

HP Data Protector 8.00

Zero Downtime Backup Concepts Guide

HP Part Number: N/A
Published: August 2013
Edition: Third



© Copyright 2013 Hewlett-Packard Development Company, L.P.

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

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

Intel®, Itanium®, Pentium®, Intel Inside®, and the Intel Inside logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

Adobe and Acrobat are trademarks of Adobe Systems Incorporated.

Java is a registered trademark of Oracle and/or its affiliates.

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

UNIX® is a registered trademark of The Open Group.

LiveVault® is a registered trademark of Autonomy Corporation plc.

Contents

Publication history.....	7
About this guide.....	8
Intended audience.....	8
Documentation set.....	8
Help.....	8
Guides.....	8
Documentation map.....	11
Abbreviations.....	11
Map.....	12
Integrations.....	12
Document conventions and symbols.....	13
Data Protector graphical user interface.....	14
General information.....	14
HP technical support.....	14
Subscription service.....	14
HP websites.....	15
Documentation feedback.....	15
1 Overview.....	16
Introduction.....	16
Zero downtime backup.....	16
Online and offline creation of replicas.....	17
Creating replicas.....	17
ZDB types.....	17
Support on disk arrays.....	18
Instant recovery and restore of ZDB data.....	19
Instant recovery.....	19
Other ZDB data restore methods.....	19
Restore possibilities for ZDB types.....	20
2 Replication techniques.....	21
Disk array basics.....	21
RAID technology.....	21
Replication techniques.....	22
Local replication.....	22
Split mirror replication.....	22
Snapshot replication.....	23
Standard snapshot.....	24
Vsnap.....	25
Snapclone.....	26
Local replication integrating with HP-UX LVM mirroring.....	27
Remote replication.....	28
Split mirror replication.....	28
Remote plus local replication.....	29
Split mirror replication.....	29
Snapshot replication.....	29
3 Using Data Protector for ZDB and instant recovery.....	31
Data Protector cells.....	31
Cell components.....	31
Cell Manager.....	32
Application systems.....	32

Backup system.....	32
ZDB database.....	32
User interfaces.....	33
GUI.....	33
CLI.....	34
Disk array integrations available with Data Protector.....	34
HP P4000 SAN Solutions.....	35
HP P6000 EVA Disk Array Family.....	35
P6000 EVA Array storage presentation.....	35
Local replication.....	35
Local replication integrating with LVM mirroring.....	35
Remote plus local replication.....	36
HP P9000 XP Disk Array Family.....	36
Local replication.....	37
Local replication integrating with LVM mirroring.....	37
Remote replication.....	37
Remote plus local replication.....	38
HP 3PAR StoreServ Storage.....	38
EMC Symmetrix.....	38
Local replication.....	39
Local replication integrating with LVM mirroring.....	39
Remote replication.....	39
Remote plus local replication.....	40
Application integrations.....	40
Application data consistency.....	41
Transaction logs.....	41
Restore.....	41
Application integrations and Microsoft Volume Shadow Copy Service.....	41
4 Replica life cycle.....	43
Overview.....	43
Creating replicas.....	43
Replica sets.....	44
Replica set rotation.....	44
Scheduling replication.....	44
Using replicas.....	44
ZDB to tape.....	45
ZDB to disk.....	45
ZDB to disk+tape.....	45
Instant recovery.....	46
Deleting replicas.....	47
5 ZDB session process.....	48
ZDB process overview.....	48
Locating data objects.....	48
Freezing operation of the application or database.....	48
Creating a replica.....	49
Replicating the data objects.....	49
Streaming the replica to tape.....	49
Backing up a replica to tape.....	50
Creating mount points.....	50
Standard data movement to tape.....	50
Incremental ZDB.....	50
The replica after creation.....	50
Mounting the replica on the backup system.....	50
Recording session information.....	50

Writing session information to the IDB.....	51
6 Instant recovery and other restore techniques from ZDB sessions.....	52
Overview.....	52
Instant recovery.....	52
Standard Data Protector restore.....	52
Split mirror restore.....	53
Instant recovery.....	53
Instant recovery process.....	54
Instant recovery and LVM mirroring.....	56
Instant recovery in a cluster.....	56
Split mirror restore.....	56
Split mirror restore process.....	57
7 Planning.....	58
Introduction.....	58
Flexibility in recovery.....	58
Split mirror disk arrays.....	58
Snapshot disk arrays.....	59
Disk array-specific considerations.....	59
Replica sets on P4000 SAN Solutions.....	59
Instant recovery on P4000 SAN Solutions.....	59
Replica creation on P6000 EVA Array.....	59
Replica set rotation on P6000 EVA Array.....	60
Instant recovery on P6000 EVA Array.....	60
Replica type selection on P9000 XP Array.....	60
Instant recovery on P9000 XP Array.....	61
Replica creation on 3PAR StoreServ system.....	61
Concurrency handling.....	61
Locking.....	61
Backup device locking.....	61
Disk locking.....	61
Backup scenarios.....	61
A Supported configurations.....	63
Introduction.....	63
Supported HP P6000 EVA Disk Array Family configurations.....	63
Local replication configurations.....	63
Local replication configurations with HP-UX LVM mirroring.....	64
Remote plus local replication configurations.....	67
Supported HP P9000 XP Disk Array Family configurations.....	68
Local replication configurations.....	68
Single-host (BC1) configuration.....	69
Cascading configurations.....	69
Local replication configurations with HP-UX LVM mirroring.....	70
Remote replication configurations.....	71
Remote plus local replication configurations.....	73
Cluster configurations.....	74
Supported EMC Symmetrix configurations.....	74
Local replication configurations.....	74
Local replication configurations with HP-UX LVM mirroring.....	75
Remote replication configurations.....	77
Remote plus local replication configurations.....	78

Cluster configurations.....	79
Glossary.....	81
Index.....	110

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
N/A	June 2013	Data Protector release 8.00
N/A	June 2013 (second edition)	Data Protector release 8.00
N/A	August 2013	Data Protector release 8.00

About this guide

This guide describes zero downtime backup and instant recovery concepts and how these are used within Data Protector.

Intended audience

This guide is intended for users interested in understanding the concepts of the Data Protector zero downtime backup and instant recovery capabilities and who wish to improve backup strategies for high-availability systems. It is recommended to use this guide together with the *HP Data Protector Concepts Guide* and the task-oriented *HP Data Protector Zero Downtime Backup Administrator's Guide* and *HP Data Protector Zero Downtime Backup Integration Guide*.

Documentation set

The Help and other guides provide related information.

NOTE: The documentation set available at the HP support website at <http://support.openview.hp.com/selfsolve/manuals> contains the latest updates and corrections.

Help

Data Protector provides Help topics and context-sensitive (F1) Help for Windows and UNIX platforms. Install the Help during the Data Protector setup procedure by selecting the installation component *English Documentation (Guides, Help)* (Windows systems) or *OB2-DOCS* (on UNIX systems). Once installed, the Help resides in the following directory:

Windows systems: `Data_Protector_home\help\enu`

UNIX systems: `/opt/omni/help/C/help_topics`

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

Windows systems: Open `DP_help.chm`.

UNIX systems: Unpack the zipped tar file `DP_help.tar.gz` and open `DP_help.htm`.

Guides

Data Protector guides are available in the electronic PDF format. Install the PDF files during the Data Protector setup procedure by selecting the installation component *English Documentation (Guides, Help)* (on Windows systems) or *OB2-DOCS* (on UNIX systems). Once installed, the guides reside in the following directory:

Windows systems: `Data_Protector_home\docs`

UNIX systems: `/opt/omni/doc/C`

You can also access the guides:

- From the **Help** menu of the Data Protector graphical user interface
- From the HP support website at <http://support.openview.hp.com/selfsolve/manuals> (where the most up-to-date guide versions are available)

Data Protector guides are:

- *HP Data Protector Getting Started Guide*
This guide contains information to get you started with using Data Protector. It lists installation prerequisites, provides instructions on installing and configuring a basic backup environment and procedures for performing backup and restore. It also lists resources for further information.
- *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 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 Command Line Interface Reference*
This guide describes the Data Protector command-line interface, command options, and their usage as well as provides some basic command-line examples. It is located in the following directory:
Windows systems: `Data_Protector_home\docs\MAN`
UNIX systems: `/opt/omni/doc/C/`
On UNIX systems, you can use the `omniintro` man page to display a list of the available Data Protector commands. You can then execute the `man CommandName` command to retrieve information about each Data Protector command.
- *HP Data Protector Product Announcements, Software Notes, and References*
This guide gives a description of new features of HP Data Protector 8.00. It also provides information on installation requirements, required patches, and limitations, as well as known issues and workarounds.
- *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 and 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 and Network Data Management Protocol Server*
This guide describes the integrations of Data Protector with Sybase Server and Network Data Management Protocol Server.
- *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 Virtualization Environments*
This guide describes the integrations of Data Protector with virtualization environments: VMware Virtual Infrastructure, VMware vSphere, VMware vCloud Director, Microsoft Hyper-V, and Citrix XenServer.
- *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 P4000 SAN Solutions, HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, HP 3PAR StoreServ Storage, 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 Exchange Server*
This guide describes how to configure and use the Data Protector Granular Recovery Extension for Microsoft Exchange Server. Graphical user interface of the Data Protector Granular Recovery Extension for Microsoft Exchange Server is integrated into the Microsoft Management Console. This guide is intended for Microsoft Exchange Server administrators and Data Protector backup administrators.
- *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 Deduplication*
This technical white paper describes basic data deduplication concepts, principles of Data Protector integration with Backup to Disk devices and its use of deduplication. It also provides instructions how to configure and use deduplication in Data Protector backup environments.
- *HP Data Protector Integration with Autonomy IDOL Server*
This technical white paper describes all aspects of integrating Data Protector with Autonomy IDOL Server: integration concepts, installation and configuration, Data Protector backup image indexing, full content search-based restore, and troubleshooting.
- *HP Data Protector Integration with Autonomy LiveVault*
This technical white paper all aspects of integrating Data Protector with Autonomy LiveVault: integration concepts, installation and configuration, backup policy management, cloud backup, cloud restore, and troubleshooting.

Documentation map

Abbreviations

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

Abbreviation	Documentation item
CLI	Command Line Interface Reference
Concepts	Concepts Guide
DR	Disaster Recovery Guide
GS	Getting Started Guide
GRE Exchange	Granular Recovery Extension User Guide for Microsoft Exchange Server
GRE SPS	Granular Recovery Extension User Guide for Microsoft SharePoint Server
GRE VMware	Granular Recovery Extension User Guide for VMware vSphere
Help	Help
Install	Installation and Licensing Guide
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 VSS	Integration Guide for Microsoft Volume Shadow Copy Service
IG O/S	Integration Guide for Oracle and SAP
IG Var	Integration Guide for Sybase and Network Data Management Protocol Server
IG VirtEnv	Integration Guide for Virtualization Environments
IG IDOL	Integration with Autonomy IDOL Server
IG LV	Integration with Autonomy LiveVault

Abbreviation	Documentation item
PA	Product Announcements, Software Notes, and References
Trouble	Troubleshooting Guide
ZDB Admin	ZDB Administrator's Guide
ZDB Concepts	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.

									Integr. guides						ZDB			GRE		
	Help	GS	Concepts	Install	Trouble	DR	CLI	PA	MS	O/S	IBM	Var	VSS	VirtEnv	Concepts	Admin	IG	Exchange	SPS	VMware
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
Disaster recovery	X		X			X														
Installation, upgrade	X	X		X				X												
Instant recovery	X		X												X	X	X			
Licensing	X			X				X												
Limitations	X				X			X	X	X	X	X	X	X			X			
New features	X							X												
Planning strategy	X		X												X					
Procedures, tasks	X			X	X	X			X	X	X	X	X	X		X	X	X	X	X
Recommendations			X					X							X					
Requirements				X				X	X	X	X	X	X	X						
Restore	X	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

Integrations

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

Software application	Guides
Autonomy IDOL Server	IG IDOL
Autonomy LiveVault	IG LV
IBM DB2 UDB	IG IBM
Informix Server	IG IBM
Lotus Notes/Domino Server	IG IBM
Microsoft Exchange Server	IG MS, ZDB IG, GRE Exchange

Software application	Guides
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
Sybase Server	IG Var
VMware vCloud Director	IG VirtEnv
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 Concepts, ZDB Admin, IG VSS
HP P6000 EVA Disk Array Family	all ZDB, IG VSS
HP P9000 XP Disk Array Family	all ZDB, IG VSS
HP 3PAR StoreServ Storage	ZDB Concepts, ZDB Admin, 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	<ul style="list-style-type: none"> 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
<i>Italic</i> text	Text emphasis
Monospace text	<ul style="list-style-type: none"> File and directory names System output Code Commands, their arguments, and argument values
<i>Monospace, italic</i> text	<ul style="list-style-type: none"> 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 graphical user interface for Microsoft Windows operating systems. For information about it, see the *HP Data Protector Help*.

Figure 1 Data Protector graphical user interface



General information

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

HP technical support

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

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

Before contacting HP, collect the following information:

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

Subscription service

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

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

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

HP websites

For additional information, see the following HP websites:

- <http://www.hp.com>
- <http://www.hp.com/go/software>
- <http://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 with the subject line Feedback on Data Protector documentation to AutonomyTPFeedback@hp.com. All submissions become the property of HP.

1 Overview

Introduction

Zero downtime backup (ZDB) and instant recovery (IR) have two great advantages over other backup and restore techniques:

- Minimal downtime or impact on the application system during backup
- Significantly shorter restore times

The growing requirement for data security for mission critical applications, together with the increasing sophistication of Storage Area Network (SAN) environments, has resulted in a rapid expansion in the use of large disk arrays containing RAID technology. These can hold large application databases, containing vast amounts of data.

By using storage virtualization techniques, disk arrays can be divided into many virtual disks. These can easily be copied within an disk array, perhaps many times dependent on disk array technology and the available storage space. This makes it possible to perform operations on copies of data without any risk to the original data. In particular, it enables effective backup solutions for applications in high-availability and mission-critical areas.

Conventional tape backup and restore techniques are not fast enough to handle the enormous amounts of data involved in a world of terabyte databases where information is expected to be available 24 hours a day.

This guide describes ZDB and instant recovery techniques that use the potential of disk arrays to streamline backup and recovery.

Zero downtime backup

Conventional methods of backing up to tape are not well suited for large database applications; either the database has to be taken offline or, if the application allows it, put into “hot-backup mode” while data in it is streamed to tape.


The first can cause major disruption to the application’s operation. The second can produce many large transaction log files, putting extra load on the application system.

Zero downtime backup (ZDB) uses disk array technology to minimize the disruption. In very general terms, a copy or **replica** of the data is created or maintained on a disk array. This is very fast and has little impact on the application’s performance. The replica can itself form the backup, or it can be streamed to tape without further interruption to the application’s use of the source database.

Depending on the hardware and software with which it is created, a replica may be an exact duplicate (mirror, snapclone), or a virtual copy (snapshot) of the data being backed up.

In ZDB, **replication** (the process of creating or maintaining a replica) is the critical factor in minimizing interruption to the application.

Figure 2 Zero downtime backup and instant recovery concept



Online and offline creation of replicas

For database applications, backup can be performed with the database online or offline:

- **Online backup**

The database is placed in hot-backup mode while a replica of sections to be backed up is created. In this mode, any changes to the database are written to transaction logs, not the database itself. When the database is fully functional again, it is updated from the transaction logs. This allows the database to be operated on without stopping the application.

- **Offline backup**

Database operation is simply stopped while a replica is created. No transactions are possible during this time.

After the replica is created, the database returns to normal operation. Any subsequent backup operations, such as streaming data to tape, are performed on the replica, leaving the database online and unaffected.

In both cases, the effect on the application is limited to the period during which the replica is created, much less than with standard tape backup techniques. For online backup, database operation is never stopped (zero downtime) and the effect on performance can be minimal, limited mainly to the effect of having to write increased information to the transaction logs.

Creating replicas

Replication creates a replica of application or filesystem data at a particular moment.

The volumes containing the source or original data objects to be replicated are referred to as **source volumes**. These are replicated to an equivalent number of **target volumes**. When the replication process is complete, the data in the target volumes constitutes the replica.

Currently there are two basic replication techniques (described in more detail in [“Replication techniques”](#) (page 21)):

- **Split mirror**

A mirror is a dynamic duplicate of the source data, synchronized with it. Any changes to the source are also applied to the mirror.

The technique allows a duplicate of filesystem or application data to be created and maintained during normal application use.

To create a replica, the mirror is temporarily split from the source. Data is backed up from the mirror and the mirror is then resynchronized with the source.

For more details, see [“Split mirror replication”](#) (page 22).

- **Snapshot**

A snapshot replica is created by making a copy of data at a particular moment. The snapshot can be a full copy, thus independent of the source volume, or a virtual copy that still depends on the source volume.

For more details, see [“Snapshot replication”](#) (page 23).

ZDB types

After a replica has been created, by whatever method, it can be backed up. It is mounted to a **backup system** connected to the disk array on which the replica was created. To take full advantage of ZDB, this should be a separate computer system. There are then three forms of ZDB:

- **ZDB to tape** – see [“ZDB to tape”](#) (page 45)

1. Data in the replica is streamed to tape according to the tape backup type you have selected:

HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, HP 3PAR StoreServ Storage (through the HP P6000 / HP 3PAR SMI-S Agent), EMC Symmetrix:

Full, Incr, Incr1-9.

HP P4000 SAN Solutions, HP 3PAR StoreServ Storage (through the HP 3PAR VSS Agent):

Full

2. The replica can then be discarded.

Data can be restored from the tape using standard Data Protector techniques.

- **ZDB to disk** – see “ZDB to disk” (page 45)

The replica is kept on the disk array and used as the backup.

Data can be restored using instant recovery (see “Instant recovery” (page 19)), which recovers the complete replica.

- **ZDB to disk+tape** – see “ZDB to disk+tape” (page 45)

1. Data in the replica is streamed to tape according to the tape backup type you have selected:

HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, EMC Symmetrix:

Full, Incr, Incr1-9.

HP P4000 SAN Solutions, HP 3PAR StoreServ Storage:

Full

2. The replica is kept on the disk array.

This provides extra flexibility because data can be restored in two ways:

- Using standard Data Protector restore from tape (allowing restore of individual backup objects)
- Directly from the replica using instant recovery (see “Instant recovery” (page 19)) of the complete replica

Support on disk arrays

Table 3 ZDB types and replication techniques versus disk array families

	Split mirror		Snapshot			
ZDB type and replication technique	HP P9000 XP Disk Array Family	EMC Symmetrix	HP P4000 SAN Solutions	HP P6000 EVA Disk Array Family	HP P9000 XP Disk Array Family	HP 3PAR StoreServ Storage
ZDB to tape, local	Yes	Yes	Yes	Yes	Yes	Yes
ZDB to tape, remote	Yes	Yes	No	No	No	No
ZDB to tape, remote + local	Yes	Yes	No	Yes	Yes	No

Table 3 ZDB types and replication techniques versus disk array families *(continued)*

ZDB type and replication technique	Split mirror		Snapshot			
	HP P9000 XP Disk Array Family	EMC Symmetrix	HP P4000 SAN Solutions	HP P6000 EVA Disk Array Family	HP P9000 XP Disk Array Family	HP 3PAR StoreServ Storage
ZDB to disk, local	Yes	No	Yes	Yes	Yes	Yes ¹
ZDB to disk+tape, local	Yes	No	Yes	Yes	Yes	Yes ²

¹ This ZDB type is only available in configurations where the application and backup systems are Windows systems.

² This ZDB type is only available in configurations where the application and backup systems are Windows systems.

Local and **remote** refer to the disk array on which the replica is made, whether it is the same disk array on which the source data resides or a separate disk array on a remote site. For details of these terms and their implications, see:

- “Local replication” (page 22)
- “Remote replication” (page 28)
- “Remote plus local replication” (page 29)

Instant recovery and restore of ZDB data

Instant recovery

Instant recovery requires a replica to exist on the same disk array to which the data is to be restored. Application and backup systems are disabled and the contents of the replica are either restored directly to their original locations or presented to the system in place of contents of the source volumes. Because the restore is performed internally within the disk array, it runs at very high speed.

Once the restore is completed, the sections of the database or filesystem concerned are returned to their states at the time the replica was created and the application system can be re-enabled.

Depending on the application/database concerned, this may be all that is required. In some cases, additional action is required for full recovery, such as applying archived transaction log files that have been backed up separately.

For details, see “Instant recovery” (page 53).

Other ZDB data restore methods

Data backed up to tape can be restored using the standard Data Protector restore procedure.

For details, see the *HP Data Protector Concepts Guide*.

However, with specific disk array families, it is possible to first restore data from tape to update a replica and *then* restore the replica contents to their original locations. This is known as **split mirror restore**. Restoring the replica contents to their original locations is a similar process to instant recovery. It is only necessary to suspend application operation during this stage, minimizing the impact on the application.

For more details, see “Split mirror restore” (page 56).

