HP Data Protector 6.20
Zero Downtime Backup Administrator's Guide

HP Part Number: N/A
Published: December 2011
Edition: Third

Contents

Publication history

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

Table 1 Edition history

Part number	Guide edition	Product
B6960-96012	July 2006	Data Protector Release A.06.00
B6960-96046	November 2008	Data Protector Release A.06.10
B6960-90162	September 2009	Data Protector Release A.06.11
N/A	March 2011	Data Protector Release 6.20
N/A	December 2011	Data Protector Release 6.20 with any of the following patches: DPWIN_00551, PHSS_42652, DPSOL_00477, DPLNX_00183
N/A	December 2011 (third edition)	Data Protector Release 6.20 with any of the following patches: DPWIN_00551, PHSS_42652, DPSOL_00477, DPLNX_00183

About this guide

This guide provides information about:

- configuring a disk array integration
- using Data Protector ZDB integrations for backing up data

Intended audience

This guide is intended for backup administrators and operators with knowledge of:

- Disk arrays (HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, HP P4000 SAN Solutions, EMC Symmetrix)
- Basic operating system commands and utilities

Documentation set

Other documents and online Help provide related information.

Guides

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

You can find these documents from the Manuals page of the HP Information Management Digital Hub website:

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

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

- *HP Data Protector Concepts Guide*

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

- *HP Data Protector Installation and Licensing Guide*

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

- *HP Data Protector Troubleshooting Guide*

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

- *HP Data Protector Disaster Recovery Guide*

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

- *HP Data Protector Integration Guides*

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

  - *HP Data Protector Integration Guide for Microsoft Applications: SQL Server, SharePoint Server, and Exchange Server*

    This guide describes the integrations of Data Protector with the following Microsoft applications: Microsoft SQL Server, Microsoft SharePoint Server, and Microsoft Exchange Server.

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

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

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

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

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

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

  - *HP Data Protector Integration Guide for Virtualization Environments*

    This guide describes the integrations of Data Protector with virtualization environments: VMware Virtual Infrastructure and VMware vSphere, Microsoft Hyper-V, and Citrix XenServer.

  - *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*

    This guide describes the integration of Data Protector with the Microsoft Volume Shadow Copy Service. This guide also documents application writer specifics.

- *HP Data Protector Integration Guide for HP Operations Manager for UNIX*

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

- *HP Data Protector Integration Guide for HP Operations Manager for Windows*

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

- *HP Data Protector Zero Downtime Backup Concepts Guide*

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

- *HP Data Protector Zero Downtime Backup Administrator's Guide*

  This guide describes how to configure and use the integration of Data Protector with HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, HP P4000 SAN Solutions, and EMC Symmetrix Remote Data Facility and TimeFinder. It is intended for backup administrators or operators. It covers the zero downtime backup, instant recovery, and the restore of filesystems and disk images.

- *HP Data Protector Zero Downtime Backup Integration Guide*

  This guide describes how to configure and use Data Protector to perform zero downtime backup, instant recovery, and standard restore of Oracle Server, SAP R/3, Microsoft Exchange Server, and Microsoft SQL Server databases.

- *HP Data Protector Granular Recovery Extension User Guide for Microsoft SharePoint Server*

  This guide describes how to configure and use the Data Protector Granular Recovery Extension for Microsoft SharePoint Server. The Data Protector Granular Recovery Extension is integrated into Microsoft SharePoint Server Central Administration and enables you to recover individual items. This guide is intended for Microsoft SharePoint Server administrators and Data Protector backup administrators.

- *HP Data Protector Granular Recovery Extension User Guide for VMware vSphere*

  This guide describes how to configure and use the Data Protector Granular Recovery Extension for VMware vSphere. The Data Protector Granular Recovery Extension is integrated into VMware vCenter Server and enables you to recover individual items. This guide is intended for VMware vCenter Server users and Data Protector backup administrators.

- *HP Data Protector Media Operations User Guide*

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

- *HP Data Protector Product Announcements, Software Notes, and References*

  This guide gives a description of new features of HP Data Protector 6.20. It also provides information on installation requirements, required patches, and limitations, as well as known issues and workarounds.

- *HP Data Protector Product Announcements, Software Notes, and References for Integrations to HP Operations Manager*

  This guide fulfills a similar function for the HP Operations Manager integration.

- *HP Data Protector Media Operations Product Announcements, Software Notes, and References*

  This guide fulfills a similar function for Media Operations.

- *HP Data Protector Command Line Interface Reference*

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

Online Help

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

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

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

Documentation map

Abbreviations

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

Abbreviation	Guide
CLI	Command Line Interface Reference
Concepts	Concepts Guide
DR	Disaster Recovery Guide
GS	Getting Started Guide
GRE-SPS	Granular Recovery Extension User Guide for Microsoft SharePoint Server
GRE-VMware	Granular Recovery Extension User Guide for VMware vSphere
Help	Online Help
IG-IBM	Integration Guide for IBM Applications: Informix, DB2, and Lotus Notes/Domino
IG-MS	Integration Guide for Microsoft Applications: SQL Server, SharePoint Server, and Exchange Server
IG-O/S	Integration Guide for Oracle and SAP
IG-OMU	Integration Guide for HP Operations Manager for UNIX
IG-OMW	Integration Guide for HP Operations Manager for Windows
IG-Var	Integration Guide for Sybase, Network Node Manager, and Network Data Management Protocol Server
IG-VirtEnv	Integration Guide for Virtualization Environments
IG-VSS	Integration Guide for Microsoft Volume Shadow Copy Service
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
PA	Product Announcements, Software Notes, and References
Trouble	Troubleshooting Guide
ZDB Admin	ZDB Administrator's Guide
ZDB Concept	ZDB Concepts Guide
ZDB IG	ZDB Integration Guide

Map

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

								Integration Guides								ZDB			GRE		MO			
	Help	GS	Concepts	Install	Trouble	DR	PA	MS	O/S	IBM	Var	VSS	VirtEnv	OMU	OMW	Concept	Admin	IG	SPS	VMware	GS	User	PA	CLI
Backup	X	X	X					X	X	X	X	X	X			X	X	X						
CLI																								X
Concepts/ techniques	X		X					X	X	X	X	X	X	X	X	X	X	X	X	X				
Disaster recovery	X		X			X																		
Installation/ upgrade	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					X	
New features	X						X																X	
Planning strategy	X		X													X								
Procedures/ tasks	X			X	X	X		X	X	X	X	X	X	X	X		X	X	X	X		X		
Recommendations			X				X									X							X	
Requirements				X			X	X	X	X	X	X	X	X	X						X	X	X	
Restore	X	X	X					X	X	X	X	X	X				X	X	X	X				
Supported configurations																X								
Troubleshooting	X			X	X			X	X	X	X	X	X	X	X		X	X	X	X				

Integrations

Look in these guides for details of the integrations with the following software applications:

Software application	Guides
HP Network Node Manager (NNM)	IG-Var
HP Operations Manager	IG-OMU, IG-OMW
IBM DB2 UDB	IG-IBM
Informix Server	IG-IBM
Lotus Notes/Domino Server	IG-IBM
Media Operations	MO User
Microsoft Exchange Server	IG-MS, ZDB IG
Microsoft Hyper-V	IG-VirtEnv
Microsoft SharePoint Server	IG-MS, ZDB IG, GRE-SPS
Microsoft SQL Server	IG-MS, ZDB IG
Microsoft Volume Shadow Copy Service (VSS)	IG-VSS
Network Data Management Protocol (NDMP) Server	IG-Var
Oracle Server	IG-O/S, ZDB IG
SAP MaxDB	IG-O/S
SAP R/3	IG-O/S, ZDB IG

Software application	Guides
Sybase Server	IG-Var
VMware vSphere	IG-VirtEnv, GRE-VMware

Look in these guides for details of the integrations with the following families of disk array systems:

Disk array family	Guides
EMC Symmetrix	all ZDB
HP P4000 SAN Solutions	ZDB Concept, ZDB Admin, IG-VSS
HP P6000 EVA Disk Array Family	all ZDB, IG-VSS
HP P9000 XP Disk Array Family	all ZDB, IG-VSS

Document conventions and symbols

Table 2 Document conventions

Convention	Element
Blue text: "Document conventions" (page 13)	Cross-reference links and e-mail addresses
Blue, underlined text: http://www.hp.com	Website addresses
Bold text	• Keys that are pressed • Text typed into a GUI element, such as a box • GUI elements that are clicked or selected, such as menu and list items, buttons, tabs, and check boxes
Italic text	Text emphasis
`Monospace` text	• File and directory names • System output • Code • Commands, their arguments, and argument values
`Monospace, italic` text	• Code variables • Command variables
`Monospace, bold` text	Emphasized monospace text

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

IMPORTANT: Provides clarifying information or specific instructions.

NOTE: Provides additional information.

TIP: Provides helpful hints and shortcuts.

Data Protector graphical user interface

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

Figure 1 Data Protector graphical user interface

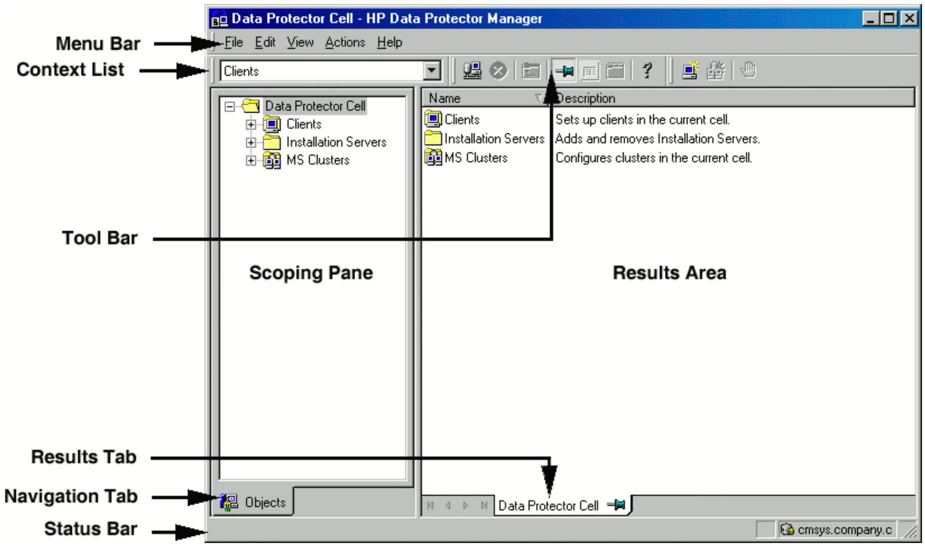

Name	Description
Clients	Sets up clients in the current cell.
Installation Servers	Adds and removes Installation Servers.
MS Clusters	Configures clusters in the current cell.

General information

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

HP technical support

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

http://www.hp.com/support

Before contacting HP, collect the following information:

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

Subscription service

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

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

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

HP websites

For additional information, see the following HP websites:

- http://www.hp.com
- http://www.hp.com/go/software
- http://www.hp.com/go/imhub
- http://support.openview.hp.com/selfsolve/manuals
- http://www.hp.com/support/downloads

Documentation feedback

HP welcomes your feedback.

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

Part I HP P6000 EVA Disk Array Family

This part describes how to configure the Data Protector HP P6000 EVA Disk Array Family integration, how to perform zero downtime backup and instant recovery using the HP P6000 EVA Disk Array Family integration, and how to resolve the integration-specific Data Protector problems.

1 Configuration and maintenance

Introduction

This chapter describes the configuration of the Data Protector HP P6000 EVA Disk Array Family integration. It also provides information on the ZDB database and on how to maintain the integration.

Prerequisites

- Obtain and/or install:

  P6000 EVA Array licenses and components:

  - HP Command View (CV) EVA and Virtual Controller Software (VCS or XCS).

    For installation instructions, see the SMI-S P6000 EVA Array provider and VCS or XCS documentation. For information on supported product versions, see the latest support matrices at http://www.hp.com/support/manuals.

  - HP Continuous Access (CA) P6000 EVA and/or HP Business Copy (BC) P6000 EVA license and microcode.

  - ***HP-UX systems:*** HP-UX MirrorDisk/UX software license.

    This license is required to enable mirroring functionality on HP-UX LVM.

  - An appropriate multi-path device management software.

    The software must be installed on the application system and the backup system.

    HP-UX systems: HP Secure Path (HP-UX)

    On HP-UX 11.31 systems, the multi-path device management software is not required since the operating system has native device multi-pathing capability.

    Linux systems: HP Device Mapper Multipath Enablement Kit for HP Disk Arrays 4.2.0 or newer version

    To configure the installed multi-path device management software:

    1. Start the multipath daemon and run the following command to configure the daemon so that it gets started during system startup:

       Red Hat Enterprise Linux: `chkconfig multipathd on`

       SUSE Linux Enterprise Server: `chkconfig boot.multipath on`

    2. Prevent the multipath device management software from queuing for unavailable disk volumes by modifying its configuration file.

       Add the following line into the `defaults` section of the file `/etc/multipath.conf`:

       ```
               no_path_retry            fail
       ```

       Ensure that this `no_path_retry` parameter value is not overridden by analogous entries in the `device` sections of the same file in which the corresponding P6000 EVA Array storage systems are configured.

    Windows systems: HP MPIO Full Featured DSM (Device Specific Module) for HP P6000 EVA Disk Array Family

  - A license for controlling the P6000 EVA Array storage system.

  - SANworks Snapshot licenses.

Data Protector licenses and components:

  - Appropriate zero downtime backup extension and instant recovery extension licenses-to-use (LTU).
  - HP StorageWorks P6000 EVA SMI-S Agent on the application system and the backup system.

  For installation and upgrade instructions and licensing information, see the *HP Data Protector Installation and Licensing Guide*.

- Make sure the same operating system version is installed on the application system and the backup system.
- If the application system and the backup system reside in a Data Protector cell with secured clients, ensure that access between both systems is allowed in both directions.
- Verify that the backup system is listed inside Command View EVA.
- Using Command View EVA, create source volumes and present them to the application system.

Prerequisites for Windows systems

- On Windows Server 2003 and Windows Server 2008 systems, disable the operating system option **Automatic mounting of new volumes**. In the Command Prompt window, run the command `mountvol /N`.
- Do not manually mount target volumes that were created by Data Protector.

For additional prerequisites for using HP P6000 EVA Disk Array Family with the Data Protector Microsoft Volume Shadow Copy Service integration, see the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

Limitations

- In cluster environments, the backup system must not be in the same cluster as the application system. Additionally, the backup system cannot be the cluster virtual server, it can only be a cluster node.
- In zero downtime backup sessions using multisnapping, only two snapshot types are supported by default: standard snapshot and snapclone. For information if your P6000 EVA Array environment supports multisnapping using vsnaps, see your Command View (CV) documentation. For instructions on how to enable support for the vsnap snapshot type in multisnapping ZDB sessions in Data Protector, contact HP technical support.

For information on either of the following items, see the *HP Data Protector Product Announcements, Software Notes, and References*:

- general Data Protector and integration-specific limitations
- supported platforms and integrations
- supported backup and connectivity topologies

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

ZDB database – SMISDB

ZDB database for the Data Protector HP P6000 EVA Disk Array Family integration is referred to as **SMISDB**. It keeps information about:

- Management systems on which Command View EVA runs. For each system, the following is stored:
  - Hostname as recognized in the IP network.
  - Port number through which the HP StorageWorks P6000 EVA SMI-S Agent communicates with the SMI-S P6000 EVA Array provider. For non-SSL connections, the default port is 5988. For SSL connections, the default port is 5989.
  - User name and encoded password for the SMI-S P6000 EVA Array provider login.
- Policies for redirecting the creation of snapclones and mirrorclones into specific disk groups.
- Information about the home (HP CA+BC P6000 EVA configurations).
- Replicas (groups of target volumes created in different backup sessions) kept on the disk array. For each target volume, the information includes:
  - ID of the ZDB session that produced the target volume
  - 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 P6000 EVA Array storage system on which the target volume resides
  - Snapshot type used for the replica (vsnap, standard snapshot, snapclone) and the type of source volumes of which the snapshots were created (original volume, mirrorclone)
  - 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 is marked for deletion)
  - Storage redundancy level (Vraid type) of the target volume
  - Exclusion flag (indicating that the replica is not involved in the replica set rotation and cannot be used for instant recovery)
  - Names of the application and backup systems involved in the session

  This information is written to the SMISDB when a replica is created, and is deleted from the database when a replica is deleted.
- Retained source volumes flag (after the instant recovery session, if the corresponding instant recovery option was selected).
- Mirrorclones created by Data Protector (tracked similarly as the replicas which cannot be used for instant recovery and are excluded from use).

SMISDB resides on the Cell Manager in:

Windows Server 2008: `Data_Protector_program_data\db40\smisdb`

Other Windows systems: `Data_Protector_home\db40\smisdb`

UNIX systems: `/var/opt/omni/server/db40/smisdb`

Configuring the integration

Before you start with the configuration, make sure the prerequisites listed in “Introduction” (page 17) are fulfilled. In addition, do the following:

HP BC P6000 EVA configurations: Connect the application and backup systems to the same P6000 EVA Array storage system. For ZDB to tape or ZDB to disk+tape, attach a backup device to the backup system.

For more information about HP BC P6000 EVA configurations, see the *HP Data Protector Zero Downtime Backup Concepts Guide*.

Combined (HP CA+BC P6000 EVA) configurations: For this configuration, you need at least two P6000 EVA Array storage systems located at different sites (with at least one HP CA P6000 EVA license, to set up the HP CA P6000 EVA links between the arrays, and at least one HP BC P6000 EVA license on the array where the replicas will be created).

Connect the application system to the P6000 EVA Array storage system containing source volumes (local disk array), and the backup system to the P6000 EVA Array storage system containing target volumes (remote disk array). Connect a backup device to the backup system.

For more information about HP CA+BC P6000 EVA configurations, see "ZDB in HP CA+BC P6000 EVA environments" (page 29) and the *HP Data Protector Zero Downtime Backup Concepts Guide*.

HP-UX LVM mirroring configurations: Group the physical volumes of a volume group into physical volume groups (PVGs). Each PVG may contain physical volumes from one or more P6000 EVA Array storage systems. All logical volumes in a volume group must be created with the `PVG-strict` allocation policy. Consequently, the mirrors will be created on different PVGs.

Before you run a backup, ensure that the mirrors of logical volumes involved in the backup are consistent. You can achieve this by running the `vgsync` command. Alternatively, specify the `vgsync` command in the **pre-exec** option in the backup specification. Consequently, Data Protector automatically runs the command before the replica is created.

For more information about LVM mirroring configurations, see "ZDB in HP-UX LVM mirroring environments" (page 33) and the *HP Data Protector Zero Downtime Backup Concepts Guide*. For more information about LVM mirroring, see the document *Managing Systems and Workgroups: A Guide for HP-UX System Administrators*.

To configure the integration:

- Provide the login information for the SMI-S P6000 EVA Array provider running on a management system. See "Setting login information for the SMI-S P6000 EVA Array provider" (page 20).
- If desired, set disk group pairs. See "P6000 EVA disk group pairs configuration file" (page 21).
- For HP CA+BC P6000 EVA configurations, set the home disk array. For details, see "HP CA P6000 EVA HOME configuration file" (page 21). If the home is not set, the HP StorageWorks P6000 EVA SMI-S Agent considers the configuration to be non-failover. In this case, replicas will always be created on the disk array remote to the current source.

Setting login information for the SMI-S P6000 EVA Array provider

Before starting ZDB sessions, provide login information for the SMI-S P6000 EVA Array provider running on a management system.

To set, delete, list, or check the login information, use the `omnidbsmis` command. For command syntax and examples, see the `omnidbsmis` man page.

If a failover from the active to the standby management system happens, proceed as follows:

- If standby and failed management systems have the same hostname, no action is needed.
- If standby and failed management systems have different hostnames, remove the failed system from the Data Protector configuration, and then add the new management system.

ⓘ **IMPORTANT:** If your SMI-S P6000 EVA Array provider is using non-default port numbers for SSL and non-SSL connections, enter the settings in the SMISDB database accordingly (use `omnidbsmis`).

To verify the configuration of SMI-S P6000 EVA Array provider, run `omnidbsmis -ompasswd -check [-host` *`ClientName`*`]`. It is recommended to run this command before backup and instant recovery sessions to check if the SMI-S P6000 EVA Array provider is operational and available on the network.

P6000 EVA disk group pairs configuration file

You can create snapclones and mirrorclones in a different disk group from that of the source volumes (original virtual disks). In this way, you help to reduce potential application performance degradation, since different physical disks are used for the source volumes and the replica. Note that standard snapshots and vsnaps are always created in the disk group of their source volumes whether the latter are original volumes or mirrorclones.

To set disk group pairs, use the `omnidbsmis` command. For command syntax and examples of manipulating the disk group pairs configuration file, see the `omnidbsmis` man page. The file template is as follows.

```
#
# HP Data Protector A.06.20
#
# StorageWorks P6000 EVA SMI-S disk group pairs configuration file
#
# Syntax:
#"EVA Node World Wide Name": "Working DG1", "Backup DG1"
#"EVA Node World Wide Name": "Working DG2", "Backup DG2"
#
# Example:
# "500508B101007000": "dg1", "dg2"
#
#
#
# End of file
```

NOTE: After the instant recovery session that uses the instant recovery method of switching the disks, the disk group of the former target volumes becomes the disk group of the new source volumes. In cases where characteristics of the two disk groups differ, the application system performance may be affected.

HP CA P6000 EVA HOME configuration file

This section is only applicable if you perform ZDB in HP CA+BC P6000 EVA configurations.

Due to HP P6000 EVA Disk Array Family hardware limitations, the concept of a defined home disk array does not exist within the HP P6000 EVA Disk Array Family. The HP StorageWorks P6000 EVA SMI-S Agent introduces this concept with the static HP CA P6000 EVA HOME configuration file. By setting the home disk array, you influence the Data Protector behavior in case of a failover. For more information, see "ZDB in HP CA+BC P6000 EVA environments" (page 29).

To create an P6000 EVA HOME configuration file template and put it into its default location (*`Data_Protector_program_data`*`\db40\smisdb`, *`Data_Protector_home`*`\db40\smisdb`, or `/var/opt/omni/server/db40/smisdb`), use the `omnidbsmis` command. This command is also used to upload the configuration file after editing (using an ASCII text editor like Notepad on Windows or VI on UNIX) back into its configuration directory. You can also list the DR groups with a specified P6000 EVA Array acting as a home and check if a specified DR group is part of an HP CA+BC P6000 EVA configuration. For command syntax and examples, see the `omnidbsmis` man page.

File template

```
#
# HP Data Protector A.06.20
#
# StorageWorks P6000 EVA SMI-S Continuous Access HOME configuration file
#
# Syntax:
# [EVA WWN]
# DRGroup1,DRGroup2
# DRGroup3
#
# Example:
# [50001FE15005DC00]
# "DRGroup 001"
#
#
#
# End of file
```

Configuration of the backup system

As part of a ZDB session, Data Protector performs necessary configuration steps, such as configuring volume groups, filesystems, mount points on the backup system. Data Protector can either create the same volume group structure on the backup system as it is on the application system and mounts the volumes to such mount points, or it can mount the volumes to the mount points specified in the backup specification.

For more information on creation of mount points on the backup system, see the *HP Data Protector Zero Downtime Backup Concepts Guide*.

Before running backup sessions, ensure that the host representing the backup system is configured on the P6000 EVA Array storage system. If it is not, configure it manually. If the hostname on the P6000 EVA Array storage system is different from the network hostname, use the `omnirc` variable `EVA_HOSTNAMEALIASES` to define the backup system object name.

Cluster environment:

If the backup system is a cluster virtual server, configure host objects using Command View in such a way that only one cluster node is configured in one host object. Additionally, set the variable `EVA_HOSTNAMEALIASES` to the appropriate host object on each cluster node.

For more information on the variable, see “ZDB omnirc variables” (page 139).

Use of mirrorclones for zero downtime backup

Specific firmware revisions of disk arrays of the HP P6000 EVA Disk Array Family support mirrorcloning, a special type of local replication. A mirrorclone is a dynamic replica of a storage volume, which is kept updated with changes made to the original storage volume via a local replication link. Replication between the original storage volume and its mirrorclone can be suspended and later re-established. For each storage volume, a single mirrorclone can be created on the disk array. Mirrorclones can be further replicated. As a result, mirrorclone snapshots are created – either standard snapshots or vsnaps. Each mirrorclone can have several snapshots attached and they can only be of the same type. For more information on the snapshot types, see “Snapshot types” (page 27).

Mirrorclone is one of the snapshot sources available for zero downtime backup in Data Protector. If selected in a ZDB backup specification, snapshots of mirrorclones of the selected storage volumes are created in the corresponding ZDB sessions, rather than snapshots of the storage volumes themselves. See “Creation of a standard snapshot of a mirrorclone” (page 23).

Figure 2 Creation of a standard snapshot of a mirrorclone

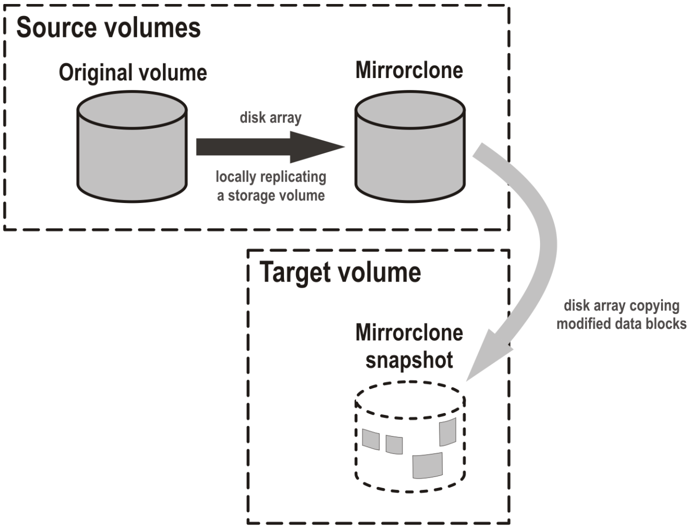

The advantage of this approach is in further shortening the backup window during which performance of the application using the source volumes for its data storage is affected. Mirrorclones can be created in advance using Command View EVA. However, if they do not exist yet when a ZDB session starts, but mirrorclone is selected as the snapshot source in the ZDB backup specification, they are automatically created by the HP StorageWorks P6000 EVA SMI-S Agent at the beginning of the session. For more information on the snapshot sources, see "Snapshot sources" (page 27). For information on creating mirrorclones outside Data Protector, see the HP P6000 EVA Disk Array Family documentation.

In Data Protector instant recovery sessions, data from the mirrorclone snapshots is restored directly to the corresponding original volumes. See "Instant recovery using a standard snapshot of a mirrorclone" (page 23).

Figure 3 Instant recovery using a standard snapshot of a mirrorclone

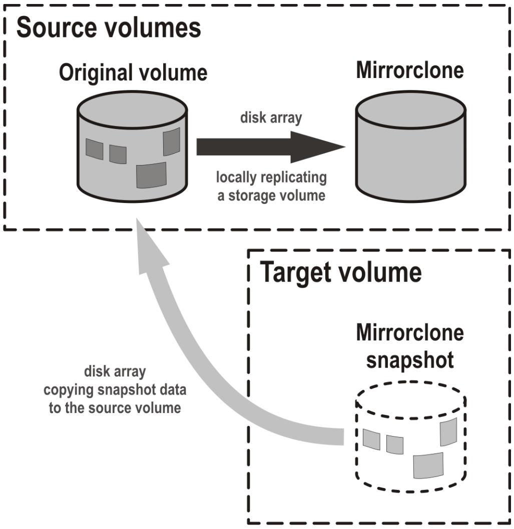
Source volumes
Original volume
Mirrorclone
disk array
locally replicating a storage volume
Target volume
Mirrorclone snapshot
disk array copying snapshot data to the source volume

You can delete mirrorclones that were created by Data Protector using the `omnidbsmis` command. For more information, see "Deleting replicas and the associated SMISDB entries" (page 25).

Maintaining the integration

Maintenance tasks are divided into the following categories:

- Querying information. See "Querying the SMISDB" (page 24).
- Checking consistency. See "Checking the SMISDB for consistency" (page 24).
- Deleting backup sessions. See "Purging the SMISDB" (page 24) and "Deleting replicas and the associated SMISDB entries" (page 25).
- Excluding and including ZDB sessions. See "Excluding and including sessions" (page 25)

Querying the SMISDB

Using the `omnidbsmis` command, you can list:

- all available zero downtime backup (ZDB) sessions
- all ZDB sessions based on a specific ZDB backup specification
- all ZDB sessions that are excluded from the replica set rotation
- obsolete volumes marked for purging
- disk group redirection configuration
- sets of retained source volumes, kept for forensic purposes after instant recovery
- details on a specific successful ZDB session and a report about all ZDB sessions based on a specific ZDB backup specification

For HP CA+BC P6000 EVA configurations, you can list data replication (DR) groups with a specified P6000 EVA Array acting as home. You can also check if a specified DR group is defined to be part of the HP CA+BC P6000 EVA HOME configuration in this cell.

For command syntax and examples, see the *HP Data Protector Command Line Interface Reference* or the `omnidbsmis` man page.

Checking the SMISDB for consistency

Data Protector can check the persistent data in the SMISDB against the P6000 EVA Array storage system and list the differences. Note that the check operation cannot detect whether the P6000 EVA Array configuration is correct or if the SMI-S P6000 EVA Array provider is currently operational. It just compares the saved data against the actual setup. This may provide misleading results, if the Command View EVA environment is not operating properly. If you use the results for an actual cleanup, verify the configuration first. The check operation also checks the entries which should be purged.

To check the SMISDB for consistency, use the `omnidbsmis` command. For details, see the *HP Data Protector Command Line Interface Reference* or the `omnidbsmis` man page.

Purging the SMISDB

During purge (normally started at the beginning of the backup session for the selected backup specification), the HP StorageWorks P6000 EVA SMI-S Agent attempts to delete storage volumes marked for purging. You can also run the SMISDB purge manually using the `omnidbsmis` command. For details, see the *HP Data Protector Command Line Interface Reference* or the `omnidbsmis` man page.

Deleting replicas and the associated SMISDB entries

Using the `omnidbsmis` command, you can delete:

- A specific ZDB session (and the replica created in it), identified by the session ID.
- ZDB sessions based on a specific ZDB backup specification (and the replicas created in them), identified by the ZDB backup specification name.
- A specific pseudo-ZDB session that tracks mirrorclone creation performed by Data Protector (and the mirrorclones created in it), identified by the ID of the associated "regular" ZDB session.

In all cases, you can either remove the corresponding replica (target volumes) or the mirrorclones from the disk array as well as delete the session information about them (the associated entries) from the SMISDB, or delete only the session information from the SMISDB.

IMPORTANT: Regardless of the chosen deletion scope, you cannot perform instant recovery using the affected replica after deletion, because the associated information is missing from the SMISDB.

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

Excluding and including sessions

After a backup session, you can leave the replica mounted on the backup system for other purposes than backup (replica set rotation) and instant recovery. For example, you can use such replica for database replication.

However, the intended use time for these replicas may exceed the time that is allowed by the current active rotation scheme, in which Data Protector automatically recycles the oldest replica. In such cases, you can exclude a session (a replica) from use (the replica set rotation and possibility to perform instant recovery) and thus preserve all target volumes of the replica.

Once you exclude a replica, the session that created the replica will not be used for replica set rotation, cannot be used for instant recovery, and cannot be deleted using the Data Protector CLI. To use an excluded session in an instant recovery or to delete the target volumes created in this session, you must first include the replica.

Excluding or including sessions can be triggered from the CLI for an individual *backup session*. To exclude a session from use (the replica set rotation and possibility to perform instant recovery), use the `omnidbsmis` command.

Using the `omnidbsmis` command, you can:

- Exclude a session (the option `-exclude`)
- Include a session (the option `-include`)
- List all excluded sessions (the options `-session –excluded`)

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

2 Backup

Introduction

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

You should be familiar with the HP P6000 EVA Disk Array Family concepts and procedures and basic Data Protector ZDB and instant recovery functionality. See the HP P6000 EVA Disk Array Family documentation and the *HP Data Protector Zero Downtime Backup Concepts Guide*.

Limitations

- The backup fails if you try to create a replica of a particular snapshot type and a replica of a different snapshot type (more specifically, standard snapshot or vsnap) for the same source volumes already exists. You must delete the existing replicas first. Snapclones are an exception. They do not block the creation of other snapshot types.
- Only one snapshot type for target volumes can be created during a ZDB session.
- When cloning process for a source volume is in progress, another snapshot (any type) of that source volume cannot be created.
- You cannot back up replicas (target volumes from existing and currently recorded backup sessions).
- If you perform ZDB in HP Continuous Access + Business Copy (CA+BC) P6000 EVA environments, note that the objects belonging to each specific data replication (DR) group are omitted from the ZDB session if:
  - the DR group write history log (DR group log) is in a state other than “not in use”.
  - the DR group is in the “suspended” state.
  - the DR group is in the “failsafe-locked” mode.

  If a DR group write mode is “asynchronous”, the HP StorageWorks P6000 EVA SMI-S Agent switches the mode into “synchronous” before starting ZDB. In this case, after ZDB is completed, the mode is reset to “asynchronous”.
- If there is not enough space for a standard snapshot or snapclone creation, the session fails.

Considerations

- If you do not select all of the filesystems on the disk for backup, Data Protector does not check if there are any filesystems that are not included in the backup specification and creates a replica of the entire disk. During instant recovery, the entire disk is restored and overrides also the filesystems that are not included in the backup specification, resulting in a possible data loss.
- If the source disks selected in a zero downtime backup specification are located on more than one P6000 EVA Array storage system, Data Protector will perform multisnapping for each unit separately, provided that it is not backing up the Oracle Server data in ASM configurations and multisnapping is not enforced by the `omnirc` variable `SMISA_ENFORCE_MULTISNAP`.

For more information on the backup-related considerations, see the *HP Data Protector Zero Downtime Backup Concepts Guide*. For detailed information on the backup-related problems and possible workarounds, see “Backup problems” (page 57).

Snapshot types

Data Protector supports the following snapshot types:

- snapshot *with* pre-allocation of disk space (**standard snapshot**).
- snapshot *without* pre-allocation of disk space (virtually capacity-free snapshot or shortly **vsnap**).
- complete copy of the source volume (the virtual disk containing original data), which is independent of the source volume (**snapclone**).

You can select the snapshot type in the GUI when creating a ZDB backup specification. For more information on snapshot types, see the *HP Data Protector Zero Downtime Backup Concepts Guide*.

NOTE: The snapclone snapshot type can only be used when the snapshot source selected in the ZDB backup specification is original volume.

Additionally, with the standard snapshot and snapclone types of snapshots, Data Protector supports multisnapping. Multisnapping is simultaneous creation of target volumes so that the backup data is consistent not only on each individual target volume, but also across all the volumes that constitute a snapshot.

Snapshot sources

Data Protector can replicate the following kinds of storage volumes which are supported with disk arrays of the HP P6000 EVA Disk Array Family:

- ordinary storage volume (**original volume**)

  This term refers to a storage volume on which original data resides and which is presented to the application system.

- **mirrorclone**

  This term refers to a mirrorclone of a storage volume on which original data resides. Mirrorclone is a particular type of local replication copy that can be created for a storage volume residing on a P6000 EVA Array. For more information on mirrorclones, see "Use of mirrorclones for zero downtime backup" (page 22).

  In a particular ZDB backup specification, when the selected snapshot source is mirrorclone, the only available snapshot types are standard snapshot and vsnap.

  Additionally, in the above circumstances, if mirrorclones of the selected storage volumes do not exist yet when the corresponding ZDB session is started, Data Protector automatically creates them first. Automatic mirrorclone creation may prolong the first ZDB session started for such a ZDB backup specification. To prevent this, create mirrorclones of the original volumes in advance using Command View EVA.

ZDB types

Using the P6000 EVA Array integration, you can perform:

- **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 ZDB 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, according to the tape backup type you have selected (Full, Incr, Incr1-9).

This replica is deleted after backup if the option **Keep the replica after the backup** is cleared for the 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 according to the tape backup type you have selected (Full, Incr, Incr1-9). 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 ZDB backup specification, and **To disk+tape** is selected when running/scheduling a backup.

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

Replica creation and reuse

On UNIX, the HP StorageWorks P6000 EVA SMI-S Agent identifies physical volumes, the volume group, and all logical volumes residing on it. This enables replication of the entire volume group on the array. On Windows, the HP StorageWorks P6000 EVA SMI-S Agent identifies partitions on a physical volume and entire disk is replicated. As a best practice, backup objects, such as filesystems or raw devices, from all logical volumes in a volume group and all partitions on physical volumes should be included in the backup. This helps in ensuring proper handling of filesystems and mount points during backup and restore.

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 or 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 first and then 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 or ZDB to disk+tape is performed and the specified **Number of replicas rotated** is reached.

If the oldest replica needs to be deleted, target volumes of the oldest replica are reused for creation of a new replica. Before such reuse, the target volumes are first converted into **containers** whenever the following prerequisites are fulfilled:

- the target volumes are standard snapshots (provided that the current ZDB session uses multisnapping), vsnaps (provided that the current ZDB session uses multisnapping), or snapclones
- the target volumes have the same size, storage redundancy level, and disk group location as required by the current ZDB session

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.

Note that for standard snapshots and snapclones, the number of replicas rotated has a significant impact on the amount of the required storage space. You should consider this storage requirement when defining your backup environment and/or backup policy.

Replica storage redundancy levels

The HP P6000 EVA Disk Array Family implements nested (hybrid) storage redundancy (RAID) technology, referred to as Vraid. P6000 EVA Array storage systems support creation of snapshots and snapclones which have a different storage redundancy level (Vraid type) than their source storage volumes. Of the supported Vraid types, **Vraid1** consumes the most storage space, followed by **Vraid6**, **Vraid5**, and finally **Vraid0**.

While you can freely select a Vraid type for snapclones, specific constraints apply to the Vraid type selection for standard snapshots and vsnaps. For details, see the table that follows.

Table 3 Allowed storage redundancy levels for standard snapshots and vsnaps

	Target volume – Vraid 6	Target volume – Vraid 1	Target volume – Vraid 5	Target volume – Vraid 0
Source volume – Vraid 6	Allowed	Allowed	Allowed	Allowed
Source volume – Vraid 1	Not allowed	Allowed	Allowed	Allowed
Source volume – Vraid 5	Not allowed	Not allowed	Allowed	Allowed
Source volume – Vraid 0	Not allowed	Not allowed	Not allowed	Allowed

If the redundancy level of source volumes is such that the specified snapshot redundancy level is not allowed, the zero downtime backup session creates snapshots with the redundancy level of their source volumes. The redundancy level is checked for each source volume separately.

Advantages

- By selecting the storage redundancy level, you can control the amount of storage space required.

NOTE: Target volumes of the **Vraid6** type can only be created in enhanced P6000 EVA disk groups.

In the Data Protector ZDB sessions during which mirrorclones are automatically created, the storage redundancy level of the source volumes (original volumes) is used for the mirrorclones. The storage redundancy level selected in the ZDB backup specification only applies to the target volumes.

ZDB in HP CA+BC P6000 EVA environments

The P6000 EVA Array storage system containing source volumes is known as a **local (source) disk array**, while the P6000 EVA Array storage system on which the replicas are created is a **remote (destination) disk array**. The mirrored source and target volumes constitute a **copy set**.

Data replication is always initiated from a local to a remote array. It is executed over a logical grouping of P6000 EVA virtual disks, known as a **data replication (DR) group**. A DR group can contain up to eight copy sets and share a common HP CA P6000 EVA log. Data replication control is always maintained at a DR group level.

The data backed up in HP CA+BC P6000 EVA configurations can be restored using either instant recovery or the standard Data Protector restore from tape procedure. After backup to tape, you can choose to keep replicas on the array for purposes other than instant recovery (by selecting **Keep the replica after the backup** in the backup specification).

DR group write history log (DR group log) states

If data replication is not possible, for example, due to the broken connection between the local disk array and the remote disk array, new data and changes to the existing data on the application system are written to the DR group log which resides on the local disk array. Each DR group configured on the disk array has its own DR group log.

During the logging process, the status of the DR group logs for the source virtual disks is set to “logging”. After the connection between the disk arrays is re-established, the contents of the DR group log are merged with the contents of the corresponding destination virtual disks on the remote disk array, so that the data redundancy is restored. For the duration of this activity, the status of the involved DR group logs is set to “merging”. After the merge is complete, the status is set back to “not in use”.

If the interruption of data replication is long-lasting, the storage space reserved for the DR group logs may run out. In this case, logs cannot hold all the changes. After the connection between the arrays is re-established, all original data in the involved DR groups has to be copied over. During this operation, the DR group log status is set to “copying”, and is re-set to “not in use” after the operation is complete.

DR groups with the DR group log state other than “not in use” are excluded from backup.

DR group states

DR group states are “normal/good”, “warning/attention”, “severe/failure”, and “unknown”. Data consistency is only guaranteed when a DR group is in the “normal/good” or “warning/attention” state. DR groups that are found in other states are excluded from the backup session.

DR group modes

DR group modes are as follows:

- Suspend

  This mode indicates that data replication is suspended and changes to the existing data are written to the log space until the replication is resumed. In the “suspend” mode, the DR group log state is set to “logging”.

  DR groups in such mode are excluded from backup.

- Failover

  This mode indicates that the replication direction is reversed after a failover.

- Failsafe-locked

  When a DR group is in this mode, write/read access to the source DR group is blocked due to the broken connection between the local disk array and the remote disk array. DR groups found in such a mode are excluded from the backup session.

HP CA+BC P6000 EVA ZDB scenarios

The HP StorageWorks P6000 EVA SMI-S Agent introduces the concept of a home disk array, which is defined inside a static HP CA P6000 EVA HOME configuration file. By setting the home disk array using the `omnidbsmis` command and specifying HP CA P6000 EVA failover handling options in the backup specification, you influence the Data Protector behavior in case of a failover. The information about home is stored in SMISDB and is used by the HP StorageWorks P6000 EVA SMI-S Agent to determine the state of a DR group (ideal or failed over).

If you intend to maintain the replica location after a failover, you must set the home disk array before creating a ZDB backup specification. If you intend to follow the replication direction, setting home is optional. For more information, see “HP CA P6000 EVA HOME configuration file” (page 21) and the `omnidbsmis` man page.

IMPORTANT: To enable proper replication handling after a failover, make sure the disk array you set as home is also your source disk array (the disk array acting as source at the time of the first ZDB session).

HP CA+BC P6000 EVA enables the following backup scenarios:

- Ideal, or non-failover scenarios, where replicas are always created on the array remote to current home.

Figure 4 A non-failover scenario

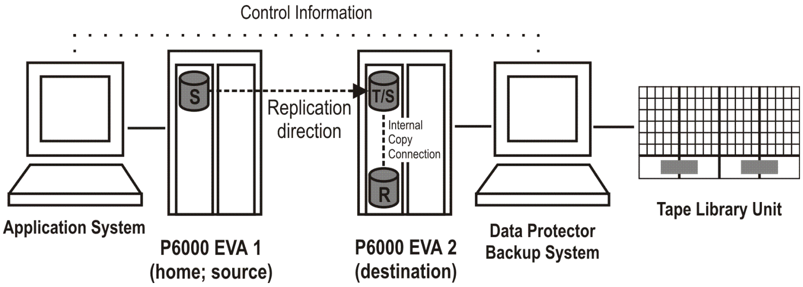

- Failover scenarios, where the roles of original source and destination are reversed after a failover. Replicas in such scenarios can be created:
  - On the disk array remote to the current source (**Follow direction of replication** backup option selected in the backup specification). It means that after a failover, the replication direction is reversed and the replicas are created on the array that was originally a source P6000 EVA Array. "Failover scenario 1" (page 31) depicts an environment where the location of replica creation was switched after a failover.

Figure 5 Failover scenario 1

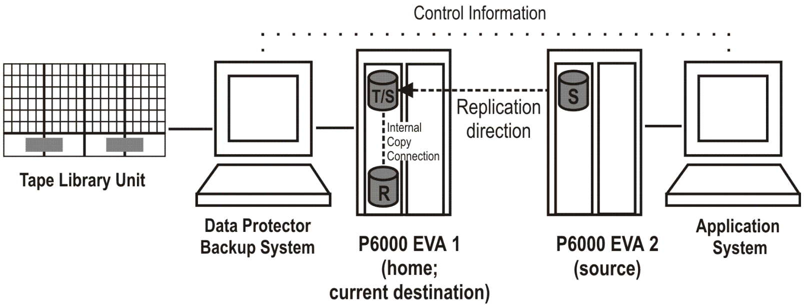

  - On the array remote to home (**Maintain replica location** backup option is selected in the backup specification). It means that after a failover, replica location is maintained and replicas continue on the destination array that has now become a source array. Note that for the time of replica creation, the source array performance may be affected.

