Microsoft Cluster Server on the application system requires some additional steps. See "Instant Recovery in a Cluster" on page A-15 for instructions.

8   Troubleshooting

Before You Begin

This chapter contains a list of general checks and verifications and a list of problems you may encounter when using the EVA integration. For general Data Protector troubleshooting information, see the *HP OpenView Storage Data Protector Troubleshooting Guide*.

✓ Ensure that the latest official Data Protector patches are installed. See the online Help index: "patches" on how to verify this.

✓ See the *HP OpenView Storage Data Protector Software Release Notes* for general Data Protector and integration-specific limitations, as well as known problem and workarounds.

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

Checks and Verifications

✓ On the application and backup systems, examine system errors reported in:

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

Windows: `<Data_Protector_home>\log\debug.log`

Backup Problems

Problem   **You cannot select StorageWorks mode in the Data Protector user interface when creating a backup specification**

Action   Check that the HP StorageWorks EVA Agent (legacy) or HP StorageWorks EVA SMI-S Agent integration module is installed on the application and backup systems. To do that, open the cell_info file located in:

Windows Cell Manager:
<Data_Protector_Home>\Config\server\cell\cell_info

UNIX Cell Manager: /etc/opt/omni/server/cell/cell_info

File contents should look similar to the following:

-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

Problem   **EVA Agent (EVA Agent (legacy) or SMI-S Agent) fails to connect to Cell Manager and retrieve configuration data**

[Major]

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

EVA Agent is always started as an administrator's process on the application and backup systems. Therefore, the user who starts it must be the member of admin or operator user groups.

Action   Using the GUI, check if the user is a member of admin or operator user groups. If not, add the user to one of these groups. In addition, ensure that administrators from both the application and backup systems belong to Data Protector admin or operator.

Problem   **No HP StorageWorks SMI-S EVA provider entries are configured within SMISDB**

Action   Add an HP StorageWorks SMI-S EVA provider entry to SMISDB:

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

Problem   **On the application system, dismounting a filesystem fails**

Action   Ensure that no other processes use the filesystem to be dismounted. If
`Stop/quiesce the application` was specified, check that it stops all
processes using the filesystem.

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

The following message appears:

`[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. ).`

When a backup with nested mountpoint objects is run, snapshots cannot
be mounted to the target mountpoint location on the backup system if
cleaning of the target mountpoint location fails.

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` (default), empty the `tmp` directory.

Problem   **Adding login information for CV EVA using omnidbeva fails**

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

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

If you use CV EVA v3.0 or v3.1, the following login schemes are
available:

• new IIS scheme

• legacy Elm scheme

Action   When running `omnidbeva -empasswd -add`
<management_appliance_hostname>, use the legacy Elm scheme. The
login information may differ from what you have entered in CV 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`".

Problem   **EVA Agent (legacy) fails to obtain EVA storage system port
information**

`[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.
```

Action I   Check that CV EVA controlling the EVA storage system is configured
within the Data Protector cell by running `omnidbeva -empasswd -list`.
This lists all configured Management Appliances.

• If the corresponding Management Appliance is missing, provide its
login information:

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

• If the corresponding Management Appliance is configured but its
password changed, update the configuration:

`omnidbeva -empasswd -update`
<management_appliance_hostname>

Action II   If a new source volume exists on the EVA storage system, or the
configuration of the EVA storage system changed, manually rescan the
storage system configuration:

`omnidbeva -hwrescan`

Problem   **An error occurs while parsing a CV EVA response (EVA Agent
(legacy) problem)**

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

This may happen if port number 2301 is used for CV EVA v3.0 (instead of
12301).

Action   Switch the port to 12301 by running on the application or backup system:

```
omnidbeva -empasswd -update <management_appliance_ hostname>
-port 12301
```

Problem   **CV EVA command fails (EVA Agent (legacy) problem)**

A Command View EVA command failed on the server side.

This may also be followed by:

- 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

If CV EVA executing Data Protector requests returns an error, Data Protector retries such requests. However, the number of retries or the time interval between the retries may be insufficient on heavily loaded EVA storage systems.

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

Problem   **Data Protector cannot find FC HBAs on the backup system (EVA Agent (legacy) problem)**

[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>

Data Protector automatically detects local FC HBAs connecting the backup system to the target EVA storage system. This is necessary for the automatic backup system preparation. Detection may fail if:

• FC HBA drivers on the backup system do not support SNIA's Common HBA API

• There is a connectivity problem between the backup system and the target EVA storage system

Action   Check the connectivity and SAN zoning between the backup system and the target storage system. Then update FC HBA drivers on the backup system.

In addition, you may disable the automatic preparation of the backup system by setting the EVA_DISABLE_AUTODETECTION variable on the backup system to 1 and ensuring that the host object representing the backup system is configured within CV EVA. If the name of this host is different from the actual backup system hostname, specify the host object name in the EVA_HOSTNAMEALIASES variable on the backup system.

For more information, see "ZDB Agents Omnirc Variables" on page A-18.

Problem   **Preparation of the backup system takes longer than usual (EVA Agent (legacy) problem)**

[Normal]

Preparing the backup system for backup to tape.

Data Protector searches through the EVA storage system to check if the host object representing the backup system is already configured. In large environments, this can take a long time.

Action   Disable automatic preparation of the backup system by setting the EVA_DISABLE_HBA_AUTODETECTION variable on the backup system. Additionally, ensure that the host object representing the backup system is configured within CV EVA. If the name of this host is different from the actual backup system hostname, specify the host object name in the EVA_HOSTNAMEALIASES variable on the backup system.

For more information, see "ZDB Agents Omnirc Variables" on page A-18.

Problem   **Preparation of the backup system fails (EVA Agent (legacy) problem)**

[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!
```

This problem appears if automatic preparation of the backup system is not done and the backup system hostname cannot be found. In this case, either the host object representing the backup system is not configured within CV EVA, or the name of this host object is different from the actual backup system hostname.

Action   To automatically prepare the backup system, ensure that the EVA_DISABLE_HBA_AUTODETECTION variable is set to 0 (not enabled). If the variable is currently not enabled, the problem may be caused by FC HBA drivers on the backup system. Automatic preparation of the backup system is only possible if FC HBA drivers support SNIA's Common HBA API. To enable automatic preparation, update these drivers.

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

For more information, see "ZDB Agents Omnirc Variables" on page A-18.

Problem   **Data Protector cannot start device scanning on the backup system**

[Major]

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

After target volumes are created and presented to the backup system, device scanning runs on the backup system to detect target volumes.

On some systems, device scanning may take longer, and when the timeout is reached, the session aborts.

Action   Increase the number of retries for device scanning by setting the EVA_SCANDEVICES_LOCK_MAX_RETRY variable on the backup system.

If the problem persists, check if there is an active session hanging during preparation or resumption of the backup system. Abort such a session. See also the next problem.

Problem   **On HP-UX, backup session freezes during either preparation or resuming of the backup system**

One of the following messages appears:

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

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

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

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

If processes cannot be stopped, restart the backup system and clean it up manually:

1. On the backup system, run spmgr display to display the target volumes (created in the failed session) left under the Secure Path control.

2. Remove such target volumes from the Secure Path control using spmgr delete.

3. Run spmgr update, and then follow reported instructions to make changes persistent across reboots.

4. Using the CV EVA user interface, delete all presentations attached to removed target volumes.

Problem   **Data Protector cannot map a target volume to a valid character device file**

[Major]

Failed to map a target volume to a valid character device file:

Virtual disk name: *<Virtual_disk_name>*

StorageWorks EVA name: *<EVA_name>*

Data Protector cannot find target volumes on the backup system. This may be caused by:

• Connectivity problem between the backup system and the EVA storage system.

Action   Manually present an EVA virtual disk from the target EVA storage system to the backup host. If the virtual disk is not visible on the backup system, there may be a connectivity or host object configuration problem. See the EVA and Secure Path documentation to troubleshoot the problem.

• Problem with device drivers on the backup system (Windows 2000). This problem may occur if detecting new devices attached to the system takes a long time.

Action   By default, Data Protector waits for 300 seconds for a new physical drive to appear before aborting a session. You can increase this time by setting the SMB_SCAN_RDSK_TIMEOUT variable (see "ZDB Agents Omnirc Variables" on page A-18).

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

• Problem with the configuration entry for the backup system within the EVA storage system (HP-UX, Solaris). When you unpresent/delete EVA virtual disks without removing them from the Secure Path control first, device instances may be left in the driver configuration, which will disable detection of new devices.

Action   1. Run spmgr display on the backup system and check if such instances exist. Typically, they are reported as failed or inactive.

2. Delete such instances using spmgr delete.

3. Run spmgr update and follow the reported instructions to make changes persistent across reboots.

Problem **On Windows, Data Protector cannot find the volume on a target volume**

[Major]

Volume *<volume_name>* on target volume has not been found.

Data Protector found target volumes on the backup system, but the volumes on target volumes were not found in the system's device list. This can be caused by:

- 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 it does not exist, install the "Platform Kit for EVA".

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

Action Increase the Data Protector timeout for detecting new volumes by setting the SMB_SCAN_FOR_VOLUME_TIMEOUT variable on the backup system (see "ZDB Agents Omnirc Variables" on page A-18).

- The backup system is in inconsistent state.

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

I

Instant Recovery Problems

Problem   **Instant recovery fails** TBD

The problem may occur if the option Force removal of all replica
presentations is not selected and a snapclone from the selected replica
is presented to a host.

Action   Select the option Force removal of all replica presentations and
restart instant recovery.

Problem   **Instant recovery with the application system in a cluster fails**

[Critical]

Data consistency check failed!

Configuration of volume group *<vg_name>* has changed since the last
backup session!

The problem occurs if the option Check data configuration
consistency is selected and may be caused by:

• Failover to a secondary node

• Changed volume group configuration on the application system

Action   Ensure that the volume group configuration on the application system
has not changed and/or deselect Check the data configuration
consistency, and then restart instant recovery.

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

[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.

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

Action   Disable the automatic preparation of the application system:

1. On the application system, enable the ZDB_IR_MANUAL_AS_PREPARATION variable (see "ZDB Agents Omnirc Variables" on page A-18) and manually dismount the volumes to be restored.

2. Start instant recovery.

3. After instant recovery, manually mount restored volumes.

- If Check the data configuration consistency is selected, the following takes place:

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

- A check is 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 session is aborted.

III HP StorageWorks Disk Array XP

9   Configuration and Maintenance

Overview

This chapter describes configuration and maintenance of the Data
Protector HP StorageWorks Disk Array XP (XP) integration. It also
provides information on the ZDB database.

Prerequisites
- Install:

  XP components:

  — RAID Manager XP and RAID Manager/LIB XP on the application
  and backup systems. See the RAID Manager/LIB XP
  documentation for installation instructions.

  RAID Manager XP and RAID Manager Library are
  firmware-dependant. Consult the HP sales representative for
  information on which version of RAID Manager and RAID
  Manager Library to use.

  — Business Copy (BC) XP or Continuous Access (CA) XP microcode
  and license.

  — The xpinfo utility on the application and backup systems.
  Contact the HP support center to obtain this utility.

  Data Protector components:

  — A license for using the XP integration.

  — HP StorageWorks XP Agent.

  For installation instructions, see the *HP OpenView Storage Data
  Protector Installation and Licensing Guide*.

- Make sure the same operating system (and its version) is installed on
  the application and backup systems.

- Connect XP to the application and backup systems.

- Assign LUNs to the respective ports.

- For ZDB to disk, configure a backup device (for example, a standalone
  file device), as you cannot configure a backup specification without
  selecting a device. For instructions, see the online Help index:
  "standalone devices".

See the *HP OpenView Storage Data Protector Software Release Notes* for information on:

- General Data Protector and integration-specific limitations

- Supported platforms and integrations

- Supported backup and connectivity topologies

- Supported cluster configurations

For information on supported configurations, see the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

ZDB Database - XPDB

ZDB database is referred to as **XPDB**. It keeps the following information about every split LDEV pair:

- Session ID of the ZDB session that involved the LDEV pair.
- LDEV, volume group, and filesystem configuration.
- CRC check information calculated during the session.
- IR flag (indicating that the pair can be used for instant recovery).

Filesystem and volume management system information is also stored in XPDB.

The information is written to XPDB when a pair is split, and is deleted from XPDB when the pair is resynchronized (prior version of data is then overwritten).

Volume group configuration and CRC check information stored in XPDB is compared to the volume group configuration and CRC check during an instant recovery session. If these items do not match, instant recovery fails.

Objects and their mirror configurations during backup and restore sessions are kept in XPDB for replica set rotation and instant recovery. Only LDEV pairs recorded in XPDB can be used for instant recovery.

XPDB resides on the Cell Manager in:

- *UNIX:* /var/opt/omni/server/db40/xpdb
- *Windows:* *<Data_Protector_home>*\db40\xpdb

Configuring the Integration

Before you start configuration, make sure you met the prerequisites described in "Overview" on page 102. In addition, do the following:

✓ **Solaris:** Run the Sun format utility to label and format mirrored LDEVs (on both the application and backup systems). See the *HP StorageWorks Disk Array XP Operating System Configuration Guide: Sun Solaris* for information.

✓ **BC configurations:** Connect the application and backup systems to the same XP. Primary LDEVs (P-VOLs) must be connected to the application system and have paired disks (S-VOLs) assigned.

When using first-level mirrors, primary LDEVs (P-VOLs) must be connected to the application system and each have 2-3 paired disks (S-VOLs) assigned. Mirrored LDEVs (S-VOLs) must be connected to the backup system.

CA configurations: Connect the application system to the Main Control Unit (MCU), and the backup system to the Remote Control Unit (RCU). ESCON links provide communication links between XP MCU and RCU.

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

Combined CA+BC configurations: Connect the application system to MCU, and the backup system to RCU.

Main LDEVs (P-VOLs) must be paired to remote volumes in the RCU (S-VOLs). S-VOLs also function as BC primary volumes (P-VOLs) and must be paired to local copies (BC S-VOLs).

Windows: Connect only BC S-VOLs to the backup system.

HP-UX: Connect only BC S-VOL to the backup system. If CA S-VOL is connected as well, special care must be taken if /etc/lvmtab is lost in this configuration: use vgscan to recreate the volume groups and vgreduce to delete potentially added pvlinks to the S-VOL. Re-import or re-create the volume group to ensure the configuration is correct.

HP-UX LVM mirroring: Use the physical volume groups mirroring of LDEVs to ensure that each logical volume is mirrored to an LDEV on a different I/O bus. This arrangement is called **PVG-strict mirroring**. Disk hardware must be already configured, so that the mirror copy disk is connected to the system on a different bus (not the bus used for the primary copy).

1. Create the volume group with the LDEV that has S-VOL assigned using vgcreate. LVM mirror primary volume must be the LDEV that has its S-VOL.

2. Extend the volume group with an LDEV that has no S-VOL assigned using vgextend. LVM mirror secondary volume must be the LDEV that has no S-VOL.

For more information on LVM mirroring, see the *HP-UX Managing Systems and Workgroups* manual.

To configure the integration: TBD

- Set XP command devices. See "Command Device Handling" on page 106.

- If needed, set the XP LDEV exclude file. See "XP LDEV Exclude File" on page 107.

Command Device Handling

XP command devices are needed by processes requiring access to XP. The information about command devices is kept in XPDB.

- Whenever a session is started, Data Protector queries XPDB for a list of command devices. If there is none (default behavior when the first session is started), Data Protector generates a list of command devices 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, Data Protector recognizes that the command device is assigned to another system; such command device-hostname combination gets the next available instance number.

Thus, every XP attached to the application and backup systems has a list of command devices and systems having access to these devices (together with an instance number).

- Whenever during a session the application or backup system needs access to XP, it uses the first assigned command device with the instance number from the list. If the command device fails, the next device from the list is used. If all devices fail, the session fails. If successful, a command device is used by a system until the end of the session, and the list of command devices is used for all consecutive sessions.

Below is an example of command device entries in XPDB:

```
Serial#CU:Ldev(LDEV)InstSystem

========================================================

3537100:67(103)301application.system1.com

3537100:67(103)302backup.system.com

3537200:68(104)301application.system2.com

3537300:69(105)301application.system3.com
```

Use the `omnidbxp` command to:

- Assign a command device (identified by XP serial number and LDEV number) to a particular system. Optionally, you can assign an instance number. If the instance number is not specified, Data Protector assigns the lowest unassigned instance number. The entire information is written to XPDB.

- List all command devices in XPDB.

- Update the information about a command device.

- Remove one or all command devices from XPDB.

See the `omnidbxp` man page for command syntax and the examples.

XP LDEV Exclude File

You can reserve certain LDEVs for purposes other than Data Protector backup and restore. A session is aborted if the participating replica set contains an excluded LDEV.

Disabled mirrors are listed in the XP LDEV exclude file on the Cell Manager:

- *UNIX:* `/var/opt/omni/server/db40/xpdb/exclude/XPexclude`

- *Windows:*
  `<Data_Protector_home>\db40\xpdb\exclude\XPexclude`

Mirrors listed in this file must be backup system LDEVs identified by the backup system LDEV#.

Use the omnidbxp command to:

- Set and change the exclude file
- Identify excluded LDEVs
- Reset the exclude file
- Delete the content of the exclude file

See the omnidbxp man page for command syntax and the examples of manipulating the exclude file. The file syntax and the example are given below:

Syntax

```
#HP OpenView Storage Data Protector A.06.00
#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