NOTE: Replicas can be used for purposes other than instant recovery, such as data mining. Although Data Protector can create and administer replicas for such purposes, replicas created for the purpose of instant recovery should only be used for instant recovery. In the opposite case, loss of the backed up data may occur.

Restore possibilities for ZDB types

Table 4 ZDB types versus restore possibilities

	Restore possibilities		
ZDB forms and techniques	Individual objects	Disaster recovery	Instant recovery
ZDB to tape, local	Yes	Yes	No
ZDB to tape, remote	Yes	Yes	No
ZDB to tape, remote + local	Yes	Yes	No
ZDB to disk, local	No	No	Yes
ZDB to disk+tape, local	Yes	Yes	Yes


2 Replication techniques

Disk array basics

The replication techniques available depend on the type of disk array and the firmware/software installed.

Disk arrays support disk virtualization techniques, which enable the creation of virtual disks, logical volumes, and so on.

Figure 3 Disk virtualization



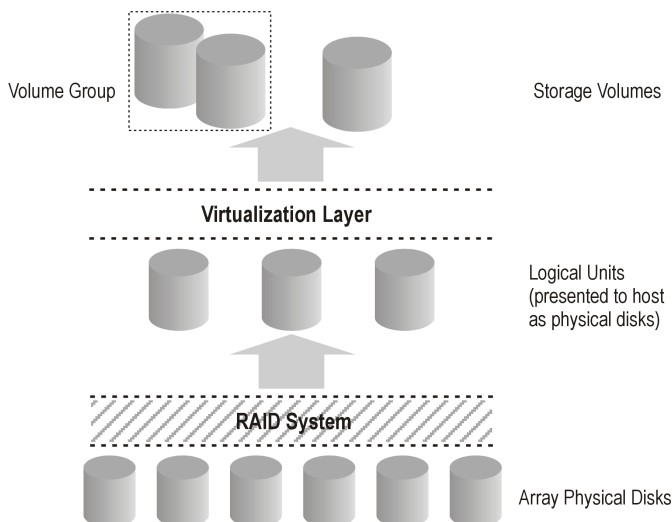
An array of physical disks is configured in such a way that it appears as one large block of data storage. This can then be divided into a number of virtual storage blocks, which are presented to the host/operating system.

These blocks can have different names, but basically the techniques for their production are similar and, for simplicity, in this guide, are considered as **storage volumes**.

RAID technology

Disk arrays use **RAID technology** which is applied to the available storage by the RAID system, to provide data redundancy and improved data protection.

Figure 4 Disk virtualization with RAID



Various RAID levels are available, providing different levels of data redundancy, speed, and access time. In some cases, it is possible to adjust the balance between these attributes according to the amount of free storage available.

RAID systems operate by distributing data across the physical disks and presenting them to the host as logical units, which, in turn, can be regarded as the physical disks considered in the

previous disk virtualization illustration. What are finally presented to the host operating system after virtualization are again virtual disks, or storage volumes.

Replication techniques

Basic replication can be performed in three contexts:

- Local (source and target volumes on the same disk array)
- Local – integrating with HP-UX LVM mirroring (source and target volumes on the same disk array, but at least two disk arrays are required)
- Remote (source and target volumes on different disk arrays)
- Remote plus local (remote plus local replication on the remote disk array)


From the operating system point of view, contents of particular source volumes and their replica are the same, whichever technique is used to produce the replica. However, the method used can affect such things as:

- the speed of replication
- the amount of storage space used
- the impact on the application involved
- data security

The following sections discuss methods of replication within each of these contexts.

Local replication

Local Disk Array



In **local replication**, data is replicated within the same disk array, that is, source and target volumes are both on the same disk array. There are two techniques:

- Split mirror
- Snapshot

Advantages of local replication

- The processes are fast.
- Disruption to the application or filesystem involved is minimized.
- All ZDB types (and therefore instant recovery) are supported, giving you flexibility in choosing your backup strategy.

Disadvantages

- Both source data and replicas are vulnerable to catastrophic failure of the disk array or the local system.


There are two styles of local replication:

- split mirror replication
- snapshot replication

Split mirror replication

In disk array terms, a **mirror**, is a dynamic copy of a source volume.

Figure 5 Split mirror replication



When a mirror is first created, data in it is synchronized until it is identical to that in the source volume. During normal application usage, the mirrors are kept synchronized with the source volumes. Any updates to the source volumes are also applied to the mirrors.

When a replica of the data at a fixed point in time is required for an administrative task (such as backup):

1. Synchronization between the mirrored volumes is stopped (the mirrors are split) leaving an independent replica of the source volumes.
2. The replica is used for the backup or other task, leaving the application to continue virtually unaffected using the source data.
3. If necessary, after the work on the replica is complete, the two sets of data can be resynchronized until mirrored data is required for another administrative task.

Splitting is very fast and has minimal impact on the application system.

Characteristics of split mirror replicas

- A split mirror replica is a complete duplicate (or clone) of the source volumes, which, from the point of view of the host/operating system, is identical to the source at the moment the duplicate was created.

At the physical disk or logical unit level, a complete physical copy of contents of the source storage blocks exists.

- It is completely independent of the original.

Because there is a separate physical copy of data, there is a higher likelihood that these target volumes will remain intact and available, if the disk array hardware experiences a partial failure that impacts the source volumes.

Snapshot replication

Snapshot replicas are created at a particular instant and are immediately available for use. Unlike split mirror replicas, no data is copied initially, but rather, a duplicate of the original storage is created through virtualization. At that moment, the replica is a virtual copy. The actual data is shared by both source and replica.

After that, when data in the source volumes is changed for the first time, the original data is first copied to the snapshot and then the source data is updated. Over time, the snapshot references partly its own independent data and partly shared data (in the form of pointers to unchanged source data). However, from the host or operating system's point of view, the snapshot always contains a full copy of the source volumes at the time it was created.

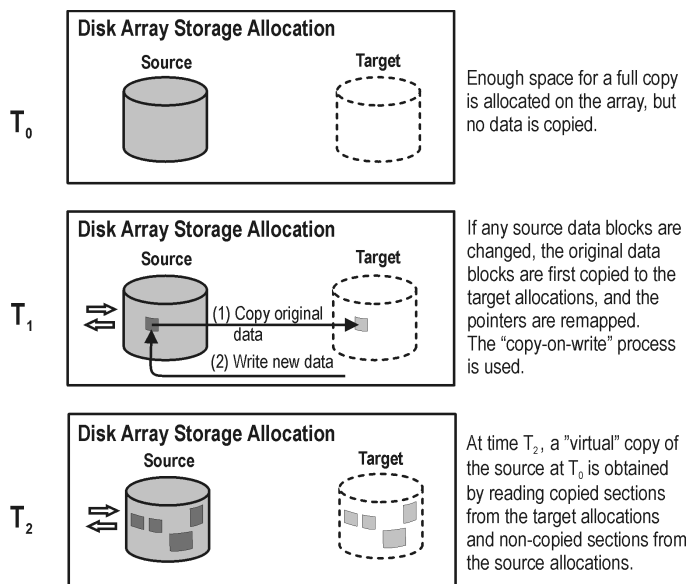
The supported integrations of arrays with Data Protector enable you to create the following types of snapshots:

- **Standard snapshot** (also known as “pre-allocated snapshot”, “fully-allocated snapshot”, or simply “snapshot”), where enough space is allocated when the snapshot is created to hold a full copy of all the source data.
- **Vsnap** (also known as “virtually capacity-free snapshot”, or “demand-allocated snapshot”), where no space is pre-allocated.
- **Snapclone**, which starts as a standard snapshot but where data is copied as a background task until the snapclone is a complete physical copy of the source volumes at the time it was created.

These are described below in more detail.

Standard snapshot

Figure 6 Creating a standard snapshot



1. At time T₀, storage capacity equal to that taken up by the source volumes concerned is allocated on the disk array for the target volumes.
No data is copied from the source storage blocks. Instead, pointers are mapped to the storage blocks holding the original data and the copy is completely virtual. From a host’s perspective, however, a complete replica of the source volumes at time T₀ exists in the target volumes and it is ready for use.
2. After snapshot creation, the first time T₀ source data needs to be updated, it is first copied to target storage blocks and pointers in the snapshot are remapped to these copies. Only then is the source data updated.
This is known as “copy-on-write”.
3. The snapshot is now partly real (where source data has been copied) and partly virtual. When the replica is accessed, any previously copied data is read from the target storage blocks and any data that has not been copied is read from the source storage blocks. From a host’s perspective, therefore, a complete replica of the source data at time T₀ still exists.

Characteristics of standard snapshots

- A standard snapshot is not an independent duplicate of the original data (it is however possible that in time, every single storage block in the source volume has been updated and therefore copied).
- Adequate space is guaranteed for the snapshot, even if all the data in the source volume changes.
- It is space-inefficient. Enough space is always reserved for all the data to be changed, though normally only part is used. While the snapshot exists, the rest of the reserved space cannot be used for any other purpose.

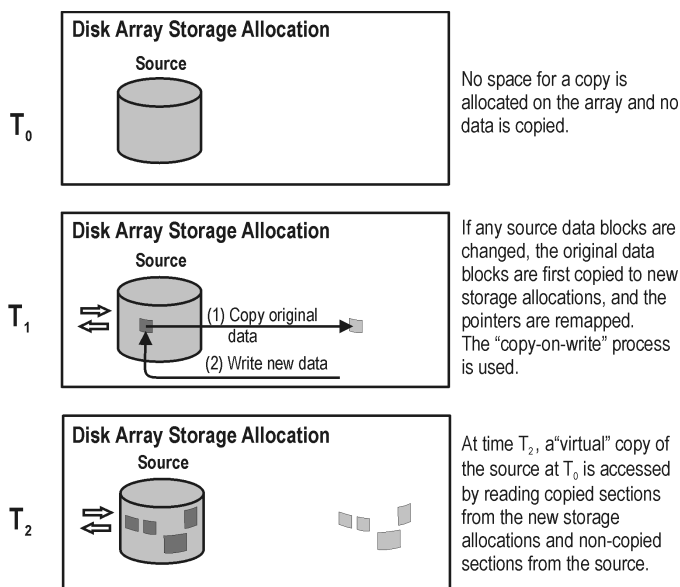
Impact on application performance

When a backup system accesses the snapshot, it reads disk blocks from both the source volumes and the replica. Consequently, both the application and the backup systems disk resources are used, which results in the application performance degradation when the disk array is excessively loaded.

Vsnap

With vsnap snapshot, no storage capacity is reserved at the start. Otherwise, the process is very similar to that for the standard snapshot:

Figure 7 Creating a vsnap



1. At time T_0 , only pointers are copied to the target, as for a standard snapshot, but no space is reserved for the target volumes. The snapshot takes up no storage space other than that required for the pointers.
2. After snapshot creation, the first time T_0 source data needs to be updated, "copy-on-write" is used, as in standard snapshots. Storage space is required only for the changed data.
3. As with standard snapshots, the snapshot is now partly real and partly virtual.

Characteristics of vsnaps

- Like the standard snapshot, a vsnap is not an independent duplicate of the original data.
- A vsnap requires independent disk capacity management to guarantee enough space for replica growth. If space on the disk array runs out, vsnap updates will fail, and it could affect general disk array operation.

- It is space-efficient. The vsnap only uses the space it needs.
- It is intended to be short-lived. Since the storage requirement for vsnaps is dynamic, the disk array may run out of space if there are many changes to the source volumes after the snapshots have been created. Other storage requests to a disk array can also cause the disk array to run out of storage.

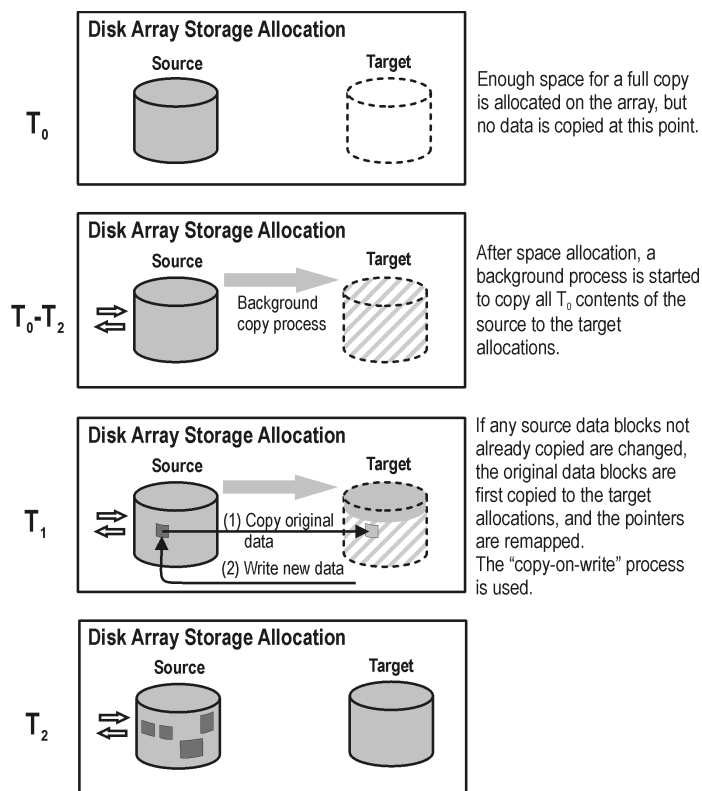
Impact on application performance

As with standard snapshots, when a backup system accesses the snapshot, it reads disk blocks from both the source volumes and the replica. Consequently both the application and the backup systems disk resources are used, which can result in the application performance degradation in cases where the disk array is excessively loaded.

Snapclone

Snapclone starts as a standard snapshot and ends up as a complete duplicate (or clone), similar to a split mirror replica.

Figure 8 Creating a snapclone



Data Protector snapclones are created in combination with a storage object called a **container** to speed up the snapclone creation process and reduce the impact on source volumes during copy of data. A container is the space on a disk array, which is pre-allocated for later use as a standard snapshot, vsnap, or snapclone. It can be either created from free disk space or converted from a storage volume that is no longer needed.

The process of a snapclone creation is as follows:

1. Containers of the same size and storage redundancy level as of the source volumes are created on the disk array if they do not exist yet.
2. The write cache policy on the source volumes is set to the write-through mode, so that all data in the cache is written to physical disks.
3. A standard snapshot is created, including allocation of enough space for a full copy.

4. A background process starts to copy all unchanged data from source storage blocks to target storage blocks. At this point, the write cache policy automatically reverts to write back mode.
5. If source data that has not already been copied by the background process needs to be updated, it is first copied (copy-on-write), as in a standard snapshot.
During execution of the background copy process, if the snapshot is required for use, the copy is partly virtual and partly real, as in a standard snapshot.
6. When all data has been copied to the target storage locations, the background process is stopped and a standalone duplicate, or clone, of the source at time T_0 remains.

Characteristics of snapclones (after copying finishes)

- A snapclone is a complete duplicate of the source volume, which, from the point of view of the host and operating system, is identical to the source at the moment the replica was created. At the physical disk, or logical unit level, a complete physical copy of contents of the source storage blocks exists.
- It is completely independent of the original.
Because the physical copy is complete, if the contents of the source volume are lost or corrupted, the contents of the target volume are not affected.
- It is intended to be long-lived.

Impact on application performance

- The background data copying process can affect application performance, through competition for resources. Copying can take a significant period of time when producing snapclones of large databases.
By using containers, the impact of the data copy process on the application performance is reduced. The time frame for which the application is required to stay in the backup mode is shortened significantly as well.
- If a system accesses a snapclone before the cloning process is finished, disk blocks not yet copied are read from the source volume. In the case of ZDB to tape or ZDB to disk+tape, data is read by using both application and backup systems disk resources; this can degrade the application's performance. To avoid this, Data Protector delays copying snapclone data to tape by up to 90 minutes if the cloning process is still in progress. This is the default; you can change it in the Data Protector GUI when configuring a backup specification.

Local replication integrating with HP-UX LVM mirroring

Local replication integrating with HP-UX LVM mirroring is a specific integration, which reduces the amount of storage which needs to be replicated in order to get a complete version. At the same time, LVM mirroring can be configured to provide functionality similar to that of HP Continuous Access (CA) or EMC Symmetrix Remote Data Facility (SRDF) in remote plus local replication environments on split mirror and snapshot arrays.


Advantages of local replication integrating with LVM mirroring

- Disk space usage is reduced by making a copy of part of the total disks used.
- It may be easier to set up and administer an LVM mirroring environment than a pure CA or SRDF environment.
- Costs for LVM mirroring environments are lower than for CA or SRDF environments because no CA/SRDF licenses are required. A BC license is only required on the system where the replicas are created.

Disadvantages

- A setup for LVM mirroring configurations can be more complex and has stricter requirements than that for BC or TimeFinder environments.
- LVM mirroring configurations introduce increased complexity in performing instant recovery. With specific disk arrays, instant recovery for the data backed up in LVM mirroring configurations is not supported.

Remote replication



In **remote replication**, data is replicated to a separate remote disk array from where it can be further backed up to locally available media. Once established, remote replication operations continue unattended, providing continuous, real-time remote data replication.

Advantages of remote replication

- Protects from catastrophic failure, such as loss of the storage system. If the remote disk array is located in a different (remote) computing center, remote replication also eliminates the risk of a fire or any other disaster damaging both the production and the backup environment at the same time, even in the event that an entire computing center is destroyed.
- Suitable for disaster recovery.
- Ensures continuous availability of important data.

Disadvantages

- Network and fibre channel connectivity transfer speeds increase the effect of replication on application or database performance.
- The need for synchronous transmission may affect application systems.
- At least two disk arrays are required, with associated licenses, increasing cost.
- The necessity for maintaining synchronization remotely can have an impact on performance and the application.

Split mirror replication


As with local mirroring, a duplicate of the source volumes is created and maintained on the target volumes, only in this case the target volumes are on a remote disk array. Once established, the target volumes are kept synchronized with the source volumes on the local disk array.

When a replica of the source volumes at a particular point in time is required, the synchronization between the mirrored volumes is stopped. The remote disk array then contains a fixed copy, or independent replica, of the source volumes on the local disk array.

However, if the arrays are installed at separate sites, the continuous remote synchronization may take place over several kilometers and this can impact performance on the application system. For Data Protector, the link to the remote system must usually be synchronous. With CA however, asynchronous communication is supported; Data Protector changes to synchronous mode for copying data to the mirror and then changes back to asynchronous.

You can choose this configuration for disaster recovery purposes (often in a cluster environment) where the potential benefits outweigh the disadvantages of maintaining the CA link. To break the link for backup purposes would reduce disaster recovery coverage and make failover impossible. Compare [“Remote plus local replication”](#) (page 29).

Remote plus local replication



Remote plus local replication uses both remote and local replication; replicas are created on a remote disk array using remote replication, and then used as the source volumes for a local replication.

This configuration is typically used if the remote site functions as a disaster recovery site and a split of the remote pairs is not possible. To automate failover, a cluster application can be used.

Advantages of remote plus local replication

As for remote replication, plus:

- Allows you to create tape backup without further affecting the application system or database.
- Maintains the possibility of automated failover.
- On P6000 EVA Array, you can influence the Data Protector behavior in case of a failover and choose to either follow the replication direction or maintain replica location.

Disadvantages

As for remote replication.

Split mirror replication

Remote part of replication

Mirrored volumes are set up with the source and target volumes on separate disk arrays as with remote replication.

Once established, the mirror volumes on the remote disk array are kept synchronized with the source volumes. For Data Protector, the link between the arrays must be synchronous.

Local part of replication

The target volumes of the remote replication stage become source volumes for local replication on the remote disk array.

When a replica is required, synchronization between the locally mirrored volumes is stopped (the mirror is split), but synchronization is still maintained between the remotely mirrored volumes. The local replica on the remote disk array (the replica of the replica) is then a fixed copy, or independent replica, of the source volumes on the local disk array.

Snapshot replication

Remote part of replication

The data is written from the application system to the source volumes on a local array, and is replicated to the target volumes on a remote disk array. Applications continue to run unaffected while data replication goes on in the background.

Local part of replication

The target volumes of the remote replication stage become source volumes for local replication on the remote disk array.

Snapshot replica volumes are created at a particular instant and are immediately available for use. For more information, see [“Snapshot replication” \(page 23\)](#).

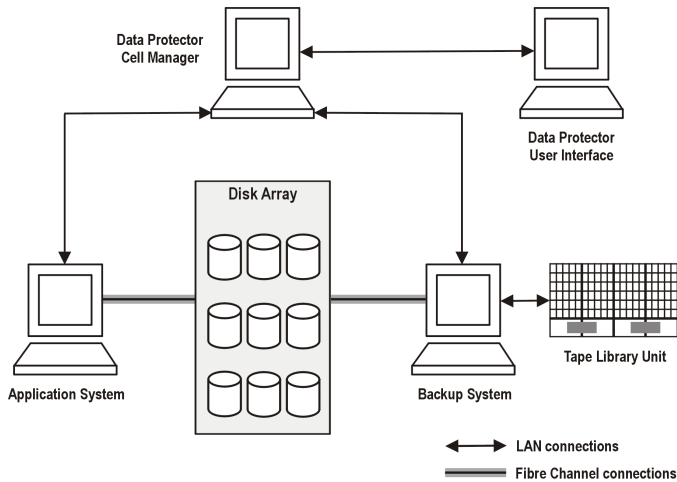
NOTE: Remote plus local replication provides a method for understanding and handling replica creation in non-failover and failover scenarios, thus enabling you to perform ZDB at either the source or destination site.