Figure 6 Failover scenario 2

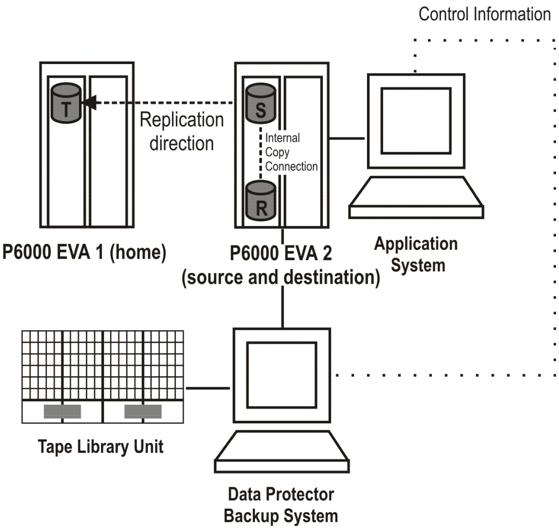

Consider the following:

- If you intend to always follow the replication direction, make sure the backup system has access to both local and remote P6000 EVA Array storage systems. Otherwise, after a failover, ZDB session fails because the replication direction switches and the backup system is no longer visible to the array where the replicas are created.
- If you intend to follow the replication direction, setting home in the HP CA P6000 EVA HOME configuration file is optional. However, if you will maintain replica location, you must set up the home before you create a ZDB backup specification. Is this is not done, the implications are as follows:
  - ***Non-failover scenarios:***

    ZDB sessions end successfully, but a warning that the home is not defined in the HP CA P6000 EVA HOME configuration file is issued.
  - ***Failover scenarios:***

    Replicas are created on the array remote to current source. However, if you maintain replication direction because your backup environment is distributed and the backup system is only accessible to one array (were the replicas were originally created), ZDB session fails as the replicas are now created on another array.

The basic HP CA+BC P6000 EVA configuration behavior is presented in the following diagram.

Figure 7 HP CA+BC P6000 EVA configuration behavior

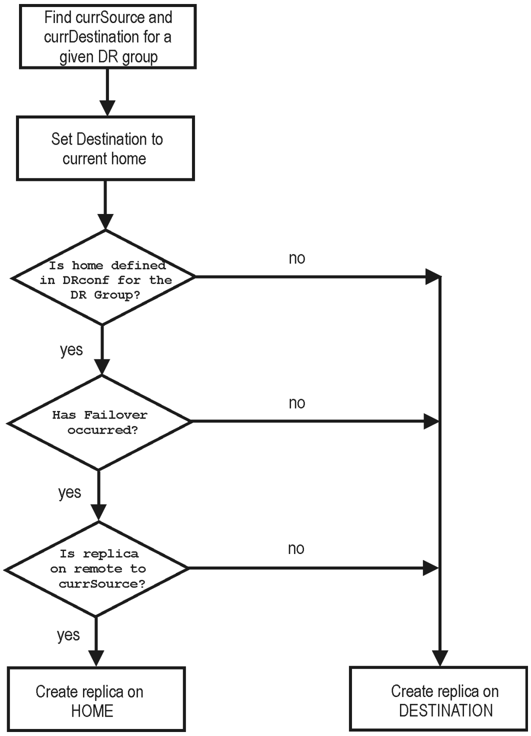

Find currSource and currDestination for a given DR group
Set Destination to current home
Is home defined in DRconf for the DR Group?
no
yes
Has Failover occurred?
no
yes
Is replica on remote to currSource?
no
yes
Create replica on HOME
Create replica on DESTINATION

Replica set rotation

In the HP CA+BC P6000 EVA non-failover scenarios, replicas are always created on the array remote to home. If the existing replica count (on the array where new replicas are) exceeds the specified number of replicas rotated, the oldest replica is deleted and the new one is created in its place (ensuring the maximum number of replicas is always within the defined rotation set).

In the HP CA+BC P6000 EVA failover scenarios, replicas are created either on:

- The array remote to current source (or on the home disk array)
- The array remote to home

In the first case, the number of replicas in a rotation set is only checked on the current destination array. The replicas created on the current source, which was a destination before a failover, are ignored. Therefore, there are situations when two replica sets are created on both the source and destination arrays.

In the second case, replica set rotation verification happens in a normal way.

NOTE: Replica rotation set is only created if you select the option **Keep the replica after the backup** and specify **Number of replicas rotated**. Without these options specified, the replica is deleted from the array after the backup to tape is completed.

For more information about replica set rotation, see the *HP Data Protector Zero Downtime Backup Concepts Guide*.

ZDB in HP-UX LVM mirroring environments

Your HP-UX LVM mirroring environment should be configured as follows:

- All logical volumes inside a volume group must be created with the `PVG-strict` allocation policy. Consequently, the mirrors will be created on different PVGs.
- As a best practice, different PVGs should be located on separate arrays. Consequently, mirrors are created on separate arrays.
- At least one PVG must contain a consistent mirror copy for all logical volumes of the volume group.

During a backup, Data Protector first checks the status of all mirror copies (see “Checking the mirrors” (page 34)). Out of all consistent mirror copies (mirrors without stale extents), one is backed up, preferably the one residing on a different array than the first mirror copy. If such a mirror copy does not exist, the first mirror copy is backed up. If the `ZDB_LVM_PREFERRED_PVG` omnirc variable is set, the mirror copy residing in the PVG specified in the variable is backed up, provided that this mirror copy does not have stale extents. Otherwise, another mirror copy is selected for backup according to the algorithm described above.

For more information on the `ZDB_LVM_PREFERRED_PVG` omnirc variable, see “ZDB omnirc variables” (page 139).

Figure 8 Checking the mirrors

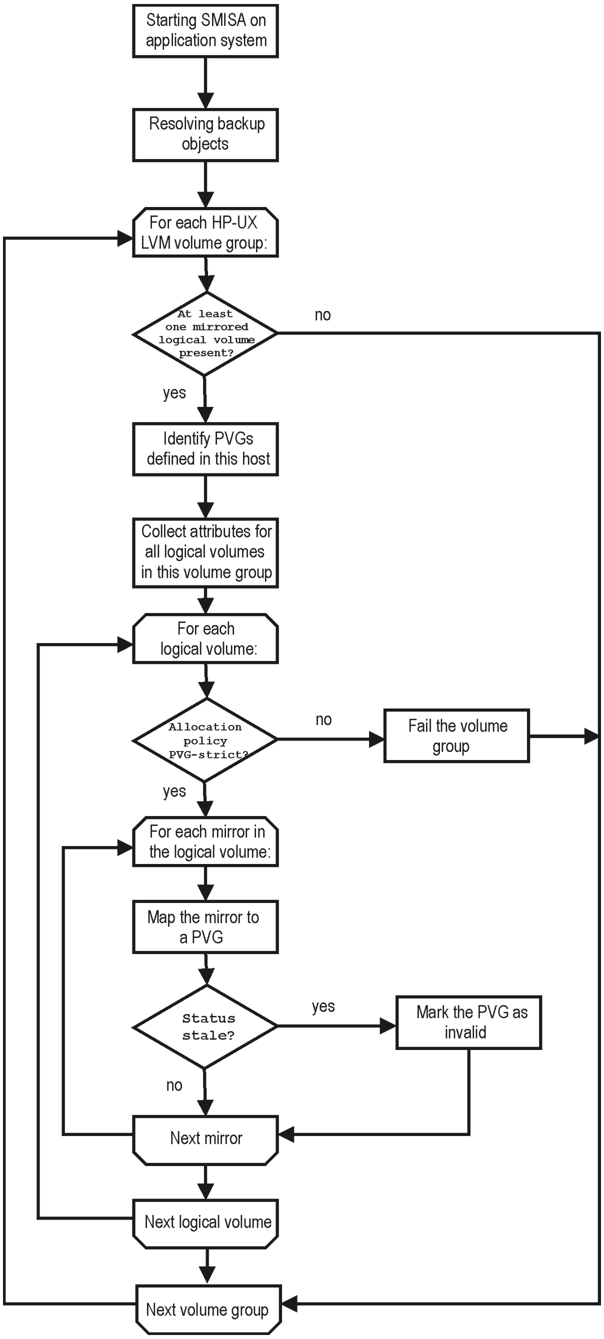
Starting SMISA on application system
Resolving backup objects
For each HP-UX LVM volume group:
At least one mirrored logical volume present?
no
yes
Identify PVGs defined in this host
Collect attributes for all logical volumes in this volume group
For each logical volume:
Allocation policy PVG-strict?
no
Fail the volume group
yes
For each mirror in the logical volume:
Map the mirror to a PVG
Status stale?
yes
Mark the PVG as invalid
no
Next mirror
Next logical volume
Next volume group

Data in replicas created using LVM mirroring can be restored in instant recovery sessions or sessions performing standard restore from tape.

Creating backup specifications

Considerations

- Consider all limitations that apply to the Data Protector P6000 EVA Array integration. See the *HP Data Protector Product Announcements, Software Notes, and References*, the *HP Data*

Protector Zero Downtime Backup Concepts Guide, and the limitation list in "Introduction" (page 17).

- If original volume is selected as the snapshot source in the ZDB backup specification, and mirrorclones of the selected storage volumes exist on the disk array when a corresponding ZDB session is started, the session fails.
- If mirrorclone is selected as the snapshot source in the ZDB backup specification, and mirrorclones of the selected storage volumes already exist when a corresponding ZDB session is started, the mirrorclones should not be presented to any system for the session to succeed.

Procedure

To create a ZDB backup specification for a disk array of the HP P6000 EVA Disk Array Family using the Data Protector GUI (**Data Protector Manager**), follow the steps:

1. In the Context List, select **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**. Right-click **Filesystem** (for both object types: filesystem and disk image) and click **Add Backup**.

   The Create New Backup dialog box appears.

   In the Filesystem pane, select the **Blank Filesystem Backup** template or some other template which you might have created. For information on templates, see the online Help index: "backup templates".

   Select **Snapshot or split mirror backup** as **Backup type** and **HP StorageWorks P6000 EVA SMI-S** as **Sub type**. For description of options, press **F1**.

   Click **OK**.
3. Under Client systems, select **Application system** and **Backup system**. If the application system is part of a server cluster, select the virtual server.
4. Under Replication mode, select the P6000 EVA Array configuration. If you select **HP Continuous Access P6000 EVA + HP Business Copy P6000 EVA**, also specify a choice for Replica handling during failover scenarios. For details about handling replica set rotation in HP CA+BC P6000 EVA configurations, see "HP CA+BC P6000 EVA ZDB scenarios" (page 30).

Figure 9 P6000 EVA Array backup options

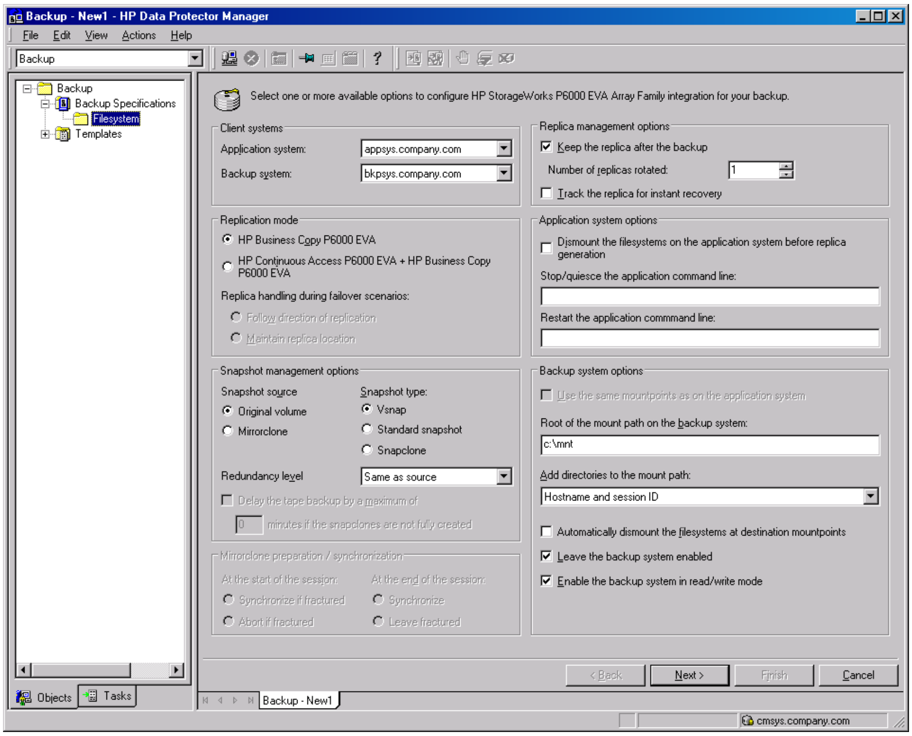

Select one or more available options to configure HP StorageWorks P6000 EVA Array Family integration for your backup.

Client systems
Application system: appsys.company.com
Backup system: bkpsys.company.com

Replication mode
● HP Business Copy P6000 EVA
○ HP Continuous Access P6000 EVA + HP Business Copy P6000 EVA
Replica handling during failover scenarios:
○ Follow direction of replication
○ Maintain replica location

Snapshot management options
Snapshot source
● Original volume
○ Mirrorclone
Snapshot type:
● Vsnap
○ Standard snapshot
○ Snapclone
Redundancy level: Same as source
□ Delay the tape backup by a maximum of 0 minutes if the snapclones are not fully created

Mirrorclone preparation / synchronization
At the start of the session:
○ Synchronize if fractured
○ Abort if fractured
At the end of the session:
○ Synchronize
○ Leave fractured

Replica management options
☑ Keep the replica after the backup
Number of replicas rotated: 1
□ Track the replica for instant recovery

Application system options
□ Dismount the filesystems on the application system before replica generation
Stop/quiesce the application command line:
Restart the application commmand line:

Backup system options
□ Use the same mountpoints as on the application system
Root of the mount path on the backup system: c:\mnt
Add directories to the mount path: Hostname and session ID
□ Automatically dismount the filesystems at destination mountpoints
☑ Leave the backup system enabled
☑ Enable the backup system in read/write mode

5. Under Snapshot management options, select the desired **Snapshot source**, **Snapshot type**, and **Redundancy level**.

TIP: For ZDB to disk+tape and ZDB to tape, and when snapclone is selected as the snapshot type, select **Delay the tape backup by a maximum of *n* minutes if the snapclones are not fully created**. In this case, backup to tape starts when the cloning process finishes, but not later than after the specified number of minutes. This helps prevent degradation of the application system performance during backup by reducing the concurrent load on the disk array.

6. If you have selected Mirrorclone as the snapshot source, under Mirrorclone preparation / synchronization, specify the options for handling local replication links between original volumes and mirrorclones during ZDB sessions. For information, see "Backup options" (page 38) or press **F1**.
7. Under Replica management options, specify a value for **Number of replicas rotated**. The number of standard snapshots or vsnaps that can exist for a specific source volume is limited by the target P6000 EVA Array storage system. The GUI does not limit the number of replicas rotated, but the session fails if the disk array-specific limit is exceeded.

   ZDB to disk, ZDB to disk+tape:

   Select the option **Track the replica for instant recovery** to enable instant recovery.

   NOTE: You can choose a ZDB-to-disk session or a ZDB-to-disk+tape session by selecting an appropriate value for the **Split mirror/snapshot backup** option when running or scheduling a ZDB session based on this ZDB backup specification. See "Scheduling ZDB sessions" (page 124).

   ZDB to tape:

   Leave the option **Track the replica for instant recovery** cleared.

   To preserve the replica on the disk array after the ZDB session, leave the option **Keep the replica after the backup** selected. To remove the replica after the session, clear this option.
8. Specify other zero downtime backup options as desired. For information, see "Backup options" (page 38) or press **F1**.

   Click **Next**.
9. Select the desired backup objects.

   Filesystem backup: Expand the application system and select the objects to be backed up. Note that all drive letters or mount points that reside on the system are displayed. You must select only the objects that reside on the disk array, otherwise the ZDB session will fail.

   IMPORTANT: To ensure that instant recovery succeeds and the environment is consistent after instant recovery, select all volumes on a disk (Windows systems) or all logical volumes of a volume group (UNIX systems) to be backed up. Even if you do not select an entire disk or volume group, the backup will succeed, but instant recovery may experience issues during configuration check of the environment. The configuration check can be disabled by clearing the option **Check the data configuration consistency** in the GUI or not specifying the option `-check_config` in the CLI when preparing for an instant recovery session. If this option is cleared (GUI) or not specified (CLI), the entire disk or volume group will be overwritten during instant recovery.

   Click **Next**.

   Disk image backup: Click **Next**.

10. 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 the backup image, specify the desired number of mirrors by clicking **Add mirror** or **Remove mirror**. Select separate devices for the backup image and each mirror.

    For information on object mirroring, see the online Help index: "object mirroring".

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

    Click **Next**.

11. In the Backup Specification Options group box, click **Advanced** and then the **HP StorageWorks P6000 EVA SMI-S** tab to open the P6000 EVA Array backup options pane.

    You can specify Application system options and modify all other options, except **Application system** and **Backup system** (note that you can change them after you save the ZDB backup specification). See "Backup options" (page 38) or press **F1**.

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

    Windows systems: To configure a ZDB backup specification for incremental ZDB sessions, select the **Do not use archive attribute** filesystem option in the WinFS Options pane to enhance the incremental ZDB behavior. For details, see "Backup options" (page 38).

    Click **Next**.

12. Following the wizard, open the scheduler to schedule the ZDB sessions. For more information, see "Scheduling ZDB sessions" (page 124) or press **F1**.

    Click **Next**.

13. In the **Backup Object Summary** page, specify additional options.

    Filesystem backup: You can modify options for the listed objects by right-clicking an object and then clicking **Properties**. For information on the object properties, press **F1**.

    Disk image backup: Follow the steps:

    a. Click **Manual add** to add disk image objects.
    b. Select **Disk image object** and click **Next**.
    c. Select the client system. Optionally, enter the description for your object. 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.

       Windows systems:

       Use the format

       `\\.\PHYSICALDRIVE#`

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

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

       UNIX systems:

       Specify a disk image section:

       `/dev/rdsk/`*`Filename`*, for example: `/dev/rdsk/c2t0d0`

       On HP-UX 11.31 systems, the new naming system can be used:

       `/dev/rdisk/disk#`, for example `/dev/rdisk/disk2`

       Specify a raw logical volume section:

       `/dev/`*`vgnumber`*`/rlvol`*`Number`*, for example: `/dev/vg01/rlvol1`

ⓘ **IMPORTANT:** To ensure that instant recovery succeeds and the environment is consistent after instant recovery, select all volumes on a disk (Windows systems) or all logical volumes of a volume group (UNIX systems) to be backed up. Even if you do not select an entire disk or volume group, the backup will succeed, but instant recovery may experience issues during configuration check of the environment. The configuration check can be disabled by clearing the option **Check the data configuration consistency** in the GUI or not specifying the option `-check_config` in the CLI when preparing for an instant recovery session. If this option is cleared (GUI) or not specified (CLI), the entire disk or volume group will be overwritten during instant recovery.

f. Click **Finish**.

Click **Next**.

14. Save your ZDB backup specification.

For information on scheduling ZDB sessions and starting interactive ZDB sessions, see "Scheduling ZDB sessions" (page 124).

NOTE: Backup preview is not supported.

Backup options

The following tables describe the P6000 EVA Array and ZDB related backup options. See also "P6000 EVA Array integration" (page 148).

Table 4 Client systems

Application system	The system on which the application runs. In cluster environments, specify the virtual server hostname (rather than the physical node hostname).
Backup system	The system to which your data will be replicated (backed up). In ZDB-to-disk+tape and ZDB-to-tape sessions, the backup data is copied from this system to a backup device.

Table 5 Replication mode

HP Business Copy P6000 EVA	Select this option to configure a ZDB backup specification for HP Business Copy (BC) P6000 EVA environments. Default: selected.
HP Continuous Access P6000 EVA + HP Business Copy P6000 EVA	Select this option to configure a ZDB backup specification for combined HP Continuous Access + Business Copy (CA+BC) P6000 EVA environments. Default: not selected.
Follow direction of replication	This option is only available if **HP Continuous Access P6000 EVA + HP Business Copy P6000 EVA** is selected as the P6000 EVA Array configuration. Select to follow the replication direction and create replicas on the disk array remote to the current source. After a failover, the replication direction is reversed and the replicas are created on the disk array that was originally a source P6000 EVA Array storage system. Default: selected.
Maintain replica location	This option is only available if **HP Continuous Access P6000 EVA + HP Business Copy P6000 EVA** is selected as the P6000 EVA Array configuration. Select to maintain replica location and create replicas on the disk array remote to home. After a failover, replicas will continue on the destination disk array that became the source P6000 EVA Array storage system during the failover. Default: not selected.

Table 6 Snapshot management options

Snapshot source	This option offers two choices: • **Original volume** Select this choice to create snapshots of the selected storage volumes. Note that the ZDB session fails if mirrorclones of the selected storage volumes (original volumes) exist on the disk array when the session is started. If this snapshot source is selected, the options **At the start of the session** and **At the end of the session** are not available. • **Mirrorclone** Select this choice to create snapshots of mirrorclones of the selected storage volumes. If no mirrorclones of the selected storage volumes (original volumes) exist when the ZDB session is started, Data Protector automatically creates them. Note that the ZDB session fails if snapshots of the original volumes exist on the disk array when the session is started. If this snapshot source is selected, the **Snapclone** snapshot type is not available. Default: Original volume.		
Snapshot type	**Vsnap** (default)	Creates snapshots without the pre-allocation of disk space.	If source volumes used in the session have existing target volumes of a different type (more specifically, vsnap or standard snapshot), the session is aborted. To successfully create a replica of a different type, first delete the existing target volumes. For more information on snapshot types, see the *HP Data Protector Zero Downtime Backup Concepts Guide*.
	Standard snapshot	Creates snapshots with the pre-allocation of disk space.	
	Snapclone	Creates clones of the source volumes. This snapshot type is only available if **Original volume** is selected as the snapshot source.	
Redundancy level	Select the storage redundancy level (Vraid type) to be used for the target volumes, or specify that the same redundancy level as for the source volumes should be used. If you create standard snapshots or vsnaps, the selected redundancy level should be the same or lower than the one used for the source volumes. Otherwise, the same redundancy level as used for the source volumes is applied. The redundancy level is checked for each source volume separately. The storage redundancy level and consequently the storage reliability of volumes using different Vraid types decreases as follows: **Vraid6** **Vraid1** **Vraid5** **Vraid0** **NOTE:** Target volumes using **Vraid6** can only be created in an enhanced P6000 EVA disk group. If **Vraid6** target volumes are in a basic P6000 EVA disk group, the effective Vraid type is reverted to **Vraid5**. Note that this options does not apply to mirrorclones. The mirrorclones that are automatically created during the Data Protector ZDB sessions always use the storage redundancy level of the source volumes (original volumes). Default: Same as source.		
Delay the tape backup by a maximum of *n* minutes if the snapclones are not fully created	This option is only available if **Snapclone** is selected as the snapshot type. Prevents degradation of the application data access times and reduces the load on the disk array by delaying the operation of copying the data to tape until the cloning process completes (ZDB to tape, ZDB to disk+tape). Defines also the maximum waiting time. When the specified time is reached, backup to tape starts in any case (even if the cloning process has not finished yet). Default: selected, 90 minutes.		

Table 7 Mirrorclone preparation / synchronization

At the start of the session	Replication links between original storage volumes and their mirrorclones can be in different states. For Data Protector to be able to created a mirrorclone snapshot, the replication link between the mirrorclone and the corresponding original storage volume must be in the "synchronized" state. This option offers two choices: • **Synchronize if fractured** Select this choice to enable running the ZDB session even when the replication link between a storage volume selected in the ZDB backup specification and its mirrorclone is fractured at the start of the session. In this case, Data Protector restores the synchronized state of each such replication link before a mirrorclone snapshot creation starts. • **Abort if fractured** Select this choice to make Data Protector abort the ZDB session in circumstances when the replication link between a storage volume selected in the ZDB backup specification and its mirrorclone is in the fractured state. If this choice is selected, the **Synchronize** choice for the **At the end of the session** option is automatically selected, and the **Leave fractured** choice is not available. Default: Synchronize if fractured.
At the end of the session	This option determines how Data Protector handles the mirrorclone replication links after the ZDB session. It offers two choices: • **Synchronize** Select this choice to make Data Protector restore the "synchronized" state of replication links between the mirrorclones involved in the ZDB session and the corresponding original volumes after the mirrorclone snapshot creation. Selecting this choice has an advantage over selecting the **Synchronize if fractured** choice for the **At the start of the session** option. The reason is that the time required for synchronization is usually much shorter if synchronization takes place immediately after the mirrorclone snapshots are created, and not only before creating the mirrorclone snapshots in the next ZDB session. • **Leave fractured** Select this choice to make Data Protector leave replication links between the mirrorclones involved in the ZDB session and the corresponding original volumes in the "fractured" state after the mirrorclone snapshot creation. If this choice is selected, the **Synchronize if fractured** choice for the **At the start of the session** option is automatically selected, and the **Abort if fractured** choice is not available. Default: Synchronize.

Table 8 Replica management options

Keep the replica after the backup	If configuring a ZDB to tape, select this option to keep the replica on the disk array after the zero downtime session. The replica becomes part of a replica set (specify a value for the option **Number of replicas rotated**). Unless the additional option **Track the replica for instant recovery** is selected, the replica is *not* available for instant recovery. If this option is not selected, the replica is removed at the end of the session. If the option **Track the replica for instant recovery** is selected, this option is automatically selected and cannot be changed. Default: selected.
Number of replicas rotated	This option is only available if the option **Keep the replica after the backup** is selected. During ZDB sessions, Data Protector creates a new replica and leaves it on the disk array until the value specified for the option **Number of replicas rotated** is reached. After that, the oldest replica is deleted and a new one created.

Table 8 Replica management options *(continued)*

	The number of standard snapshots or vsnaps is limited by the P6000 EVA Array storage system. Data Protector does not limit the number of replicas rotated, but the session fails if the limit is exceeded. Default: 1.
Track the replica for instant recovery	This option is only available if the option **Keep the replica after the backup** is selected. Select this option to perform a ZDB-to-disk or ZDB-to-disk+tape session and leave the replica on the disk array to enable instant recovery. Specify also a value for the option **Number of replicas rotated**. If this option is not selected, you cannot perform instant recovery using the replica created or reused in this session. Default: not selected.

Table 9 Application system options

Dismount the filesystems on the application system before replica generation	Select this option to dismount the filesystems on the application system before replica creation and remount them afterwards. Additionally, when entire physical drives (on Windows systems) or entire disks or logical volumes (on UNIX systems) are selected as backup objects in a disk image backup specification, selecting this option will dismount and later remount all filesystems on these objects. If any of these filesystems cannot be dismounted, the backup session fails. If an integrated application (for example, Oracle Server) exclusively controls data I/O on each physical drive, disk, or logical volume that will be backed up, the dismount operation is not needed. In such a case, you can leave this option cleared. Default: not selected.
Stop/quiesce the application command line	If a command is specified in this option, it is invoked on the application system immediately before replica creation. A use example is to stop applications not integrated with Data Protector. The command must reside on the application system in the directory `*Data_Protector_home*\bin` (Windows systems) or `/opt/omni/lbin` (UNIX systems). Do not specify the path to the command in this option. If the command fails, the command specified in the option **Restart the application command line** is not invoked. Thus, you may need to implement a cleanup procedure in the command specified in **Stop/quiesce the application command line**. If the `omnirc` variable `ZDB_ALWAYS_POST_SCRIPT` is set to `1`, the command specified in the option **Restart the application command line** is always invoked.
Restart the application command line	If a command is specified in this option, it is invoked on the application system immediately after replica creation. A use example is to resume operation of applications not integrated with Data Protector. The command must reside on the application system in the directory `*Data_Protector_home*\bin` (Windows systems) or `/opt/omni/lbin` (UNIX systems). Do not specify the path to the command in this option.

Table 10 Backup system options

Use the same mountpoints as on the application system	This option is not available if the application system is also the backup system (a single-host configuration). If this option is selected, the paths to mount points used for mounting the filesystems of the replica on the backup system are the same as paths to mount points where source volume filesystems were mounted on the application system. If the mount points are already in use, the session fails. For such circumstances, you must select the option **Automatically dismount the filesystems at destination mountpoints** in order for the session to succeed. **Windows systems:** The drive letters must be available, otherwise the session fails.

Table 10 Backup system options *(continued)*

	Default: not selected.
Root of the mount path on the backup system	This option is only available if the option **Use the same mountpoints as on the application system** is not selected. Specifies the root directory under which the filesystems of the replica are mounted. Where exactly the filesystems are mounted depends on how you define the option **Add directories to the mount path**. **NOTE:** For the SAP R/3 integration, the option is not applicable (the mount points created are always the same as on the application system). Defaults: **Windows systems:** `c:\mnt` **UNIX systems:** `/mnt`

Add directories to the mount path	This option is only available if the option **Use the same mountpoints as on the application system** is not selected.
	This option enables control over the created mount points. It defines which subdirectories will be created in the directory defined with the **Root of the mount path on the backup system** option. When Session ID is used in path composition, this guarantees unique mount points.
	Example for **Windows systems**:
	Root directory: `C:\mnt`
	Application system: `applsys.company.com`
	Backup session ID: `2008-02-22-4`
	Mount path on the application system: `E:\disk1`
	If **Hostname** is selected:
	`C:\mnt\applsys.company.com\E\disk1`
	If **Hostname and session ID** is selected:
	`C:\mnt\applsys.company.com\2008-02-22-4\E\disk1`
	If **Session ID** is selected:
	`C:\mnt\2008-02-22-4\E\disk1`
	If **Session ID and hostname** is selected:
	`C:\mnt\2008-02-22-4\applsys.company.com\E\disk1`
	NOTE:
	For the SAP R/3 integration, the option is not applicable (the mount points created are always the same as on the application system).
	Default: Hostname and session ID.

Automatically dismount the filesystems at destination mountpoints	If the mount points are in use (for example, volumes involved in the previous session may still be mounted) and this option is selected, Data Protector attempts to dismount the mounted filesystems. If the option is not selected and the mount points are in use, or if the option is selected and the dismount operation fails, the session fails. Default: not selected.
Leave the backup system enabled	This option is only available if the option **Keep the replica after the backup** is selected. If this option is selected, the filesystems remain mounted, the volume groups remain imported and active (UNIX systems), and the target volumes remain presented after the session. In this case, you can use the backup system for data warehousing purposes, but *not* for instant recovery. If the replica has to be reused later on (deleted, rotated out, or used for instant recovery), Data Protector automatically connects to the backup system, dismounts the filesystems, unpresents the target volumes, and clears the related logical structures on the backup system. At that point in time, if the filesystems are not mounted to the current backup system, Data

Table 10 Backup system options *(continued)*

	Protector cannot perform a proper cleanup, and aborts the operation or the instant recovery session. If this option is not selected, Data Protector dismounts filesystems, exports volume groups (UNIX systems), and unpresents the target volumes on the backup system at the end of the ZDB session. Default: selected.
Enable the backup system in read/write mode	This option is applicable to and can only be changed for UNIX systems only. On Windows systems, filesystems cannot be mounted in the read-only mode. Select this option to enable write access to volume groups and filesystems on the backup system. For backup purposes, it is sufficient to activate the backup system volume groups and mount the filesystems in the read-only mode. For other tasks, the read/write mode may be needed. Note that when this option is selected, the replica is open to modifications while the backup system is online. Consequently, data restored from such a replica includes all potential modifications. Defaults: **Windows systems:** selected. **UNIX systems:** not selected.

NOTE: In a particular ZDB session, the mount point paths to which filesystems of the replica are mounted on the backup system correspond the mount point paths to which source volumes were mounted on the application system if at least one of the following conditions is met:

- The GUI option **Use the same mountpoints as on the application system** is selected.
- The `omnirc` variable `ZDB_PRESERVE_MOUNTPOINTS` is set to `1`.

If the option **Use the same mountpoints as on the application system** is not selected, and the `omnirc` variable `ZDB_PRESERVE_MOUNTPOINTS` is set to `0`, the mount point paths are determined by the GUI options **Root of the mount path on the backup system** and **Add directories to the mount path**, and the `omnirc` variables `ZDB_MULTI_MOUNT` and `ZDB_MOUNT_PATH` are ignored.

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

Figure 10 ZDB-to-disk session flow for filesystem backup objects

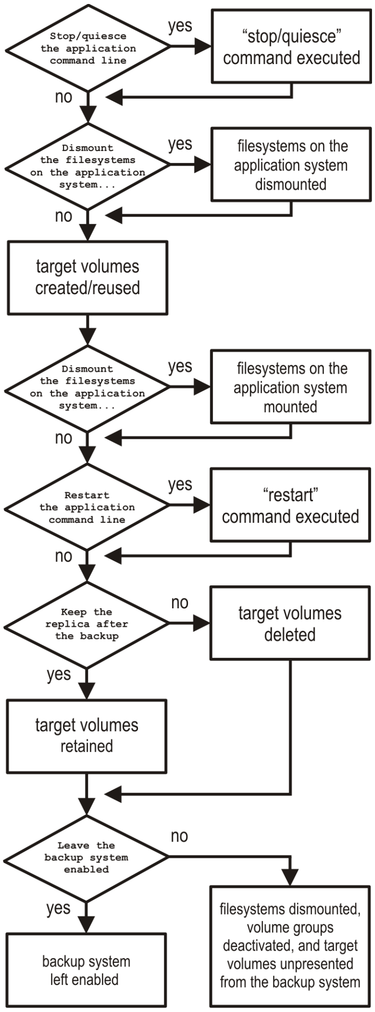
Stop/quiesce the application command line
yes
"stop/quiesce" command executed
no
Dismount the filesystems on the application system...
yes
filesystems on the application system dismounted
no
target volumes created/reused
Dismount the filesystems on the application system...
yes
filesystems on the application system mounted
no
Restart the application command line
yes
"restart" command executed
no
Keep the replica after the backup
no
target volumes deleted
yes
target volumes retained
Leave the backup system enabled
no
filesystems dismounted, volume groups deactivated, and target volumes unpresented from the backup system
yes
backup system left enabled

Figure 11 ZDB-to-tape and ZDB-to-disk+tape session flow for filesystem backup objects

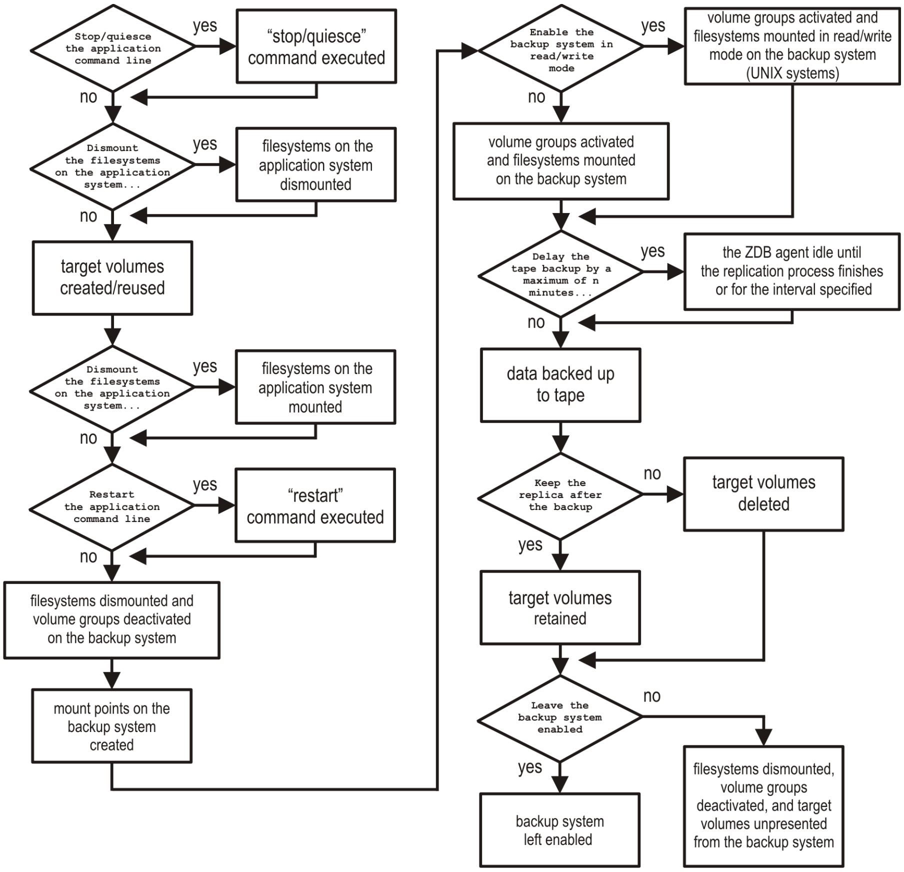
Stop/quiesce the application command line
yes
“stop/quiesce” command executed
no
Dismount the filesystems on the application system...
yes
filesystems on the application system dismounted
no
target volumes created/reused
Dismount the filesystems on the application system...
yes
filesystems on the application system mounted
no
Restart the application command line
yes
“restart” command executed
no
filesystems dismounted and volume groups deactivated on the backup system
mount points on the backup system created
Enable the backup system in read/write mode
yes
volume groups activated and filesystems mounted in read/write mode on the backup system (UNIX systems)
no
volume groups activated and filesystems mounted on the backup system
Delay the tape backup by a maximum of n minutes...
yes
the ZDB agent idle until the replication process finishes or for the interval specified
no
data backed up to tape
Keep the replica after the backup
no
target volumes deleted
yes
target volumes retained
Leave the backup system enabled
no
filesystems dismounted, volume groups deactivated, and target volumes unpresented from the backup system
yes
backup system left enabled

- *“Reuse”* means that target volumes from the oldest replica are deleted and a new replica is created.
- Due to an HP P6000 EVA Disk Array Family limitation, if a standard snapshot or vsnap exists on a disk array for a particular source volume, creation of another target volume of some other snapshot type fails, even if a different ZDB specification is used. To enable the creation of such a target volume, existing standard snapshots or vsnaps of the source volume must be deleted first.
- In ZDB-to-disk sessions, the backup option **Enable the backup system in read/write mode** is ignored.
- When configuring a ZDB backup specification for ZDB-to-tape sessions, you can select the option **Keep the replica after the backup**. When configuring a ZDB backup specification for ZDB-to-disk+tape sessions, this option is selected by default and cannot be deselected.

3 Restore

Introduction

This chapter describes configuring and running a filesystem or disk image restore of the data backed up using the P6000 EVA Array 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 only (ZDB to disk), on backup media only (ZDB to tape), or at both locations (ZDB to disk+tape).

Available restore types are:

- Restore from backup media on a LAN (standard restore). See "Standard restore" (page 46).
- Instant recovery. See "Instant recovery" (page 46).

Table 11 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. For more information on this restore type, see the online Help index: "restore".

TIP: You can improve the data transfer rate by connecting a backup device directly to the application system. For information on configuring backup devices, see the online Help index: "backups devices: configuring". For information on performing a restore using another device, see the online Help index: "selecting, devices for restore".

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, including filesystems or other objects which were not explicitly selected for backup. For instant recovery concepts, see the *HP Data Protector Zero Downtime Backup Concepts Guide*.

You can perform instant recovery using:

- the Data Protector GUI (

  See "Instant recovery using the GUI" (page 50).
- the Data Protector CLI

  See "Instant recovery using the CLI" (page 51)).

The number of replicas available for instant recovery is limited by the value of the option **Number of replicas rotated**, which determines the size of the replica set. You can view these replicas in the GUI in the Instant Recovery context by expanding Restore Sessions. Replicas are identified by the backup specification name and the session ID. Other information, such as time when the replica was created, is also provided. Alternately, you can use the Data Protector command `omnidbsmis` to list sessions. For more information, see the *HP Data Protector Command Line Interface Reference* or the `omnidbsmis` man page.

When instant recovery starts, Data Protector disables the application system. This includes dismounting filesystems and deactivating or exporting volume groups (UNIX). Before this is done, filesystems' and volume groups' status is checked, and only mounted filesystems are dismounted and active volume groups are deactivated or exported. At the end of the session, volume groups are reactivated and dismounted filesystems are mounted to the same mount points as were used during backup.

Limitations

- Instant recovery fails in the following situations:
  - The source volumes do not exist on the disk array any more.
  - The source volumes are not presented to the application system.
  - If the current configuration of the participating volumes (on Windows systems) or volume groups (on UNIX systems) is different from the volume/volume group configuration that existed at the time of the ZDB session and which was recorded in the SMISDB.
  - After instant recovery, restored filesystems are mounted to the same mount points or drive letters on the application system as they were at the backup time, but these mount points or drive letters have other filesystems mounted.
- While an instant recovery session is in progress, you cannot perform a zero downtime backup session that involves the source volumes to which the data is being restored.

For the P6000 EVA Array instant recovery-related limitations and considerations, see the *HP Data Protector Product Announcements, Software Notes, and References* and the *HP Data Protector Zero Downtime Backup Concepts Guide*.

Instant recovery methods

Depending on the snapshot type of the selected replica and the instant recovery options you select in the GUI or CLI, instant recovery can be performed using one of the following methods:

- by switching the disks

  This instant recovery method (also referred to as the "switch" method) is available only for replicas which consist of snapclones.

- by copying replica data and retaining the source volumes

  This instant recovery method (also referred to as the "copy-back" method) is available for replicas of any snapshot type.

- by copying replica data without retaining the source volumes

  This instant recovery method (also referred to as the "copy-back" method) is available for replicas of any snapshot type.

Switching the disks

With this instant recovery method, the source volumes are unpresented and the target volumes (replica from the selected session) are presented in the place of the source volumes. During this action, which is called identity exchange, information such as the volume names and comments are also exchanged between the source and target volumes. You can select to retain the old source volumes. However, you cannot retain the replica and cannot perform another instant recovery using the same backup data.

This method may change the physical location of the application data in production. After instant recovery, target volumes become the source volumes, therefore, the performance characteristics of the replica now become the characteristics of the application data. The application starts using the physical disks that were previously used for storing backup data.

Advantages

- Instant recovery is very fast, regardless of the amount of data that was backed up.
- The old source volumes can be retained after instant recovery.

Disadvantages

- It is not possible to perform another instant recovery using the same backup data, because the target volumes have become the new source volumes.
- If a replica to be used for instant recovery belongs to a different disk group than its source volumes, the disk group of the replica becomes the disk group of the source volumes after the instant recovery session. In this case, depending on the target disk group characteristics, the source storage volume performance may decrease.
- Storage reliability of the source storage volumes may decrease if the replica to be used for instant recovery has a lower redundancy level.

To perform this type of instant recovery, select the option **Switch to the replica** in the Data Protector GUI.

Copying replica data and retaining the source volume

With this instant recovery method, the process depends on the snapshot type used for the target volumes:

- If the target volumes are standard snapshots or vsnaps, new snapshots of the source volumes are created inside the same P6000 EVA disk group first, and the source volumes are overwritten with data from the existing replica afterwards. Original data is retained in the newly created snapshots.
- If the target volumes are snapclones, containers are created in the disk group of the source volumes first, the data from the existing replica is restored to the containers, and finally the source volumes are switched with the containers.

The replica is also retained in the replica set and another instant recovery using the same backup data can be run.

NOTE: If snapshots of mirrorclones were created in the ZDB session, during instant recovery, data from mirrorclone snapshots is restored directly to the corresponding original volumes.

Advantages

- Another instant recovery using the same backup data is possible.
- The old source volumes are retained after instant recovery.
- The performance characteristics of the restored volumes remain the same. This is because the physical disks and the characteristics of the source storage volumes (size, storage redundancy level) do not change.

Disadvantages

- Instant recovery is not as fast as with the "switching the disks" method.
- Instant recovery requires additional storage space in the disk group of the source volumes.
- When a replica that consists of standard snapshots or vsnaps is used, if newer replicas than the selected replica exist in the replica set, the instant recovery process lasts longer because not only the source volumes, but all newer replicas must be updated during the session as well.

To perform this type of instant recovery, select the options **Copy replica data to the source location** and **Retain source for forensics** in the Data Protector GUI.

Copying replica data without retaining the source volume

With this instant recovery method, the source volumes are directly overwritten with data from the replica. If the replica consists of snapclones, the source volumes are converted into containers before being overwritten. The source volumes are not retained and if the instant recovery session fails, the original application data residing on the source volumes is lost.

The replica is retained in the replica set and another instant recovery using the same backup data can be run.

NOTE: If snapshots of mirrorclones were created in the ZDB session, during instant recovery, data from mirrorclone snapshots is restored directly to the corresponding original volumes.

Advantages

- Another instant recovery using the same backup data is possible.
- The performance characteristics of the restored volumes remain the same. This is because the physical disks and the characteristics of the source storage volumes (size, storage redundancy level) do not change.
- No additional storage space is required for instant recovery.

Disadvantages

- Instant recovery is not as fast as with the "switching the disks" method.
- Data in the source volumes is lost during instant recovery.
- If the instant recovery session fails, the original data in the source volumes to be restored is lost.
- When a replica that consists of standard snapshots or vsnaps is used, if newer replicas than the selected replica exist in the replica set, the instant recovery process lasts longer because not only the source volumes, but all newer replicas must be updated during the session as well.

To perform this type of instant recovery, select the option **Copy replica data to the source location** and clear the option **Retain source for forensics** in the Data Protector GUI.

Instant recovery procedure

Prerequisites

- Target volumes used in an instant recovery session should not be presented to any system other than the backup system. You can make Data Protector automatically remove any disallowed target volume presentations by selecting the option **Force the removal of all replica presentations** in the GUI or by specifying the `omnir` option `-force_prp_replica` in the CLI.
- If a disk image backup with filesystems mounted on the selected disks was performed, manually dismount the filesystems on the disks to be restored before disk image instant recovery. If the option **Check the data configuration consistency** is cleared in the GUI or the `omnir` option `-check_config` is not specified in the CLI, the disks are dismounted automatically. In any case, re-mount the filesystems back after instant recovery.
- In HP Continuous Access + Business Copy (CA+BC) P6000 EVA environments, if the source volumes are included in a data replication (DR) group, instant recovery requires prior manipulation of the DR group and other steps that need to be followed before and after the instant recovery session. For details, see “Instant recovery for P6000 EVA Array in HP CA+BC P6000 EVA configurations” (page 133).