```
#HP OpenView Storage Data Protector A.06.00
#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
```

Automatic Configuration of Backup System

Necessary configuration steps, such as configuring volume groups and filesystems on the backup system, are performed 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 and mounts these filesystems during ZDB-to-tape or ZDB-to-disk+tape sessions. For more information on the backup system mountpoint creation, see the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

Maintaining the Integration TBD

Maintenance tasks include querying the information kept in XPDB, in particular:

- All available backup sessions
- All backup system LDEVs involved in a particular session
- All backup system LDEVs stored in XPDB
- XPDB information about a particular LDEV pair

See the omnidbxp man page for command syntax and the examples.

10 Backup

Overview

This chapter describes configuring a filesystem or disk image ZDB using the Data Protector GUI.

You should be familiar with XP concepts and procedures and basic Data Protector ZDB and instant recovery functionality. See the XP-related documentation and the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

ZDB Types

You can classify ZDB sessions in two ways:

- According to the media to which data is backed up:

    - **ZDB to disk**

        The replica produced is kept on a disk array until reused. This replica becomes part of the replica set and can be used for instant recovery.

        ZDB to disk is performed when the option `Track the replica for instant recovery` is selected in a backup specification, and `To disk` is selected when running/scheduling a backup.

        ZDB to disk is only possible using the BC configuration.

    - **ZDB to tape**

        The replica produced is streamed to backup media, typically tape.

        This replica is deleted after backup if the option Keep the replica after the backup *is not* selected in a backup specification. If this option is selected, the replica remains on a disk array until reused and becomes part of the replica set TBD. However, it cannot be used for instant recovery.

        ZDB to tape is performed when the option `Track the replica for instant recovery` *is not* selected in a backup specification.

    - **ZDB to disk+tape**

        The replica produced is kept on a disk array until reused and is also streamed to backup media. This replica becomes part of the replica set and can be used for instant recovery.

ZDB to disk+tape is performed when the option `Track the replica for instant recovery` is selected in a backup specification, and `To disk+tape` is selected when running/scheduling a backup.

ZDB to disk+tape is only possible using the BC configuration.

• According to the amount of data that is backed up to tape:

Table 10-1   **Backup Types**

Full	Backs up all selected objects, even if there are no changes from the previous backup.
Incremental	Backs up changes from the last protected backup.

NOTE   The replica is created in the same way for both full and incremental ZDB. The difference between the two ZDB types is only in the amount of files that are streamed from the replica to tape.

Limitations   — Incremental ZDB is not supported for ZDB to disk.

— Incremental ZDB is supported only for filesystem backups.

Backup Concepts

XP backup consists of two phases:

1. Application system data gets synchronized to the backup system.

   During this phase, the synchronization is performed on the level of participating volume groups (UNIX) or disks (Windows). Therefore, if multiple filesystems/disk images are configured in the same volume group or on the same disk, 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 objects selected for backup.

2. Synchronized backup system data is backed up to a backup device.

   During this phase, only the objects selected for backup are backed up.

NOTE   With ZDB to disk, the second phase does not occur. Backed up data can only be restored using instant recovery.

This concept enables a restore of selected objects for a split mirror restore and restore from backup media on LAN, but not for instant recovery.

With instant recovery, the links from the application to 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 selected objects to the backup system and resynchronizing the backup system to the application system.

Configuring a Backup Specification

IMPORTANT   Before you begin, consider all limitations regarding the XP integration. For more information, see the *HP OpenView Storage Data Protector Software Release Notes* and the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

1. In the Context List, select `Backup`.

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

   The `Create New Backup` dialog box appears.

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

   Select `Split mirror backup` as `Backup type` and `HP StorageWorks XP` as `Sub type`. For descriptions of options, press **F1**.

   Click `OK`.

3. Under `Client systems`, select the application and backup systems.

   Under `Mirror type`, specify the XP configuration.

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

   For information on options, see "Backup Options" on page 119.

   Click `Next`.

4. *Filesystem backup:* Expand application systems and select the objects to be backed up.

IMPORTANT   To perform instant recovery on UNIX, select all filesystems inside the volume group. Otherwise, instant recovery may fail or require workarounds and there is a risk of data corruption.

Click Next.

Disk image backup: click Next.

5. Select devices. Click Properties to set device concurrency, media pool, and preallocation policy. For information on these options, click Help.

To create additional copies (mirrors) of backup, specify the number of mirrors by clicking Add mirror/Remove mirror. Select separate devices for each mirror backup.

For information on object mirroring, see *HP OpenView Storage Data Protector Administrator's Guide*.

NOTE   Object mirroring and object copying are not supported for ZDB to disk.

Click Next.

6. Under Backup Specification Options, click Advanced and then HP StorageWorks XP to open XP backup options.

Here, you can specify Application options, and modify all other options, except Application system and Backup system. See "Backup Options" on page 119.

In the Filesystem Options group box, click Advanced and specify filesystem options as desired. For information, press **F1**.

Windows only: If you plan to do incremental ZDB, select the Do not use archive attribute filesystem option in the WinFSOptions page.

Figure 10-1   **XP Backup Options**
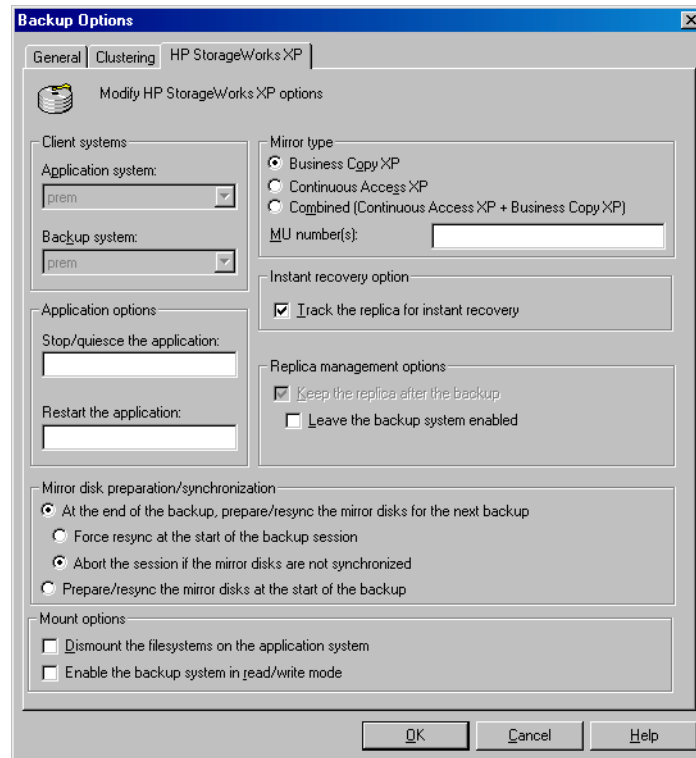
7. Follow the wizard to open the scheduler (for information, press **F1** or see "Running and Scheduling ZDB Sessions" on page A-2), and then the backup summary.

8. *Filesystem backup:* Click Next.

   Disk image backup:

   a. Click Manual add to add disk image objects.

   b. Select Disk image object and click Next.

   c. Select the client and click Next.

   d. Specify General Object Options and Advanced Object Options. For descriptions of these options, press **F1**.

e. In the `Disk Image Object Options` window, specify disk image sections.

UNIX:

Specify a rawdisk section:

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

Specify a raw logical volume section:

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

IMPORTANT   To perform instant recovery, specify all raw logical volumes inside the volume group. Otherwise, instant recovery may fail or require workarounds, and there is a risk of data corruption.

Windows:

Use the following format:

\\.\PHYSICALDRIVE#

where # is the current number of the disk to be backed up.

For information on finding current disk numbers (physical drive numbers), see the online Help index: "disk image backups".

f. Click `Finish` and `Next`.

9. Save your backup specification. For information on starting and scheduling ZDB sessions, see "Running and Scheduling ZDB Sessions" on page A-2.

NOTE   Preview backup is not supported.

Backup Options

The following tables describe XP backup options. See also "XP Integration" on page A-30

Table 10-2   **Client Systems Options**

Data Protector GUI	Function
Application system	System on which the application runs. In cluster environments, specify the virtual server hostname (rather than the physical node hostname).
Backup system	System to which your data will be backed up. In cluster environments, specify the virtual server hostname (rather than the physical node hostname).

Table 10-3   **Mirror Type Options**

Mirror Type	XP mirror configuration: Business copy XP, Continuous Access XP, or Combined (Continuous Access XP + Business Copy XP).
MU number(s)	Enabled if Business copy XP is used.
	Sets the number of replicas in the replica set. Enter an integer number from 0 to 2, or any range/combination of integer numbers from 0 to 2, separated by a comma, for example:
	1
	1-2
	2,0,1
	A sequence does not set the order in which the replicas are used. The algorithm of using replicas is described in the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.
	A range must be specified in ascending order.
	If this option is not specified, MU# 0 is set.

Table 10-4 **Application Options**

Stop/quiesce the application	Create the optional `Stop/quiesce the application` command in /opt/omni/lbin (UNIX) or `<Data_Protector_home>`\bin (Windows) on the application system, and specify the filename in the backup specification.
	This command is executed on the application system before splitting the links. It is mainly used to stop applications not integrated with Data Protector.
	If this command fails, `Restart the application` (see below) is not executed. Therefore, you need to implement a cleanup procedure in `Stop/quiesce the application`. Note that if the `ZDB_ALWAYS_POST_SCRIPT` variable is set to 1, `Restart the application` is always executed if set (default is 0). See "ZDB Agents Omnirc Variables" on page A-18 for details.
Restart the application	Create the optional `Restart the application` command in /opt/omni/lbin (UNIX) or `<Data_Protector_home>`\bin (Windows) on the application system, and specify the filename in the backup specification.
	This command is executed on the application system after split. It is mainly used to restart applications not integrated with Data Protector.

Table 10-5 **Instant Recovery Option**

Track the replica for instant recovery	Enabled if Business copy XP is used.
	Select this option to perform ZDB to disk/ZDB to disk+tape and leave the replica on a disk array for instant recovery. If this option is not set, you cannot perform instant recovery from the replica created or reused in this session.
	If you select this option, do not manually resynchronize the affected mirrors. Otherwise, instant recovery will not be possible.
	Default: selected.

Table 10-6   **Replica Management Options**

Keep the replica after the backup	By default, this option is automatically selected if `Track the replica for instant recovery` is set, and cannot be deselected.
	If this option is selected, participating pairs remain split after backup, enabling you to restore directly from the replica.
	If this option is not selected, participating disks are resynchronized after backup if one or no replica is set by `MU Number(s)`. If more than one replica is set, the disks remain split.
Leave the backup system enabled	Available if `Keep the replica after the backup` is selected.
	By default, Data Protector dismounts filesystems (all platforms), deactives volume groups (HP-UX, Solaris), and removes volume groups from `/etc/lvmtab` (HP-UX) on the backup system after each backup.
	If this option is selected, filesystems remain mounted (all platforms), volume/disk groups remain activated (HP-UX, Solaris), and volume groups are not removed from `/etc/lvmtab` (HP-UX) after backup.
	Thus, you can use the backup system for data warehouse activities, *but not for instant recovery*.
	Default: not selected.

NOTE   By selecting `Leave the backup system enabled` you cannot use the replica for instant recovery unless you also select `Track the replica for instant recovery`.

Table 10-7   **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 (default), the next replica is prepared according to replica set rotation (resynchronized with P-VOLs) for the next backup at the end of the current backup.
	If this option is not selected, next two options are disabled.

Table 10-7   **Mirror Disk Preparation/Synchronization Options**
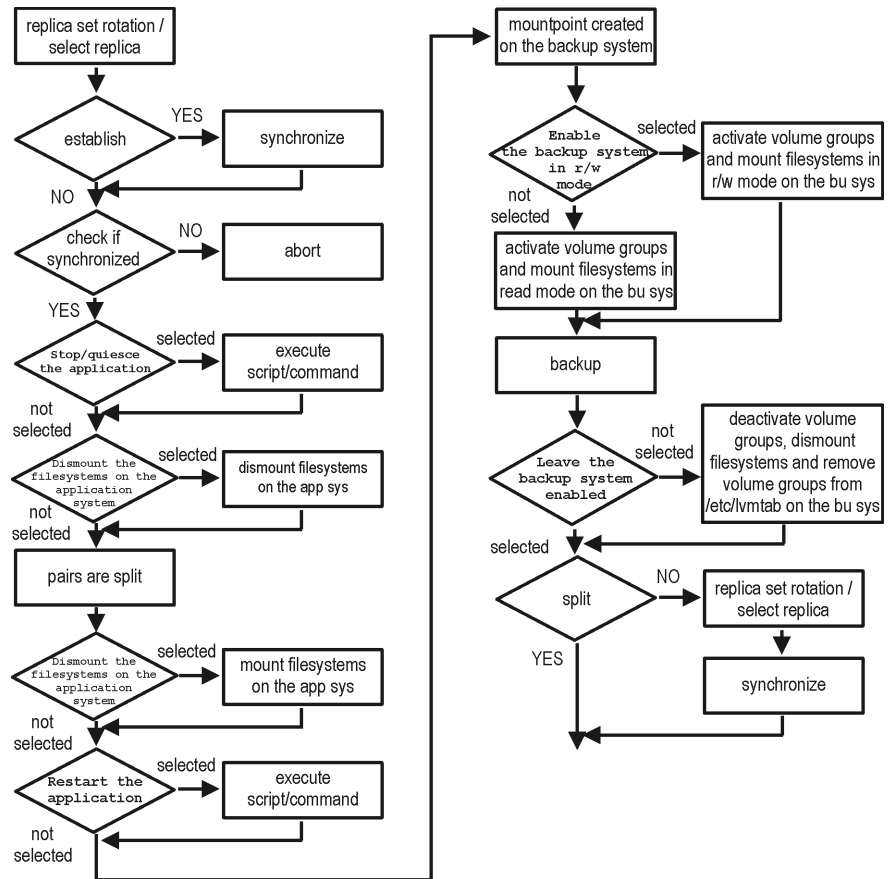
Force resync at the start of the backup session	A resync is initiated before backup. Default: not selected.
Abort the session if the mirror disks are not synchronized	If mirror disks are not synchronized when the backup starts, the session is aborted. Defaul: selected.
Prepare/resync the mirror disks at the start of the backup	Mirror disks in the replica selected for the current session are resynchronized with P-VOLs at the beginning of the session. Default: not selected.

Table 10-8   **Mount Options**

Dismount the filesystems on the application system	Select this option to dismount filesystems on the application system before the split and remount them afterwards. 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. If integrated applications (for example, Oracle) run on the filesystem, they control the I/O to the disk, so, it is not necessary to dismount the filesystem before the split. Default: not selected.
Enable the backup system in read/write mode	HP-UX and Solaris only (on Windows, filesystems are always mounted in read/write mode). Select this option to have read/write access to 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. For other tasks, read/write mode may be needed. Note that if this option is selected, the replica used for instant recovery includes all modifications made when the backup system was online. Default: not selected.

The chart and table below provide detailed backup flow according to the backup options selected.

Figure 10-2   **Filesystem Split Mirror Backup Flow**

The "establish" and "split" checks depend on the following XP backup options:

Table 10-9

At the end of the backup, prepare/resync the mirror disks for the next backup	split = NO
Force resync at the start of the backup session	establish = YES
Abort the session if the mirror disks are not synchronized	establish = NO

Table 10-9

`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   Simultaneous selection of options in the first and the last row of Table 10-9 is conflicting. In such a situation, the "split" check is set to YES.

---

11 Restore

Overview

This chapter describes configuring and running a filesystem or disk image restore of the data backed up using the XP integration. The sections describe restore procedures using the Data Protector GUI and CLI.

The data backed up in a ZDB session can be stored on a disk array (ZDB to disk, ZDB to disk+tape), or on backup media (ZDB to tape, ZDB to disk+tape).

Available restore types are:

- Restore from backup media on LAN (standard restore). See "Standard Restore" on page 127.

- Split mirror restore. See "Split Mirror Restore" on page 129.

- Instant recovery. See "Instant Recovery" on page 135.

Table 11-1   **Restore Types**

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

Standard Restore

Data is restored from the backup media to the application system through a LAN. Only selected backed up objects are restored. For more information on this restore type, see the *HP OpenView Storage Data Protector Administrator's Guide*.

TIP   You can improve the data transfer rate by connecting a backup device to the application system. See the *HP OpenView Storage Data Protector Administrator's Guide* TBD for more information on configuring backup devices. See "Restoring Under Another Device" in the same guide for more information on restoring using another device.

The procedure below is a general description of restoring the objects backed up in a ZDB session.

1. In the Context List, select Restore.

2. Select the objects for restore and click them to display their properties.

   In the Scoping Pane, select the application system as Target client under the Destination tab.

   For information on restore options, press **F1**.

3. Click Restore. The Start Restore Session dialog box appears.

4. Click Next to specify the report level and network load. Click Next.

5. Select StorageWorks XP restore (relevant only if the EMC Symmetrix component is installed on the application system):

Figure 11-1 **XP Restore**
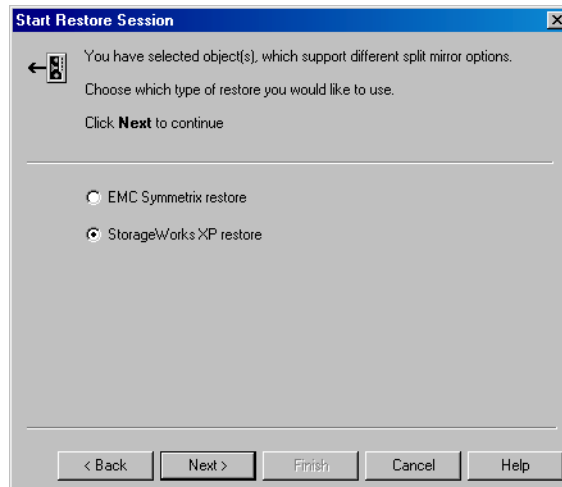
Click Next.

6. In the Start Restore Session window, select Disabled as Mirror mode. This sets a direct restore to the application system.

7. Click Finish to start restore.

Split Mirror Restore

Considerations
- Using split mirror restore, you can restore filesystems and disk images backed up in BC and BC1 configurations. Database applications restore is not supported.

- You must start split mirror restore only after the preceding session using the same disk on the application system finishes synchronization.

Data is restored from backup media on LAN to the mirror LDEVs (S-VOLs), and then moved to the original LDEVs (P-VOLs). The procedure consists of the following automated steps:

1. Applying replica set rotation (if a replica set is defined) to the specified replica set to select the replica for restore. See the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide* for more information.

2. Preparing the backup and application systems.

3. Restoring data from backup media on LAN to the backup system and synchronizing this data to the application system.

Split Mirror Restore Procedure

1. In the `Context List`, select `Restore`.

2. Select the objects for restore and click them to display their properties.

NOTE
Select the application system as `Target client` under the `Destination` tab. If the backup system is selected, standard restore to the backup system is performed.

3. Click `Restore`. The `Start Restore Session` dialog box appears.

4. Click `Next` to specify the report level and network load.

5. Select `StorageWorks XP restore` (relevant only if the EMC Symmetrix component is installed on the *target* client). Click `Next`.

6. Specify split mirror restore options. See "Split Mirror Restore Options" on page 130 for information.

Figure 11-2   Split Mirror Restore Options
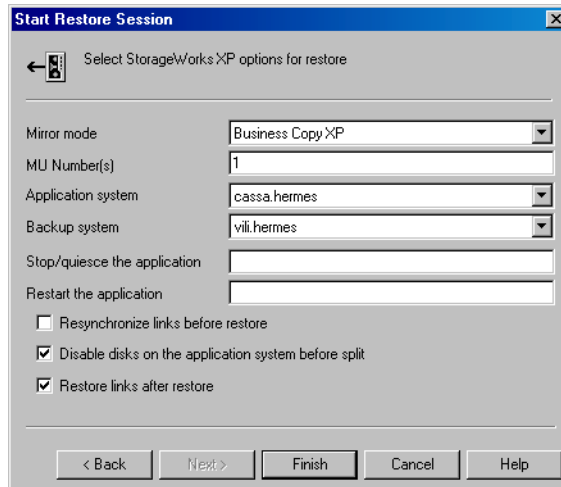
7. Click Finish to start restore.

NOTE   If LVM Mirroring is used, a warning appears during the restore, since the volume group LDEVs in the physical volume group on the application system do not have BC pairs assigned. This warning should be ignored.

See the *HP OpenView Storage Data Protector Administrator's Guide* for information on general restore process.

Split Mirror Restore Options

The following table explains split mirror restore options.

Table 11-2   Split Mirror Restore Options

Data Protector GUI	Function
Mirror mode	XP configuration. Only BC configuration is supported.

Table 11-2   **Split Mirror Restore Options**

Data Protector GUI	Function
MU Number(s)	Sets the number of replicas in the replica set. If this option is not specified, MU# 0 is set.
Application system	System to which your data will be restored. In cluster environments, specify the virtual server hostname (rather than the physical node hostname).
Backup system	System to which your data will be backed up. In cluster environments, specify the virtual server hostname (rather than the physical node hostname).
Stop/quiesce the application	Create the optional Stop/quiesce the application command in /opt/omni/lbin (HP-UX, Solaris) or *<Data_Protector_home>*\bin (Windows) on the application system. This command is executed on the application system before splitting the links. It is mainly used to stop applications or dismount filesystems that are not to be restored in the active session and are mounted to the same volume or disk group as the filesystem to be restored. If this command fails, Restart the application (see below) is not executed. Therefore, you need to implement a cleanup procedure in Stop/quiesce the application. Note that if the ZDB_ALWAYS_POST_SCRIPT variable is set to 1, Restart the application is always executed if set (default is 0). See "ZDB Agents Omnirc Variables" on page A-18 for details.
Restart the application	Create the optional Restart the application command in /opt/omni/lbin (HP-UX, Solaris) or in *<Data_Protector_home>*\bin (Windows) on the application system. The command is executed on the application system immediately after the split. It is mainly used to restart applications or mount filesystems.

Table 11-2   **Split Mirror Restore Options**

Data Protector GUI	Function
`Resynchronize links before restore`	Synchronizes pairs (moves data to backup disks) thus preparing the disks for restore. If the pairs were split before restore, and only some files need to be restored, use this option to update the backup system. This ensures that correct data is resynchronized to the application system. Default: not selected.
`Disable disks on the application system before split`	Disables disks on the application system by dismounting filesystems (HP-UX, Solaris) and deactivating volume groups (HP-UX) before the split. The disks are enabled after restoring the links. If other filesystems exist in the volume or disk group, use `Stop/quiesce the application` and `Restart the application` commands to dismount these filesystems. Always select this option if you want to move data from the backup to the application system (to incrementally restore links). Application system disks must be disabled to provide data integrity after the links are restored.
`Restore links after restore`	Incrementally restores links for LDEVs, successfully restored from the backup media on LAN. The links for LDEVs not successfully restored are incrementally re-established.

The chart below provides detailed split mirror restore flow depending on the options selected.

Figure 11-3 **Filesystem Split Mirror Restore Flow**
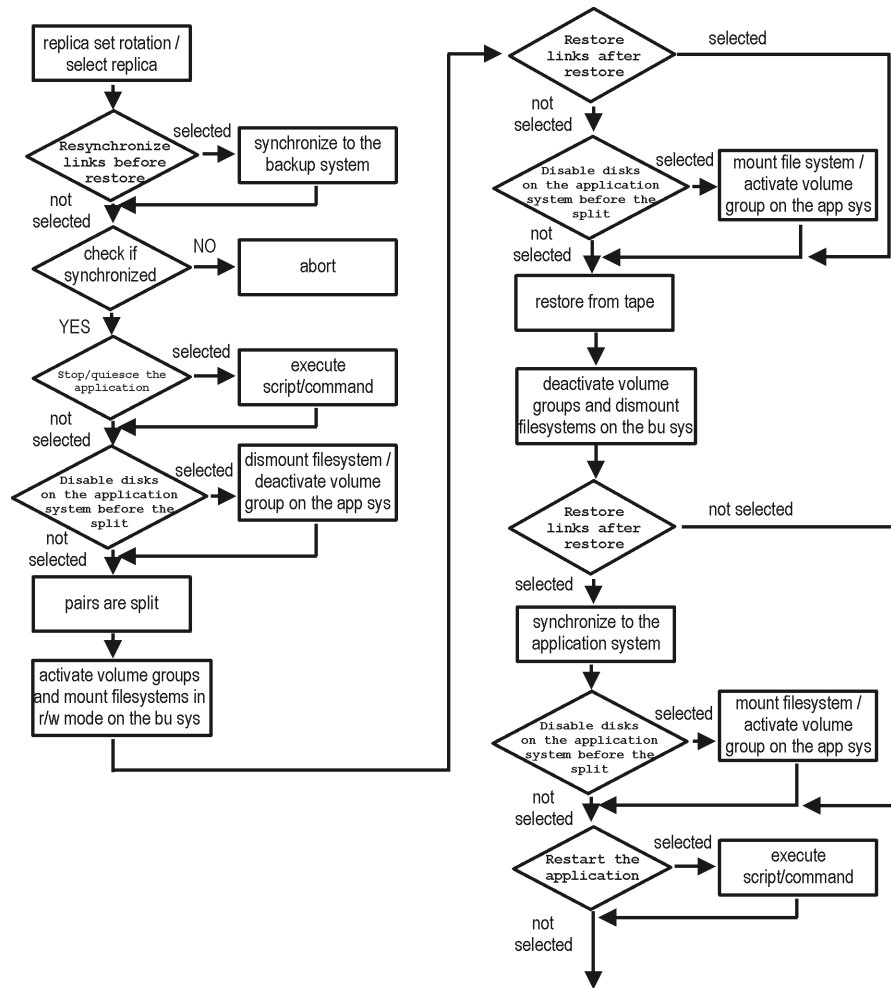
Split Mirror Restore in a Cluster

Split mirror restore with a filesystem running in MC/ServiceGuard or
Microsoft Cluster Server on the application system requires some
additional steps.

MC/ServiceGuard Procedure

1. Stop the filesystem cluster package:

   cmhaltpkg *<app_pkg_name>*

   This stops filesystem services and dismounts the mirrored volume group filesystem.

2. Deactivate the mirrored volume group from cluster mode and activate it in 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. Start split mirror restore (see "Split Mirror Restore Procedure" on page 129).

IMPORTANT   When specifying the application system, specify the hostname of the application system *node* on which the mirrored volume group was activated in the normal mode (step 2 of this procedure).

---

5. After restore, dismount the mirrored volume group filesystem:

   umount /*<mountpoint>*

6. Deactivate the mirrored volume group in normal mode and activate it in 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

Instant recovery restores data directly from a replica to source volumes, without involving a backup device. All data in the replica is restored (regardless of selections during backup). For instant recovery concepts, see the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

You can perform instant recovery using Data Protector GUI (see "Instant Recovery Using the GUI" on page 136) or CLI (see "Instant Recovery Using the CLI" on page 138).

Considerations

- Only three first-level mirrors can be used for instant recovery. Six additional (cascading) copies are not supported.

- Instant recovery restores the data backed up in BC and BC1 configurations.

- When instant recovery starts, the application system needs to be disabled. This includes dismounting filesystems and deactivating volume groups (HP-UX, Solaris). Before this is done, filesystems' and volume groups' status is checked, and only mounted filesystems and activated volume groups are dismounted and deactivated. At the end of the session, only previously dismounted filesystems are mounted, and only previously deactivated volume groups are activated.

- You cannot start instant recovery using the same disk on the application system at the same time. A session can be started only after the preceding session using the same disk on the application system finishes synchronization.

- In a ZDB to disk+tape session, the streaming of files from the replica to tape impacts the files access time, change time (UNIX), and archive (Windows) attributes, depending on how you select the Do not preserve access time attributes and Do not use archive attribute filesystem options. For more information on the options, see online Help.

IMPORTANT After instant recovery, restored filesystems are mounted to the same mount points/drive letters as they were at the backup time. If these mount points/drive letters have other filesystems mounted, these filesystems are automatically dismounted before instant recovery, and the restored filesystems are mounted afterwards.

For more information about XP instant recovery considerations and limitations, see the *HP OpenView Storage Data Protector Software Release Notes* and the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

IMPORTANT Instant recovery does not recover databases/applications. It only synchronizes the application system LDEVs from their mirrors on the backup system. To recover a database/application, you need to perform additional steps.

Prior to instant recovery, Data Protector checks:

- Volume group configuration
- CRC of the mirror copy

These checks assure that data in the replica is left intact since backup. If any of these checks fail, the session fails.

Once the replica is restored, it can be left unchanged or resynchronized, depending on selected options. See "Instant Recovery Options" on page 139 for information.

Instant Recovery Procedure

Prerequisites • When performing a disk image instant recovery, manually dismount the disks before instant recovery, and re-mount them afterwards.
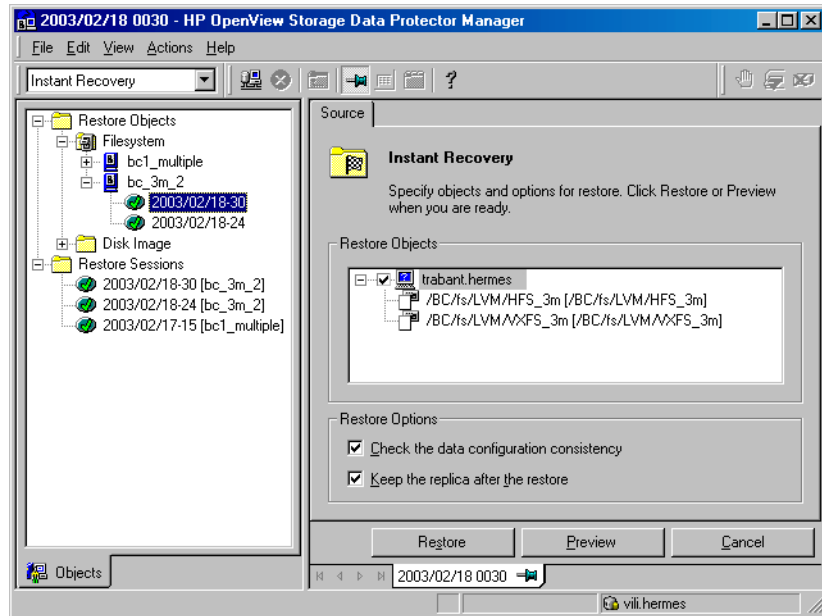
Instant Recovery Using the GUI

1. In the Context List, select Instant Recovery.

2. Select the backup session (replica) from which you want to perform the recovery. This can be done by selecting:

- Backup session ID and name (in the Scoping Pane, expand `Restore Sessions` and select the session from a list of ZDB-to-disk/ZDB-to-disk+tape sessions)

- Backup type (filesystem, Oracle, SAP R/3,...) and backup session name and ID:

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

     Backed up object types (Filesystem, Disk Image, SAP R/3, Microsoft SQL Server, ...) are displayed.

  b. Expand the object type you want to restore.

     Available backup specification used in ZDB-to-disk/ZDB-to-disk+tape sessions for the selected object type are displayed.

  c. Expand the backup specification containing the required objects. Available sessions are displayed:

Figure 11-4 **Selecting a Session**

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

The application system and its mount points/drive letters backed up during the selected session are displayed.

4. Select the application system and specify instant recovery options (see "Instant Recovery Options" on page 139).

5. Click Restore to start instant recovery, or Preview to preview it (filesystem backup only).

6. Select Start Restore Session to start instant recovery, or Start Preview Session to start the preview. Click OK.

IMPORTANT   You cannot use the CLI to perform instant recovery from ZDB to disk+tape after exporting or overwriting the media used in the session. Use the CLI instead. Note that backup media must not be exported or overwritten even after an object copy session.

Instant Recovery Using the CLI

1. List all available ZDB-to-disk or ZDB-to-disk+tape sessions, identified by the session ID:

   omnidbxp -ir -session -list

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

2. Execute:

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

   Where:

   <application_system_name> Application system hostname.

   <SessionID>   Backup session ID (step 1 of this procedure).

   For *<INSTANT RECOVERY OPTIONS>*, see Table 11-3 on page 139.

See omnidbxp and omnir man pages for details.

Instant Recovery Options

Table 11-3   **Instant Recovery Options**

Data Protector GUI	Function
`Check the data configuration consistency/-check_config`	The current configuration of the participating volume groups is compared with the volume group configuration during the backup session kept in XPDB. If the configuration changed since then, the restore fails.
	When performing instant recovery in an MC/ServiceGuard cluster to another node (not the one that was backed up), leave this option selected. The current volume group configuration on the node to which instant recovery is being performed is different from the configuration kept in XPDB. With this option selected, XPDB volume group configuration is replaced by the current configuration on the node to which instant recovery is being performed, and the session is not aborted.
	CRC check information for the selected LDEV pairs stored in XPDB is compared to the current CRC check information. If these items do not match, the session fails.
	Default: selected.
`Keep the replica after the restore/-keep_version`	Participating LDEV pairs are split after restore.
	Default: selected.

Instant Recovery and LVM Mirroring

1. Reduce the logical volume containing LVM mirroring disks (using the `lvreduce -m 0` command) by excluding the LVM mirror disk (the one having no S-VOL assigned) from the logical volume.

   This assures that restored data is not overwritten by the data on the LVM mirror disk.

2. Start instant recovery.

3. Extend the logical volume containing LVM mirroring disks (using the `lvextend -m` command) with the LVM mirror disk that was previously excluded from the logical volume.

Instant Recovery in a Cluster

Instant recovery with an application/filesystem running in an MC/ServiceGuard or Microsoft Cluster Server on the application system requires some additional steps. See "Instant Recovery in a Cluster" on page A-15 for instructions.

12 Troubleshooting

Before You Begin

This chapter contains a list of general checks and verifications and a list of problems you may encounter when using the XP integration. For general Data Protector troubleshooting information, see the *HP OpenView Storage Data Protector Troubleshooting Guide*.

✓ Ensure that the latest official Data Protector patches are installed. See the online Help index: "patches" on how to verify this.

✓ See the *HP OpenView Storage Data Protector Software Release Notes* for general Data Protector and integration-specific limitations, as well as known problem and workarounds.

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

Checks and Verifications

✓ On the application and backup systems, examine system errors reported in:

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

Windows: `<Data_Protector_home>\log\debug.log`

✓ Ensure that RAID Manager/LIB XP is correctly installed on both the application and backup systems.

HP-UX, Solaris: The `/usr/omni/bin/ssea` command should return the following output:

```
HP OpenView Storage Data Protector A.06.00: SSEA, internal build
<build_ID>, built on <date>
```

In addition:

— ***HP-UX:*** `ls` `<RMLIB_home>`/libsvrrm.sl should return:

  `<RMLIB_home>`/libsvrrm.sl

— ***Solaris:*** `ls` `<RMLIB_home>`/libsvrrm.so should return:

  `<RMLIB_home>`/libsvrrm.so

Windows: Check if the libsvrrm.dll file exists in the `<RMLIB_home>` directory.

Backup Problems

Problem   **You cannot select StorageWorks mode in the Data Protector user interface when creating a backup specification**

Action   Check that the `HP StorageWorks XP Agent` integration module is installed on the application and backup systems. To do that, open the `cell_info` file located in:

Windows Cell Manager:
`<Data_Protector_Home>\Config\server\cell\cell_info`

UNIX Cell Manager: `/etc/opt/omni/server/cell/cell_info`

File contents should look similar to:

```
-host "hpsap001.bbn.hp.com" -os "hp s800 hp-ux-11.00" -cc A.06.00
-da A.06.00 -ssea A.06.00

-host "hpsap002.bbn.hp.com" -os "hp s800 hp-ux-11.00" -cc A.06.00
-da A.06.00 -ma A.06.00 -ssea A.06.00
```

Problem   **On the application system, dismounting a filesystem fails**

Action   In the `Stop/quiesce the application` script, stop all processes using the filesystem.

Problem   **On the backup system, mounting a filesystem fails**

Action   Check that the mountpoint directory exists on the backup system.

Problem   **Pair synchronization fails (the split fails)**

To successfully split the pair, XP Agent first checks its status. Pairs can only be split (in PSUS/SSUS status) after they are synchronized (in PAIR status). XP Agent checks the status of links after every 2 seconds and retries 10 times.

Action   Increase the time frame for synchronization by setting `SSEA_SYNC_RETRY` and `SSEA_SYNC_SLEEP_TIME` variables.

See "ZDB Agents Omnirc Variables" on page A-18 for more information.

Problem **P-VOL has no paired S-VOL**

Action Check the XP configuration as follows:

BC: All P-VOLs on the application system must have associated BC S-VOLs on the backup system.

CA: All P-VOLs on the application system must have associated CA S-VOLs on the backup system.

CA+BC: 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.

Problem **Invalid pair state of LDEVs**

Action Check the link state. If the link is split, use the `Prepare/resync the mirror disks at the start of the backup` option.

Configure and start RAID Manager XP instances manually. You can get a list of LDEVs from the backup session report.

Problem **Missing details for a specific LDEV/MU# are reported:**

[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 ensure that LDEV/MU# stated in the output is not set in the XP LDEV exclude file.

Restart the session.

Split Mirror Restore Problems

Problem **Session fails with the following 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 Ensure that the filesystem/volume group is not in use (you are positioned in the filesystem mountpoint directory), and then restart the session.

Problem **LDEV pair is in "STAT_COPY" state when split mirror restore starts, and the session fails with:**

[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 session.

Instant Recovery Problems

Problem   **LDEV pair is in "STAT_COPY" state when split mirror restore starts, and the session fails with:**

```
[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 session.