3 Using Data Protector for ZDB and instant recovery

Data Protector cells

Data Protector uses the concept of the **managed cell**. The following figure shows how a cell is set up for ZDB and IR purposes:

Figure 9 Data Protector cell set up for ZDB and IR




To be able to use ZDB and IR techniques, the application database or filesystem data to be backed up must be on a disk array to which the application and backup systems are both directly attached. The library or other storage device is optional for ZDB and IR applications.

Cell components

For a typical Data Protector cell, operational software components should be installed on the hardware as shown in the figure that follows.

Figure 10 Location of software components for ZDB and IR



Cell Manager

The Cell Manager is the main system of the cell. For information about the functions the Cell Manager performs in a Data Protector cell, how to access the Cell Manager, and coexistence of the Cell Manager with other Data Protector components, see the *HP Data Protector Concepts Guide*.

Application systems

Each application system for which replicas are to be created must have the following Data Protector components installed:

- A **disk array agent** or **ZDB agent**, which controls interaction between the Data Protector Cell Manager and the disk array on which the application database/filesystem is installed. Each supported type of disk array has its own dedicated agent.
- An **application integration agent**, which controls interaction between the Data Protector Cell Manager and the application. Data Protector requires the agent to perform functions such as controlling the state of the database during the backup and restore sessions for database applications. Without this agent, only filesystem backup is available.

Backup system

It is the system to which a replica is presented after it is created, so it is the system by which the replica can be accessed for subsequent processing, whether or not the data contained in it is to be backed up to media. The connection between the backup system and the application system is only used to coordinate processes that are involved in ZDB and IR sessions. The backup system can serve multiple application systems running different applications. It also performs various checks and administration functions.

The backup system must:

- Be capable of performing backup within a reasonable time.
- Have a relevant Data Protector ZDB agent installed. In some cases, it may also require an application integration agent.
- Use the same operating system version as the application system.

Generally, the backup system should not be the same as the application system.

ZDB database

The ZDB database is an extension to the Data Protector Internal Database (IDB) on the Cell Manager. It holds array-specific information about replicas needed for instant recovery purposes.

The ZDB database has a separate section for each disk array family that natively supports ZDB (and with most families also IR) within Data Protector:

- SMISDB for HP P6000 EVA Disk Array Family and HP 3PAR StoreServ Storage
- XPDB for HP P9000 XP Disk Array Family

Additionally, a separate section contains operating system information such as file system or volume management configurations:

- SYSDB

The exact information stored in the ZDB depends on the disk array. Generally speaking, each section contains the following types of information:

- Information on replicas kept on disk arrays, including:
 - Backup session ID
 - When the backup session was performed
 - Name of the backup specification used in the backup session

- Name, ID, and WWN of the target volume created in the session
- **HP P6000 EVA Disk Array Family:** Name and ID of the disk array unit on which the target volume resides
- **HP P6000 EVA Disk Array Family:** Information on the target volume type (standard snapshot, vsnap, or snapclone)
- Information about home (CA+BC configurations)
- ID of the source volume used in the backup session
- Whether the target volume can be used for instant recovery (IR flag)
- Whether the target volume can be deleted (purge flag)
- The application and backup systems involved in the session
- Disk array volumes excluded from the replica set rotation and other kinds of use.
- Additional configuration information:
 - **HP P6000 EVA Disk Array Family:** defined disk group pair relationships
 - **HP P9000 XP Disk Array Family:** detected P9000 XP Array command devices

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

The ZDB database stores information only about ZDB sessions that have the **Keep the replica after the backup** option selected in the backup specification. Replicas created in ZDB-to-tape sessions without this option selected are deleted from the ZDB database after the backup.

Information on ZDB-to-tape sessions and some information on ZDB-to-disk+tape sessions is also stored in other parts of the IDB.

The sections of the ZDB database and their use are fully described in the *HP Data Protector Zero Downtime Backup Administrator's Guide*.

User interfaces

You can use either the Data Protector graphical user interface (GUI) or command-line interface (CLI) to perform ZDB and IR operations.

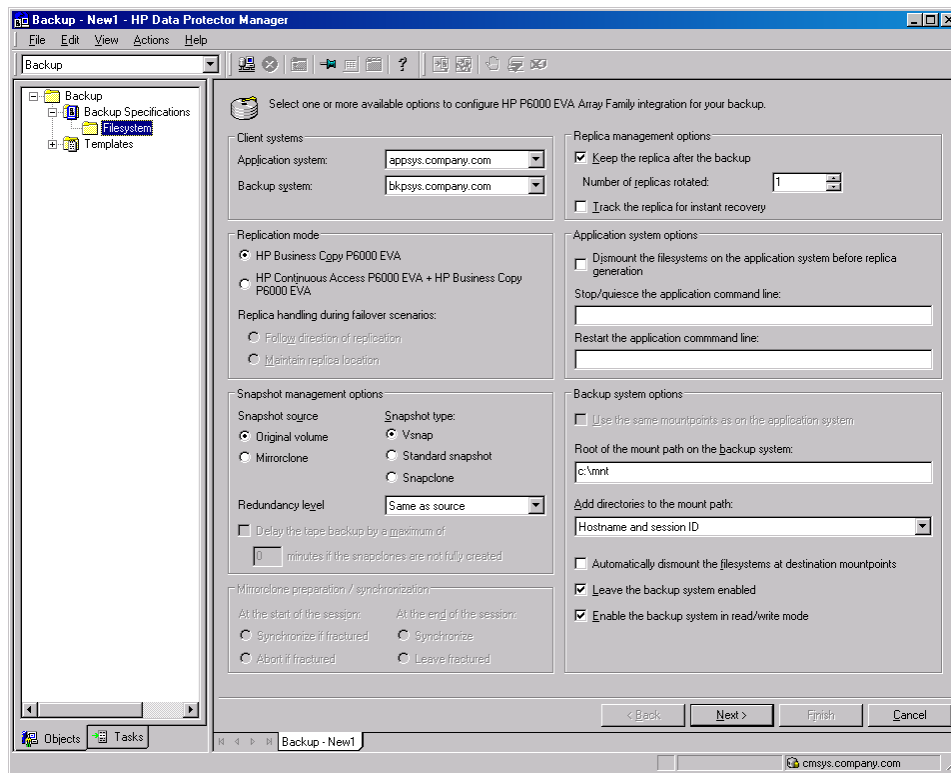
GUI

The GUI enables you to administer your ZDB environment from a single system. You can:

- Create backup specifications for ZDB, schedule them, and start ZDB sessions.
- Monitor active operations.
- Use Data Protector reporting and notification capabilities.
- In the **Instant Recovery** context, browse for sessions marked for instant recovery, define necessary options, and start an instant recovery session.
- In the **Restore** context, browse for sessions stored on a backup medium, define necessary options, and start the standard Data Protector restore procedure from tape.

The following is an example of the GUI window, where the backup options for a ZDB session running on P6000 EVA Array are selected.

Figure 11 Data Protector GUI



CLI

You can use the CLI to perform most ZDB and IR operations available in the GUI, but some administrative tasks can only be done using the CLI:

- Querying, synchronizing, and purging the ZDB database
- Checking the consistency of the ZDB database
- Manually deleting a replica or replica set when it is no longer needed, together with information on it stored in the ZDB database
- Excluding or including replicas from use with Data Protector.
- **HP P6000 EVA Disk Array Family:** Setting disk group pairs.

For details on available commands, see the *HP Data Protector Command Line Interface Reference*.

Disk array integrations available with Data Protector

Data Protector supports the following disk arrays capable of creating replicas and, in most cases, replica sets:

Table 5 Disk arrays integrating with Data Protector

Replica type	Supported disk arrays	Abbreviations
Split mirror	HP P9000 XP Disk Array Family	P9000 XP Array
	EMC Symmetrix Disk Array	EMC
Snapshot	HP P4000 SAN Solutions	P4000 SAN Solutions
	HP P6000 EVA Disk Array Family	P6000 EVA Array
	HP P9000 XP Disk Array Family	P9000 XP Array
	HP 3PAR StoreServ Storage	3PAR StoreServ

For the current list of configurations supported by HP, see <http://support.openview.hp.com/selfsolve/manuals>.

HP P4000 SAN Solutions

HP P4000 SAN Solutions support the creation of snapshots which use demand-allocated storage space and are based on the “redirect on write” technique. With this disk array family, Data Protector only supports local replication.

HP P6000 EVA Disk Array Family

The Data Protector P6000 EVA Array integration supports the creation of standard snapshots, vsnaps, and snapclones.

The following configurations are possible using the Data Protector P6000 EVA Array integration:

- Local replication
- Local replication integrating with LVM mirroring
- Remote plus local replication (giving the greatest level of data protection)

For further examples of P6000 EVA Array configurations, see “Supported configurations” (page 63).

P6000 EVA Array storage presentation

P6000 EVA Array uses virtualization technology, which organizes physical disks into **disk groups**. Each disk group is a storage pool from which **virtual disks** are allocated. A virtual disk is limited by the boundaries of a disk group, but may span over any number of physical disks within one disk group. You cannot control the exact allocation of virtual disks on physical disks, but you can influence it by choosing different protection characteristics. For that, RAID technology is used, which provides various levels of data redundancy, speed, and access time.

Local replication

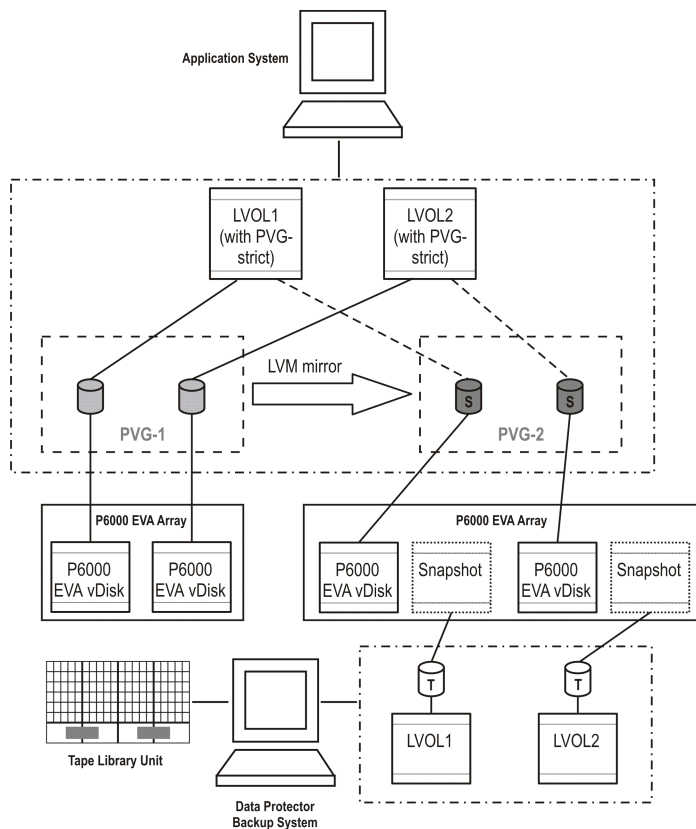
For local replication, the **HP Business Copy (BC) P6000 EVA configuration** is used. This enables you to create replicas that can be used for instant recovery purposes, regardless of the snapshot type used. Large replica sets can be created on the disk array. While the maximum number of replicas in a replica set consisting of standard snapshots and vsnaps is limited by the firmware revision of the P6000 EVA storage system, the maximum number of replicas in a replica set consisting of snapclones is limited only by the remaining storage capacity of the disk array.

Local replication integrating with LVM mirroring

The Data Protector P6000 EVA Array integration supports LVM mirroring in configurations where volume groups are LVM-mirrored from one P6000 EVA Array (or more P6000 EVA Array units) to another P6000 EVA Array (or other P6000 EVA Array units). The LVM-mirrored source volumes and their LVM mirrors belong to the same logical volume.

For this configuration, you need at least two disk arrays located in physically separate sites.

Figure 12 Example LVM mirroring configuration – P6000 EVA Array




Remote plus local replication

For remote plus local replication, a combination of HP BC P6000 EVA and HP **Continuous Access (CA)** P6000 EVA is used. This enables creation of snapshot replicas on a remote machine, and then creation of local replicas of those replicas on the remote machine.

For this configuration, you need at least two disk arrays located in physically separate sites.

Figure 13 Example HP CA+BC P6000 EVA configuration



HP P9000 XP Disk Array Family

The following configurations are possible using the Data Protector P9000 XP Array integration:

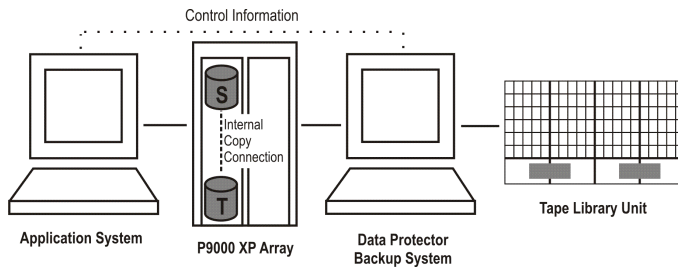
- Local replication
- Local replication – integrating with LVM mirroring
- Remote replication
- Remote plus local replication (giving the greatest level of data protection)

A separate backup system is connected to the disk array containing the target volumes, while the source volumes are connected to the application system. Data can be streamed to tape from the replica after the mirrors have been split or snapshots have been created, so that during the backup, the application system remains online and available for use.

Local replication

For local replication, the **HP Business Copy (BC) P9000 XP configuration** is used. This enables you to create either **first-level mirrors** or **volumes to be used for snapshot storage** for instant recovery purposes, in other words, a replica set.

Figure 14 Example HP BC P9000 XP configuration

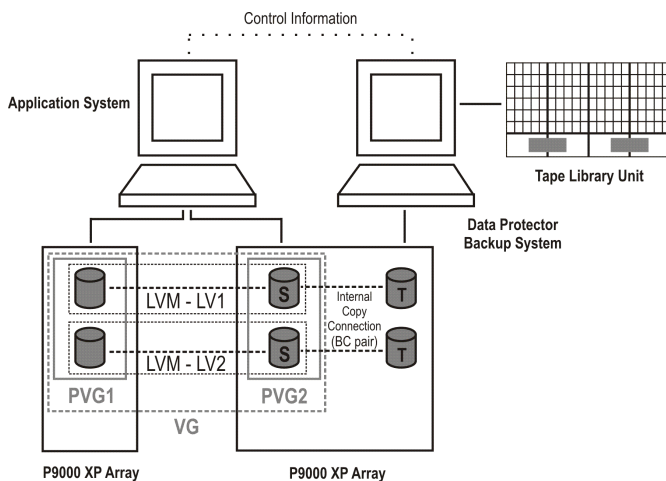


For further examples of P9000 XP Array configurations, see [“Supported HP P9000 XP Disk Array Family configurations” \(page 68\)](#).

Local replication integrating with LVM mirroring

The Data Protector P9000 XP Array integration supports HP-UX Logical Volume Manager mirroring (**LVM mirroring**) in configurations where one logical volume on one physical disk (LDEV) is mirrored onto a logical volume on another physical disk (LDEV).

Figure 15 Example LVM mirroring configuration – P9000 XP Array



Remote replication


For remote replication, the **HP Continuous Access (CA) P9000 XP configuration** is used. This enables you to create remote split mirror replicas on a remote system a considerable distance away.

The following two types of interfaces are supported for HP CA P9000 XP:

- Extended Serial Adapter (ESCON) for large distances
- Fibre Channel (FC) for distances up to 2 km

You can increase the Fibre Channel distance by using FC switches with built-in single-mode fibre multiplexors.

Figure 16 Example HP CA P9000 XP configuration




Remote plus local replication

For remote plus local replication, a **combination of HP CA P9000 XP and HP BC P9000 XP configurations** is used. This enables creation of split mirror replicas on a remote system, and then creation of local split mirror or snapshot replicas of those replicas on the remote system.

You need at least two disk arrays, located in physically separate sites.

When a replica is required, the integration splits the BC pair. To ensure data consistency, the CA pair status is checked before the BC pair split is executed. This ensures that all data from the Main Control Unit is in the Remote Control Unit.

Figure 17 HP CA P9000 XP configuration in a cluster



For more information about cluster configurations, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*.

HP 3PAR StoreServ Storage

HP 3PAR StoreServ Storage supports the creation of snapshots that use demand-allocated storage space and are based on the “copy on write” technique. With this disk array family, Data Protector only supports local replication.

EMC Symmetrix

The following configurations are possible using the Data Protector EMC integration:

- Local replication
- Local replication integrating with LVM mirroring
- Remote replication
- Remote plus local replication

The integration enables you to create single split mirror replicas that can be used for ZDB to tape and split mirror restore purposes.

NOTE: Instant recovery is not supported.


A separate backup system is connected to the disk array containing the target volumes, while the source volumes are connected to the application system. Data from the replica is streamed to tape after the pair has been split, so that during the backup, the application system remains online and available for use.

For further examples of EMC Symmetrix configurations, see [“Supported EMC Symmetrix configurations”](#) (page 74).

Local replication

For local replication, the **EMC Symmetrix TimeFinder configuration** is used.

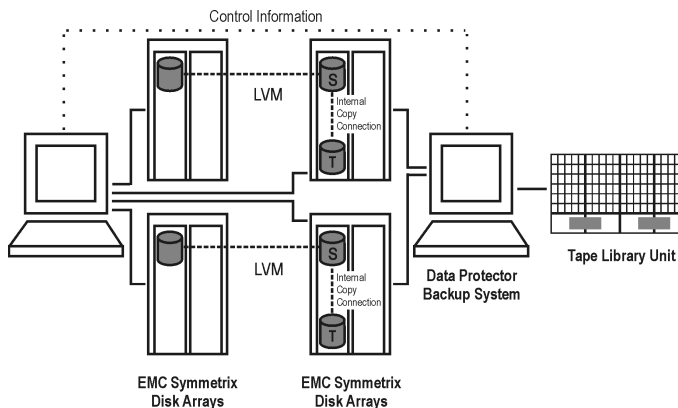
Figure 18 Example TimeFinder configuration



Local replication integrating with LVM mirroring

The Data Protector EMC integration supports LVM mirroring in configurations where one logical volume on one physical disk is mirrored onto a logical volume on another physical disk.

Figure 19 Example LVM mirroring configuration - EMC



Remote replication


For remote replication, the **EMC Symmetrix Remote Data Facility (SRDF) configuration** is used. This enables you to create split mirror replicas on a remote system.

Limitation

A cluster configuration is not supported in this environment.

At least two disk arrays, located in physically separate sites, are needed for such a configuration.

Figure 20 Example SRDF configuration



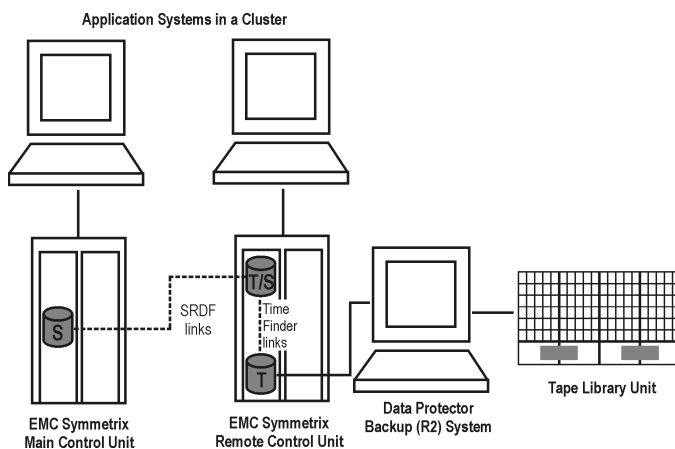
Remote plus local replication

For remote plus local replication, a combination of SRDF and TimeFinder configurations is used. This enables the creation of split mirror replicas on a remote system, and then creation of local replicas of those replicas on the remote system. At least two disk arrays, located in physically separate sites, are needed for such a configuration.

When a replica is required, the integration splits the TimeFinder pair. To ensure data consistency, the SRDF pair status is checked before the TimeFinder pair split is executed. This ensures that all data from the EMC Symmetrix Main Control Unit is in the EMC Symmetrix Remote Control Unit.

Typically, this configuration is used if the remote site functions as a disaster recovery site and a split of the SRDF pairs is not possible.

Figure 21 Example SRDF+TimeFinder configuration in a cluster



For more information about cluster configurations, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*.