Considerations

- If mirrorclones were used in the corresponding zero downtime backup session, in an instant recovery session the data from the replica is restored directly to the original volumes.

Instant recovery using the GUI

Follow the steps:

1. In the Context List, select **Instant Recovery**.
2. In the Results Area, 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 a session from the list of ZDB-to-disk and ZDB-to-disk+tape sessions)
   - Backup object type (Filesystem, Disk Image, SAP R/3, ...) and backup session name and ID:

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

     Backed up object types are displayed.

     b. Expand the object type you want to restore.

     All available backup specification 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 replica set. Available sessions IDs (replicas) are displayed.

Figure 12 Selecting a session

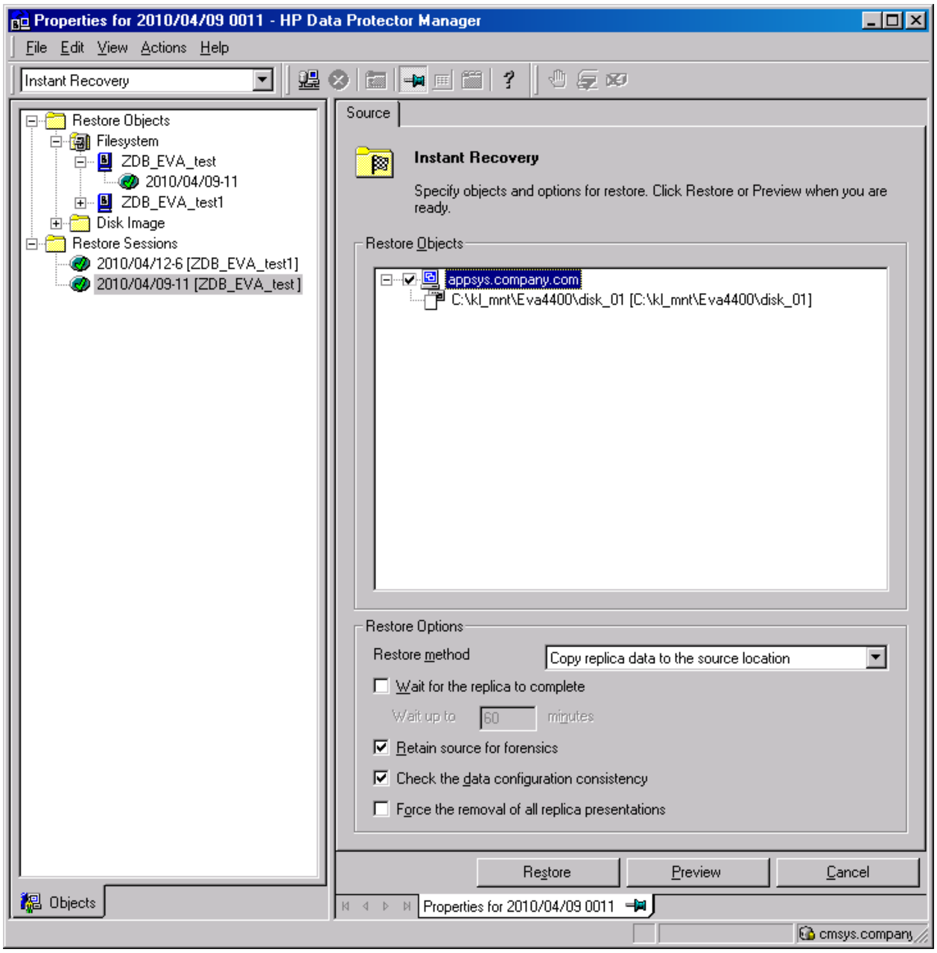

Properties for 2010/04/09 0011 - HP Data Protector Manager

- Restore Objects
  - Filesystem
    - ZDB_EVA_test
      - 2010/04/09-11
    - ZDB_EVA_test1
  - Disk Image
- Restore Sessions
  - 2010/04/12-6 [ZDB_EVA_test1]
  - 2010/04/09-11 [ZDB_EVA_test]

Instant Recovery

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

Restore Objects

- ☑ appsys.company.com
  - C:\kl_mnt\Eva4400\disk_01 [C:\kl_mnt\Eva4400\disk_01]

Restore Options

Restore method: Copy replica data to the source location

□ Wait for the replica to complete

Wait up to 60 minutes

☑ Retain source for forensics

☑ Check the data configuration consistency

□ Force the removal of all replica presentations

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

   The application system and its mount points or drive letters representing source volumes backed up during the selected session are displayed. Note that on UNIX all logical volumes inside a volume group and on Windows all partitions on a disk were backed up and if you did not select them all, they are not displayed here.
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 replica.
5. Specify other instant recovery options as desired. For information, see "Selecting a session" (page 50) and "Instant recovery options" (page 51), or press **F1**.
6. Click **Restore** to start the instant recovery session or **Preview** to start the instant recovery preview.

IMPORTANT: You cannot use the Data Protector GUI to perform instant recovery using backup data crated in a ZDB-to-disk+tape session after the media used in the session has been exported or overwritten. In such circumstances, use the Data Protector CLI instead. Note that the 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):

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

   From the output, select the backup session you want to restore.
2. Run the following command:

   ```
   omnir -host ClientName -session SessionID -instant_restore
   [INSTANT_RECOVERY_OPTIONS]
   ```

   where the meaning of the options is as follows:

ClientName	Application system name.
SessionID	Backup session ID (Step 1 of this procedure).

   For *INSTANT_RECOVERY_OPTIONS*, see "Instant recovery options" (page 51).

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

Instant recovery options

Table 12 Instant recovery options

Data Protector GUI/CLI	Function	
Restore method	**Copy replica data to the source location** / `-copyback`	Select this method to copy the replica data of the specified ZDB session to the original storage as follows: • With the **Retain source for forensics** option selected in the GUI or with the option `-leave_source` specified in the CLI, the process depends on the snapshot type used for the target volumes. If the target volumes are standard snapshots or vsnaps, new snapshots of the source volumes are created inside their P6000 EVA disk group first, the data from the existing replica is restored to the source volumes afterwards. Original data is retained in the newly created snapshots. If the target volumes are snapclones, containers are created in the disk group of the source volumes first, the data from the existing replica is restored to the containers,

Table 12 Instant recovery options *(continued)*

Data Protector GUI/CLI	Function	
		and finally the source volumes are switched with the containers. • With the **Retain source for forensics** option not selected in the GUI or with the option `-no_leave_source` specified in the CLI, the data from the existing replica is restored to the source volumes without prior operations. **CAUTION:** If the instant recovery session fails, a data loss on the source volumes may occur. **NOTE:** If mirrorclones were used in the corresponding zero downtime backup session, the term “source volumes” as used in both described cases refers to the original volumes, not the mirrorclones. After the instant recovery session, the replica is not deleted from the replica set, and the information about it is not deleted from the SMISDB. Therefore, the replica is available for another instant recovery session until it is rotated out from the replica set or deleted manually. This instant recovery method takes about as much time as the replica creation did, but the storage redundancy level is preserved and the source volumes remain in their P6000 EVA disk group. Default (GUI): selected
	Switch to the replica / `-switch`	This method can only be selected if the target volumes created in the corresponding zero downtime backup session are snapclones. Select this method to switch the target volumes of the specified ZDB session with the corresponding source volumes. Identifiers (WWNs) of the source volumes are assigned to the target volumes and these volumes are presented in the place of the source volumes. This action is called identity exchange. After the instant recovery session, the replica is not available for another instant recovery session, and the information about it is deleted from the SMISDB. Unless the option **Retain source for forensics** is selected in the GUI or unless the option `-leave_source` is specified in the CLI, the source volumes are removed. This instant recovery method is much faster than the method of copying replica data, since no data needs to be copied. However, after the instant recovery session, the new source volumes may have a lower storage redundancy level and may be located in a different P6000 EVA disk group. Default (GUI): not selected.
Wait for the replica to complete / `wait_clonecopy` *`n`*	This option is only available if **Copy replica data to the source location** is selected as the restore method in the GUI or the option `-copyback` is specified in the CLI. Before the actual data copy operation, storage space is allocated for replica restoration. Although the copy of the replica is only virtual at that time, it is immediately available for use. In the background, however, a process is still copying data from the replica to the source location (the replica normalization process). This copy process may degrade the disk array performance, and indirectly the application system performance as well. To reduce a potential performance degradation, select this option to make Data Protector wait for the copy to complete before proceeding with the session. In the GUI, you can set the maximum delay with the option **Wait up to *n* minutes**.	

Table 12 Instant recovery options *(continued)*

Data Protector GUI/CLI	Function
	Additionally, you can control the copy process by setting appropriate `omnirc` variables. See "P6000 EVA Array specific variables" (page 141). Default (GUI): not selected.
Wait up to *n* minutes	This option is only available if the option **Wait for the replica to complete** is selected. This option defines the maximum time that Data Protector waits for the replica data to be copied to the source location before proceeding with the instant recovery session. If the copy process completes before the time period expires, the session continues immediately. Default (GUI): 60 minutes.
Retain source for forensics / `-leave_source` | `—no_leave_source`	If this option is selected in the GUI or the `-leave_source` option is specified in the CLI, Data Protector preserves original data from the source volumes on the disk array after instant recovery. The original data resides in the source volume snapshots in the same P6000 EVA disk group as the source volumes. For example, you can use this option to investigate why the original data got corrupted. If this option is not selected in the GUI or the `-no_leave_source` option is specified in the CLI, the source volumes are either overwritten with data from the replica (with the "copy-back" instant recovery method) or deleted (with the "switch" instant recovery method) during the instant recovery session. In case of the "copy-back" instant recovery method in which the replica used consists of snapclones, the source volumes are converted into containers before being overwritten, provided that the source and target volumes match in size, redundancy level, and belong to the same P6000 EVA disk group. **CAUTION:** If you decide to perform instant recovery by copying replica data and not to preserve source volumes after the session (the option **Copy replica data to the source location** is selected and the option **Retain source for forensics** is cleared), and then the instant recovery session fails, a data loss on the source volumes may occur. Default (GUI): selected.
Check the data configuration consistency / `-check_config` | `—no_check_config`	If this option is selected in the GUI or the `-check_config` option is specified in the CLI, Data Protector performs a sanity check and a comparison of current volume group configuration of the volume groups participating in the instant recovery session and the volume group configuration information kept in the SMISDB after the corresponding zero downtime backup session. If the sanity check fails or the volume group configuration has changed since the zero downtime backup session, the instant recovery session aborts. Additionally, if the "switch" restore method is chosen and this option is selected for an instant recovery session, and storage redundancy levels of a particular target volume in the replica and its corresponding source volume differ, the session fails. **MC/ServiceGuard clusters:** When performing instant recovery to some other node than the one from which data was backed up, you must select this option in the GUI or specify the `-check_config` option in the CLI. In such circumstances, the current volume group configuration on the node to which data is to be restored differs from the volume group configuration kept in the SMISDB. Consequently, the SMISDB volume group configuration data is replaced by the current volume group configuration data on the node to which data is to be restored, and the instant recovery session succeeds. Default (GUI): selected.
Force the removal of all replica presentations / `-force_prp_replica`	If this option is selected in the GUI or specified in the CLI, and a target volume containing data to be restored is presented to a system other than the backup system, the HP StorageWorks P6000 EVA SMI-S Agent removes such presentation. If the option is not selected in the GUI or not specified in the CLI, the instant recovery session fails in such circumstances. If this option is selected in the GUI or specified in the CLI, and any target volume containing data to be restored is presented to the backup system, but cannot be dismounted in an operating system-compliant way, the HP StorageWorks P6000

Table 12 Instant recovery options *(continued)*

Data Protector GUI/CLI	Function
	EVA SMI-S Agent performs a forced dismount. If the option is not selected in the GUI or not specified in the CLI, the instant recovery session fails in such circumstances. Default (GUI): not selected.

Instant recovery in HP CA+BC P6000 EVA configurations

You can perform instant recovery to restore the data backed up in HP CA+BC P6000 EVA configurations. For detailed information, see "Instant recovery for P6000 EVA Array in HP CA+BC P6000 EVA configurations" (page 133).

Instant recovery and LVM mirroring

Method 1 – instant recovery with reducing and extending the mirrors

Using this method, you reduce the mirrors to include only the PVG from which the backup was taken. Instant recovery is performed after the volume is reduced, and then the logical volume is mirrored again to include all PVGs.

CAUTION: Before reducing the mirrors, verify that the mirror which is being reduced is the correct one. Otherwise, depending on the restore options selected, irrecoverable loss of data may happen. It is recommended to record and verify mirroring settings and the output of `lvdisplay` and `vgdisplay` commands.

1. Reduce the mirrors using the `lvreduce` command. Only the mirror copy that was backed up should remain.

   #### Example

   If the VG01 volume group contains a logical volume `lvol1`, which contains the disks `/dev/dsk/c12t0d0` and `/dev/dsk/c12t0d1` (belonging to PVG-2), and `/dev/dsk/c15t0d0` and `/dev/dsk/c15t0d1` (belonging to PVG-1), reduce the volume to contain only disks from PVG-2:

   ```
   lvreduce -m 0 /dev/vg01/lvol1 /dev/dsk/c15t0d0
   ```

   ```
   lvreduce -m 0 /dev/vg01/lvol2 /dev/dsk/c15t0d1
   ```

   You can also check the output using the `lvdisplay` command.

2. Perform instant recovery using the Data Protector GUI or CLI. For instructions, see "Instant recovery procedure" (page 49).

   NOTE: If the option **Check the data configuration consistency** option is selected in the GUI or the option `-check_config` is specified in the CLI, instant recovery fails, as the configuration of the volume group changed. Therefore, clear (GUI) or do not specify (CLI) this option before instant recovery.

3. Extend the mirror to include PVG-1 in the logical volume. The mirror is created again to include both volume groups.

   Example

   To extend the logical volume to contain two mirrors as in the original setup, execute:

   ```
   lvextend -m 1 /dev/vg01/lvol1 /dev/dsk/c15t0d0
   ```

   ```
   lvextend -m 1 /dev/vg01/lvol1 /dev/dsk/c15t0d1
   ```

   This way, `lvol1` contains the disks `/dev/dsk/c15t0d0` and `lvol2` contains the disks `/dev/dsk/c15t0d1` as a mirrored copy.

Method 2 – instant recovery with splitting and merging the mirrors

This method uses the splitting functionality of LVM mirroring. Logical volumes are first split to create backup volumes. These backup volumes can be overwritten by the data from the replica created. Later, the backup volumes are merged back.

△ **CAUTION:** Before splitting the mirrors, verify that the mirror which is being split is the correct one. Otherwise, irrecoverable loss of data may happen. It is recommended to record mirroring settings and the output of `lvdisplay` and `vgdisplay` commands.

1. Split the mirrors using the `lvsplit` command. Specify the group where the replica will not be restored by checking `vgdisplay` and `lvdisplay` outputs. After the split, volumes in the PVGs are no longer in the mirror, and their backup copies are present.

   Example

   A volume group VG01 contains logical volumes `lvol1` and `lvol2`, which contain the disks belonging to PVG-1 and PVG-2. To split the logical volume to contain the disks from PVG-2 only, execute:

   ```
   lvsplit -s back -g PVG1 /dev/vg01/lvol1 /dev/vg01/lvol2
   ```

   The disks from PGV-1 are split and a new logical volume with the suffix `back` is created. This logical (backup) volume can be accessed at `/dev/vg01/lvol1back` and `/dev/vg01/lvol2back`.

   You can check this using the `vgdisplay` command, which shows that another pair of logical volumes is now present in the volume group `vg01`. Similarly, the `lvdisplay` command shows that the physical disks from PVG-1 are no longer part of `lvol1` (they belong to `lvol1back`).

2. Perform instant recovery using the Data Protector GUI or CLI. For instructions, see “Instant recovery procedure” (page 49).

   NOTE: If the option **Check the data configuration consistency** option is selected in the GUI or the option `-check_config` is specified in the CLI, instant recovery fails, as the configuration of the volume group changed. Therefore, clear (GUI) or do not specify (CLI) this option before instant recovery.

3. Merge the mirrors back to their original logical volume using the `lvmerge` command (the newly created logical volumes, which are merged back, have the `back` suffix). This way, the mirror is created again to include both volume groups.

   Example

   The logical volume `lvol1` was split before instant recovery. After instant recovery, execute:

   ```
   lvmerge /dev/vg01/lvol1back /dev/vg01/lvol1
   ```

   ```
   lvmerge /dev/vg01/lvol2back /dev/vg01/lvol2
   ```

Instant recovery in a cluster

For information on instant recovery with an application running on or a filesystems residing in MC/ServiceGuard or Microsoft server cluster, see "Cluster configurations" (page 126).

4 Troubleshooting

Before you begin

This chapter lists general checks and verifications plus problems you may encounter when using the P6000 EVA Array integration. For general Data Protector troubleshooting information, see the *HP Data Protector Troubleshooting Guide*.

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

Checks and verifications

- On the application and backup systems, examine system errors logged into:

  Windows Server 2008: `Data_Protector_program_data\log\debug.log`

  Other Windows systems: `Data_Protector_home\log\debug.log`

  UNIX systems: `/var/opt/omni/log/debug.log`

Backup problems

Problem

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

Action

Check that the `HP StorageWorks P6000 EVA SMI-S Agent` integration module is installed on the application system and the backup system. To do that, open the `cell_info` file located on the Cell Manager in the following directory:

Windows Server 2008: `Data_Protector_program_data\Config\server\cell\cell_info`

Other Windows systems: `Data_Protector_home\Config\server\cell\cell_info`

UNIX systems: `/etc/opt/omni/server/cell/cell_info`

File contents should look similar to the following:

```
-host "sap002.company.com" -os "HP s800 HP-ux-11.00" -cc A.06.20 -da
A.06.20 -ma A.06.20 -SMISA A.06.20
```

Problem

The HP StorageWorks P6000 EVA SMI-S Agent fails to connect to the Cell Manager and retrieve configuration data

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

The HP StorageWorks P6000 EVA SMI-S 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

On an HP-UX system, the HP StorageWorks P6000 EVA SMI-S Agent fails to communicate with the HP SMI-S P6000 EVA Array provider using SSL

```
[Warning]
The SSL connection to the StorageWorks P6000 EVA Array Family SMI-S \
provider has failed.
The error description returned is:
SSL Exception: Random seed file required
```

On HP-UX systems, Pegasus libraries require the random number generator pseudo device for its SSL-based communication with the SMI-S P6000 EVA Array provider. If the pseudo device is not present, the warning appears.

Action

1. Install the pseudo device in `/dev/random` on the HP-UX backup system.
2. Re-run the session.

Problem

No HP SMI-S CIMOM login entries are configured within SMISDB

Action

Add an HP SMI-S CIMOM login information to SMISDB:

```
omnidbsmis -ompasswd -add ClientName [-ssl] [-port PortNumber] [-user
Username] [-passwd Password]
```

Problem

On a UNIX system, ZDB backup sessions stop responding for a long time during the resolving of the backup objects on the application system

When resolving the backup objects on the application system, Data Protector sends SCSI inquiries to identify the vendor-specific details of the virtual disk to be replicated. If this virtual disk belongs to a DR group that is in the “failsafe-locked” mode, SCSI inquiries do not return at all. As a result, the session stops responding.

Action

1. Abort the session and stop the ZDB agent processes that stopped responding on the application system.
2. Identify the root cause for the “failsafe-locked” mode of the DR group and fix it by bringing the DR group back into normal operational mode.

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 command line` was specified, check that it stops all processes using the filesystem.

Problem

On a Windows system, replica cannot be mounted to the target location on the backup system

```
[Major]
Filesystem \\.\Volume{9640da9a-6f36-11d7-bd7a-000347add7ba} could not
be mounted to C:\mnt.
([145] The directory is not empty.).
```

When a backup with nested mountpoint objects is run, replica 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 empty the directory where filesystems are to be mounted or select the backup option **Automatically dismount the filesystems at destination mountpoints**. If you choose manual action, and leave the default root mount path `c:\mnt` in the ZDB backup specification, you should empty the `mnt` directory.

Problem

On a Windows Server 2003 system, the HP StorageWorks P6000 EVA SMI-S Agent fails to resolve filesystem objects during the backup system preparation

```
[Major]
Resolving of filesystem G:\ failed. (Details unknown.)
[Minor]
Preparation of the backup system failed.
```

The HP StorageWorks P6000 EVA SMI-S Agent, after presenting a replicated disk and finishing rescan, starts searching filesystem volumes attached to this disk. However, on some Windows 2003 systems, more time is needed to recognize filesystem volumes and make them available for mount operations. As a result of this delay, the HP StorageWorks P6000 EVA SMI-S Agent fails to resolve filesystem objects on the backup system.

Action

Enable volume rescan retries with a defined delay period as follows:

1. On the backup system, set `ZDB_VOLUMESCAN_RETRIES` and `ZDB_POST_RESCAN_DELAY` `omnirc` variables to moderately higher values (defaults are 5 retries and 30 seconds delay time).
2. Restart the backup session.

Problem

Data Protector fails to delete a replica from the replica set in a cluster environment

A ZDB session reports the following major error and message:

```
[Major]
     Resolving of storage volume TargetVolumeID has failed.
...
[Normal]
  Some disks are still in use. They will be moved in purge bucket.
```

This error may occur in a cluster environment with the backup system which is a cluster virtual server. In such circumstances, after a failover, new backup sessions cannot rotate out the replicas on the active node because the presentations match the passive node. The replicas to be removed are marked with the purge flag in the SMISDB, and you are advised to delete such replicas.

Action

To delete the replicas with the purge flag from the disk array and the SMISDB, perform one of the following actions:

- Manually delete all storage volumes that are marked for purging by running:

  ```
  omnidbsmis -purge [-force] -host ClientName
  ```

  where *`hostname`* is the name of the node on which you want to perform the purge operation.

Use the `-force` option to remove the volumes marked for purging even if they are presented to a system.

- Perform manual failover and run another ZDB session. The session will delete all the volumes marked for purging on the new active node.

Problem

On an HP-UX system, 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 SMI-S P6000 EVA Array provider user interface, delete all presentations attached to removed target volumes.

Instant recovery problems

Problem

Instant recovery fails

The problem may occur if the option **Force the removal of all replica presentations** is not selected and a target volume from the selected replica is presented to some system other than the backup system or the target volume cannot be dismounted.

Action

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

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 omnirc variables" (page 139)) and manually dismount the volumes to be restored.
2. Start instant recovery.
3. After instant recovery, manually mount restored volumes.

Problem

Instant recovery fails on HP-UX 11.31 in LVM mirroring environments

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

This problem occurs if the option **Check the data configuration consistency** is selected in the GUI or the option `-check_config` is specified in the CLI, and may be caused by operating system upgrade. When upgrading the operating system on your application system from HP-UX 11.23 to HP-UX 11.31, device special files (DSFs) change from the legacy format to the new persistent DSF format. As a result, your LVM configuration now refers to physical volumes in the new format, which is checked during instant recovery.

Action

Clear the option **Check the data configuration consistency** in the GUI or do not specify the option `-check_config` in the CLI for the backup objects that are part of the LVM mirroring configuration, and restart the instant recovery session.

Part II HP P9000 XP Disk Array Family

This part describes how to configure the Data Protector HP P9000 XP Disk Array Family integration, how to perform zero downtime backup and instant recovery using the HP P9000 XP Disk Array Family integration, and how to resolve the integration-specific Data Protector problems.

5 Configuration and maintenance

Introduction

This chapter describes the configuration of the Data Protector HP P9000 XP Disk Array Family (HP P9000 XP Disk Array Family) integration. It also provides information on the ZDB database and on how to maintain the integration.

Prerequisites

- Obtain and/or install:

  P9000 XP Array components:

  - RAID Manager Library on the application system and the backup system.

    RAID Manager Library is disk array firmware-dependent. For information on which version of RAID Manager Library to use, see the latest support matrices at http://www.hp.com/support/manuals. For installation instructions, see the RAID Manager Library documentation.

    Note that snapshots are supported only by the disk array microcode 50-04-20 and newer versions, and only by RAID Manager Library 01.15.00 and newer versions.

    To enable Data Protector to use a disk array through a command device which is operating in the user authentication mode (available only with specific disk array models), you must use a specific RAID Manager Library version. For the version number and additional information, see the latest support matrices at http://www.hp.com/support/manuals.

  - HP Continuous Access (CA) P9000 XP and/or HP Business Copy (BC) P9000 XP license and microcode.

  - An appropriate multi-path device management software.

    The software must be installed on the application system and the backup system.

    HP-UX systems: HP Secure Path (HP-UX)

    On HP-UX 11.31 systems, the multi-path device management software is not required since the operating system has native device multi-pathing capability.

    Linux systems: HP Device Mapper Multipath Enablement Kit for HP Disk Arrays 4.2.0 or newer version

    To configure the installed multi-path device management software:

    1. Start the multipath daemon.
    2. Run the following command to configure the daemon so that it gets started during system startup:

       Red Hat Enterprise Linux: `chkconfig multipathd on`

       SUSE Linux Enterprise Server: `chkconfig boot.multipath on`

    Windows systems: HP MPIO Full Featured DSM (Device Specific Module) for HP P9000 XP Disk Array Family

  Data Protector licenses and components:

  - Appropriate zero downtime backup extension and instant recovery extension licenses-to-use (LTU).
  - HP StorageWorks P9000 XP Agent on the application system and the backup system.

For installation and upgrade instructions and licensing information, see the *HP Data Protector Installation and Licensing Guide*.

- Make sure that the same operating system version is installed on the application system and the backup system.
- If the application system and the backup system reside in a Data Protector cell with secured clients, ensure that access between both systems is allowed in both directions.
- Make sure the SAN environment and the P9000 XP Array storage system are properly configured:
  - Primary volumes (P-VOLs) are available to the application system.
  - Secondary volumes (S-VOLs) of the desired type (mirror, snapshot storage volume) are available to the backup system.
  - A pair relationship is defined between both sets of volumes (LDEVs) with HP P9000 XP Remote Web Console (formerly known as HP Command View XP).
  - LUNs are assigned to the respective ports.
- On HP-UX 11.31 systems, if you use VxVM disk groups, enable the legacy Device Special Files format.

Prerequisites for Linux systems

- For each configured S-VOL, follow the steps:
  1. Put the corresponding LDEV pair into the SUSPENDED state, that is, suspend the pair relationships between the S-VOL and its P-VOL.
  2. If multiple S-VOLs are in a pair relationship with its P-VOL, change the UUID of the S-VOL by running the command `pvchange -u` *`PVName`* on the backup system, where *`PVName`* is the LVM physical volume name of the S-VOL.

Prerequisites for Windows systems

- On Windows Server 2003 and Windows Server 2008 systems, disable the operating system option **Automatic mounting of new volumes**. In the Command Prompt window, run the command `mountvol /N`.
- Do not manually mount target volumes that were created by Data Protector.

For additional prerequisites for using HP P9000 XP Disk Array Family with the Data Protector Microsoft Volume Shadow Copy Service integration, see the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

Limitations

- In cluster environments, the backup system must not be in the same cluster with the application system. Additionally, the backup system cannot be the cluster virtual server, it can only be a cluster node.
- Zero downtime backup using snapshots is only supported in HP Business Copy (BC) P9000 XP and HP Continuous Access + Business Copy (CA+BC) P9000 XP configurations.
- Instant recovery is only supported in HP Business Copy (BC) P9000 XP configurations.
- Using split mirror restore, you can only restore filesystems and disk images backed up in HP Business Copy (BC) P9000 XP configurations, including their single-host (BC1) implementations. Other Data Protector backup object types are not supported.

For information on any of the following items, see the *HP Data Protector Product Announcements, Software Notes, and References*:

- general Data Protector and integration-specific limitations
- supported platforms and integrations
- supported backup and connectivity topologies

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

ZDB database – XPDB

ZDB database for the Data Protector HP P9000 XP Disk Array Family integration is referred to as **XPDB**. It keeps information about:

- LDEV pairs that are split (put into the SUSPENDED state). This information includes:
  - ID of the ZDB session that involved handling the LDEV pair.
  - LDEV, volume group, and filesystem configuration.
  - CRC information calculated during the session.
  - IR flag (indicating that the target volume can be used for instant recovery)
- Filesystem and volume management system information.

The information is written to the XPDB when a LDEV pair is put into the SUSPENDED state, and is deleted from the XPDB when a LDEV pair is resynchronized (is put into the PAIR state). During resynchronization, prior version of data is overwritten.

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

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

XPDB resides on the Cell Manager in:

Windows Server 2008: `Data_Protector_program_data\db40\xpdb`

Other Windows systems: `Data_Protector_home\db40\xpdb`

UNIX systems: `/var/opt/omni/server/db40/xpdb`

Configuring the integration

Before you start with the configuration, make sure the prerequisites listed in “Introduction” (page 63) are fulfilled. In addition, do the following:

Solaris systems: Run the Sun format utility to label and format the paired LDEVs (on both the application system and the backup systems). For information, see the *HP Disk Array XP Operating System Configuration Guide: Sun Solaris*.

HP BC P9000 XP configurations: Connect the application system and the backup system to the same disk array unit.

When using first-level mirrors or snapshot volumes, primary LDEVs (P-VOLs) must be connected to the application system and the paired secondary LDEVs (S-VOLs) must be connected to the backup system.

HP CA P9000 XP 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 the P9000 XP Array 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 (HP CA+BC P9000 XP) configurations: Connect the application system to the MCU, and the backup system to the RCU.

Main LDEVs (P-VOLs) must be paired to remote volumes in the RCU (S-VOLs). S-VOLs also function as HP BC P9000 XP primary volumes (P-VOLs) and must be paired to local copies (HP BC P9000 XP S-VOLs):

- ***Windows systems:*** Connect only HP BC P9000 XP S-VOLs to the backup system.
- ***HP-UX systems:*** Connect only HP BC P9000 XP S-VOLs to the backup system. If HP CA P9000 XP S-VOLs are 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-VOLs. Re-import or re-create the volume groups to ensure the configuration is correct.
- ***Linux systems:*** Connect only HP BC P9000 XP S-VOLs to the backup system.

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 each secondary LDEV is connected to the system on a different bus (not the bus used for the primary LDEV).

1. Create the volume group with the LDEVs that have S-VOLs assigned using `vgcreate`. LVM mirror primary volumes must be the LDEVs that have their S-VOLs.
2. Extend the volume group with LDEVs that have no S-VOLs assigned using `vgextend`. LVM mirror secondary volumes must be the LDEVs that have no S-VOLs.

For more information on LVM mirroring, see the document *Managing Systems and Workgroups: A Guide for HP-UX System Administrators*.

To configure the integration:

- Set the P9000 XP Array command devices. See "Command device handling" (page 66).
- If needed, set the P9000 XP LDEV exclude file. See "P9000 XP LDEV exclude file" (page 68).
- To enable zero downtime backup and instant recovery sessions that involve a disk array which is operating in the user authentication mode, configure the user authentication data. See "Configuring the user authentication data" (page 67).

Command device handling

HP P9000 XP Disk Array Family **command devices** are dedicated volumes in the disk arrays which act as the interface between management applications and the storage systems. They cannot be used for data storage and only accept requests for operations that are then executed by the disk arrays. A command device is needed and used by any process requiring access to a P9000 XP Array. Data about all command devices detected by Data Protector is stored in the XPDB for the purpose of avoiding concurrent overallocation of each particular command device.

Whenever a ZDB session is started, the Data Protector HP StorageWorks P9000 XP Agent queries the XPDB for a list of command devices, and updates it if needed. When the first ZDB session is started, the HP StorageWorks P9000 XP Agent generates a list of command devices connected to every application and backup system in the cell. All subsequent sessions automatically update the list if the configuration of command devices has changed.

Every command device is assigned an instance number (starting from 301) and the system name (hostname) having access to it. If a command device can be accessed from more than one system, the HP StorageWorks P9000 XP Agent recognizes that the command device is assigned to another system; such a command device-hostname combination gets the next available instance number.

Thus, every P9000 XP Array storage system attached to the application and backup systems has a list of command devices and systems having access to them (together with an instance number).

Below is an example of command device entries in the 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
```

To be able to control which command device and instance number should be used on a specific system, you can disable the automatic update of the command device list in the XPDB. To disable the automatic update:

1. Set the `SSEA_QUERY_STORED_CMDDEVS` `omnirc` variable to `1`.
2. Use the `omnidbxp` command to manually add, list, remove, and update the command devices.

For the command syntax and examples, see the *HP Data Protector Command Line Interface Reference* or the `omnidbxp` man page

If you decide to disable the automatic update, the initial list of command devices is still created in the XPDB during the first ZDB session. For subsequent backup sessions, the Data Protector HP StorageWorks P9000 XP Agent behavior is as follows:

- Whenever an application system or a backup system needs access to the P9000 XP Array storage system during a session, 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, the command device is used by the system until the end of the session, and the list of command devices is used for all consecutive sessions.

Configuring the user authentication data

Specific disk array models of the HP P9000 XP Disk Array Family provide increased security with authorization verification that involves user and resource groups, roles, and user authenticity verification. Authorization verification is enabled by a special operating mode called **user authentication mode**. When an application, for example the Data Protector HP StorageWorks P9000 XP Agent, communicates with a disk array which is operating in this mode, the application must supply appropriate user credentials in order for queries and modifications of the disk array configuration or its resources to succeed. On the disk arrays on which the conventional operating mode is still available for compatibility reasons, the user authentication mode is disabled by default.

Authorization system of a disk array on which the user authentication mode is available defines a fixed set of roles that belong to different task groups: security-related, storage-related, and maintenance-related tasks. It assigns a particular subset of roles and a particular set of resource groups to each user group. While there are several preconfigured user groups, which can be used immediately, you can easily create additional ones. Each disk array user account can belong to multiple user groups. Similarly, each user group can have multiple resource groups assigned, and each resource group can belong to multiple user groups. Each time an application attempts to start a specific operation on a specific resource of the disk array, the authorization system first determines the user account based on the supplied user credentials. It then checks if any user group the user account belongs to is allowed to perform the operation on the resource. If user credentials are not supplied, the disk array always rejects to execute the operation.

IMPORTANT: The operating mode setting is actually a command device property. For example, if a particular backward compatible disk array has two command devices configured, one can operate in the conventional mode and the other in the user authentication mode. It therefore depends on the configuration of the command device used whether the application should supply user credentials to successfully start the requested operation.

For more information on using the HP P9000 XP Disk Array Family authorization system, see the *HP P9000 Remote Web Console User Guide* and other parts of the HP P9000 XP Disk Array Family documentation set.

User authentication data and the XPDB

To enable the HP StorageWorks P9000 XP Agent to perform zero downtime backup (ZDB) and instant recovery (IR) sessions using a command device for which the user authentication mode is enabled, you must add appropriate user credentials to the ZDB database (XPDB) in advance. The credentials must belong to a disk array user account which has the *Local Copy*, the *Remote Copy*, or both roles assigned, depending on the HP P9000 XP Disk Array Family configuration. The HP StorageWorks P9000 XP Agent then reads the credentials from the XPDB each time such a ZDB or IR session is started. User credentials are bound to a specific disk array serial number. For each particular serial number, you can add user credentials of a single disk array user account. To add and manage user credentials in the XPDB, use the Data Protector `omnidbxp` command.

Configuration procedure

To properly add the required user credentials for a specific disk array that will be involved in the ZDB and IR sessions, follow the steps:

1. Identify the serial number of the disk array.
2. Identify which disk array volumes (LDEVs) will be involved in the ZDB and IR sessions.
3. Identify which disk array user group has been granted adequate access to all volumes that you identified in the previous step
4. Choose a disk array user account that belongs to the disk array user group from the previous step. Identify and write down its user name and password that you will need in the next step.
5. Using the `omnidbxp -user -add` command, add the user name and password that you acquired in the previous step to the XPDB, providing the disk array serial number you identified in step 1 of this procedure.

   For command syntax and usage examples, see the `omnidbxp` reference page in the *HP Data Protector Command Line Interface Reference* or the `omnidbxp` man page.
6. Using the `omnidbxp -user -check` command, verify that the HP StorageWorks P9000 XP Agent can connect to the disk array using the configured user authentication data.

For more information on performing other tasks related to management of user credentials in the XPDB, see the `omnidbxp` reference page in the *HP Data Protector Command Line Interface Reference* or the `omnidbxp` man page.

P9000 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 contains an excluded LDEV.

Disabled secondary LDEVs (S-VOLs) are listed in the P9000 XP LDEV exclude file on the Cell Manager:

Windows Server 2008: `Data_Protector_program_data\db40\xpdb\exclude\XPexclude`

Other Windows systems: `Data_Protector_home\db40\xpdb\exclude\XPexclude`

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

Secondary LDEVs (S-VOLs) listed in this file must be the 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

For the command syntax and the examples of manipulating the exclude file, see the *HP Data Protector Command Line Interface Reference* or the omnidbxp man page. The file syntax and the example are as follows.

Syntax

