

HP Data Protector A.06.10

Zero downtime backup concepts guide

[Build 500]



T B D

Part number: TBD
First edition: TBD Month 2008



i n v e n t

Legal and notice information

© Copyright 2004, 2008 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 US trademark of Sun Microsystems, Inc.

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

UNIX is a registered trademark of The Open Group.

Printed in the US

Contents

Publication history	13
About this guide	15
Intended audience	15
Documentation set	15
Guides	15
Online help	18
Documentation map	19
Abbreviations	19
Map	20
Integrations	21
Document conventions and symbols	22
Data Protector graphical user interface	23
General Information	24
HP technical support	24
Subscription service	25
HP websites	25
Documentation feedback	25
1 Overview	27
Introduction	27
Zero downtime backup	27
Online and offline creation of replicas	28
Creating replicas	29
ZDB types	29
Support on arrays	30
Instant recovery and restoring data from ZDB	31
Instant recovery	31
Other methods of restore from ZDB	31
Restore possibilities for ZDB types	32
2 Replication techniques	33
Disk array basics	33

RAID technology	33
Replication techniques	34
Local replication	35
Split mirror replication	35
Snapshot replication	37
Standard snapshot	38
Vsnap	39
Snapclone	41
Local replication integrating with HP-UX LVM mirroring	42
Remote replication	43
Split mirror replication	44
Remote plus local replication	44
Split mirror replication	45
Snapshot replication	45

3 Using Data Protector for ZDB and instant recovery 47

Data Protector cells	47
Cell components	47
Application systems	48
Backup system	48
ZDB database	49
User interfaces	49
GUI	50
CLI	51
Disk array integrations available with Data Protector	52
HP StorageWorks Disk Array XP	52
Local replication	52
Local replication integrating with LVM mirroring	53
Remote replication	53
Remote plus local replication	54
EMC Symmetrix	55
Local replication	56
Local replication integrating with LVM mirroring	56
Remote replication	56
Remote plus local replication	57
HP StorageWorks Virtual Array	58
VA storage presentation	58
Local replication	58
Local replication integrating with LVM mirroring	59
HP StorageWorks Enterprise Virtual Array	60
EVA storage presentation	60
Local replication	61

Local replication integrating with LVM mirroring	61
Remote plus local replication	62
Application integrations	63
Application data consistency	63
Transaction logs	64
Restore	64
4 The replica life-cycle	65
Overview	65
Creating replicas	67
Replica sets	68
Replica set rotation	68
Scheduling replication	69
Using replicas	69
ZDB to tape	69
ZDB to disk	70
ZDB to disk+tape	70
Instant recovery	71
Deleting replicas	72
5 ZDB session process	73
ZDB process overview	73
Locating data objects	73
Freezing operation of the application or database	74
Creating a replica	75
Replicating the data objects	75
Streaming the replica to tape	75
Backing up a replica to tape	76
Creating mount points	76
Standard data movement to tape	76
Incremental ZDB	76
Direct backup	76
The replica after creation	76
Recording session information	77
Writing session information to the IDB	77
6 Instant recovery and other restore techniques from ZDB sessions	79
Restore process overview	79
Instant recovery	79
Standard Data Protector restore	80

Split mirror restore	81
Instant recovery	81
Instant recovery process	82
Instant recovery and LVM mirroring	83
Instant recovery in a cluster	84
Split mirror restore	84
Split mirror process	84
7 Planning	87
Introduction	87
Flexibility in recovery	87
Split mirror disk arrays	87
Snapshot disk arrays—VA and EVA	88
Snapshot types	88
Standard snapshots	88
Vsnap snapshots	88
Snapclones	89
Other array-specific considerations	89
Replica creation on EVA	89
Replica sets on EVA	89
Concurrency handling	90
Locking	90
Backup device locking	90
Disk locking	90
Backup scenarios	91
A Supported configurations	93
Introduction	93
Supported HP StorageWorks Disk Array XP configurations	94
Local replication configurations	94
Single-Host (BC1) configuration	96
Cascading configurations	96
Local replication configurations with HP-UX LVM mirroring	97
Remote replication configurations	100
Remote plus local replication configurations	102
Cluster configurations	104
Supported EMC Symmetrix configurations	105
Local replication configurations	105
Local replication configurations with HP-UX LVM mirroring	107
Remote replication configurations	110
Remote plus local replication configurations	112
Cluster configurations	114

Supported snapshot configurations	115
Local replication configurations on VA and EVA	115
Local replication configurations with HP-UX LVM mirroring on EVA	117
Local replication configurations with HP-UX LVM mirroring on VA	120
Remote plus local replication configurations on EVA	123
Glossary	125
Index	181

Figures

1 Data Protector graphical user interface	24
2 Zero downtime backup and instant recovery concept	28
3 Disk virtualization	33
4 Disk virtualization with RAID	34
5 Split mirror replication	36
6 Creating a standard snapshot	38
7 Creating a vsnap	40
8 Creating a snapclone	41
9 Data Protector cell set up for ZDB+IR	47
10 Location of software components for ZDB+IR	48
11 Data Protector GUI	51
12 Example BC XP configuration	53
13 Example LVM mirroring configuration XP	53
14 Example CA XP configuration	54
15 CA+BC Configuration in a cluster	55
16 Example TimeFinder configuration	56
17 Example LVM mirroring configuration EMC	56
18 Example SRDF configuration	57
19 Example SRDF+TimeFinder configuration in a cluster	58
20 Example BC snapshot configuration	59
21 Example LVM mirroring configuration VA	60
22 Example LVM mirroring configuration EVA	62
23 Example CA+BC EVA configuration	63
24 Replica life-cycle	66