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

```
[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!
```

The problem occurs if the option Check data configuration consistency is selected and may be caused by:

• Failover to a secondary node

• Changed volume group configuration on the application system

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

- If `Check the data configuration consistency` is selected, the
  following happens:

  - ✓ Volume group configuration for the selected LDEV pairs stored in
    XPDB is compared to the current volume group configuration. If
    these items do not match, the session is aborted.

  - ✓ CRC check information for selected LDEV pairs stored in the
    XPDB is compared to the current CRC check information. If these
    items 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.

IV   EMC Symmetrix

13 Configuration

Overview

This chapter describes configuring the Data Protector EMC Symmetrix (EMC) integration.

Prerequisites

- Install:

  EMC components:

  — EMC Solution Enabler

  — EMC Symmetrix TimeFinder or EMC Symmetrix Remote Data Facility (SRDF) microcode and license. TBD

  Data Protector components:

  — A license for using the EMC integration.

  — EMC Symmetrix Agent.

  For installation instructions, see the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

- You should be familiar with:

  — EMC Command-Line Interface

  — Logical Volume Manager concepts

- Make sure the same operating system (and its version) is installed on the application and backup systems.

- Connect EMC to the application and backup systems.

See the *HP OpenView Storage Data Protector Software Release Notes* for information on:

- General Data Protector and integration-specific limitations

- Supported platforms and integrations

- Supported backup and connectivity topologies

- Supported cluster configurations

For information on supported configurations, see the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

EMC Symmetrix Database File and Data Protector EMC Log File

EMC Symmetrix Database File

EMC Symmetrix database file contains the physical configuration information of SCSI parameters that define your storage complex. It is located in:

HP-UX: /var/symapi/db/symapi_db.bin

Windows: *<symapi_home>*\db\symapi_db.bin

Data Protector EMC Log File

EMC log file keeps information about objects, devices, and device groups. It is located in:

HP-UX: /var/opt/omni/tmp/emc

Windows: *<Data_Protector_home>*\Config\client\tmp\emc

on the application and backup systems. Log files are named as R1_*<session_name>*.log or R2_*<session_name>*.log, where *<session_name>* is composed of the sessionID, the forward slashes "/" replaced with dashes "-." For example:

R1_1999-09-13-3.log

R2_1999-09-13-3.log

The log contains:

- Resolved EMC configuration (mapping to EMC devices).
- Created and deleted device groups, and the devices added to device groups.
- Operations on device groups (splitting links, incremental establish, incremental restore, ...).
- Status of backup and restore objects.

Check both log files if you encounter any problems. The logs can also be useful if you leave the links split after backup/restore.

Configuring the Integration

Before you start configuration, make sure you met the prerequisites described in "Overview" on page 152. In addition, do the following:

Symmetrix Remote Data Facility (SRDF) configurations: Connect the application system to Application (R1) Symmetrix, and the backup system - to Backup (R2) Symmetrix.

Main Source (R1) Devices must be connected to the application system and have paired disks assigned. Paired Target (R2) Devices in the remote disk array must be connected to the backup system.

TimeFinder configurations: Connect the application and backup systems to the same disk array.

Standard Devices must be connected to the application system and have paired disks assigned. BCV Devices must be connected to the backup system.

Combined SRDF+TimeFinder configurations: Connect the application system to Application (R1) Symmetrix, and the backup system - to Backup (R2) Symmetrix.

Main Source (R1) Devices must be paired to Target (R2) Devices in Backup (R2) Symmetrix. Backup (R2) Symmetrix Target (R2) Devices also function as TimeFinder Standard Devices. They must be paired to BCV (R2) Devices.

It is recommended that only TimeFinder BCV (R2) Devices be connected to the backup system. If SRDF Target (R2) Devices are connected as well, /etc/lvmtab may get lost in this configuration. To ensure the configuration is correct, re-create volume groups using vgscan, and delete potentially added pvlinks to SRDF Target (R2) Devices using vgreduce.

To configure the integration: TBD

• Create the Data Protector EMC database file. See "Creating Data Protector EMC Database File" on page 155.

• If needed, rebuild the EMC Symmetrix database file. See "Rebuilding EMC Symmetrix Database File" on page 155.

Creating Data Protector EMC Database File

Data Protector EMC database file, used to store configuration information, is the same as the EMC Symmetrix database file. Create this file:

- Prior to starting Data Protector backups

- Each time your disk configuration changes

Alternately, you can set the `Run discovery of Symmetrix environment` option in the backup specification. However, this operation may be time-consuming because it checks disk configuration through low-level SCSI commands.

To create the Data Protector EMC database file, run:

HP-UX: `/opt/omni/lbin/syma -init`

Windows: `<Data_Protector_home>\bin\syma -init`

This command creates the `/var/opt/omni/client/emc/symm.bin` (HP-UX) or `<Data_Protector_home>\Config\Client\EMC\symm.bin` (Windows) Data Protector EMC database file on both the application and backup systems.

Rebuilding EMC Symmetrix Database File

Rebuild the EMC Symmetrix database file with the current information about physical devices connected through SCSI buses to your system if:

- Your configuration changes

- You run the first command-line session

To scan the hardware and rebuild the database, execute:

`symcfg discover`

This command scans all SCSI buses on the system (not only those connected to EMC arrays).

To display the contents of the EMC Symmetrix database file, run:

- `syminq -sym` (displays all EMC devices).

- `symbcv list dev` (lists all BCV devices configured on EMC).

- `symrdf list` (lists all RDF disk devices known to the system).

See "EMC - Obtaining Disk Configuration Data" on page A-33 for more information.

Automatic Configuration of Backup System

The necessary configuration steps, such as configuring volume groups and filesystems on the backup system, are performed 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 and mounts these filesystems during ZDB sessions. For more information on the mountpoint creation, see the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

14 Backup

Overview

This chapter describes configuring a filesystem or disk image ZDB using the Data Protector GUI.

You should be familiar with the EMC concepts and procedures and basic Data Protector ZDB functionality. See the EMC-related documentation and the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

ZDB Types

You can classify ZDB sessions in two ways:

- According to the media to which data is backed up:

  With EMC, only ZDB to tape is supported. With ZDB to tape, mirrors are created, and data from the replica is moved to backup media.

  If the option Re-establish links after backup is not selected, the replica remains on a disk array until reused in the next backup session using the same EMC device pairs.

  If the option Re-establish links after backup is selected, the replica is synchronized with the original after backup.

  See the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide* for more information on ZDB-to-tape process.

- According to the amount of data that is backed up to tape:

Table 14-1   **Backup Types**

Full	Backs up all selected objects, even if there are no changes from the previous backup.
Incremental	Backs up changes from the last protected backup.

NOTE   The replica is created in the same way for both full and incremental ZDB. The difference between the two ZDB types is only in the amount of files that are streamed from the replica to tape.

Limitations — Incremental ZDB is supported only for filesystem backups.

Backup Concepts

EMC backup consists of two phases:

1. Application system data gets synchronized to the backup system.

   During this phase, the synchronization is performed on the level of participating volume groups (HP-UX) or disks (Windows). Therefore, if multiple filesystems/disk images are configured in the same volume group or on the same disk, 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 objects selected for backup.

2. Synchronized backup system data is backed up to a backup device.

   During this phase, only the objects selected for backup are backed up.

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.

Configuring a Backup Specification

IMPORTANT   Before you begin, consider all limitations regarding the EMC
integration. For more information, see the *HP OpenView Storage Data
Protector Software Release Notes* and the *HP OpenView Storage Data
Protector Zero Downtime Backup Concepts Guide*.

1. In the Context List, select Backup.

2. In the Scoping Pane, expand Backup and Backup Specifications.
    Right-click Filesystem (for a filesystem backup TBD), and click Add
    Backup.

    The Create New Backup dialog box appears.

    Select Split mirror backup as Backup type and EMC Symmetrix
    as Sub type. See online Help for options' descriptions. Click OK.

3. Select the application and backup systems. Also, specify the desired
    EMC configuration - TimeFinder, SRDF, or Combined (SRDF +
    TimeFinder).

    See "Backup Options" on page 164 for information on options.

IMPORTANT   In EMC GeoSpan for Microsoft Cluster Service environments, 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. *Filesystem backup:* Expand application systems and select the
    objects to be backed up.

    Click Next.

    Disk image backup: Click Next.

5. Select devices. Click `Properties` to set the device concurrency, media pool, and preallocation policy. For more information on these options, click `Help`.

   To create additional copies (mirrors) of backup, specify the number of mirrors by clicking `Add mirror`/`Remove mirror`. Select separate devices for each mirror backup.

   For information on object mirroring, see *HP OpenView Storage Data Protector Administrator's Guide*.

   Click `Next`.

6. Under `Backup Specification Options`, click the `Advanced` and then the `EMC Symmetirx` tab to open EMC backup options.

   Here, you can modify all options, except `Application system` and `Backup system`, as shown in Figure 14-1 on page 162. See also "Backup Options" on page 164.

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

Figure 14-1 **Backup Options**
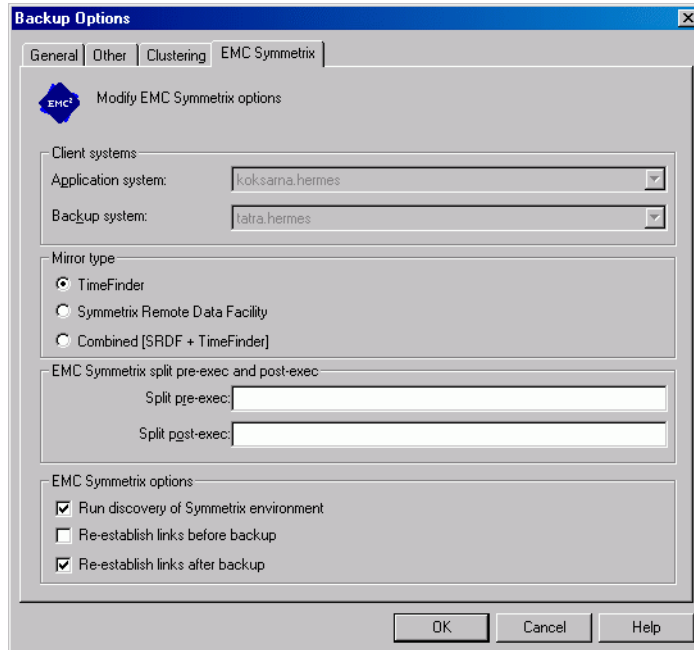
7. Following the wizard, open the scheduler (for information, press **F1** or see "Running and Scheduling ZDB Sessions" on page A-2), and then the backup summary.

8. *Filesystem backup:* click Next.

   Disk image backup:

   a. Click Manual add to add disk image objects.

   b. Select Disk image object and click Next.

   c. Select the client and click Next.

   d. Specify General Object Options and Advanced Object Options. For information on these options, press **F1**.

   e. In the Disk Image Object Options window, specify disk image sections.

   HP-UX:

   Specify a rawdisk section:

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

Specify a raw logical volume section:

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

Windows:

Use the following format:

\\.\PHYSICALDRIVE<#>

Where *<#>* is the current number of the disk to be backed up.

For information on finding current disk numbers (physical drive numbers), see the online Help index: "disk image backups".

  f. Click Finish and Next.

9. Save your backup specification. For information on starting and scheduling backup sessions, see "Running and Scheduling ZDB Sessions" on page A-2.

NOTE   Preview backup is not supported.

Backup Options

The following tables describe EMC backup options. See also "EMC Integration" on page A-32

Table 14-2   **EMC Backup Options**

Data Protector GUI	Function
Application system	System on which the application runs. In cluster environments, specify the virtual server hostname (rather than the physical node hostname).
Backup system	System to which the data will be backed up. In cluster environments, specify the virtual server hostname (rather than the physical node hostname).
	In EMC GeoSpan for MSCS environments, 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	EMC configuration: TimeFinder, Symmetrix Remote Data Facility, or Combined (SRDF + TimeFinder).
	In EMC GeoSpan for MSCS environments, specify the TimeFinder configuration.
Split pre-exec	Create the optional Split pre-exec command in /opt/omni/lbin (HP-UX) or *<Data_Protector_home>*\bin (Windows) on the application system. This command is executed on the application system before the split and is mainly used to stop applications not integrated with Data Protector.
	If Split pre-exec fails, Split post-exec is also not executed. Therefore, you need to implement a cleanup procedure in Split pre-exec.
	If the ZDB_ALWAYS_POST_SCRIPT file variable is set to 1, Split post-exec is always executed if set (default is 0). See "ZDB Agents Omnirc Variables" on page A-18 for more information.
	Backup session is not aborted if the command set by Split pre-exec is not executed.

Table 14-2   **EMC Backup Options**

Data Protector GUI	Function
Split post-exec	Create the optional `Split post-exec` command in `/opt/omni/lbin` (HP-UX) or `<Data_Protector_home>\bin` (Windows) on the application system. This command is executed on the application system after split and is mainly used to restart applications not integrated with Data Protector.
Run discovery of Symmetrix environment	Builds/re-builds the Data Protector EMC database on both the application and backup systems. See "Creating Data Protector EMC Database File" on page 155 for more information. Default: selected.
Re-establish links before backup	Synchronizes disks before backup to maintain data integrity (may be necessary if you disabled `Re-establish links after backup` or used EMC commands that left the links split). Default: not selected.
Re-establish links after backup	Re-establishes links between the application and mirrored devices after backup. If this option is disabled, the links remain split after backup (in this case, you can use the mirrored devices on the backup system). Default: selected.

The chart and table below provide detailed backup flow according to the backup options selected.

Figure 14-2   **Filesystem Split Mirror Backup Flow**
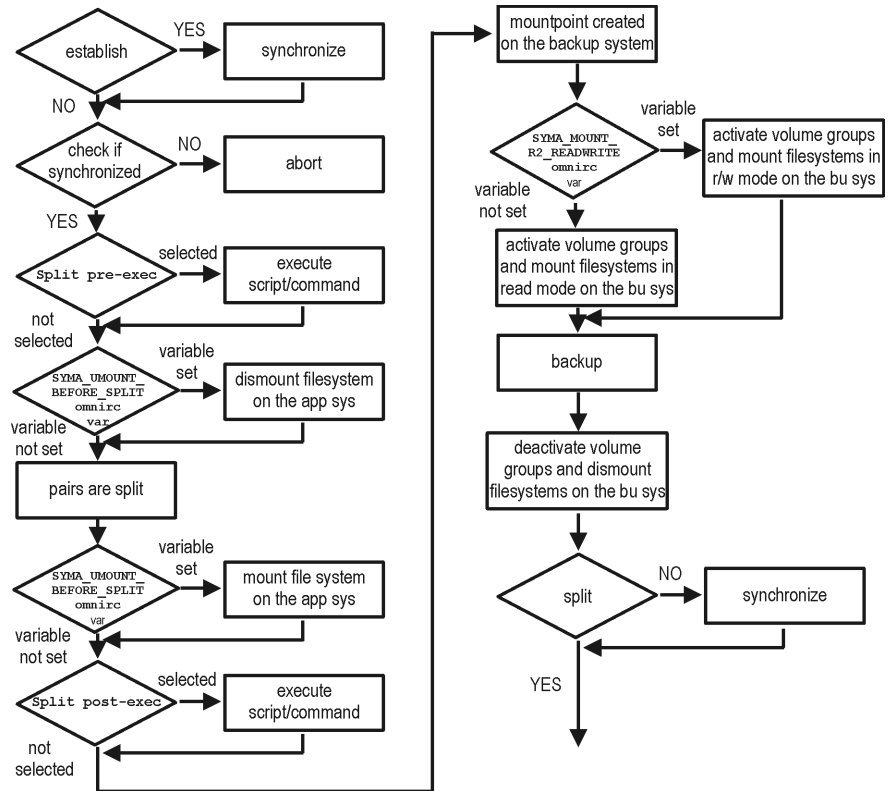
The "establish" and "split" checks depend on the following XP backup options:

Table 14-3

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

Backup Disk Usage

If mirrored devices are not 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 only be restored using EMC device mirroring facilities.

To view this data, enable mirrored devices by activating volume groups (HP-UX) and mounting filesystems. The log file containing information about volume groups and filesystems is located in:

HP-UX: `/var/opt/omni/tmp/emc/R2_<session_name>.log`

Windows:
`<Data_Protector_home>\Config\client\tmp\emc\R2_<session_na me>.log`

where `<session_name>` is composed of the sessionID, forward slashes "/" replaced with dashes "-".

Testing Backed Up Data

To test your backed up data:

1. Restore the data to the backup system or use mirrored devices not re-established after backup. Meanwhile, your applications run uninterrupted on the application system.

2. Test data integrity.

To restore to the backup system, follow the steps described in "Split Mirror Restore Procedure" on page 175 and set EMC split mirror restore options as explained in "EMC Test Restore Options" on page 168.

EMC Test Options

NOTE   For testing, set the SYMA_UMOUNT_BEFORE_SPLIT variable to 0 (default), and SYMA_MOUNT_R2_READWRITE to 1. For details, see "ZDB Agents Omnirc Variables" on page A-18.

Table 14-4   **EMC Test Restore Options**

Data Protector GUI	Function
EMC Symmetrix mode	EMC configuration for test backup: TimeFinder, SRDF, or Combined (SRDF+TimeFinder).
	In EMC GeoSpan for MSCS environments, specify the TimeFinder configuration.
Application system	System on which the application runs. In cluster environments, specify the virtual server hostname (rather than the physical node hostname).
Backup system	System to which your data will be restored. In cluster environments, specify the virtual server hostname (rather than the physical node hostname).
	In EMC GeoSpan for MSCS environments, 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.

Table 14-4   **EMC Test Restore Options**

Data Protector GUI	Function
`Run discovery of the Symmetrix environment`	Disable this option.
`Re-establish links before restore`	Either enable or disable this option.
`Disable disks on application client before split`	Disable this option upon testing your backup (disks on the application system *must not* be disabled). `Restore links after restore` is also disabled, so applications on the application system run uninterrupted. Do not move restored data to the application system for test purposes. This can cause integrity problems.
`Restore links after restore`	Disable this option, leaving the links split. You can then check the integrity of restored data on the backup system.

See "Split Mirror Restore Options" on page 176 for more information about options.

Checking Your Restored Data

If `Restore links after restore` is disabled, mirrored devices contain the restored version of data. To view this data, enable mirrored devices and mount filesystems.

Manually re-establish links using the appropriate EMC CLI command (`symrdf` or `symmir`), or enable the option `Re-establish links before backup`/`Re-establish links before restore` for the next backup/restore.

CAUTION   Do not restore data to the application system for test purposes. Otherwise, you will lose all data written to mirrored devices on the application system.

15 Restore

Overview

This chapter describes configuring and running a filesystem or disk image restore of the data backed up using the EMC integration. The sections describe restore procedures using the Data Protector GUI.

Available restore types are:

- Restore from backup media on LAN (standard restore). See "Standard Restore" on page 173.

- Split mirror restore. See "Split Mirror Restore" on page 175.

Standard Restore

Data is restored from the backup media to the application system through a LAN. Only selected backed up objects are restored. For more information on this restore type, see the *HP OpenView Storage Data Protector Administrator's Guide*

TIP   You can improve the data transfer rate by connecting a backup device to the application system. See the *HP OpenView Storage Data Protector Administrator's Guide* TBD for more information on configuring backup devices. See "Restoring Under Another Device" in the same guide for information on restoring using another device.

The procedure below is a general description of restoring the objects backed up in a ZDB session.

1. In the Context List, select Restore.

2. Select the objects for restore and click them to display their properties.

    In the Scoping Pane, select the application system as Target client under the Destination tab.

    For information on restore options, press **F1**.

3. Click Restore. The Start Restore Session dialog box appears.

4. Click Next to specify the report level and network load. Click Next.

5. In the Start Backup Session window, select Disabled as EMC Symmetrix mode. This sets a restore from backup media on LAN. See Figure 15-1 on page 174.

Figure 15-1   **Restore from Backup Media on LAN**
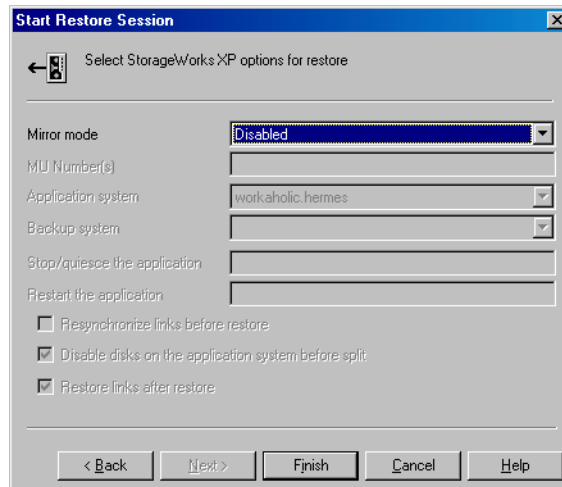
6. Click Finish to start restore.

Split Mirror Restore

Split mirror restore consists of the following automated steps:

1. Preparing the backup and application systems.

2. Restoring data from backup media on LAN to the backup system and synchronizing this data to the application system.

For a description of a split mirror restore process, see the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

Split Mirror Restore Procedure

1. In the Context List, select Restore.

2. Select the objects for restore and click them to display their properties.

NOTE   Select the application system as Target client under the Destination tab. If the backup system is selected, standard restore to the backup system is performed.

3. Click Restore. The Start Restore Session dialog box appears.

4. Click Next to specify the report level and network load. Click Next.

5. Specify EMC split mirror restore options presented in Figure 15-2 on page 176. See "Split Mirror Restore Options" on page 176 for information.

Figure 15-2   **EMC Split Mirror Restore Options**
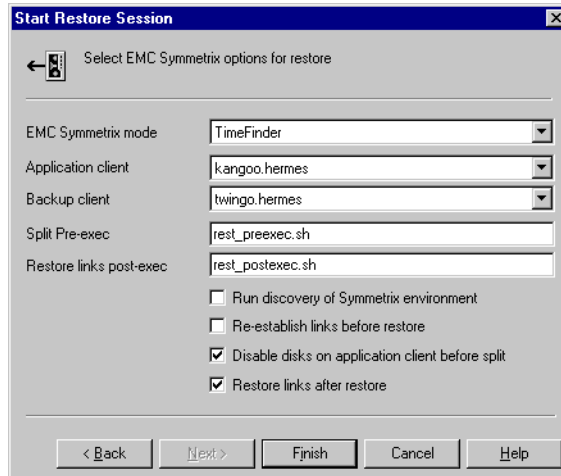
6. Click Finish to start the restore session.

IMPORTANT   You cannot start split mirror backup/restore using the same disk on the application system at the same time. A split mirror session must be started only after the preceding session using the same disk on the application system finishes synchronization; otherwise, the session fails.

Split Mirror Restore Options

The following table explains split mirror restore options.

Table 15-1   **EMC Split Mirror Restore Options**

Data Protector GUI	Function
EMC Symmetrix mode	EMC Symmetrix configuration: TimeFinder, SRDF, or Combined (SRDF + TimeFinder).
Application system	System on which the application runs. In cluster environments, specify the virtual server hostname (rather than the physical node hostname).

Table 15-1   **EMC Split Mirror Restore Options**

Data Protector GUI	Function
Backup system	System to which your data is first restored. In cluster environments, specify the virtual server hostname (rather than the physical node hostname).
Split pre-exec	Specify the Split pre-exec command, executed before the split. Create the command in /opt/omni/lbin (HP-UX) or *<Data_Protector_home>*\bin (Windows) on the application system. This command can be used to stop applications and dismounting filesystems (HP-UX 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 volume groups for de-activation. Restore session is not aborted if the command set by this option is not executed. If Split pre-exec fails, Restore links post-exec (see below) is also not executed. Therefore, you need to implement a cleanup procedure in Restore links post-exec. If the ZDB_ALWAYS_POST_SCRIPT file variable is set to 1, Restore links post-exec is always executed if set (default is 0). See "ZDB Agents Omnirc Variables" on page A-18 for more information.
Restore links post-exec	Specify the Restore links post-exec command, executed after the links are restored. Create the command in /opt/omni/lbin (HP-UX) or *<Data_Protector_home>*\bin (Windows) on the application system. It is used to remount filesystems (HP-UX only) and restart applications. Do not use this command to enable applications if you disabled Re-establish links after restore. Applications using restored disks must not be restarted until the links are manually established.
Run discovery of Symmetrix environment	Builds/re-builds the Data Protector EMC database on both the application and backup systems. See "Creating Data Protector EMC Database File" on page 155 for more information. Default: not selected.