```
#
#     HP Data Protector A.06.20
#     HP P9000 XP Disk Array Family 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 Data Protector A.06.20
#     HP P9000 XP Disk Array Family LDEV Exclude File
#
[35241]
3603, 3610, 3620-3625 # Some excluded LDEVs
2577 #
2864-3527 #
#
#     End of file
```

Automatic configuration of the 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 a ZDB-to-tape or ZDB-to-disk+tape session.

For more information on the backup system mountpoint creation, see "Backup system mount point creation" (page 150).

Maintaining the integration

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

- available zero downtime backup sessions
- backup system LDEVs involved in a particular session
- backup system LDEVs stored in the XPDB
- XPDB information about particular LDEV pairs

You can retrieve the information stored in the XPDB using the omnidbxp command. For the command syntax and usage examples, see the *HP Data Protector Command Line Interface Reference* or the omnidbxp man page.

6 Backup

Introduction

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

You should be familiar with the HP P9000 XP Disk Array Family concepts and procedures and basic Data Protector ZDB and instant recovery functionality. See the HP P9000 XP Disk Array Family documentation and the *HP Data Protector Zero Downtime Backup Concepts Guide*.

ZDB types

Using the P9000 XP Array integration, you can perform:

- **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 ZDB backup specification, and **To disk** is selected when running/scheduling a backup.

  ZDB to disk is only possible using the HP BC P9000 XP configuration.

- **ZDB to tape**

  The replica produced is streamed to backup media, typically tape, according to the tape backup type you have selected (Full, Incr, Incr1-9).

  This replica is deleted after backup if the option **Keep the replica after the backup** is *not* selected in a ZDB 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 according to the tape backup type you have selected (Full, Incr, Incr1-9). 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 ZDB backup specification, and **To disk+tape** is selected when running/scheduling a backup.

  ZDB to disk+tape is only possible using the HP BC P9000 XP configuration.

Replica types

Using the P9000 XP Array integration, you can create the following replica types:

- split mirror

  This replica type is supported by all disk array models of the HP P9000 XP Disk Array Family that are officially supported by Data Protector.

- snapshot

  This replica type is supported by specific P9000 XP disk array microcode versions and specific RAID Manager Library versions only. For details, see the prerequisite list in "Introduction" (page 63).

You cannot directly select a replica type when configuring a Data Protector ZDB backup specification. You must chose the replica type in advance when creating secondary LDEVs (S-VOLs) with HP P9000 XP Remote Web Console (formerly known as HP Command View XP). During ZDB sessions, the Data Protector HP StorageWorks P9000 XP Agent always uses the S-VOLs (the target volumes specified in the ZDB backup specification) in the same way, regardless of their type – mirror or snapshot storage volume. Thus, you can even create replica sets of which specific replicas are mirror copies and others are snapshots.

In general, both replica types are available for all ZDB types, for instant recovery, and for split mirror restore. However, a specific limitation applies to the HP Continuous Access (CA) P9000 XP configurations. See the limitation list in "Introduction" (page 63).

Backup concepts

P9000 XP Array zero downtime backup consists of two phases:

1. The data from P-VOLs presented to the application system is synchronized with the S-VOLs presented to the backup system.

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

Creating backup specifications

Considerations

- Consider all limitations that apply to the Data Protector P9000 XP Array integration. See the *HP Data Protector Product Announcements, Software Notes, and References*, the *HP Data Protector Zero Downtime Backup Concepts Guide*, and the limitation list in "Introduction" (page 63).

Procedure

To create a ZDB backup specification for a disk array of the HP P9000 XP Disk Array Family using the Data Protector GUI (**Data Protector Manager**), follow the steps:

1. In the Context List, select **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**. Right-click **Filesystem** (for both object types: filesystem and disk image) and click **Add Backup**.

   The Create New Backup dialog box appears.

   In the Filesystem pane, select the **Blank Filesystem Backup** template or some other template which you might have created. For information on templates, see the online Help index: "backup templates".

   Select **Snapshot or split mirror backup** as **Backup type** and **HP StorageWorks P9000 XP** as **Sub type**. For descriptions of options, press **F1**.

   Click **OK**.

3. Under Client systems, select **Application system** and **Backup system**. If the application system is part of a server cluster, select the virtual server.
4. Under Mirror type, select the P9000 XP Array configuration, and specify a value for **MU number(s)**. The maximum number of replicas that can be created for the same source volumes is different for mirror copies and snapshots. Both limitations are imposed by the P9000 XP Array storage system.

   Figure 13 P9000 XP Array backup options

   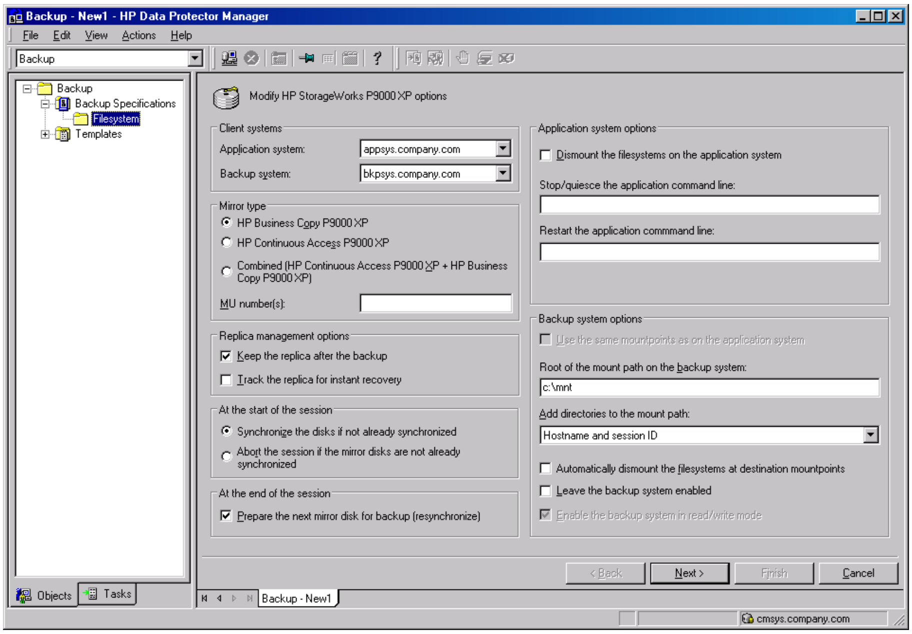

   Modify HP StorageWorks P9000 XP options

   Client systems

   Application system: appsys.company.com

   Backup system: bkpsys.company.com

   Mirror type

   ● HP Business Copy P9000 XP

   ○ HP Continuous Access P9000 XP

   ○ Combined (HP Continuous Access P9000 XP + HP Business Copy P9000 XP)

   MU number(s):

   Replica management options

   ☑ Keep the replica after the backup

   □ Track the replica for instant recovery

   At the start of the session

   ● Synchronize the disks if not already synchronized

   ○ Abort the session if the mirror disks are not already synchronized

   At the end of the session

   ☑ Prepare the next mirror disk for backup (resynchronize)

   Application system options

   □ Dismount the filesystems on the application system

   Stop/quiesce the application command line:

   Restart the application commmand line:

   Backup system options

   □ Use the same mountpoints as on the application system

   Root of the mount path on the backup system: c:\mnt

   Add directories to the mount path: Hostname and session ID

   □ Automatically dismount the filesystems at destination mountpoints

   □ Leave the backup system enabled

   ☑ Enable the backup system in read/write mode

5. Under Replica management options, select the desired options.

   ZDB to disk, ZDB to disk+tape:

   Select the option **Track the replica for instant recovery** to enable instant recovery.

   NOTE: You can choose a ZDB-to-disk session or a ZDB-to-disk+tape session by selecting an appropriate value for the **Split mirror/snapshot backup** option when running or scheduling a ZDB session based on this ZDB backup specification. See "Scheduling ZDB sessions" (page 124).

   ZDB to tape:

   Leave the option **Track the replica for instant recovery** cleared.

   To preserve the replica on the disk array after the ZDB session, leave the option **Keep the replica after the backup** selected. To remove the replica after the session, clear this option.
6. Under At the start of the session and At the end of the session, specify how states of the source volumes and the corresponding target volumes are handled during zero downtime backup sessions.
7. Specify other zero downtime backup options as desired. For information, see "Backup options" (page 74) or press **F1**.
8. Select the desired backup objects.

   Filesystem backup: Expand the application system and select the objects to be backed up. Note that all drive letters or mount points that reside on the system are displayed. You must select only the objects that reside on the disk array, otherwise the ZDB session will fail.

IMPORTANT: To ensure that instant recovery succeeds and the environment is consistent after instant recovery, select all volumes on a disk (Windows systems) or all logical volumes of a volume group (UNIX systems) to be backed up. Even if you do not select an entire disk or volume group, the backup will succeed, but instant recovery may experience issues during configuration check of the environment. The configuration check can be disabled by clearing the option **Check the data configuration consistency** in the GUI or not specifying the option `-check_config` in the CLI when preparing for an instant recovery session. If this option is cleared (GUI) or not specified (CLI), the entire disk or volume group will be overwritten during instant recovery.

Click **Next**.

Disk image backup: Click **Next**.

9. 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 desired number of mirrors by clicking **Add mirror** or **Remove mirror**. Select separate devices for the backup image and each mirror.

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

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

   Click **Next**.

10. In the Backup Specification Options group box, click **Advanced** and then the **HP StorageWorks P9000 XP** tab to open the P9000 XP Array backup options pane.

    You can specify Application system options and modify all other options, except **Application system** and **Backup system** (note that you can change them after you save the ZDB backup specification). See “Backup options” (page 74) or press **F1**.

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

    Windows systems: To configure a ZDB backup specification for incremental ZDB sessions, select the **Do not use archive attribute** filesystem option in the WinFS Options pane to enhance the incremental ZDB behavior. For details, see “Backup options” (page 74)

    Click **Next**.

11. Following the wizard, open the scheduler to schedule the ZDB sessions. For more information, see “Scheduling ZDB sessions” (page 124) or press **F1**.

    Click **Next**.

12. In the Backup Object Summary page, specify additional options.

    Filesystem backup: You can modify options for the listed objects by right-clicking an object and then clicking **Properties**. For information on the object properties, press **F1**.

    Disk image backup: Follow the steps:

    a. Click **Manual add** to add disk image objects.
    b. Select **Disk image object** and click **Next**.
    c. Select the client system. Optionally, enter the description for your object. 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.

       Windows systems:

       Use the format

       ```
       \\.\PHYSICALDRIVE#
       ```

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

For information on how to identify current disk numbers (physical drive numbers), see the online Help index: “disk image backups”.

HP-UX and Solaris systems:

Specify a disk image section:

`/dev/rdsk/`*`filename`*, for example: `/dev/rdsk/c2t0d0`

On HP-UX 11.31 systems, the new naming system can be used:

`/dev/rdisk/disk#`, for example `/dev/rdisk/disk2`

Specify a raw logical volume section:

`/dev/`*`vgnumber`*`/rlvol`*`Number`*, for example: `/dev/vg01/rlvol1`

Linux systems:

Specify a disk image section:

`/dev/`*`Filename`*, for example: `/dev/dm-10`

IMPORTANT: To ensure that instant recovery succeeds and the environment is consistent after instant recovery, select all volumes on a disk (Windows systems) or all logical volumes of a volume group (UNIX systems) to be backed up. Even if you do not select an entire disk or volume group, the backup will succeed, but instant recovery may experience issues during configuration check of the environment. The configuration check can be disabled by clearing the option **Check the data configuration consistency** in the GUI or not specifying the option `-check_config` in the CLI when preparing for an instant recovery session. If this option is cleared (GUI) or not specified (CLI), the entire disk or volume group will be overwritten during instant recovery.

f. Click **Finish**.

Click **Next**.

13. Save your ZDB backup specification.

For information on scheduling ZDB sessions and starting interactive ZDB sessions, see “Scheduling ZDB sessions” (page 124).

NOTE: Backup preview is not supported.

Backup options

The following tables describe the P9000 XP Array and ZDB related backup options. See also “P9000 XP Array integration” (page 149).

Table 13 Client systems

Application system	The system on which the application runs. In cluster environments, specify the virtual server hostname (rather than the physical node hostname).
Backup system	The system to which your data will be replicated (backed up). In ZDB-to-disk+tape and ZDB-to-tape sessions, the backup data is copied from this system to a backup device.

Table 14 Mirror type

HP Business Copy P9000 XP	Select this option to configure a ZDB backup specification for the HP P9000 XP Disk Array Family configuration HP Business Copy P9000 XP. Default: selected.
HP Continuous Access P9000 XP	Select this option to configure a ZDB backup specification for the HP P9000 XP Disk Array Family configuration HP Continuous Access P9000 XP.

Table 14 Mirror type *(continued)*

	Default: not selected.
Combined (HP Continuous Access P9000 XP + HP Business Copy P9000 XP)	Select this option to configure a ZDB backup specification for the HP P9000 XP Disk Array Family combined configuration HP Continuous Access P9000 XP + HP Business Copy P9000 XP. Default: not selected.

MU number(s)	This option is only available if the HP P9000 XP Disk Array Family configuration HP Business Copy P9000 XP is selected.
	This option defines the mirror unit (MU) number(s) of a replica or a replica set from which the Data Protector HP StorageWorks P9000 XP Agent, according to the replica set rotation, selects the replica to be used in the zero downtime backup session. The replica selection rule is described in the *HP Data Protector Zero Downtime Backup Concepts Guide*. The maximum number of replicas that can be created for the same source volumes is different for mirror copies and snapshots. Both limitations are imposed by the HP P9000 XP Disk Array Family storage system.
	You can specify one or more non-negative integer numbers, one or more ascending ranges of such numbers, or any combination of both. Use a comma as the separator character. Examples:
	`5`
	`7-9`
	`4,0,2-3`
	When a sequence is specified, it does not define the order in which the replicas are used.
	Default: `0` (nothing is specified).

Table 15 Replica management options

Keep the replica after the backup	If configuring a ZDB to tape, select this option to keep the replica on the disk array after the zero downtime backup session. The replica becomes part of a replica set (specify a value for the option **MU number(s)**). Unless the additional option **Track the replica for instant recovery** is selected, the replica is *not* available for instant recovery. If this option is not selected, the replica is removed at the end of the session. If the option **Track the replica for instant recovery** is selected, this option is automatically selected and cannot be changed. Default: selected.
Track the replica for instant recovery	This option is only available if the HP P9000 XP Disk Array Family configuration **HP Business Copy P9000 XP** is selected. Select this option to perform a ZDB-to-disk or ZDB-to-disk+tape session and leave the replica on the disk array to enable instant recovery. If this option is not selected, you cannot perform instant recovery using the replica created or reused in this session. **CAUTION:** If you select this option, do not manually resynchronize the affected mirrors and do not empty the volumes used for snapshot storage. Otherwise, instant recovery will not be possible. Default: not selected.

Table 16 At the start of the session

Synchronize the disks if not already synchronized	On the P9000 XP Array, primary volumes (source volumes) and their corresponding secondary volumes (target volumes) must be in the PAIR state to enable Data Protector zero downtime backup: mirrors must be synchronized and volumes to be used for snapshot storage must be empty.

Table 16 At the start of the session *(continued)*

	This option is automatically selected and cannot be changed if the option **Prepare the next mirror disk for backup (resynchronize)** is cleared. If this option is selected, all volumes of the replica to be used in the current ZDB session are put into the PAIR state with the corresponding source volumes at the start of the session: mirrors are resynchronized and volumes to be used for snapshot storage are made empty. Default: selected.
Abort the session if the mirror disks are not already synchronized	Available only if the option **Prepare the next mirror disk for backup (resynchronize)** is selected. The option is only applicable if at least one volume of the replica to be used in the current ZDB session is a mirror (or mirror copy). In the opposite case, Data Protector treats as if the option **Synchronize the disks if not already synchronized** is selected instead. If this option is selected and at least one volume of the replica to be used in the current ZDB session is not in the PAIR state with the corresponding source volume, the session fails. Default: not selected.

Table 17 At the end of the session

Prepare the next mirror disk for backup (resynchronize)	This option is only applicable if at least one volume of the replica to be used in the next ZDB session is a mirror (or mirror copy). In the opposite case, Data Protector behaves as if the option is not selected. If this option is selected, all volumes of the replica to be used in the next ZDB session are put into the PAIR state with the corresponding source volumes at the end of the current ZDB session: mirrors are resynchronized and volumes to be used for snapshot storage are made empty. If this option is not selected, the volumes of the replica to be used in the next ZDB session are left intact at the end of the current ZDB session. If this option is not selected, the **Synchronize the disks if not already synchronized** option is automatically selected, and the **Abort the session if the mirror disks are not already synchronized** option is not available. Default: selected.

Table 18 Application system options

Dismount the filesystems on the application system	Select this option to dismount the filesystems on the application system before replica creation and remount them afterwards. Additionally, when entire physical drives (on Windows systems) or entire disks or logical volumes (on UNIX systems) are selected as backup objects in a disk image backup specification, selecting this option will dismount and later remount all filesystems on these objects. If any of these filesystems cannot be dismounted, the backup session fails. If an integrated application (for example, Oracle Server) exclusively controls the data I/O on each physical drive, disk, or logical volume that will be backed up, the dismount operation is not needed. In such a case, you can leave this option cleared. Default: not selected.
Stop/quiesce the application command line	If a command is specified in this option, it is invoked on the application system immediately before replica creation. A use example is to stop applications not integrated with Data Protector. The command must reside on the application system in the directory `*Data_Protector_home*\bin` (Windows systems) or `/opt/omni/lbin` (UNIX systems). Do not specify the path to the command in this option. If the command fails, the command specified in the option **Restart the application command line** is not invoked. Thus, you may need to implement a cleanup procedure in the command specified in **Stop/quiesce the**

Table 18 Application system options *(continued)*

	application command line. If the omnirc variable ZDB_ALWAYS_POST_SCRIPT is set to 1, the command specified in the option **Restart the application command line** is always invoked.
Restart the application command line	If a command is specified in this option, it is invoked on the application system immediately after replica creation. A use example is to resume operation of applications not integrated with Data Protector. The command must reside on the application system in the directory *Data_Protector_home*\bin (Windows systems) or /opt/omni/lbin (UNIX systems). Do not specify the path to the command in this option.

Table 19 Backup system options

Use the same mountpoints as on the application system	This option is not available if the application system is also the backup system (a single-host configuration). If this option is selected, the paths to mount points used for mounting the filesystems of the replica on the backup system are the same as paths to mount points where source volume filesystems were mounted on the application system. If the mount points are already in use, the session fails. For such circumstances, you must select the option **Automatically dismount the filesystems at destination mountpoints** in order for the session to succeed. **Windows systems:** The drive letters must be available, otherwise the session fails. Default: not selected.
Root of the mount path on the backup system	This option is only available if the option **Use the same mountpoints as on the application system** is not selected. Specifies the root directory under which the filesystems of the replica are mounted. Where exactly the filesystems are mounted depends on how you define the option **Add directories to the mount path**. **NOTE:** For the SAP R/3 integration, the option is not applicable (the mount points created are always the same as on the application system). Defaults: **Windows systems:** c:\mnt **UNIX systems:** /mnt
Add directories to the mount path	This option is only available if the option **Use the same mountpoints as on the application system** is not selected. This option enables control over the created mount points. It defines which subdirectories will be created in the directory defined with the **Root of the mount path on the backup system** option. When Session ID is used in path composition, this guarantees unique mount points. Example for **Windows systems**: Root directory: C:\mnt Application system: applsys.company.com Backup session ID: 2008-02-22-4 Mount path on the application system: E:\disk1 If **Hostname** is selected: C:\mnt\applsys.company.com\E\disk1 If **Hostname and session ID** is selected: C:\mnt\applsys.company.com\2008-02-22-4\E\disk1 If **Session ID** is selected:

Table 19 Backup system options *(continued)*

	`C:\mnt\2008-02-22-4\E\disk1`
	If **Session ID and hostname** is selected:
	`C:\mnt\2008-02-22-4\applsys.company.com\E\disk1`
	NOTE:
	For the SAP R/3 integration, the option is not applicable (the mount points created are always the same as on the application system).
	Default: Hostname and session ID.

Automatically dismount the filesystems at destination mountpoints	If the mount points are in use (for example, volumes involved in the previous session may still be mounted) and this option is selected, Data Protector attempts to dismount the mounted filesystems. If the option is not selected and the mount points are in use, or if the option is selected and the dismount operation fails, the session fails. Default: not selected.
Leave the backup system enabled	This option is only available if the option **Keep the replica after the backup** is selected. If this option is selected, the filesystems remain mounted, the volume groups remain imported and active (UNIX systems), and the target volumes remain presented after the session. In this case, you can use the backup system for data warehousing purposes, but *not* for instant recovery. If the replica has to be reused later on (deleted, rotated out, or used for instant recovery), Data Protector automatically connects to the backup system, dismounts the filesystems, unpresents the target volumes, and clears the related logical structures on the backup system. At that point in time, if the filesystems are not mounted to the current backup system, Data Protector cannot perform a proper cleanup, and aborts the operation or the instant recovery session. If this option is not selected, Data Protector dismounts filesystems, exports volume groups (UNIX systems), and unpresents the target volumes on the backup system at the end of the ZDB session. Default: not selected.
Enable the backup system in read/write mode	This option is applicable to and can only be changed for UNIX systems only. On Windows systems, filesystems cannot be mounted in the read-only mode. Select this option to enable write access to volume groups and filesystems on the backup system. For backup purposes, it is sufficient to activate the backup system volume groups and mount the filesystems in the read-only mode. For other tasks, the read/write mode may be needed. Note that when this option is selected, the replica is open to modifications while the backup system is online. Consequently, data restored from such a replica includes all potential modifications. Defaults: **Windows systems:** selected. **UNIX systems:** not selected.

NOTE: In a particular ZDB session, the mount point paths to which filesystems of the replica are mounted on the backup system correspond the mount point paths to which source volumes were mounted on the application system if at least one of the following conditions is met:

- The GUI option **Use the same mountpoints as on the application system** is selected.
- The `omnirc` variable `ZDB_PRESERVE_MOUNTPOINTS` is set to `1`.

If the option **Use the same mountpoints as on the application system** is not selected, and the `omnirc` variable `ZDB_PRESERVE_MOUNTPOINTS` is set to `0`, the mount point paths are determined by the GUI options **Root of the mount path on the backup system** and **Add directories to the mount path**, and the `omnirc` variables `ZDB_MULTI_MOUNT` and `ZDB_MOUNT_PATH` are ignored.

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

Figure 14 ZDB session flow for filesystem backup objects

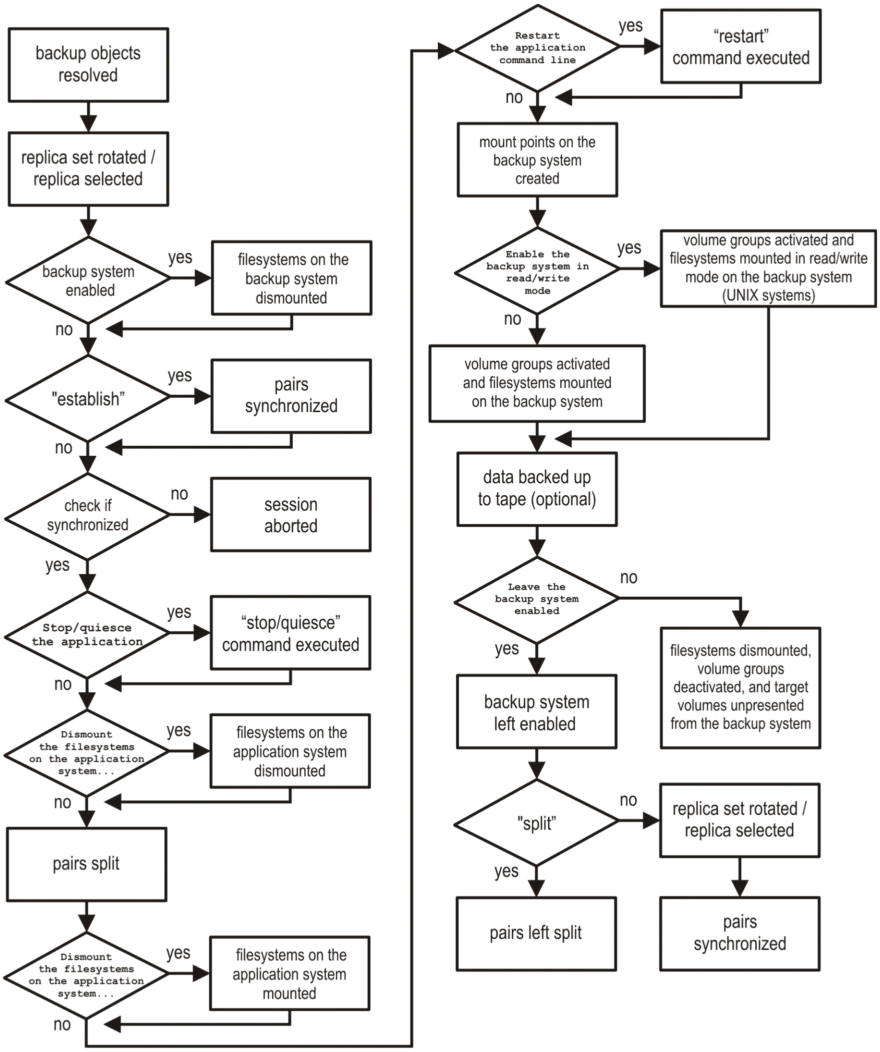
backup objects resolved
replica set rotated / replica selected
backup system enabled
yes
filesystems on the backup system dismounted
no
"establish"
yes
pairs synchronized
no
check if synchronized
no
session aborted
yes
Stop/quiesce the application
yes
"stop/quiesce" command executed
no
Dismount the filesystems on the application system...
yes
filesystems on the application system dismounted
no
pairs split
Dismount the filesystems on the application system...
yes
filesystems on the application system mounted
no
Restart the application command line
yes
"restart" command executed
no
mount points on the backup system created
Enable the backup system in read/write mode
yes
volume groups activated and filesystems mounted in read/write mode on the backup system (UNIX systems)
no
volume groups activated and filesystems mounted on the backup system
data backed up to tape (optional)
Leave the backup system enabled
no
filesystems dismounted, volume groups deactivated, and target volumes unpresented from the backup system
yes
backup system left enabled
"split"
no
replica set rotated / replica selected
yes
pairs left split
pairs synchronized

The "establish" and "split" checks depend on the P9000 XP Array zero downtime backup options listed in the table "Relation between particular zero downtime backup options and the "establish" and "split" checks" (page 79).

Table 20 Relation between particular zero downtime backup options and the "establish" and "split" checks

The option **Synchronize the disks if not already synchronized** is selected.	"establish" = yes
The option **Abort the session if the mirror disks are not already synchronized** is selected.	"establish" = no
The option **Prepare the next mirror disk for backup (resynchronize)** is selected.	"split" = no
The option **Prepare the next mirror disk for backup (resynchronize)** is cleared.	"split" = yes
No value or a single number is specified for the option **MU number(s)** or the option **Keep the replica after the backup** is selected.	"split" = yes

7 Restore

Introduction

This chapter describes configuring and running a filesystem or disk image restore of the data backed up using the P9000 XP Array 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" (page 110).
- Split mirror restore. See "Split mirror restore" (page 81).
- Instant recovery. See "Instant recovery" (page 85).

Table 21 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 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 more information on this restore type, see the online Help index: "restore".

TIP: You can improve the data transfer rate by connecting a backup device to the application system. For information on configuring backup devices, see the online Help index: "backups devices, configuring". For information on performing a restore using another device, see the online Help index: "selecting, devices for restore".

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. If the Data Protector EMC Symmetrix Agent is also installed on the *target* client system, select **StorageWorks P9000 XP restore**. Click **Next**.

Figure 15 StorageWorks P9000 XP restore

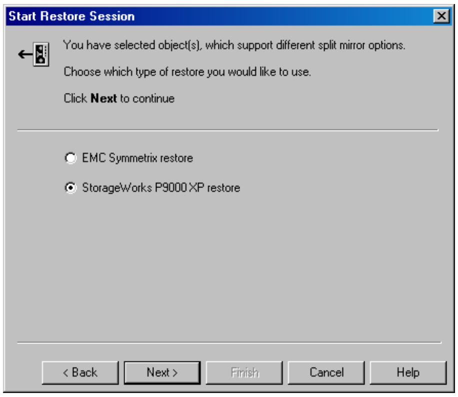

You have selected object(s), which support different split mirror options.

Choose which type of restore you would like to use.

Click **Next** to continue

○ EMC Symmetrix restore

● StorageWorks P9000 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 the restore.

Split mirror restore

Considerations

- Split mirror restore can be run with both replica types: split mirror and snapshot. The same split mirror restore procedure applies in both cases.
- You can start a split mirror restore session only after the preceding session using the same internal disks on the application system finishes with the disk pairs synchronization (the transition of the LDEV pairs into the PAIR state).

Split mirror restore process

Data is restored from backup media on LAN to the secondary LDEVs (S-VOLs), and then copied to the primary LDEVs (P-VOLs). The process 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. For more information, see the *HP Data Protector Zero Downtime Backup Concepts Guide*.
2. Preparing the application system and the backup system.
3. Restoring data from the backup media on LAN to the backup system and copying 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**.
5. Specify the report level and network load. Click **Next**.
6. If the Data Protector EMC Symmetrix Agent is also installed on the *target* client system, select **StorageWorks P9000 XP restore**. Click **Next**.
7. Specify the split mirror restore options. See “StorageWorks P9000 XP split mirror restore options” (page 82). For more information, see “Split mirror restore options” (page 82).

Figure 16 StorageWorks P9000 XP split mirror restore options

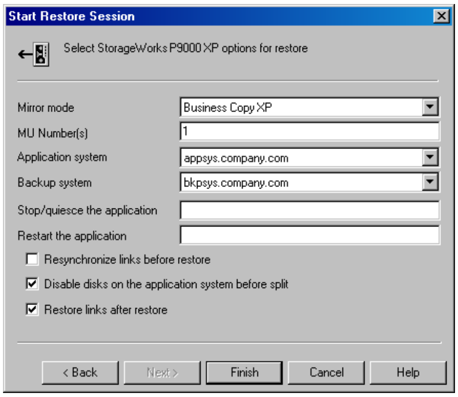

Start Restore Session

Select StorageWorks P9000 XP options for restore

Mirror mode	Business Copy XP
MU Number(s)	1
Application system	appsys.company.com
Backup system	bkpsys.company.com
Stop/quiesce the application	
Restart the application	

□ Resynchronize links before restore

☑ Disable disks on the application system before split

☑ Restore links after restore

8. Click **Finish** to start the split mirror restore.

NOTE: If LVM mirroring is used, a warning appears during the session, since the volume group LDEVs in the physical volume group on the application system do not have HP BC P9000 XP pairs assigned. This warning should be ignored.

For information on the general restore process, see the online Help index: “restore”.

Split mirror restore options

The following table explains the split mirror restore options.

Table 22 Split mirror restore options

Data Protector GUI	Function
Mirror mode	Selects a P9000 XP Array configuration. Only the HP Business Copy P9000 XP configuration is supported.

MU Number(s)	This option defines the mirror unit (MU) number(s) of a replica or a replica set from which the Data Protector HP StorageWorks P9000 XP Agent, according to the replica set rotation, selects the replica to be used in the restore session. The replica selection rule is described in the *HP Data Protector Zero Downtime Backup Concepts Guide*.
	You can specify one or more non-negative integer numbers, one or more ascending ranges of such numbers, or any combination of both. Use a comma as the separator character. Examples:
	`5`
	`7-9`
	`4,0,2-3`
	When a sequence is specified, it does not define the order in which the replicas are used.
	Default: `0` (nothing is specified).

Application system	Specifies the system to which your data will be restored. In cluster environments, specify the virtual server hostname (rather than the physical node hostname).
Backup system	Specifies the system to which your data will be restored from the backup media on LAN.
Stop/quiesce the application	Optionally specifies the command/script to be run before the LDEV pairs are split (put into the SUSPENDED state). The command/script must reside on the application system in the directory *`Data_Protector_home`*`\bin` (Windows systems) or `/opt/omni/lbin` (HP-UX, Solaris, and Linux systems). It can be used, for example, for stopping the application, dismounting the file systems that are not

Table 22 Split mirror restore options *(continued)*

Data Protector GUI	Function
	to be restored in the active session, but belong to the same volume group or disk, or preparing the volume group for deactivation. If this command/script fails, the command/script specified with the option **Restart the application** is not executed. Therefore, you need to implement a cleanup procedure in this command/script. Note that if the `omnirc` variable `ZDB_ALWAYS_POST_SCRIPT` is set to `1`, the command/script specified with the option **Restart the application** is always executed. For details, see "ZDB omnirc variables" (page 139).
Restart the application	Specifies the command/script to be run immediately after the LDEV pairs are resynchronized (put into the PAIR state). The command/script must reside on the application system in the directory *`Data_Protector_home`*`\bin` (Windows systems) or `/opt/omni/lbin` (HP-UX, Solaris, and Linux systems). It can be used, for example, for restarting the application or mounting the filesystems.
Resynchronize links before restore	Directs the Data Protector disk array agent to synchronize the LDEV pairs, that is, to copy the application data to the disks which store backup data. This is necessary to prepare the disks for restore and to enable consistent data restore. If the paired LDEVs have been split (put into the SUSPENDED state) before the restore, and only some files need to be restored, then this option updates the backup system. This will ensure that the correct data is resynchronized to the application system. If this option is not selected, the synchronization is not performed. Default: not selected.
Disable disks on the application system before split	Directs the Data Protector disk array agent to disable disks on the application system, that is, dismount the filesystems and deactivate the volume groups. This is performed before the LDEV pairs are split. The disks are enabled after the links are restored. Note that only filesystems selected for restore are dismounted. If other filesystems exist in the volume group or on the disk, appropriate commands/scripts must be used to dismount these filesystems (specified with the options **Stop/quiesce the application** and **Restart the application**). You must always select this option for restore when you want to copy data from the backup system to the application system, that is, to incrementally restore links. The application system disks have to be disabled to provide data integrity after the links are restored, that is, data is copied. Default: selected.
Restore links after restore	Directs the Data Protector disk array agent to incrementally restore the links for the LDEVs that Data Protector has successfully restored to the backup system. The HP StorageWorks P9000 XP Agent also incrementally re-establishes links for the LDEVs for which the Data Protector restore failed. Default: selected.

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

Figure 17 Filesystem split mirror restore flow

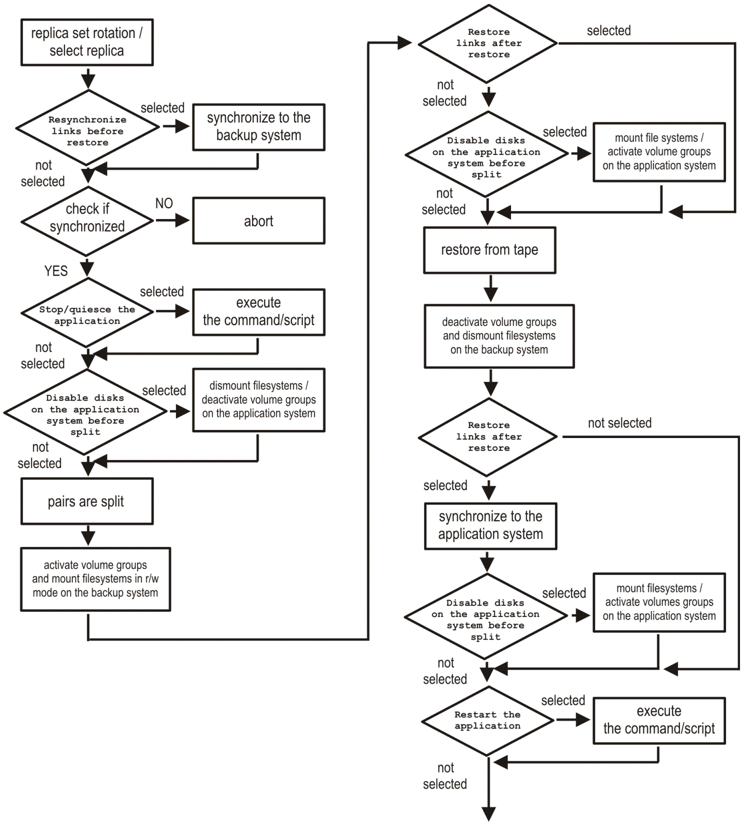
replica set rotation / select replica
Resynchronize links before restore
selected
synchronize to the backup system
not selected
check if synchronized
NO
abort
YES
Stop/quiesce the application
selected
execute the command/script
not selected
Disable disks on the application system before split
selected
dismount filesystems / deactivate volume groups on the application system
not selected
pairs are split
activate volume groups and mount filesystems in r/w mode on the backup system
Restore links after restore
selected
not selected
Disable disks on the application system before split
selected
mount file systems / activate volume groups on the application system
not selected
restore from tape
deactivate volume groups and dismount filesystems on the backup system
Restore links after restore
not selected
selected
synchronize to the application system
Disable disks on the application system before split
selected
mount filesystems / activate volumes groups on the application system
not selected
Restart the application
selected
execute the command/script
not selected

Split mirror restore in a cluster

Split mirror restore in configurations with the application system in MC/ServiceGuard or a Microsoft server cluster requires 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. For details, see "Split mirror restore procedure" (page 111).

   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 the 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 the source volumes, without involving a backup device. All data (entire volume group on UNIX systems or entire disk on Windows systems) in the replica is restored. For instant recovery concepts, see the *HP Data Protector Zero Downtime Backup Concepts Guide*.

You can perform instant recovery using the Data Protector GUI (see "Instant recovery using the GUI" (page 86)) or CLI (see "Instant recovery using the CLI" (page 87)).

Considerations

- Only first-level mirrors or snapshot volumes can be used for instant recovery. Second-level (cascading) mirrors and snapshot volumes are not supported.
- Instant recovery can be run with both replica types: split mirror and snapshot. The same instant recovery procedure applies in both cases.
- When instant recovery starts, Data Protector disables the application system. This includes dismounting filesystems and exporting volume groups (on UNIX systems only). Before this is done, filesystems' and volume groups' status is checked, and only mounted filesystems and imported volume groups are dismounted and exported. At the end of the session, dismounted filesystems are mounted and exported volume groups are imported to the same mount points as were used during backup.
- You cannot start several instant recovery sessions using the same disk on the application system at once. A session can be started only after the preceding session using the same source volume on the application system finishes synchronization.

ⓘ **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 P9000 XP Array instant recovery considerations and limitations, see the *HP Data Protector Product Announcements, Software Notes, and References* and the *HP Data Protector Zero Downtime Backup Concepts Guide*.

ⓘ **IMPORTANT:** Instant recovery does not recover databases or applications. It only synchronizes the primary LDEVs on the application system with the secondary LDEVs on the backup system. To recover a database or application data, you need to perform additional steps.

Prior to instant recovery, Data Protector:

- checks the volume group configuration (on UNIX systems only)
- verifies the replica

These steps assure that data in the replica has been left intact after the replica was created. If either of these steps fails, the instant recovery session fails.

Once the replica is restored, it can be left unchanged or resynchronized, depending on the selected instant recovery options. For information, see "Instant recovery options" (page 87).

Instant recovery procedure

Prerequisites

- Before performing a disk image instant recovery, manually dismount the disks before the 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 whose replica you want to use for instant recovery. This can be done by selecting:
   - a zero downtime backup session ID and the corresponding ZDB backup specification name:

     In the Scoping Pane, expand **Restore Sessions** and select the session from a list of ZDB-to-disk and ZDB-to-disk+tape sessions.
   - a backup object type, a ZDB backup specification name, and a ZDB session ID:

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

     Backup object types are displayed. Examples of backup object types are filesystem, disk Image, SAP R/3, and Microsoft SQL Server.

     b. Expand the backup object type for which you want to perform instant recovery.

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

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

Figure 18 Selecting a session for instant recovery

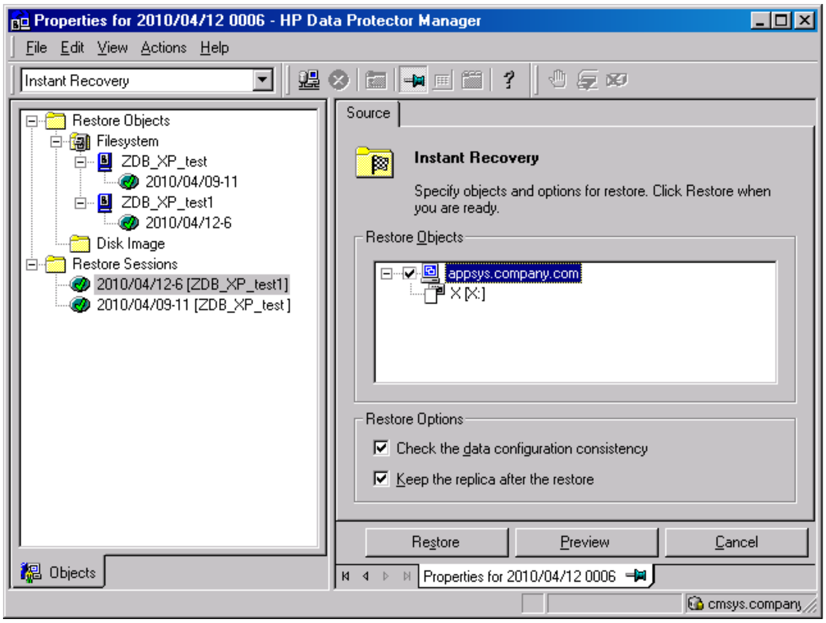

Properties for 2010/04/12 0006 - HP Data Protector Manager

Instant Recovery

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

Restore Options

☑ Check the data configuration consistency

☑ Keep the replica after the restore

3. In the Scoping Pane, click the desired ZDB session.

   The application system and its mount points/drive letters backed up during the selected session are displayed.
4. Select the application system and specify the instant recovery options. For details, see "Instant recovery options" (page 87).

5. Click **Restore** to start the instant recovery, or **Preview** to preview it. Note that preview is available only for filesystem backup objects.
6. Select **Start Restore Session** to start instant recovery, or **Start Preview Session** to start the preview. Click **OK**.

NOTE: 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 GUI 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 and ZDB-to-disk+tape sessions, identified by the session ID:

   ```
   omnidbxp -ir -session -list
   ```

   From the output, select the backup session whose replica you want to use for instant recovery.

2. Run:

   ```
   omnir -host ClientName -session SessionID -instant_restore
   [INSTANT_RECOVERY_OPTIONS]
   ```

   where the meaning of the options is as follows:

ClientName	The application system name.
SessionID	The backup session ID (Step 1 of this procedure).

   For *INSTANT_RECOVERY_OPTIONS*, see "Instant recovery options" (page 87).

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

Instant recovery options

Table 23 Instant recovery options

Data Protector GUI/CLI	Function
Check the data configuration consistency / `-check_config`	If this option is selected in the GUI or specified in the CLI, the current configuration of the participating volume groups is compared with the volume group configuration as it was during the ZDB session and which is stored in the XPDB. If the configuration has changed since the ZDB session, the instant recovery session aborts. Additionally, the CRC information for the selected LDEV pairs stored in the XPDB is compared to the current CRC information. If the items compared do not match, the instant recovery session aborts. A RAID Manager Library flag, which is set whenever the selected secondary LDEV is accessed/changed by any process (including non-Data Protector processes) is checked. If the flag is set, the session fails with an appropriate warning. **MC/ServiceGuard clusters:** When instant recovery is performed to some other node than the one from where the volumes were backed up, the current volume group configuration on the target node is different from the volume group configuration kept in the XPDB. In such a case, the XPDB volume group configuration data is replaced by the current volume group configuration data on the target node, and the session does not abort. When performing instant recovery to some other node than the one that was backed up, select (GUI) or specify (CLI) this option. Default (GUI): selected.
Keep the replica after the restore / `-keep_version`	If this option is selected in the GUI or specified in the CLI, the LDEV pairs involved in the current instant recovery session are split and left in the SUSPENDED state after the restore of data is complete. In the opposite case, the LDEV pairs are left in the PAIR state.

Table 23 Instant recovery options *(continued)*

Data Protector GUI/CLI	Function
	Even if the instant recovery is successful, it is recommended to keep the replica until the next ZDB session. **Linux systems:** This option must be selected (GUI) or specified (CLI) if the replica set consists of more than a single replica. Default (GUI): selected.

Instant recovery and LVM mirroring

If you use an LVM mirroring configuration, perform the following instant recovery steps:

1. Reduce all logical volumes which have LVM mirrors, specifically, reduce or remove the mirrors that reside on primary LDEVs that are not paired with secondary LDEVs on the P9000 XP Array. This ensures that restored data cannot be accidentally overwritten by a synchronization of the LVM mirror.

   Rebuild the LVM mirroring environment to the previous configuration.
2. Start the instant recovery session.
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

For information about and instructions for instant recovery in configurations with the application system in MC/ServiceGuard or a Microsoft server cluster, see "Instant recovery in a cluster" (page 131).

8 Troubleshooting

Before you begin

This chapter lists general checks and verifications plus problems you may encounter when using the P9000 XP Array integration. For general Data Protector troubleshooting information, see the *HP Data Protector Troubleshooting Guide*.

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

Checks and verifications

- On the application and backup systems, examine system errors logged into:

  Windows Server 2008: `Data_Protector_program_data\log\debug.log`

  Other Windows systems: `Data_Protector_home\log\debug.log`

  HP-UX, Solaris, and Linux systems: `/var/opt/omni/log/debug.log`
- Ensure that RAID Manager Library is correctly installed on the application system and the backup system and is accessible by the HP StorageWorks P9000 XP Agent, that is, listed in the library path.

General problems

Problem

A process stops responding when attempting to read data from a secondary LDEV (SVOL) in the PAIR state

When a secondary LDEV (S-VOL) is presented to the backup system, the LDEV pair it belongs to is in the PAIR state, and a process attempts to read the data from the secondary LDEV, the process stops responding. If such a problem occurs, all other processes attempting to read the data from such a secondary LDEV, for example the `pvscan` command, are affected, too.

Action

Unpresent the secondary LDEV from the backup system or unzone them.

△ **CAUTION:** Under the described circumstances, you should not try to restart the backup system before resolving the issue. Doing so may result in a data loss due to corruption of the involved file system. If the file system is corrupt, the backup system may even not be able to start up.

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 P9000 XP Agent` integration module is installed on the application and backup systems. To do that, open the `cell_info` file located on the Cell Manager in the following directory:

Windows Server 2008: `*Data_Protector_program_data*\Config\server\cell\cell_info`

Other Windows systems: `*Data_Protector_home*\Config\server\cell\cell_info`

UNIX systems: `/etc/opt/omni/server/cell/cell_info`

File contents should look similar to:

```
-host "sap001.company.com" -os "HP s800 HP-ux-11.10" -cc A.06.20 -da
A.06.20 -ssea A.06.20
```

```
-host "sap002.company.com" -os "HP s800 HP-ux-11.10" -cc A.06.20 -da
A.06.20 -ma A.06.20 -ssea A.06.20
```

Problem

On the application system, dismounting of a filesystem fails

Action

In the `Stop/quiesce the application command line` or `Stop/quiesce the application` script, stop all processes using the filesystem.

Use appropriate operating system tools or utilities to get a list of processes that are using the filesystem in order to identify any processes that lock the filesystem. For example, `lsof` on HP-UX.

Problem

On the backup system, mounting of a filesystem fails

Action

Check that the mountpoint directory exists on the backup system and that it is writable. On Windows Server 2003 and Windows Server 2008 systems, if the option **Automatically dismount the file systems on the application system** is selected, check if any processes are locking the filesystem.

Problem

Pair synchronization fails (the split fails)

To successfully split the pair, the HP StorageWorks P9000 XP Agent first checks its status. Pairs can only be split (in PSUS/SSUS status) after they are synchronized (in PAIR status). HP StorageWorks P9000 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.

For more information, see "ZDB omnirc variables" (page 139).

Problem

P-VOL has no paired S-VOL

Action

Check the P9000 XP Array configuration as follows:

HP BC P9000 XP: All P-VOLs on the application system must have associated HP BC P9000 XP S-VOLs on the backup system.

HP CA P9000 XP: All P-VOLs on the application system must have associated HP CA P9000 XP S-VOLs on the backup system.

HP CA+BC P9000 XP: All P-VOLs on the application system must have associated HP CA P9000 XP S-VOLs on the backup system and all S/P-VOLs must have HP BC P9000 XP 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 P9000 XP instances manually. You can get a list of LDEVs from the backup session report. Alternatively, with newer models of the HP P9000 XP Disk Array Family, you can use also HP P9000 XP Remote Web Console (formerly known as HP Command View XP).

Problem

Missing details for a specific LDEV/MU# are reported

```
[Warning] From: SSEA@machine_app.company.com ""
Time: 17.10.2008. 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: 17.10.2008.
10:41:27
Resolving of backup objects on the application system completed.
[Normal] From: SSEA@machine_bu.company.com ""  Time: 17.10.2008. 10:41:27
Resolving backup objects on the backup system.
[Critical] From: SSEA@machine_bu.company.com ""  Time: 17.10.2008. 10:41:29
Resolving of backup objects on the backup system failed.
```

Action

1. 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 P9000 XP LDEV exclude file.
2. Restart the session.

Problem

Filesystems not resolved on the backup system

On Windows systems, in some initial configurations filesystems may not be resolved on the backup system. The filesystems do not show up at all, even after a manual pair or split operation is performed on the disk array.

Action

Using the device manager, remove the problematic disks from the disk array and rescan the backup system.

Problem

During a zero downtime backup session, when a second replica is selected from the replica set specified by the ZDB backup specification option MU number(s), the session fails

If more than one replica is specified in the ZDB backup specification option **MU number(s)**, and a ZDB session is run which, according to the replica selection rule, selects the second or any subsequent replica, the session fails.

Action

The problem may be related to the duplicate disk signatures assigned to the target volumes by the Windows operating system.

Perform the following:

1. Unpresent all involved target volumes from the backup system.
2. On the backup system, clean the Registry.

   Windows Server 2003: Run the Scrubber utility by invoking the `scrubber2003` command. You can download Scrubber from the website http://support.microsoft.com/kb/277222/.

   Windows Server 2008: Run the DiskPart utility by invoking the `diskpart` command. Inside the DiskPart shell, run the command `automount scrub`.
3. Put all involved P-VOL - S-VOL pairs into the SUSPENDED state.
4. Present the target volumes to the backup system.
5. Start the ZDB session once again.

Split mirror restore problems

Problem

Session fails with the following message:

```
[Major] From: SSEA@machine.company.com ""  Time: 17.10.2008. 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: 17.10.2008. 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: 16.10.2008. 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: 16.10.2008. 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: 16.10.2008. 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: 16.10.2008. 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 fails on HP-UX 11.31 in LVM mirroring environments

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

This problem occurs if the option `Check the data configuration consistency` is selected and is caused by the following:

If your HP-UX 11.23 application system was migrated to HP-UX 11.31, Device Special Files (DSFs) change from the legacy format to the new persistent DSF format. As a result of this change, your LVM configuration now refers to physical volumes in new format, which is checked during instant recovery.

Action

Disable the `Check the data configuration consistency` option for the backup objects that are part of the LVM mirroring configuration and restart the instant recovery session.

Part III HP P4000 SAN Solutions

This part describes how to configure the Data Protector HP P4000 SAN Solutions integration. For information on how to perform zero downtime backup and instant recovery using the HP P4000 SAN Solutions integration, see the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

9 Configuration

Introduction

This chapter describes the configuration of the Data Protector HP P4000 SAN Solutions integration.

Prerequisites

- Obtain and/or install:

  P4000 SAN Solutions licenses and components:

  - HP P4000 SAN/iQ software
  - HP P4000 Virtual SAN Appliance Software / HP P4000 Centralized Management Console
  - HP P4000 SAN Solutions DSM (Device Specific Module) for MPIO

  For installation instructions, see the HP P4000 SAN Solutions documentation. For information on supported product versions, see the latest support matrices at http://www.hp.com/support/manuals.

  Data Protector licenses and components:

  - Appropriate zero downtime backup extension and instant recovery extension licenses-to-use (LTU).
  - HP StorageWorks P4000 Agent on the application system and the backup system.

  For licensing information and installation and upgrade instructions, see the *HP Data Protector Installation and Licensing Guide*.

- Make sure the same operating system version is installed on the application system and the backup system.
- If the application system and the backup system reside in a Data Protector cell with secured clients, ensure that access between both systems is allowed in both directions.
- Source volumes must be created and presented to the application system and the backup system.

For additional prerequisites for using HP P4000 SAN Solutions with the Data Protector Microsoft Volume Shadow Copy Service integration, see the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

Limitations

- In cluster environments, the backup system must not be in the same cluster with the application system. Additionally, the backup system cannot be the cluster virtual server, it can only be a cluster node.

For information on either of the following items, see the *HP Data Protector Product Announcements, Software Notes, and References*:

- general Data Protector and integration-specific limitations
- supported platforms and integrations
- supported backup and connectivity topologies

Configuring the integration

Before you start with the configuration, make sure the prerequisites listed in “Introduction” (page 95) are fulfilled.

To be able to use the Data Protector HP P4000 SAN Solutions integration with a storage system of the HP P4000 SAN Solutions family, you must perform the mandatory configuration step. In this step, you need to provide the Data Protector HP StorageWorks P4000 Agent the data which the ZDB agent will use to establish connection to a Common Information Model Object Manager (CIMOM) provider of your choice.