Application integrations

Data Protector supports integration of supported disk arrays with the following database applications and replication types (online or offline):

- Oracle – online and offline backup
- SAP R/3 – online and offline backup
- Microsoft SQL Server – online backup
- Microsoft Exchange Server – filesystem-based offline backup

Microsoft SQL Server and Microsoft Exchange Server are also supported through the Data Protector Microsoft Volume Shadow Copy Service integration. For details, see the *HP Data Protector Zero Downtime Backup Integration Guide*.

For information on online and offline backup, see [“Freezing operation of the application or database”](#) (page 48).

All replication techniques (local, remote, remote plus local) are available for all database applications supported by Data Protector. However, not all application integrations are supported for all ZDB agents or their platforms. For details, see the latest support matrices at <http://support.openview.hp.com/selfsolve/manuals>.

Application data consistency

A simple ZDB of logical volumes or disks guarantees only filesystem consistency, but not application data consistency. After an instant recovery of such a backup, the database may not recover properly. For supported integrations, Data Protector ensures that the application is set to the backup mode (online backup – for the duration of which the set of operations the application can execute may be reduced) or shut down (offline backup), but you must back up transaction logs separately. For non-integrated applications, you must ensure that the backup is usable for database recovery. You can either shut down the application or set it to an appropriate mode by using pre-exec scripts.

Transaction logs

When backing up database applications online, you need to back up separately any archived database transaction logs in order to be able to perform a complete database recovery. The transaction logs should not be backed up in the same zero downtime backup session as the rest of the database data.

You can only back up the archived database transaction logs to disk or tape by running a separate ordinary Data Protector backup session after the ZDB session. The script that starts the backup session can be specified in the **Post-exec** option in the Data Protector ZDB backup specification. This way, backup of the transaction logs is started automatically after the replica creation completes.

Restore

For details of restore methods available with supported database applications, see the latest support matrices at <http://support.openview.hp.com/selfsolve/manuals>.

With instant recovery, you can recover a database to the point in time at which the replica was created. In most cases however, to fully recover the database, the transaction logs must be applied afterwards. Using these logs, you can also roll forward the database to a certain point in time. A shorter ZDB backup window results in a smaller amount of archived log file data that need to be applied during a full database recovery.

For detailed instructions on how to use the Data Protector disk array integrations with the database applications, see the *HP Data Protector Zero Downtime Backup Integration Guide*.

Application integrations and Microsoft Volume Shadow Copy Service

In the traditional backup model, the backup application coordinates various systems and components involved in the backup process: the application and backup systems, and the disk array. This is the case with the Data Protector HP P9000 XP Disk Array Family and HP P6000 EVA Disk Array Family integrations, where HP P9000 XP Agent and HP P6000 / HP 3PAR SMI-S Agent control the disk array and the Data Protector integrations interact with the database applications.

On Windows systems, a unified backup and restore service—the Microsoft Volume Shadow Copy Service (**VSS**)—coordinates the components involved in the backup process. The VSS model provides a standardized interface to the applications (**writers**) and disk arrays (**providers**).

The writers interact with the applications, providing a list of items that can be backed up. Data integrity is provided by the writers on the operating system and application levels.

The hardware providers replace the disk array agent functionality and behave from the Data Protector point of view similarly to disk array agents.

When performing an instant recovery of the data that were backed up in a zero downtime backup session with the Data Protector Microsoft Volume Shadow Copy Service integration, you can select

to use the Microsoft Virtual Disk Service or the disk array agent. The selection also depends on the way the backup was made.

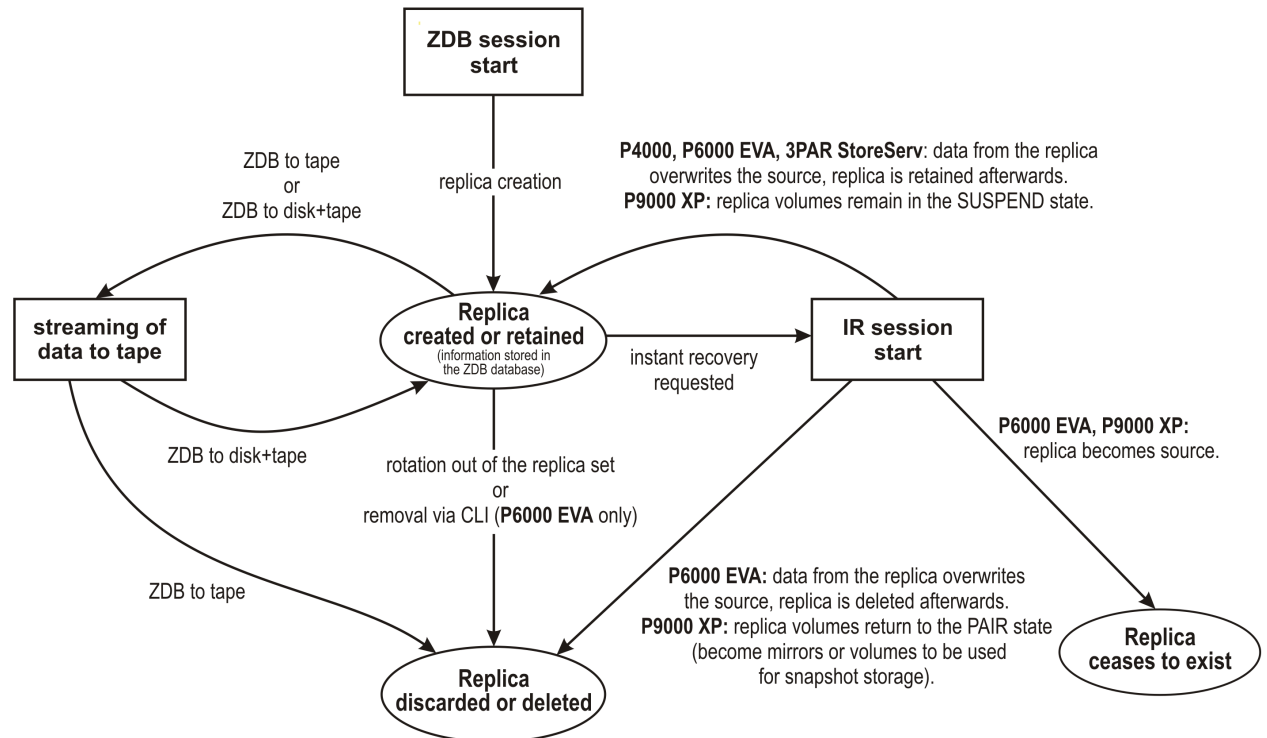
For detailed instructions on how to use the Data Protector Microsoft Volume Shadow Copy Service integration, see the *HP Data Protector Zero Downtime Backup Integration Guide*.

4 Replica life cycle

Overview

This chapter describes the life cycle of replicas, summarized in the following diagram.

Figure 22 Replica life cycle



A replica life cycle depends on the following:

- the disk array model
- the Data Protector components involved in the ZDB and IR sessions
- the options selected for the zero downtime backup session
- the instant recovery method which is either selected from the available ones or enforced by specific replica types
- other options selected for the instant recovery session

Creating replicas

With both split mirror and snapshot replication techniques, the basic idea is the same: to produce copies or images of the storage volumes (source volumes) containing the specified data objects. These copies are created in other storage volumes (target volumes) on the same disk array, and which can then be presented to a host system.

In all cases, only complete source volumes on the disk array can be replicated. Even if the data selected for replication only take up a small part of a source volume, the full source volume is replicated.

ZDB sessions that create replicas are defined by **backup specifications**, which contain all the information required to run a ZDB session:

- The type of application or filesystem data to be backed up
- The source data to be backed up

- The type of replica (or replica set – see “[Replica set rotation](#)” (page 44)) to be created
- The type of disk array on which the data resides
- The application and backup systems to be used
- Replica management and replica mounting options

For applications not fully integrated with Data Protector, you can also set options to stop the application before replication and restart it afterwards.

After you have created a backup specification, it is stored on the Cell Manager and can be reviewed or updated at any time.

A backup session can then be started interactively by an operator using the Data Protector user interface, or scheduled to start automatically at specified times.

NOTE: With some database applications, when an online backup session is run, it is also necessary to back up the log file currently in use by the database. This is done by backing up the log to a file, which can then be streamed to tape if required.

It is generally *not* recommended to include the log file in the volumes to be replicated. With some integration agents, this is not allowed. With others, it reduces or limits some restore scenarios.

After successful backup, details of the backup session are saved to the ZDB part of the IDB.

Replica sets

A **replica set** is a collection of replicas created at different times using the same backup specification. Replica sets are normally used when creating replicas for instant recovery purposes. The maximum number of replicas that you can define for a replica set depends on one or more of the following factors: replica type, the disk array model, the installed disk array firmware revision, snapshot type used for the target volumes (only with snapshot replicas).

In Data Protector, the members of a set can undergo **replica set rotation**, either interactively or at times specified in the scheduler. Note that specific disk array models do not support replica set rotation.

Replica set rotation

When you create a backup specification for ZDB and instant recovery purposes, you need to specify the maximum number of replicas in the replica set. Each time the backup is run, a new replica is created and added to the set. When the specified maximum number of replicas is reached, the next replica to be created replaces the oldest replica in the set. With some replica types, this is achieved by directly overwriting the oldest replica, in other cases, the oldest replica must be deleted before the new replica is created.

Scheduling replication

If you want replication sessions to be run automatically, enter details of required times into the Data Protector **scheduler** when creating or modifying the backup specification. You can either schedule a single session at a specific time, or regular sessions, repeated over periods of days, weeks, or months.

Using replicas

Once you have created replicas or replica sets, what happens to them depends on the form of ZDB used:

- **ZDB to tape:** Stream the data in the replica to tape. After that the replica is discarded.
- **ZDB to disk:** Keep it on the disk array for instant recovery purposes.
- **ZDB to disk+tape:** Stream the data in the replica to tape and keep it on the disk array for instant recovery purposes.

After ZDB to disk and ZDB to disk+tape sessions, one or more replicas can be kept on an disk array. You can use replica set rotation to maintain a set of replicas created at different times using the same backup specification, where each new replica replaces the oldest replica in the set. Each replica continues to exist until it is either rotated out of the replica set, you delete it using the Data Protector CLI, or it is “consumed” in sessions using a specific instant recovery method.

ZDB to tape

With ZDB to tape, a replica is normally only kept on an disk array temporarily. It effectively allows a staged backup-to-tape process.

After creation, the replica is mounted on the backup system and backup objects specified in the backup specification are streamed to tape (or other backup medium).

After the backup is complete, the replica is no longer required for backup purposes, so by default it is automatically deleted from the disk array. You can, however, opt to keep the replica on the disk array to reserve space on the disk array for future ZDB-to-tape sessions using the same backup specification. This way, you guarantee there is enough space on the disk array for your backup.



IMPORTANT: The replica is *not* available for instant recovery.

Advantages	Disadvantages
Suitable for backup and disaster recovery.	For disaster recovery, restore of a complete session for a large database would take a very long time for a high-availability system.
Individual data objects can be restored from the tape backup.	
The replica is by default deleted from the disk array, freeing the space.	Instant recovery is not possible.
Extensive tape library support.	

ZDB to disk

With ZDB to disk, the replica is kept on the disk array and used as the backup image for instant recovery purposes.

One or more replicas can be kept on a disk array. You can use replica set rotation to maintain a set of replicas created at different times, where each new replica replaces the oldest replica in the set.

Advantages	Disadvantages
Suitable for backup and instant recovery.	Disk space is permanently required for replicas.
	Limited disk array support compared with ZDB to tape.

ZDB to disk+tape

ZDB to disk+tape is basically a combination of ZDB to disk and ZDB to tape.

A replica is created on disk, exactly as in ZDB to disk, and then the replica is streamed to tape other backup medium. The disk replica is retained and, unlike in ZDB to tape, *can* be used for instant recovery.

Replication method/disk array support is the same as for ZDB to disk.

It is possible to specify ZDB-to-disk+tape sessions in the same schedule as ZDB-to-disk sessions, using the same backup specification. This means you can set up more sophisticated backup arrangements, such as performing ZDB to disk for six days per week and ZDB to disk+tape for the

seventh day, using the same backup specification. This enables greater flexibility for restore. Note that the same replica set will be used for both types of session.

Advantages	Disadvantages
Suitable for backup and instant recovery.	Disk space is permanently required for replicas.
Individual data objects can be restored from the tape backup.	Limited disk array support compared with ZDB to tape.
Sophisticated combinations of backup using ZDB to disk and ZDB to disk+tape are possible.	
Replica set rotation is available, even for tape.	

Instant recovery

Using a replica created in a ZDB to disk or ZDB to disk+tape session, instant recovery enables you to restore data objects to their states at a particular point in time. For details of the process, see [“Instant recovery” \(page 53\)](#).

What happens to the replica after an instant recovery session depends on the disk array model, the selected available instant recovery method, and other options selected (GUI) or specified (CLI) for the instant recovery session:

- With HP P4000 SAN Solutions, the replica data is copied back to the source volumes, and the replica is retained on the disk array. However, for each target volume that is selected for instant recovery, if newer target volumes exist for the same source volume, they are removed from the disk array automatically, regardless of the replica set they belong to.
- With HP P6000 EVA Disk Array Family:
 - By switching the disks, the replica ceases to exist as a replica.
 - By copying the replica data back to the source volumes, the replica is retained on the disk array.
- With HP P9000 XP Disk Array Family:
 - By switching the disks (in case of split-mirror replicas), the replica ceases to exist as a replica.
 - By resynchronizing the source volumes (in case of split-mirror replicas) or restoring the data from the replica to the source volumes (in case of snapshot replicas):
 - If only the Data Protector HP P9000 XP Agent is used, the replica can be retained on the disk array or not, depending on the options selected (GUI) or specified (CLI) for the instant recovery session.
 - If the Data Protector MS Volume Shadow Copy Integration and the Data Protector HP P9000 XP Agent are used, the replica is retained on the disk array
- With HP 3PAR StoreServ Storage (in configurations where instant recovery is available), the replica data is copied back to the source volumes, and the replica is retained on the disk array.

Deleting replicas

Replicas can be deleted automatically or manually:

- *Automatically:*
 - When a replica becomes the oldest member of a replica set, it is automatically overwritten (or deleted) when a new replica is created in the set.
You can however exclude replicas from use to protect them. For more information, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*.
 - If a replica is used for a ZDB-to-tape session, it is automatically deleted after the session unless you explicitly specify that it should be retained.
 - A replica is deleted after instant recovery if the instant recovery options are configured so.
 - A replica can no longer be used as a replica after a session using a specific instant recovery method: with HP P9000 XP Disk Array Family or HP P6000 EVA Disk Array Family, and with the instant recovery method of switching the disks, the replica ceases to exist as a replica after it becomes the recovered source.
 - With HP P4000 SAN Solutions, a target volume (not the entire replica) is automatically removed from the disk array when an older target volume created for the same source volume is used for instant recovery. Such a target volume is removed even if it belongs to another replica set. The automatic removal of newer target volumes is invoked by the disk array itself.
- *Manually:*

When replicas are no longer required to be used by Data Protector, you can delete them from the disk array using the Data Protector CLI.

5 ZDB session process

ZDB process overview

With conventional Data Protector backup, application operation is affected for the whole of the backup session, until streaming of the data to backup medium is complete. With Data Protector zero downtime backup, application operation is only affected during the creation of a replica.

The principal steps in a ZDB process are:

1. Locate the data objects for backup. See [“Locating data objects” \(page 48\)](#).
2. Freeze operation of the application database. See [“Freezing operation of the application or database” \(page 48\)](#)
3. Create a replica containing the specified data objects. See [“Creating a replica” \(page 49\)](#).
4. If backup to tape is required, stream the replica to tape. See [“Streaming the replica to tape” \(page 49\)](#).
5. If the ability to perform instant recovery is required, record information about the session. See [“Recording session information” \(page 50\)](#).

Locating data objects

Data that will be backed up is located and prepared as follows:

1. Data Protector starts processes on the application and backup systems.
2. The Backup Session Manager reads the backup specification for ZDB and passes the necessary instructions to the application integration agent and the disk array agent on the application system, and to the disk array agent on the backup system.

The ZDB agent on the application system resolves data objects to filesystems (if any), volume groups (if any), and the underlying storage volumes. These data objects may come directly from a backup specification or may be provided by one of supported application integrations.

For details, see the *HP Data Protector Concepts Guide*.

3. The application system is prepared, bringing data into a consistent state. For online backup, the database is quiesced. For offline backup, the database is taken offline. If the ZDB option **Dismount the filesystems on the application system before replica generation** (HP P6000 EVA Disk Array Family, HP 3PAR StoreServ Storage) or **Dismount the filesystems on the application system** (HP P9000 XP Disk Array Family) is selected, the filesystems involved are dismounted.

Freezing operation of the application or database

While a replica is being created, operation of the application or section of the database concerned must be frozen.

The application integration agent puts the application database or filesystem into the required state. This could be with all database updates stopped for an offline replication, or with all database updates re-routed to log files in the case of an online replication:

- In **offline** replication, the database is taken offline, so that all file I/O is stopped while the replica is created. The database is usually placed into a consistent state, for instance by applying any previously unapplied redo logs.

Although creating a replica is very fast, the application is offline for a short time, so this method is less suitable for high availability applications.

- In **online** replication, the database is placed into **hot-backup mode** while the replica is created. In this mode, the database remains online, but all database I/O is diverted to transaction log

files instead of updating the database. After the replica is made, the transaction log files are applied to the database to bring it up-to-date.

This method of replication reduces impact on the application to a minimum, making it suitable for uninterrupted operations.

The steps concerned in these operations can be controlled automatically when backing up database applications supported by Data Protector. However, it is also possible to set up similar behavior when backing up other applications or filesystems; pre- and post-exec options enable you to specify scripts to run before and after replication.

In both cases, the effect of the backup process on the application is limited to the period during which the replica is created. In the “online” case, database operation is never stopped (zero downtime) and the effect on performance is minimal, limited mainly to the effect of having to write increased information to the transaction logs.

Both online and offline backup are also available within Data Protector without using ZDB replication techniques. However, there is a much greater impact on application/database operation since with conventional backup to tape, a database has to be put into hot-backup mode or taken offline for the whole of a backup session.

Creating a replica

1. A replica is created.
2. The application system is resumed. Any dismounted filesystem is remounted.
In the case of an offline backup, the database can be brought back online and normal operation started again.
In the case of an online backup, transaction log files and cached information from the replica creation period are applied to the database.
3. The backup system environment is prepared for the replica’s disks and data. New devices are detected by scanning. Any volume groups are imported and activated. Filesystems are mounted.

Replicating the data objects

With the database/filesystem in the required state, the disk array agents on the application system and the backup system are triggered to perform the replication.

The two disk array agents act as a pair:

- On the application system, the agent resolves the specified data to the volumes containing it.
- On the backup system, the agent allocates the volumes required for the replica.

The disk array then creates the replica on its disks.

The replication method depends on the type of disk array being used, whether the disk array is configured for local or remote replication, whether LVM mirroring is required or not, and so on. For information on how split mirror and snapshot replication is performed, see [“Replication techniques”](#) (page 21).

Streaming the replica to tape

1. In ZDB to tape and ZDB to disk+tape, the replica is streamed to tape.
2. The backup system is cleaned. Filesystems are dismounted. New volume management systems are deactivated and removed.

Backing up a replica to tape

Creating mount points

Before data can be moved from the replica to tape or other backup medium, the replica must first be mounted on the backup system.

Data Protector creates mount points on the backup system and mounts filesystems in the replica to them. The process depends on whether an application, disk image, or filesystem backup is being performed.

Standard data movement to tape

As specified in the backup specification, data objects are streamed to tape using the Data Protector Media Agent.

Data Protector writes the information to tape as though the data objects are coming from their original locations, rather than the replica, so that the session information on tape and in the IDB are as if a conventional backup to tape has been performed. This means that data objects from ZDB-to-tape and ZDB-to-disk+tape sessions can be restored directly to the application system, using the standard restore procedure.

Incremental ZDB

Incremental ZDB is a *filesystem* ZDB to tape or ZDB to disk+tape session in which Data Protector streams to tape only files that fit the incremental backup criteria, the same criteria that are used for incremental non-ZDB sessions. Note that the replica is created in the same way for both full and incremental ZDB sessions.

The replica after creation

After the replica is created:

- With *ZDB to disk* and *ZDB to disk+tape*, the replica remains on the disk array for instant recovery purposes. If it is part of a replica set, it remains on the disk array until it is the oldest replica in the set. After that, it is replaced by the replica created in the first next ZDB-to-disk or ZDB-to-disk+tape session performed using the same backup specification (except if it is excluded from use).
- After a *ZDB to tape* session, when the data has been backed up to tape, the replica is automatically deleted by default. You can opt to keep the replica on the disk array, but it cannot be used for instant recovery.

For information on ZDB options, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*.

Mounting the replica on the backup system

Data Protector creates mount points on the backup system and mounts the filesystems of the replica to them. The mount point paths depend on whether an application or filesystem backup is being performed and the backup specification options you select in the GUI. You can also choose to leave the filesystems mounted on the mount point paths after the ZDB session completes.