Table 15-1   **EMC Split Mirror Restore Options**

Data Protector GUI	Function
Re-establish links before restore	Synchronizes split disks (moves data to backup disks) thus preparing disks for restore. Default: not selected.
Disable disks on application client before split	Disables disks on the application system by dismounting filesystems and de-activating volume groups (HP-UX) before the split. The disks are enabled after restore. Always select this option when you want to move data from the backup to the application system, that is, to incrementally restore links. Application system disks must be disabled to provide data integrity after restore.
Restore links after restore	Incrementally restores links of devices, successfully restored to the backup system. Links of devices that were not successfully restored are incrementally re-established.

The chart below provides detailed split mirror restore flow according to the options selected.

Figure 15-3 **Split Mirror Restore Flow**
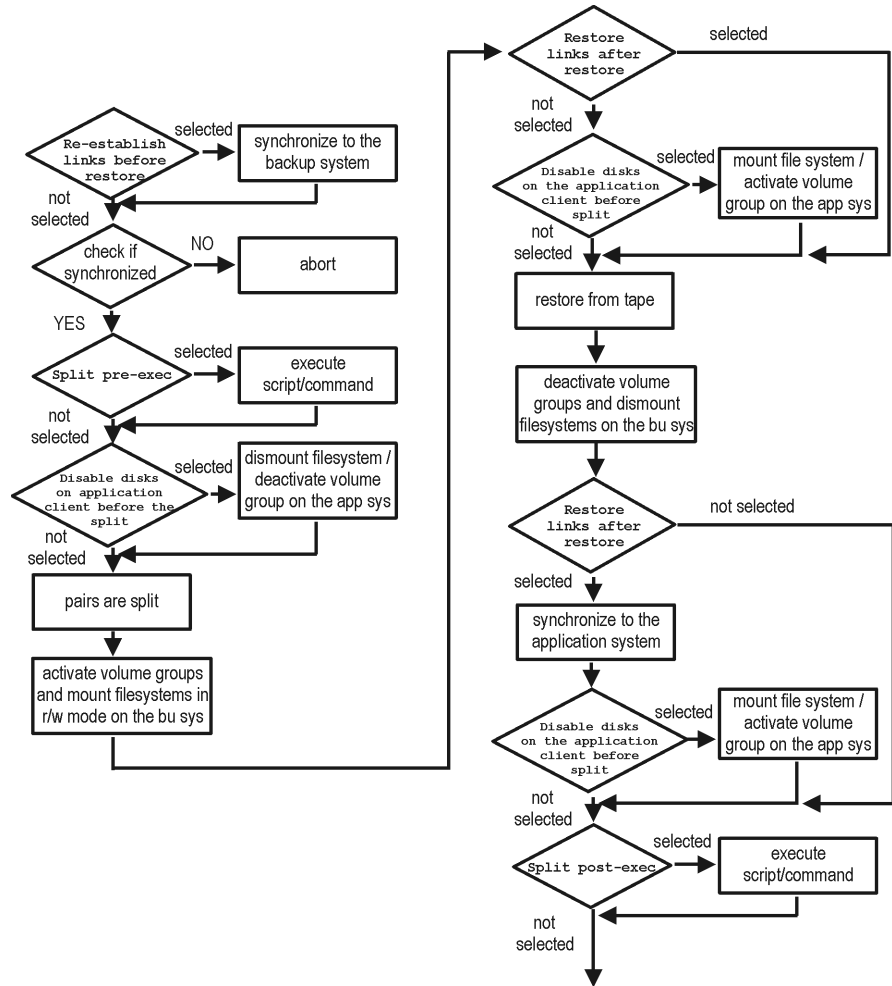
Split Mirror Restore in a Cluster

Split mirror restore with a filesystem running in an MC/ServiceGuard or
Microsoft Cluster Server on the application system requires some
additional steps. See the below sections.

MC/ServiceGuard Procedure

1. Stop the filesystem cluster package:

   cmhaltpkg *<app_pkg_name>*

   This stops filesystem services and dismounts 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. Start split mirror restore (see "Split Mirror Restore Procedure" on page 175).

IMPORTANT　　When specifying the application system, specify the hostname of the application system *node* on which the mirrored volume group was activated in the normal mode (step 2 of this procedure).

---

5. After restore, 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>*

16   Troubleshooting

Before You Begin

This chapter contains a list of general checks and verifications and a list of problems you may encounter when using the EMC integration. For general Data Protector troubleshooting information, see the *HP OpenView Storage Data Protector Troubleshooting Guide*.

✓ Ensure that the latest official Data Protector patches are installed. See the online Help index: "patches" on how to verify this.

✓ See the *HP OpenView Storage Data Protector Software Release Notes* for general Data Protector and integration-specific limitations, as well as known problem and workarounds.

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

Checks and Verifications

✓ On the application and backup systems, examine system errors reported in:

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

Windows: `<Data_Protector_home>\log\debug.log`

Backup Problems

Problem   **You cannot select EMC mode in the Data Protector GUI when creating a backup specification**

Action   Check that the EMC Symmetrix Agent integration module is installed on the application and backup systems. To do that, open the cell info file located in:

Windows Cell Manager:
<Data_Protector_Home>\Config\server\cell\cell_info

UNIX Cell Manager: /etc/opt/omni/server/cell/cell_info

File contents should look similar to the following:

```
-host "hpsap001.bbn.hp.com" -os "hp s800 hp-ux-11.00" -cc A.06.00
-da A.06.00 -emc A.06.00

-host "hpsap002.bbn.hp.com" -os "hp s800 hp-ux-11.00" -cc A.06.00
-da A.06.00 -ma A.06.00 -emc A.06.00
```

Problem   **On the application system, dismounting a filesystem fails**

Action   In Split pre-exec script, stop all processes using the filesystem.

Problem   **Disks synchronization fails (split fails)**

To successfully split the disks, EMC Agent first checks the status of the links. Links can only be split after all devices are synchronized. EMC Agent checks the status of links every 30 seconds and retries 15 times.

Action   Increase the time frame for synchronization by setting SYMA_SYNC_RETRY and SYMA_SLEEP_FOR SYNC variables.

See "ZDB Agents Omnirc Variables" on page A-18 for more information.

Problem   **EMC device is not part of a BCV pair**

Action   If the TimeFinder or SRDF + TimeFinder configuration is used, check that all backup disks on the application system have an associated BCV device on the backup system.

Problem	**Device group cannot be created**
Action	Check if any of the previous sessions was improperly stopped, and run EMC Agent recovery for this session on the backup system. See "Recovery Using the EMC Agent" on page 194 for instructions.
Problem	**Adding a device into a device group/associating BCV to a device group fails**
Action	Check if any of the previous backups was improperly stopped, and run EMC Agent recovery for this session on the backup system. See "Recovery Using the EMC Agent" on page 194 for instructions.
Problem	**Volume group on the backup system cannot be de-activated.**
Action	Stop the processes that run on the volume group filesystem.
Problem	**Rebuilding the Data Protector EMC database fails**
Action	Run a discovery from:

Windows: `<Data_Protector_home>\bin\syma -init`

UNIX: `/opt/omni/lbin/syma -init`

on both the application and backup systems. If the operation succeeds, disable the `Run discovery of Symmetrix environment` option and restart the backup.

If discovery fails, run the `symcfg -discover` command.

Problem	**Resolving an object fails**
Action	Check the EMC Agent log file on the application system and ensure that all objects logged into this file are created on the mirrored EMC devices.
Problem	**Invalid link state on the EMC device**
Action	Check the link state. If it is split, set the `Re-establish links before backup` option.

Error Messages

This section provides information on error messages.

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)

One of previous sessions failed.

Actions   • Run a recovery of the failed session to create a consistent
environment.

• Check that the /var directory is not full (if it is full, EMC Agent does
not have enough space to write its record into the file; the session
then fails). Clean the directory and restart the session. TBD

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

Backup volume group is not deactivated or there is a problem with
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 to split links and activate the backup volume group manually. If
this is not done, 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.

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, create
pseudo-devices for all gatekeepers and BCV devices.

• See 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 the links or use the option Establish Links Before
Backup. If the problem persists, run:

symrdf -g *<Dg_name>* establish -bypass

CAUTION   See the symrdf man page about the bypass option before running this
command.

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 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)

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)

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 backup specification.

Split Mirror Restore Problems

Problem   **Deactivating volume groups during restore fails (HP-UX only)**

Action   In the Split pre-exec 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.

Problem   **Disks synchronization fails (split fails)**

To successfully split the disks, EMC Agent first checks the status of the
links. Links can only be split after all devices are synchronized. EMC
Agent checks the status of links every 30 seconds and retries 15 times.

Action   Increase the time frame for synchronization by setting SYMA_SYNC_RETRY
and SYMA_SLEEP_FOR SYNC variables.

See "ZDB Agents Omnirc Variables" on page A-18 for more information.

Problem   **EMC device is not part of a BCV pair**

Action   If the TimeFinder or SRDF + TimeFinder configuration is used, check
that all backup disks on the application system have an associated BCV
device on the backup system.

Problem   **Device group cannot be created**

Action   Check if any of the previous sessions was improperly stopped, and run
EMC Agent recovery for this session on the backup system. See
"Recovery Using the EMC Agent" on page 194 for instructions.

Problem   **Adding a device into a device group/associating BCV to a device
group fails**

Action   Check if any of the previous backups was improperly stopped, and run
EMC Agent recovery for this session on the backup system. See
"Recovery Using the EMC Agent" on page 194 for instructions.

Problem	**Rebuilding the Data Protector EMC database fails**
Action	Run a discovery from:

Windows: `<Data_Protector_home>\bin\syma -init`

UNIX: `/opt/omni/lbin/syma -init`

on both the application and backup systems. If the operation succeeds, disable the `Run discovery of Symmetrix environment` option and restart the backup.

If discovery fails, run the `symcfg -discover` command.

Problem	**Resolving an object fails**
Action	Check the EMC Agent log file on the application system and ensure that all objects logged into this file are created on the mirrored EMC devices.
Problem	**Invalid link state on the EMC device**
Action	Check the state of the link. If it is split, set the `Re-establish links before backup` option.

Error Messages

This section provides information on error messages.

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)

One of previous sessions failed.

Actions
- Run a recovery of the failed session to create a consistent environment.

- Check that the `/var` directory is not full (if it is full, EMC Agent does not have enough space to write its record into the file; the session then fails). Clean the directory and restart the session. TBD

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

Backup volume group is not deactivated or there is a problem with 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 to split links and activate the backup volume group manually. If this is not done, 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.

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 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, create pseudo-devices for all gatekeepers and BCV devices.

- See 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:

symrdf -g *<Dg_name>* establish -bypass

CAUTION   See the symrdf man page about the bypass option before running this
command.

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)

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)`

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 backup specification.