CIMOM provider connection configuration

In order to be able to connect to a CIMOM provider, the Data Protector HP StorageWorks P4000 Agent needs the following information:

- fully qualified domain name or IP address of the system where the CIMOM service is running

  In case the system has multiple IP addresses configured, the address by which the system can be accessed by the Data Protector ZDB agent should be used.
- whether the connection uses Secure Sockets Layer (SSL)
- port number of the port on which the CIMOM service is accepting requests
- user name and password

  This data must belong to a user account which has administrative privileges on the P4000 SAN Solutions storage system.

This information should be provided for each CIMOM provider that the Data Protector HP StorageWorks P4000 Agent should connect to. Once added, the connection configuration data for a particular CIMOM provider is stored in a separate configuration file located on the Cell Manager in the directory:

Windows Server 2008: `Data_Protector_program_data\db40\smisdb\p4000\login`

Other Windows systems: `Data_Protector_home\db40\smisdb\p4000\login`

UNIX systems: `/var/opt/omni/server/db40/smisdb/p4000/login`

To add the connection configuration data, use the Data Protector `omnidbp4000` command. With `omnidbp4000`, you can also update or remove the configuration data, list the contents of the configuration files, and check if the connection to a particular CIMOM provider can be established. For these purposes, the `omnidbp4000` command provides the basic options `--add`, `--remove`, `--list`, and `--check`. For command syntax and usage examples, see the `omnidbp4000` reference page in the *HP Data Protector Command Line Interface Reference* or the `omnidbp4000` man page.

10 Backup

Zero downtime backup sessions that involve a storage system of the HP P4000 SAN Solutions family can only be initiated through the Data Protector Microsoft Volume Shadow Copy Service integration.

For information about the supported configurations, ZDB types and replication techniques available on this storage system family, and storage system-specific ZDB considerations, see the *HP Data Protector Zero Downtime Backup Concepts Guide*.

For additional storage system-specific ZDB considerations, procedure for configuring ZDB backup specifications, and instructions for running ZDB sessions, see the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

11 Restore

Instant recovery sessions that involve a storage system of the HP P4000 SAN Solutions family can only be initiated through the Data Protector Microsoft Volume Shadow Copy Service integration.

For information on replica handling during instant recovery, description of the instant recovery process, and storage system-specific instant recovery considerations, see the *HP Data Protector Zero Downtime Backup Concepts Guide*.

For additional storage system-specific instant recovery considerations and instructions for running instant recovery sessions, see the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

12 Troubleshooting

Before you begin

This chapter lists general checks and verifications that you may need to perform when you encounter problems with the P4000 SAN Solutions integration. For general Data Protector troubleshooting information, see the *HP Data Protector Troubleshooting Guide*.

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

Checks and verifications

- On the application and backup systems, examine system errors logged into:

  Windows Server 2008: `Data_Protector_program_data\log\debug.log`

  Other Windows systems: `Data_Protector_home\log\debug.log`

Part IV EMC Symmetrix

This part describes how to configure the Data Protector EMC Symmetrix integration, how to perform zero downtime backup and instant recovery using the EMC Symmetrix integration, and how to resolve the integration-specific Data Protector problems.

13 Configuration

Introduction

This chapter describes the configuration of the Data Protector EMC Symmetrix (EMC) integration. It also provides information on the EMC Symmetrix database file and Data Protector EMC log file.

Prerequisites

- Install:

  EMC licenses and components:

  - EMC Solution Enabler
  - EMC Symmetrix TimeFinder or EMC Symmetrix Remote Data Facility (SRDF) microcode and license.

  Data Protector licenses and components:

  - Appropriate zero downtime backup extension and instant recovery extension licenses-to-use (LTU).
  - EMC Symmetrix Agent.

  For installation instructions, see the *HP 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.
- If the application system and the backup system reside in a Data Protector cell with secured clients, ensure that access between both systems is allowed in both directions.
- Connect EMC to the application and backup systems.

See the *HP Data Protector Product Announcements, Software Notes, and References* for information on:

- General Data Protector and integration-specific limitations
- Supported platforms and integrations
- Supported backup and connectivity topologies

For information on supported configurations, see the *HP 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:

Windows: `symapi_home\db\symapi_db.bin`

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

Data Protector EMC log file

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

Windows: `Data_Protector_home\Config\client\tmp\emc`

HP-UX: `/var/opt/omni/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 with the configuration, make sure the prerequisites listed in "Introduction" (page 101) are fulfilled. 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:

- Create the Data Protector EMC database file. See "Creating Data Protector EMC database file" (page 102).
- If needed, rebuild the EMC Symmetrix database file. See "Rebuilding EMC Symmetrix database file" (page 103).

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:

Windows: *`Data_Protector_home`*`\bin\syma -init`

HP-UX: `/opt/omni/lbin/syma -init`

This command creates the `Data Protector\Config\Client\EMC\symm.bin` (Windows) or `/var/opt/omni/client/emc/symm.bin` (HP-UX) 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” (page 152) for more information.

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 sessions.

For more information on the mountpoint creation, see the *HP Data Protector Zero Downtime Backup Concepts Guide*.

14 Backup

Introduction

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 Data Protector Zero Downtime Backup Concepts Guide*.

ZDB types

The only supported ZDB type is ZDB to tape.

With ZDB to tape, mirrors are created, and data from the replica is moved to backup media according to the tape backup type you have selected (Full, Incr, Incr1-9).

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 Data Protector Zero Downtime Backup Concepts Guide* for more information on ZDB-to-tape process.

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.

Backup in LVM mirroring configurations

Consider the following:

- Only the physical volumes that contain the logical volumes selected for backup will be considered for replication.

  Example

  ◦ A Volume Group (VG01) is made up of two physical volumes (PV1 and PV2)

  ◦ VG01 has two logical volumes (`lvol1` and `lvol2`)

  ◦ The `lvol1` has its logical extents on PV1, and `lvol2` - on PV2

  ◦ A backup object belonging to `lvol1` is selected in the backup specification

  PV1 will be selected for replication.

Creating backup specifications

ⓘ **IMPORTANT:** Before you begin, consider all limitations regarding the EMC integration. For more information, see the *HP Data Protector Product Announcements, Software Notes, and References* and the *HP 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**, and click **Add Backup**.

   The Create New Backup dialog box appears.

   In the Filesystem pane, select the **Blank Filesystem Backup** template or some other template which you might have created. For information on templates, see the online Help index: “backup templates”.

   Select **Split mirror backup** as **Backup type** and **EMC Symmetrix** as **Sub type**. See online Help for options’ descriptions. Click **OK**.
3. Under Client systems, select **Application system** and **Backup system**. Also, specify the desired EMC configuration - TimeFinder, SRDF, or Combined (SRDF + TimeFinder).

   See “Backup options” (page 107) 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 the application system and select the objects to be backed up. Note that all drive letters or mount points that reside on the system are displayed. You must select only objects that reside on the disk array, otherwise the backup session fails.

   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** or **Remove mirror**. Select separate devices for each mirror backup.

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

   Click **Next**.

6. Under Backup Specification Options, click **Advanced** and then the **EMC Symmetrix** tab to open the EMC backup options pane.

   Here, you can modify all options, except **Application system** and **Backup system**, as shown in "Backup options" (page 106). See also "Backup options" (page 107).

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

   ### Figure 19 Backup options

   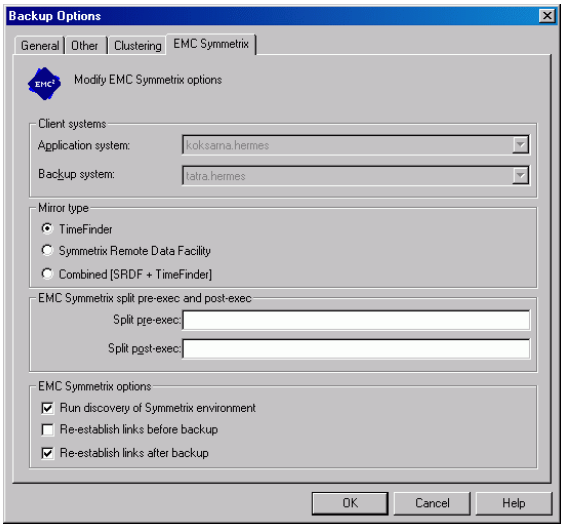

   Modify EMC Symmetrix options

   Client systems

   Application system: koksarna.hermes

   Backup system: tatra.hermes

   Mirror type

   ● TimeFinder

   ○ Symmetrix Remote Data Facility

   ○ Combined [SRDF + TimeFinder]

   EMC Symmetrix split pre-exec and post-exec

   Split pre-exec:

   Split post-exec:

   EMC Symmetrix options

   ☑ Run discovery of Symmetrix environment

   □ Re-establish links before backup

   ☑ Re-establish links after backup

7. Following the wizard, open the scheduler (for information, press **F1** or see "Scheduling ZDB sessions" (page 124)), 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/`*`vgnumber`*`/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 "Scheduling ZDB sessions" (page 124).

NOTE: Backup preview is not supported.

Backup options

The following tables describe EMC backup options. See also “EMC integration” (page 150).

Table 24 EMC backup options

Data Protector GUI	Function
Application system	The system on which the application runs. In cluster environments, specify the virtual server hostname (rather than the physical node hostname).
Backup system	The 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 omnirc variables” (page 139) for more information. Backup session is not aborted if the command set by `Split pre-exec` is not executed.
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” (page 102) 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 20 Filesystem split mirror backup flow

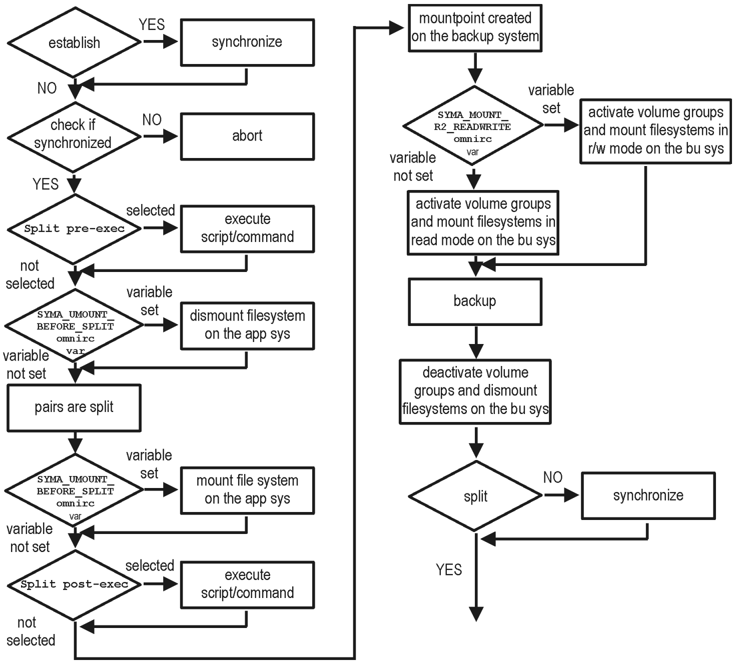
establish
YES
synchronize
NO
check if synchronized
NO
abort
YES
Split pre-exec
selected
execute script/command
not selected
SYMA_UMOUNT_ BEFORE_SPLIT omnirc var
variable set
dismount filesystem on the app sys
variable not set
pairs are split
SYMA_UMOUNT_ BEFORE_SPLIT omnirc var
variable set
mount file system on the app sys
variable not set
Split post-exec
selected
execute script/command
not selected
mountpoint created on the backup system
SYMA_MOUNT_ R2_READWRITE omnirc var
variable set
activate volume groups and mount filesystems in r/w mode on the bu sys
variable not set
activate volume groups and mount filesystems in read mode on the bu sys
backup
deactivate volume groups and dismount filesystems on the bu sys
split
NO
synchronize
YES

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

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:

Windows: `*Data_Protector_home*\Config\client\tmp\emc\R2_*session_name*.log`

HP-UX: `/var/opt/omni/tmp/emc/R2_*session_name*.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" (page 111) and set EMC split mirror restore options as explained in "EMC test restore options" (page 109).

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 omnirc variables" (page 139).

Table 25 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	The system on which the application runs. In cluster environments, specify the virtual server hostname (rather than the physical node hostname).
Backup system	The 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.
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" (page 111) 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

Introduction

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” (page 110).
- Split mirror restore. See “Split mirror restore” (page 110).

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 online Help index: “restore”.

TIP: You can improve the data transfer rate by connecting a backup device to the application system. For information on configuring backup devices, see the online Help index: “backups devices: configuring”. For information on performing a restore using another device, see the online Help index: “selecting, devices for restore”.

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 “Restore from backup media on LAN” (page 110).

   Figure 21 Restore from backup media on LAN

   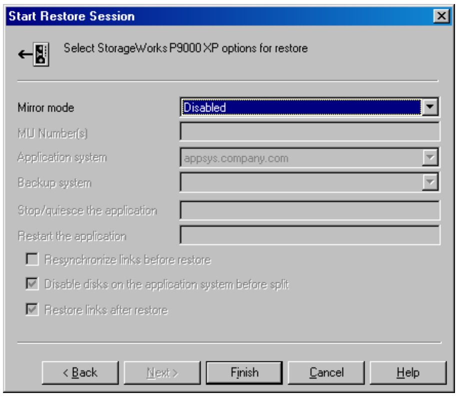

   Start Restore Session

   Select StorageWorks P9000 XP options for restore

   Mirror mode: Disabled

   MU Number(s):

   Application system: appsys.company.com

   Backup system:

   Stop/quiesce the application:

   Restart the application:

   ☐ Resynchronize links before restore

   ☑ Disable disks on the application system before split

   ☑ Restore links after restore

6. Click **Finish** to start the 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 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**.
5. Specify the report level and network load. Click **Next**.
6. Select **EMC Symmetrix restore**. Click **Next**.
7. Specify the split mirror restore options. See "EMC Symmetrix split mirror restore options" (page 111). For more information, see "Split mirror restore options" (page 111).

   Figure 22 EMC Symmetrix split mirror restore options

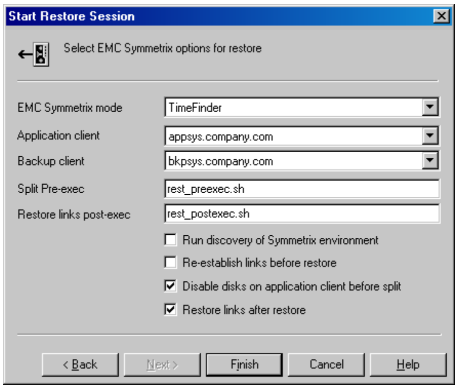

Start Restore Session

Select EMC Symmetrix options for restore

EMC Symmetrix mode: TimeFinder

Application client: appsys.company.com

Backup client: bkpsys.company.com

Split Pre-exec: rest_preexec.sh

Restore links post-exec: rest_postexec.sh

☐ Run discovery of Symmetrix environment

☐ Re-establish links before restore

☑ Disable disks on application client before split

☑ Restore links after restore

8. Click **Finish** to start the split mirror restore.

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 26 EMC split mirror restore options

Data Protector GUI	Function
EMC Symmetrix mode	EMC Symmetrix configuration: TimeFinder, SRDF, or Combined (SRDF + TimeFinder).
Application system	The system on which the application runs. In cluster environments, specify the virtual server hostname (rather than the physical node hostname).
Backup system	The system to which your data is first restored. In cluster environments, specify the virtual server hostname (rather than the physical node hostname).

Table 26 EMC split mirror restore options *(continued)*

Data Protector GUI	Function
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 omnirc variables" (page 139) 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 restor**e. 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" (page 102) for more information. Default: not selected.
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 23 Split mirror restore flow

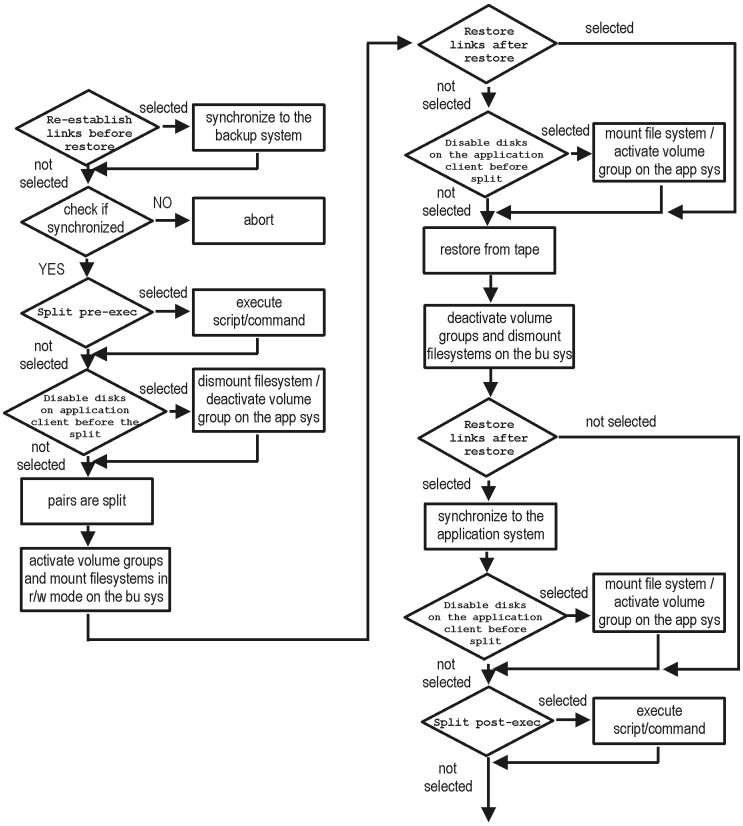
Re-establish links before restore
selected
synchronize to the backup system
not selected
check if synchronized
NO
abort
YES
Split pre-exec
selected
execute script/command
not selected
Disable disks on application client before the split
selected
dismount filesystem / deactivate volume group on the app sys
not selected
pairs are split
activate volume groups and mount filesystems in r/w mode on the bu sys
Restore links after restore
selected
not selected
Disable disks on the application client before split
selected
mount file system / activate volume group on the app sys
not selected
restore from tape
deactivate volume groups and dismount filesystems on the bu sys
Restore links after restore
not selected
selected
synchronize to the application system
Disable disks on the application client before split
selected
mount file system / activate volume group on the app sys
not selected
Split post-exec
selected
execute script/command
not selected

Split mirror restore in a cluster

Split mirror restore in configurations with the application system in MC/ServiceGuard or a Microsoft server cluster requires additional steps. For details, see the sections that follow.

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" (page 111)).

   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 lists general checks and verifications, and problems you may encounter when using the EMC integration. For general Data Protector troubleshooting information, see the *HP Data Protector Troubleshooting Guide*.

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

Checks and verifications

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

  Windows: `Data_Protector_home\log\debug.log`

  HP-UX: `/var/opt/omni/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 on the Cell Manager in the following directory:

Windows Server 2008: `Data_Protector_program_data\Config\server\cell\cell_info`

Other Windows system: `Data_Protector_home\Config\server\cell\cell_info`

UNIX system: `/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.06.10 -da A.06.10 -emc A.06.10
-host "hpsap002.bbn.hp.com" -os "hp s800 hp-ux-11.00"
-cc A.06.10 -da A.06.010 -ma A.06.10 -emc A.06.10
```

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 omnirc variables" (page 139) 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" (page 122) 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" (page 122) 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.

Problem

Preparation of the backup system fails when VxVM is used

This problem may be caused by the following:

- If a backup specification involves VxVM volume groups, EMC arrays do not support I/O on a BCV device in a synchronized state.
- The information about volume groups is not added to the VxVM configuration.

Action

1. Check if any backup objects in the backup specification belong to VxVM disk groups.
2. If there are objects belonging to VxVM volume groups, proceed as follows:
   a. Check if a BCV is visible on the backup system.
   b. Check the synchronization state of the BCV devices. If the BCV devices are synchronized, split them.
   c. Run `vxdisk scandisks`.
   d. Re-establish the mirror.

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.

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 omnirc variables" (page 139) 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" (page 122) 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" (page 122) 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 00028260031 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.

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:

Windows:

`Data_Protector_home``\Config\Emc\symmR1.rec`

`Data_Protector_home``\Config\Emc\symmR2.rec`

HP-UX:

```
/var/opt/omni/emc/symmR1.rec
/var/opt/omni/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:

  Windows: *Data_Protector_home*\bin\syma -r1 -session *sessionID* -recovery

  HP-UX: /opt/omni/lbin/syma -r1 -session *sessionID* -recovery

- On the backup system:

  Windows: *Data_Protector_home*\bin\syma -no_r1 -session *sessionID* -recovery [-split]

  HP-UX: /opt/omni/lbin/syma -no_r1 -session *sessionID* -recovery [-split]

  You can obtain *sessionID* from the Data Protector GUI as shown in “Obtaining session ID” (page 123).

Figure 24 Obtaining session ID

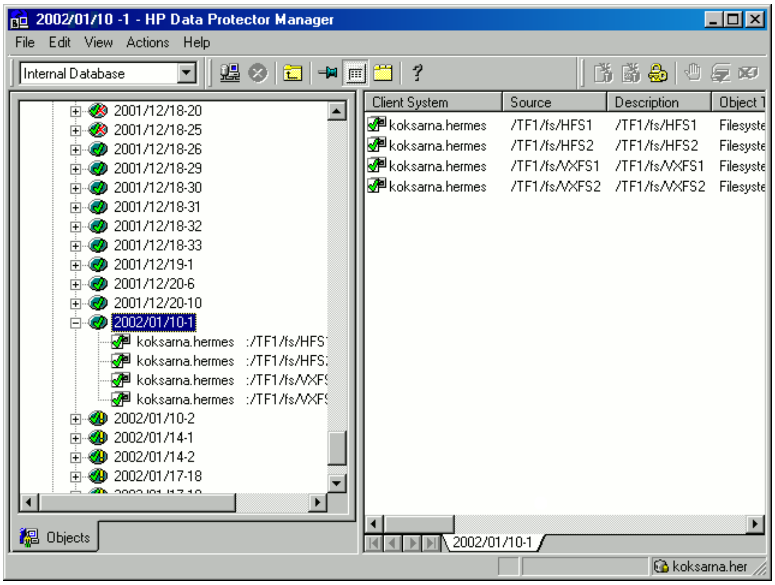

Client System	Source	Description	Object T
koksarna.hermes	/TF1/fs/HFS1	/TF1/fs/HFS1	Filesyste
koksarna.hermes	/TF1/fs/HFS2	/TF1/fs/HFS2	Filesyste
koksarna.hermes	/TF1/fs/VXFS1	/TF1/fs/VXFS1	Filesyste
koksarna.hermes	/TF1/fs/VXFS2	/TF1/fs/VXFS2	Filesyste

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

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 the online Help index: “scheduled backups”.

Figure 25 Scheduling ZDB to disk/disk+tape

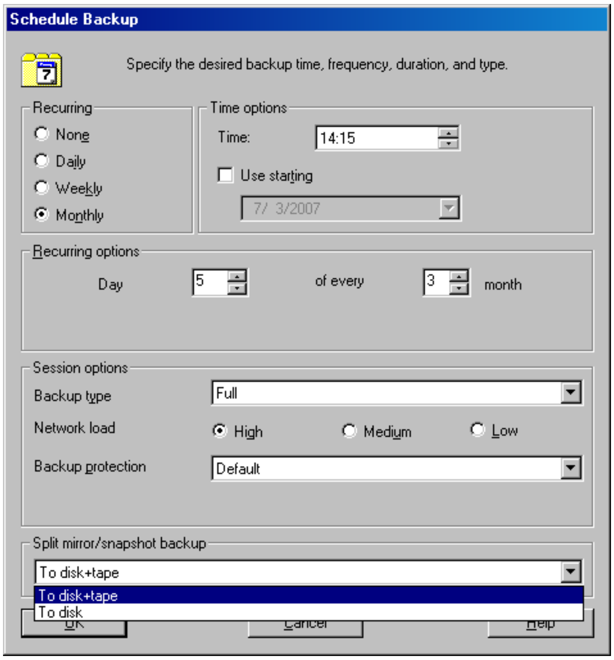

Schedule Backup

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

Recurring
○ None
○ Daily
○ Weekly
● Monthly

Time options
Time: 14:15
□ Use starting
7/ 3/2007

Recurring options
Day 5 of every 3 month

Session options
Backup type: Full
Network load: ● High ○ Medium ○ Low
Backup protection: Default

Split mirror/snapshot backup
To disk+tape
To disk+tape
To disk

Starting interactive ZDB sessions

Prerequisites

- In a Microsoft Cluster Service configuration, if a cluster resource disk is to be backed up, it should not be in a maintenance mode before the backup.

NOTE: When running concurrent ZDB sessions using one or several application systems, consider the limitations described in the *HP 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.

   For ZDB to tape and ZDB to disk+tape, specify **Backup Type**.

   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.

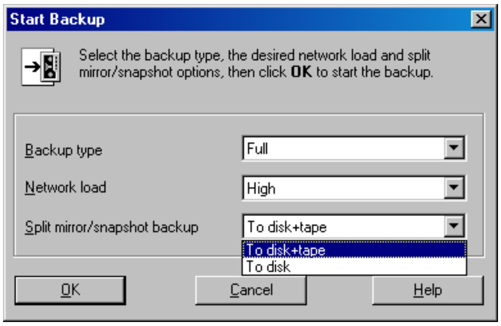

Start Backup

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

Backup type: Full

Network load: High

Split mirror/snapshot backup: To disk+tape

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

Alternate paths support

For systems with multiple host adapters and connections to a disk array, the multi-path device management solution performs dynamic load balancing and monitors each path to ensure that the I/O subsystem 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. Failover is transparent to applications, so they continue unaffected.

NOTE: On UNIX systems, the multi-path device management software used for the Data Protector HP P6000 EVA Disk Array Family integration to import volume groups on the backup system should be limited to the maximum supported number of paths.

NOTE: On HP-UX 11.31 systems, the multi-path device management software is not supported since the operating system has native device multi-pathing capability.

For information on which multi-path device management solutions are supported by specific Data Protector ZDB agents and disk array models, see the latest support matrices at http://www.hp.com/support/manuals

With the HP P9000 XP Disk Array Family, you can control AutoPath load balancing using the `OB2AUTOPATH_BALANCING_POLICY` variable (by default, AutoPath Round Robin load balancing policy is used). For more information, see “ZDB omnirc variables” (page 139).

When using AutoPath, consider the following:

- During a ZDB-to-tape session, if a failover to an alternate path occurs and the 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: When using a disk array of the HP P6000 EVA Disk Array Family together with the multi-path device management software HP Secure Path, load balancing as configured by HP Secure Path is used; you cannot change the load balancing policy using Data Protector.

Cluster configurations

Data Protector ZDB agents support:

- MC/ServiceGuard (on HP-UX systems) with all disk array models supported by Data Protector
- Veritas Cluster (on Solaris systems) with disk arrays of the HP P6000 EVA Disk Array Family and the HP P9000 XP Disk Array Family
- Microsoft Cluster Server (on Windows systems) with disk arrays of the HP P6000 EVA Disk Array Family and the HP P9000 XP Disk Array Family
- EMC GeoSpan for Microsoft Cluster Service (on Windows systems) with EMC Symmetrix disk arrays

If the application system is in a server cluster, the backup system must be outside this cluster: it may run in a different cluster or may not be part of a cluster at all. During backup sessions, the filesystem or the database structure (filesystems and the volume/disk group) is active on the backup system and would prevent activation during failover if both systems were part of the same cluster.

IMPORTANT: If the backup system is running in a server cluster, target volumes on this system must not be configured as cluster resources.

IMPORTANT: If a failover to the remote site happens, the disk array configuration changes from the combined HP CA+BC P9000 XP (HP P9000 XP Disk Array Family) or SRDF+TimeFinder (EMC Symmetrix) to HP BC P9000 XP (HP P9000 XP Disk Array Family) or TimeFinder (EMC Symmetrix). This means that the next ZDB session can no longer start automatically, so the ZDB backup specification must be updated to reflect the configuration change.

For more information on cluster support, see the *HP Data Protector Product Announcements, Software Notes, and References* and the online Help index: “cluster”.

Sections below discuss supported ZDB cluster configurations.

“Client on the application system in a cluster” (page 129) through “EMC GeoSpan for Microsoft Cluster Service” (page 131) illustrate Data Protector *application* backup disk array configurations and scenarios. For *filesystem and disk image* backup, only a Data Protector ZDB agent is needed; an application database and binaries are not installed as presented in the figures. On Windows systems, to perform zero downtime backup and instant recovery using Microsoft Volume Shadow Copy Service, the Data Protector component `MS Volume Shadow Copy Integration` must be installed.

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 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 all objects** is selected.
- Cell Manager failover before backup: session completes successfully.

Limitations

- Not supported in Veritas Cluster.

Install:

- On the application system on all cluster nodes on local disks: application binaries, Data Protector integration software component, Data Protector ZDB agent.
- On the application system cluster shared disk: application database. Note that this shared disk must be on a disk array.
- On any system cluster shared disk: Cell Manager.
- On the backup system on local disks: Data Protector integration software component, Data Protector ZDB agent, Data Protector General Media Agent.

Figure 26 Client on the application system in a cluster, Cell Manager in a cluster

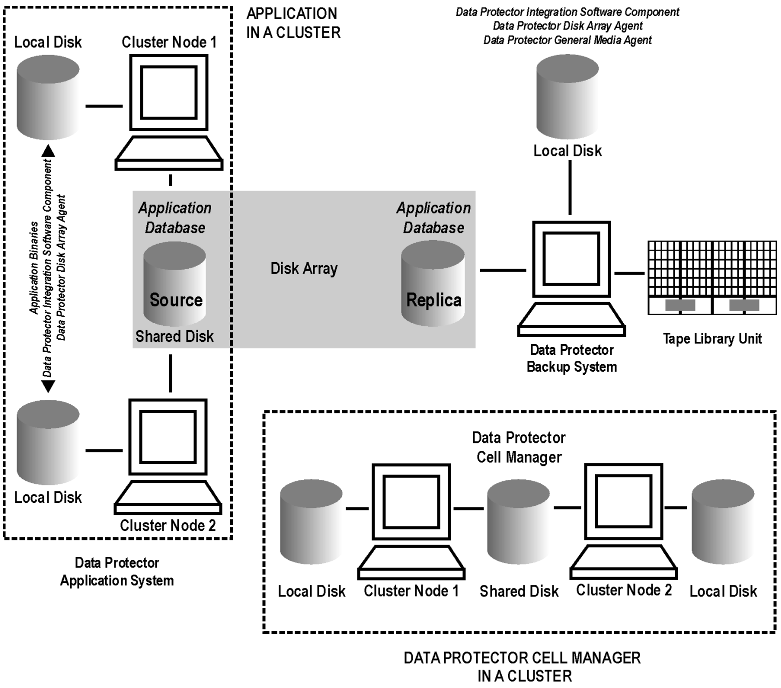
APPLICATION IN A CLUSTER
Local Disk
Cluster Node 1
Application Binaries
Data Protector Integration Software Component
Data Protector Disk Array Agent
Application Database
Source
Shared Disk
Local Disk
Cluster Node 2
Data Protector Application System
Disk Array
Application Database
Replica
Data Protector Integration Software Component
Data Protector Disk Array Agent
Data Protector General Media Agent
Local Disk
Data Protector Backup System
Tape Library Unit
Data Protector Cell Manager
Local Disk
Cluster Node 1
Shared Disk
Cluster Node 2
Local Disk
DATA PROTECTOR CELL MANAGER 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 all objects** is selected.
- Cell Manager failover in between backups: session completes successfully.

Limitations

- Not supported in Veritas Cluster.

Install:

- On the application system: application binaries, Data Protector integration software component, Data Protector ZDB 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 ZDB agent, Data Protector General Media Agent.

Figure 27 Cell Manager on the backup system in a cluster

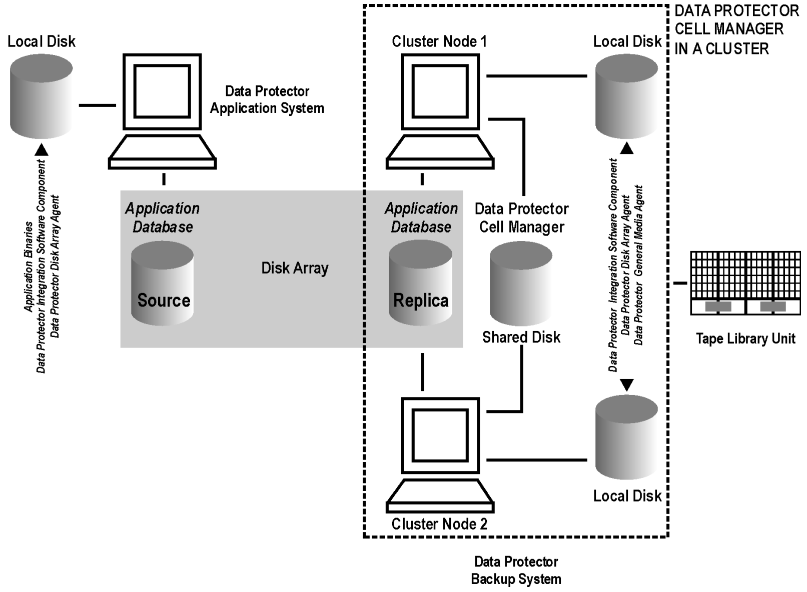

Cell Manager and client on the application system in a cluster

Scenarios

- Application or Data Protector failover during backup: session is restarted automatically.
- Application or Data Protector failover in between backups: session completes successfully.

Limitations

- Not supported in Veritas Cluster.
- Split mirror restore is not possible (HP P9000 XP Disk Array Family, EMC Symmetrix).

Install:

- On the application system on all cluster nodes on local disks: application binaries, Data Protector integration software component, Data Protector ZDB 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 ZDB 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 28 Cell Manager and client on the application system in a cluster

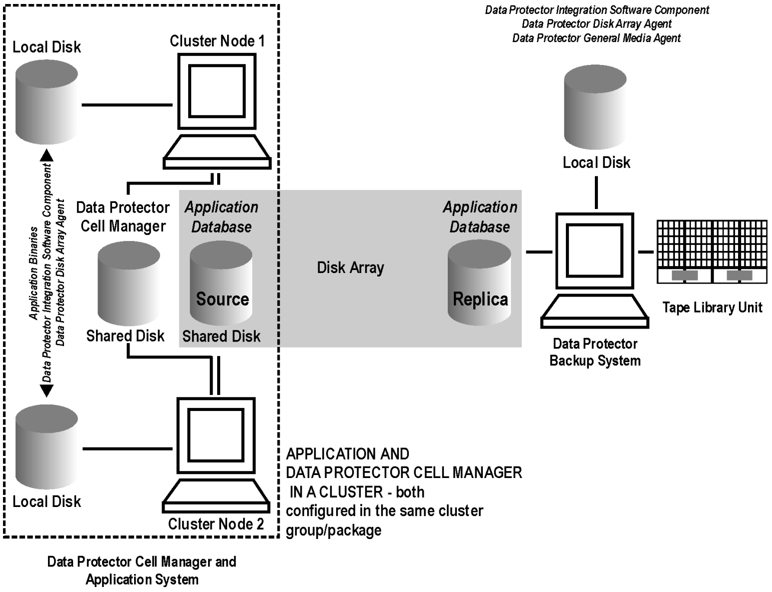

Client on the application system in a cluster, Cell Manager not in a cluster

Scenarios

- Application failover during backup: session fails and must be restarted manually.
- Application failover in between backups: session completes successfully.

Install:

- On the application system on all cluster nodes on local disks: application binaries, Data Protector integration software component, Data Protector ZDB agent.
- On the application system cluster shared disk: application database. Note that this shared disk must be on a disk array.
- On the backup system on local disks: Data Protector integration software component, Data Protector ZDB agent, Data Protector General Media Agent.

Figure 29 Client on the application system in a cluster

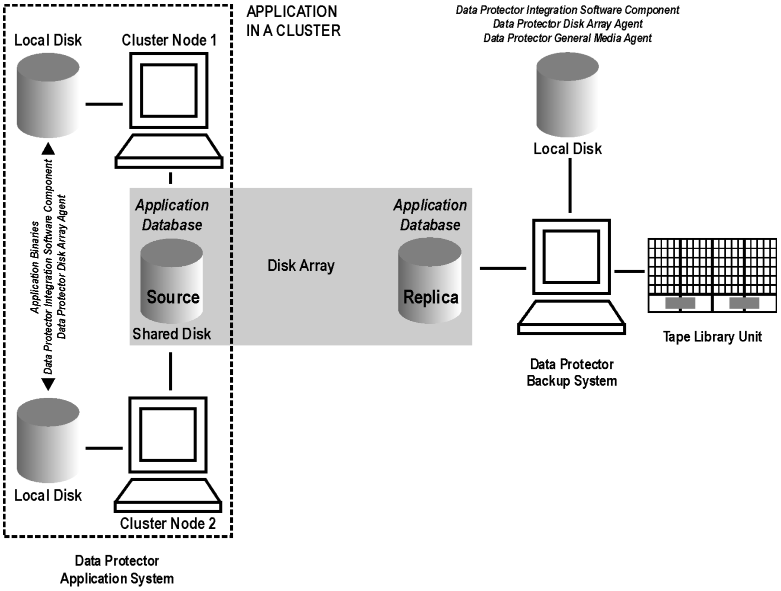

Client on the application system in a cluster, Cell Manager on the backup system in a cluster

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 all objects** is selected.
- Cell Manager failover before backup: session completes successfully.

Limitations

- Not supported in Veritas Cluster.

Install:

- On the application system on all cluster nodes on local disks: application binaries, Data Protector integration software component, Data Protector ZDB 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 ZDB agent, Data Protector General Media Agent.

Figure 30 Client on the application system in a cluster, Cell Manager on the backup system in a cluster

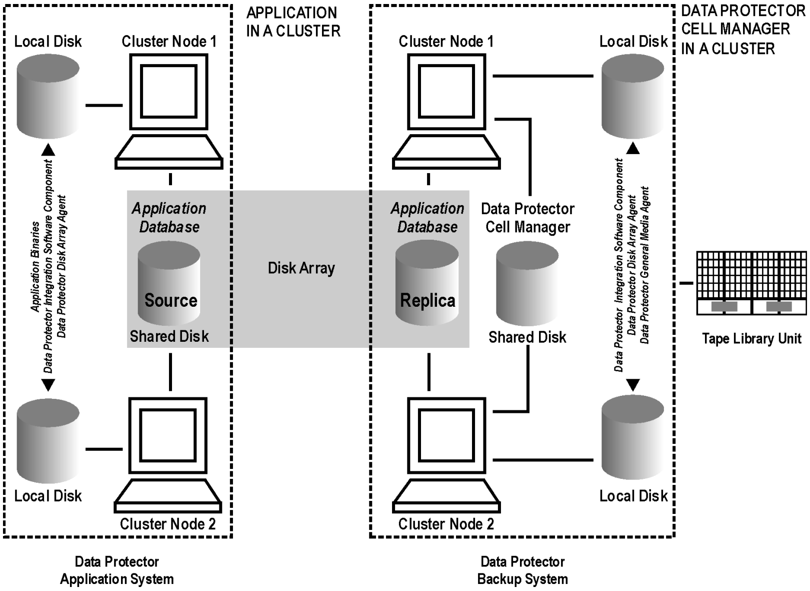

APPLICATION IN A CLUSTER
Local Disk
Cluster Node 1
Application Binaries
Data Protector Integration Software Component
Data Protector Disk Array Agent
Application Database
Source
Shared Disk
Local Disk
Cluster Node 2
Disk Array
DATA PROTECTOR CELL MANAGER IN A CLUSTER
Cluster Node 1
Local Disk
Application Database
Replica
Data Protector Cell Manager
Shared Disk
Data Protector Integration Software Component
Data Protector Disk Array Agent
Data Protector General Media Agent
Local Disk
Cluster Node 2
Tape Library Unit
Data Protector Application System
Data Protector Backup System

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, Data Protector EMC Agent, Data Protector General Media Agent.

Figure 31 EMC GeoSpan for Microsoft Cluster Service

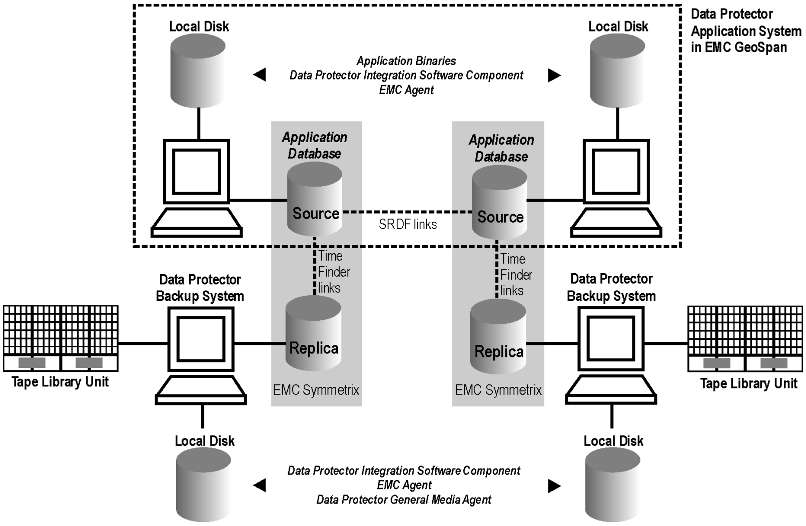

Instant recovery in a cluster

With an application or filesystem running on MC/ServiceGuard or Microsoft Cluster Server on the application system, instant recovery requires some *additional* steps. Additionally, there are limitations regarding instant recovery on Microsoft Cluster Server.

IMPORTANT: If HP-UX LVM mirroring is used, see also "Instant recovery and LVM mirroring" (page 88).

MC/ServiceGuard

Cluster File System

During an instant recovery procedure, Data Protector cannot automatically unmount and mount Cluster File System (CFS) volumes that will be recovered. You must manually unmount the filesystem before starting the instant recovery and manually mount it after instant recovery finishes. To enable manual mounting on the application system, set the variable `ZDB_IR_MANUAL_AS_PREPARATION` to `1`. The instant recovery session will finish with warnings. See "Common ZDB variables" (page 140).

Procedure

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).

   Oracle 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" (page 49) (P6000 EVA Array)
   - "Instant recovery procedure" (page 86) (P9000 XP Array)

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. After instant recovery, recover the database. For detailed procedures, see the database documentation.

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`. For more information, see "ZDB omnirc variables" (page 139).

Microsoft Cluster Server

Limitations

- Instant recovery of a cluster quorum disk is not supported because the cluster service must never lose the connection with the quorum disk, which happens during instant recovery (when disks are unpresented).
- In the configuration where a local disk is mounted to a cluster resource disk, instant recovery of a such disk is not supported.
- Any target cluster disk resource must be owned by the currently active node. Instant recovery is not supported if the disk resource is owned by the non-active node.
- Instant recovery of combination of cluster and non-cluster disks is not supported.

Considerations

- In a Microsoft Cluster Server environment, disks are distinguished by their disk signature. Because two disks cannot have the same signature, the operating system dynamically changes the signature once it detects the replica on the backup system. During the instant recovery procedure, Data Protector restores the disk signature to ensure that the recovered disk will have the same signature as the original disk on the application system. Data Protector will display notifications, informing you about the changed signature.

Prerequisites

- On Windows Server 2008 systems, before running an instant recovery session, you need to bring the original disks online.

Procedure

1. Using the Cluster Administrator utility or Cluster CLI, take the application cluster resource offline. For detailed instructions, see the Microsoft Cluster Server documentation.
2. Shut down the application (database).
3. Start instant recovery. For instructions, see:
    - “Instant recovery procedure” (page 49) (P6000 EVA Array)
    - “Instant recovery procedure” (page 86) (P9000 XP Array)
4. Restart the application (database).
5. Recover the database. For detailed procedures, see the database documentation.
6. Using the Cluster Administrator utility or CLI, put the application cluster resource online.

Instant recovery for P6000 EVA Array in HP CA+BC P6000 EVA configurations

Introduction

This section describes the steps to be followed for executing instant recovery in HP Continuous Access + Business Copy (CA+BC) P6000 EVA environments of the HP P6000 EVA Disk Array Family using Data Protector.

The section gives details of the following:

- The different situations where HP CA+BC P6000 EVA impacts instant recovery
- Instant recovery concepts
- HP CA+BC P6000 EVA configurations supported for instant recovery
- How to plan and perform instant recovery in HP CA+BC P6000 EVA configurations

Prerequisites

You should be familiar with the following:

- *HP Data Protector Zero Downtime Backup Concepts Guide*
- HP storage management appliance (SMA) documentation
- HP P6000 EVA Disk Array Family documentation
- Failover or cluster-failover documentation, such as the *HP Cluster Extension EVA user guide*

Overview

With instant recovery, lost or corrupted data (or rather, the whole volumes containing it) is replaced with known good data. This good data resides on whole storage volumes, or virtual disks, which have been created previously as an HP BC P6000 EVA during a ZDB. These replicated target volumes are used for restores internally within the array, involving no other backup medium or device.

The general SMISA instant recovery flow is as follows:

1. The application system is prepared for restore by dismounting filesystems and taking volume groups offline.
2. Source volumes are masked or unpresented from the application system.
3. The identities of each matched pair of source and target storage volumes are exchanged. This involves the WWN, the name, and the comments of each volume.
4. The exchanged storage volume is unmasked or presented to the application system.
5. Volume groups are put online and filesystems remounted.

Each source storage volume that is backed up using ZDB has a matching target storage volume in the replica.

NOTE: To enable instant recovery, each pair of matched replica and source storage volumes must reside on the same disk array. This is required for a valid exchange of identities (step 3 in the general instant recovery flow above).