With the VSS integration, the backup specification options selected in the GUI determine whether mount points are created on the backup system, and if the replica filesystems are mounted to the mount point paths in the read-write or read-only mode.

Recording session information

At this stage, created replicas can be recycled for the next session. If instant recovery has been enabled, additional IR session information is stored in the IDB, and the replicas retained in case IR is required.

Writing session information to the IDB

As with a conventional Data Protector backup, ZDB session information is written to the IDB throughout the session, including information on the backup medium and data objects available for restore.

- For *ZDB to disk* and *ZDB to disk+tape*, disk array-specific information about the replica is also written to the ZDB database for instant recovery purposes.
- For *ZDB to tape*, no instant recovery information is recorded in the ZDB database even if the replica is kept on the disk array after a backup.

The **ZDB database** is an extension of the IDB on the Cell Manager. It has separate sections for each disk array that natively supports ZDB and IR within Data Protector:

- SMISDB for disk arrays of the HP P6000 EVA Disk Array Family and the HP 3PAR StoreServ Storage family
- XPDB for disk arrays of the HP P9000 XP Disk Array Family

Information is written to the ZDB database whenever a replica is created, and deleted when the replica is deleted.

For details on the sections of the ZDB database and their use, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*.

6 Instant recovery and other restore techniques from ZDB sessions

Overview

With instant recovery, you restore complete replicas, at high speed, with minimum impact on the application system. All volumes containing the data objects specified in the backup specification are returned to their states at a specific point in time.

After a ZDB session, you can view the associated restore objects and restore sessions in the following GUI contexts:

- After ZDB to tape or ZDB to disk+tape, in the **Restore** context, enabling restore of data objects from tape.
- After ZDB to disk or ZDB to disk+tape, in the **Instant Recovery** context, enabling restore from replicas.

Alternatively, you can use the Data Protector CLI.

The restore methods depend on the type of ZDB session performed and the type of disk array being used. The available methods are described in the sections that follow.

Instant recovery

Availability

In local replications:

- from ZDB to disk
- from ZDB to disk+tape

NOTE: Instant recovery is not supported on EMC arrays; for them, only ZDB to tape is possible.

Features

You can restore complete replicas, at high speed, with minimum impact on the application system. All volumes containing the data objects specified in the backup specification are returned to their states at a specific point in time.

More information

See “Instant recovery” (page 19).

Because of the different types of replicas involved and various disk array limitations, the detailed restore process is different for each disk array type. For more information, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*.

Standard Data Protector restore

Availability

In local and remote replication:

- from ZDB to tape
- from ZDB to disk+tape

Features

You can restore individual backup objects directly from tape to the application system.

What is available for standard restore depends on what is actually streamed to tape. This, in turn, depends on how the ZDB-to-tape or ZDB-to-disk+tape backup specification is created. If the complete contents of the source volumes are selected in the backup specification, all objects will be streamed to tape. If not, only the selected backup objects will be streamed to tape, even though the whole of the source volumes are replicated.

More information

See the *HP Data Protector Help* index: “standard restore procedure”.

Split mirror restore

NOTE: With the speed of contemporary SAN-attached, very fast tape drives, it is usually quicker to restore directly to the application system than use split mirror restore.

Availability

In local replications on specific disk array models:

- from ZDB to tape
- from ZDB to disk+tape

Available for disk image, filesystem, and filesystem-based application backups.

Features

You can potentially restore anything from an individual backup object to the whole contents of the replica, with minimum impact on the application system. Split mirror restore can be used to perform a low impact restore for a system that is partially corrupted, but still operational.

What is available for split mirror restore depends on what is actually streamed to tape, as for standard restore above.

More information

See “Split mirror restore” (page 56).

Instant recovery

With instant recovery, lost or corrupt data is replaced with known good data, which was previously replicated to other volumes on a disk array. This previously replicated data is handled on the complete storage volume level. The remainder of the process depends on the application being recovered:

- Where a *filesystem* has been replicated, this step is all that is required to return the data to its state at the moment the replica was created.
- For a *database application*, you may need to perform additional operations to fully recover the database after performing instant recovery, such as restoring and applying transaction log files. In this way, you may be able to recover the database to a later point in time than that at which the replica was created, if log files for that time exist (commonly known as **roll forward**). This usually involves the use of another backup medium or device. For more information, see the *HP Data Protector Zero Downtime Backup Integration Guide*.

During instant recovery, either target volumes are presented to the system in place of source volumes (this instant recovery method is only available with snapclones) or a data copy operation is performed in which data located on the source volume is replaced by data located on the target volumes. These operations are performed internally within the disk array, involving no other backup medium or device. This makes instant recovery very fast.

In instant recovery sessions that only use Data Protector disk array agents, you cannot define which backup objects specified in the backup specification should be restored; only a complete backup object set can be selected for instant recovery and, hence, only the complete replica can be

restored. Additionally, on UNIX systems with configured LVM, not only the volumes constituting the replica are restored, but entire volume groups in which these volumes reside are returned to the state they were in when the replica was created.

In instant recovery sessions that use the Data Protector Microsoft Volume Shadow Copy Service integration, you can individually select backup objects specified in the backup specification for instant recovery, as long as all backup objects stored on each individual volume to be involved in the instant recovery session are selected. Only the volumes containing the objects selected for instant recovery are restored, and other volumes of the same volume groups are left intact.


Replicas cannot be displayed or selected directly in the Data Protector GUI, but the sessions that created replicas available for instant recovery can.

Because of the different types of replicas involved and various disk array limitations, the restore process details are different for each disk array type and also depend on the involvement of Data Protector Microsoft Volume Shadow Copy Service integration. For more information on the HP P4000 SAN Solutions, HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, and HP 3PAR StoreServ Storage, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*. For more information on the Data Protector Microsoft Volume Shadow Copy Service integration, see the *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

Instant recovery process

The following is an example of instant recovery:

Figure 23 Instant recovery example



1. Decide which replica you want to restore and select the ZDB session that created it.
2. Select the instant recovery options, which are primarily provided for selecting the instant recovery method and the data safety level.

Depending on the operating system, selected instant recovery method, and disk array model, these options enable you to:

- **UNIX systems with configured LVM:** Check if configurations of the volume groups involved in the instant recovery have not changed since the replica to be restored was created. This check also verifies that CRCs performed on the data in the replica to be restored match those produced when the replica was created.
 - With specific instant recovery methods, keep the replica on the disk array after the instant recovery session – for situations with potential problems with any step after data restore.
 - **HP P6000 EVA Disk Array Family:** Remove the presentation of replica volumes to systems other than the backup system.
3. Optionally, perform a preview of the instant recovery session to provide an extra level of security.

NOTE: Instant recovery preview is not available in instant recovery sessions that use the Data Protector Microsoft Volume Shadow Copy Service integration.

4. Start the instant recovery.

Data Protector then:

1. Starts processes on the application system and the backup system.
2. Extracts the session information from the IDB and the array-specific information associated with the session from the ZDB database.
3. Performs the necessary checks to verify that all required conditions for a successful instant recovery are met (including any instant recovery options specified).
4. Prepares the application system by deactivating any volume groups (on UNIX systems with configured LVM) and dismounts any filesystems associated with the replica.
5. Restores the original data.

Depending on the disk array model, the instant recovery method – either selected from the available ones or enforced by specific replica types, and other options selected for the instant recovery session, the following instant recovery methods are available:

- With HP P4000 SAN Solutions, only one instant recovery method is possible:
 - Copying the replica data back to the source volumes
Data from the replica is copied back to the original storage, and the source volumes are not retained. The replica is retained, but if newer replicas than the one selected for instant recovery exist in the replica set, they are removed from the disk array.
For this method, both the Data Protector Microsoft Volume Shadow Copy Service integration and the Data Protector HP P4000 Agent are used.
- With HP P6000 EVA Disk Array Family, two instant recovery methods are possible:
 - Switching the disks
The selected snapclone replica is substituted with the original source volumes. Any host presentations that were created for the original source volumes are then created for the restored snapclone volumes which effectively become the new source volumes. As far as Data Protector is concerned, the snapclone replica is deleted from the associated replica set. Another instant recovery is not possible. The old source volumes can be retained or removed.
For this method, depending on the Data Protector components involved in the zero downtime backup session, either only the Data Protector HP P6000 / HP 3PAR SMI-S Agent is used or the Data Protector Microsoft Volume Shadow Copy Service integration and the Microsoft Virtual Disk Service are used.
 - Copying the replica data back to the source volumes
Data from the replica is copied back to the original storage. You can retain the source volumes or not.
If you choose to retain the source volumes, 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 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, and finally the source volumes are switched with the containers.If you choose not to retain the source volumes, the data from the existing replica is restored to the source volumes without prior operations.
For this method, depending on the Data Protector components involved in the zero downtime backup session, either only the Data Protector HP P6000 / HP 3PAR SMI-S

Agent is used or the Data Protector Microsoft Volume Shadow Copy Service integration and the Data Protector HP P6000 / HP 3PAR SMI-S Agent are used.

- With HP P9000 XP Disk Array Family, two instant recovery methods are possible:
 - Switching the disks:

The selected split mirror replica is substituted with the original source volumes. Any host presentations that were created for the original source volumes are then created for the restored replica volumes which effectively become the new source volumes. As far as Data Protector is concerned, the replica is deleted from the associated replica set. Another instant recovery is not possible. The old source volumes can be retained or removed.

For this method, the Data Protector Microsoft Volume Shadow Copy Service integration and the Microsoft Virtual Disk Service are used.
 - Resynchronizing the source volumes (with split mirror replicas) or restoring the data from snapshots to the source volumes (with snapshot replicas):

If a split mirror replica is used, the source volumes are resynchronized with those of the selected replica. If a snapshot replica is used, data from the selected replica is copied to the source volumes.

For this method, depending on the Data Protector components involved in the zero downtime backup session, either only the Data Protector HP P9000 XP Agent is used or the Data Protector Microsoft Volume Shadow Copy Service integration and the HP P9000 XP Agent are used.
- With HP 3PAR StoreServ Storage (in configurations where instant recovery is available), only one instant recovery method is possible:
 - Copying the replica data back to the source volumes

Data from the replica is copied back to the original storage, and the source volumes are not retained. The replica is retained in the associated replica set.

For this method, both the Data Protector Microsoft Volume Shadow Copy Service integration and the Data Protector HP 3PAR VSS Agent are used.

6. Re-enables any volume groups that it disabled and re-mounts any filesystems that it dismounted. After instant recovery, the contents of the source volumes are returned to the state they were in when the replica was created.

Instant recovery and LVM mirroring

Instant recovery is supported for ZDB sessions produced on HP-UX systems with an LVM mirroring plus HP BC P6000 EVA or HP BC P9000 XP configurations. However, it is necessary to perform additional manual steps. For information, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*.

Instant recovery in a cluster

Instant recovery is supported for an application or a filesystem running in a cluster environment on the application system. However, you need to perform additional steps. For information, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*. For VSS integration-specific information, see the *HP Data Protector Zero Downtime Backup Integration Guide*.

Split mirror restore

NOTE: With the speed of contemporary SAN-attached tape drives, it is usually quicker to restore directly to the application system than use split mirror restore.

In split mirror restore, backup objects are first moved from tape to a replica (either existing or newly created for the purpose) on the backup system. Data from the replica is then restored to the source volumes available to the application system, effectively replacing the existing contents of the source volume. It can be used to restore complete sessions or individual backup objects.


The method can be used to restore data from filesystem or disk image ZDB-to-tape or ZDB-to-disk+tape sessions produced under the following conditions:

- On P9000 XP Array, using the HP Business Copy (BC) P9000 XP configuration.
- On EMC, using the Symmetrix TimeFinder, SRDF, or combined (SRDF+TimeFinder) configurations.

Split mirror restore process

The following is an example of a split mirror restore process on P9000 XP Array:

Figure 24 Split mirror restore example



1. Select a replica to use for the restore or create a new replica to produce an up-to-date duplicate of the source volumes.
2. Restore the required objects from tape to the replica through the backup system.
3. Restore data from the replica, effectively replacing the data located on the source volumes with the data stored in the replica.

After the process is complete, the contents of the selected replica replace those of the source volumes:

- The backup objects restored from tape to the replica are returned to their states at the time the ZDB session was performed.
- The rest of the contents are returned to their states at the time the replica was created.

7 Planning

Introduction

To plan your ZDB strategy, you need to consider the following steps:

1. Define the requirements and constraints for backups, such as:
 - How often does your data need to be backed up?
 - Do you need additional copies of the backed up data on additional media sets?
2. Understand the factors that affect disk array performance.
3. Prepare a backup strategy that supports your backup concept and how it is implemented.

This chapter provides some important information and considerations that will help you plan your backup solution and improve ZDB performance.

Flexibility in recovery

For maximum flexibility in recovery to a point in time:

- Create replicas regularly and keep them on the disk array.
- Back up log files regularly.

To control disk array space usage:

- By defining a backup policy based on scheduled ZDB backup sessions, set up a time-based series of replicas, with each replica corresponding to a particular point in time. The number of replicas in such a replica set depends on the available disk array space and the desired time range.

Note that with specific types of snapshot replicas, the maximum number of replicas in the set may be limited by the disk array model and/or the installed disk array firmware revision.

- **HP P6000 EVA Disk Array Family:** Choose an appropriate snapshot type.

Split mirror disk arrays

The HP P9000 XP Disk Array Family and EMC Symmetrix Disk Array integrations provide options enabling you to define your backup policy, such as:

- Move the mirror copy of the original data to tape.
- Leave the mirror split or resynchronize it.
- Prepare the next disk for backup.

For example backup policies, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*.

You can find general recommendations and limitations about the split mirror disk array's performance discussed in this chapter.

Snapshot disk arrays

If you use the Data Protector HP P4000 SAN Solutions integration, consider the following when planning your backup strategy:

- instant recovery – see [“Disk array-specific considerations” \(page 59\)](#).

If you use the Data Protector HP P6000 EVA Disk Array Family integration, consider the following when planning your backup strategy:

- type of snapshot (standard snapshot, vsnap, or snapclone)
- replica redundancy level – see the *HP Data Protector Zero Downtime Backup Administrator's Guide*
- other disk array-specific considerations – see [“Disk array-specific considerations” \(page 59\)](#)
- instant recovery – see [“Disk array-specific considerations” \(page 59\)](#) and the *HP Data Protector Zero Downtime Backup Administrator's Guide*

If you use the Data Protector HP P9000 XP Disk Array Family integration, consider the following when planning your backup strategy:

- replica type (split mirror or snapshot) – see [“Disk array-specific considerations” \(page 59\)](#)
- instant recovery – see [“Disk array-specific considerations” \(page 59\)](#) and the *HP Data Protector Zero Downtime Backup Administrator's Guide*

If you use the Data Protector HP 3PAR StoreServ Storage integration, consider the following when planning your backup strategy:

- replica creation – see [“Disk array-specific considerations” \(page 59\)](#).

Disk array-specific considerations

Replica sets on P4000 SAN Solutions

Although you can create replica sets, replica set rotation is not supported with this disk array family.

Instant recovery on P4000 SAN Solutions

When a target volume is selected for instant recovery, if newer target volumes than the selected one exist for the same source volume, they are removed from the disk array automatically, regardless of the replica set they belong to. If a particular newer target volume cannot be removed, for example, because its smartclone exists on the disk array, the instant recovery session fails. The instant recovery session also fails if a newer snapshot not created by Data Protector exists for the source volume selected for instant recovery.

When the same source volume is included in several ZDB backup specifications, running an instant recovery session based on a particular ZDB backup specification may result in the inability to perform instant recovery sessions based on other ZDB backup specifications. Such a problem occurs if the following operations take place in the presented order:

1. An instant recovery session based on a particular ZDB backup specification (specification A) is run, and a newer target volume from the one selected for instant recovery is removed from the disk array. The removed target volume was created in a ZDB session (session B) based on another ZDB backup specification (specification B).
2. An instant recovery session corresponding to the ZDB session (session B) is started.

Replica creation on P6000 EVA Array

A new snapclone for a particular source volume can only be created if creation of the previous snapclone for the same volume has finished. If it has not, Data Protector automatically retries the operation a configurable number of times at configurable intervals. Standard snapshots and vsnaps are not subjected to this constraint.

You can shorten the time frame during which the performance of the application system is affected during zero downtime backup sessions by using mirrorclones, at the expense of the disk array storage space:

1. Using HP Command View (CV) EVA, create mirrorclones of the original storage volumes on which your application data resides.

Mirrorclone creation may be time-consuming and makes shortening the backup window impossible if it is triggered while a Data Protector ZDB session is already running. This step helps you avoid such situations.

2. In the ZDB backup specification that will be used in ZDB sessions, select mirrorclone as the snapshot source.

For details, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*.

Replica set rotation on P6000 EVA Array

A replica cannot be reused in the following cases:

- One of the target volumes that is a snapclone has a snapshot attached to it.
- One of the target volumes to be reused is presented to a system.

“Reuse” means that a replica in the replica set is removed and a new one created. This typically happens with the oldest replica when the specified maximum number of replicas in the replica set is reached and a new replica is required.

If the replica to be reused is in use and therefore locked by another session, Data Protector HP P6000 / HP 3PAR SMI-S Agent creates a new replica, and marks the existing replica for deletion. You can manually remove such residual replicas using the `omnidbsmis` command at a later time. For details, see the *HP Data Protector Command Line Interface Reference*.

Mirrorclones that are automatically created by Data Protector in particular ZDB sessions cannot be used for instant recovery and are therefore excluded from the replica set rotation.

For details, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*.

Instant recovery on P6000 EVA Array

Instant recovery can be performed regardless of the snapshot type used for the target volumes. If newer replicas than the one selected for instant recovery exist in the replica set, they are preserved regardless of the snapshot type they use: standard snapshot, vsnap, or snapclone.

Before choosing a snapshot type for the zero downtime backup session, consider the following:

- The fastest instant recovery method is switching the disks, which is only available for the snapclone snapshot type.
- When a replica consisting of standard snapshots or vsnaps is selected for instant recovery, and newer replicas from the selected replica exist in the replica set, the instant recovery process lasts longer than usual. The reason is that not only the source volumes, but all newer replicas must be updated during the session as well. In such circumstances, you can influence the time required for instant recovery by carefully defining the number of replicas in the replica set.

If snapshots of mirrorclones were created in the corresponding zero downtime backup session, during instant recovery, the data from mirrorclone snapshots is restored to the original volumes, rather than the mirrorclones themselves.

Replica type selection on P9000 XP Array

When creating a ZDB backup specification, you cannot directly select a desired replica type in the Data Protector GUI. You can ensure that Data Protector uses a specific replica type by specifying the appropriate mirror unit (MU) numbers or number ranges. When a source volume belonging to a particular MU number is used in a zero downtime backup session, the Data Protector HP

P9000 XP Agent chooses the replica type according to the type of the paired virtual disk which must be pre-configured using HP P9000 XP Remote Web Console.

Instant recovery on P9000 XP Array

When a replica is selected for instant recovery, if newer replicas than the selected replica exist in the replica set, they are preserved after the session regardless of their type: split mirror or snapshot.

Before choosing a replica type to be used in the ZDB session running in scope of your backup policy, consider the following: the instant recovery process runs fastest when a split mirror replica is selected for instant recovery, and the P9000 XP Array feature Quick Restore mode is enabled during pre-configuration of replica volumes on the disk array.

Replica creation on 3PAR StoreServ system

Each time replica creation is invoked on a 3PAR StoreServ system, two snapshots are actually created for each source volume: one read-only snapshot and one read-write snapshot. Only the read-write snapshot is presented to the outside application, and the read-only snapshot is internally used by the storage system. For more information about storage space consumption, see the HP 3PAR StoreServ Storage documentation.

Concurrency handling

Locking

Backup device locking

Regular (non-ZDB) Data Protector backup and restore sessions lock a tape device used in the session at the beginning of a backup or restore session and unlock it at the end of the session. The Data Protector tape device locking is described in detail in the *HP Data Protector Help*. With ZDB integrations, the tape device locking is changed so that a device is locked only for the time needed to transfer data to or from a tape device:

- During a ZDB-to-tape session or a ZDB-to-disk+tape session, the lock occurs after the replica is created but before the replicated data is streamed to tape.
- During a split mirror restore session (supported on specific disk array families), the lock occurs after the replica is created, but before the backup data is moved from a tape device to the replica.

A device is released when the transfer of data to or from a tape device is finished.

During a ZDB-to-disk or instant recovery session, tape devices are not used, so there is no tape device locking with these two operations.

Disk locking

To prevent a ZDB or instant recovery session from accessing storage volumes that may still be in use by another session, an internal disk locking mechanism is introduced by Data Protector. With this, storage volumes are locked for the time during which they are being used by another operation.

Data Protector issues a warning and aborts a session if it cannot lock storage volumes needed for the required operation (because they are already locked by another process).