Recovery Using the EMC Agent

If a backup or other operation did not finish successfully, 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 in 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
```

When a record is entered, it is marked as valid. If the session is not successful, the record is marked as invalid. 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 that re-establish links and delete device groups. Next split mirror backup or split mirror restore 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 *<sessionID>* from the Data Protector GUI as shown in Figure 16-1 on page 195.

Figure 16-1   **Obtaining Session ID**
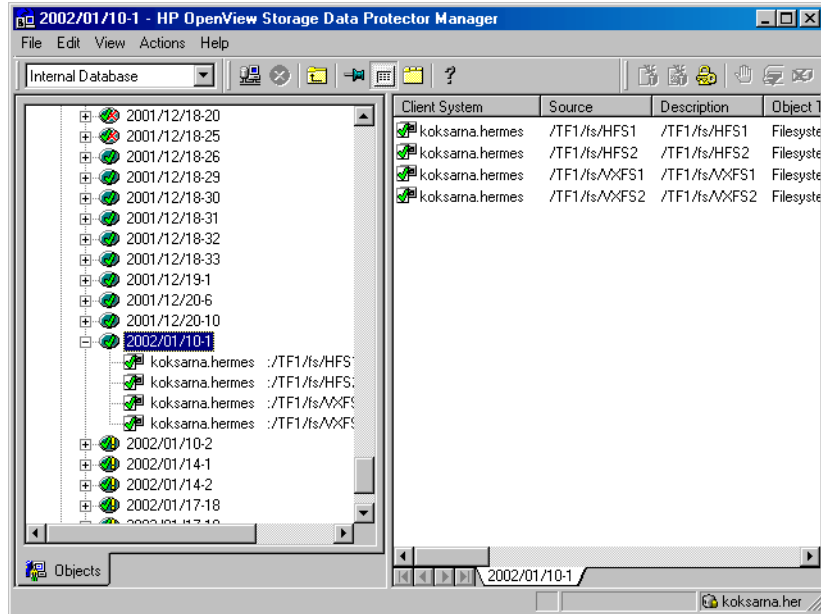
The split option disables synchronization of links.

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.

A   Appendix

Running and Scheduling ZDB Sessions

You run ZDB by scheduling a backup or starting an interactive backup session.

Scheduling ZDB Sessions

To schedule a filesystem or disk image ZDB, create a new or modify an existing backup specification. For detailed steps, see the online Help index: "scheduling backups on specific dates and times".

For general information on scheduling, see "Scheduling Unattended Backups" in the *HP OpenView Storage Data Protector Administrator's Guide* TBD.

Figure A-1   **Scheduling ZDB to Disk/Disk+Tape**

I
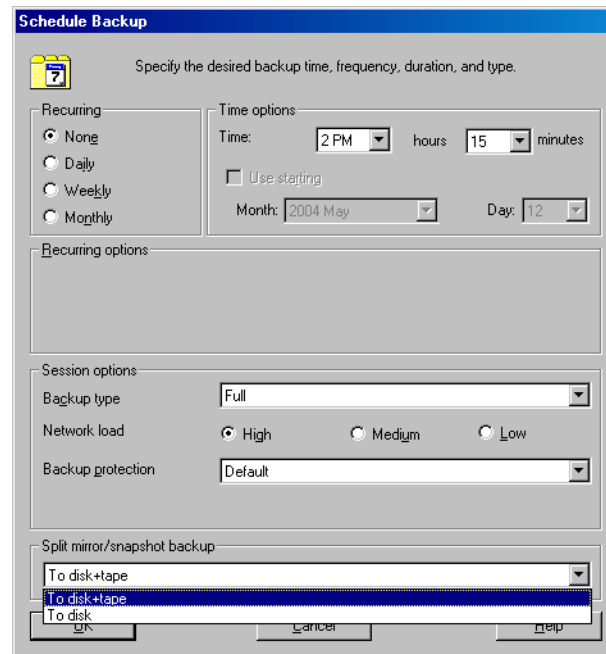
Starting Interactive ZDB Sessions

NOTE   When running concurrent ZDB sessions using one or several application systems, consider the limitations described in the *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide*.

Using the GUI

1. In the Context List, select Backup.

2. In the Scoping Pane, expand Backup, Backup Specification, and Filesystem. Right-click the required backup specification, and select Start Backup.

3. The Start Backup dialog box appears.

   With ZDB to tape and ZDB to disk, specify Backup Type. Incremental ZDB is not supported for ZDB to disk.

   To run ZDB to disk or ZDB to disk+tape (Track the replica for instant recovery selected), select To disk or To disk+tape in the Split mirror/snapshot backup drop-down list.
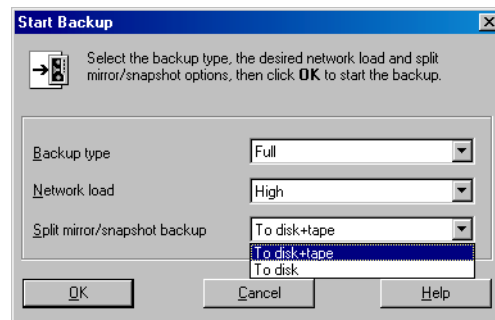
   For information on options, press **F1**.

4. Click OK.

Using the CLI

Run:

ZDB to tape, ZDB to disk+tape: omnib -datalist *<Name>*

ZDB to disk: `omnib -datalist` *<Name>* `-disk_only`

where *<Name>* is the backup specification name. See the `omnib` man page for details.

Alternate Paths Support

For systems with multiple host adapters and connections to a disk array, the alternate paths solution, configured on the backup system, performs dynamic load balancing and monitors each path to ensure that I/O completes its transactions. If a path between a disk array and a server fails, alternate path software automatically switches to an alternate path, removing the failed path from I/O rotation without data loss. Switchover is transparent to applications, so they continue unaffected.

Alternate paths load balancing is enabled for the newly created target volumes in a replica for HP StorageWorks AutoPath (VA, XP) and HP StorageWorks Secure Path (EVA). TBD

You can control AutoPath load balancing using the OB2AUTOPATH_BALANCING_POLICY variable (by default, AutoPath Round Robin load balancing policy is used). See "ZDB Agents Omnirc Variables" on page A-18 for more information.

When using AutoPath, consider the following:

• You need to perform additional steps to ensure that the VA integration using AutoPath works properly. See "Configuring VA with HP StorageWorks AutoPath Installed" on page 7 for instructions.

• If a failover to an alternate path occurs and AutoPath Shortest Queue Length load balancing is set, the session completes with errors.

• If a failover to an alternate path occurs during disk image backup without using raw logical volumes (rlvols), the session completes with errors. If rlvols are used, the session completes successfully.

NOTE   On EVA with HP StorageWorks Secure Path, load balancing configured by Secure Path is used; you cannot change the load balancing policy using Data Protector.

For information on which alternate paths solutions are supported by which ZDB integration, see supported matrices in the *HP OpenView Storage Data Protector Software Release Notes*.

Cluster Configurations

Data Protector ZDB integrations support:

- MC/ServiceGuard (HP-UX, all disk arrays)

- Veritas Cluster on Solaris and Microsoft Cluster Server (VA, EVA, XP)

- EMC GeoSpan for Microsoft Cluster Service (EMC)

For more information on cluster support, see the *HP OpenView Storage Data Protector Software Release Notes* and the *HP OpenView Storage Data Protector Administrator's Guide*.

Sections below discuss supported ZDB cluster configurations.

Figures A-2 to A-7 illustrate Data Protector *application* backup configurations. For *filesystem and disk image* backup, the Disk Agent must be installed instead of a Data Protector integration software component; an application database and binaries is not installed as presented in the figures.

NOTE   For applications in a cluster, use a floating IP address rather than a static one. This allows a successful backup to start even after a local failover.

Client on the Application System in a Cluster, Cell Manager not in a Cluster

Scenarios
- Application failover during backup: session fails and must be restarted manually.
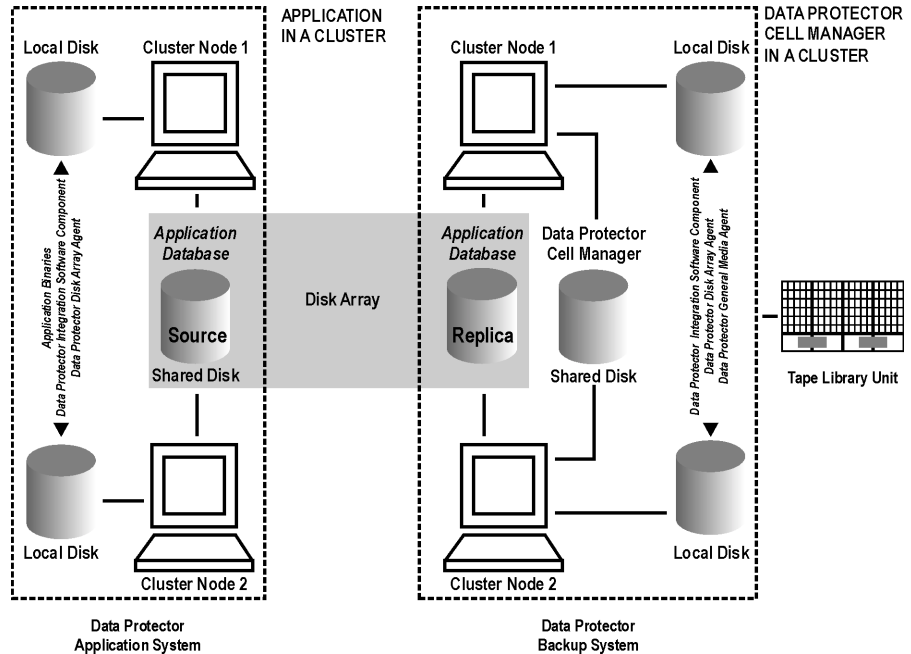
- Application failover before backup: session completes successfully.

Install:

- On the application system on all cluster nodes on local disks: application binaries, Data Protector integration software component, Data Protector disk array agent.

- On the application system cluster shared disk: application database. Note that this shared disk must be a disk array replicated disk.

- On the backup system on local disks: Data Protector integration software component, Data Protector disk array agent, Data Protector General Media Agent.

Figure A-2   **Client on the Application System in a Cluster**
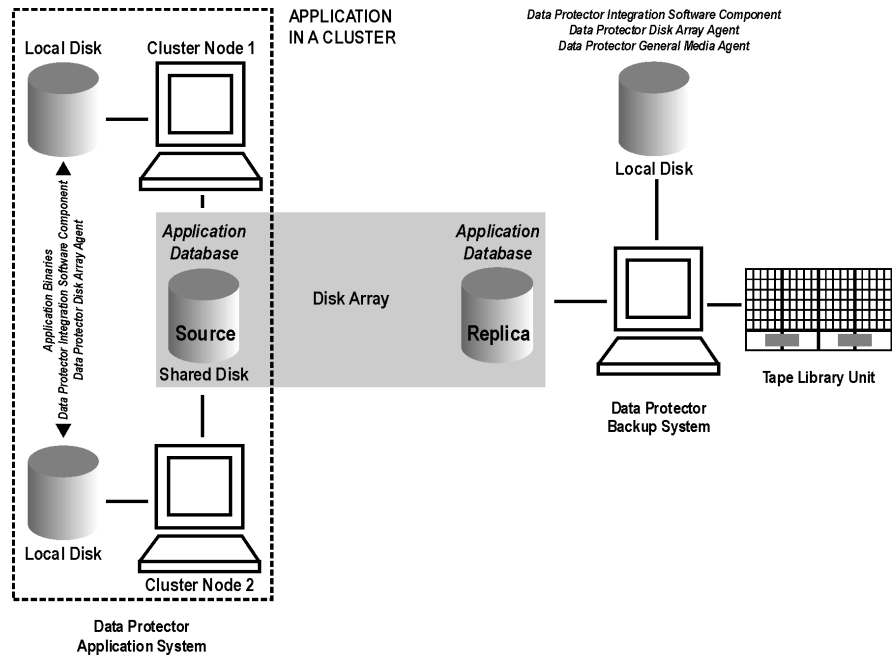
Cell Manager and Client on the Application System in a Cluster

Scenarios
- Application/Data Protector failover during backup: session is restarted automatically.

- Application/Data Protector failover before backup: session completes successfully.

Limitations
- Not supported on Veritas Cluster.

- Split mirror restore is not possible (XP, EMC).

Install:

- On the application system on all cluster nodes on local disks: application binaries, Data Protector integration software component, Data Protector disk array agent.

- On the application system cluster shared disk: application database. Note that this shared disk must be a disk array replicated disk.

- On the application system cluster shared disk: Cell Manager.

- On the backup system on local disks: Data Protector integration software component, Data Protector disk array agent, Data Protector General Media Agent.

Configure Cell Manager cluster's critical resources in the same cluster group/package as those for the application being backed up.

Figure A-3   **Cell Manager and 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

Scenarios
- Application failover during backup: session fails and must be restarted manually.

- Application failover before backup: session completes successfully.

- Cell Manager failover during backup: failed session is automatically restarted, provided the option `Restart backup of failed objects` is set.

- Cell Manager failover before backup: session completes successfully.

See the *HP OpenView Storage Data Protector Administrator's Guide* for information on defining cluster options for Cell Manager failover.

Limitations
- Not supported on Veritas Cluster.

Install:

- On the application system on all cluster nodes on local disks: application binaries, Data Protector integration software component, Data Protector disk array agent.

- On the application system cluster shared disk: application database. Note that this shared disk must be a disk array replicated disk.

- On the backup system cluster shared disk: Cell Manager.

- On the backup system on all cluster nodes on local disks: Data Protector integration software component, Data Protector disk array agent, Data Protector General Media Agent.

Figure A-4   **Client on the Application System in a Cluster, Cell Manager on the Backup System in a Cluster**

Cell Manager on the Backup System in a Cluster

Scenarios

- Cell Manager failover during backup: session is automatically restarted, provided the option `Restart backup of failed objects` is set.

- Cell Manager failover before backup: session completes successfully.

See the *HP OpenView Storage Data Protector Administrator's Guide* for information on defining cluster options for Cell Manager failover.

Limitations

- Not supported on Veritas Cluster.

Install:

- On the application system: application binaries, Data Protector integration software component, Data Protector disk array agent.

- On the application system cluster shared disk: application database. Note that this shared disk must be a disk array replicated disk.

- On the backup system cluster shared disk: Cell Manager. Note that this shared disk must be a disk array replicated disk.

- On the backup system on all cluster nodes on local disks: Data Protector integration software component, Data Protector disk array agent, Data Protector General Media Agent.

Figure A-5   **Cell Manager on the Backup System in a Cluster**
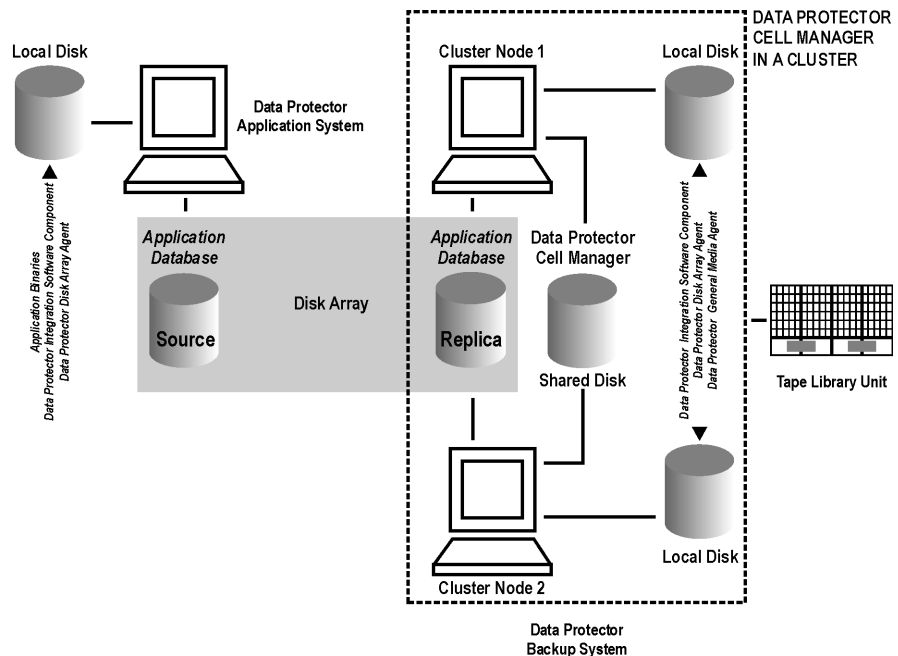
Client on the Application System in a Cluster, Cell Manager in a Cluster

Cell Manager is installed in a cluster on any system that is not a backup or application system.

Scenarios
- Application failover during backup: session fails and must be restarted manually.

- Application failover before backup: session completes successfully.

- Cell Manager failover during backup: failed session is automatically restarted, provided the option `Restart backup of failed objects` is set.

- Cell Manager failover before backup: session completes successfully.

See the *HP OpenView Storage Data Protector Administrator's Guide* for information on defining cluster options for Cell Manager failover.

Limitations
- Not supported on Veritas Cluster.

Install:

- On the application system on all cluster nodes on local disks: application binaries, Data Protector integration software component, Data Protector disk array agent.

- On the application system cluster shared disk: application database. Note that this shared disk must be a disk array replicated disk.

- On any system cluster shared disk: Cell Manager.

- On the backup system on local disks: Data Protector integration software component, Data Protector disk array agent, Data Protector General Media Agent.

Figure A-6   **Client on the Application System in a Cluster, Cell Manager in a Cluster**
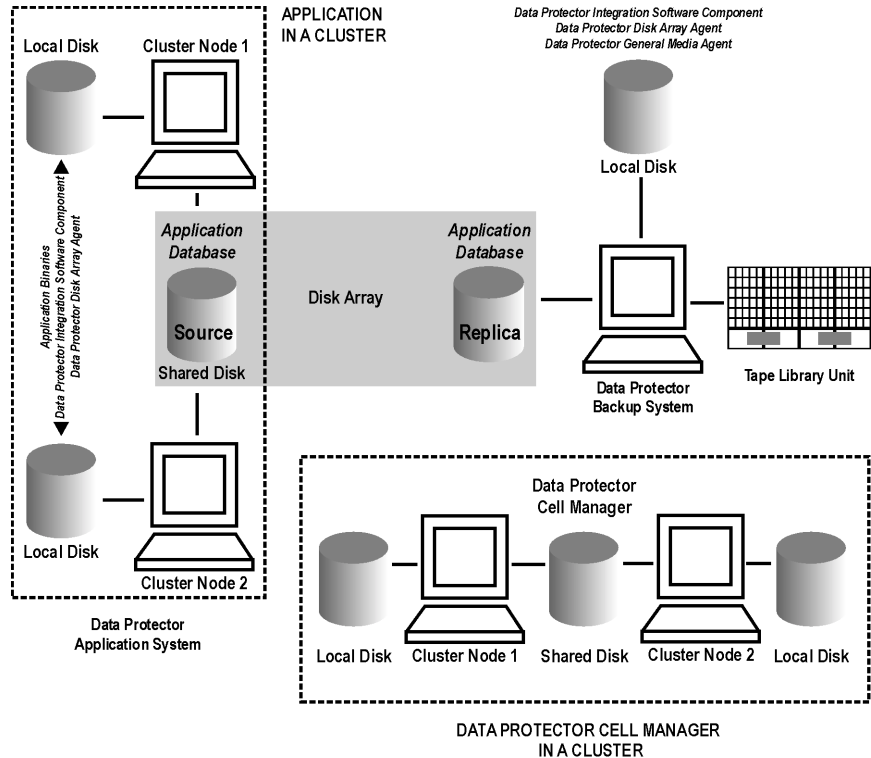
EMC GeoSpan for Microsoft Cluster Service

Cell Manager is not in a cluster; application client is in a cluster on the application system.

EMC Symmetrix SRDF links are controlled by EMC GeoSpan, EMC Symmetrix TF links are controlled by Data Protector.

Scenarios
- Application/hardware failover during backup: session fails and must be restarted manually. The backup system in the backup specification must be set as the backup system for the active node.

- Application failover before backup: session completes successfully if the backup system is set as the backup system for the active node.

Install:

- On the application system on all cluster nodes on local disks: application binaries, Data Protector integration software component, EMC Agent.

- On the application system cluster shared disk: application database. Note that this shared disk must be a disk array replicated disk.

- On the backup system on local disks: Data Protector integration software component, EMC Agent, Data Protector General Media Agent.

Figure A-7   **EMC GeoSpan for Microsoft Cluster Service**
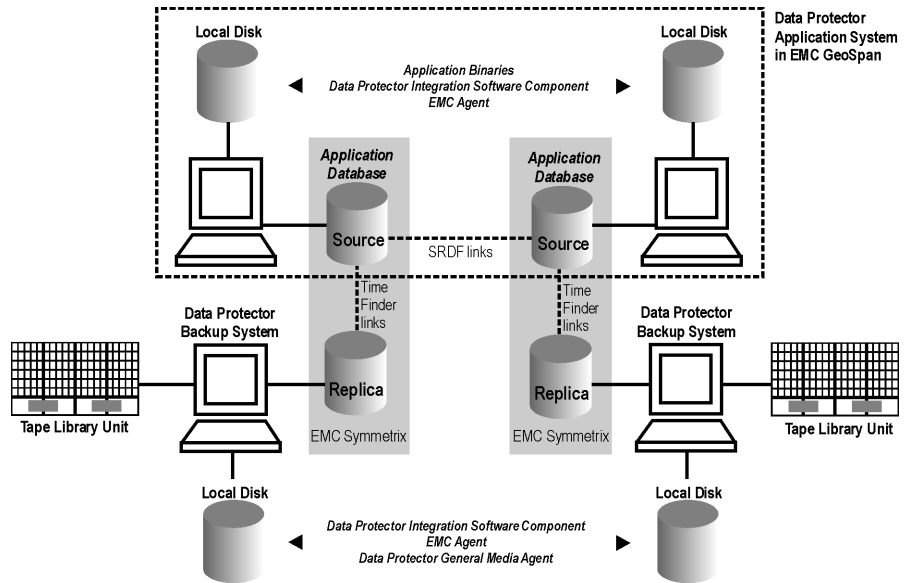
Instant Recovery in a Cluster

With an application/filesystem running in MC/ServiceGuard or Microsoft
Cluster Server on the application system, instant recovery requires some
additional steps.

IMPORTANT   If HP-UX LVM mirroring is used, see also "Instant Recovery and LVM
Mirroring" on page 139.

MC/ServiceGuard

1. Stop the application cluster package:

   cmhaltpkg *<app_pkg_name>*

2. In the *shell script for starting, shutting down and monitoring the
   database*, comment the lines that monitor application processes (by
   putting # at the beginning of the line).

Oracle8 Example   `#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}`

This shuts down the application (database) running in the cluster
without causing a failover.

3. Restart the application cluster package:

   cmrunpkg *<app_pkg_name>*

4. Shut down the application (database).

5. Start instant recovery. For instructions, see:

   - "Instant Recovery Procedure" on page 35 (VA)

   - "Instant Recovery Procedure" on page 81 (EVA)

   - "Instant Recovery Procedure" on page 136 (XP)

IMPORTANT   When performing instant recovery to the node other than that backed up, select the `Check the data configuration consistency` instant recovery option.

6. When the session finished, stop the application cluster package:

   cmhaltpkg *<app_pkg_name>*

7. Uncomment the lines (delete #) commented in step 2 of this procedure to re-enable an application failover.

8. Restart the application cluster package:

   cmrunpkg *<app_pkg_name>*

9. Recover the database. See the database documentation for detailed procedures.

NOTE   After resynchronization with the application system finishes, enable replicated volume groups on the application system in the exclusive mode by setting the `ZDB_IR_VGCHANGE_A` variable on the application system to `vgchange -a e`. See "ZDB Agents Omnirc Variables" on page 18 for more information.

Microsoft Cluster Server

1. Using the Cluster Administrator utility or Cluster CLI, take the application cluster resource offline. See the Microsoft Cluster Server documentation for detailed instructions.

2. Shut down the application (database).

3. Start instant recovery. For instructions, see:

   - "Instant Recovery Procedure" on page 35 (VA)

   - "Instant Recovery Procedure" on page 81 (EVA)

   - "Instant Recovery Procedure" on page 136 (XP)

4. Restart the application (database).

5. Recover the database. See the database documentation for detailed procedures.

6. Using the Cluster Administrator utility or CLI, put the application cluster resource online.

ZDB Agents Omnirc Variables

To customize ZDB agents, you can set environment variables in the
/opt/omni/.omnirc (UNIX) or *<Data_Protector_home>*\omnirc
(Windows) file, on both application and backup systems. See the *HP
OpenView Storage Data Protector Administrator's Guide* for information
on the omnirc file.

Common ZDB Agents Variables

This section explains omnirc variables that can be set for all ZDB agents.

ZDB_PRESERVE_MOUNTPOINTS: Controls the mount point
creation on the backup system. Possible values are 0 (default) and 1.

NOTE   The creation of mount points is also influenced by ZDB_MULTI_MOUNT and
ZDB_MOUNT_PATH variables.

If ZDB_PRESERVE_MOUNTPOINTS = 0, the mount point for a backed up
filesystem is created as follows:

- ZDB_MULTI_MOUNT = 1:

  — *VA, EVA:*

    *<BU_MOUNT_PATH>/<application_system_name>/<mountpoint_
    name_on_application_system>_<SessionID>*

  — *XP:*

    *<BU_MOUNT_PATH>/<application_system_name>/<mountpoint_
    name_on_application_system>_<LDEV_MU#>*

- ZDB_MULTI_MOUNT = 0 or not used:

  *<BU_MOUNT_PATH>/<application_system_name>/<mountpoint_nam
  e_on_application_system>*

where *<BU_MOUNT_PATH>* is:

  UNIX client:

  /var/opt/omni/tmp (ZDB_MOUNT_PATH not set), or

<ZDB_MOUNT_PATH> (ZDB_MOUNT_PATH set).

Windows client:

<Data_Protector_home>\tmp (ZDB_MOUNT_PATH not set), or

<ZDB_MOUNT_PATH> (ZDB_MOUNT_PATH set).

If ZDB_PRESERVE_MOUNTPOINTS = 1, the mount point for a backed up filesystem is created on the backup system in:

> *UNIX client:* /*<mountpoint_name_on_application_system>*

> *Windows client:* *<mountpoint_name_on_application_system>* or *<Drive_letter_on_the_app_system>*

IMPORTANT   For disk image, Oracle8/9, SAP R/3, and MS SQL Server 2000 ZDB, ZDB_PRESERVE_MOUNTPOINTS = 1 and its override is ignored; ZDB_MULTI_MOUNT and ZDB_MOUNT_PATH variables are ignored.

ZDB_MULTI_MOUNT: Specifies, together with ZDB_PRESERVE_MOUNTPOINTS and ZDB_MOUNT_PATH, the mount point creation 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 ZDB_MULTI_MOUNT = 1 (default), SessionID (VA, EVA) or LDEV MU# (XP) is appended at the end of the mount point path, thus enabling every group of mount points for one replica in the replica set to be mounted to their own mount points.

If ZDB_MULTI_MOUNT = 0, the selected group of mount points for one replica in the replica set is mounted to the same mount points.

ZDB_MULTI_MOUNT is ignored if ZDB_PRESERVE_MOUNTPOINTS = 1.

ZDB_MOUNT_PATH: Specifies, together with ZDB_PRESERVE_MOUNTPOINTS and ZDB_MULTI_MOUNT, the mount point creation on the backup system.

By default, this variable is not set. In this case, the first part of the mount point path is set as:

- *UNIX client:* /var/opt/omni/tmp

- *Windows client:* <Data_Protector_home>\tmp

To set this variable, specify the first part of the mount point path.

NOTE This variable is ignored if ZDB_PRESERVE_MOUNTPOINTS = 1.

ZDB_ALWAYS_POST_SCRIPT: By default, the command set by the Restart the application option is not executed if Stop/quiesce the application fails.

If this variable = 1, the command set by Restart the application is always executed.

Default: 0.

ZDB_BACKUP_VG_EXIST: For systems configured with multiple HBAs and connections to a disk array, the alternate paths solution performs dynamic load balancing. By default, during preparation for backup and restore, Data Protector creates a volume group with the disk on the first HBA as the primary path.

To disable volume group autoconfiguration on the backup host and load balance the data across multiple paths manually, set this variable to 1. The existing backup volume group will be used in the next backup/restore session.

NOTE If this variable is set on HP-UX, volume groups are not removed from /etc/lvmtab on the backup system after each backup. See "Backup Options" on page 119 for information.

Default: 0.

ZDB_IR_VGCHANGE: Specifies the mode in which replicated volume groups on the application system are activated after restore (HP-UX). The variable can be set on the application system only.

NOTE Not supported on EMC.

Select from the following modes:

- Exclusive: ZDB_IR_VGCHANGE_A=vgchange -a e
- Shared: ZDB_IR_VGCHANGE_A=vgchange -a s
- Normal (default): ZDB_IR_VGCHANGE_A=vgchange -q n -a y

IMPORTANT Use exclusive mode to enable instant recovery if an application/filesystem runs in the MC/ServiceGuard cluster on the application system.

SMB_SCAN_RDSK_TIMEOUT: During backup system preparation, the system is scanned for new devices (Windows). When new devices are detected, they appear on the backup system as new physical drives. This variable sets the maximum time (in seconds) for which a ZDB Agent on the backup system waits for a new physical drive to appear.

Default: 30 seconds. Usually, it is sufficient, unless there are configuration problems on the backup system.

SMB_SCAN_FOR_VOLUME_TIMEOUT: Sets the maximum time (in seconds) for which a ZDB Agent on the backup system waits for new volumes to appear on the backup system (Windows). This happens after a physical drive is detected during backup system preparation.

Default: 300 seconds. Usually, it is sufficient, unless there are configuration problems on the backup system.

OB2AUTOPATH_BALANCING_POLICY: Specifies the HP StorageWorks 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 dynamically balances data load between the alternate paths to achieve optimum performance.

NOTE Not supported on EMC and EVA.

Possible values are:

- 0 [none]–No policy

- `1` [RR]–Round Robin policy (default)
- `2` [SQL]–Shortest Queue Length policy (default) TBD

IMPORTANT   If the AutoPath `Shortest Queue Length` load balance policy is set and failover to an alternate path occurs, the session is aborted.

- `3` [SST]–Shortest Service Time policy

See the AutoPath documentation for more information.

ZDB_IR_MANUAL_AS_PREPARATION: To manually prepare the application system for instant recovery (dismounting filesystems and disabling volume groups), set this variable to `1`. After instant recovery, manually enable volume groups and mount filesystems again.

Use this variable also if automatic preparation of the application system fails because the application data configuration changed after backup. For example, if a failover to a secondary cluster node occurred between backup and instant recovery, Data Protector may have difficulty matching the secondary node resources to resources that existed on the primary node during backup.

Default: `0`.

VA Agent Specific Variables

There are no VA-specific `omnirc` variables.

See "Common ZDB Agents Variables" on page A-18 for descriptions of common `omnirc` variables.

EVA Agent Specific Variables

This section explains EVA-specific `omnirc` variables.

See also "Common ZDB Agents Variables" on page A-18.

EVA_HOSTNAMEALIASES: Use when the autodetection of FC HBAs on the backup system fails or when EVA_DISABLE_HBA_AUTODETECTION = 1.

During ZDB sessions, EVA Agent searches CV EVA for hostname objects matching the backup system hostname. By default, it searches for full and short backup system hostnames as they seen on the IP network.

Default: no hostnames specified. To add more hostnames to the search, specify hostname object names for this variable.

Example   Your backup host is represented within CV EVA by:

- /Hosts/Backup hosts/MyHost_Port1

- /Hosts/Backup hosts/MyHost_Port2

To force EVA client to find these host objects, set:

EVA_HOSTNAMEALIASES=MyHost_Port1,MyHost_Port2

EVA_DISABLE_HBA_AUTODETECTION: By default, EVA Agent (legacy) TBD attempts to perform a fully automatic backup system preparation during ZDB-to-tape/ZDB-to-disk+tape sessions. This is based on detection of FC HBAs on the backup system, followed by a search through the EVA storage system configuration. If this variable is set to 1, the automatic configuration of the backup system is skipped. In this case, configure the backup system manually and use the EVA_HOSTNAMEALIASES variable to define the backup system object name.

Default: 0.

EVA_EMAPI_MAX_RETRY: Some CV EVA commands may fail when the target EVA system is under heavy load. In this case, Data Protector waits for the period set by the EVA_EMAPI_RETRY_DELAY variable and tries to execute the command again. EVA_EMAPI_MAX_RETRY defines the number of retries before concluding that a command failed.

Default: 10 retries. Increase this value when your EVA system is under heavy load.

EVA_EMAPI_RETRY_DELAY: Some CV EVA commands may fail when the target EVA system is under heavy load. In this case, Data Protector waits for the period set by the this variable and tries to execute the command again. See also EVA_EMAPI_MAX_RETRY, which sets the maximum number of retries.

Default: 20 seconds.

TIP   Since each EVA command increases the load on your EVA system, it is better to increase the delay than the number of retries.

EVA_GETOBJID_MAX_RETRY: When creating new CV EVA objects, the EVA client sends creation requests to CV EVA and then collects data about created objects. When the target EVA system is under heavy load, object creation may be time-consuming. Therefore, Data Protector retries collecting created objects until all objects are found. This variable defines the number of retries before concluding that a command failed.

Default: 10 retries. Increase this value when your EVA system is under heavy load.

See also EVA_GETOBJID_RETRY_DELAY, which sets the delay between the retries.

EVA_GETOBJID_RETRY_DELAY: When creating new CV EVA objects, the EVA client sends the creation requests to CV EVA and then collects data about created objects. When the target EVA system is under heavy load, the object creation may be time-consuming. Therefore, Data Protector retries collecting created objects until all objects are found. This variable sets the delay between the retries.

Default: 10 seconds. TBD

See also EVA_GETOBJID_MAX_RETRY, which sets the number of retries.

TIP   Since each EVA command increases the load on your EVA system, it is recommended to increase the delay than the number of retries.

EVA_MSGWAITING_INTERVAL: Specifies the time interval between messages reporting the snapclone creation progress (monitored during ZDB-to-tape and ZDB-to-disk+tape sessions immediately after the backup system preparation). The backup option Delay the tape backup by a maximum of [XX] minutes if the snapclones are not fully created must be selected.

Default: 10 minutes.

EVA_CLONECREATION_QUERY_INTERVAL: Specifies the time interval between queries checking the snapclone creation progress (appears during ZDB-to-tape and ZDB-to-disk+tape sessions immediately after backup system preparation). The backup option `Delay the tape backup by a maximum of [XX] minutes if the snapclones are not fully created` must be selected. A shorter time interval ensures that snapclone completion is detected more promptly, but also increases the load on the EVA system.

Default: 5 minutes.

EVA_SCANDEVICES_LOCK_MAX_RETRY: After target volumes are created and presented to the backup system, the backup system scans devices to detect target volumes. Only one device scan runs at a time. In the case of parallel backups, one session scans, and the rest wait for this to finish before starting another scan.

This variable defines the number of retries for running device scanning before aborting the session. TBD EVA Agent waits for 10 seconds between the retries.

Default: 100 retries. If device scanning takes longer, increase the default setting.

XP Agent Specific Variables

This section explains XP-specific `omnirc` variables.

See also "Common ZDB Agents Variables" on page A-18.

SSEA_SPLIT_REPORT_RATE: During the split, XP Agent checks the status of mirrored disks within an interval specified by `SSEA_SPLIT_SLEEP_TIME` for the number of times specified by `SSEA_SPLIT_RETRY`. `SSEA_SPLIT_REPORT RATE` specifies the frequency of displaying the mirrored disks status to the Data Protector Monitor. For example, if `SSEA_SPLIT_SLEEP_TIME` is 2 seconds and `SSEA_SPLIT_REPORT_RATE` - 5, the status is displayed for every fifth check (every 10 seconds).

Default: 5.

SSEA_SPLIT_RETRY: During the split, XP Agent checks the mirrored disks status within an interval specified by `SSEA_SPLIT_SLEEP_TIME`. `SSEA_SPLIT_RETRY` sets the number of retries for the checks. If there is no progress after that, the split is aborted.

Default: 120 retries.

SSEA_SPLIT_SLEEP_TIME: During the split, XP Agent checks the mirrored disks' status for the number of times specified by SSEA_SPLIT_RETRY. SSEA_SPLIT_SLEEP_TIME sets the time interval between the checks.

Default: 2 seconds.

SSEA_SYNC_REPORT_RATE: During the disks' resynchronization, XP Agent checks the mirrored disks' status within an interval specified by SSEA_SYNC_SLEEP_TIME for the number of times specified by SSEA_SYNC_RETRY.

SSEA_SYNC_REPORT_RATE specifies the rate of displaying the mirrored disks status. For example, if SSEA_SYNC_SLEEP_TIME is 5 seconds and SSEA_SPLIT_REPORT_RATE - 2, the status is displayed for every second check (every 10 seconds).

Default: 2.

SSEA_SYNC_RETRY: During the disks' resynchronization, XP Agent checks the mirrored disks' status within an interval specified by SSEA_SYNC_SLEEP_TIME. SSEA_SYNC_RETRY sets the number of retries for these checks. If there is no progress after that, the resynchronization is aborted.

Default: 10 retries.

SSEA_SYNC_SLEEP_TIME: During the disks' resynchronization, XP Agent checks the mirrored disks' status for the number of times specified by SSEA_SYNC_RETRY. SSEA_SYNC_SLEEP_TIME sets the time interval between these status checks.

Default: 5 seconds.

SSEA_WAIT_PAIRS_PROPER_STATUS: All disk pairs must be in proper status (either STAT_PSUS/SSUS or STAT_PAIR) before a process continues. This variable defines the maximum waiting period for disk pairs to change to proper status.

Default: 120 minutes.

EMC Agent Specific Variables

This section explains EMC-specific omnirc variables.

See also "Common ZDB Agents Variables" on page A-18.

SYMA_LOCK_RETRY, **SYMA_SLEEP_FOR_LOCK**: Each time EMC Agent calls the WideSky library, it initiates the WideSky session, which locks the EMC Symmetrix database file. Other sessions must wait to get the lock.

Default: 15 retries, 30 seconds sleep time.

SYMA_SYNC_RETRY, **SYMA_SLEEP_FOR_SYNC**: To successfully split the disks, EMC Agent first checks the links' status (links can be split only after all devices are synchronized).

Default: 15 retries, 30 seconds sleep time.

These two variables are also used for incremental restore of device groups. EMC Agent starts the incremental restore only when there are no write pending tracks to devices in the restore device group.

Default: 15 retries; checking the number of write pending track - every 30 seconds.

SYMA_REC_FILE_LIMIT: Invalid records are automatically deleted when the EMC Agent recovery file exceeds a certain size.

Default: 102400 bytes.

SYMA_MOUNT_R2_READWRITE:

Determines the mode in which volume groups and filesystems are activated and mounted:

- `0`: read-only mode (default)
- `1`: read/write mode

For backup, it is sufficient to activate volume groups and filesystems in read-only mode. If you use the mirror for DSS or other tasks after backup, this may not be sufficient.

SYMA_UMOUNT_BEFORE_SPLIT:

Determines whether filesystems on the application system are dismounted before the split:

- `0`: not dismounted (default)
- `1`: dismounted before the split, remounted after (to ensure filesystem data is consistent)

A filesystem does not have a stop I/O to flush data from the filesystem cache to disk and stop I/O during the split. The only way to back up filesystems in split mirror mode is to dismount the mount point on the application system. If applications run on the filesystem, they control I/O to the disk. In this case, it is not necessary to dismount the filesystem before the split.

User Scenarios - Examples of ZDB Options

This section gives examples of backup policies with appropriate ZDB options.

VA and EVA Integrations

NOTE   If HP StorageWorks Secure Manager (VA) is used, set the password information using the `omnidbva` command and select the option `Integrate with VA LUN security`. For more information, see "LUN Security" on page 9.

Example 1   Backup to tape must be performed once a day (during the night). During the day, three copies must be available for instant recovery.

To implement such policy:

- Select `Track the replica for instant recovery`
- Set `Number of replicas rotated` to 3.

The following options are set automatically:

- `Use an existing replica` (VA)
- `Keep the replica after the backup`
- `Snapclone` (EVA)
- `Strict` (EVA)

Then, schedule the backup specification to start three ZDB-to-disk sessions during the day and one ZDB-to-disk+tape session during the night.

Example 2   Backup to tape must be performed every three hours. Replicas created are used for data mining (not for instant recovery), and can be obsoleted after three hours.

To implement such policy:

- Do not select `Track the replica for instant recovery`.

- Select `Keep the replica after the backup`.

- *EVA:* Set `Number of replicas rotated` to 1.

  VA: Select `Leave the backup system enabled`.

- Optionally, select `Enable the backup system in read/write mode` (UNIX).

- Set `ZDB_ORA_INCLUDE_CF_OLF` to 1 (see "ZDB Agents Omnirc Variables" on page A-18).

Then, schedule the backup specification to start one ZDB-to-tape session every three hours.

On VA, after the first backup, change the backup specification by selecting the option `Use an existing replica`. If this option is not set, every new session for this backup specification creates a new replica without removing it.

Example 3   Backup to tape must be performed every three hours. The replica created must be available for instant recovery for 12 hours.

To implement such policy:

- Select `Track the replica for instant recovery`

- Set `Number of replicas rotated` to 4.

The following options are set automatically:

- `Use an existing replica` (VA)

- `Keep the replica after the backup`

- `Snapclone` (EVA)

- `Strict` (EVA)

Then, schedule the backup specification to start eight ZDB-to-disk+tape sessions every three hours.

XP Integration

Example 1   A replica set is configured, with all replicas available for instant recovery. The next replica must be prepared according to replica set rotation after backup and force-synchronized before the next backup.

To implement such policy, select the following options:

- `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, with all replicas available for offline data processing after backup. The next replica must be prepared according to replica set rotation after backup, and the next backup must be aborted if data processing is not finished.

NOTE   This example assumes that offline data processing involves splitting links before data processing and resynchronizing links afterwards.

To implement such policy, select the following options:

- `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, with versions on replicas TBD available for on-demand offline data processing (links are split on demand and the backup system is prepared for offline data processing manually), but not for instant recovery. The replica must be prepared at the start of a backup session.

To implement such policy:

- Select `Prepare/resync the mirror disks at the start of the backup`.

- Do not select `Keep the replica after the backup`.