However, when a source volume is attached to a DR group and so participates in remote replication, the P6000 EVA Array does not allow the WWN of that virtual disk to be modified. Therefore, to prepare your environment for instant recovery and successfully recover your data, you need to carry out the following steps:

1. Manually prepare the storage volumes and the storage environment for instant recovery, as described in "Instant recovery in CA+BC environments" (page 136).
2. Perform instant recovery with the Data Protector HP StorageWorks P6000 EVA SMI-S Agent.
3. Optionally, return the storage volumes and storage environment to the state they were in before instant recovery.

The following sections outline different HP CA+BC P6000 EVA configurations and the manual steps you need to follow for successful instant recovery.

Supported instant recovery configurations

The manual steps needed to prepare the environment for instant recovery and bring it back after instant recovery differ depending on the current configuration of HP CA+BC P6000 EVA or DR group connections.

Identifying the setup depends on the following environment information:

- The current site for the source side of any DR groups that include the source storage volumes
- Whether the HP BC P6000 EVA or target storage volumes are on the same array as the source storage volumes (*local*), or on the remote side of the DR group (*remote*)

From this information, there are two possible configurations:

- Configuration I – HP Business Copy P6000 EVA is on the local side of the HP CA P6000 EVA link
- Configuration II – HP Business Copy P6000 EVA is on the remote side of the HP CA P6000 EVA link

Configuration I – local HP Business Copy P6000 EVA

Figure 32 Replicas on the local site

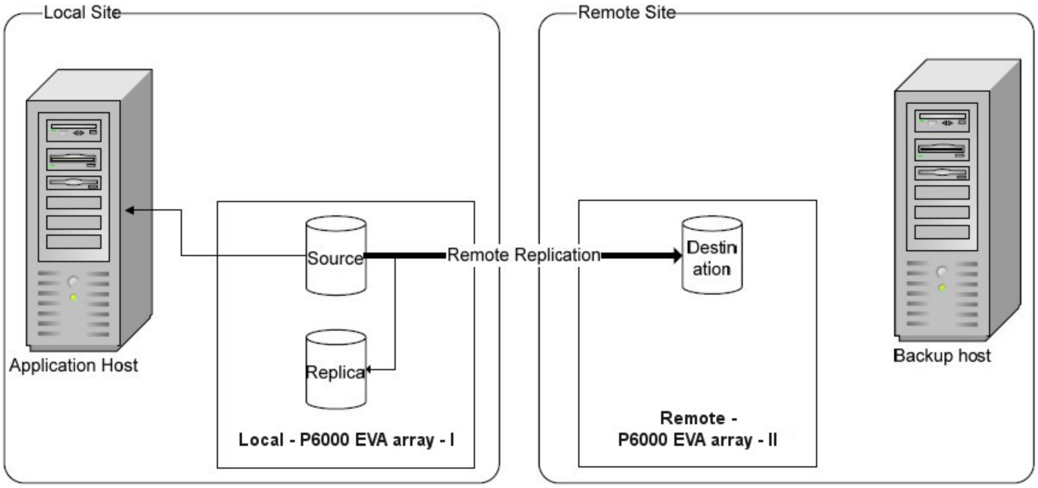

In this configuration, at the time of instant recovery, the source and replica storage volumes reside on the current local site.

NOTE: The source storage volume ("Source" in the diagram) acts as both the source of the replica storage volume and the source for the remotely replicated storage volume ("Destination" in the diagram).

The configuration may be a result of any of the following:

- Performing an HP BC P6000 EVA backup of a volume that is remotely replicated at backup time.
- Adding remote replication to a storage volume that was previously backed up by an HP BC P6000 EVA backup.
- Performing an HP CA+BC P6000 EVA backup with the HP BC P6000 EVA on the current local site.

Configuration II – remote HP Business Copy P6000 EVA

Figure 33 Replicas on the remote site

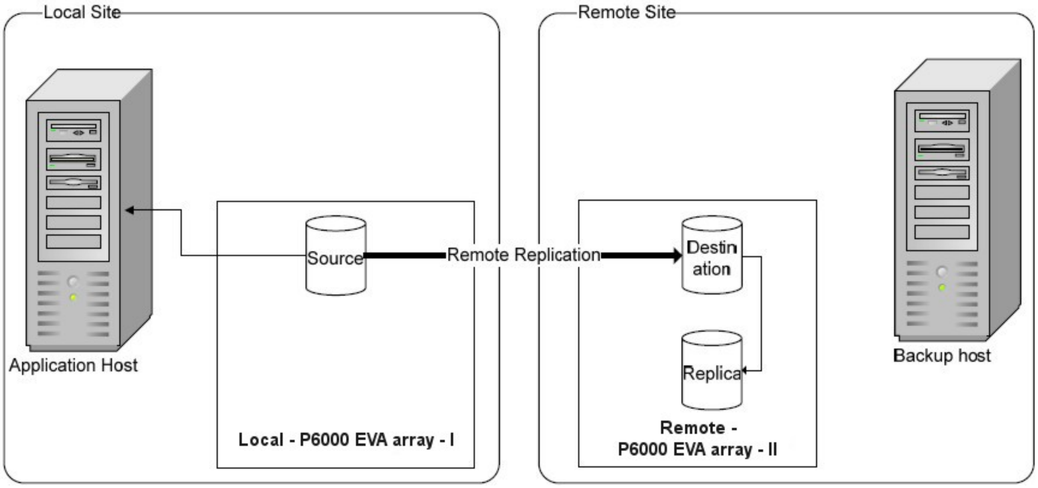

In this configuration, at the time of instant recovery, the customer environment has the source virtual disk residing on the local site. The remote replica (the replica of the source virtual disk replicated using HP CA P6000 EVA) and its local replica are both on the remote site.

NOTE: The storage volume marked "Destination" in the diagram is both the *destination* of the remote replication link and the *source* of the replica storage volume.

Such a configuration may be a result of any of the following:

- Performing an HP BC P6000 EVA backup of a storage volume that is remotely replicated, and then failing over the environment.
- Adding remote replication to a volume that was previously backed up by an HP BC P6000 EVA backup, and then failing over the environment.
- Performing an HP CA+BC P6000 EVA backup with the HP BC P6000 EVA on the current remote site.

Instant recovery in HP CA+BC P6000 EVA environments

The initial steps of preparation for instant recovery are as follows:

1. Understand the current configuration of the environment.
2. Optionally, perform a DR group failover.
3. Modify the DR group so that the source storage volumes involved in restore no longer participate in the DR group.

The following flow chart summarizes this general process.

Figure 34 General instant recovery flow in HP CA+BC P6000 EVA environments

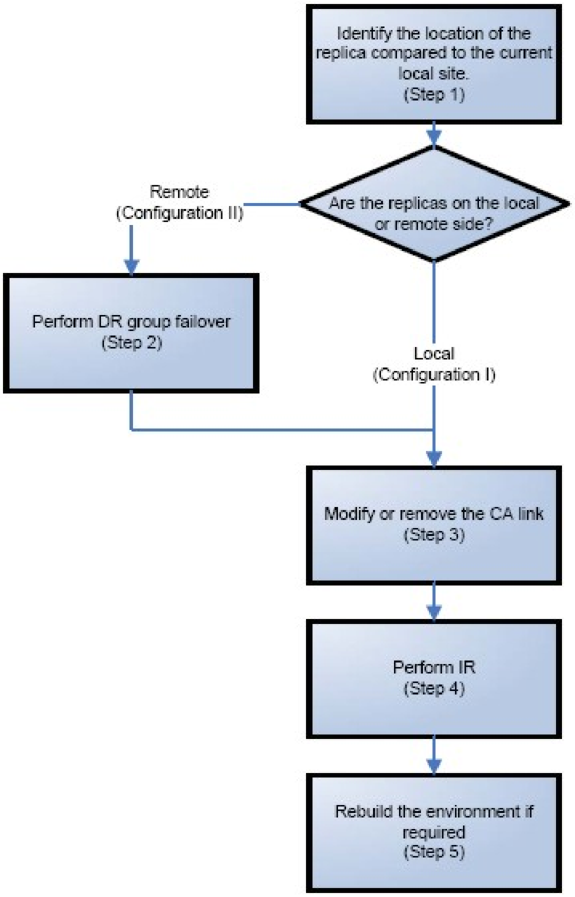
Identify the location of the replica compared to the current local site. (Step 1)
Are the replicas on the local or remote side?
Remote (Configuration II)
Perform DR group failover (Step 2)
Local (Configuration I)
Modify or remove the CA link (Step 3)
Perform IR (Step 4)
Rebuild the environment if required (Step 5)

Step 1: Identifying the current configuration

The following steps help identify the location of the source and target volumes:

1. Select the session for which instant recovery will be performed.

   List the sessions available for instant recovery using the Data Protector GUI (the **Instant Recovery** context) or the Data Protector CLI (the `omnidbsmis` command)

```
# omnidbsmis -list -session -ir
Found 2 StorageWorks P6000 EVA SMI-S session(s) in the internal \
database:

Session ID    IR   Type       Excluded  Backup Specification
============================================================
2010/06/20-5  Yes  Snapclone  No        VolDpdevpa2
2010/06/20-6  Yes  Snapclone  No        VolDpdevpa2
#
```

2. Identify the source objects and the HP CA P6000 EVA link information.

   Query the objects of the specific session using the `omnidbsmis` command. The following example is for a session with ID 2010/06/20-5.

   ```
   # omnidbsmis -show -session 2010/06/20-5

   Info on session "2010/06/20-5":

   Target volume virtual disk name : \Virtual Disks\SNEHA\DP-200
   8.06.20-5-04497CA1A\ACTIVE
   Target volume virtual disk ID   : 6005-08b4-0010-3a70-0000-90
   00-0661-0000
   Target volume virtual disk WWN  : 6005-08b4-0010-3a70-0000-90
   00-0661-0000
   StorageWorks P6000 EVA Array Family name : DPCA
   StorageWorks P6000 EVA Array Family ID   : 5000-1fe1-5005-dc0
   0
   Target volume snapshot type     : Snapclone
   Source volume virtual disk ID   : 6005-08b4-0010-3a70-0000-90
   00-0042-0000
   Session ID                      : 2010/06/20-5
   Creation Date                   : Sun Jun 20 15:42:42 2010

   IR flag                         : 1
   Excluded                        : 0
   Source disk version             : 0
   Backup specification            : VolDpdevpa2
   Application System              : dpdevpa2.hp.com
   Backup System                   : dpdevpa2.hp.com
   #
   ```

   From this output, you can find the following information:

   - The target/replica virtual disk WWN, UUID, and the name:
     - *WWN and UUID:* 6005-08b4-0010-3a70-0000-9000-0661-0000,
     - *Name:* \Virtual Disks\SNEHA\DP-2010.06.20-5-04497CA1A\ACTIVE
   - The source (of the replica) virtual disk UUID:
     - *UUID:* 6005-08b4-0010-3a70-0000-9000-0042-0000
   - The P6000 EVA Array name and the WWN where the matched source and target volumes exist:
     - *Name:* DPCA
     - *WWN:* 5000-1fe1-5005-dc00

3. Use this information to locate the source storage volume and the P6000 EVA Array where it resides. You can also locate the target storage volume or the target virtual disk to verify that it still exists:

   a. Connect to the Storage Management Appliance or any other CV EVA management host that manages the specific P6000 EVA Array storage system.

   b. Browse through the **Virtual Disk** folder until the virtual disk with a matching UUID is found.

      In the following figure, the source virtual disk with a UUID of 6005-08b4-0010-3a70-0000-9000-0042-0000 has been located:

Figure 35 Locating the source virtual disk

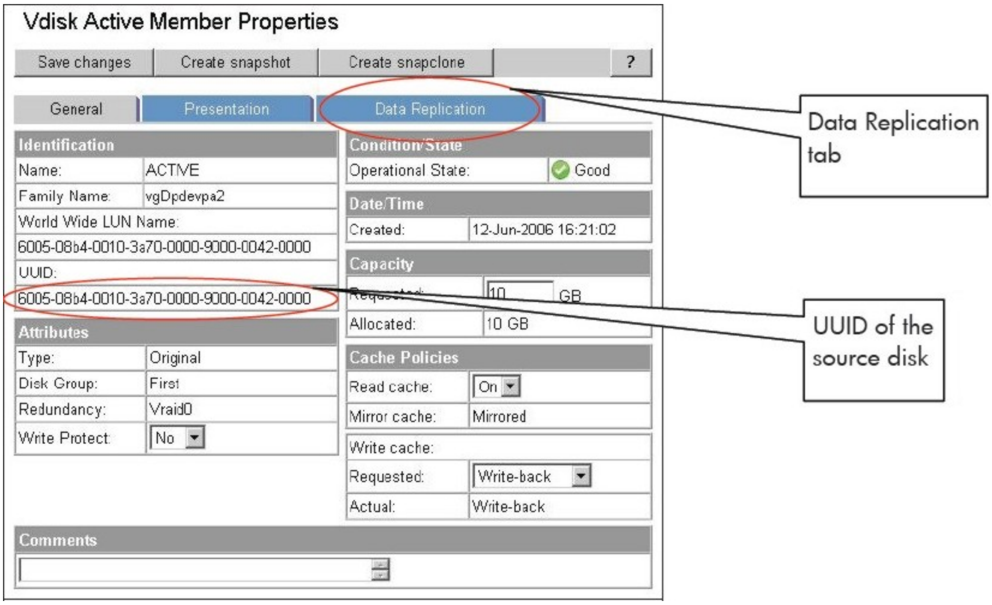

Vdisk Active Member Properties

Identification

Name:	ACTIVE
Family Name:	vgDpdevpa2
World Wide LUN Name:	
6005-08b4-0010-3a70-0000-9000-0042-0000	
UUID:	
6005-08b4-0010-3a70-0000-9000-0042-0000	

Attributes

Type:	Original
Disk Group:	First
Redundancy:	Vraid0
Write Protect:	No

Condition/State

Operational State:	Good

Date/Time

Created:	12-Jun-2006 16:21:02

Capacity

Requested:	10 GB
Allocated:	10 GB

Cache Policies

Read cache:	On
Mirror cache:	Mirrored
Write cache:	
Requested:	Write-back
Actual:	Write-back

Comments

c. Select the **Data Replication** tab to identify HP CA P6000 EVA link properties for this virtual disk. The following information should be gathered from this panel:

- DR group name
- DR mode

Figure 36 Checking the DR mode

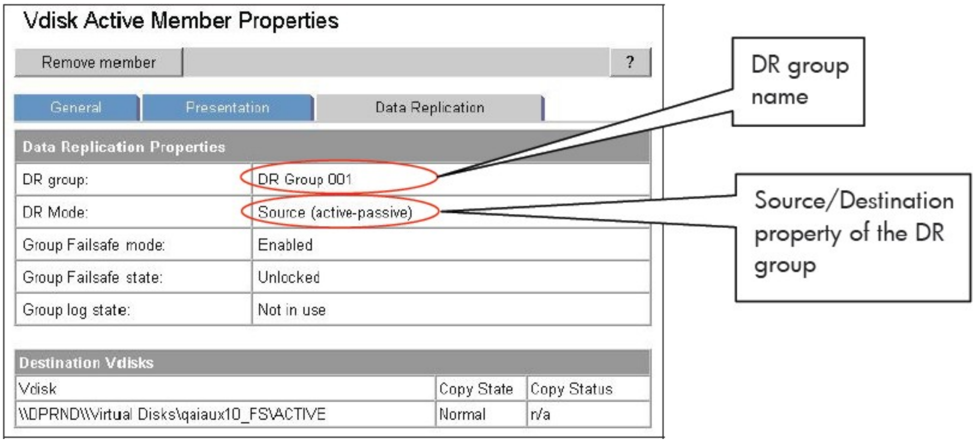

Vdisk Active Member Properties

Data Replication Properties

DR group:	DR Group 001
DR Mode:	Source (active-passive)
Group Failsafe mode:	Enabled
Group Failsafe state:	Unlocked
Group log state:	Not in use

Destination Vdisks

Vdisk	Copy State	Copy Status
\\DPRND\\Virtual Disks\qaiaux10_FS\ACTIVE	Normal	n/a

The DR mode is used to identify the configuration of the current environment:

- If the DR mode is "Source", the current environment is Configuration I – local HP Business Copy P6000 EVA. In this case, proceed to "Step 3: Modifying or removing the CA link" (page 139).
- If the DR mode is "Destination", the current configuration is Configuration II – remote HP Business Copy P6000 EVA. In this case, proceed to "Step 2: Performing failover" (page 139).

NOTE: Complex environments may include a mixture of Configuration I and Configuration II. In this scenario, business copies exist that are both local and remote in relation to the source storage volumes. To handle this, perform the actions stated in Step 2: Performing Failover only to the DR groups with the "Destination" DR mode.

Step 2: Performing failover

Use the information you have gathered regarding DR groups to perform failover as appropriate for the environment. For simple environments, this may include interactions with CV EVA or HP CA P6000 EVA GUI, as well as some configuration steps on the application system. Refer to the appropriate documentation for full details before taking such actions.

For more complex environments, including clusters or other high-availability solutions, refer to the appropriate documentation for that solution before performing any failover actions.

After performing the failover, proceed to step 3 to modify or remove the HP CA P6000 EVA link.

Step 3: Modifying or removing the HP CA P6000 EVA link

NOTE: Before taking any action, record the information relating to the DR groups. This includes such things as the virtual disks participating in the DR group, which P6000 EVA Array storage systems are being replicated to, the mode of operation, and other specific details.

Modify the environment so that the source virtual disks no longer participate in a DR group. You can do this is either of two ways:

- Reduce the DR group by removing each source virtual disk. Do this if the environment is complex and simplifying the DR groups will make it easier to reconfigure the environment.
- Delete the DR group completely. Do this if the HP CA P6000 EVA links are no longer needed or are easily reconfigurable.

Refer to the CV EVA user documentation or other specific documentation for details of these methods.

In the case of a DR group reduction, there may only be source storage volumes inside the DR group. A DR group must always have at least one virtual disk participating. In this case, it is advisable to create a temporary virtual disk and add it to the DR group. With this temporary storage volume, all source storage volumes may be removed from the DR group, and the DR group will still persist.

When this has been completed, proceed to step 4 to perform the instant recovery.

Step 4: Performing instant recovery

Using the Data Protector GUI or CLI, perform instant recovery with the selected session. This should complete successfully with the appropriately reconfigured environment.

When this has been completed, optionally proceed to rebuilding the HP CA P6000 EVA link.

Step 5: Rebuilding the HP CA P6000 EVA link (optional)

If required, return the new source virtual disks to the specific DR groups. Using the information you recorded in step 3 regarding the environment and specific DR groups, either rebuild or recreate the DR groups.

NOTE: Ensure that you use the newly-recovered storage volumes for this rebuild of the HP CA P6000 EVA links. These storage volumes should have the same names and the WWNs as the storage volumes used previously. However, as these are different virtual disks, the UUIDs will be different from those used by the application system before for the virtual disks.

Refer to the CV EVA documentation for details. You may also need to perform additional steps to bring the environment to the same initial state, including failing over the HP CA P6000 EVA links again, to return operation to the correct P6000 EVA Array storage systems and application servers.

ZDB omnirc variables

To customize operation of the ZDB agents, you can set environment variables in the file *Data_Protector_program_data*\omnirc (Windows Vista, Windows 7, Windows Server 2008), *Data_Protector_home*\omnirc (other Windows systems), or /opt/omni/.omnirc (UNIX systems). The variables must be set on the application system and the backup system. Changes to the environment variables in the omnirc file on a particular system do not affect the agents that are already running on the system at the moment the changes are made. For information on the omnirc file, see the *HP Data Protector Troubleshooting Guide* or the online Help index: “omnirc”. Instructions on how to set the variables are provided in the file itself.

Common ZDB variables

This section explains omnirc variables that can be set for all ZDB agents.

ZDB_PRESERVE_MOUNTPOINTS: Determines, together with ZDB_MULTI_MOUNT and ZDB_MOUNT_PATH, the mount point creation on the backup system.

If ZDB_PRESERVE_MOUNTPOINTS is set to 0 (default value), the mount point for a backed up filesystem is created as follows:

- When ZDB_MULTI_MOUNT is set to 1:
  - ***P6000 EVA Array:***

    BU_MOUNT_PATH/Application_System_Name/Mount_Point_Name_on_Application_System_SessionID
  - ***P9000 XP Array:***

    BU_MOUNT_PATH/Application_System_Name/Mount_Point_Name_on_Application_System_LDEV_MU#
- When ZDB_MULTI_MOUNT is set to 0 or not set:

  BU_MOUNT_PATH/Application_System_Name/Mount_Point_Name_on_Application_System

where *BU_MOUNT_PATH* corresponds to one of the following locations on a Data Protector client:

- With ZDB_MOUNT_PATH set: *ZDB_MOUNT_PATH*
- With ZDB_MOUNT_PATH not set:
  - ***Windows Server 2008:*** *Data_Protector_program_data*\tmp
  - ***Other Windows system:*** *Data_Protector_home*\tmp
  - ***UNIX systems:*** /var/opt/omni/tmp

If ZDB_PRESERVE_MOUNTPOINTS is set to 1, the mount point for a backed up filesystem is created on the backup system at:

- ***Windows systems:*** \Mount_Point_Name_on_Application_System or *Drive_Letter_on_Application_System*:\
- ***UNIX systems:*** /*Mount_Point_Name_on_Application_System*

IMPORTANT: For zero downtime backup of disk images, Oracle 8/9/10 databases, SAP R/3 databases, and Microsoft SQL Server 2000 databases, Data Protector adopts that ZDB_PRESERVE_MOUNTPOINTS is set to 1, and ignores its override and the variables ZDB_MULTI_MOUNT and ZDB_MOUNT_PATH.

ZDB_MULTI_MOUNT: Determines, together with ZDB_PRESERVE_MOUNTPOINTS and ZDB_MOUNT_PATH, the mount point creation on the backup system.

ZDB_MULTI_MOUNT is ignored if ZDB_PRESERVE_MOUNTPOINTS is set to 1.

If ZDB_MULTI_MOUNT is set to 1 (default value), *SessionID* (P6000 EVA Array) or *LDEV_MU#* (P9000 XP Array) 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 is set to 0, the selected group of mount points for one replica in the replica set is mounted to the same mount points.

IMPORTANT: With EMC Symmterix, this variable is ignored and Data Protector adopts that it is set to 0.

ZDB_MOUNT_PATH: Determines, together with ZDB_PRESERVE_MOUNTPOINTS and ZDB_MULTI_MOUNT, the mount point creation on the backup system.

`ZDB_MOUNT_PATH` is ignored if `ZDB_PRESERVE_MOUNTPOINTS` is set to `1`.

By default, this variable is not set. In this case, the first part of the mount point path is defined as:

Windows Server 2008: *`Data_Protector_program_data`*`\tmp`

Other Windows systems: *`Data_Protector_home`*`\tmp`

UNIX systems: `/var/opt/omni/tmp`

To set this variable, specify the first part of the mount point path.

NOTE: If the option **Use the same mountpoints as on the application system** is not selected in the GUI, the variable `ZDB_MOUNT_PATH` is ignored and values of the ZDB options **Root of the mount path on the backup system** and **Add directories to the mount path** specified in the Data Protector GUI are used for mount point creation in the ZDB session instead.

ZDB_ALWAYS_POST_SCRIPT: By default, the command specified in the option **Restart the application command line** is not executed if the command specified in the option **Stop/quiesce the application command line** fails.

If this variable is set to `1`, the command specified in the option **Restart the application command line** is always executed.

Default: `0`.

ZDB_IR_VGCHANGE: On HP-UX platform, determines the mode in which replicated volume groups on the application system are activated after restore. The variable can be set on the application system only.

NOTE: This variable is 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.

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`.

P6000 EVA Array specific variables

This section explains P6000 EVA Array-specific `omnirc` variables.

See also “Common ZDB variables” (page 140).

EVA_HOSTNAMEALIASES: Allows a given ID to match the P6000 EVA Array host objects.

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 P6000 EVA Array client to find these host objects, set:

```
EVA_HOSTNAMEALIASES=MyHost_Port1,MyHost_Port2
```

EVA_MSGWAITING_INTERVAL: Determines 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 n minutes if the snapclones are not fully created` must be selected.

Default: 10 minutes.

EVA_CLONECREATION_QUERY_INTERVAL: Determines 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 n 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 P6000 EVA Array storage system.

Default: 5 minutes.

ZDB_VOLUMESCAN_RETRIES: During the backup system preparation, the system is scanned for new filesystem volumes. This variable determines the number of scans required to identify the new volumes.

The variable is only applicable on Windows.

Default: 5 retries. If scanning takes longer (a known problem on Windows Server 2003), increase the default setting.

ZDB_POST_RESCAN_INIT_DELAY: During the backup system preparation, the system is scanned for new filesystem volumes. This variable sets the time period to wait before initiating next scan of new filesystem volumes.

The variable is only applicable on Windows.

Default: 30 seconds.

ZDB_LVM_PREFERRED_PVG: With HP-UX LVM mirroring, determines the physical volume group (PVG) to be selected for HP BC P6000 EVA pair replication. Data Protector checks the value of this variable when at least one logical volume identified as a backup object is mirrored.

The variable format is as follows:

```
ZDB_LVM_PREFERRED_PVG=VGNAME1:PVG_NAME;VGNAME2:PVG_NAME; ...
```

For example, if three volume groups are participating in backup, you can define the following in your application system configuration file:

```
ZDB_LVM_PREFERRED_PVG=/dev/vg01:PVG-0;/vgAppln1:PVG-1;
/dev/vgAppln2:PVG-1
```

When the backup objects are from the volume group `/dev/vg01`, the HP StorageWorks P6000 EVA SMI-S Agent applies the mirror selection rules and prefers `PVG-0` over any other valid PVG defined for if when the same disk array is used. On other disk array, any valid PVG will be used.

ZDB_SMISA_LVM_MIRRORING_DISABLED: Data Protector versions prior A.06.00 do not support LVM mirroring using the HP StorageWorks P6000 EVA SMI-S Agent. Starting with Data Protector A.06.00, LVM mirroring has been a default feature available with the HP StorageWorks P6000 EVA SMI-S Agent which has applied certain configuration restrictions. As a consequence, some configurations may become unsupported after upgrade of the HP StorageWorks P6000 EVA SMI-S Agent from the version A.05.50. Use `ZDB_SMISA_LVM_MIRRORING_DISABLED` to disable LVM mirroring and continue performing ZDB sessions in your older configurations, unsupported by LVM mirroring.

Default: `0` (LVM mirroring enabled). Possible: `0|1`.

EVACA_QUERY_INTERVAL: Determines the time interval (in minutes) between queries of the P6000 EVA Array storage system for checking the progress of the logging and/or copying process on HP CA+BC P6000 EVA. Such querying occurs during a ZDB-to-tape session immediately after backup system resolving.

Default: 5 minutes.

EVACA_WAIT_FOR_NORMAL_STATE: Determines if the HP StorageWorks P6000 EVA SMI-S Agent has to wait for the DR group write history log (DR group log) to move out of the "logging", "copying", or "merging" state back to the "not in use" state. With this variable set, the DR group log state is monitored for the time period set by `EVACA_LOGGINGSTATE_TIMEOUT`, `EVACA_COPYSTATE_TIMEOUT`, or `EVACA_MERGINGSTATE_TIMEOUT`. If at the end of the period the DR group log state has not returned to "not in use", the backup for the objects belonging to that DR group is aborted.

Default: `0` (not set). Possible: `0|1`.

EVACA_WAIT_FOR_NORMAL_STATE_TIMEOUT: Determines the time interval (in minutes) to wait for the DR group write history log (DR group log) found in the "logging", "copying", or "merging" state to move to the "not in use" state. After the timeout, the backup process skips the objects belonging to DR groups whose log is a state other than "not in use", and continues with backup of other objects specified in a ZDB backup specification. The variable is only considered when `EVACA_WAIT_FOR_NORMAL_STATE` is set to `1`.

Default: 15 minutes.

EVACA_LOGGINGSTATE_TIMEOUT: Determines the time interval (in minutes) to wait for the DR group write history log (DR group log) found in the "logging" state to move to the "not in use" state. After the timeout, backup process skips the objects belonging to DR groups whose log is in the "logging" state, and continues with backup of other objects specified in a ZDB backup specification.

Default: 10 minutes.

EVACA_MSGWAITING_INTERVAL: Determines the time interval (in minutes) between messages that report the progress of the logging and/or copying process on HP CA+BC P6000 EVA. This progress is monitored during a ZDB-to-tape session immediately after backup system resolving.

Default: 10 minutes.

EVACA_COPYSTATE_TIMEOUT: Determines the time interval (in minutes) after which the backup process stops waiting for the DR group write history log (DR group log) found in the "copying" state to move to the "not in use" state. Backup process skips the source virtual disks (in case of the source virtual disks backup) or the destination virtual disks (in case of the destination virtual disks backup) belonging to the DR groups whose log is in the "copying" state, and continues with backup of other objects specified in a ZDB backup specification.

Default: 15 minutes.

EVACA_MERGINGSTATE_TIMEOUT: Determines the time interval (in minutes) after which the backup process stops waiting for the DR group write history log (DR group log) found in the "merging" state to move to the "not in use" state. Backup process skips the source virtual disks (in case of the source virtual disks backup) or the destination virtual disks (in case of the destination virtual disks backup) belonging to the DR groups whose log is in the "merging" state, and continues with backup of other objects specified in a ZDB backup specification.

Default: 15 minutes.

SMISA_BACKUPPREPARE_RETRY: Determines the number of the HP StorageWorks P6000 EVA SMI-S Agent queries checking for completion of container allocation or creation and setting the write cache policy on the source volumes to the write-through mode during zero downtime backup sessions. If the operations do not complete by the time the last query is made, the HP StorageWorks P6000 EVA SMI-S Agent aborts the currently running session.

Default: 10 queries.

SMISA_BACKUPPREPARE_DELAY: Determines the interval (specified in seconds) between the HP StorageWorks P6000 EVA SMI-S Agent queries checking for completion of container allocation or creation and setting write cache policy on the source volumes to the write-through mode during zero downtime backup sessions.

Default: 120 seconds.

SMISA_CONTAINERCREATION_RETRY: Determines the number of the HP StorageWorks P6000 EVA SMI-S Agent queries checking for completion of container allocation or creation during instant recovery sessions. If the operation does not complete by the time the last query is made, the HP StorageWorks P6000 EVA SMI-S Agent aborts the currently running session.

Default: 10 queries.

SMISA_CONTAINERCREATION_DELAY: Determines the interval (specified in seconds) between the HP StorageWorks P6000 EVA SMI-S Agent queries checking for completion of container allocation or creation during instant recovery sessions.

Default: 120 seconds.

SMISA_MSGWAITING_INTERVAL: Determines the interval (specified in seconds) between session messages reporting the progress of container allocation or creation during zero downtime backup and instant recovery sessions, between session messages reporting the progress of setting the write cache policy on the source volumes to the write-through mode during zero downtime backup sessions, and between session messages reporting the progress of deleting storage volumes from the disk array.

Default: 300 seconds.

SMISA_CHECKFORABORT_DELAY: Many operations that the HP StorageWorks P6000 EVA SMI-S Agent triggers are long-lasting. During the wait for their completion, the agent periodically checks whether an abort request was issued. This variable determines the interval (specified in seconds) between each pair of checks for the abort request.

Default: 2 seconds.

SMISA_FORCE_DISMOUNT: On Windows Server 2008 systems, determines whether the Data Protector HP StorageWorks P6000 EVA SMI-S Agent performs forced dismount of the volumes which are locked by the Windows system processes and cannot be dismounted using the ordinary dismount operation. You can enable forced dismount operation by setting this variable to `1`.

Default: `0` (disabled). Possible: `0`| `1`.

ZDB_DONOT_PRESENT_DISKS: During ZDB-to-disk sessions, if this variable is set to `1`, the HP StorageWorks P6000 EVA SMI-S Agent does not present volumes to the backup system. ZDB-to-disk+tape and ZDB-to-tape sessions are not affected by the variable.

Default: `0` (disabled). Possible: `0`| `1`.

ZDB_SKIP_LOCK_AT_DISMOUNT: On Windows systems, the HP StorageWorks P6000 EVA SMI-S Agent locks the volumes on the application system prior to dismounting them. If this variable is set to `1`, the volumes do not get locked. Other operating systems are not affected by the variable.

Default: `0` (disabled). Possible: `0`| `1`.

EVA_CIMOM_CONNECTION_TIMEOUT: Determines the interval (specified in seconds) for which the HP StorageWorks P6000 EVA SMI-S Agent waits for a response to an outstanding request from the CIMOM.

Default: 900 seconds.

EVA_CIMOM_QUERY_RETRIES: CIMOM operations include communication over the network and may fail unexpectedly. This variable determines the maximum number of retried attempts the HP StorageWorks P6000 EVA SMI-S Agent performs if the CIMOM returns an unexpected response.

Default: 10 attempts.

EVA_CIMOM_QUERY_INTERVAL: Determines the interval (specified in seconds) between each pair of attempts, whose maximum number is defined by `EVA_CIMOM_QUERY_RETRIES`.

Default: 10 seconds.

SMISA_ENFORCE_MULTISNAP: Determines how Data Protector behaves in either of the following cases:

- multisnapping is not supported by the current P6000 EVA Array configuration
- disk array limitation on the number of source disks that can be involved in multisnapping is exceeded
- source disks are located on more than one P6000 EVA Array storage system

If multisnapping is enforced by `SMISA_ENFORCE_MULTISNAP`, the zero downtime backup session is aborted.

If multisnapping is not enforced, Data Protector creates target volumes sequentially (in the first case) or attempts to create target volumes with several multisnapping operations instead of only one (in the other two cases).

Note that `SMISA_ENFORCE_MULTISNAP` should not be used to enforce multisnapping in zero downtime backup sessions for backing up the Oracle Server data in ASM configurations, since the HP StorageWorks P6000 EVA SMI-S Agent detects such sessions automatically.

Default: `0` (multisnapping not enforced). Possible: `0| 1`.

SMISA_WAIT_MIRRORCLONE_PENDING_TIMEOUT: Determines the time period (specified in minutes) for which the HP StorageWorks P6000 EVA SMI-S Agent waits for the mirrorclone link to transition from some other state into the synchronized state. If the time period expires before the mirrorclone link gets into the synchronized state, the HP StorageWorks P6000 EVA SMI-S Agent aborts the session. The variable affects only the zero downtime backup sessions for which the selected snapshot source is mirrorclone.

Default: 60 minutes.

SMISA_WAIT_MIRRORCLONE_PENDING_RETRY: Determines the interval (specified in seconds) between each pair of checks of the mirrorclone link state when waiting for the link to transition into the synchronized state. The variable affects only the zero downtime backup sessions for which the selected snapshot source is mirrorclone.

Default: 30 seconds.

P9000 XP Array specific variables

This section explains P9000 XP Array-specific `omnirc` variables.

See also "Common ZDB variables" (page 140).

ZDB_BACKUP_VG_EXIST: On HP-UX platform, 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 or restore session.

NOTE: If this variable is set, volume groups are not removed from `/etc/lvmtab` on the backup system after each backup. For more information, see "Backup options" (page 74).

Default: `0`.

OB2AUTOPATH_BALANCING_POLICY: Determines the HP 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.

Possible values are:

- `0` [none] – No policy
- `1` [RR] – Round Robin policy (default)
- `2` [SQL] – Shortest Queue Length policy

IMPORTANT: During a ZDB-to-tape session, 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

For more information, see the AutoPath documentation.

SSEA_SPLIT_REPORT_RATE: During the split, the HP StorageWorks P9000 XP Agent checks the status of mirrored disks within an interval determined by `SSEA_SPLIT_SLEEP_TIME` for the number of times determined by `SSEA_SPLIT_RETRY`.

`SSEA_SPLIT_REPORT_RATE` determines 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` is `5`, the status is displayed for every fifth check (every 10 seconds).

Default: `5`.

SSEA_SPLIT_RETRY: During the split, the HP StorageWorks P9000 XP Agent checks the mirrored disks' status within an interval determined by `SSEA_SPLIT_SLEEP_TIME`. `SSEA_SPLIT_RETRY` determines 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, the HP StorageWorks P9000 XP Agent checks the mirrored disks status for the number of times determined by `SSEA_SPLIT_RETRY`. `SSEA_SPLIT_SLEEP_TIME` determines the time interval between the status checks.

Default: 2 seconds.

SSEA_SYNC_REPORT_RATE: During the disks' resynchronization, the HP StorageWorks P9000 XP Agent checks the mirrored disks' status within an interval determined by `SSEA_SYNC_SLEEP_TIME` for the number of times determined by `SSEA_SYNC_RETRY`.

`SSEA_SYNC_REPORT_RATE` determines the rate of displaying the mirrored disks status. For example, if `SSEA_SYNC_SLEEP_TIME` is 5 seconds and `SSEA_SPLIT_REPORT_RATE` is `2`, the status is displayed for every second check (every 10 seconds).

Default: `2`.

SSEA_SYNC_RETRY: During the disks' resynchronization, the HP StorageWorks P9000 XP Agent checks the mirrored disks' status within an interval specified by `SSEA_SYNC_SLEEP_TIME`. `SSEA_SYNC_RETRY` determines 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, the HP StorageWorks P9000 XP Agent checks the mirrored disks' status for the number of times determined by `SSEA_SYNC_RETRY`. `SSEA_SYNC_SLEEP_TIME` determines the time interval between the 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 determines the maximum waiting period for disk pairs to change to proper status.

SMB_SCAN_RDSK_TIMEOUT: On Windows, during backup system preparation, the system is scanned for new devices. 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: On Windows, 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. 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.

Default: 120 minutes.

SSEA_FORCE_DISMOUNT: On Windows Server 2008 systems, determines whether the HP StorageWorks P9000 XP Agent will perform forced dismount of the volumes which are locked by the Windows system processes and cannot be dismounted using the ordinary dismount operation. You can enable forced dismount operation by setting this variable to `1`.

Default: `0` (disabled). Possible: `0| 1`.

MAXIMUM_HOST_LOCKING_RETRY: The HP StorageWorks P9000 XP Agent will lock the backup system during the backup system preparation. The lock operation may fail due to concurrent ZDB sessions or similar actions. This variable determines the maximum number of attempts by the HP StorageWorks P9000 XP Agent at locking the backup system.

Default: 60 attempts.

SSEA_ATTACH_RETRY: Prior to manipulating volumes on a disk array, the HP StorageWorks P9000 XP Agent must connect to an appropriate command device. In case of a problem with the SAN connectivity, establishing such a connection may fail. This variable determines the number of attempts made by the HP StorageWorks P9000 XP Agent at connecting to the command device.

Default: 5 attempts.

SSEA_ATTACH_SLEEP_TIME: Determines the interval (specified in seconds) between each pair of attempts of HP StorageWorks P9000 XP Agent at connecting to the command device.

Default: 10 seconds.

EMC specific variables

This section explains EMC-specific `omnirc` variables.

See also "Common ZDB variables" (page 140).

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.

P6000 EVA Array integration

Example 1

ZDB 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`.
- Select the desired snapshot source.
- Select the desired snapshot type.
- Select **Same as source** for the redundancy level.

The following option is selected automatically:

- **Keep the replica after the backup**

Then, schedule the ZDB backup specification to start three ZDB-to-disk sessions during the day and one ZDB-to-disk+tape session during the night.

Example 2

ZDB to tape must be performed every three hours. Replicas created are used for data mining (not for instant recovery) for the time period of three hours.

To implement such policy:

- Clear **Track the replica for instant recovery**.
- Select **Keep the replica after the backup**.
- Set **Number of replicas rotated** to `1`.
- Select the desired snapshot source.
- Select the desired snapshot type.
- Select **Leave the backup system enabled**.
- On UNIX systems, optionally select **Enable the backup system in read/write mode**.
- Set the `omnirc` variable `ZDB_ORA_INCLUDE_CF_OLF` to `1`. For more information, see the *HP Data Protector Zero Downtime Backup Integration Guide*

Then, schedule the ZDB backup specification to start one ZDB-to-tape session every three hours.

Example 3

ZDB 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`.
- Select the desired snapshot source.
- Select the desired snapshot type.
- Select **Same as source** for the redundancy level.

The following option is selected automatically:

- **Keep the replica after the backup**

Then, schedule the ZDB backup specification to start eight ZDB-to-disk+tape sessions every three hours.

P9000 XP Array 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 zero downtime backup and forcibly synchronized before the next zero downtime backup.

To implement such policy, select the following options:

- **Track the replica for instant recovery**
- **Synchronize the disks if not already synchronized**
- **Prepare the next mirror disks for the backup (resynchronize)**

The following option is selected automatically:

- **Keep the replica after the backup**

Example 2

A replica set is configured, with all replicas available for offline data processing after the ZDB session. The next replica must be prepared according to replica set rotation after the zero downtime backup, and the next ZDB session 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**
- **Abort the session if the mirror disks are not synchronized**
- **Prepare the next mirror disks for the backup (resynchronize)**
- **Leave the backup system enabled**

Example 3

A replica set is configured, with versions on replicas 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 ZDB session.

To implement such policy:

- Select **Synchronize the disks if not already synchronized**.
- Clear **Keep the replica after the backup**.

Example 4

A single replica is configured, with the version on the replica available for offline data processing. The replica must be prepared at the start of a ZDB session.

To implement such policy, select the following options:

- **Keep the replica after the backup**
- **Synchronize the disks if not already synchronized**
- **Leave the backup system enabled**

Conflicting Options

If a single replica is configured and the following options are selected, the second option is ignored, since the replica to be kept is at the same time the replica to be prepared for the next zero downtime backup:

- **Keep the replica after the backup**
- **Prepare the next mirror disks for the backup (resynchronize)**

NOTE: A conflict may also occur when a replica set is configured, depending on the replica set selection and the P9000 XP LDEV exclude file.

EMC integration

Example 1

After zero downtime backup, the replica must be discarded and prepared for the next zero downtime backup at the end of the ZDB session.

To implement such backup policy:

- Select **Re-establish links after backup**.
- Do not select **Re-establish links before backup**.

Example 2

After zero downtime backup, the replica must be used for offline data processing and prepared at the start of the next ZDB session.

To implement such backup policy:

- Select **Re-establish links before backup**.
- Do not select **Re-establish links after backup**.

Backup system mount point creation

Data Protector disk array integrations support configurations where multiple application systems are connected to a disk array and one system (the backup system) is responsible for backing up these applications. Local, remote, or remote plus local replication configuration (if supported on a particular array) can be used for ZDB in such a configuration. For more information on supported configurations, see the *HP Data Protector Zero Downtime Backup Concepts Guide*.

Each application system uses its own original storage, from which replicas are created; in case of ZDB to tape and ZDB to disk+tape, filesystems are mounted on the backup system.

Filesystem and Microsoft Exchange Server backup

To perform a concurrent backup of multiple application systems, the mount points assigned to the filesystems in the original storage *do not need to be* different for each application system. The backup of the Microsoft Exchange Server application is performed as *filesystem* backup. With filesystem backup, Data Protector, during a ZDB session, creates or reuses unique mount points on the backup system. Data Protector then mounts filesystems to these mount points.

Figure 37 Backup system mount point creation: filesystem and Microsoft Exchange Server backup

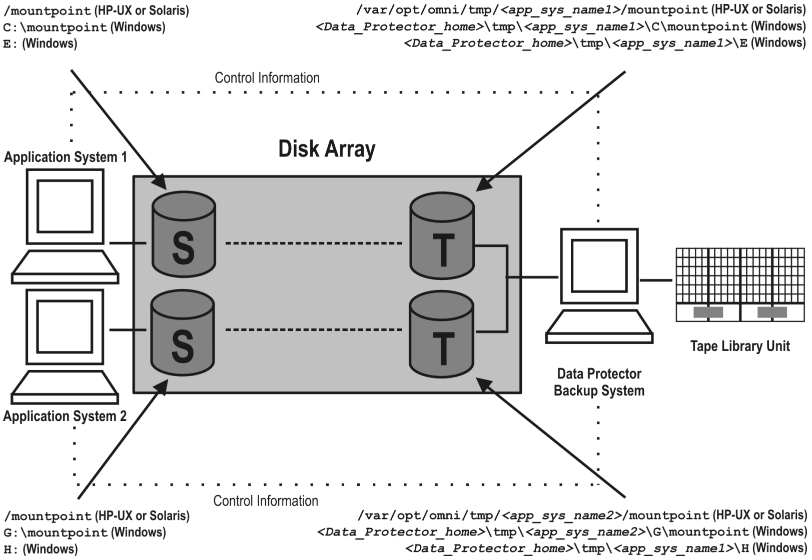
/mountpoint (HP-UX or Solaris)
C:\mountpoint (Windows)
E: (Windows)
/var/opt/omni/tmp/<app_sys_name1>/mountpoint (HP-UX or Solaris)
<Data_Protector_home>\tmp\<app_sys_name1>\C\mountpoint (Windows)
<Data_Protector_home>\tmp\<app_sys_name1>\E (Windows)
Control Information
Application System 1
Disk Array
S
T
S
T
Tape Library Unit
Data Protector Backup System
Application System 2
Control Information
/mountpoint (HP-UX or Solaris)
G:\mountpoint (Windows)
H: (Windows)
/var/opt/omni/tmp/<app_sys_name2>/mountpoint (HP-UX or Solaris)
<Data_Protector_home>\tmp\<app_sys_name2>\G\mountpoint (Windows)
<Data_Protector_home>\tmp\<app_sys_name1>\H (Windows)

NOTE: The above example depicts the default Data Protector behavior. You can change the backup system mount point pathname creation by setting the `ZDB_PRESERVE_MOUNTPOINTS`, `ZDB_MOUNT_PATH` and `ZDB_MULTI_MOUNT` variables in the `.omnirc` file.