Backup scenarios

Your backup strategy may consist of full and incremental backups. These sessions are not necessarily exclusively ZDB or non-ZDB. You can combine them in various ways. The following combinations are supported:

Table 6 Backup scenarios

Full backup	Incremental backups
ZDB	ZDB
ZDB	non-ZDB
ZDB	non-ZDB and ZDB
non-ZDB	ZDB
non-ZDB	ZDB and non-ZDB

NOTE: If you want to back up the same objects in ZDB and non-ZDB sessions, create separate backup specifications for each backup type. For example, one for ZDB to disk+tape, one for ZDB to tape, and one for non-ZDB session.

Ensure that the selected backup objects in the backup specifications match (the same client, mount point, and description). Otherwise, during restore, incremental and full backups from the tape cannot be included in the same restore chain because Data Protector treats these backups as separate objects.

Here are some advantages of incremental ZDB sessions:

- Good instant recovery granularity (provided that you have selected the `Track the replica for instant recovery` option in the backup specification)
- Reduced impact on the application system performance during backup
- Reduced amount of data that is streamed to tape

Example

To provide good instant recovery granularity, by creating replicas every two or three days and keeping them for instant recovery purposes, and to reduce the amount of data that is streamed to tape, you can decide for the following backup strategy:

- Full ZDB to disk+tape sessions on Sundays
- Incremental ZDB to disk+tape sessions on Tuesdays and Thursdays
- Incremental ZDB to tape sessions on other weekdays

In this scenario, configure the backups as follows:

- Create a ZDB to disk+tape backup specification and schedule full backups on Sundays and incremental backups on Tuesdays and Thursdays.
- Create a ZDB to tape backup specification and schedule incremental backups on Mondays, Wednesdays, Fridays, and Saturdays.

To restore your data, you can then use either replicas (quick restore) or backups from the tape. You can also combine the two restore types by restoring replicas first and then restoring individual files from a specific backup image from the tape.

A Supported configurations

Introduction

This appendix gives you information on the configurations supported on different disk arrays. The configurations described are supported by Hewlett-Packard. For an up-to-date list of supported configurations, see the latest support matrices at <http://support.openview.hp.com/selfsolve/manuals>. If you want to back up data in a configuration not listed, this does not mean that it cannot be supported. Contact your local HP representative or HP consulting to investigate the supportability of additional configurations.

The single-host (BC1) configuration, where a single system acts as the application system and the backup system, is not recommended because of performance issues. With the BC1 configuration, only filesystem and disk image backups are possible.

The HP P6000 EVA Disk Array Family single-host (BC1) configuration based on Linux platform is not supported. In such a configuration, a single Linux system acts as the application system and the backup system.

In the following table, disk arrays supported by Data Protector and capable of creating replicas and, in most cases, replica sets are listed.

Table 7 Disk arrays integrating with Data Protector

Disk array family	Abbreviations	Supported replication technologies available
HP P4000 SAN Solutions	P4000 SAN Solutions	Snapshot
HP P6000 EVA Disk Array Family	P6000 EVA Array	Snapshot
HP P9000 XP Disk Array Family	P9000 XP Array	Split mirror, snapshot
HP 3PAR StoreServ Storage	3PAR StoreServ	Snapshot
EMC Symmetrix Disk Array	EMC	Split mirror

For all supported configurations, a ZDB backup specification can only include one application system and one backup system. You can, however, have multiple ZDB backup specifications for each application system, and you can use these to back up the same application system simultaneously to different filesystems. For information on configurations with multiple application systems, see “Creating mount points” (page 50). With all configurations, original data and backup data can be spread across multiple disk arrays of the same type.

Note that each configuration has a specific behavioral pattern imposing specific requirements on the control functions to guarantee backup and recovery functionality.

Supported HP P6000 EVA Disk Array Family configurations

Local replication configurations

For local replication, HP BC P6000 EVA configuration is used.

A separate backup system needs to be connected to a disk array. After the replica is created, Data Protector scans for new disks on the backup system, creates device files (UNIX systems), and performs all other necessary steps to mount the filesystems on the backup system so that it can access the replicated data. Data is streamed to tape from the replica, while the application system continues with operations.

“HP BC P6000 EVA snapshot configuration 1” (page 64) through “HP BC P6000 EVA snapshot configuration 3” (page 64) are examples of supported local replication configurations.

Figure 25 HP BC P6000 EVA snapshot configuration 1




Figure 26 HP BC P6000 EVA snapshot configuration 2


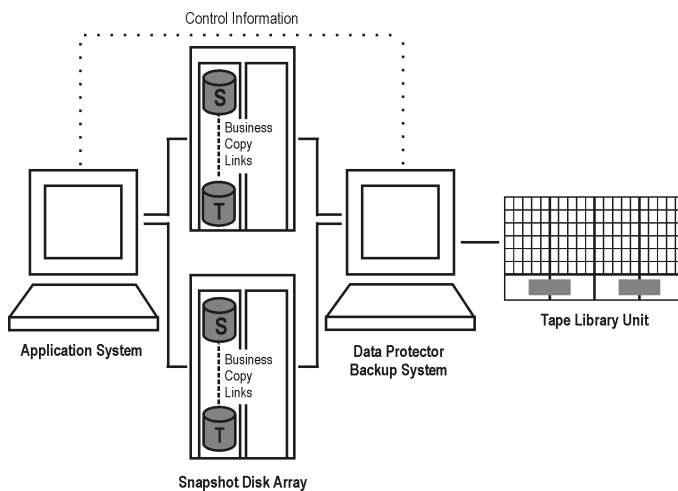


Figure 27 HP BC P6000 EVA snapshot configuration 3




Local replication configurations with HP-UX LVM mirroring

It is recommended to group the physical volumes of a volume group into physical volume groups (PVGs) and specify the `PVG-strict` policy for the mirror creation. With that, the mirrors of one logical volume will belong to different PVGs, which helps avoid certain situations, such as mirroring a logical volume onto the same disk.


“Supported LVM mirroring configuration 1” (page 65) through “Supported LVM mirroring configuration 3” (page 67) are examples of supported LVM mirroring configurations on P6000 EVA Array.

Figure 28 Supported LVM mirroring configuration 1



All logical volumes in a volume group are specified as backup objects in a backup specification. All logical volumes (with their extent distributions) are on different physical volumes within a PVG. Replicas are only created for those storage volumes that are found in that PVG. Later, these replicas are presented to the backup system for further backup of the selected backup objects. Both PVG-1 and PVG-2 satisfy the mirror selection rules. However, as the HP P6000 / HP 3PAR SMI-S Agent always attempts to select a secondary mirror, it will choose PVG-2 for the HP BC P6000 EVA pair replication.


Figure 29 Supported LVM mirroring configuration 2



Only selected logical volumes are included in a backup specification. Still, the PVG selected is the one that hosts all logical volumes of that volume group.

In this configuration, only PVG-2 can satisfy the mirror set selection rules, so it is selected for the BC pair replication.

Figure 30 Supported LVM mirroring configuration 3




Some of the secondary mirror members are hosted by the primary mirror disk array, so they cannot be replication candidates. The primary mirror set is therefore selected for the BC pair replication. For more information about LVM mirroring and mirror selection rules, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*.

Remote plus local replication configurations

For remote plus local replication on P6000 EVA Array, HP CA+BC P6000 EVA configuration is used.


"HP CA+BC P6000 EVA configuration 1" (page 67) through "HP CA+BC P6000 EVA configuration 3" (page 68) are examples of supported remote plus local configurations on P6000 EVA Array.

Figure 31 HP CA+BC P6000 EVA configuration 1




This configuration represents an ideal (non-failover) scenario.

Figure 32 HP CA+BC P6000 EVA configuration 2



This configuration represents a failover scenario with a reversed replication direction.

Figure 33 HP CA+BC P6000 EVA configuration 3



This configuration represents a failover scenario with maintained replica location.

Supported HP P9000 XP Disk Array Family configurations

Local replication configurations

"HP BC P9000 XP configuration 1" (page 68) through "HP BC P9000 XP configuration 3" (page 69) are examples of supported local replication configurations on P9000 XP Array.

Figure 34 HP BC P9000 XP configuration 1

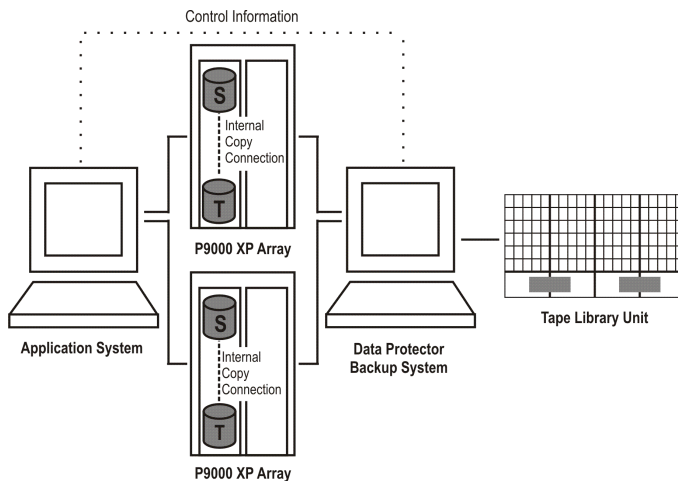


Figure 35 HP BC P9000 XP configuration 2

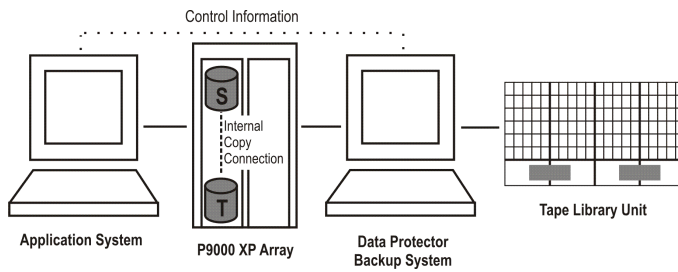
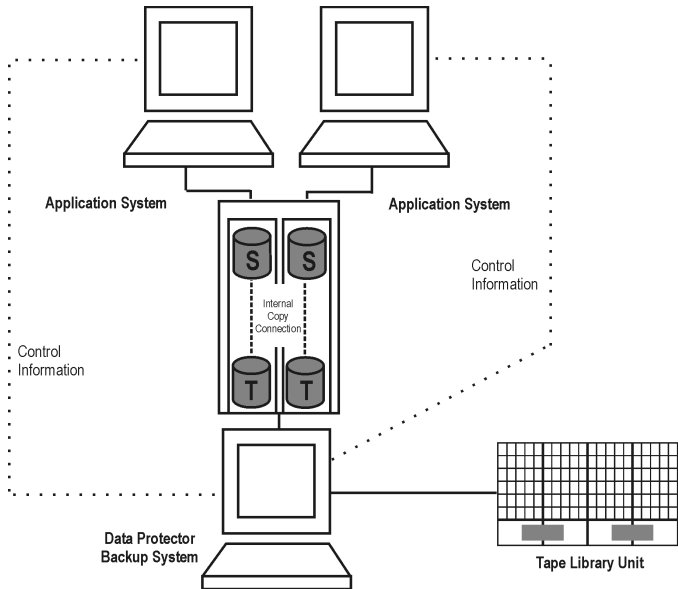



Figure 36 HP BC P9000 XP configuration 3



Single-host (BC1) configuration

The figure that follows shows a single-host configuration, also called **BC1 configuration**.

Figure 37 HP BC1 P9000 XP configuration

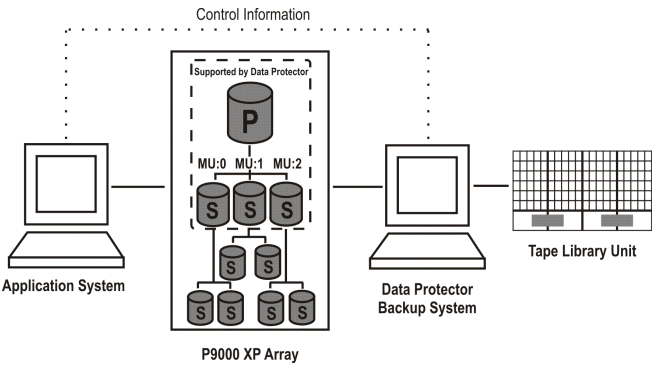


Cascading configurations

The HP P9000 XP Disk Array Family allows you to configure additional second-level mirrors or snapshot volumes for each first-level mirror or snapshot volume. This is referred to as a **cascading configuration**. However, Data Protector only uses first-level mirrors or snapshot volumes in zero downtime backup, instant recovery, and split mirror restore sessions.

The figure that follows is an example of the cascading configuration, where MU:0, MU:1 and MU:2 are first-level mirrors supported by Data Protector, and the six mirrors underneath are the second-level mirrors.

Figure 38 Cascading configuration



Local replication configurations with HP-UX LVM mirroring

“LVM mirroring configuration 1” (page 70) through “LVM mirroring configuration in a cluster” (page 71) are examples of supported LVM mirroring configurations on P9000 XP Array.

Figure 39 LVM mirroring configuration 1

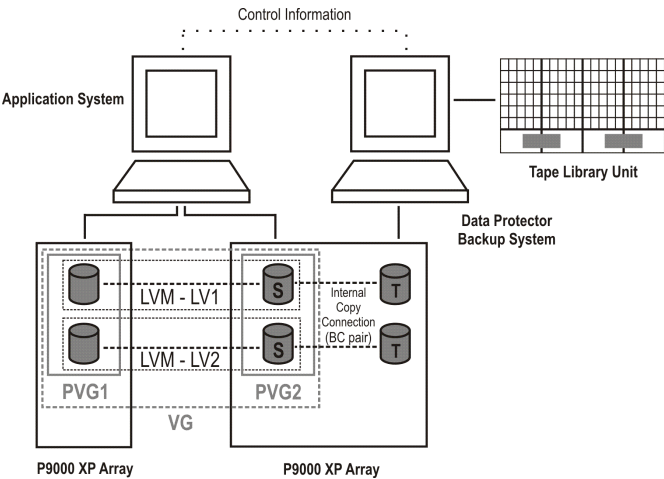


Figure 40 LVM mirroring configuration 2

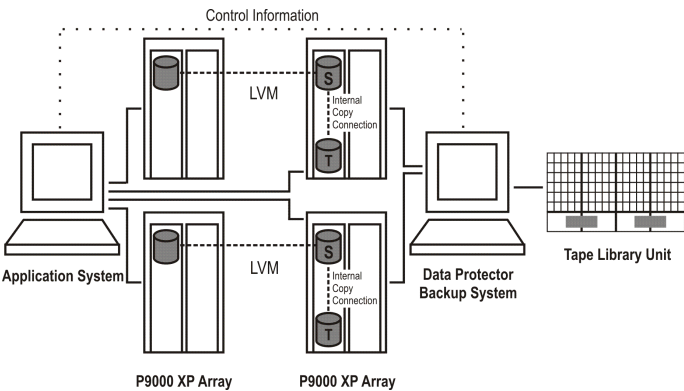


Figure 41 LVM mirroring configuration 3

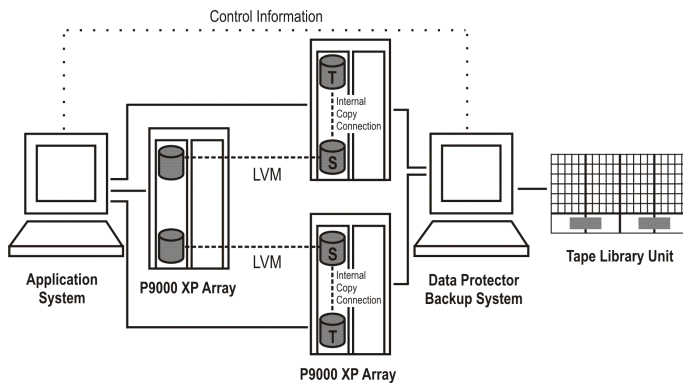


Figure 42 LVM mirroring configuration 4

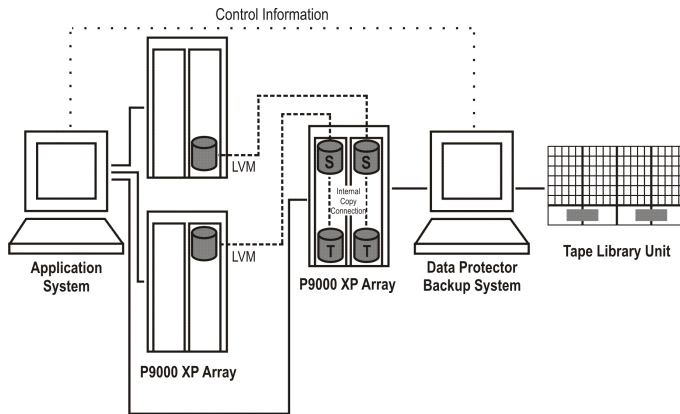
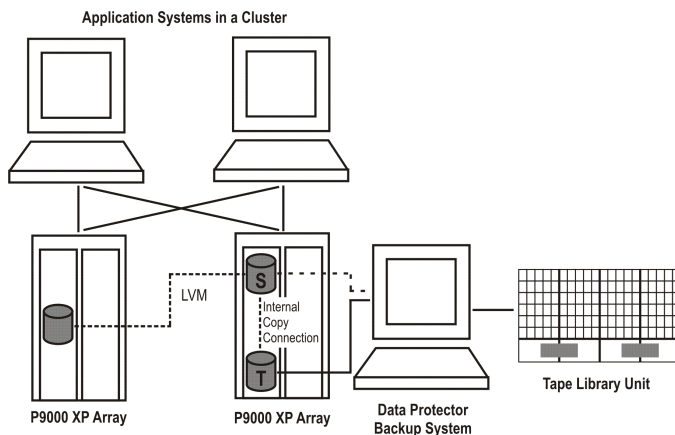


Figure 43 LVM mirroring configuration in a cluster



Remote replication configurations

A single backup system and a single P9000 XP Array can be used to back up multiple main disk arrays. See [“HP CA P9000 XP configuration 4” \(page 72\)](#). With this approach, you can build a central backup site. At least two disk arrays, located in physically separate sites, are needed for such a configuration.

During Data Protector sessions, the mirroring (CA) link between the disk arrays is used for zero downtime backup purposes. Additional mirroring (CA) links are needed to retain full high availability of data at the same time.

[“HP CA P9000 XP configuration 1” \(page 72\)](#) through [“HP CA P9000 XP configuration 4” \(page 72\)](#) are examples of supported remote replication configurations on P9000 XP Array.

Figure 44 HP CA P9000 XP configuration 1

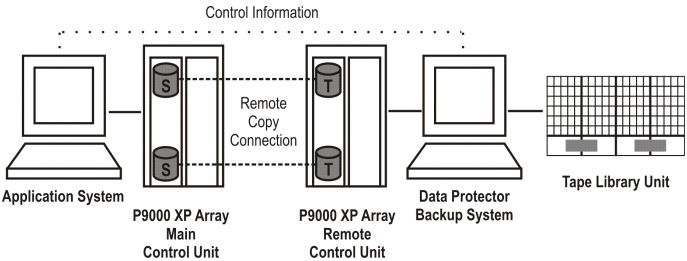


Figure 45 HP CA P9000 XP configuration 2

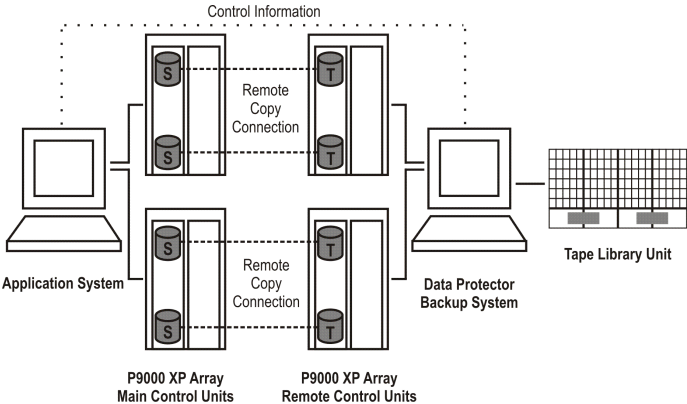


Figure 46 HP CA P9000 XP configuration 3


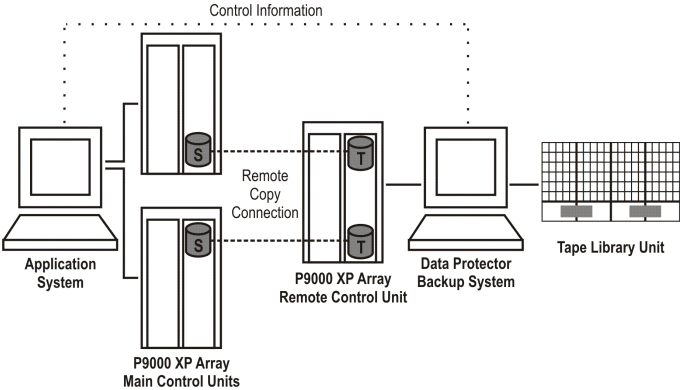


Figure 47 HP CA P9000 XP configuration 4



Remote plus local replication configurations

Limitations

- On HP-UX, it is recommended that only the BC target volume is connected to the backup system. If for any reason the CA target volume is connected as well, special care must be taken. For more information, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*.
- The asynchronous CA configuration as a part of the combined CA+BC configuration is not supported.