Example 4   A single replica is configured, with the version on the replica TBD available for offline data processing. The replica must be prepared at the start of a backup session.

To implement such policy, select the following options:

• `Keep the replica after the backup`

• `Leave the backup system enabled`

• `Prepare/resync the mirror disks at the start of the backup`

Conflicting Options   If a single replica is configured and the following 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.

NOTE   A conflict can also happen when a replica set is configured, depending on the replica set selection and the XP LDEV exclude file.

EMC Integration

Example 1   After backup, the replica must be discarded and prepared for the next backup at the end of the backup session.

To implement such backup policy:

• Select `Re-establish links after backup`

• Do not select `Re-establish links before backup`

Example 2   After backup, the replica must be used for offline data processing and prepared at the start of the next backup session.

To implement such backup policy:

• Select `Re-establish links before backup`

• Do not select `Re-establish links after backup`

EMC - Obtaining Disk Configuration Data

Obtaining disk information is necessary during installation and configuration. The examples below describe choosing and checking EMC devices (disks) for the correct connection type (TimeFinder, SRDF, SRDF+TimeFinder).

To check if the EMC configuration is correct, run:

- syminq to display disk type (blank, R1, R2, or BCV).
- symbcv list to display SLD-BCV pairs.
- symrdf list to display RDF1 - RDF2 pairs.

Example 1 The application system is connected to Primary (R1) Symmetrix and the backup system to Secondary (R2) Symmetrix. Disks 008 and 009 on the application system can be used for SRDF or SRDF+TimeFinder. To verify the configuration:

1. Run syminq on the application system and search for disk numbers in the Ser Num column.

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 Type column, you see that the disks are R1 (required for SRDF and SRDF+TimeFinder).

2. To check if the disks have the same serial number on the backup system, run `symrdf list` on the backup system.

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 see from the first two columns that the disks have the same numbers on both hosts.

3. Query additional information by running `syminq` and look for disks 008 and 009.

4. If you have SRDF+TimeFinder:

   a. Run `symbcv list` on the backup system to find associated BCVs.

BCV Device				Standard Device		Status	
Physical	**Sym**	**RDF Att.**	**Inv. Tracks**	**Physical**	**Sym**	**Inv. Tracks**	**BCV<=>STD TBD**
HP-UX: c1t8d0 Windows: DRIVE5	038	+	0	HP-UX: c1t1d0 Windows: Not Visible	008	0	Synch
HP-UX: c1t8d1 Windows: DRIVE6	039	+	0	HP-UX: c1t1d1 Windows: Not Visible	009	0	Synch

You can see which BCV belongs to which SLD. The first four columns contain information about BCVs, the last four about SLDs.

b. To ensure that the disks are correct, run syminq on the backup system and search for BCVs under disk numbers 038 and 039. The disk you find should be BCV.

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

Example2 Both application and backup systems are connected to the same EMC. Disks 048 and 049 on the application system can be used for TimeFinder. To check the configuration:

1. Run syminq on the application system and search for disk numbers in the Ser Num column.

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 Type column, you see that the disk type 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.

BCV Device				Standard Device		Status	
Physical	Sym	RDF att.	Inv. Tracks	Physical	Sym	Inv. Tracks	BCV<=>STD TBD
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

You can see which BCV belongs to which SLD. The first four columns contain information about BCVs, the last four about SLDs

You can double-check BCV by running `syminq` on the backup system. The disk you find should be BCV.

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

Listing and Unlocking Locked Backup Devices and Target Volumes

If backup/restore fails, devices and target volumes remain locked and are unlocked after 60 minutes (default).

To list locked devices and target volumes, run:

omnidbutil -show_locked_devs

To unlock all devices and target volumes, run:

omnidbutil -free_locked_devs

To unlock a specific device, run:

omnidbutil -free_locked_devs *<devname>*

To unlock a specific target volume, run:

VA, EVA:

omnidbutil -free_locked_devs *<wwn_lun>*, where:

<wwn>   VA/EVA world-wide-name.

<lun>   Logical unit number.

XP:

omnidbutil -free_locked_devs *<serial_ldev>*, where:

<serial>   XP serial number.

<ldev>   LDEV number.

For more information on the omnidbutil command, see its man page.

Additional Information for Troubleshooting

HP-UX

To identify physical devices belonging to a particular volume group, run:

On the application system:

- `strings /etc/lvmtab`

  All volume groups and devices belonging to volume groups are displayed.

- `vgdisplay -v /dev/VG_name`

  Logical volumes and devices for a specified volume group are displayed.

On the backup system:

- `/usr/symcli/bin/symdg list`

  Device group names and additional information about devices is displayed.

- `/usr/symcli/bin/symdg show DgName`

  Detailed information about devices and associated BCVs is displayed.

Windows

Run `symntctl` with additional parameters to get information about disks, signatures, and drives. See the EMC documentation for more information.

On the backup system, run:

- `symdg list` to display device group names and additional information about devices.

- `symdg show DgName` to display detailed information about devices and associated BCVs.

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

Glossary

(disk partitioning and logical volume configuration) and automatic recovery of the original system configuration and user data that was backed up during the full client backup.
These files are stored as an ASR archive file on the Cell Manager (in *<Data_Protector_home>*\Config\Server\dr\asr on a Windows Cell Manager or in /etc/opt/omni/server/dr/asr/ on a UNIX Cell Manager) as well as on the backup medium. The ASR archive file is extracted to three diskettes for 32-bit Windows systems or four diskettes for 64-bit Windows systems after a disaster occurs. You need these diskettes to perform ASR.

autochanger
See **library**

autoloader
See **library**

BACKINT *(SAP R/3 specific term)*
SAP R/3 backup programs can call the Data Protector backint interface program via an open interface, which enables them to communicate with Data Protector software. For backup and restore, SAP R/3 programs issue orders for the Data Protector backint interface.

backup API
The Oracle interface between the Oracle backup/restore utility and the backup/restore media management layer. The

interface defines a set of routines to allow the reading and writing of data to the backup media, the creation, searching and removing the backup files.

backup chain
This relates to a situation where full and incremental backups are performed. Based on the level of the incremental backups used (Incr, Incr 1, Incr 2, and so on), simple or rather complex dependencies of incrementals to previous incrementals can exist. The backup chain are all backups, starting from the full backup plus all the dependent incrementals up to the desired point in time.

backup device
A device configured for use with Data Protector, which can write data to and read data from storage media. This can be, for example, a standalone DDS/DAT drive or a library.

backup generation
One backup generation includes one full backup and all incremental backups until the next full backup.

backup ID
An identifier of an integration object that equals the session ID of the backup of this object. The backup ID is preserved when an object is copied, exported, or imported.

Glossary

backup object
A backup unit that contains all items backed up from one disk volume (logical disk or mount point). The backed up items can be any number of files, directories, or the entire disk or mount point. Additionally, a backup object can be a database entity or a disk image (rawdisk).

A backup object is defined by:

- Client name: hostname of the Data Protector client where the backup object resides.

- Mount point: the access point in a directory structure (drive on Windows and mount point on UNIX) on the client where the backup object is located.

- Description: uniquely defines backup objects with identical client name and mount point.

- Type: backup object type (for example filesystem or Oracle).

backup owner
Each backup object in the IDB has an owner. The default owner of a backup is the user who starts the backup session.

backup session
A process that creates a copy of data on storage media. The activities are

specified in a backup specification or an interactive session. All clients configured in one backup specification are backed up together in one backup session using the same backup type (full or incremental). The result of a backup session is a set of media, which was written to, also called the backup or media set.
See also **incremental backup** and **full backup.**

backup set
A complete set of integration objects associated with a backup.

backup set *(Oracle specific term)*
A logical grouping of backed up files that are created using the RMAN backup command. A backup set is a complete set of files associated with a backup. The files can be multiplexed to increase performance. A backup set contains either datafiles or archived logs, but not both together.

backup specification
A list of objects to be backed up, together with a set of devices or drives to be used, backup options for all objects in the specification, days and time that you want backups to be performed. The objects are entire disks/ volumes or parts of them such as files, directories, or even the Windows

Glossary

Registry for example. File selection lists such as include-lists and exclude-lists can be specified.

backup system *(ZDB specific term)*
A system connected to target volumes of one or multiple application systems. The backup system is typically connected to a backup device to perform the backup of the data in a replica.
See also **application system**, **target volume**, and **replica**.

backup types
See **incremental backup**, **differential backup**, **transaction backup**, **full backup** and **delta backup**.

backup view
Data Protector provides different views for backup specifications: By Type - according to the type of data available for backups/templates. Default view. By Group - according to the group to which backup specifications/templates belong. By Name - according to the name of backup specifications/templates. By Manager - if you are running MoM, you can also set the Backup view according to the Cell Manager to which backup specifications/templates belong.

BC *(EMC Symmetrix specific term)*
Business Continuance are processes that allow customers to access and manage

instant copies of EMC Symmetrix standard devices.
See also **BCV**.

BC *(HP StorageWorks Disk Array XP specific term)*
The Business Copy XP allows to maintain internal copies of HP StorageWorks Disk Array XP LDEVs for purposes such as data backup or data duplication. The copies (secondary volumes or S-VOLs) can be separated from the primary volumes (P-VOLs) and connected to a different system for various purposes, such as backup and development. For backup purposes, P-VOLs should be connected to the application system, and one of the S-VOL mirror sets to the backup system.
See also **HP StorageWorks Disk Array XP LDEV**, **CA**, **Main Control Unit**, **application system**, and **backup system**.

BC Process *(EMC Symmetrix specific term)*
A protected storage environment solution that has defined specially configured EMC Symmetrix devices as mirrors or Business Continuance Volumes to protect data on EMC Symmetrix standard devices.
See also **BCV**.

BC VA *(HP StorageWorks Virtual Array specific term)*
Business Copy VA allows you to

Glossary

maintain internal copies of HP StorageWorks Virtual Array LUNs for data backup or data duplication within the same virtual array. The copies (child or Business Copy LUNs) can be used for various purposes, such as backup, data analysis or development. When used for backup purposes, the original (parent) LUNs are connected to the application system and the Business Copy (child) LUNs are connected to the backup system.
See also **HP StorageWorks Virtual Array LUN**, **application system**, and **backup system**.

BCV *(EMC Symmetrix specific term)*
Business Continuance Volumes, or BCV devices, are dedicated SLDs that are pre-configured in the ICDA on which the business continuation operation runs. BCV devices are assigned separate SCSI addresses, differing from the addresses used by the SLDs they mirror. The BCV devices are used as splittable mirrors of the primary EMC Symmetrix SLDs that need to be protected.
See also **BC** and **BC Process**.

Boolean operators
The Boolean operators for the full text search functionality of the online Help system are AND, OR, NOT, and NEAR. Used when searching, they enable you to define your query precisely by creating a relationship between search terms. If no operator is specified in a

multi-word search, AND is used by default. For example, the query manual disaster recovery is equivalent to manual AND disaster AND recovery.

boot volume/disk/partition
A volume/disk/partition with files required for the initial step of the boot process. Microsoft terminology defines the boot volume/disk/partition as a volume/disk/partition containing the operating system files.

BRARCHIVE *(SAP R/3 specific term)*
An SAP R/3 backup tool that allows you to archive redo log files. BRARCHIVE also saves all the logs and profiles of the archiving process.
See also **SAPDBA**, **BRBACKUP** and **BRRESTORE**.

BRBACKUP *(SAP R/3 specific term)*
An SAP R/3 backup tool that allows an online or offline backup of the control file, of individual data files, or of all tablespaces and, if necessary, of the online redo log files.
See also **SAPDBA**, **BRARCHIVE** and **BRRESTORE**.

BRRESTORE *(SAP R/3 specific term)*
An SAP R/3 tool that can be used to restore files of the following type:

- Database data files, control files, and online redo log files saved with BRBACKUP

Glossary

- Redo log files archived with BRARCHIVE

- Non-database files saved with BRBACKUP

You can specify files, tablespaces, complete backups, log sequence numbers of redo log files, or the session ID of the backup.
See also **SAPDBA**, **BRBACKUP** and **BRARCHIVE**.

BSM
The Data Protector Backup Session Manager controls the backup session. This process always runs on the Cell Manager system.

CA *(HP StorageWorks Disk Array XP specific term)*
Continuous Access XP allows you to create and maintain remote copies of HP StorageWorks Disk Array XP LDEVs for purposes such as data duplication, backup, and disaster recovery. CA operations involve the main (primary) disk arrays and the remote (secondary) disk arrays. The main disk arrays contain the CA primary volumes (P-VOLs), which contain the original data and are connected to the application system. The remote disk arrays contain the CA secondary volumes (S-VOLs) connected to the backup system.
See also **BC** *(HP StorageWorks Disk*

Array XP specific term), **Main Control Unit** and **HP StorageWorks Disk Array XP LDEV**.

CAP *(StorageTek specific term)*
Cartridge Access Port is a port built into the door panel of a library. The purpose is to enter or eject media.

catalog protection
Defines how long information about backed up data (such as file names and file versions) is kept in the IDB.
See also **data protection**.

CDB
The Catalog Database is a part of the IDB that contains information about backups, object copies, restores, media management sessions, and backed up data. Depending on the selected logging level, it also contains file names and file versions. This part of the database is always local to the cell.
See also **MMDB**.

CDF file *(UNIX specific term)*
A Context Dependent File is a file consisting of several files grouped under the same pathname. The system ordinarily selects one of the files using the context of the process. This mechanism allows machine dependent executables, system data, and device files to work correctly from all hosts in a cluster while using the same pathname.

Glossary

cell
A set of systems that are under the control of a Cell Manager. The cell typically represents the systems on a site or an organizational entity, which are connected to the same LAN. Central control is available to administer the backup and restore policies and tasks.

Cell Manager
The main system in the cell where the essential Data Protector software is installed and from which all backup and restore activities are managed. The GUI used for management tasks can be located on a different system. Each cell has one Cell Manager system.

centralized licensing
Data Protector allows you to configure centralized licensing for the whole enterprise environment consisting of several cells. All Data Protector licenses are installed and kept on the Enterprise Cell Manager system. You can then allocate licenses to specific cells to suit your needs.
See also **MoM**.

Centralized Media Management Database (CMMDB)
See **CMMDB**.

channel *(Oracle specific term)*
An Oracle Recovery Manager resource allocation. Every allocated channel starts a new Oracle process, which performs backup, restore, and recovery actions. The type of channel allocated determines the type of media used:

- type "disk"

- type 'SBT_TAPE'

If the specified channel is type 'SBT_TAPE' and Oracle is integrated with Data Protector, the server process will attempt to read backups from or write data files to Data Protector.

circular logging *(Microsoft Exchange Server and Lotus Domino Server specific term)*
Circular logging is a Microsoft Exchange Server database and Lotus Domino Server database mode, in which the transaction log file contents are periodically overwritten after the corresponding data is committed to the database. Circular logging reduces disk storage space requirements.

client backup
A backup of all filesystems mounted on a client. Filesystems mounted on the client after the backup specification was created are not automatically detected.

client backup with disk discovery
A backup of all filesystems mounted on a client. When the backup starts, Data Protector discovers the disks on the clients. Client backup with disk

Glossary

discovery simplifies backup configuration and improves backup coverage of systems that often mount or dismount disks.

client or **client system**
Any system configured with any Data Protector functionality and configured in a cell.

cluster-aware application
It is an application that supports the cluster Application Programming Interface. Each cluster-aware application declares its own critical resources (disk volumes (on Microsoft Cluster Server), volume groups (on MC/ServiceGuard), application services, IP names and addresses ...).

CMD Script for OnLine Server
(Informix specific term)
Windows CMD script that is created in INFORMIXDIR when Informix OnLine Server is configured. The CMD script is a set of system commands that export environment variables for OnLine Server.

CMMDB
The Data Protector Centralized Media Management Database (CMMDB) is the result of merging MMDBs from several cells in the MoM environment. It allows you to share high-end devices and media across multiple cells in a MoM environment. One cell can control the robotics, serving the devices that are connected to systems in other cells. The CMMDB must reside on the MoM Manager. A reliable network connection between the MoM cell and the other Data Protector cells is highly recommended
See also **MoM**.

COM+ Registration Database
(Windows specific term)
The COM+ Registration Database and the Windows Registry store COM+ application attributes, class attributes, and computer-level attributes. This guarantees consistency among these attributes and provides common operation on top of these attributes.

command-line interface
A set of DOS and UNIX like commands that you can use in shell scripts to perform Data Protector configuration, backup, restore, and management tasks.

Command View (CV) EVA *(HP StorageWorks EVA specific term)*
The user interface that allows you to configure, manage, and monitor your HP StorageWorks EVA storage system. It is used to perform various storage management tasks, for example, creating virtual disk families, managing storage system hardware, creating snapclones and snapshots of virtual disks. The Command View EVA software runs on the HP OpenView

Glossary

Storage Management Appliance, and is accessed by a Web browser.
See also **HP StorageWorks EVA Agent (legacy)** and **HP StorageWorks EVA SMI-S Agent**.

concurrency
See **Disk Agent concurrency**.

control file *(Oracle and SAP R/3 specific term)*
An Oracle data file that contains entries specifying the physical structure of the database. It provides database consistency information used for recovery.

CRS
The Cell Request Server process (service) runs on the Data Protector Cell Manager. It starts and controls the backup and restore sessions. The service is started as soon as Data Protector is installed on the Cell Manager.
CRS runs under the account root on UNIX systems, and under any Windows account. By default, it runs under the account of the user, specified at installation time.

CSM
The Data Protector Copy Session Manager process controls the object copy session and runs on the Cell Manager system.

data file *(Oracle and SAP R/3* specific term)
A physical file created by Oracle that contains data structures such as tables and indexes. A data file can only belong to one Oracle database.

data protection
Defines how long the backed up data on media remains protected, that is, Data Protector will not overwrite it. When the protection expires, Data Protector will be able to reuse the media in one of the next backup sessions.
See also **catalog protection**.

Data Protector Event Log
A central repository of all Data Protector related notifications. By default, all notifications are sent to the Event Log. The Event Log is accessible only to Data Protector users in the Admin group and to Data Protector users who are granted the Reporting and notifications user rights. You can view or delete all events in the Event Log.

Data Protector user account
You can use Data Protector only if you have a Data Protector user account, which restricts unauthorized access to Data Protector and to backed up data. Data Protector administrators create this account specifying a user logon name, the systems from which the user can log on, and a Data Protector user group

Glossary

membership. This is checked whenever the user starts the Data Protector user interface or performs specific tasks.

data stream
Sequence of data transferred over the communication channel.

database library
A Data Protector set of routines that enables data transfer between Data Protector and a server of an online database integration, for example, the Oracle Server.

database parallelism
More than one database is backed up at a time if the number of available devices allows you to perform backups in parallel.

database server
A computer with a large database stored on it, such as the SAP R/3 or Microsoft SQL database. A server has a database that can be accessed by clients.

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

Glossary

device gets full, the backup automatically continues on a medium in the next device in the device chain.

device group *(EMC Symmetrix specific term)*
A logical unit representing several EMC Symmetrix devices. A device cannot belong to more than a single device group. All devices in a device group must be on the same EMC Symmetrix unit. You can use a device group to identify and work with a subset of the available EMC Symmetrix devices.

device streaming
A device is streaming if it can feed enough data to the medium to keep it moving forward continuously. Otherwise, the tape has to be stopped, the device waits for more data, reverses the tape a little and resumes to write to the tape, and so on. In other words, if the data rate written to the tape is less or equal the data rate which can be delivered to the device by the computer system, then the device is streaming. Streaming significantly improves the performance of the device and use of space.

DHCP server
A system running the Dynamic Host Configuration Protocol (DHCP) providing dynamic configuration of IP addresses and related information.

differential backup
An incremental backup (incr) based on any previous Data Protector backup (full or any incremental), which must still be protected.
See **incremental backup**.

differential backup *(MS SQL specific term)*
A database backup that records only the data changes made to the database after the last full database backup.
See also **backup types**.

differential database backup
A differential database backup records only those data changes made to the database after the last full database backup.

direct backup
A SAN-based backup solution in which data movement directly from disk to tape (or to other secondary storage) is facilitated by the SCSI Extended Copy (Xcopy) command. Direct backup lessens the backup I/O load on systems in a SAN environment. The data movement is facilitated directly from disk to tape (or to other secondary storage) by the SCSI Extended Copy (XCopy) command. The command is provided by any element of the infrastructure including bridges, switches, tape libraries, and disk subsystems.
See also **XCopy engine**.

Glossary

directory junction *(Windows specific term)*
Directory junctions use the reparse point concept of Windows. An NTFS 5 directory junction allows you to redirect a directory/file request to another location.

disaster recovery
A process to restore a client's main system disk to a state close to the time when a (full) backup was performed.

Disk Agent
A component needed on a client to back it up and restore it. The Disk Agent controls reading from and writing to a disk. During a backup session, the Disk Agent reads data from a disk and sends it to the Media Agent, which then moves it to the device. During a restore session the Disk Agent receives data from the Media Agent and writes it to the disk.

Disk Agent concurrency
The number of Disk Agents that are allowed to send data to one Media Agent concurrently.

disk discovery
The detection of disks during client backup with disk discovery. During this backup, Data Protector discovers (detects) the disks that are present on the client — even though they might not have been present on the system when the backup was configured — and backs

them up. This is particularly useful in dynamic environments, where configurations change rapidly. After the disks are expanded, each inherits all options from its master client object. Even if pre- and post-exec commands are specified once, they are started many times, once per each object.

disk group *(Veritas Volume Manager specific term)*
The basic unit of data storage in VxVM system. A disk group can consist of one or more physical volumes. There can be more than one disk group on the system.

disk image (rawdisk) backup
A high-speed backup where Data Protector backs up files as bitmap images. A disk image (rawdisk) backup does not track the files and directory structure stored on the disk, but stores a disk image structure on byte level. You can perform a disk image backup of either specific disk sections or a complete disk.

disk quota
A concept to manage disk space consumption for all or a subset of users on a computer system. This concept is used by several operating system platforms.

disk staging
The process of backing up data in several phases to improve the

Glossary

performance of backups and restores, reduce costs of storing the backed up data, and increase the data availability and accessibility for restore. The backup stages consist of backing up data to one media type first (for example disk) and later copying it to a different media type (for example tape).

Distributed File System (DFS)
A service that connects file shares into a single namespace. The file shares can reside on the same or on different computers. DFS provides client access to the resources in a location-transparent manner.

DMZ
The Demilitarized Zone (DMZ) is a network inserted as a "neutral zone" between a company's private network (intranet) and the outside public network (Internet). It prevents outside users from getting direct access to company servers in the intranet.

DNS server
In the DNS client-server model, this is the server containing information about a portion of the DNS database that makes computer names available to client resolvers querying for name resolution across the Internet.

domain controller
A server in a network that is responsible for user security and verifying passwords within a group of other servers.

DR image
Data required for temporary disaster recovery operating system (DR OS) installation and configuration.

DR OS
A disaster recovery operating system is an operating system environment in which disaster recovery runs. It provides Data Protector a basic runtime environment (disk, network, tape, and filesystem access). The OS has to be installed and configured before the Data Protector disaster recovery can be performed. DR OS not only hosts the Data Protector disaster recovery process but is also a part of the restored system because it replaces its own configuration data with the original configuration data.

drive
A physical unit that receives data from a computer system and can write it onto a magnetic medium (typically a tape drive). It can also read the data from the medium and send it to the computer system.

Glossary

drive index
A number that identifies the mechanical position of a drive inside a library device. This number is used by the robotic control to access a drive.

dynamic client
See **client backup with disk discovery**.

EMC Symmetrix Agent (SYMA)

(EMC Symmetrix specific term)
See **Symmetrix Agent (SYMA)**

emergency boot file *(Informix specific term)*
An Informix configuration file that resides in the *<INFORMIXDIR>*\etc directory (on HP-UX) or *<INFORMIXDIR>*/etc directory (on Windows) and is called ixbar.*<server_id>*, where *<INFORMIXDIR>* is the OnLine Server home directory and *<server_id>* is the value of the SERVERNUM configuration parameter. Each line of the emergency boot file corresponds to one backup object.

Enterprise Backup Environment
Several cells can be grouped together and managed from a central cell. The enterprise backup environment includes all clients located in several Data Protector cells which are managed and administered from a central cell using the Manager-of-Managers concept.
See also **MoM**.

Event Logs
Files in which Windows logs all events, such as the starting or stopping of services and the logging on and off of users. Data Protector can back up Windows Event Logs as part of the Windows configuration backup.

exchanger
Also referred to as SCSI Exchanger.
See also **library**.

exporting media
A process that removes all data about backup sessions, such as systems, objects, and file names, which reside on the media from the IDB. Information about the media and their relation to a pool is also removed from the IDB. The data on the media remains unchanged.
See also **importing media.**

Extensible Storage Engine (ESE)

(Microsoft Exchange Server specific term)
A database technology used as a storage system for information exchange in Microsoft Exchange Server.

failover
Transferring of the most important cluster data, called group (on Windows)

Glossary

or package (on Unix) from one cluster node to another. A failover can occur mostly because of software or hardware failures or maintenance on the primary node.

FC bridge
See **Fibre Channel bridge**

Fibre Channel
An ANSI standard for high-speed computer interconnection. Using either optical or copper cables, it allows the high speed bidirectional transmission of large data files and can be deployed between sites kilometers apart. Fibre Channel connects nodes using three physical topologies: point-to-point, loop, and switched.

Fibre Channel bridge
A Fibre Channel bridge or multiplexer provides the ability to migrate existing parallel SCSI devices, like RAID arrays, solid state disks (SSD), and tape libraries to a Fibre Channel environment. On one side of the bridge or multiplexer there is a Fibre Channel interface while on the other side there are parallel SCSI ports. The bridge or multiplexer enables SCSI packets to be moved between the Fibre Channel and parallel SCSI devices.

file depot
A file containing the data from a backup to a file library device.

file jukebox device
A device residing on disk consisting of multiple slots used to store file media.

file library device
A device which resides on a disk emulating a library with several media, hence containing multiple files, referred to as file depots.

File Replication Service (FRS)
A Windows service that replicates the domain controller store logon scripts and group policies. FRS also enables replication of Distributed File System (DFS) shares between systems and allows any server to perform replication activity.

file version
The same file can be backed up multiple times in case of full backups and incremental backups (if the file changed). If the log level ALL is selected for backup, Data Protector retains one entry in the IDB for the filename itself and one for each version (date/time) of the file.

filesystem
The organization of files on a hard disk. A filesystem is backed up so that the file attributes and the file contents are stored on the backup media.

Glossary

first level mirror (*HP StorageWorks Disk Array XP specific term*)