Application and disk image backup

The information in this section applies only for the backup of the following:

- Disk images
- Oracle
- SAP R/3
- Microsoft SQL Server

For a list of applications, supported for a particular type of a disk array, see the *HP Data Protector Product Announcements, Software Notes, and References*.

Applications on filesystems

To perform a concurrent backup of multiple application systems, the mount points or drive letters assigned to the original storage *must be* different for each application system. Data Protector, during a ZDB session, creates mount points or drive letters with the same names as on the application system. Data Protector then mounts filesystems in a replica to these mount points.

If the mount points or drive letters are the same for different application systems, concurrent backup of such systems is not possible; backup of objects that belong to these mount points or drive letters must be run sequentially.

Applications on disk images + disk image backup

If your application uses raw disk images as the data source, or if you are performing a disk image backup without an application, the following applies: Data Protector, during a ZDB session, finds and uses raw device files (UNIX systems) or physical drive numbers (Windows systems) for the replica created from the original storage raw device files (UNIX systems) or physical drive numbers (Windows systems) on the backup system. Therefore, make sure the device file names and physical drive numbers are the same on the application and the backup systems.

Note that due to the limitation described above, snapshot integrations are not suitable for such backups (with snapshot integrations, Data Protector cannot guarantee that after presentation to the backup system replicas are assigned the same raw device files or physical drive numbers as on the application system).

NOTE: With the HP P9000 XP Disk Array Family, if the HP Business Copy (BC) P6000 EVA first-level mirrors or snapshot volumes are configured, the integration always mounts the selected first-level mirror or snapshot volume to the same mount point.

Figure 38 Backup system mount point creation: application or disk image backup

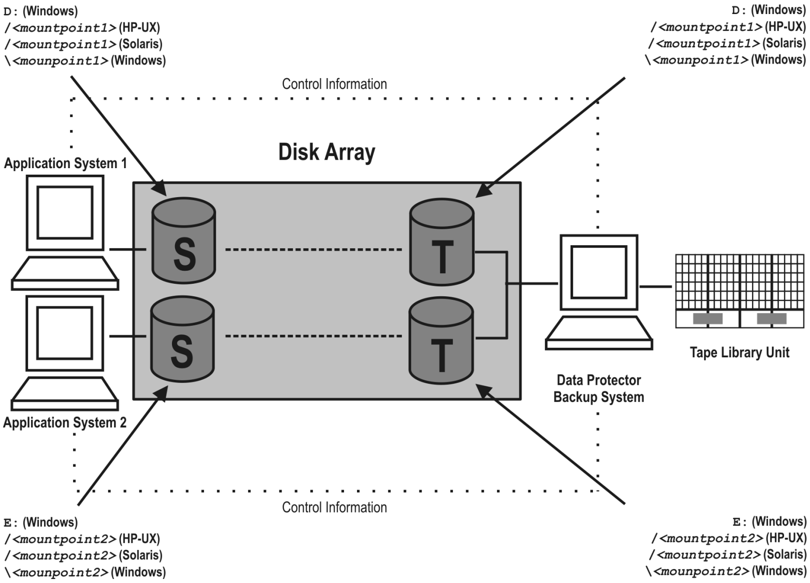

EMC – obtaining disk configuration data

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**
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

Example 2

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**
HP-UX: c0t5d0 Windows: DRIVE13	028	+	0	HP-UX: c0t10d0 Windows: Not Visible	048	0	Synch
HP-UX: c0t5d1 Windows: DRIVE14	029	+ p	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

Additional information for troubleshooting

HP-UX systems

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 systems

Run `symntctl` with additional parameters to get information about disks, signatures, and drives. For more information, see the EMC documentation.

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

A

access rights *See* user rights.

ACSLS *(StorageTek specific term)* The Automated Cartridge System Library Server (ACSLS) software that manages the Automated Cartridge System (ACS).

Active Directory *(Windows specific term)* The directory service in a Windows network. It contains information about resources on the network and makes them accessible to users and applications. The directory services provide a consistent way to name, describe, locate, access and manage resources regardless of the physical system they reside on.

AES 256-bit encryption Data Protector software encryption, based on the AES-CTR (Advanced Encryption Standard in Counter Mode) encryption algorithm that uses random keys of 256-bit length. The same key is used for both encryption and decryption. With AES 256-bit encryption, data is encrypted before it is transferred over a network and before it is written to media.

AML *(ADIC/GRAU specific term)* Automated Mixed-Media library.

AMU *(ADIC/GRAU specific term)* Archive Management Unit.

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.

archive logging *(Lotus Domino Server specific term)* Lotus Domino Server database mode where transaction log files are overwritten only after they have been backed up.

archived redo log *(Oracle specific term)* Also called offline redo log. If the Oracle database operates in the ARCHIVELOG mode, as each online redo log is filled, it is copied to an archived log destination. This copy is the archived redo log. The presence or absence of an archived redo log is determined by the mode the database is using:

- ARCHIVELOG - The filled online redo log files are archived before they are reused. The database can be recovered if an instance or a disk fails. A “hot” backup can be performed only when the database is running in this mode.
- NOARCHIVELOG - The filled online redo log files are not archived.

See also online redo log.

ASR set A collection of files stored on several diskettes, required for proper reconfiguration of the replacement disk (disk partitioning and logical volume configuration) and automatic recovery of the original system configuration and user data that was backed up during the full client backup. These files are stored as an ASR archive file on the Cell Manager in the directory `*Data_Protector_program_data*\Config\Server\dr\asr` (Windows Server 2008), `*Data_Protector_home*\Config\Server\dr\asr` (other Windows systems), or `/etc/opt/omni/server/dr/asr` (UNIX systems) as well as on the backup medium. After a disaster occurs, the ASR archive file is extracted to diskettes which you need to perform ASR.

audit logs Data files to which auditing information is stored.

audit report User-readable output of auditing information created from data stored in audit log files.

auditing information Data about every backup session that was performed over an extended, user-defined period for the whole Data Protector cell.

autochanger *See* library.

autoloader *See* library.

Automatic Storage Management (ASM) *(Oracle specific term)* A filesystem and volume manager integrated into Oracle which manages Oracle database files. It eliminates complexity associated with data and disk management and optimizes performance by providing striping and mirroring capabilities.

automigration *(VLS specific term)* The functionality that allows data backups to be first made to the VLS' virtual tapes and then migrated to physical tapes (one virtual tape emulating one physical tape) without using an intermediate backup application.
See also Virtual Library System (VLS) and virtual tape.

auxiliary disk A bootable disk that has a minimal operating system with networking and Data Protector Disk Agent installed. It can be carried around and used to boot the target system in Phase 1 of Disk Delivery Disaster Recovery of UNIX clients.

B

BACKINT *(SAP R/3 specific term)* SAP R/3 backup programs can call the Data Protector backint interface program via an open interface, which enables them to communicate with Data Protector software. For backup and restore, SAP R/3 programs issue orders for the Data Protector backint interface.

backup API The Oracle interface between the Oracle backup/restore utility and the backup/restore media management layer. The interface defines a set of routines to allow the reading and writing of data to the backup media, the creation, searching and removing the backup files.

backup chain *See* restore chain.

backup device A device configured for use with Data Protector, which can write data to and read data from storage media. This can be, for example, a standalone DDS/DAT drive or a library.

backup generation One backup generation includes one full backup and all incremental backups until the next full backup.

backup ID An identifier of an integration object that equals the session ID of the backup of this object. The backup ID is preserved when an object is copied, exported, or imported.

backup object A backup unit that contains all items backed up from one disk volume (logical disk or mount point). The backed up items can be any number of files, directories, or the entire disk or mount point. Additionally, a backup object can be a database/application entity or a disk image (rawdisk).

A backup object is defined by:

- Client name: Hostname of the Data Protector client where the backup object resides.
- Mount point: For filesystem objects — the access point in a directory structure on the client where the backup object is located (drive on Windows and mount point on UNIX). For integration objects — backup stream identification, indicating the backed up database/application items.
- Description: For filesystem objects — uniquely defines objects with identical client name and mount point. For integration objects — displays the integration type (for example, SAP or Lotus).
- Type: Backup object type. For filesystem objects — filesystem type (for example, WinFS). For integration objects — "Bar".

backup owner Each backup object in the IDB has an owner. The default owner of a backup is the user who starts the backup session.

backup session A process that creates a copy of data on storage media. The activities are specified in a backup specification or an interactive session. All clients configured in one backup specification are backed up together in one backup session using the same backup type. The result of a backup session is a set of media, which was written to, also called the backup or media set.
See also backup specification, full backup, and incremental backup.

backup set A complete set of integration objects associated with a backup.

backup set *(Oracle specific term)* A logical grouping of backed up files that are created using the RMAN backup command. A backup set is a complete set of files associated with a backup. The files can be multiplexed to increase performance. A backup set contains either datafiles or archived logs, but not both together.

backup specification A list of objects to be backed up, together with a set of devices or drives to be used, backup options for all objects in the specification, and days and time that you want backups to be performed. The objects are entire disks/volumes or parts of them such as files, directories, or even the Windows Registry for example. File selection lists such as include-lists and exclude-lists can be specified.

backup system *(ZDB specific term)* A system connected to a disk array together with one or multiple application systems. The backup system is typically connected to a disk array to create target volumes (a replica) and is used for mounting the target volumes (the 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 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.

BCV *(EMC Symmetrix specific term)* Business Continuance Volumes, or BCV devices, are dedicated SLDs that are pre-configured in the ICDA on which the business continuation operation runs. BCV devices are assigned separate SCSI addresses, differing from the addresses used by the SLDs they mirror. The BCV devices are used as splittable mirrors of the primary EMC Symmetrix SLDs that need to be protected.
See also BC and BC Process.

Boolean operators The Boolean operators for the full text search functionality of the online Help system are AND, OR, NOT, and NEAR. Used when searching, they enable you to define your query precisely by creating a relationship between search terms. If no operator is specified in a multi-word search, AND is used by default. For example, the query manual disaster recovery is equivalent to manual AND disaster AND recovery.

boot volume/disk/partition A volume/disk/partition with files required for the initial step of the boot process. Microsoft terminology defines the boot volume/disk/partition as a volume/disk/partition containing the operating system files.

BRARCHIVE *(SAP R/3 specific term)* An SAP R/3 backup tool that allows you to archive redo log files. BRARCHIVE also saves all the logs and profiles of the archiving process.
See also BRBACKUP and BRRESTORE.

BRBACKUP *(SAP R/3 specific term)* An SAP R/3 backup tool that allows an online or offline backup of the control file, of individual data files, or of all tablespaces and, if necessary, of the online redo log files.
See also BRARCHIVE and BRRESTORE.

BRRESTORE *(SAP R/3 specific term)* An SAP R/3 tool that can be used to restore files of the following type:

- Database data files, control files, and online redo log files saved with BRBACKUP
- Redo log files archived with BRARCHIVE
- Non-database files saved with BRBACKUP

You can specify files, tablespaces, complete backups, log sequence numbers of redo log files, or the session ID of the backup.
See also BRBACKUP and BRARCHIVE.

BSM The Data Protector Backup Session Manager controls the backup session. This process always runs on the Cell Manager system.

C

CAP *(StorageTek specific term)* Cartridge Access Port is a port built into the door panel of a library. The purpose is to enter or eject media.

catalog protection Defines how long information about backed up data (such as file names and file versions) is kept in the IDB.
See also data protection.

CDB The Catalog Database is a part of the IDB that contains information about backup, restore, object copy, object consolidation, object verification, and media management sessions. Depending on the selected logging level, it also contains file names and file versions. This part of the database is always local to the cell.
See also MMDB.

CDF file *(UNIX specific term)* A Context Dependent File is a file consisting of several files grouped under the same pathname. The system ordinarily selects one of the files using the context of the process. This mechanism allows machine dependent executables, system data, and device files to work correctly from all hosts in a cluster while using the same pathname.

cell A set of systems that are under the control of a Cell Manager. The cell typically represents the systems on a site or an organizational entity, which are connected to the same LAN or SAN. Central control is available to administer the backup and restore policies and tasks.

Cell Manager The main system in the cell where the essential Data Protector software is installed and from which all backup and restore activities are managed. The GUI used for management tasks can be located on a different system. Each cell has one Cell Manager system.

centralized licensing Data Protector allows you to configure centralized licensing for the whole enterprise environment consisting of several cells. All Data Protector licenses are installed and kept on the Enterprise Cell Manager system. You can then allocate licenses to specific cells to suit your needs.
See also MoM.

Centralized Media Management Database (CMMDB) *See* CMMDB.

Certificate Server A Windows Certificate Server can be installed and configured to provide certificates for clients. It provides customizable services for issuing and managing certificates for the enterprise. These services issue, revoke, and manage certificates employed in public key-based cryptography technologies.

Change Journal *(Windows specific term)* A Windows filesystem feature that logs a record of each change as it occurs to the files and directories on a local NTFS volume.

Change Log Provider *(Windows specific term)* A module that can be queried to determine which objects on a filesystem have been created, modified, or deleted.

channel *(Oracle specific term)* An Oracle Recovery Manager resource allocation. Every allocated channel starts a new Oracle process, which performs backup, restore, and recovery actions. The type of channel allocated determines the type of media used:

- type 'disk'
- type ‘sbt_tape’

If the specified channel is of type ‘sbt_tape’ and Oracle is integrated with Data Protector, the server process will attempt to read backups from or write data files to Data Protector.

circular logging *(Microsoft Exchange Server and Lotus Domino Server specific term)* Circular logging is a Microsoft Exchange Server database and Lotus Domino Server database mode, in which the transaction log file contents are periodically overwritten after the corresponding data is committed to the database. Circular logging reduces disk storage space requirements.

client backup A backup of all volumes (filesystems) mounted on a Data Protector client. What is actually backed up depends on how you select objects in a backup specification:

- If you select the check box next to the client system name, a single backup object of the `Client System` type is created. As a result, at the time of the backup, Data Protector first detects all volumes that are mounted on the selected client and then backs them up. On Windows clients, `CONFIGURATION` is also backed up.

- If you individually select all volumes that are mounted on the client system, a separate backup object of the `Filesystem` type is created for each volume. As a result, at the time of the backup, only the selected volumes are backed up. Volumes that have been potentially mounted on the client after the backup specification was created are not backed up.

client or client system
Any system configured with any Data Protector functionality and configured in a cell.

cluster continuous replication
(Microsoft Exchange Server specific term) Cluster continuous replication (CCR) is a high availability solution that uses cluster management and failover options to create and maintain an exact copy (CCR copy) of a storage group. A storage group is replicated to a separate server. CCR removes any single point of failure in your Exchange back-end servers. You can perform backups using VSS on your passive Exchange Server node where a CCR copy is located and thus reducing the load on the active node.

A CCR copy is used for disaster recovery since you can switch to the CCR copy in a few seconds. A replicated storage group is represented as a new instance of Exchange writer called Exchange Replication Service and can be backed up (using VSS) like an ordinary storage group.

See also Exchange Replication Service and local continuous replication.

cluster-aware application
It is an application that supports the cluster Application Programming Interface. Each cluster-aware application declares its own critical resources (disk volumes (on Microsoft Cluster Server), volume groups (on MC/ServiceGuard), application services, IP names and addresses, and so on).

CMD script for Informix Server
(Informix Server specific term) A Windows CMD script that is created in INFORMIXDIR when an Informix Server database is configured. The CMD script is a set of system commands that export environment variables for Informix Server.

CMMDB
The Data Protector Centralized Media Management Database (CMMDB) is the result of merging MMDBs from several cells in the MoM environment. It allows you to share high-end devices and media across multiple cells in a MoM environment. One cell can control the robotics, serving the devices that are connected to systems in other cells. The CMMDB must reside on the Manager-of-Managers. A reliable network connection between the MoM cell and the other Data Protector cells is highly recommended
See also MoM.

COM+ Class Registration Database
(Windows specific term) The COM+ Class Registration Database and the Windows Registry store application attributes, class attributes, and computer-level attributes. This guarantees consistency among these attributes and provides common operation on top of these attributes.

command device
(HP P9000 XP Disk Array Family specific term) A dedicated volume in the disk array which acts as the interface between a management application and the disk array's storage system. It cannot be used for data storage and only accepts requests for operations that are then executed by the disk array.

Command View VLS
(VLS specific term) A web browser-based GUI that is used to configure, manage, and monitor the VLS through a LAN.
See also Virtual Library System (VLS).

command-line interface (CLI)
A set of DOS and UNIX like commands that you can use in shell scripts to perform Data Protector configuration, backup, restore, and management tasks.

concurrency
See Disk Agent concurrency.

container
(HP P6000 EVA Disk Array Family specific term) Space on a disk array, which is pre-allocated for later use as a standard snapshot, vsnap, or snapclone.

control file
(Oracle and SAP R/3 specific term) An Oracle data file that contains entries specifying the physical structure of the database. It provides database consistency information used for recovery.

copy set
(HP P6000 EVA Disk Array Family specific term) A pair that consists of the source volumes on a local P6000 EVA and their replica on a remote P6000 EVA.
See also source volume, replica, and HP Continuous Access + Business Copy (CA+BC) P6000 EVA.

CRS
The Cell Request Server process (service), which runs on the Data Protector Cell Manager, and starts and controls the backup and restore sessions. The service is started as soon as Data Protector

is installed on the Cell Manager. On Windows systems, the CRS runs under the account of the user specified at installation time. On UNIX systems, it runs under the account `root`.

CSM The Data Protector Copy and Consolidation Session Manager process controls the object copy and object consolidation sessions and runs on the Cell Manager system.

D

data file *(Oracle and SAP R/3 specific term)* A physical file created by Oracle that contains data structures such as tables and indexes. A data file can only belong to one Oracle database.

data protection Defines how long the backed up data on media remains protected, that is, Data Protector will not overwrite it. When the protection expires, Data Protector will be able to reuse the media in one of the next backup sessions.
See also catalog protection.

data replication (DR) group *(HP P6000 EVA Disk Array Family specific term)* A logical grouping of HP P6000 EVA Disk Array Family virtual disks. It can contain up to eight copy sets provided they have common characteristics and share a common HP CA P6000 EVA log.
See also copy set.

data stream Sequence of data transferred over the communication channel.

Data_Protector_home A reference to the directory containing Data Protector program files (on Windows Vista, Windows 7, and Windows Server 2008) or the directory containing Data Protector program files and data files (on other Windows operating systems). Its default path is *`%ProgramFiles%`*`\OmniBack`, but the path can be changed in the Data Protector Setup Wizard at installation time.
See also Data_Protector_program_data.

Data_Protector_program_data A reference to the directory containing Data Protector data files on Windows Vista, Windows 7, and Windows Server 2008. Its default path is *`%ProgramData%`*`\OmniBack`, but the path can be changed in the Data Protector Setup Wizard at installation time.
See also Data_Protector_home.

database library A Data Protector set of routines that enables data transfer between Data Protector and a server of an online database integration, for example, Oracle Server.

database parallelism More than one database is backed up at a time if the number of available devices allows you to perform backups in parallel.

database server A computer with a large database stored on it, such as the SAP R/3 or Microsoft SQL database. A server has a database that can be accessed by clients.

Dbobject *(Informix Server specific term)* An Informix Server physical database object. It can be a blobspace, dbspace, or logical log file.

DC directory The Detail Catalog (DC) directory contains DC binary files, which store information about file versions. It represents the DCBF part of the IDB, which occupies approximately 80% of the IDB. The default DC directory is called the `dcbf` directory and is located on the Cell Manager in the directory *`Data_Protector_program_data`*`\db40` (Windows Server 2008), *`Data_Protector_home`*`\db40` (other Windows systems), or `/var/opt/omni/server/db40` (UNIX systems). You can create more DC directories and use a custom location. Up to 50 DC directories are supported per cell. The default maximum size of a DC directory is 16 GB.

DCBF The Detail Catalog Binary Files (DCBF) part of the IDB stores information about file versions and attributes. It occupies approximately 80% of the IDB. One DC binary file is created for each Data Protector medium used for backup. Its maximum size is limited by the filesystem settings.

delta backup A delta backup is a backup containing all the changes made to the database from the last backup of any type.
See also backup types.

device A physical unit which contains either just a drive or a more complex unit such as a library.

device chain A device chain consists of several standalone devices configured for sequential use. When a medium in one device gets full, the backup automatically continues on a medium in the next device in the device chain.

device group *(EMC Symmetrix specific term)* A logical unit representing several EMC Symmetrix devices. A device cannot belong to more than a single device group. All devices in a device group must be on the same EMC Symmetrix unit. You can use a device group to identify and work with a subset of the available EMC Symmetrix devices.

device streaming A device is streaming if it can feed enough data to the medium to keep it moving forward continuously. Otherwise, the tape has to be stopped, the device waits for more data, reverses the tape a little and resumes to write to the tape, and so on. In other words, if the data rate written to the tape is less or equal the data rate which can be delivered to the device by the computer system, then the device is streaming. Streaming significantly improves the performance of the device and use of space.

DHCP server A system running the Dynamic Host Configuration Protocol (DHCP) providing dynamic IP address assignment and network configuration for DHCP clients.

differential backup An incremental backup that backs up changes made since the last full backup. To perform this type of backup, specify the Incr1 backup type.
See also incremental backup.

differential backup *(Microsoft SQL Server specific term)* A database backup that records only the data changes made to the database after the last full database backup.
See also backup types.

differential database backup A differential database backup records only those data changes made to the database after the last full database backup.

directory junction *(Windows specific term)* Directory junctions use the reparse point concept of Windows. An NTFS 5 directory junction allows you to redirect a directory/file request to another location.

disaster recovery A process to restore a client's main system disk to a state close to the time when a (full) backup was performed.

disaster recovery operating system *See* DR OS.

Disk Agent A component needed on a client to back it up and restore it. The Disk Agent controls reading from and writing to a disk. During a backup session, the Disk Agent reads data from a disk and sends it to the Media Agent, which then moves it to the device. During a restore session the Disk Agent receives data from the Media Agent and writes it to the disk. During an object verification session the Disk Agent receives data from the Media Agent and performs the verification process, but no data is written to disk.

Disk Agent concurrency The number of Disk Agents that are allowed to send data to one Media Agent concurrently.

disk group *(Veritas Volume Manager specific term)* The basic unit of data storage in VxVM system. A disk group can consist of one or more physical volumes. There can be more than one disk group on the system.

disk image (rawdisk) backup A high-speed backup where Data Protector backs up files as bitmap images. A disk image (rawdisk) backup does not track the files and directory structure stored on the disk, but stores a disk image structure on byte level. You can perform a disk image backup of either specific disk sections or a complete disk.

disk quota A concept to manage disk space consumption for all or a subset of users on a computer system. This concept is used by several operating system platforms.

disk staging The process of backing up data in several phases to improve the performance of backups and restores, reduce costs of storing the backed up data, and increase the data availability and accessibility for restore. The backup stages consist of backing up data to one media type first (for example disk) and later copying it to a different media type (for example tape).

distributed file media format A media format, available with the file library, which supports a space efficient type of synthetic backup called virtual full backup. Using this format is a prerequisite for virtual full backup.
See also virtual full backup.

Distributed File System (DFS) A service that connects file shares into a single namespace. The file shares can reside on the same or on different computers. DFS provides client access to the resources in a location-transparent manner.

DMZ The Demilitarized Zone (DMZ) is a network inserted as a "neutral zone" between a company's private network (intranet) and the outside public network (Internet). It prevents outside users from getting direct access to company servers in the intranet.

DNS server In the DNS client-server model, this is the server containing information about a portion of the DNS database that makes computer names available to client resolvers querying for name resolution across the Internet.

domain controller A server in a network that is responsible for user security and verifying passwords within a group of other servers.

DR image Data required for temporary disaster recovery operating system (DR OS) installation and configuration.

DR OS An operating system environment in which disaster recovery runs. It provides Data Protector with a basic runtime environment (disk, network, tape, and filesystem access). It has to be installed on disk or loaded into memory and configured before the Data Protector disaster recovery can be performed. DR OS can be either temporary or active. A temporary DR OS is used exclusively as a host environment for some other operating system restore along with the target operating system configuration data. It is deleted after the target system is restored to the original system configuration. An active DR OS not only hosts the Data Protector disaster recovery process but can also be a part of the restored system because it replaces its own configuration data with the original configuration data.

drive A physical unit that receives data from a computer system and can write it onto a magnetic medium (typically a tape drive). It can also read the data from the medium and send it to the computer system.

drive index A number that identifies the mechanical position of a drive inside a library device. This number is used by the robotic control to access a drive.

drive-based encryption Data Protector drive-based encryption uses the encryption functionality of the drive. While performing the backup, the drive encrypts both the data and the meta-data that is written to the medium.

E

EMC Symmetrix Agent A Data Protector software module that prepares the EMC Symmetrix environment for backup and restore operations.

emergency boot file *(Informix Server specific term)* The Informix Server configuration file `ixbar.`*`server_id`* that resides in the directory *`INFORMIXDIR`*`/etc` (on Windows) or *`INFORMIXDIR`*`\etc` (on UNIX). *`INFORMIXDIR`* is the Informix Server home directory and *`server_id`* is the value of the `SERVERNUM` configuration parameter. Each line of the emergency boot file corresponds to one backup object.

encrypted control communication Data Protector secure communication between the clients in the Data Protector cell is based on Secure Socket Layer (SSL) that uses SSLv3 algorithms to encrypt control communication. Control communication in a Data Protector cell is all communication between Data Protector processes, except the data transfer from Disk Agent (and Integrations) to Media Agent, and the other way round.

encryption key A 256-bit randomly generated number used by the Data Protector encryption algorithm to encode information during backups for which AES 256-bit software encryption or drive-based encryption has been specified. The same key is used for subsequent decryption of the information. Encryption keys for a Data Protector cell are stored in a central keystore on the Cell Manager.

encryption KeyID-StoreID Combined identifier used by the Data Protector Key Management Server to identify and administer encryption keys used by Data Protector. `KeyID` identifies the key within the keystore. `StoreID` identifies the keystore on the Cell Manager. If Data Protector has been upgraded from an earlier version with encryption functionality, there may several `StoreID`s used on the same Cell Manager.

enhanced incremental backup Conventional incremental backup backs up files that have changed since a previous backup, but has certain limitations in detection of changes. Unlike conventional incremental backup, enhanced incremental backup reliably detects and backs up also renamed and moved files, as well as files with changes in attributes.

enterprise backup environment Several cells can be grouped together and managed from a central cell. The enterprise backup environment includes all clients located in several Data Protector cells which are managed and administered from a central cell using the Manager-of-Managers concept.
See also MoM.

Event Log (Data Protector Event Log) A central repository of all Data Protector-related notifications. By default, all notifications are sent to the Event Log. The events are logged on the Cell Manager into the file `Data_Protector_program_data\log\server\Ob2EventLog.txt` (Windows Server 2008), `Data_Protector_home\log\server\Ob2EventLog.txt` (other Windows systems), or `/var/opt/omni/server/log/Ob2EventLog.txt` (UNIX systems). The Event Log is accessible only to users of the Data Protector `Admin` user group and to users who are granted the Data Protector `Reporting and notifications` user rights. You can view or delete all events in the Event Log.

Event Logs *(Windows specific term)* Files in which Windows logs all events, such as the starting or stopping of services and the logging on and off of users. Data Protector can back up Windows Event Logs as part of the Windows configuration backup.

Exchange Replication Service *(Microsoft Exchange Server specific term)* The Microsoft Exchange Server service that represents storage groups that were replicated using either local continuous replication (LCR) or cluster continuous replication (CCR) technology.
See also cluster continuous replication and local continuous replication.

exchanger Also referred to as SCSI Exchanger.
See also library.

exporting media A process that removes all data about backup sessions, such as systems, objects, and file names, which reside on the media from the IDB. Information about the media and their relation to a pool is also removed from the IDB. The data on the media remains unchanged.
See also importing media.

Extensible Storage Engine (ESE) *(Microsoft Exchange Server specific term)* A database technology used as a storage system for information exchange in Microsoft Exchange Server.

F

failover Transferring of the most important cluster data, called group (on Windows) or package (on UNIX) from one cluster node to another. A failover can occur mostly because of software or hardware failures or maintenance on the primary node.

failover *(HP P6000 EVA Disk Array Family specific term)* An operation that reverses the roles of source and destination in HP Continuous Access + Business Copy (CA+BC) P6000 EVA configurations.
See also HP Continuous Access + Business Copy (CA+BC) P6000 EVA.

FC bridge *See* Fibre Channel bridge.

Fibre Channel An ANSI standard for high-speed computer interconnection. Using either optical or copper cables, it allows the high speed bi-directional transmission of large data files and can be deployed between sites kilometers apart. Fibre Channel connects nodes using three physical topologies: point-to-point, loop, and switched.

Fibre Channel bridge A Fibre Channel bridge or multiplexer provides the ability to migrate existing parallel SCSI devices, like RAID arrays, solid state disks (SSD), and tape libraries to a Fibre Channel environment. On one side of the bridge or multiplexer there is a Fibre Channel interface while on the other side there are parallel SCSI ports. The bridge or multiplexer enables SCSI packets to be moved between the Fibre Channel and parallel SCSI devices.

file depot A file containing the data from a backup to a file library device.

file jukebox device A device residing on disk consisting of multiple slots used to store file media.

file library device A device which resides on a disk emulating a library with several media, hence containing multiple files, referred to as file depots.

File Replication Service (FRS) A Windows service that replicates the domain controller store logon scripts and group policies. FRS also enables replication of Distributed File System (DFS) shares between systems and allows any server to perform replication activity.

file tree walk *(Windows specific term)* The process of traversing a filesystem to determine which objects have been created, modified, or deleted.

file version The same file can be backed up multiple times in case of full backups and incremental backups (if the file changed). If the log level ALL is selected for backup, Data Protector retains one entry in the IDB for the filename itself and one for each version (date/time) of the file.

filesystem The organization of files on a hard disk. A filesystem is backed up so that the file attributes and the file contents are stored on the backup media.

first-level mirror *(HP P9000 XP Disk Array Family specific term)* A mirror of an internal disk (LDEV) of a disk array of the HP P9000 XP Disk Array Family which can be further mirrored itself, producing second-level mirrors. For Data Protector zero downtime backup and instant recovery purposes, only first-level mirrors can be used.
See also primary volume and mirror unit (MU) number.

flash recovery area *(Oracle specific term)* A directory, filesystem, or Automatic Storage Management (ASM) disk group managed by Oracle that serves as a centralized storage area for files related to backup, restore, and database recovery (recovery files).
See also recovery files.

fnames.dat The `fnames.dat` files of the IDB contain information on the names of the backed up files. Typically, these files occupy about 20% of the IDB, if filenames are stored.

formatting A process that erases any data contained on a medium and prepares it for use with Data Protector. Information about media (medium ID, description, and location) is saved in the IDB as well as on the respective media (media header). Data Protector media with protected data are not formatted until the protection expires or the media are unprotected/recycled.

free pool An auxiliary source of media for use by media pools when they run out of media. The media pools must be configured to use free pools.

full backup A backup in which all selected objects are backed up, whether or not they have been recently modified.
See also backup types.

full database backup A backup of all data in a database, not only the data that has been changed after the last (full or incremental) backup of the database. A full database backup does not depend on any other backup.

full mailbox backup A full mailbox backup is a backup of the entire mailbox content.

full ZDB A ZDB-to-tape or ZDB-to-disk+tape session in which all selected objects are streamed to tape, even if there are no changes from the previous backup.
See also incremental ZDB.

G

global options file A file that allows you to customize Data Protector. It explains the global options, which cover various aspects of Data Protector, typically time-outs and limits, and affect the entire Data Protector cell. The file is located on the Cell Manager in the directory *`Data_Protector_program_data`*`\Config\Server\Options` (Windows Server 2008), *`Data_Protector_home`*`\Config\Server\Options` (other Windows systems), or `/etc/opt/omni/server/options` (HP-UX, Solaris, and Linux systems).

group *(Microsoft Cluster Server specific term)* A collection of resources (for example disk volumes, application services, IP names, and addresses) that are needed to run a specific cluster-aware applications.

GUI A graphical user interface provided by Data Protector for easy access to all configuration, administration, and operation tasks. Besides the original Data Protector GUI that runs on Windows, Data Protector also provides a Java-based graphical user interface with the same look and feel, which runs on numerous platforms.

H

hard recovery *(Microsoft Exchange Server specific term)* A Microsoft Exchange Server database recovery that is performed after a restore by the database engine, using transaction log files.

heartbeat A cluster data set with a time stamp carrying information about the operational status of a particular cluster node. This data set or packet is distributed among all cluster nodes.

Hierarchical Storage Management (HSM) A method for optimizing the use of expensive hard disk storage by migrating less frequently used data to less expensive optical platters. When needed, the data is migrated back to hard disk storage. This balances the need for fast retrieval from hard disk with the lower cost of optical platters.

Holidays file A file that contains information about holidays. You can set different holidays by editing the Holidays file on the Cell Manager in the directory `Data_Protector_program_data\Config\Server\holidays` (Windows Server 2008), `Data_Protector_home\Config\Server\holidays` (other Windows systems), or `/etc/opt/omni/server/Holidays` (UNIX systems).

hosting system A working Data Protector client used for Disk Delivery Disaster Recovery with a Data Protector Disk Agent installed.

HP Business Copy (BC) P6000 EVA *(HP P6000 EVA Disk Array Family specific term)* A local replication software solution that enables creation of point-in-time copies (replicas) of the source volumes using the snapshot and clone capabilities of the P6000 EVA firmware.
See also replica, source volume, snapshot, and HP Continuous Access + Business Copy (CA+BC) P6000 EVA.

HP Business Copy (BC) P9000 XP *(HP P9000 XP Disk Array Family specific term)* An HP P9000 XP Disk Array Family configuration that enables creation and maintenance of internal copies of LDEVs for various purposes, such as data duplication and backup. The copies (secondary volumes or S-VOLs) can be separated from the primary volumes (P-VOLs) and connected to a different system. For Data Protector zero downtime backup purposes, P-VOLs should be available to the application system, and one of the S-VOL sets should be available to the backup system.
See also LDEV, HP Continuous Access (CA) P9000 XP, Main Control Unit, application system, and backup system.

HP Command View (CV) EVA *(HP P6000 EVA Disk Array Family specific term)* The user interface that enables you to configure, manage, and monitor your P6000 EVA storage system. It is used to perform various storage management tasks, for example, creating virtual disk families, managing storage system hardware, and creating snapshots, snapclones, and mirrorclones of virtual disks. The HP Command View EVA software runs on the HP Storage Management Appliance, and is accessed by a Web browser.
See also HP StorageWorks P6000 EVA SMI-S Agent and HP StorageWorks SMI-S P6000 EVA Array provider.

HP Continuous Access (CA) P9000 XP *(HP P9000 XP Disk Array Family specific term)* An HP P9000 XP Disk Array Family configuration that enables creation and maintenance of remote copies of LDEVs for purposes such as data duplication, backup, and disaster recovery. HP CA P9000 XP operations involve main (primary) disk array units and remote (secondary) disk array units. The main disk array units are connected to the application system and contain primary volumes (P-VOLs), which store original data. The remote disk array units are connected to the backup system and contain secondary volumes (S-VOLs).
See also HP Business Copy (BC) P9000 XP, Main Control Unit, and LDEV.

HP Continuous Access + Business Copy (CA+BC) P6000 EVA *(HP P6000 EVA Disk Array Family specific term)* An HP P6000 EVA Disk Array Family configuration that enables creation and maintenance of copies (replicas) of the source volumes on a remote P6000 EVA, and later use of these copies as the source for local replication on this remote array.
See also HP Business Copy (BC) P6000 EVA, replica, and source volume.

HP SMI-S P6000 EVA Array provider An interface used for controlling HP P6000 EVA Disk Array Family. SMI-S P6000 EVA Array provider runs as a separate service on the HP Storage Management Appliance system and acts as a gateway between incoming requests and HP Command View EVA. With the Data Protector HP P6000 EVA Disk Array Family integration, SMI-S P6000 EVA Array provider accepts standardized requests from the P6000 EVA SMI-S Agent, communicates with HP Command View EVA for information or method invocation, and returns standardized responses.
See also HP StorageWorks P6000 EVA SMI-S Agent and HP Command View (CV) EVA.

HP StorageWorks P6000 EVA SMI-S Agent A Data Protector software module that executes all tasks required for the HP P6000 EVA Disk Array Family integration. With the P6000 EVA SMI-S Agent, the control over the array is established through HP SMI-S P6000 EVA Array provider, which directs communication between incoming requests and HP CV EVA.

See also HP Command View (CV) EVA and HP SMI-S P6000 EVA Array provider.

HP StorageWorks P9000 XP Agent
A Data Protector component that executes all tasks needed by the Data Protector HP P9000 XP Disk Array Family integration. It uses RAID Manager Library for communication with a P9000 XP Array storage system.
See also RAID Manager Library.

HP Operations Manager
HP Operations Manager provides powerful capabilities for operations management of a large number of systems and applications in a network. Data Protector provides an integration into this management product. This integration is implemented as a SMART Plug-In for HP Operations Manager management servers on Windows, HP-UX, Solaris, and Linux. Earlier versions of HP Operations Manager were called IT/Operations, Operations Center, Vantage Point Operations, and OpenView Operations.

HP Operations Manager SMART Plug-In (SPI)
A fully integrated, out-of-the-box solution which "plugs into" HP Operations Manager, extending the managed domain. Through the Data Protector integration, which is implemented as an HP Operations Manager SMART Plug-In, a user can have an arbitrary number of Data Protector Cell Managers monitored as an extension to HP Operations Manager.

I

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. IDB is an embedded database located on the Cell Manager and keeps information regarding which data was backed up, to which media it was backed up, how backup, restore, and other sessions were run, which devices, libraries, and disk arrays are configured, and so on.

IDB recovery file
An IDB file (obrindex.dat) with information about IDB backups, media, and devices used for the backup. This information can significantly simplify IDB recovery. It is recommended to relocate the file, together with IDB transaction logs, to a separate physical disk from other IDB directories, and, additionally, to make an additional copy of the file.

importing media
A process that re-reads all data about backup sessions which are on the medium back into the IDB. This then allows for fast and convenient access to the data on the media.
See also exporting media.

incremental (re)-establish
(EMC Symmetrix specific term) A BCV or SRDF control operation. In BCV control operations, an incremental establish causes the BCV device to be synchronized incrementally and to function as an EMC Symmetrix mirrored medium. The EMC Symmetrix devices must have been previously paired. In SRDF control operations, an incremental establish causes the target (R2) device to be synchronized incrementally and to function as an EMC Symmetrix mirrored medium. The EMC Symmetrix devices must have been previously paired.

incremental backup
A backup that selects only files that have changed since a previous backup. Several levels of incremental backup are available, which enables detailed control of restore chain length.
See also backup types.

incremental backup
(Microsoft Exchange Server specific term) A backup of the Microsoft Exchange Server data that has changed since the last full or incremental backup. With the incremental backup, only the transaction log files are backed up.
See also backup types.

incremental mailbox backup
An incremental mailbox backup backs up all the changes made to the mailbox after the last backup of any type.

incremental restore
(EMC Symmetrix specific term) A BCV or SRDF control operation. In BCV control operations, an incremental restore reassigns a BCV device as the next available mirror of the standard device in the pair. However, the standard devices are updated with only the data that was written to the BCV device during the time of the original pair split, and the data that was written to the standard device during the split is overwritten with data from the BCV mirror. In SRDF control operations, an incremental restore reassigns a target (R2) device as the next available mirror of the source (R1) device in the pair. However, the source (R1) devices are updated with only the data that was written to the target (R2) device during the time of the original pair split, and the

data that was written to the source (R1) device during the split is overwritten with data from the target (R2) mirror.

incremental ZDB A filesystem ZDB-to-tape or ZDB-to-disk+tape session in which only changes from the last protected full or incremental backup are streamed to tape.
See also full ZDB.

incremental1 mailbox backup An incremental1 mailbox backup backs up all the changes made to the mailbox after the last full backup.

Inet A process that runs on each UNIX system or service that runs on each Windows system in the Data Protector cell. It is responsible for communication between systems in the cell and for starting other processes needed for backup and restore. The Inet service is started as soon as Data Protector is installed on a system. The Inet process is started by the inetd daemon.

Information Store *(Microsoft Exchange Server specific term)* The Microsoft Exchange Server service that is responsible for storage management. Information Store in Microsoft Exchange Server manages two kinds of stores: mailbox stores and public folder stores. A mailbox store consists of mailboxes that belong to individual users. A public folder store contains public folders and messages that are shared among several users.
See also Key Management Service and Site Replication Service.

Informix Server *(Informix Server specific term)* Refers to Informix Dynamic Server.

initializing *See* formatting.

Installation Server A computer system that holds a repository of the Data Protector installation packages for a specific architecture. The Installation Server is used for remote installation of Data Protector clients. In mixed environments at least two Installation Servers are needed: one for UNIX systems and one for Windows systems.

instant recovery *(ZDB specific term)* A process in which a replica, produced by a ZDB-to-disk or a ZDB-to-disk+tape session, is used to restore the contents of the source volumes to their states at the time at which the replica was created, avoiding the need to perform a restore from tape. Depending on the application or database concerned, this may be all that is required, or other steps, such as the application of transaction log files, may be required for full recovery.
See also replica, zero downtime backup (ZDB), ZDB to disk, and ZDB to disk+tape.

integration object A backup object of a Data Protector integration, such as Oracle or SAP DB.

Internet Information Services (IIS) *(Windows specific term)* Microsoft Internet Information Services is a network file and application server that supports multiple protocols. Primarily, IIS transmits information in Hypertext Markup Language (HTML) pages by using the Hypertext Transport Protocol (HTTP).

ISQL *(Sybase specific term)* A Sybase utility used to perform system administration tasks on Sybase SQL Server.

J

Java GUI Client The Java GUI Client is a component of the Java GUI that contains only user interface related functionalities (the Cell Manager graphical user interface and the Manager-of-Managers (MoM) graphical user interface) and requires connection to the Java GUI Server to function.

Java GUI Server The Java GUI Server is a component of the Java GUI that is installed on the Data Protector Cell Manager system. The Java GUI Server receives requests from the Java GUI Client, processes them and then sends the responses back to the Java GUI Client. The communication is done through Hypertext Transfer Protocol (HTTP) on port 5556.

jukebox *See* library.

jukebox device A device consisting of multiple slots used to store either optical or file media. When being used to store file media, the jukebox device is known as the “file jukebox device”.

K

Key Management Service *(Microsoft Exchange Server specific term)* The Microsoft Exchange Server service that provides encryption functionality for enhanced security.
See also Information Store and Site Replication Service.

keychain A tool that eliminates the supply of a passphrase manually when decrypting the private key. It needs to be installed and configured on the Installation Server if you perform remote installation using secure shell.

keystore All encryption keys are centrally stored in the keystore on the Cell Manager and administered by the Key Management Server (KMS).

KMS Key Management Server (KMS) is a centralized service that runs on the Cell Manager and provides key management for the Data Protector encryption functionality. The service is started as soon as Data Protector is installed on the Cell Manager.

L

LBO *(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.

LDEV *(HP P9000 XP Disk Array Family specific term)* A logical partition of a physical disk of a disk array of the HP P9000 XP Disk Array Family. An LDEV is the entity that can be replicated using the split-mirror or snapshot functionality of such disk array.
See also HP Business Copy (BC) P9000 XP, HP Continuous Access (CA) P9000 XP, and replica.

library Also called autochanger, jukebox, autoloader, or exchanger. A library contains media in repository slots. Each slot holds one medium (for example, DDS/DAT). Media are moved between slots and drives by a robotic mechanism, allowing random access to media. The library can contain multiple drives.

lights-out operation or unattended operation A backup or restore operation that takes place outside of normal business hours without an operator. This implies that no operator personnel is present to work with the backup application or service mount requests, for example.

LISTENER.ORA *(Oracle specific term)* An Oracle configuration file that describes one or more Transparent Network Substrate (TNS) listeners on a server.

load balancing By default, Data Protector automatically balances the usage of devices selected for backup, so that they are used evenly. Load balancing optimizes the device usage by balancing the number of objects written to each device. Since load balancing is done automatically during backup time, you do not need to manage how the data is actually backed up. You just specify the devices to be used. If you do not want to use load balancing, you can select which device will be used for each object in the backup specification. Data Protector will access the devices in the specified order.

local and remote recovery Remote recovery is performed if all Media Agent hosts specified in the SRD file are accessible. If any of them fails, the disaster recovery process fails over to the local mode. This means that the target system is searched for locally attached devices. If only one device is found, it is automatically used. Otherwise, Data Protector prompts you to select the device, which will be used for restore.

local continuous replication *(Microsoft Exchange Server specific term)* Local continuous replication (LCR) is a single-server solution that creates and maintains an exact copy (LCR copy) of a storage group. An LCR copy is located on the same server as the original storage group. When an LCR copy is created, it is kept up to date through change propagation (log replay) technology. The replication feature in LCR guarantees that logs that have not been replicated are not deleted. The implication of this behavior is that running backups in a mode that deletes logs may not actually free space if replication is sufficiently far behind in its log copying.