“HP CA+BC P9000 XP configuration 1” (page 73) through “HP CA+BC P9000 XP configuration 4” (page 74) are examples of supported remote plus local replication configurations on P9000 XP Array.

Figure 48 HP CA+BC P9000 XP configuration 1

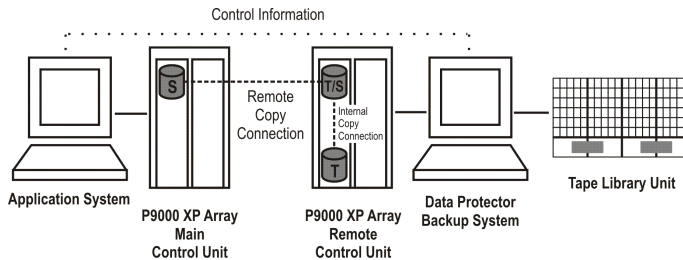


Figure 49 HP CA+BC P9000 XP configuration 2




Figure 50 HP CA+BC P9000 XP configuration 3

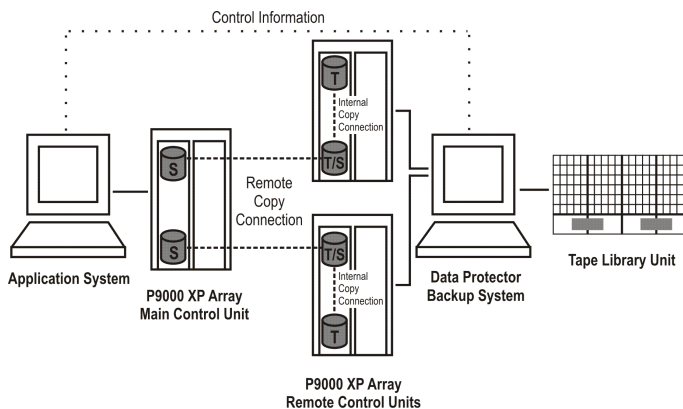



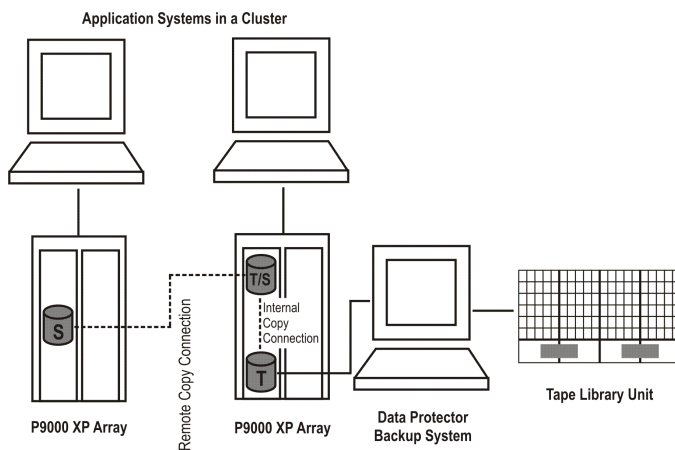
Figure 51 HP CA+BC P9000 XP configuration 4



Cluster configurations

The figure that follows is an example of an HP CA+BC P9000 XP Array configuration in a cluster.

Figure 52 HP CA+BC P9000 XP configuration in a cluster



For more information about cluster configurations, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*.

Supported EMC Symmetrix configurations

Local replication configurations

For local replication, the **EMC Symmetrix TimeFinder configuration** is used.

"TimeFinder configuration 1" (page 75) through "TimeFinder configuration 3" (page 75) are examples of supported local replication configurations on EMC.

Figure 53 TimeFinder configuration 1

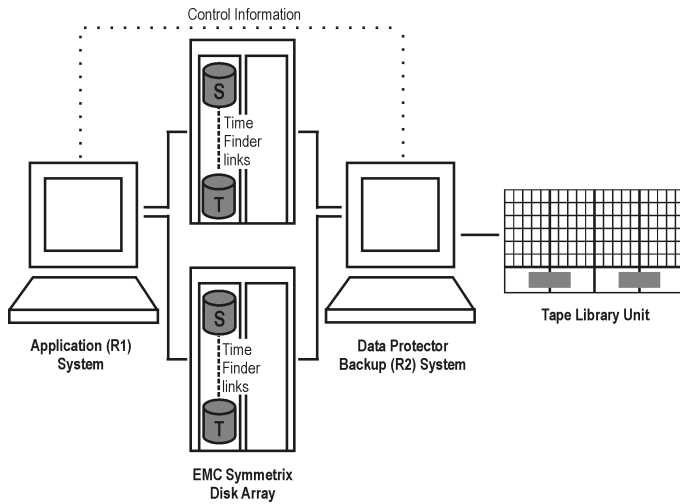


Figure 54 TimeFinder configuration 2


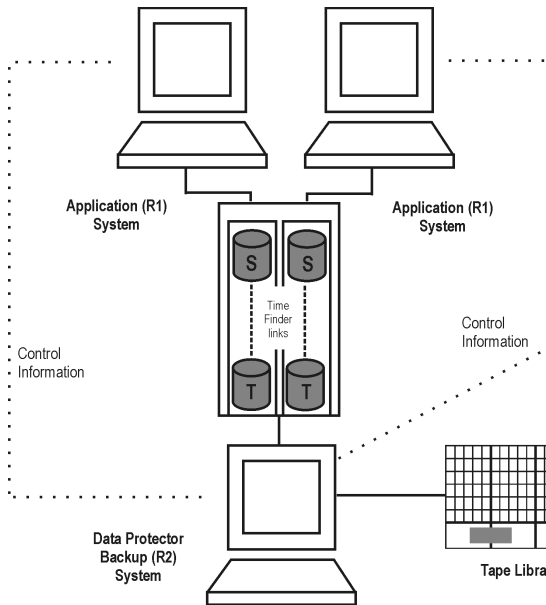


Figure 55 TimeFinder configuration 3



Local replication configurations with HP-UX LVM mirroring

“LVM mirroring configuration 1” (page 76) through “LVM mirroring configuration 5” (page 77) are examples of supported LVM mirroring configurations on EMC.

Figure 56 LVM mirroring configuration 1




Figure 57 LVM mirroring configuration 2




Figure 58 LVM mirroring configuration 3




Figure 59 LVM mirroring configuration 4


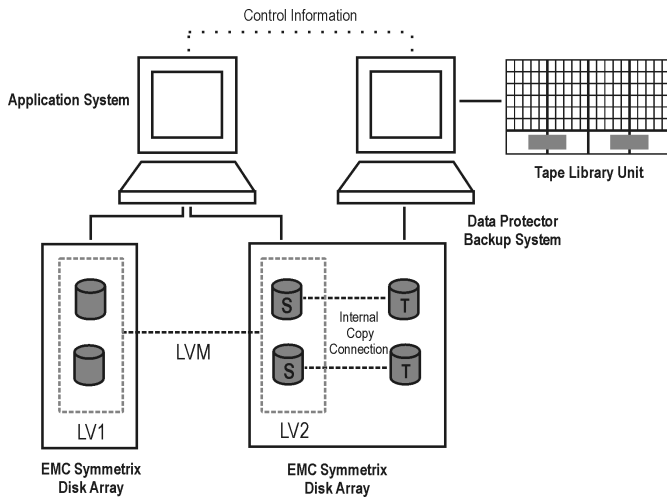


Figure 60 LVM mirroring configuration 5



Remote replication configurations

“SRDF configuration 1” (page 77) through “SRDF configuration 4” (page 78) are examples of supported remote replication configurations on EMC.

Figure 61 SRDF configuration 1




Figure 62 SRDF configuration 2




Figure 63 SRDF configuration 3





Figure 64 SRDF configuration 4



Remote plus local replication configurations

It is recommended that only the TimeFinder target volume is connected to the backup system. If for any reason the SRDF target volume is connected as well, special care must be taken. For more information, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*.

"SRDF+TimeFinder configuration 1" (page 79) through "SRDF+TimeFinder configuration 4" (page 79) are examples of supported remote plus local replication configurations on EMC.

Figure 65 SRDF+TimeFinder configuration 1




Figure 66 SRDF+TimeFinder configuration 2

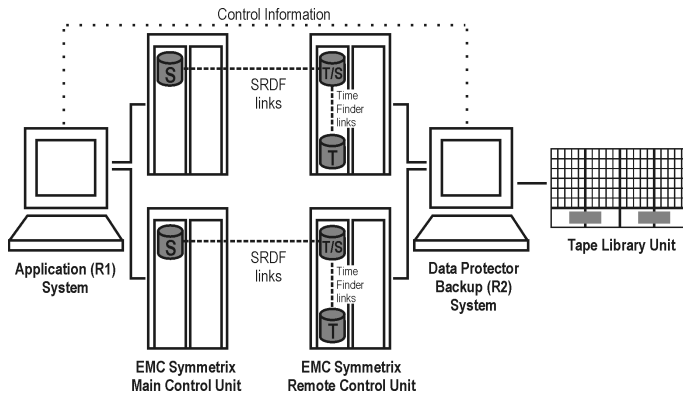


Figure 67 SRDF+TimeFinder configuration 3


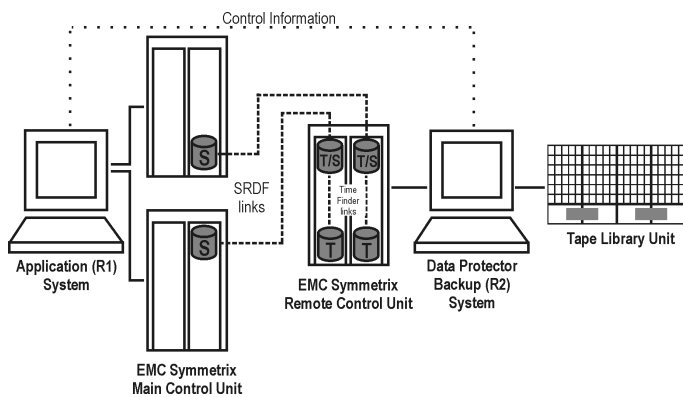


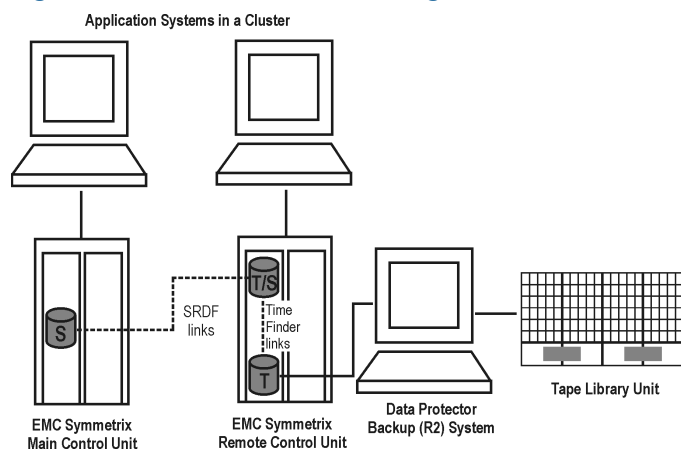
Figure 68 SRDF+TimeFinder configuration 4



Cluster configurations

The figure that follows is an example of an SRDF+TimeFinder configuration in a cluster.

Figure 69 SRDF+TimeFinder configuration in a cluster



For more information about cluster configurations, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*.

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 log files	(Data Protector specific term) Files that keep track of changes made to the Data Protector Internal Database (IDB). They are used for online and offline IDB restore and recovery where the IDB needs to be recreated either in its latest possible state, beyond the time of the most recent IDB backup session, or in a state between the times of two consecutive IDB backup sessions.
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: <ul style="list-style-type: none">• 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 <code>Data_Protector_program_data\Config\Server\dr\asr</code> (Windows systems) or <code>/etc/opt/omni/server/dr/asr</code> (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)	(<i>Oracle specific term</i>) 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.
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	(<i>SAP R/3 specific term</i>) 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	<p>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.</p> <p>A backup object is defined by:</p> <ul style="list-style-type: none"> • 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 systems and mount point on UNIX systems). 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	<p>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.</p> <p>See also backup specification, full backup, and incremental backup.</p>
backup set	A complete set of integration objects associated with a backup.
backup set	(<i>Oracle specific term</i>) 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	<i>(ZDB specific term)</i> 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	<i>(EMC Symmetrix specific term)</i> Business Continuance are processes that allow customers to access and manage instant copies of EMC Symmetrix standard devices. See also BCV.
BC Process	<i>(EMC Symmetrix specific term)</i> 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	<i>(EMC Symmetrix specific term)</i> 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 Data Protector 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	<i>(SAP R/3 specific term)</i> 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	<i>(SAP R/3 specific term)</i> 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	<i>(SAP R/3 specific term)</i> An SAP R/3 tool that can be used to restore files of the following type: <ul style="list-style-type: none"> • 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	<i>(StorageTek specific term)</i> Cartridge Access Port is a port built into the door panel of a library. The purpose is to enter or eject media.

Catalog Database (CDB)	A part of the Data Protector Internal Database (IDB) that contains information about backup, restore, object copy, object consolidation, object verification, and media management sessions. This part of the IDB is always local to the cell. It is stored in the embedded database. See also MMDB.
catalog protection	Defines how long information about backed up data (such as filenames and file attributes) is kept in the IDB. See also data protection.
CDB	See Catalog Database (CDB).
CDF file	(UNIX systems 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	<p>(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:</p> <ul style="list-style-type: none"> • type 'disk' • type 'sbt_tape' <p>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.</p>
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	<p>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:</p> <ul style="list-style-type: none"> • If you select the check box next to the client system name, a single backup object of the <code>Client System</code> 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, <code>CONFIGURATION</code> is also backed up. • If you individually select all volumes that are mounted on the client system, a separate backup object of the <code>Filesystem</code> 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	<p>(<i>Microsoft Exchange Server specific term</i>) 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.</p> <p>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.</p> <p>See also Exchange Replication Service and local continuous replication.</p>
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	(<i>Informix Server specific term</i>) 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	<p>The Data Protector Centralized Media Management Database (CMMDB) is the result of merging MMDs 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</p> <p>See also MoM.</p>
COM+ Class Registration Database	(<i>Windows specific term</i>) 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	(<i>HP P9000 XP Disk Array Family specific term</i>) 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-line interface (CLI)	A set of commands that you can use in shell scripts to perform Data Protector configuration, backup, restore, and management tasks.
concurrency	See Disk Agent concurrency.
container	(<i>HP P6000 EVA Disk Array Family specific term</i>) Space on a disk array, which is pre-allocated for later use as a standard snapshot, vsnap, or snapclone.
control file	(<i>Oracle and SAP R/3 specific term</i>) An Oracle data file that contains entries specifying the physical structure of the database. It provides database consistency information used for recovery.
copy set	<p>(<i>HP P6000 EVA Disk Array Family specific term</i>) A pair that consists of the source volumes on a local P6000 EVA and their replica on a remote P6000 EVA.</p> <p>See also source volume, replica, and HP Continuous Access + Business Copy (CA+BC) P6000 EVA.</p>
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 <code>root</code> .
CSM	The Data Protector Copy and Consolidation Session Manager process controls the object copy and object consolidation sessions and runs on the Cell Manager system.

data file	(Oracle and SAP R/3 specific term) A physical file created by Oracle that contains data structures such as tables and indexes. A data file can only belong to one Oracle database.
data protection	Defines how long the backed up data on media remains protected, that is, Data Protector will not overwrite it. When the protection expires, Data Protector will be able to reuse the media in one of the next backup sessions. See also catalog protection.
data 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, Windows 8, Windows Server 2008, and Windows Server 2012) or the directory containing Data Protector program files and data files (on other Windows operating systems). Its default path is <code>%ProgramFiles%\OmniBack</code> , but the path can be changed in the Data Protector Setup Wizard at installation time. See also <code>Data_Protector_program_data</code> .
Data_Protector_program_data	A reference to the directory containing Data Protector data files on Windows Vista, Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012. Its default path is <code>%ProgramData%\OmniBack</code> , but the path can be changed in the Data Protector Setup Wizard at installation time. See also <code>Data_Protector_home</code> .
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 blob space, db space, or logical log file.
DC directory	A directory that contains DC binary files, one for each configured Data Protector backup medium. DC directories constitute the Detail Catalog Binary Files part of the Data Protector Internal Database. See also Detail Catalog Binary Files (DCBF) and Internal Database (IDB).
DCBF	See Detail Catalog Binary Files (DCBF).
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.
Detail Catalog Binary Files (DCBF)	A part of the Data Protector Internal Database that stores names, versions, and metadata of the backed up items. It consists of DC directories with DC binary files. See also DC directory and Internal Database (IDB).
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	(<i>Microsoft SQL Server specific term</i>) 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	(<i>Windows specific term</i>) 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	(<i>Veritas Volume Manager specific term</i>) 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 backup	A high-speed backup where Data Protector backs up files as bitmap images. A disk image 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 metadata 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	<i>(Informix Server specific term)</i> The Informix Server configuration file <code>ixbar.server_id</code> that resides in the directory <code>INFORMIXDIR/etc</code> (on Windows systems) or <code>INFORMIXDIR/etc</code> (on UNIX systems). <code>INFORMIXDIR</code> is the Informix Server home directory and <code>server_id</code> is the value of the <code>SERVERNUM</code> 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. <code>KeyID</code> identifies the key within the keystore. <code>StoreID</code> identifies the keystore on the Cell Manager. If Data Protector has been upgraded from an earlier version with encryption functionality, there may be several <code>StoreIDs</code> 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. <i>See also MoM.</i>
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 <code>Data_Protector_program_data\log\server\Ob2EventLog.txt</code> (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	(<i>Windows specific term</i>) 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	(<i>Microsoft Exchange Server specific term</i>) 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 filenames, 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)	(<i>Microsoft Exchange Server specific term</i>) 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 systems) or package (on UNIX systems) from one cluster node to another. A failover can occur mostly because of software or hardware failures or maintenance on the primary node.
failover	(<i>HP P6000 EVA Disk Array Family specific term</i>) 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	(<i>Windows specific term</i>) 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	(<i>HP P9000 XP Disk Array Family specific term</i>) 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	(<i>Oracle specific term</i>) 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.
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	A set of options that define behavior of the entire Data Protector cell. The options are stored in a plain text file on the Cell Manager
group	(<i>Microsoft Cluster Server specific term</i>) 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. It is available for Microsoft Windows operating systems.

H

hard recovery	(<i>Microsoft Exchange Server specific term</i>) 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 <code>Data_Protector_program_data\Config\Server\holidays</code> (Windows systems), or <code>/etc/opt/omni/server/Holidays</code> (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	(<i>HP P6000 EVA Disk Array Family specific term</i>) 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 P6000 / HP 3PAR SMI-S Agent and HP 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 P6000 / HP 3PAR SMI-S Agent

A Data Protector software module that executes all tasks required for the HP P6000 EVA Disk Array Family integration. With the HP P6000 / HP 3PAR SMI-S Agent, the control over the array is established through an appropriate SMI-S provider, which directs communication between incoming requests and the storage system's native interface.

See also HP Command View (CV) EVA and HP SMI-S P6000 EVA Array provider.

HP 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 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 HP P6000 / HP 3PAR SMI-S Agent, communicates with HP Command View EVA for information or method invocation, and returns standardized responses. See also HP P6000 / HP 3PAR SMI-S Agent and HP Command View (CV) EVA.

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

See Internal Database (IDB).

IDB recovery file

A file that maintains information about completed IDB backup sessions and the backup media and backup devices used in them. If available, it significantly simplifies and speeds up offline recovery of the Internal Database in case of a Cell Manager disaster. Its filename is obdrindex.dat.