HP StorageWorks Disk Array XP allows up to three mirror copies of a Primary Volume and each of these copies can have additional two copies. The three mirror copies are called first level mirrors.
See also **Primary Volume**, and **MU numbers**.

fnames.dat
The fnames.dat files of the IDB contain information on the names of the backed up files. Typically, these files occupy about 20% of the IDB, if filenames are stored.

formatting
A process that erases any data contained on a medium and prepares it for use with Data Protector. Information about media (media ID, description, and location) is saved in the IDB as well as on the respective media (media header). Data Protector media with protected data are not formatted until the protection expires or the media are unprotected/recycled.

free pool
An auxiliary source of media for use by media pools when they run out of media. The media pools must be configured to use free pools.

full backup
A backup in which all selected objects are backed up, whether or not they have been recently modified.
See also **backup types**.

full database backup
A backup of all data in a database, not only the data that has been changed after the last (full or incremental) backup of the database. A full database backup does not depend on any other backup.

full mailbox backup
A full mailbox backup is a backup of the entire mailbox content.

full ZDB
A ZDB backup in which all selected objects are backed up, even if there are no changes from the previous backup.
See also **incremental ZDB**.

global options file
A file that allows you to customize Data Protector. It explains the global options, which cover various aspects of Data Protector, typically time-outs and limits, and affect the entire Data Protector cell. The file is located in the /etc/opt/omni/server/options directory on HP-UX and Solaris systems and in the *<Data_Protector_home>*\Config\Server\Options directory on Windows systems.

Glossary

group (*Microsoft Cluster Server specific term*)
A collection of resources (for example disk volumes, application services, IP names and addresses) that are needed to run a specific cluster-aware applications.

GUI
A cross-platform (HP-UX, Solaris, and Windows) graphical user interface, provided by Data Protector for easy access to all configuration, administration, and operation tasks.

hard recovery (*Microsoft Exchange Server specific term*)
A Microsoft Exchange Server database recovery that is performed after a restore by the database engine, using transaction log files.

heartbeat
A cluster data set with a time stamp carrying information about the operational status of a particular cluster node. This data set or packet is distributed among all cluster nodes.

Hierarchical Storage Management (HSM)
A method for optimizing the use of expensive hard disk storage by migrating less frequently used data to less expensive optical platters. When needed, the data is migrated back to

hard disk storage. This balances the need for fast retrieval from hard disk with the lower cost of optical platters.

Holidays file
A file that contains information about holidays. You can set different holidays by editing the Holidays file: /etc/opt/omni/server/Holidays on the UNIX Cell Manager and *<Data_Protector_home>*\Config\Server\holidays on the Windows Cell Manager.

host backup
See **client backup with disk discovery**.

hosting system
A working Data Protector client used for Disk Delivery Disaster Recovery with a Data Protector Disk Agent installed.

HP ITO
See **OVO**.

HP OpC
See **OVO**.

HP OpenView SMART Plug-In (SPI)
A fully integrated, out-of-the-box solution which "plugs into" HP OpenView Operations, extending the managed domain. Through the Data Protector integration, which is implemented as an HP OpenView SMART Plug-In, a user can have an

Glossary

arbitrary number of Data Protector Cell Managers monitored as an extension to HP OpenView Operations (OVO).

HP OVO
See **OVO**.

HP StorageWorks Disk Array XP LDEV
A logical partition of a physical disk within an HP StorageWorks Disk Array XP. LDEVs are entities that can be replicated in the Continuous Access XP (CA) and Business Copy XP (BC) configurations, or can be used as standalone entities.
See also **BC** *(HP StorageWorks Disk Array XP specific term)*, **CA** *(HP StorageWorks Disk Array XP specific term)*, and **replica**.

HP StorageWorks EVA Agent (legacy)
A Data Protector software module that executes all tasks required for the HP StorageWorks Enterprise Virtual Array integration operating on HP StorageWorks EVA with Command View (CV) EVA software v3.1 or lower, and the EVA VCS firmware v3.01x or lower.
See also **Command View (CV) EVA** and **HP StorageWorks EVA SMI-S Agent**.

HP StorageWorks EVA SMI-S Agent
A Data Protector software module that executes all tasks required for the HP StorageWorks Enterprise Virtual Array integration operating on HP StorageWorks EVA with Command View (CV) EVA software starting with v3.2. With the EVA SMI-S Agent, the control over the array is established through HP StorageWorks SMI-S EVA provider, which directs communication between incoming requests and CV EVA.
See also **Command View (CV) EVA**, **HP StorageWorks SMI-S EVA provider**, and **HP StorageWorks EVA Agent (legacy)**.

HP StorageWorks SMI-S EVA provider
An interface used for controlling HP StorageWorks Enterprise Virtual Array. SMI-S EVA provider runs as a separate service on the HP OpenView Storage Management Appliance system and acts as a gateway between incoming requests and Command View EVA. With the Data Protector HP StorageWorks EVA integration, SMI-S EVA provider accepts standardized requests from the EVA SMI-S Agent, communicates with Command View EVA for information or method invocation, and returns standardized responses.
See also **HP StorageWorks EVA SMI-**

Glossary

S Agent and **Command View (CV) EVA**.

HP StorageWorks Virtual Array LUN
A logical partition of a physical disk within an HP StorageWorks Virtual Array. LUNs are entities that can be replicated in the HP StorageWorks Business Copy VA configuration, or can be used as standalone entities.
See also **BC VA** and **replica**.

HP VPO
See **OVO**.

ICDA *(EMC Symmetrix specific term)*
EMC's Symmetrix Integrated Cached Disk Arrays (ICDA) is a disk array device that combines a set of physical disks, a number of FWD SCSI channels, an internal cache memory, and control and diagnostic software commonly referred to as the microcode.

IDB
The Data Protector Internal Database is an embedded database located on the Cell Manager that keeps information regarding which data is backed up, on which media it is backed up, how backup and restore sessions are run, and which devices and libraries are configured.

importing media
A process that re-reads all data about backup sessions which are on the medium back into the IDB. This then allows for fast and convenient access to the data on the media.
See also **exporting media.**

incremental backup
A backup that selects only files that have changed since a previous backup. Several levels of incremental backup are available, allowing selective backup of only files that have changed since the last incremental backup.
See also **backup types**.

incremental backup *(Microsoft Exchange Server specific term)*
A backup of the Microsoft Exchange Server data that has changed since the last full or incremental backup. With the incremental backup, only the transaction log files are backed up.
See also **backup types**.

incremental mailbox backup
An incremental mailbox backup backs up all the changes made to the mailbox after the last backup of any type.

incremental1 mailbox backup
An incremental1 mailbox backup backs up all the changes made to the mailbox after the last full backup.

Glossary

incremental (re)-establish *(EMC Symmetrix specific term)*
A BCV or SRDF control operation. In BCV control operations, an incremental establish causes the BCV device to be synchronized incrementally and to function as an EMC Symmetrix mirrored medium. The EMC Symmetrix devices must have been previously paired.
In SRDF control operations, an incremental establish causes the target (R2) device to be synchronized incrementally and to function as an EMC Symmetrix mirrored medium. The EMC Symmetrix devices must have been previously paired.

incremental restore *(EMC Symmetrix specific term)*
A BCV or SRDF control operation.
In BCV control operations, an incremental restore reassigns a BCV device as the next available mirror of the standard device in the pair. However, the standard devices are updated with only the data that was written to the BCV device during the time of the original pair split, and the data that was written to the standard device during the split is overwritten with data from the BCV mirror. In SRDF control operations, an incremental restore reassigns a target (R2) device as the next available mirror of the source (R1) device in the pair. However, the source (R1) devices are updated with only the data that was

written to the target (R2) device during the time of the original pair split, and the data that was written to the source (R1) device during the split is overwritten with data from the target (R2) mirror.

incremental ZDB
A ZDB to tape or ZDB to disk+tape session in which only changes from the last full or incremental protected backup are streamed to tape.
See also **full ZDB**.

Inet
A process that runs on each UNIX system or service that runs on each Windows system in the Data Protector cell. It is responsible for communication between systems in the cell and for starting other processes needed for backup and restore. The Inet service is started as soon as Data Protector is installed on a system. The Inet process is started by the inetd daemon.

Information Store *(Microsoft Exchange Server specific term)*
The Microsoft Exchange Server service that is responsible for storage management. Information Store in Microsoft Exchange Server manages two kinds of stores: mailbox stores and public folder stores. A mailbox store consists of mailboxes that belong to individual users. A public folder store contains public folders and messages

Glossary

that are shared among several users.
See also **Key Management Service** and
Site Replication Service.

initializing
See **formatting**.

Installation Server
A computer system that holds a
repository of the Data Protector
software packages for a specific
architecture. The Installation Server is
used for remote installation of Data
Protector clients. In mixed environments
at least two Installation Servers are
needed: one for UNIX systems and one
for Windows systems.

instant recovery (*ZDB specific term*)
A process in which a replica, produced
by a ZDB-to-disk or a ZDB-to-
disk+tape sessions, is used to restore the
contents of the source volumes to their
states at the time at which the replica
was created, avoiding the need to
perform a restore from tape. Depending
on the application/database concerned,
this may be all that is required, or other
steps, such as the application of
transaction log files, may be required for
full recovery.
See also **replica**, **zero downtime
backup (ZDB)**, **ZDB to disk**, and **ZDB
to disk+tape**.

integrated security *(MS SQL specific
term)*
Integrated security allows the Microsoft
SQL Server to use Windows
authentication mechanisms to validate
Microsoft SQL Server logins for all
connections. Using integrated security
means that users have one password for
both Windows and Microsoft SQL
Server. Integrated security should be
used in environments where all clients
support trusted connections.
Connections validated by Windows
Server and accepted by Microsoft SQL
Server are referred to as trusted
connections. Only trusted connections
are allowed.

integration object
A backup object of a Data Protector
integration, such as Oracle or SAP DB.

Internet Information Server (IIS)

(Windows specific term)
Microsoft Internet Information Server is
a network file and application server that
supports multiple protocols. Primarily,
IIS transmits information in Hypertext
Markup Language (HTML) pages by
using the Hypertext Transport Protocol
(HTTP).

IP address
Internet Protocol address is a numeric
address of a system used to uniquely
identify the system on the network. The

Glossary

IP address consists of four groups of numbers separated by periods (full stops).

ISQL *(Sybase specific term)*
A Sybase utility used to perform system administration tasks on Sybase SQL Server.

ITO
See **OVO**.

jukebox
See **library**.

jukebox device
A device consisting of multiple slots used to store either optical or file media. When being used to store file media, the jukebox device is known as the "file jukebox device".

Key Management Service *(Microsoft Exchange Server specific term)*
The Microsoft Exchange Server service that provides encryption functionality for enhanced security.
See also **Information Store** and **Site Replication Service**.

LBO *(EMC Symmetrix specific term)*
A Logical Backup Object (LBO) is an object of data storage/retrieval in the EMC Symmetrix environment. It is stored/retrieved by EMC Symmetrix as one entity and can only be restored as a whole.

library
Also called autochanger, jukebox, autoloader, or exchanger. A library contains media in repository slots. Each slot holds one medium (for example, DDS/DAT). Media are moved between slots and drives by a robotic mechanism, allowing random access to media. The library can contain multiple drives.

lights-out operation or **unattended operation**
A backup or restore operation that takes place outside of normal business hours without an operator. This implies that no operator personnel is present to work with the backup application or service mount requests, for example.

LISTENER.ORA *(Oracle specific term)*
An Oracle configuration file that describes one or more Transparent Network Substrate (TNS) listeners on a server.

load balancing
By default, Data Protector automatically balances the usage of devices selected for backup, so that they are used evenly. Load balancing optimizes the device usage by balancing the number of objects written to each device. Since load balancing is done automatically during backup time, you do not need to manage how the data is actually backed up. You just specify the devices to be

Glossary

used. If you do not want to use load balancing, you can select which device will be used for each object in the backup specification. Data Protector will access the devices in the specified order.

local and remote recovery
Remote recovery is performed if all Media Agent hosts specified in the SRD file are accessible. If any of them fails, the disaster recovery process fails over to the local mode. This means that the target system is searched for locally attached devices. If only one device is found, it is automatically used. Otherwise, Data Protector prompts you to select the device, which will be used for restore.

lock name
You can configure the same physical device several times with different characteristics, by using different device names.

The lock name is a user specified string that is used for locking all such device configurations to prevent collision if several such devices (device names) are used concurrently. Use an identical lock name for all device definitions which use the same physical device.

log_full shell script *(Informix UNIX specific term)*
A script provided by ON-Bar that you

can use to start backing up logical-log files when OnLine Server issues a log-full event alarm. The Informix ALARMPROGRAM configuration parameter defaults to the *<INFORMIXDIR>*/etc/log_full.sh, where *<INFORMIXDIR>* is the OnLine Server home directory. If you do not want logical logs to be backed up continuously, set the ALARMPROGRAM configuration parameter to *<INFORMIXDIR>*/etc/ no_log.sh.

logging level
The logging level determines the amount of details on files and directories written to the IDB during backup or object copying. You can always restore your data, regardless of the logging level used during backup. Data Protector provides four logging levels: Log All, Log Directories, Log Files, and No Log. The different logging level settings influence the IDB growth, backup speed, and the convenience of browsing data for restore.

logical-log files
This applies to online database backup. Logical-log files are files in which modified data is first stored before being flushed to disk. In the event of a failure, these logical-log files are used to roll forward all transactions that have been

Glossary

committed as well as roll back any transactions that have not been committed.

login ID *(MS SQL Server specific term)*
The name a user uses to log on to Microsoft SQL Server. A login ID is valid if Microsoft SQL Server has an entry for that user in the system table syslogin.

login information to the Oracle Target Database *(Oracle and SAP R/3 specific term)*
The format of the login information is <user_name>/<password>@<service>, where:

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

- <password> is a string used for data security and known only to its owner. Passwords are entered to connect to an operating system or software application. The password has to be the same as the password specified in the Oracle password file (orapwd), which is used for authentication of users performing database administration.

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

login information to the Recovery Catalog Database *(Oracle specific term)*
The format of the login information to the Recovery (Oracle) Catalog Database is <user_name>/ <password>@<service>, where the description of the user name, password, and service name is the same as in the Oracle SQL*Net V2 login information to the Oracle target database. In this case, <service> is the name of the service to the Recovery Catalog Database, not the Oracle target database.

Note that the Oracle user specified here has to be the owner of the Oracle Recovery (Oracle) Catalog.

Lotus C API *(Lotus Domino Server specific term)*
An interface for the exchange of backup and recovery information between Lotus Domino Server and a backup solution, like Data Protector.

LVM
A Logical Volume Manager is a subsystem for structuring and mapping physical disk space to logical volumes on UNIX systems. An LVM system

consists of several volume groups, where each volume group has several volumes.

Magic Packet
See **Wake ONLAN**.

mailbox (*Microsoft Exchange Server specific term*)
The location to which e-mail is delivered, which is set up by the administrator for each user. If a set of personal folders is designated as the e-mail delivery location, e-mail is routed from the mailbox to this location.

Mailbox Store (*Microsoft Exchange Server specific term*)
A part of the Information Store that maintains information about user mailboxes. A mailbox store consists of a binary rich-text .edb file and a streaming native internet content .stm file.

Main Control Unit (MCU) (*HP StorageWorks Disk Array XP specific term*)
An HP StorageWorks XP disk array that contains the primary volumes for the Continuous Access configuration and acts as a master device.
See also **BC** (*HP StorageWorks Disk Array XP specific term*), **CA** (*HP StorageWorks Disk Array XP specific term*), and **HP StorageWorks Disk Array XP LDEV**.

Manager-of-Managers (MoM)
See **Enterprise Cell Manager**.

Media Agent
A process that controls reading from and writing to a device, which reads from or writes to a medium (typically a tape). During a backup session, a Media Agent receives data from the Disk Agent and sends it to the device for writing it to the medium. During a restore session, a Media Agent locates data on the backup medium and sends it to the Disk Agent. The Disk Agent then writes the data to the disk. A Media Agent also manages the robotics control of a library.

MAPI (*Microsoft Exchange specific term*)
The MAPI (Messaging Application Programming Interface) is the programming interface that lets applications and messaging clients interact with messaging and information systems.

media allocation policy
Determines in which sequence media are used for backup. The Strict allocation policy directs Data Protector to prompt for a specific medium. The Loose policy directs Data Protector to prompt for any suitable medium. The Formatted First policy directs Data Protector to give preference to unknown media, even if unprotected media are available in the library.

Glossary

media condition
The quality of a medium as derived from the media condition factors. Heavy usage and age result in an increased number of read and write errors with tape media. Media need to be replaced when they are marked as POOR.

media condition factors
The user-assigned age threshold and overwrite threshold used to determine the state of a medium.

media ID
A unique identifier assigned to a medium by Data Protector.

media label
A user-defined identifier used to describe a medium.

media location
A user-defined physical location of a medium, such as "building 4" or "off-site storage".

media management session
A session performing some action on a medium, such as initializing, scanning the content, verifying data on a medium, or copying a medium.

media pool
A set of media of the same type (such as DDS) used and tracked as a group. Media are formatted and assigned to a media pool.

media set
The result of a backup session is data backed up on a group of media called media set. Depending on the media usage policy, several sessions can share the same media.

media type
The physical type of media, such as DDS or DLT.

media usage policy
The media usage policy controls how new backups are added to the already used media. It can be Appendable, Non-Appendable, or Appendable for incrementals only.

merging
This defines one mode to resolve file conflicts during restore. If the file to be restored already exists at the destination, the one with the more recent modification date is kept. Files not present on the disk are always restored. *See also* **overwrite**.

MFS
The Migrating File System enables a standard JFS filesystem with migration capabilities (on HP-UX 11.00). The MFS is accessed via a standard filesystem interface (DMAPI), it is mounted to a directory the same way as any HP-UX filesystem. In an MFS, only the superblock, the inode and the 'extended attribute' information remain

Glossary

permanently on the hard disk and are never migrated.
See also **VBFS**.

Microsoft Exchange Server
A "client-server" messaging and a workgroup system that offers a transparent connection to many different communication systems. It provides users with an electronic mail system, individual and group scheduling, online forms, and workflow automation tools. It provides a developer with a platform on which to build custom information-sharing and messaging-service applications.

Microsoft Management Console (MMC) *(Windows specific term)*
An administration model for Windows-based environments. It provides a simple, consistent, and integrated administration user interface allowing management of many applications through the same GUI, provided that the applications adhere to the MMC model.

Microsoft SQL Server
A database management system designed to meet the requirements of distributed "client-server" computing.

Microsoft Volume Shadow Copy service (VSS)
A software service that provides a unified communication interface to coordinate backup and restore of a VSS-aware application regardless of its specific features. This service collaborates with the backup application, writers, shadow copy providers, and the operating system kernel to implement the management of volume shadow copies and shadow copy sets.
See also **shadow copy, shadow copy provider, writer**.

mirror *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)*
See **target volume**.

mirror rotation *(HP StorageWorks Disk Array XP specific term)*
See **replica set rotation**.

MMD
The Media Management Daemon process (service) runs on the Data Protector Cell Manager and controls media management and device operations. The process is started when Data Protector is installed on the Cell Manager.

MMDB
The Media Management Database (MMDB) is a part of the IDB that contains information about media, media pools, devices, libraries, library drives, and slots configured in the cell, as well as the Data Protector media used for backup. In an enterprise backup

Glossary

environment, this part of the database can be common to all cells.
See also **CMMDB**, **CDB**.

MoM
Several cells can be grouped together and managed from a central cell. The management system of the central cell is the Manager-of-Managers (MoM). The MoM allows you to configure and manage multiple cells from a central point.

mount request
A screen prompt that tells you to insert a specific medium into a device. Once you respond to the mount request by providing the required medium and confirm the mount request, the session continues.

mount point
The access point in a directory structure for a disk or logical volume, for example /opt or d:. On UNIX, the mount points are displayed using the bdf or df command.

MSM
The Data Protector Media Session Manager, which runs on the Cell Manager and controls media sessions, such as copying media.

MU number (*HP StorageWorks Disk Array XP specific term*)
A Mirror Unit number is an integer

number (0, 1 or 2), used to indicate a first level mirror.
See also **first level mirror**.

multi-drive server
A license that allows you to run an unlimited number of Media Agents on a single system. This license, which is bound to the IP address of the Cell Manager, is no longer available.

obdrindex.dat
An IDB file with information about IDB backups, media, and devices used for the backup. This information can significantly simplify IDB recovery. It is recommended to relocate the file, together with IDB transaction logs, on a separate physical disk from other IDB directories, and, additionally, to make a copy of the file and locate it where you want.

OBDR capable device
A device that can emulate a CD-ROM drive loaded with a bootable disk and can thus be used as a backup or boot device for disaster recovery purposes.

object
See **backup object**

object copy
A copy of a specific object version that is created during an object copy session or a backup session with object mirroring.

Glossary

object copy session
A process that creates an additional copy of the backed up data on a different media set. During an object copy session, the selected backed up objects are copied from the source to the target media.

object copying
The process of copying selected object versions to a specific media set. You can select object versions from one or several backup sessions to be copied.

Object ID *(Windows specific term)*
The object IDs (OIDs) enable access to NTFS 5 files no matter where in the system the files reside. Data Protector treats the OIDs as alternate streams of the files.

object mirror
A copy of a backup object created using object mirroring. Object mirrors are often referred to as object copies.

object mirroring
The process of writing the same data to several media sets during a backup session. Data Protector enables you to mirror all or some backup objects to one or more media sets.

offline backup
A backup during which an application database cannot be used by the application.

- For simple backup methods (non ZDB), the database is generally put into a quiescent state that allows use by the backup system, but not the application, for the whole backup period (~minutes/hours). For instance, for backup to tape, until streaming of data to the tape is finished.

- For ZDB methods, the database is also put into the quiescent state, but for the period of the data replication process only (~seconds). Normal database operation can then be resumed for the rest of the backup process.

See also **zero downtime backup (ZDB)** and **online backup**.

offline recovery
Offline recovery is performed if the Cell Manager is not accessible, for example, due to network problems. Only standalone and SCSI library devices can be used for offline recovery. Recovery of the Cell Manager is always offline.

offline redo log
See **archived redo log**

OmniStorage
Software providing transparent migration of less frequently used data to the optical library while keeping more

Glossary

frequently used data on the hard disk. HP OmniStorage runs on HP-UX systems.

On-Bar *(Informix specific term)*
A backup and restore system for OnLine Server. ON-Bar enables you to create a copy of your OnLine Server data and later restore the data. The ON-Bar backup and restore system involves the following components:

- onbar utility

- Data Protector, as the backup solution

- XBSA interface

- ON-Bar catalog tables, which are used to back up dbobjects and track instances of dbobjects through multiple backups.

onbar utility *(Informix specific term)*
The Informix utility that communicates backup and restore requests to OnLine Server. The utility uses XBSA to exchange control data and back up and restore data with Data Protector.

ONCONFIG *(Informix specific term)*
An environment variable that specifies the name of the active ONCONFIG configuration file. If the ONCONFIG environment variable is not present, OnLine uses the configuration values

from the file *<INFORMIXDIR>*/etc/ onconfig (on HP-UX) or *<INFORMIXDIR>*\etc\onconfig (on Windows).

online backup
A backup performed while a database application remains available for use. The database is placed into a special backup mode of operation for the time period that the backup application requires access to the original data objects. During this period, the database is fully operational, but there may be a small performance impact and log files may grow very quickly.

- For simple backup methods (non ZDB), backup mode is required for the whole backup period (~minutes/ hours). For instance, for backup to tape, until streaming of data to tape is finished.

- For ZDB methods, backup mode is required for the short period of the data replication process only (~seconds). Normal database operation can then be resumed for the rest of the backup process.

In some cases, transaction logs may also have to be backed up to allow a consistent database to be restored. *See also* **zero downtime backup (ZDB)** and **offline backup**.

Glossary

online redo log *(Oracle specific term)*
Redo logs that have not been archived, but are either available to the instance for recording database activity or are filled and waiting to be archived or reused.
See also **archived redo log**.

OnLine Server *(Informix specific term)*
Refers to INFORMIX-OnLine Dynamic Server.

OpC
See **OVO**.

Oracle instance *(Oracle specific term)*
Each installation of an Oracle database on one or more systems. One computer system can have several instances of a database running.

ORACLE_SID *(Oracle specific term)*
A unique name for an Oracle Server instance. To switch among Oracle Servers, specify the desired *<ORACLE_SID>*. The *<ORACLE_SID>* is included in the CONNECT DATA parts of the connect descriptor in a TNSNAMES.ORA file and in the definition of the TNS listener in the LISTENER.ORA file.

original system
The system configuration backed up by Data Protector before a computer disaster hits the system.

overwrite
An option that defines one mode to resolve file conflicts during restore. All files are restored from a backup even if they are older than existing files.
See also **merging**.

OVO
HP OpenView Operations for Unix provides powerful capabilities for operations management of a large number of systems and applications on in a network. Data Protector provides an integration into this management product. This integration is implemented as a SMART Plug-In for OVO management servers on HP-UX and Solaris. Earlier versions of OVO were called IT/Operation, Operations Center and Vantage Point Operations.
See also **merging**.

ownership
The ownership of a backup determines who can restore from the backup. The user who starts an interactive backup is the session owner. If a user starts an existing backup specification without modifying it, the session is not considered interactive. In that case, if the backup owner has been defined in the backup specification, they remain the session owner. Otherwise, the session owner becomes the user who started the backup in question. For the scheduled backups, by default, the session owner is for the UNIX Cell

Glossary

Manager: root.sys@<Cell Manager>, and for the Windows Cell Manager, the user that was specified during the installation of the Cell Manager. It is possible to modify the ownership, so that the specific user becomes the session owner.

P1S file
P1S file contains information on how to format and partition all disks installed in the system during Enhanced Automated Disaster Recovery (EADR). It is created during full backup and is saved on backup medium and on Cell Manager into
<*Data_Protector_home*>\Config\Server\dr\p1s directory on a Windows Cell Manager or in /etc/opt/omni/server/dr/p1s directory on a UNIX Cell Manager with the filename recovery.p1s.

package (*MC/ServiceGuard and Veritas Cluster specific term*)
A collection of resources (for example volume groups, application services, IP names and addresses) that are needed to run a specific cluster-aware application.

pair status (*HP StorageWorks Disk Array XP specific term*)
A mirrored pair of disks can have various status values depending on the action performed on it. The three most important status values are:

- COPY - The mirrored pair is currently resynchronizing. Data is transferred from one disk to the other. The disks do not contain the same data.

- PAIR - The mirrored pair is completely synchronized and both disks (the primary volume and the mirrored volume) contain identical data.

- SUSPENDED - The link between the mirrored disks is suspended. That means that both disks are accessed and updated independently. However, the mirror relationship is still maintained and the pair can be resynchronized without transferring the complete disk.

parallel restore