An LCR copy is used for disaster recovery because you can switch to the LCR copy in a few seconds. If an LCR copy is used for backup and if it is located on a different disk than the original data, then the I/O load on a production database is minimal.

A replicated storage group is represented as a new instance of Exchange writer called Exchange Replication Service and can be backed up (using VSS) as a normal storage group.

See also cluster continuous replication and Exchange Replication Service.

lock name You can configure the same physical device several times with different characteristics, by using different device names. The lock name is a user specified string that is used for locking all such device configurations to prevent collision if several such devices (device names) are used concurrently. Use an identical lock name for all device definitions which use the same physical device.

log_full shell script

(Informix Server UNIX specific term) A script provided by ON-Bar that you can use to start backing up logical log files when Informix Server issues a logfull event alarm. The Informix Server `ALARMPROGRAM` configuration parameter defaults to the *`INFORMIXDIR`*`/etc/log_full.sh`, where *`INFORMIXDIR`* is the Informix Server home directory. If you do not want logical logs to be backed up continuously, set the `ALARMPROGRAM` configuration parameter to *`INFORMIXDIR`*`/etc/no_log.sh`.

logging level

The logging level determines the amount of details on files and directories written to the IDB during backup, object copying, or object consolidation. You can always restore your data, regardless of the logging level used during backup. Data Protector provides four logging levels: Log All, Log Directories, Log Files, and No Log. The different logging level settings influence the IDB growth, backup speed, and the convenience of browsing data for restore.

logical-log files

This applies to online database backup. Logical-log files are files in which modified data is first stored before being flushed to disk. In the event of a failure, these logical-log files are used to roll forward all transactions that have been committed as well as roll back any transactions that have not been committed.

login ID

(Microsoft SQL Server specific term) The name a user uses to log on to Microsoft SQL Server. A login ID is valid if Microsoft SQL Server has an entry for that user in the system table syslogin.

login information to the Oracle Target Database

(Oracle and SAP R/3 specific term) The format of the login information is *`user_name/password@service`*, where:

- *`user_name`* is the name by which a user is known to Oracle Server and to other users. Every user name is associated with a password and both have to be entered to connect to an Oracle Target Database. This user must have Oracle SYSDBA or SYSOPER rights.
- *`password`* must be the same as the password specified in the Oracle password file (`orapwd`), which is used for authentication of users performing database administration.
- *`service`* is the name used to identify an SQL*Net server process for the target database.

login information to the Recovery Catalog Database

(Oracle specific term) The format of the login information to the Recovery (Oracle) Catalog Database is *`user_name/password@service`*, where the description of the user name, password, and service name is the same as in the Oracle SQL*Net V2 login information to the Oracle target database. In this case, *`service`* is the name of the service to the Recovery Catalog Database, not the Oracle target database.

Note that the Oracle user specified here must be the owner of the Oracle Recovery Catalog.

Lotus C API

(Lotus Domino Server specific term) An interface for the exchange of backup and recovery information between Lotus Domino Server and a backup solution, like Data Protector.

LVM

A Logical Volume Manager is a subsystem for structuring and mapping physical disk space to logical volumes on UNIX systems. An LVM system consists of several volume groups, where each volume group has several volumes.

M

Magic Packet

See Wake ONLAN.

mailbox

(Microsoft Exchange Server specific term) The location to which e-mail is delivered, which is set up by the administrator for each user. If a set of personal folders is designated as the e-mail delivery location, e-mail is routed from the mailbox to this location.

mailbox store

(Microsoft Exchange Server specific term) A part of the Information Store that maintains information in user mailboxes. A mailbox store consists of a binary rich-text `.edb` file and a streaming native internet content `.stm` file.

Main Control Unit (MCU)

(HP P9000 XP Disk Array Family specific term) An HP P9000 XP Disk Array Family unit that contains primary volumes (P-VOLs) for the HP CA P9000 XP or HP CA+BC P9000 XP configuration and acts as a master device.
See also HP Business Copy (BC) P9000 XP, HP Continuous Access (CA) P9000 XP, and LDEV.

make_net_recovery `make_net_recovery` is an Ignite-UX command, which allows you to create a recovery archive over the network onto the Ignite-UX server or any other specified system. The target system can be recovered across subnets after booting either from a bootable tape created by the Ignite-UX `make_boot_tape` command or the system boots directly from the Ignite-UX server. Booting directly from the Ignite-UX server can be automated with the Ignite-UX `bootsys` command or interactively specified on the boot console.

make_tape_recovery `make_tape_recovery` is a command on Ignite-UX which creates a bootable recovery (installation) tape, customized for your system and enables you unattended disaster recovery by connecting the backup device directly to the target system and booting the target system from the bootable recovery tape. The backup device has to be locally connected to the client during the creation of the archive and recovery of the client.

Manager-of-Managers (MoM) *See* MoM.

MAPI *(Microsoft Exchange Server specific term)* The MAPI (Messaging Application Programming Interface) is the programming interface that lets applications and messaging clients interact with messaging and information systems.

MCU *See* Main Control Unit (MCU).

Media Agent A process that controls reading from and writing to a device, which reads from or writes to a medium (typically a tape). During a backup session, a Media Agent receives data from the Disk Agent and sends it to the device for writing it to the medium. During a restore or object verification session, a Media Agent locates data on the backup medium and sends it to the Disk Agent for processing. For a restore session, the Disk Agent then writes the data to the disk. A Media Agent also manages the robotics control of a library.

media allocation policy Determines in which sequence media are used for backup. The Strict allocation policy directs Data Protector to prompt for a specific medium. The Loose policy directs Data Protector to prompt for any suitable medium. The Formatted First policy directs Data Protector to give preference to unknown media, even if unprotected media are available in the library.

media condition The quality of a medium as derived from the media condition factors. Heavy usage and age result in an increased number of read and write errors with tape media. Media need to be replaced when they are marked as POOR.

media condition factors The user-assigned age threshold and overwrite threshold used to determine the state of a medium.

media label A user-defined identifier used to describe a medium.

media location A user-defined physical location of a medium, such as "building 4" or "off-site storage".

media management session A session performing some action on a medium, such as initializing, scanning the content, verifying data on a medium, or copying a medium.

media pool A set of media of the same type (such as DDS) used and tracked as a group. Media are formatted and assigned to a media pool.

media set The result of a backup session is data backed up on a group of media called media set. Depending on the media usage policy, several sessions can share the same media.

media type The physical type of media, such as DDS or DLT.

media usage policy The media usage policy controls how new backups are added to the already used media. It can be `Appendable`, `Non-Appendable`, or `Appendable for incrementals only`.

medium ID A unique identifier assigned to a medium by Data Protector.

merging This defines one mode to resolve file conflicts during restore. If the file to be restored already exists at the destination, the one with the more recent modification date is kept. Files not present on the disk are always restored.
See also overwrite.

Microsoft Exchange Server A “client-server” messaging and a workgroup system that offers a transparent connection to many different communication systems. It provides users with an electronic mail system, individual and group scheduling, online forms, and workflow automation tools. It provides a developer with a platform on which to build custom information-sharing and messaging-service applications.

Microsoft Management Console (MMC) *(Windows specific term)* An administration model for Windows-based environments. It provides a simple, consistent, and integrated administration user interface allowing management of many applications through the same GUI, provided that the applications adhere to the MMC model.

Microsoft SQL Server A database management system designed to meet the requirements of distributed "client-server" computing.

Microsoft Volume Shadow Copy Service (VSS) A software service that provides a unified communication interface to coordinate backup and restore of a VSS-aware application regardless of its specific features. This service collaborates with the backup application, writers, shadow copy providers, and the operating system kernel to implement the management of volume shadow copies and shadow copy sets.
See also shadow copy, shadow copy provider, replica, and writer.

mirror *(EMC Symmetrix and HP P9000 XP Disk Array Family specific term)* *See* target volume.

mirror rotation *(HP P9000 XP Disk Array Family specific term)* *See* replica set rotation.

mirror unit (MU) number *(HP P9000 XP Disk Array Family specific term)* A non-negative integer number that determines a secondary volume (S-VOL) of an internal disk (LDEV) located on a disk array of the HP P9000 XP Disk Array Family.
See also first-level mirror.

mirrorclone *(HP P6000 EVA Disk Array Family specific term)* A dynamic replica of a storage volume, which is kept updated with changes made to the original storage volume via a local replication link. Replication between the original storage volume and its mirrorclone can be suspended. For each storage volume, a single mirrorclone can be created on the disk array.

MMD The Media Management Daemon process (service) runs on the Data Protector Cell Manager and controls media management and device operations. The process is started when Data Protector is installed on the Cell Manager.

MMDB The Media Management Database (MMDB) is a part of the IDB that contains information about media, media pools, devices, libraries, library drives, and slots configured in the cell, as well as the Data Protector media used for backup. In an enterprise backup environment, this part of the database can be common to all cells.
See also CMMDB and CDB.

MoM Several cells can be grouped together and managed from a central cell. The management system of the central cell is the Manager-of-Managers (MoM). The cells are called MoM clients. The MoM enables you to configure and manage multiple cells from a central point.

mount point The access point in a directory structure for a disk or logical volume, for example`/opt` or `d:`. On UNIX, the mount points are displayed using the `bdf` or `df` command.

mount request A screen prompt that tells you to insert a specific medium into a device. Once you respond to the mount request by providing the required medium and confirm the mount request, the session continues.

MSM The Data Protector Media Session Manager, which runs on the Cell Manager and controls media sessions, such as copying media.

multisnapping *(HP P6000 EVA Disk Array Family specific term)* Simultaneous creation of target volumes so that the backup data is consistent not only on each individual target volume, but also across all the volumes that constitute a snapshot.
See also snapshot.

O

OBDR capable device A device that can emulate a CD-ROM drive loaded with a bootable disk and can thus be used as a backup or boot device for disaster recovery purposes.

obdrindex.dat *See* IDB recovery file.

object *See* backup object.

object consolidation The process of merging a restore chain of a backup object, consisting of a full backup and at least one incremental backup, into a new, consolidated version of this object. The process is a part of the synthetic backup procedure. The result is a synthetic full backup of the specified backup object.

object consolidation session A process that merges a restore chain of a backup object, consisting of a full backup and at least one incremental backup, into a new, consolidated version of this object.

object copy A copy of a specific object version that is created during an object copy session or a backup session with object mirroring.

object copy session A process that creates an additional copy of the backed up data on a different media set. During an object copy session, the selected backed up objects are copied from the source to the target media.

object copying The process of copying selected object versions to a specific media set. You can select object versions from one or several backup sessions to be copied.

object ID *(Windows specific term)* The object IDs (OIDs) enable access to NTFS 5 files no matter where in the system the files reside. Data Protector treats the OIDs as alternate streams of the files.

object mirror A copy of a backup object created using object mirroring. Object mirrors are often referred to as object copies.

object mirroring The process of writing the same data to several media sets during a backup session. Data Protector enables you to mirror all or some backup objects to one or more media sets.

object verification The process of verifying the data integrity of backup objects, from the Data Protector point of view, and the ability of Data Protector to deliver them to the required destination. The process can be used to provide a level of confidence in the ability to restore object versions created by backup, object copy, or object consolidation sessions.

object verification session A process that verifies the data integrity of specified backup objects or object versions and the ability of selected Data Protector network components to deliver them to a specified host. Object verification sessions can be run interactively, or as specified in automated post-backup, or scheduled specifications.

offline backup A backup during which an application database cannot be used by the application. In an offline backup session, the database is generally put into a quiescent state that allows use by the backup system, but not the application, for the time period of the data replication process. For instance, for backup to tape, until streaming of data to the tape is finished. Normal database operation is resumed before potential post-backup operations are started.
See also zero downtime backup (ZDB) and online backup.

offline recovery Offline recovery is performed if the Cell Manager is not accessible, for example, due to network problems. Only standalone and SCSI library devices can be used for offline recovery. Recovery of the Cell Manager is always offline.

offline redo log *See* archived redo log.

ON-Bar *(Informix Server specific term)* A backup and restore system for Informix Server. ON-Bar enables you to create a copy of your Informix Server data and later restore the data. The ON-Bar backup and restore system involves the following components:

- the `onbar` command
- Data Protector as the backup solution
- the XBSA interface
- ON-Bar catalog tables, which are used to back up dbobjects and track instances of dbobjects through multiple backups.

ONCONFIG *(Informix Server specific term)* An environment variable that specifies the name of the active ONCONFIG configuration file. If the ONCONFIG environment variable is not present, Informix Server uses the configuration values from the `onconfig` file in the directory `INFORMIXDIR\etc` (on Windows) or `INFORMIXDIR/etc/` (on UNIX).

online backup A backup performed while a database application remains available for use. The database is placed into a special backup mode of operation for the time period of the data replication process. For instance, for backup to tape, until streaming of data to tape is finished. During this period, the database is fully operational, but there may be a small performance impact and log files may grow very quickly. Normal database operation is resumed before potential post-backup operations are started.

In some cases, transaction logs may also have to be backed up to allow a consistent database to be restored.

See also zero downtime backup (ZDB) and offline backup.

online recovery Online recovery is performed when Cell Manager is accessible. In this case, most of the Data Protector] functionalities are available (Cell Manager runs the session, restore sessions are logged in the IDB, you can monitor the restore progress using the GUI, and so on).

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.

Oracle Data Guard *(Oracle specific term)* Oracle Data Guard is Oracle's primary disaster recovery solution. Oracle Data Guard is able to maintain up to nine standby databases, each of which is a real-time copy of the production (primary) database, to protect against corruptions, data failures, human errors, and disasters. If a failure occurs on the production (primary) database, then a failover to one of the standby databases which becomes the new primary database is possible. In addition, planned downtime for maintenance can be reduced because the production processing can be moved from the current primary database to a standby database and back quickly.

Oracle instance *(Oracle specific term)* Each installation of an Oracle database on one or more systems. One computer system can have several instances of a database running.

ORACLE_SID *(Oracle specific term)* A unique name for an Oracle Server instance. To switch among Oracle Servers, specify the desired *`ORACLE_SID`*. The *`ORACLE_SID`* is included in the CONNECT DATA parts of the connect descriptor in a `TNSNAMES.ORA` file and in the definition of the TNS listener in the `LISTENER.ORA` file.

original system The system configuration backed up by Data Protector before a computer disaster hits the system.

overwrite An option that defines one mode to resolve file conflicts during restore. All files are restored from a backup even if they are older than existing files.
See also merging.

ownership Backup ownership affects the ability of users to see and restore data. Each backup session and all the data backed up within it is assigned an owner. The owner can be the user that starts an interactive backup, the account under which the CRS process is running, or the user specified as the owner in the backup specification options.

If a user starts an existing backup specification without modifying it, the backup session is not considered as interactive.

If a modified backup specification is started by a user, the user is the owner unless the following is true:

- The user has the Switch Session Ownership user right.
- The backup session owner is explicitly defined in the backup specification, where the username, group or domain name, and the system name are specified.

If a backup is scheduled on a UNIX Cell Manager, the session owner is root:sys unless the above conditions are true.

If a backup is scheduled on a Windows Cell Manager, the session owner is the user specified during the installation, unless the above conditions are true.

When copying or consolidating objects, by default the owner is the user who starts the operation, unless a different owner is specified in the copy or consolidation specification.

P

P1S file P1S file contains information on how to format and partition all disks installed in the system during Enhanced Automated Disaster Recovery (EADR). It is created during full backup and is saved on backup medium and on Cell Manager into the directory `Data_Protector_program_data\Config\Server\dr\p1s` (Windows Server 2008), `Data_Protector_home\Config\Server\dr\p1s` (other Windows systems), or `/etc/opt/omni/server/dr/p1s` (UNIX systems) with the filename `recovery.p1s`.

package *(MC/ServiceGuard and Veritas Cluster specific term)* A collection of resources (for example volume groups, application services, IP names, and addresses) that are needed to run a specific cluster-aware application.

pair status *(HP P9000 XP Disk Array Family specific term)* The status of a disk pair (secondary volume and its corresponding primary volume) of a disk array of the HP P9000 XP Disk Array Family. Depending on the circumstances, the paired disks can be in various states. The following states are particularly important for the operation of the Data Protector HP StorageWorks P9000 XP Agent:

- PAIR – The secondary volume is prepared for zero downtime backup. If it is a mirror, it is completely synchronized, and if it is a volume to be used for snapshot storage, it is empty.
- SUSPENDED – The link between the disks is suspended. However, the pair relationship is still maintained, and the secondary disk can be prepared for zero downtime backup again at a later time.
- COPY – The disk pair is currently busy and making a transition into the PAIR state. If the secondary volume is a mirror, it is re-synchronizing with the primary volume, and if it is a volume to be used for snapshot storage, its contents are getting cleared.

parallel restore Restoring backed up data to multiple disks at the same time (that is, in parallel) by running multiple Disk Agents, that receive data from one Media Agent. For the parallel restore to work, select data that is located on different disks or logical volumes and during backup, the data from the different objects must have been sent to the same device using a concurrency of 2 or more. During a parallel restore, the data for multiple objects selected for restore is read from media at the same time, thereby improving performance.

parallelism The concept of reading multiple data streams from an online database.

phase 0 of disaster recovery Preparation for disaster recovery - the prerequisite condition for a successful disaster recovery.

phase 1 of disaster recovery Installation and configuration of DR OS, establishing previous storage structure.

phase 2 of disaster recovery Restoration of operating system (with all the configuration information that defines the environment) and Data Protector.

phase 3 of disaster recovery Restoration of user and application data.

physical device A physical unit that contains either a drive or a more complex unit such as a library.

post-exec A backup option that executes a command or script after the backup of an object or after the entire session completes. Post-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows 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.

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.

prealloc list A subset of media in a media pool that specifies the order in which media are used for backup.

primary volume (P-VOL) *(HP P9000 XP Disk Array Family specific term)* An internal disk (LDEV) of a disk array of the HP P9000 XP Disk Array Family for which a secondary volume (S-VOL), either its mirror or a volume

to be used for its snapshot storage, exists. In the HP CA P9000 XP and HP CA+BC P9000 XP configurations, primary volumes are located in the Main Control Unit (MCU).
See also secondary volume (S-VOL) and Main Control Unit (MCU).

protection *See* data protection and also catalog protection.

public folder store *(Microsoft Exchange Server specific term)* The part of the Information Store that maintains information in public folders. A public folder store consists of a binary rich-text `.edb` file and a streaming native internet content `.stm` file.

public/private backed up data When configuring a backup, you can select whether the backed up data will be:

- public, that is visible (and accessible for restore) to all Data Protector users
- private, that is, visible (and accessible for restore) only to the owner of the backup and administrators

R

RAID Redundant Array of Independent Disks.

RAID Manager Library *(HP P9000 XP Disk Array Family specific term)* A software library that is used for accessing the configuration, status, and performance measurement data of a P9000 XP Array storage system, and for invoking operations on the disk array. It translates function calls into sequences of low-level SCSI commands.
See also HP StorageWorks P9000 XP Agent.

RAID Manager P9000 XP *(HP P9000 XP Disk Array Family specific term)* A software application that provides a command-line interface to disk arrays of the HP P9000 XP Disk Array Family. It offers an extensive set of commands for reporting and controlling the status of a P9000 XP Array storage system, and for performing various operations on the disk array.

rawdisk backup *See* disk image backup.

RCU *See* Remote Control Unit (RCU).

RDBMS Relational Database Management System.

RDF1/RDF2 *(EMC Symmetrix specific term)* A type of SRDF device group. Only RDF devices can be assigned to an RDF group. An RDF1 group type contains source (R1) devices and an RDF2 group type contains target (R2) devices.

RDS The Raima Database Server process (service) runs on the Data Protector Cell Manager and manages the IDB. The process is started when Data Protector is installed on the Cell Manager.

Recovery Catalog *(Oracle specific term)* A set of Oracle tables and views that are used by Recovery Manager to store information about Oracle databases. This information is used by Recovery Manager to manage the backup, restore, and recovery of Oracle databases. The recovery catalog contains information about:

- The physical schema of the Oracle target database
- Data file and archived log backup sets
- Data file copies
- Archived Redo Logs
- Stored scripts

Recovery Catalog Database *(Oracle specific term)* An Oracle database that contains a recovery catalog schema. You should not store the recovery catalog in your target database.

recovery files *(Oracle specific term)* Recovery files are Oracle specific files that reside in the flash recovery area: the current control file, online redo logs, archived redo logs, flashback logs, control file autobackups, datafile copies, and backup pieces.
See also flash recovery area.

Recovery Manager (RMAN) *(Oracle specific term)* An Oracle command-line interface that directs an Oracle Server process to back up, restore, or recover the database it is connected to. RMAN uses either the recovery catalog or the control file to store information about backups. This information can be used later in restore sessions.

RecoveryInfo When backing up Windows configuration files, Data Protector collects the information about the current system configuration (information on disk layout, volume, and network configuration). This information is needed for disaster recovery.

recycle or unprotect A process that removes the data protection from all backed up data on a medium, allowing Data Protector to overwrite it during one of the next backups. Data that belongs to the same session(s) but resides on other media is also unprotected. Recycling does not actually alter the data on the medium.

redo log *(Oracle specific term)* Every Oracle database has a set of two or more redo log files. The set of redo log files for the database is known as the database's redo log. Oracle uses the redo log to record all changes made to data.

Remote Control Unit (RCU) *(HP P9000 XP Disk Array Family specific term)* An HP P9000 XP Disk Array Family unit that acts as a slave device to the Main Control Unit (MCU) in the HP CA P9000 XP or HP CA+BC P9000 XP configuration. In bidirectional configurations, the RCU can also act as an MCU.

Removable Storage Management Database *(Windows specific term)* A Windows service used for managing removable media (such as tapes and disks) and storage devices (libraries). Removable Storage allows applications to access and share the same media resources.

reparse point *(Windows specific term)* A system-controlled attribute that can be associated with any directory or file. The value of a reparse attribute can have user-controlled data. The format of the data is understood by the application that stored the data and a filesystem filter that was installed to interpret the data and process such files. Whenever the filesystem encounters a file with a reparse point, it attempts to find the filesystem filter associated with the data format.

replica *(ZDB specific term)* An image, at a particular point in time, of the data in source volumes that contain user-specified backup objects. Depending on the hardware or software with which it is created, the image may be an independent exact duplicate (clone) of the storage blocks at the physical disk level (for example, a split mirror or snapclone), or a virtual copy (for example, a snapshot). From perspective of a basic operating system, the complete physical disk containing backup objects is replicated. However, if a volume manager is used on UNIX, the whole volume or disk group containing a backup object (logical volume) is replicated. If partitions are used on Windows, the whole physical volume containing the selected partition is replicated.
See also snapshot, snapshot creation, split mirror, and split mirror creation.

replica set *(ZDB specific term)* A group of replicas, all created using the same backup specification.
See also replica and replica set rotation.

replica set rotation *(ZDB specific term)* The use of a replica set for regular backup production: Each time the same backup specification requiring the use of a replica set is run, a new replica is created and added to the set, until the maximum number for the set is reached. After that, the oldest replica in the set is replaced, maintaining the maximum number of replicas in the set.
See also replica and replica set.

restore chain All backups that are necessary for a restore of a backup object to a certain point in time. A restore chain consists of a full backup of the object and any number of related incremental backups.

restore session A process that copies data from backup media to a client.

resync mode *(HP P9000 XP Disk Array Family VSS provider specific term)* One of two P9000 XP Array VSS hardware provider operation modes. When the P9000 XP Array provider is in the resync mode, the source volume (P-VOL) and its replica (S-VOL) are in the suspended mirror relationship after a backup. The maximum number of replicas (S-VOLs per a P-VOL) rotated is three provided that MU range is 0-2 or 0, 1, 2. Restore from a backup in such a configuration is possible only by re-synchronization of an S-VOL with its P-VOL.
See also VSS compliant mode, source volume, primary volume (P-VOL), replica, secondary volume (S-VOL), mirror unit (MU) number, and replica set rotation.

RMAN *(Oracle specific term)* *See* Recovery Manager.

RSM The Data Protector Restore Session Manager controls restore and object verification sessions. This process always runs on the Cell Manager system.

RSM *(Windows specific term)* Removable Storage Manager (RSM) includes a media management service that facilitates communication among applications, robotic changers, and media libraries. It enables multiple applications to share local robotic media libraries and tape or disk drives and to manage removable media.

S

SAPDBA *(SAP R/3 specific term)* An SAP R/3 user interface that integrates the BRBACKUP, BRARCHIVE, and BRRESTORE tools.

scanning A function which identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library). It is useful to perform a scan and check the actual media in the device if someone has manually manipulated media without using Data Protector to eject or enter, for example.

Scheduler A function that controls when and how often automatic backups occur. By setting up a schedule, you automate the start of backups.

secondary volume (S-VOL) *(HP P9000 XP Disk Array Family specific term)* An internal disk (LDEV) of a disk array of the HP P9000 XP Disk Array Family which is paired with another LDEV: a primary volume (P-VOL). It can act as a mirror of the P-VOL or as a volume to be used for the P-VOL's snapshot storage. An S-VOL is assigned a SCSI address different from the one used for the P-VOL. In an HP CA P9000 XP configuration, the S-VOLs acting as mirrors can be used as failover devices in a MetroCluster configuration.
See also primary volume (P-VOL) and Main Control Unit (MCU).

session *See* backup session, media management session, and restore session.

session ID An identifier of a backup, restore, object copy, object consolidation, object verification, or media management session, consisting of the date when the session ran and a unique number.

session key This environment variable for the pre-exec and post-exec script is a Data Protector unique identification of any session, including preview sessions. The session key is not recorded in the database, and it is used for specifying options for the `omnimnt`, `omnistat`, and `omniabort` commands.

shadow copy *(Microsoft VSS specific term)* A volume that represents a duplicate of the original volume at a certain point in time. The data is then backed up from the shadow copy and not from the original volume. The original volume continues to change as the backup process continues, but the shadow copy of the volume remains constant.
See also Microsoft Volume Shadow Copy Service and replica.

shadow copy provider *(Microsoft VSS specific term)* An entity that performs the work on creating and representing the volume shadow copies. Providers own the shadow copy data and expose the shadow copies. Providers can be software (for example, system providers) or hardware (local disks, disk arrays).
See also shadow copy.

shadow copy set *(Microsoft VSS specific term)* A collection of shadow copies created at the same point in time.
See also shadow copy and replica set.

shared disks A Windows disk on another system that has been made available to other users on the network. Systems with shared disks can be backed up without a Data Protector Disk Agent installed.

SIBF The Serverless Integrations Binary Files (SIBF) is a part of the IDB that stores raw NDMP meta data. This data is necessary to perform restore of NDMP objects.

Site Replication Service *(Microsoft Exchange Server specific term)* The Microsoft Exchange Server 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.

smart copy *(VLS specific term)* A copy of the backed up data created from the virtual tape to the physical tape library. The smart copy process allows Data Protector to distinguish between the source and the target medium thus enabling media management.
See also Virtual Library System (VLS).

smart copy pool *(VLS specific term)* A pool that defines which destination library slots are available as smart copy targets for a specified source virtual library.
See also Virtual Library System (VLS) and smart copy.

SMB *See* split mirror backup.

SMBF The Session Messages Binary Files (SMBF) part of the IDB stores session messages generated during backup, restore, object copy, object consolidation, object verification, and media management sessions. One binary file is created per session. The files are grouped by year and month.

SMI-S Agent (SMISA) *See* HP StorageWorks P6000 EVA SMI-S Agent.

snapshot *(HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, and HP P4000 SAN Solutions specific term)* A type of target volumes created using a specific replication technology. Depending on the disk array model and the chosen replication technique, a range of snapshot types with different characteristics is available. Basically, each snapshot may be either a virtual copy, still reliant upon the contents of the source volume, or an independent duplicate (clone) of the source volume.
See also replica and snapshot creation.

snapshot backup *See* ZDB to tape, ZDB to disk, and ZDB to disk+tape.

snapshot creation *(HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, and HP P4000 SAN Solutions specific term)* A replica creation process in which copies of the selected source volumes are created using storage virtualization technology. Such a replica is considered to be created at a particular point in time, and is immediately available for use. However, with certain snapshot types, a background data copying process continues to run on the disk array after the moment of the replica creation.
See also snapshot.

source (R1) device *(EMC Symmetrix specific term)* An EMC Symmetrix device that participates in SRDF operations with a target (R2) device. All writes to this device are mirrored to a target (R2) device in a remote EMC Symmetrix unit. An R1 device must be assigned to an RDF1 group type.
See also target (R2) device.

source volume *(ZDB specific term)* A storage volume containing data to be replicated.

sparse file A file that contains data with portions of empty blocks. Examples are: a matrix in which some or much of the data contains zeros, files from image applications, and high-speed databases. If sparse file processing is not enabled during restore, it might be impossible to restore this file.

split mirror *(EMC Symmetrix Disk Array and HP P9000 XP Disk Array Family specific term)* A type of target volumes created using a specific replication technology. A split-mirror replica provides independent duplicates (clones) of the source volumes.
See also replica and split mirror creation.

split mirror backup *(EMC Symmetrix specific term)* *See* ZDB to tape.

split mirror backup *(HP P9000 XP Disk Array Family specific term)* *See* ZDB to tape, ZDB to disk, and ZDB to disk+tape.

split mirror creation *(EMC Symmetrix and HP P9000 XP Disk Array Family specific term)* A replica creation technique, in which a pre-configured set of target volumes (a mirror) is kept synchronized with a set of source volumes until the time at which a replica of the contents of the source volumes is required. Then, the synchronization is stopped (the mirror is split) and a split mirror replica of the source volumes at the time of the split remains in the target volumes.
See also split mirror.

split mirror restore *(EMC Symmetrix and HP P9000 XP Disk Array Family specific term)* A process in which data backed up in a ZDB-to-tape or a ZDB-to-disk+tape session is first copied from the backup media to a replica, and from the replica to the source volumes afterwards. Individual backup objects or complete sessions can be restored using this method.
See also ZDB to tape, ZDB to disk+tape, and replica.

sqlhosts file or registry
(Informix Server specific term) An Informix Server connectivity information file (on UNIX) or registry (on Windows) that contains the names of each of the database servers and any aliases to which the clients on a host computer can connect.

SRD file
(disaster recovery specific term) A text file in the Unicode (UTF-16) format, generated during CONFIGURATION backup of a Windows or Linux system and stored on the Cell Manager. It contains system information required for installing and configuring the operating system on the target system in the case of a disaster.
See also target system.

SRDF
(EMC Symmetrix specific term) The EMC Symmetrix Remote Data Facility is a business continuation process that enables effective, real-time data replication of SLDs between dislocated processing environments. These environments could be situated within the same root computer environment or separated by long distances.

SSE Agent (SSEA)
See HP StorageWorks P9000 XP Agent.

sst.conf file
The file `/usr/kernel/drv/sst.conf` is required on each Data Protector Sun Solaris client to which a multi-drive library device is connected. It must contain an entry for the SCSI address of the robotic mechanism of each library device connected to the client.

st.conf file
The file `/kernel/drv/st.conf` is required on each Data Protector Solaris client with a backup device connected. It must contain device information and a SCSI address for each backup drive connected to that client. A single SCSI entry is required for a single-drive device and multiple SCSI entries are required for a multi-drive library device.

stackers
Devices with multiple slots for media storage usually with only one drive. A stacker selects media from the stack sequentially. In contrast, a library can randomly select media from its repository.

standalone file device
A file device is a file in a specified directory to which you back up data.

Storage Group
(Microsoft Exchange Server specific term) A collection of mailbox stores and public folder stores that share a set of transaction log files. Exchange Server manages each storage group with a separate server process.

storage volume
(ZDB specific term) An object that may be presented to an operating system or some other entity (for example, a virtualization mechanism) upon which volume management systems, filesystems, or other objects may exist. The volume management systems, file systems are built on this storage. Typically, these can be created or exist within a storage system such as a disk array.

StorageTek ACS library
(StorageTek specific term) Automated Cartridge System is a library system (also known as Silo) consisting of one Library Management Unit (LMU) and one to 24 Library Storage Modules (LSM) connected to the unit.

switchover
See failover.

Sybase Backup Server API
(Sybase specific term) An industry-standard interface developed for the exchange of backup and recovery information between a Sybase SQL Server and a backup solution like Data Protector.

Sybase SQL Server
(Sybase specific term) The server in the Sybase “client-server” architecture. Sybase SQL Server manages multiple databases and multiple users, keeps track of the actual location of data on disks, maintains mapping of logical data description to physical data storage, and maintains data and procedure caches in memory.

SYMA
See EMC Symmetrix Agent.

synthetic backup
A backup solution that produces a synthetic full backup, an equivalent to a conventional full backup in terms of data, without putting stress on the production servers or the network. A synthetic full backup is created from a previous full backup and any number of incremental backups.

synthetic full backup
The result of an object consolidation operation, where a restore chain of a backup objects is merged into a new, synthetic full version of this object. A synthetic full backup is equivalent to a conventional full backup in terms of restore speed.

System Backup to Tape
(Oracle specific term) An Oracle interface that handles the actions required to load, label, and unload correct backup devices when Oracle issues a backup or restore request.

system databases *(Sybase specific term)* The four system databases on a newly installed Sybase SQL Server are the:

- master database (master)
- temporary database (tempdb)
- system procedure database (sybsystemprocs)
- model database (model).

System Recovery Data file *See* SRD file.

System State *(Windows specific term)* The System State data comprises the Registry, COM+ Class Registration database, system startup files, and the Certificate Services database (if the server is a certificate server). If the server is a domain controller, Active Directory services and the SYSVOL directory are also contained in the System State data. If the server is running the Cluster service, the System State data also includes resource registry checkpoints and the quorum resource recovery log, which contains the most recent cluster database information.

system volume/disk/partition A volume/disk/partition containing operating system files. Microsoft terminology defines the system volume/disk/partition as the volume/disk/partition containing files required for the initial step of the boot process.

SysVol *(Windows specific term)* A shared directory that stores the server copy of the domain's public files, which are replicated among all domain controllers in the domain.

T

tablespace A part of a database structure. Each database is logically divided into one or more tablespaces. Each tablespace has data files or raw volumes exclusively associated with it.

tapeless backup *(ZDB specific term)* *See* ZDB to disk.

target (R2) device *(EMC Symmetrix specific term)* An EMC Symmetrix device that participates in SRDF operations with a source (R1) device. It resides in the remote EMC Symmetrix unit. It is paired with a source (R1) device in the local EMC Symmetrix unit and receives all write data from its mirrored pair. This device is not accessed by user applications during normal I/O operations. An R2 device must be assigned to an RDF2 group type.
See also source (R1) device.

target database *(Oracle specific term)* In RMAN, the target database is the database that you are backing up or restoring.

target system *(disaster recovery specific term)* A system after a computer disaster has occurred. The target system is typically in a non-bootable state and the goal of disaster recovery is to restore this system to the original system configuration. The difference between a faulty system and a target system is that a target system has all faulty hardware replaced.

target volume *(ZDB specific term)* A storage volume to which data is replicated.

Terminal Services *(Windows specific term)* Windows Terminal Services provide a multi-session environment that allows clients to access a virtual Windows desktop session and Windows-based programs running on the server.

thread *(Microsoft SQL Server specific term)* An executable entity that belongs to only one process. It comprises a program counter, a user-mode stack, a kernel-mode stack, and a set of register values. Several threads can run at a time within one process.

TimeFinder *(EMC Symmetrix specific term)* A business continuation process that creates an instant copy of single or multiple Symmetrix Logical Devices (SLDs). The instant copy is created on specially preconfigured SLDs called BCVs and is accessible via a separate device address to the system(s).

TLU Tape Library Unit.

TNSNAMES.ORA *(Oracle and SAP R/3 specific term)* A network configuration file that contains connect descriptors mapped to service names. The file may be maintained centrally or locally, for use by all or individual clients.

transaction A mechanism for ensuring that a set of actions is treated as a single unit of work. Databases use transactions to keep track of database changes.

transaction backup Transaction backups generally use fewer resources than database backups, so they can be created more frequently than database backups. By applying transaction backups, you can recover the database to a specific point in time prior to when a problem occurred.

transaction backup *(Sybase and SQL specific term)* A backup of the transaction log providing a record of changes made since the last full or transaction backup.

transaction log backup Transaction log backups generally use fewer resources than database backups so they can be created more frequently than database backups. By applying transaction log backups, you can recover the database to a specific point in time.

transaction log files Files that record transactions of the database modifications, and provide fault tolerance in case of a database disaster.

transaction log table *(Sybase specific term)* A system table in which all changes to the database are automatically recorded.

transaction logs *(Data Protector specific term)* Keep track of IDB changes. The archiving of transaction logs should be enabled to prevent you from losing the transaction log files that are created after the last IDB backup and are necessary for IDB recovery.

transportable snapshot *(Microsoft VSS specific term)* A shadow copy that is created on the application system and can be presented to the backup system where a backup can be performed.
See also Microsoft Volume Shadow Copy Service (VSS).

TSANDS.CFG file *(Novell NetWare specific term)* A file that allows you to specify the names of containers where you want backups to begin. It is a text file located in the `SYS:SYSTEM\TSA` directory on the server where `TSANDS.NLM` is loaded.

U

UIProxy The Java GUI Server (`UIProxy` service) runs on the Data Protector Cell Manager. It is responsible for communication between the Java GUI Client and the Cell Manager, moreover, it performs business logic operations and sends only important information to the client. The service is started as soon as Data Protector is installed on the Cell Manager.

unattended operation *See* lights-out operation.

user account (Data Protector user account) You can use Data Protector only if you have a Data Protector user account, which restricts unauthorized access to Data Protector and to backed up data. Data Protector administrators create this account specifying a user logon name, the systems from which the user can log on, and a Data Protector user group membership. This is checked whenever the user starts the Data Protector user interface or performs specific tasks.

User Account Control (UAC) A security component in Windows Vista, Windows 7, and Windows Server 2008 that limits application software to standard user privileges until an administrator authorizes an increase in privilege level.

user disk quotas NTFS quota-management support enables an enhanced tracking mechanism and control over disk space usage on shared storage volumes. Data Protector backs up user disk quotas on the whole system and for all configured users at a time.

user group Each Data Protector user is member of a User Group. Each User Group has a set of user rights that are given to every user in that User Group. The number of User Groups with their associated user rights can be defined as desired. Data Protector provides three default user groups: admin, operator, and user.

user profile *(Windows specific term)* Configuration information retained on a user basis. This information includes desktop settings, screen colors, network connections, and so on. When the user logs on, the user profile is loaded and the Windows environment is set accordingly.

user rights User rights or access rights are the permissions needed to perform specific Data Protector tasks. Configuring a backup, starting a backup session, or starting a restore session are typical user rights. Users have the access rights of the user group to which they belong.

user_restrictions file A file that restricts specific user actions, which are available to Data Protector user groups according to the user rights assigned to them, to be performed only on specific systems of the Data Protector cell. Such restrictions apply only to Data Protector user groups other than *admin* and *operator*.

V

vaulting media The process of storing media to a safe and remote place. The media are brought back to the data center when they are needed for restore or are ready for reuse in subsequent backups. The vaulting procedure depends on your company's backup strategy and policies for data protection/reliability.

verify A function that lets you check whether the Data Protector data on a specified medium is readable. Additionally, consistency within each block can be checked if the backup was performed with the cyclic redundancy check (CRC) option ON.

Virtual Controller Software (VCS) *(HP P6000 EVA Disk Array Family specific term)* The firmware that manages all aspects of storage system operation, including communication with HP Command View EVA through the HSV controllers.
See also HP Command View (CV) EVA.

Virtual Device Interface *(Microsoft SQL Server specific term)* This is a Microsoft SQL Server programming interface that allows fast backup and restore of large databases.

virtual disk *(HP P6000 EVA Disk Array Family specific term)* A unit of storage allocated from a storage pool of a disk array of the HP P6000 EVA Disk Array Family. A virtual disk is the entity that can be replicated using the snapshot functionality of such disk array.
See also source volume and target volume.

virtual full backup An efficient type of synthetic backup where data is consolidated using pointers instead of being copied. It is performed if all the backups (the full backup, incremental backups, and the resulting virtual full backup) are written to a single file library that uses distributed file medium format.

Virtual Library System (VLS) A disk-based data storage device hosting one or more virtual tape libraries (VTLs).

virtual server A virtual machine in a cluster environment defined in a domain by a network IP name and address. Its address is cached by the cluster software and mapped to the cluster node that is currently running the virtual server resources. This way all requests for a particular virtual server are cached by a specific cluster node.

virtual tape *(VLS specific term)* An archival storage technology that backs up data to disk drives in the same way as if it were being stored on tape. Benefits of virtual tape systems include improved backup and recovery speed and lower operating costs.
See also Virtual Library System (VLS) and Virtual Tape Library (VTL).

Virtual Tape Library (VTL) *(VLS specific term)* An emulated tape library that provides the functionality of traditional tape-based storage.
See also Virtual Library System (VLS).

VMware management client *(VMware (Legacy) integration specific term)* The client that Data Protector uses to communicate with VMware Virtual Infrastructure. This can be a VirtualCenter Server system (VirtualCenter environment) or an ESX Server system (standalone ESX Server environment).

volser *(ADIC and STK specific term)* A VOLume SERial number is a label on the medium to identify the physical tape used in very large libraries. A volser is a naming convention specific to ADIC/GRAU and StorageTek devices.

volume group A unit of data storage in an LVM system. A volume group can consist of one or more physical volumes. There can be more than one volume group on the system.

volume mountpoint *(Windows specific term)* An empty directory on a volume that can be used to mount another volume. The volume mount point acts as a gateway to the target volume. Provided that the volume is mounted, users and applications can refer to the data on the mounted volume by the full (merged) filesystem path as if both volumes are one integral part.

Volume Shadow Copy Service *See* Microsoft Volume Shadow Copy Service (VSS).

VSS *See* Microsoft Volume Shadow Copy Service (VSS).

VSS compliant mode *(HP P9000 XP Disk Array Family VSS provider specific term)* One of two P9000 XP Array VSS hardware provider operation modes. When the P9000 XP Array provider is in the VSS compliant mode, the source volume (P-VOL) and its replica (S-VOL) are in simplex, unpaired state after a backup. Therefore the number of replicas (S-VOLs per a P-VOL) rotated is not limited. Restore from a backup in such a configuration is possible only by switching the disks.
See also resync mode, source volume, primary volume (P-VOL), replica, secondary volume (S-VOL), and replica set rotation.

VxFS Veritas Journal Filesystem.

VxVM (Veritas Volume Manager) A Veritas Volume Manager is a system for managing disk space on Solaris platforms. A VxVM system consists of an arbitrary group of one or more physical volumes organized into logical disk groups.

W

Wake ONLAN Remote power-up support for systems running in power-save mode from some other system on the same LAN.

Web reporting The Data Protector functionality that allows you to view reports on backup, object copy, and object consolidation status and Data Protector configuration using the Web interface.

wildcard character A keyboard character that can be used to represent one or many characters. The asterisk (*), for example, typically represents one or more characters, and the question mark (?) typically represents a single character. Wildcard characters are often used in operating systems as a means of specifying more than one file by name.

Windows configuration backup Data Protector allows you to back up Windows CONFIGURATION, including Windows Registry, user profiles, Event Logs, and WINS and DHCP server data (if configured on a system) in one step.

Windows Registry A centralized database used by Windows to store configuration information for the operating system and the installed applications.

WINS server A system running Windows Internet Name Service software that resolves Windows networking computer names to IP addresses. Data Protector can back up WINS server data as part of the Windows configuration.

writer *(Microsoft VSS specific term)* A process that initiates change of data on the original volume. Writers are typically applications or system services that write persistent information on a volume. Writers also participate in the shadow copy synchronization process by assuring data consistency.

X

XBSA interface *(Informix Server specific term)* ON-Bar and Data Protector communicate with each other through the X/Open Backup Services Application Programmer's Interface (XBSA).

Z

ZDB *See* zero downtime backup (ZDB).

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 in zero downtime backup, instant recovery, and split mirror restore sessions.
See also zero downtime backup (ZDB).

ZDB to disk *(ZDB specific term)* A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time. Multiple replicas, produced using the same backup specification at different times, can be kept in a replica set. A replica from ZDB to disk can be restored using the instant recovery process.
See also zero downtime backup (ZDB), ZDB to tape, ZDB to disk+tape, instant recovery, and replica set rotation.

ZDB to disk+tape *(ZDB specific term)* A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time, in the same way as with ZDB to disk. However, data in the replica is also streamed to a backup medium, as with ZDB to tape. If this backup method is used, data backed up in the same session can be restored using

the instant recovery process, the standard Data Protector restore from tape, or with specific disk array families, split mirror restore.
See also zero downtime backup (ZDB), ZDB to disk, ZDB to tape, instant recovery, replica, and replica set rotation.

ZDB to tape
(ZDB specific term) A form of zero downtime backup where data in the replica produced is streamed to a backup medium, typically tape. Instant recovery is not possible from such a backup, so the replica need not be retained on the disk array after backup completion. The backed up data can be restored using standard Data Protector restore from tape. With specific disk array families, split mirror restore can also be used.
See also zero downtime backup (ZDB), ZDB to disk, ZDB to disk+tape, instant recovery, and replica.

zero downtime backup (ZDB)
A backup approach in which data replication techniques, provided by a disk array, are used to minimize the impact of backup operations on an application system. A replica of the data to be backed up is created first. All subsequent backup operations are performed on the replicated data rather than the original data, while the application system can be returned to normal operation.
See also ZDB to disk, ZDB to tape, ZDB to disk+tape, and instant recovery.

Index