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. <i>See also</i> exporting media.
incremental (re)-establish	<i>(EMC Symmetrix specific term)</i> 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. <i>See also</i> backup types.
incremental backup	<i>(Microsoft Exchange Server specific term)</i> 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. <i>See also</i> 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	<i>(EMC Symmetrix specific term)</i> 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. <i>See also</i> 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	<i>(Microsoft Exchange Server specific term)</i> 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. <i>See also</i> Key Management Service and Site Replication Service.
Informix Server initializing	<i>(Informix Server specific term)</i> Refers to Informix Dynamic Server. <i>See</i> 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	<i>(ZDB specific term)</i> 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. <i>See also</i> 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.
Internal Database (IDB)	An entity in Data Protector that keeps information regarding which data was backed up, to which media it was backed up, how and when backup, restore, and other sessions were run, which devices, libraries, and disk arrays are configured, and so on. It stores its data in an embedded database and a collection of proprietary data files which reside on the Cell Manager. <i>See also</i> DC directory and Detail Catalog Binary Files (DBCF).
Internet Information Services (IIS)	<i>(Windows specific term)</i> 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	<i>(Sybase specific term)</i> A Sybase utility used to perform system administration tasks on Sybase SQL Server.
J	
jukebox	<i>See</i> 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	<i>(Microsoft Exchange Server specific term)</i> The Microsoft Exchange Server service that provides encryption functionality for enhanced security. <i>See also</i> 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	<i>(EMC Symmetrix specific term)</i> 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	<i>(HP P9000 XP Disk Array Family specific term)</i> 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. <i>See also</i> 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	<i>(Oracle specific term)</i> 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	<p>(<i>Microsoft Exchange Server specific term</i>) 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.</p> <p>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.</p> <p>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.</p> <p>See also cluster continuous replication and Exchange Replication Service.</p>
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	(<i>Informix Server UNIX systems specific term</i>) 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 <code>ALARMPROGRAM</code> configuration parameter defaults to the <code>INFORMIXDIR/etc/log_full.sh</code> , where <code>INFORMIXDIR</code> is the Informix Server home directory. If you do not want logical logs to be backed up continuously, set the <code>ALARMPROGRAM</code> configuration parameter to <code>INFORMIXDIR/etc/no_log.sh</code> .
logging level	An option that 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 mainly influence the IDB growth 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	(<i>Microsoft SQL Server specific term</i>) 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 <code>syslogin</code> .
login information to the Oracle Target Database	<p>(<i>Oracle and SAP R/3 specific term</i>) The format of the login information is <code>user_name/password@service</code>, where:</p> <ul style="list-style-type: none"> <code>user_name</code> 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 <code>SYSDBA</code> or <code>SYSOPER</code> rights. <code>password</code> must be the same as the password specified in the Oracle password file (<code>orapwd</code>), which is used for authentication of users performing database administration. <code>service</code> is the name used to identify an SQL*Net server process for the target database.
login information to the Recovery Catalog Database	(<i>Oracle specific term</i>) The format of the login information to the Recovery (Oracle) Catalog Database is <code>user_name/password@service</code> , 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	(<i>Lotus Domino Server specific term</i>) 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	(<i>Microsoft Exchange Server specific term</i>) 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	(<i>Microsoft Exchange Server specific term</i>) 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)	(<i>HP P9000 XP Disk Array Family specific term</i>) 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.
maintenance mode	An operating mode you can initiate on the Cell Manager to prevent changes in the Internal Database. It enables you to perform various maintenance tasks, such as upgrading and patching the Data Protector installation.
make_net_recovery	<code>make_net_recovery</code> 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 <code>make_boot_tape</code> 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 <code>bootsys</code> command or interactively specified on the boot console.
make_tape_recovery	<code>make_tape_recovery</code> 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	(<i>Microsoft Exchange Server specific term</i>) 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. <i>See also</i> 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)	<i>(Windows specific term)</i> 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. <i>See also</i> shadow copy, shadow copy provider, replica, and writer.
mirror (EMC Symmetrix and HP P9000 XP Disk Array Family specific term)	<i>See</i> target volume.
mirror rotation (HP P9000 XP Disk Array Family specific term)	<i>See</i> replica set rotation.
mirror unit (MU) number	<i>(HP P9000 XP Disk Array Family specific term)</i> 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. <i>See also</i> first-level mirror.
mirrorclone	<i>(HP P6000 EVA Disk Array Family specific term)</i> 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 Catalog Database (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 <code>/opt</code> or <code>d:</code> . On UNIX systems, the mount points are displayed using the <code>bdf</code> or <code>df</code> 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	(<i>HP P6000 EVA Disk Array Family specific term</i>) 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.
○	
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	(<i>Windows specific term</i>) 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. Cell Manager can only be recovered offline.
offline redo log	See archived redo log.
ON-Bar	<i>(Informix Server specific term)</i> 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: <ul style="list-style-type: none"> • 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	<i>(Informix Server specific term)</i> 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 <i>INFORMIXDIR/etc</i> (on Windows systems or <i>INFORMIXDIR/etc/</i> (on UNIX systems).
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	A type of the Internal Database recovery that you can use when the Cell Manager is accessible. In this case, the Cell Manager runs sessions, the sessions are logged into the IDB, and you can monitor the session progress using the GUI.
online redo log	<i>(Oracle specific term)</i> 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	<i>(Oracle specific term)</i> 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	<i>(Oracle specific term)</i> Each installation of an Oracle database on one or more systems. One computer system can have several instances of a database running.
ORACLE_SID	<i>(Oracle specific term)</i> A unique name for an Oracle Server instance. To switch among Oracle Servers, specify the desired <i>ORACLE_SID</i> . The <i>ORACLE_SID</i> 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. <i>See also</i> merging.
ownership	<p>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.</p> <p>If a user starts an existing backup specification without modifying it, the backup session is not considered as interactive.</p> <p>If a modified backup specification is started by a user, the user is the owner unless the following is true:</p> <ul style="list-style-type: none">• 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. <p>If a backup is scheduled on a UNIX Cell Manager, the session owner is root:sys unless the above conditions are true.</p> <p>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.</p> <p>When copying or consolidating backup objects, the owner of the resulting objects is the user who started the original backup session.</p>

P

P1S file	<p>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 <code>Data_Protector_program_data\Config\Server\dr\p1s</code> (Windows systems), or <code>/etc/opt/omni/server/dr/p1s</code> (UNIX systems) with the filename <code>recovery.p1s</code>.</p>
package	<p>(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.</p>
pair status	<p>(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 P9000 XP Agent:</p> <ul style="list-style-type: none">• 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 systems and as shell scripts on UNIX systems. <i>See also</i> 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 systems and as shell scripts on UNIX systems.
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 systems and as shell scripts on UNIX systems. <i>See also</i> 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)	<i>(HP P9000 XP Disk Array Family specific term)</i> 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). <i>See also</i> secondary volume (S-VOL) and Main Control Unit (MCU).
protection	<i>See</i> data protection and also catalog protection.
public folder store	<i>(Microsoft Exchange Server specific term)</i> 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: <ul style="list-style-type: none"> • 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	<i>(HP P9000 XP Disk Array Family specific term)</i> 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. <i>See also</i> HP P9000 XP Agent.
RAID Manager P9000 XP	<i>(HP P9000 XP Disk Array Family specific term)</i> 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	<i>See</i> disk image backup.
RCU	<i>See</i> Remote Control Unit (RCU).
RDBMS	Relational Database Management System.

RDF1/RDF2	<i>(EMC Symmetrix specific term)</i> 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.
Recovery Catalog	<p><i>(Oracle specific term)</i> 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:</p> <ul style="list-style-type: none"> • 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	<i>(Oracle specific term)</i> An Oracle database that contains a recovery catalog schema. You should not store the recovery catalog in your target database.
recovery files	<p><i>(Oracle specific term)</i> 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.</p> <p>See also flash recovery area.</p>
Recovery Manager (RMAN)	<i>(Oracle specific term)</i> 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	<i>(Oracle specific term)</i> 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)	<i>(HP P9000 XP Disk Array Family specific term)</i> 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	<i>(Windows specific term)</i> 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	<i>(Windows specific term)</i> 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	<i>(ZDB specific term)</i> 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 a UNIX system, the whole volume or disk group containing a backup object (logical volume) is replicated. If partitions are used on a Windows system, 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	Backup images that are needed to restore a backed up object to the state it was in at the selected point in time. In general, a restore chain of an object consists of its full backup image and one or more related incremental backup images.
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

	<p>volume. The original volume continues to change as the backup process continues, but the shadow copy of the volume remains constant.</p> <p>See also Microsoft Volume Shadow Copy Service and replica.</p>
shadow copy provider	<p>(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.</p>
shadow copy set	<p>(Microsoft VSS specific term) A collection of shadow copies created at the same point in time. See also shadow copy and replica set.</p>
shared disks	<p>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.</p>
Site Replication Service	<p>(Microsoft Exchange Server specific term) The Microsoft Exchange Server service that permits compatibility with Microsoft Exchange Server 5.5 by emulating the Exchange Server 5.5 directory service.</p> <p>See also Information Store and Key Management Service.</p>
slot	<p>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.</p>
SMB	<p>See split mirror backup.</p>
SMBF	<p>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.</p>
SMI-S Agent (SMISA)	<p>See HP P6000 / HP 3PAR SMI-S Agent.</p>
snapshot	<p>(HP P4000 SAN Solutions, HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, and HP 3PAR StoreServ Storage 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.</p> <p>See also replica and snapshot creation.</p>
snapshot backup	<p>See ZDB to tape, ZDB to disk, and ZDB to disk+tape.</p>
snapshot creation	<p>(HP P4000 SAN Solutions, HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, and HP 3PAR StoreServ Storage 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.</p> <p>See also snapshot.</p>
source (R1) device	<p>(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.</p> <p>See also target (R2) device.</p>
source volume	<p>(ZDB specific term) A storage volume containing data to be replicated.</p>
sparse file	<p>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.</p>
split mirror	<p>(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.</p> <p>See also replica and split mirror creation.</p>

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 systems) or registry (on Windows systems) 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 P9000 XP Agent.
sst.conf file	The file <code>/usr/kernel/drv/sst.conf</code> is required on each Data Protector 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 <code>/kernel/drv/st.conf</code> 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	(<i>Sybase specific term</i>) 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	(<i>Sybase specific term</i>) 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	(<i>Oracle specific term</i>) 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	<p>(<i>Sybase specific term</i>) The four system databases on a newly installed Sybase SQL Server are the:</p> <ul style="list-style-type: none"> • master database (master) • temporary database (tempdb) • system procedure database (sybsystemprocs) • model database (model).
System Recovery Data file	See SRD file.
System State	(<i>Windows specific term</i>) 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	(<i>Windows specific term</i>) 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	<p>(<i>EMC Symmetrix specific term</i>) 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.</p> <p>See also source (R1) device.</p>
target database	(<i>Oracle specific term</i>) In RMAN, the target database is the database that you are backing up or restoring.
target system	(<i>disaster recovery specific term</i>) 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	<i>(ZDB specific term)</i> A storage volume to which data is replicated.
Terminal Services	<i>(Windows specific term)</i> 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	<i>(Microsoft SQL Server specific term)</i> 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	<i>(EMC Symmetrix specific term)</i> 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	<i>(Oracle and SAP R/3 specific term)</i> 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	<i>(Sybase and SQL specific term)</i> 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	<i>(Sybase specific term)</i> A system table in which all changes to the database are automatically recorded.
transportable snapshot	<i>(Microsoft VSS specific term)</i> 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).

U

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, Windows 8, Windows Server 2008, and Windows Server 2012 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	<i>(Windows specific term)</i> 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 <i>admin</i> and <i>operator</i> .
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)	<i>(HP P6000 EVA Disk Array Family specific term)</i> The firmware that manages all aspects of storage system operation, including communication with HP Command View EVA through the HSV controllers. <i>See also</i> HP Command View (CV) EVA.
Virtual Device Interface	<i>(Microsoft SQL Server specific term)</i> This is a Microsoft SQL Server programming interface that allows fast backup and restore of large databases.
virtual disk	<i>(HP P6000 EVA Disk Array Family specific term)</i> 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. <i>See also</i> 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	<i>(VLS specific term)</i> 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. <i>See also</i> Virtual Library System (VLS) and virtual tape library (VTL).
virtual tape library (VTL)	<i>(VLS specific term)</i> An emulated tape library that provides the functionality of traditional tape-based storage. <i>See also</i> Virtual Library System (VLS).
VMware management client	<i>(VMware (Legacy) integration specific term)</i> 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	<i>(ADIC and STK specific term)</i> 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	<i>(Windows specific term)</i> 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	<i>(HP P9000 XP Disk Array Family VSS provider specific term)</i> 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	<i>(Microsoft VSS specific term)</i> 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	<i>(Informix Server specific term)</i> 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	<i>(ZDB specific term)</i> 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	<i>(ZDB specific term)</i> 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	<p><i>(ZDB specific term)</i> 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.</p> <p>See also zero downtime backup (ZDB), ZDB to disk, ZDB to tape, instant recovery, replica, and replica set rotation.</p>
ZDB to tape	<p><i>(ZDB specific term)</i> 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.</p> <p>See also zero downtime backup (ZDB), ZDB to disk, ZDB to disk+tape, instant recovery, and replica.</p>
zero downtime backup (ZDB)	<p>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.</p> <p>See also ZDB to disk, ZDB to tape, ZDB to disk+tape, and instant recovery.</p>

Index

Symbols

3PAR, restore, 46

A

application integration agent, 32

Application integrations

VSS, 41

application systems, 32

audience, 8

B

backup scenarios, 61

backup specifications, 43

backup systems, 32

backup types, 17–18

incremental ZDB, 50

ZDB to disk, 18, 45

ZDB to disk+tape, 18, 45

ZDB to tape, 17, 45

BC configuration

P6000 EVA Array, 35, 63

P9000 XP Array, 37, 68

BC1 configuration

P9000 XP Array, 69

Business Copy configuration see BC

C

CA configuration

P9000 XP Array, 37, 71

CA+BC configuration

P6000 EVA Array, 36, 67

P9000 XP Array, 38, 73

cascading configuration

P9000 XP Array, 69

Cell Manager, 32

clusters

CA+BC P9000 XP Array, 74

instant recovery, 56

LVM mirroring EMC, 77

LVM mirroring P9000 XP Array, 71

SRDF+TimeFinder EMC, 79

concurrency handling

device locking, 61

disk locking, 61

configurations

BC, P6000 EVA Array, 63

BC, P9000 XP Array, 68

BC1, P9000 XP Array, 69

CA+BC, P6000 EVA Array, 67

CA+BC, P9000 XP Array, 73

CA, P9000 XP Array, 71

cascading, P9000 XP Array, 69

LVM mirroring, EMC, 75

LVM mirroring, P6000 EVA Array, 64

LVM mirroring, P9000 XP Array, 70

SRDF+TimeFinder, EMC, 78

SRDF, EMC, 77

TimeFinder, EMC, 74

Continuous Access configuration see CA

conventions

document, 13

creating replicas, 17, 43, 49

D

Data Facility configuration see RDF

Data Protector cell, 31–34

application systems, 32

backup systems, 32

Cell Manager, 32

components, 31

ZDB database, 32

database applications, 40–41

Microsoft Exchange Server, 40

Microsoft SQL Server, 40

offline backup, 17, 48

online backup, 17, 48

Oracle, 40

restore, 41

SAP R/3, 40

supported database applications, 40

transaction log backup, 41

deleting replicas, 47

device locking, 61

disk array agent, 32

disk arrays, introduction, 21–22

disk virtualization, 21

RAID technology, 21

storage volumes, 21

disk arrays, supported configurations, 34–36

3PAR, 38

EMC, 38, 74

P6000 EVA Array, 35, 63

P9000 XP Array, 36, 68

disk arrays, supported ZDB techniques, 18, 63

disk locking, 61

disk virtualization, 16, 21

document

conventions, 13

related documentation, 8

documentation

HP website, 8

providing feedback, 15

E

EMC Symmetrix see EMC

EMC, backup

local replication, 39, 74–75

local replication integrating with LVM mirroring, 39

local replication using LVM mirroring, 75–77

remote plus local replication, 40, 78–80

remote replication, 39, 77–78

EMC, configurations
 LVM mirroring, 75
 SRDF, 39–40, 77
 SRDF+TimeFinder, 78
 TimeFinder, 39–40, 74
EMC, restore
 split mirror restore, 57

F
full ZDB, 61

H
help
 obtaining, 14
hot-backup mode, 16–17, 48
HP

 technical support, 14
HP 3PAR StoreServ Storage see 3PAR
HP P4000 SAN Solutions see P4000 SAN Solutions
HP P6000 EVA Disk Array Family see P6000 EVA Array
HP P9000 XP Disk Array Family see P9000 XP Array

I
incremental ZDB, 50, 61
instant recovery, 46, 53–56
 advantages, 16
 clusters, 56
 introduction, 19
 LVM mirroring, 56
 overview, 52
 process, 54
IR see instant recovery

L
local replication, 22–27
 advantages, 22
 disadvantages, 22
 integrating with LVM mirroring, 27
 snapshot replication, 23
 split mirror replication, 22
locking
 devices, 61
 disks, 61

Logical Volume Manager mirroring see LVM mirroring
LVM mirroring
 EMC, 39, 75
 instant recovery, 56
 local replication, 27
 P6000 EVA Array, 35, 64
 P9000 XP Array, 37, 70

M
Microsoft Exchange Server integration, 40
Microsoft SQL Server integration, 40
mirrors, 22

O
offline backup, 17, 48
online backup, 17, 48

 hot-backup mode, 17, 48
Oracle integration, 40

P
P4000 SAN Solutions, restore, 46
P6000 EVA Array, backup
 local replication, 35, 63
 local replication integrating with LVM mirroring, 35
 mirrorclones, 59
 planning ZDB strategy, 59
 remote plus local replication, 36, 67
 remote plus local replication using LVM mirroring, 64
P6000 EVA Array, configurations
 BC, 35, 63
 CA+BC, 36, 67
 LVM mirroring, 64
P6000 EVA Array, introduction, 35
P6000 EVA Array, restore, 46
 instant recovery, 53
P9000 XP Array, backup, 36–38
 local replication, 37, 68–70
 local replication integrating with LVM mirroring, 37
 local replication using LVM mirroring, 70–71
 planning ZDB strategy, 59
 remote plus local replication, 38, 73–74
 remote replication, 37, 71–72
P9000 XP Array, configurations
 BC, 37, 68
 BC1, 69
 CA, 37, 71
 CA+BC, 38, 73
 cascading, 69
 LVM mirroring, 70
P9000 XP Array, restore
 instant recovery, 46, 53
 split mirror restore, 57
planning ZDB strategy, 58–62
 backup scenarios, 61
 concurrency handling, 61
 flexibility in recovery, 58
 introduction, 58
 snapshot disk arrays, 59
 split mirror disk arrays, 58

R
RAID technology, 21
related documentation, 8
Remote Data Facility configuration see SRDF
remote plus local replication, 29–30
 advantages, 29
 disadvantages, 29
 snapshot replication, 29
 split mirror replication, 29
remote replication, 28
 advantages, 28
 disadvantages, 28
 split mirror replication, 28
replica sets, 44
 rotation, 44

- replicas
 - creating, 17, 43, 49
 - deleting, 47
 - introduction, 17
 - life cycle, 43
 - streaming to tape, 49
 - using, 44, 50
- replication
 - local, 22–27
 - remote, 28
 - remote plus local, 29–30
 - scheduling, 44
 - techniques, 22
- restore from ZDB, 20, 52–57
 - instant recovery, 19, 53–56
 - split mirror restore, 19, 56–57
 - standard Data Protector restore, 19, 52
- roll forward, 41, 53

S

- SAP R/3 integration, 40
- scheduling replication, 44
- single-host configuration, 63, 69
- snapclones, 24, 26
- snapshot replication
 - local, 23–27
 - planning, 59
 - remote plus local, 29–30
- snapshot types
 - snapclones, 24, 26
 - standard snapshots, 24
 - vsnap, 24–25
- source volumes, 17
- split mirror replication
 - local, 22–23
 - mirrors, 22
 - planning, 58
 - remote, 28
 - remote plus local, 29
- split mirror restore, 56–57
 - overview, 53
 - process, 57
- SRDF configuration
 - EMC, 39–40, 77
- SRDF+TimeFinder configuration
 - EMC, 78
- standard Data Protector restore
 - overview, 52
- standard snapshots, snapshot replication, 24
- storage volumes, 21
- Subscriber's Choice, HP, 14
- supported database applications, 40
- supported disk arrays, 18, 63
 - configurations, 63–68

T

- target volumes, 17
- technical support
 - HP, 14

- service locator website, 15
- TimeFinder configuration
 - EMC, 39–40, 74
- transaction logs, 16–17, 19, 41

U

- user interfaces, 33
 - Data Protector CLI, 34
 - Data Protector GUI, 33

V

- virtualization, 16, 21
- virtually capacity-free snapshots see vsnap;
- Volume Shadow Copy Service, 41
- vsnap, 24–25

W

- websites
 - HP, 15
 - HP Subscriber's Choice for Business, 14
 - product manuals, 8

Z

- ZDB agent, 32
- ZDB database, 32, 51
- ZDB to disk, 45
- ZDB to disk+tape, 45
- ZDB to tape, 45
- ZDB, backup process, 48–51
 - creating replicas, 49
 - freezing database applications, 48
 - locating data objects, 48
 - overview, 48
 - recording session information, 50
 - streaming replicas to tape, 49
- ZDB, backup types, 17–18
 - incremental ZDB, 50
 - ZDB to disk, 18, 45
 - ZDB to disk+tape, 18, 45
 - ZDB to tape, 17, 45
- ZDB, introduction, 16–20
 - advantages, 16
 - backup types, 17
 - concepts, 16
 - database application backup, 17
 - replicas, 16
 - replication, 16
 - snapshot backup, 17
 - source volumes, 17
 - split mirror backup, 17
 - target volumes, 17
- ZDB, planning backup strategy, 58–62
 - backup scenarios, 61
 - concurrency handling, 61
 - flexibility in recovery, 58
 - introduction, 58
 - snapshot disk arrays, 59
 - split mirror disk arrays, 58
- zero downtime backup