Restoring backed up data to multiple disks at the same time (that is, in parallel) by running multiple Disk Agents, that receive data from one Media Agent. For the parallel restore to work, select data that is located on different disks or logical volumes and during backup, the data from the different objects must have been sent to the same device using a concurrency of 2 or more. During a parallel restore, the data for multiple objects selected for restore is read from media at the same time, thereby improving performance.

Glossary

parallelism
The concept of reading multiple data streams from an online database.

physical device
A physical unit that contains either a drive or a more complex unit such as a library.

post-exec
A backup option that executes a command or script after the backup of an object or after the entire session completes. Post-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows and as shell scripts on UNIX.
See also **pre-exec**.

pre- and post-exec commands
Pre- and post-exec commands are used to perform additional action before and after a backup or restore session. They are not supplied by Data Protector. You need to create your own commands. They can be written as executables or batch files on Windows and as shell scripts on UNIX.

prealloc list
A subset of media in a media pool that specifies the order in which media are used for backup.

pre-exec
A backup option that executes a command or script before the backup of an object or before the entire session is started. Pre-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows and as shell scripts on UNIX.
See also **post-exec**.

Primary Volume (P-VOL) *(HP StorageWorks Disk Array XP specific term)*
Standard HP StorageWorks Disk Array XP LDEVs that act as a primary volume for the CA and BC configurations. The P-VOL is located in the MCU.
See also **Secondary Volume (S-VOL)**.

protection
See **data protection** and also **catalog protection**.

public folder store (*Microsoft Exchange Server specific term)*
The part of the Information Store that maintains information in public folders. A public folder store consists of a binary rich-text .edb file and a streaming native internet content .stm file.

public/private backed up data
When configuring a backup, you can select whether the backed up data will be:

Glossary

- public, that is visible (and accessible for restore) to all Data Protector users

- private, that is, visible (and accessible for restore) only to the owner of the backup and administrators

RAID
Redundant Array of Inexpensive Disks.

RAID Manager Library *(HP StorageWorks Disk Array XP specific term)*
The RAID Manager Library is used internally by Data Protector on Solaris systems to allow access to HP StorageWorks Disk Array XP configuration, status, and performance data and to key HP StorageWorks Disk Array XP features through the use of function calls translated into a sequence of low level SCSI commands.

RAID Manager XP *(HP StorageWorks Disk Array XP specific term)*
The RAID Manager XP application provides an extensive list of commands to report and control the status of the CA and BC applications. The commands communicate through a RAID Manager instance with the HP StorageWorks Disk Array XP Disk Control Unit. This instance translates the commands into a sequence of low level SCSI commands.

rawdisk backup
See **disk image backup**.

RCU *(HP StorageWorks specific term)*
The Remote Control Unit acts as a slave of an MCU in a CA configuration. In bidirectional configurations, the RCU can act as an MCU.

RDBMS
Relational Database Management System.

RDF1/RDF2 *(EMC Symmetrix specific term)*
A type of SRDF device group. Only RDF devices can be assigned to an RDF group. An RDF1 group type contains source (R1) devices and an RDF2 group type contains target (R2) devices.

RDS
The Raima Database Server process (service) runs on the Data Protector Cell Manager and manages the IDB. The process is started when Data Protector is installed on the Cell Manager.

Recovery Catalog *(Oracle specific term)*
A set of Oracle tables and views that are used by Recovery Manager to store information about Oracle databases. This information is used by Recovery Manager to manage the backup, restore,

Glossary

and recovery of Oracle databases. The recovery catalog contains information about:

- The physical schema of the Oracle target database

- Data file and archived log backup sets

- Data file copies

- Archived Redo Logs

- Stored scripts.

Recovery Catalog Database *(Oracle specific term)*
An Oracle database that contains a recovery catalog schema. You should not store the recovery catalog in your target database.

RecoveryInfo
When backing up Windows configuration files, Data Protector collects the information about the current system configuration (information on disk layout, volume, and network configuration). This information is needed for disaster recovery.

Recovery Manager (RMAN) *(Oracle specific term)*
An Oracle command-line interface that directs an Oracle Server process to back

up, restore, or recover the database it is connected to. RMAN uses either the recovery catalog or the control file to store information about backups. This information can be used later in restore sessions.

recycle
A process that removes the data protection from all backed up data on a medium, allowing Data Protector to overwrite it during one of the next backups. Data that belongs to the same session(s) but resides on other media is also unprotected. Recycling does not actually alter the data on the medium.

redo log *(Oracle specific term)*
Every Oracle database has a set of two or more redo log files. The set of redo log files for the database is known as the database's redo log. Oracle uses the redo log to record all changes made to data.

Remote Control Unit *(HP StorageWorks Disk Array XP specific term)*
The Remote Control Unit (RCU) acts as a slave of an MCU in a CA configuration. In bidirectional configurations, the RCU can act as an MCU.

Removable Storage Management Database *(Windows specific term)*
A Windows service used for managing removable media (such as tapes and

Glossary

disks) and storage devices (libraries). Removable Storage allows applications to access and share the same media resources.

reparse point *(Windows specific term)*
A system-controlled attribute that can be associated with any directory or file. The value of a reparse attribute can have user-controlled data. The format of the data is understood by the application that stored the data and a filesystem filter that was installed to interpret the data and process such files. Whenever the filesystem encounters a file with a reparse point, it attempts to find the filesystem filter associated with the data format.

replica *(ZDB specific term)*
An image, at a particular point in time, of the data in source volumes that contain user-specified backup objects. Depending on the hardware/software with which it is created, the image may be an independent exact duplicate (clone) of the storage blocks at the physical disk level (for example, a split mirror), or a virtual copy (for example, a snapshot). From a host's perspective, on a basic UNIX or Windows system, the complete physical disk containing a backup object is replicated. However, if a volume manager is used on UNIX, the whole volume/disk group containing a

backup object is replicated.
See also **snapshot**, **snapshot creation**, **split mirror**, and **split mirror creation**.

replica set *(ZDB specific term)*
A group of replicas, all created using the same backup specification.
See also **replica** and **replica set rotation**.

replica set rotation *(ZDB specific term)*
The use of a replica set for regular backup production: Each time the same backup specification requiring the use of a replica set is run, a new replica is created and added to the set, until the maximum number for the set is reached. After that, the oldest replica in the set is replaced, maintaining the maximum number of replicas in the set.
See also **replica** and **replica set**.

restore session
A process that copies data from backup media to a client.

RMAN *(Oracle specific term)*
See **Recovery Manager**.

RSM
The Data Protector Restore Session Manager controls the restore session. This process always runs on the Cell Manager system.

Glossary

RSM *(Windows specific term)*
Removable Storage Manager (RSM) includes a media management service that facilitates communication among applications, robotic changers, and media libraries. It enables multiple applications to share local robotic media libraries and tape or disk drives and to manage removable media.

SAPDBA *(SAP R/3 specific term)*
An SAP R/3 user interface that integrates the BRBACKUP, BRARCHIVE, and BRRESTORE tools.

scan
A function that identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library).

scanning
A function which identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library). It is useful to perform a scan and check the actual media in the device if someone has manually manipulated media without using Data Protector to eject or enter, for example.

Scheduler
A function that controls when and how often automatic backups occur. By setting up a schedule, you automate the start of backups.

Secondary Volume (S-VOL) *(HP StorageWorks Disk Array XP specific term)*
Secondary Volumes, or S-VOLs, are XP LDEVs that act as a secondary CA or BC mirror of another LDEV (a P-VOL). In the case of CA, S-VOLs can be used as failover devices in a MetroCluster configuration. The S-VOLs are assigned separate SCSI addresses, different from the addresses used by the P-VOLs.
See also **Primary Volume (P-VOL).**

session
See **backup session**, **media management session,** and **restore session**.

session ID
An identifier of a backup, restore, object copy, or media management session, consisting of the date when the session ran and a unique number.

session key
This environment variable for the Pre- and Post-exec script is a Data Protector unique identification of any session, including preview sessions. The session key is not recorded in the database, and

Glossary

it is used for specifying options for the omnimnt, omnistat and omniabort CLI commands.

shadow copy *(MS VSS specific term)* A volume that represents a duplicate of the original volume at a certain point in time. The data is then backed up from the shadow copy and not from the original volume. The original volume continues to 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 specific term)* The Microsoft Exchange Server 2000/ 2003 service that permits compatibility with Microsoft Exchange Server 5.5 by emulating the Exchange Server 5.5 directory service.
See also **Information Store** and **Key Management Service**.

slot A mechanical position in a library. Each slot can hold a medium, such as a DLT tape. Data Protector references each slot with a number. To read a medium, a robotic mechanism moves the medium from a slot into the drive.

SMB *See* **split mirror backup**.

SMBF The Session Messages Binary Files (SMBF) part of the IDB stores session messages generated during backup, object copy, restore, and media

Glossary

management sessions. One binary file is created per session. The files are grouped by year and month.

snapshot *(HP StorageWorks VA and HP StorageWorks EVA specific term)*
A form of replica produced using snapshot creation techniques. A range of snapshot types is available, with different characteristics, depending on the arrays/techniques used. Such replicas are dynamic and may be either virtual copies, still reliant upon the contents of the source volumes, or independent exact duplicates (clones), depending on the snapshot type and the time since creation.
See also **replica** and **snapshot creation**.

snapshot backup *(HP StorageWorks VA and HP StorageWorks EVA specific term)*
See **ZDB to tape**, **ZDB to disk**, and **ZDB to disk+tape**.

snapshot creation *(HP StorageWorks VA and HP StorageWorks EVA specific term)*
A replica creation technique, in which copies of source volumes are created using storage virtualization techniques. The replicas are considered to be created at one particular point-in-time, without pre-configuration, and are immediately available for use. However background copying processes normally continue after creation.
See also **snapshot**.

source (R1) device *(EMC Symmetrix specific term)*
An EMC Symmetrix device that participates in SRDF operations with a target (R2) device. All writes to this device are mirrored to a target (R2) device in a remote EMC Symmetrix unit. An R1 device must be assigned to an RDF1 group type.
See also **target (R2) device**.

source volume *(ZDB specific term)*
A storage volume containing data to be replicated.

sparse file A file that contains data with portions of empty blocks. Examples are: -A matrix in which some or much of the data contains zeros -files from image applications -high-speed databases If sparse file processing is not enabled during restore, it might be impossible to restore this file.

split mirror *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)*
A replica created using split mirror techniques. Such a replica provides an independent, exact duplicate, or clone, of the contents of the source volumes.
See also **replica** and **split mirror creation**.

Glossary

split mirror backup *(EMC Symmetrix specific term)*
See **ZDB to tape**.

split mirror backup *(HP StorageWorks Disk Array XP specific term)*
See **ZDB to tape**, **ZDB to disk**, and **ZDB to disk+tape**.

split mirror creation *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)*
A replica creation technique, in which a pre-configured set of target volumes (a mirror) is kept synchronized with a set of source volumes until the time at which a replica of the contents of the source volumes is required. Then, the synchronization is stopped (the mirror is split) and a split mirror replica of the source volumes at the time of the split remains in the target volumes.
See also **split mirror**.

split mirror restore *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)*
A process in which data backed up in a ZDB-to-tape or a ZDB-to-disk+tape session is restored from tape media to a split mirror replica, which is then synchronized to the source volumes. Individual backup objects or complete sessions can be restored using this method.
See also **ZDB to tape**, **ZDB to disk+tape**, and **replica**.

sqlhosts file *(Informix specific term)*
An Informix connectivity-information file that contains the names of each of the database servers and any aliases to which the clients on a host computer can connect.

SRD file
The Data Protector System Recovery Data (SRD) file contains system information required for installing and configuring the operating system in case of a disaster. The SRD file is an ASCII file, generated when a CONFIGURATION backup is performed on a Windows client and stored on the Cell Manager.

SRDF *(EMC Symmetrix specific term)*
The EMC Symmetrix Remote Data Facility is a business continuation process that enables effective, real-time data replication of SLDs between dislocated processing environments. These environments could be situated within the same root computer environment or separated by long distances.

SSE Agent *(HP StorageWorks Disk Array XP specific term)*
A Data Protector software module that executes all tasks required for a split mirror backup integration. It communicates with the HP StorageWorks Disk Array XP storing system using the RAID Manager XP

Glossary

utility (HP-UX and Windows systems) or RAID Manager Library (Solaris systems).

sst.conf file
The file /usr/kernel/drv/sst.conf is required on each Data Protector Sun Solaris client to which a multi-drive library device is connected. It must contain an entry for the SCSI address of the robotic mechanism of each library device connected to the client.

st.conf file
The file /kernel/drv/st.conf is required on each Data Protector Solaris client with a backup device connected. It must contain device information and a SCSI address for each backup drive connected to that client. A single SCSI entry is required for a single-drive device and multiple SCSI entries are required for a multi-drive library device.

stackers
Devices with multiple slots for media storage usually with only one drive. A stacker selects media from the stack sequentially. In contrast, a library can randomly select media from its repository.

standalone file device
A file device is a file in a specified directory to which you back up data.

standard security *(MS SQL specific term)*
Standard security uses the login validation process of the Microsoft SQL Server for all connections. Standard security is useful in network environments with a variety of clients, some of which may not support trusted connections. It also provides backward compatibility for older versions of the Microsoft SQL Server.
See also **integrated security**.

Storage Group
(Microsoft Exchange Server specific term)
A collection of databases (stores) that share a common set of transaction log files. Exchange manages each storage group with a separate server process.

StorageTek ACS library
(StorageTek specific term)
Automated Cartridge System is a library system (also known as Silo) consisting of one Library Management Unit (LMU) and one to 24 Library Storage Modules (LSM) connected to the unit.

storage volume *(ZDB specific term)*
A storage volume represents an object that may be presented to an operating system or some other entity (for example, a virtualization mechanism) upon which volume management systems, file systems, or other objects may exist. The volume management

Glossary

systems, file systems are built on this storage. Typically, these can be created or exist within a storage system such as a disk array.

switchover
See **failover**

Sybase Backup Server API *(Sybase specific term)*
An industry-standard interface developed for the exchange of backup and recovery information between a Sybase SQL Server and a backup solution like Data Protector.

Sybase SQL Server *(Sybase specific term)*
The server in the Sybase "client-server" architecture. Sybase SQL Server manages multiple databases and multiple users, keeps track of the actual location of data on disks, maintains mapping of logical data description to physical data storage, and maintains data and procedure caches in memory.

Symmetrix Agent (SYMA) *(EMC Symmetrix specific term)*
The Data Protector software module that prepares the EMC Symmetrix environment for backup and restore operations.

System Backup to Tape *(Oracle specific term)*
An Oracle interface that handles the actions required to load, label, and unload correct backup devices when Oracle issues a backup or restore request.

system databases *(Sybase specific term)*
The four system databases on a newly installed Sybase SQL Server are the:

- master database (master)

- temporary database (tempdb)

- system procedure database (sybsystemprocs)

- model database (model).

system disk
A system disk is a disk containing operating system files. Microsoft terminology defines the system disk as a disk containing the files required for initial step of boot process.

system partition
A system partition is a partition containing operating system files. Microsoft terminology defines a system partition as a partition containing the files required for initial step of boot process.

System State *(Windows specific term)*
The System State data comprises the Registry, COM+ Class Registration

database, system startup files, and the Certificate Services database (if the server is a certificate server). If the server is a domain controller, Active Directory directory services and the Sysvol directory are also contained in the System State data. If the server is running the Cluster service, the System State data also includes resource registry checkpoints and the quorum resource recovery log, which contains the most recent cluster database information.

system volume/disk/partition
A volume/disk/partition containing operating system files. Microsoft terminology defines the system volume/disk/partition as the volume/disk/partition containing files required for the initial step of the boot process.

SysVol (*Windows specific term*)
A shared directory that stores the server copy of the domain's public files, which are replicated among all domain controllers in the domain.

tablespace
A part of a database structure. Each database is logically divided into one or more tablespaces. Each tablespace has data files or raw volumes exclusively associated with it.

tapeless backup *(ZDB specific term)*
See **ZDB to disk**.

target database *(Oracle specific term)*
In RMAN, the target database is the database that you are backing up or restoring.

target (R2) device *(EMC Symmetrix specific term)*
An EMC Symmetrix device that participates in SRDF operations with a source (R1) device. It resides in the remote EMC Symmetrix unit. It is paired with a source (R1) device in the local EMC Symmetrix unit and receives all write data from its mirrored pair. This device is not accessed by user applications during normal I/O operations. An R2 device must be assigned to an RDF2 group type.
See also **source (R1) device**

target system *(Disaster Recovery specific term)*
A system after a computer disaster has occurred. The target system is typically in a non-bootable state and the goal of disaster recovery is to restore this system to the original system configuration. The difference between a crashed system and a target system is that a target system has all faulty hardware replaced.

target volume *(ZDB specific term)*
A storage volume to which data is replicated.

Glossary

Terminal Services *(Windows specific term)*
Windows Terminal Services provide a multi-session environment that allows clients to access a virtual Windows desktop session and Windows-based programs running on the server.

thread *(MS SQL Server specific term)*
An executable entity that belongs to only one process. It comprises a program counter, a user-mode stack, a kernel-mode stack, and a set of register values. Several threads can run at a time within one process.

TimeFinder *(EMC Symmetrix specific term)*
A business continuation process that creates an instant copy of single or multiple Symmetrix Logical Devices (SLDs). The instant copy is created on specially preconfigured SLDs called BCVs and is accessible via a separate device address to the system(s).

TLU
Tape Library Unit.

TNSNAMES.ORA *(Oracle and SAP R/3 specific term)*
A network configuration file that contains connect descriptors mapped to service names. The file may be maintained centrally or locally, for use by all or individual clients.

transaction
A mechanism for ensuring that a set of actions is treated as a single unit of work. Databases use transactions to keep track of database changes.

transaction backup
Transaction backups generally use fewer resources than database backups, so they can be created more frequently than database backups. By applying transaction backups, you can recover the database to a specific point in time prior to when a problem occurred.

transaction backup *(Sybase and SQL specific term)*
A backup of the transaction log providing a record of changes made since the last full or transaction backup.

transaction log backup
Transaction log backups generally use fewer resources than database backups so they can be created more frequently than database backups. By applying transaction log backups, you can recover the database to a specific point in time.

transaction log files
Files that record transactions of the database modifications, and provide fault tolerance in case of a database disaster.

Glossary

transaction logs (*Data Protector specific term*)
Keeps track of IDB changes. The archiving of transaction logs should be enabled to prevent you from losing the transaction log files that are created after the last IDB backup and are necessary for IDB recovery.

transaction log table (*Sybase specific term*)
A system table in which all changes to the database are automatically recorded.

transportable snapshot (*MS VSS specific term*)
A shadow copy that is created on the application system and can be presented to the backup system which performs the backup.
See also **Microsoft Volume Shadow Copy service (VSS)**.

TSANDS.CFG file (*Novell NetWare specific term*)
A file that allows you to specify the names of containers where you want backups to begin. It is text file located in the SYS:SYSTEM\TSA directory on the server where TSANDS.NLM is loaded.

unattended operation
See **lights-out operation**.

user account
You can use Data Protector only if you have a Data Protector user account, which restricts unauthorized access to Data Protector and to backed up data. Data Protector administrators create this account specifying a user logon name, the systems from which the user can log on, and a Data Protector user group membership. This is checked whenever the user starts the Data Protector user interface or performs specific tasks.

user disk quotas
NTFS quota-management support enables an enhanced tracking mechanism and control over disk space usage on shared storage volumes. Data Protector backs up user disk quotas on the whole system and for all configured users at a time.

user group
Each Data Protector user is member of a User Group. Each User Group has a set of user rights that are given to every user in that User Group. The number of User Groups with their associated user rights can be defined as desired. Data Protector provides three default user groups: admin, operator, and user.

user profile (*Windows specific term*)
Configuration information retained on a user basis. This information includes desktop settings, screen colors, network connections, and so on. When the user logs on, the user profile is loaded and the Windows environment is set accordingly.

Glossary

user rights
User rights or access rights are the permissions needed to perform specific Data Protector tasks. Configuring a backup, starting a backup session, or starting a restore session are typical user rights. Users have the access rights of the user group to which they belong.

vaulting media
The process of storing media to a safe and remote place. The media are brought back to the data center when they are needed for restore or are ready for reuse in subsequent backups. The vaulting procedure depends on your company's backup strategy and policies for data protection/reliability.

VBFS *(OmniStorage specific term)*
A Very Big File System is an extension of the standard HP-UX file system on HP-UX 9.x. It is mounted to a directory the same way as any HP-UX file system. In a VBFS, only the superblock, the inode and the 'extended attribute' information remain permanently on the hard disk and are never migrated.
See also **MFS**.

verify
A function that lets you check whether the Data Protector data on a specified medium is readable. Additionally, consistency within each block can be checked if the backup was performed with the cyclic redundancy check (CRC) option ON.

Virtual Controller Software (VCS)
(HP StorageWorks EVA specific term)
The firmware that manages all aspects of storage system operation, including communication with Command View EVA through the HSV controllers.
See also **Command View (CV) EVA**.

Virtual Device Interface *(MS SQL Server specific term)*
This is a SQL Server programming interface that allows fast backup and restore of large databases.

virtual disk *(HP StorageWorks EVA specific term)*
A unit of storage allocated from an HP StorageWorks Enterprise Virtual Array storage pool. Virtual disks are the entities that are replicated using the HP StorageWorks Enterprise Virtual Array snapshot functionality.
See also **source volume** and **target volume**.

virtual server
A virtual machine in a cluster environment defined in a domain by a network IP name and address. Its address is cached by the cluster software and mapped to the cluster node that is currently running the virtual server

Glossary

resources. This way all requests for a particular virtual server are cached by a specific cluster node.

volser *(ADIC and STK specific term)*
A VOLume SERial number is a label on the medium to identify the physical tape used in very large libraries. A volser is a naming convention specific to ADIC/GRAU and StorageTek devices.

volume group
A unit of data storage in an LVM system. A volume group can consist of one or more physical volumes. There can be more than one volume group on the system.

volume mountpoint (*Windows specific term*)
An empty directory on a volume that can be used to mount another volume. The volume mount point acts as a gateway to the target volume. Provided that the volume is mounted, users and applications can refer to the data on the mounted volume by the full (merged) filesystem path as if both volumes are one integral part.

Volume Shadow Copy service
See **Microsoft Volume Shadow Copy service**.

VPO
See **OVO**.

VSS
See **Microsoft Volume Shadow Copy service**.

VxFS
Veritas Journal Filesystem.

VxVM (Veritas Volume Manager)
A Veritas Volume Manager is a system for managing disk space on Solaris platforms. A VxVM system consists of an arbitrary group of one or more physical volumes organized into logical disk groups.

Wake ONLAN
Remote power-up support for systems running in power-save mode from some other system on the same LAN.

Web reporting
The Data Protector functionality that allows you to view reports on backup status and Data Protector configuration using the Web interface.

wildcard character
A keyboard character that can be used to represent one or many characters. The asterisk (*), for example, typically represents one or more characters, and the question mark (?) typically represents a single character. Wildcard characters are often used in operating systems as a means of specifying more than one file by name.

Glossary

Windows CONFIGURATION backup
Data Protector allows you to back up Windows CONFIGURATION, including Windows Registry, user profiles, Event Logs, and WINS and DHCP server data (if configured on a system) in one step.

Windows Registry
A centralized database used by Windows to store configuration information for the operating system and the installed applications.

WINS server A system running Windows Internet Name Service software that resolves Windows networking computer names to IP addresses. Data Protector can back up WINS server data as part of the Windows configuration.

writer
(MS VSS specific term)
A process that initiates change of data on the original volume. Writers are typically applications or system services that write persistent information on a volume. Writers also participate in the shadow copy synchronization process by assuring data consistency.

XBSA interface *(Informix specific term)*
The onbar utility and Data Protector communicate with each other through the X/Open Backup Specification Services Programmer's Interface (XBSA).

XCopy engine *(direct backup specific term)*
A SCSI-3 copy command that allows you to copy data from a storage device having a SCSI source address to a backup device having a SCSI destination address, thus enabling direct backup. The data flows from a source device (either block or streaming, that is, disk or tape) to the destination device (either block or streaming) through XCopy. This releases the controlling server of reading the data from the storage device into memory and then writing the information to the destination device.
See also **direct backup**.

ZDB
See **zero downtime backup (ZDB)**.

ZDB database *(ZDB specific term)*
A part of the IDB, storing ZDB related information such as source volumes, replicas and security information. The ZDB database is used for ZDB, instant recovery, and split mirror restore.
See also **zero downtime backup (ZDB)**.

ZDB to disk *(ZDB specific term)*
A form of zero downtime backup where the replica produced is kept on the disk

Glossary

array as a backup of the source volumes at a specific point in time. Multiple replicas, produced using the same backup specification at different times, can be kept in a replica set. A replica from ZDB to disk can be restored using the instant recovery process.
See also **zero downtime backup (ZDB)**, **ZDB to tape**, **ZDB to disk+tape**, **instant recovery**, and **replica set rotation**.

ZDB to disk+tape *(ZDB specific term)*
A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time, in the same way as with ZDB to disk. However, data in the replica is also streamed to a backup medium, as with ZDB to tape. If this backup method is used, data backed up in the same session can be restored using the instant recovery process, the standard Data Protector restore from tape, or on split mirror arrays, split mirror restore.
See also **zero downtime backup (ZDB)**, **ZDB to disk**, **ZDB to tape**, **instant recovery**, **replica**, and **replica set rotation**.

ZDB to tape *(ZDB specific term)*
A form of zero downtime backup where data in the replica produced is streamed to a backup medium, typically tape. Instant recovery is not possible from such a backup, so the replica need not be

retained on the disk array after backup completion. The backed-up data can be restored using standard Data Protector restore from tape. On split mirror arrays, split mirror restore can also be used.
See also **zero downtime backup (ZDB)**, **ZDB to disk**, **instant recovery**, **ZDB to disk+tape**, and **replica**.

zero downtime backup (ZDB)
A backup approach in which data replication techniques, provided by a disk array, are used to minimize the impact of backup operations on an application system. A replica of the data to be backed up is created first. All subsequent backup operations are performed on the replicated data rather than the original data, while the application system can be returned to normal operation.
See also **ZDB to disk**, **ZDB to tape**, **ZDB to disk+tape**, and **instant recovery**.

Glossary

Index

Index

Index

Index

Index