25	Replica life-cycle	67
26	Instant recovery example	82
27	Split mirror restore example	85
28	BC XP configuration 1	95
29	BC XP configuration 2	95
30	BC XP configuration 3	96
31	BC1 XP configuration	96
32	Cascading configuration	97
33	LVM mirroring configuration 1	98
34	LVM mirroring configuration 2	99
35	LVM mirroring configuration 3	99
36	LVM mirroring configuration 4	100
37	LVM mirroring configuration in a cluster	100
38	CA XP configuration 1	101
39	CA XP configuration 2	101
40	CA XP configuration 3	101
41	CA XP configuration 4	102
42	CA+BC XP configuration 1	103
43	CA+BC XP configuration 2	103
44	CA+BC XP configuration 3	103
45	CA+BC XP configuration 4	104
46	CA+BC configuration in a cluster	104
47	TimeFinder configuration 1	106
48	TimeFinder configuration 2	106
49	TimeFinder configuration 3	107
50	LVM mirroring configuration 1	108
51	LVM mirroring configuration 2	108
52	LVM mirroring configuration 3	109
53	LVM mirroring configuration 4	109

54 LVM mirroring configuration 5	110
55 SRDF configuration 1	111
56 SRDF configuration 2	111
57 SRDF configuration 3	111
58 SRDF configuration 4	112
59 SRDF+TimeFinder configuration 1	113
60 SRDF+TimeFinder configuration 2	113
61 SRDF+TimeFinder configuration 3	113
62 SRDF+TimeFinder configuration 4	114
63 SRDF+TimeFinder configuration in a cluster	114
64 BC snapshot configuration 1	116
65 BC snapshot configuration 2	116
66 BC snapshot configuration 3	117
67 Supported LVM mirroring configuration 1	118
68 Supported LVM mirroring configuration 2	119
69 LVM mirroring configuration 3	120
70 Supported LVM mirroring configuration 1	121
71 Supported LVM mirroring configuration 2	121
72 Supported LVM mirroring configuration 3	122
73 Supported LVM mirroring configuration 4	122
74 LVM mirroring configuration in a cluster	123
75 CA+BC EVA configuration 1	123
76 CA+BC EVA configuration 2	124
77 CA+BC EVA configuration 3	124

Tables

1	Edition history	13
2	Document conventions	22
3	ZDB types v. arrays	30
4	ZDB types v. restore	32
5	Disk arrays integrating with Data Protector	52
6	Backup scenarios	91
7	Disk arrays integrating with Data Protector	93

Publication history

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

Table 1 Edition history

Part number	Guide edition	Product
B6960-90112	October 2004	Data Protector Release A.05.50
B6960-96011	July 2006	Data Protector Release A.06.00
TBD	TBD	Data Protector Release A.06.10

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 manual 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

Other documents and online Help provide related information.

Guides

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

You can find these documents from the Manuals page of the HP Business Support Center website:

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

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

- *HP Data Protector concepts guide*

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

- *HP Data Protector installation and licensing guide*

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

- *HP Data Protector troubleshooting guide*

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

- *HP Data Protector disaster recovery guide*

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

- *HP Data Protector integration guides*

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

- *HP Data Protector integration guide for Microsoft applications: SQL Server, SharePoint Portal Server, Exchange Server, and Volume Shadow Copy Service*

This guide describes the integrations of Data Protector with the following Microsoft applications: Microsoft Exchange Server, Microsoft SQL Server, and Volume Shadow Copy Service.

- *HP Data Protector integration guide for Oracle and SAP*

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

- *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, and Lotus Notes/Domino Server.

- *HP Data Protector integration guide for Sybase, Network Node Manager, Network Data Management Protocol, and VMware*

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

- *HP Data Protector integration guide for HP Service Information Portal*

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

- *HP Data Protector integration guide for HP Reporter*

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

- *HP Data Protector integration guide for HP Operations Manager for UNIX*

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

- *HP Data Protector integration guide for HP Operations Manager for Windows*

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

There are two versions of the guide:

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

- *HP Data Protector software integration guide for HP Performance Manager software and HP Performance Agent software*

This guide provides information about how to monitor and manage the health and performance of the Data Protector environment with HP Performance Manager (PM) software and HP Performance Agent (PA) software on Windows, HP-UX, Solaris and Linux.

- *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 StorageWorks Virtual Array, HP StorageWorks Enterprise Virtual Array, EMC Symmetrix Remote Data Facility and TimeFinder, and HP StorageWorks Disk Array XP. It is intended for backup administrators or operators. It covers the zero downtime backup, instant recovery, and the restore of filesystems and disk images.

- *HP 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, SAP R/3, Microsoft Exchange Server, and Microsoft SQL Server databases. The guide also describes how to configure and use Data Protector to perform backup and restore using the Microsoft Volume Shadow Copy Service.
- *HP Data Protector MPE/iX system user guide*
This guide describes how to configure MPE/iX clients and how to back up and restore MPE/iX data.
- *HP Data Protector Media Operations user guide*
This guide provides tracking and management of offline storage media. It is intended for network administrators responsible for maintaining and backing up systems. It describes the tasks of installing and configuring the application, performing daily media operations and producing reports.
- *HP Data Protector product announcements, software notes, and references*
This guide gives a description of new features of HP Data Protector A.06.10. It also provides information on supported configurations (devices, platforms and online database integrations, SAN, and ZDB), required patches, and limitations, as well as known problems and workarounds. An updated version of the supported configurations is available at <http://www.hp.com/support/manuals>
There are also four other *Product announcements, software notes and references*, which serve a similar purpose for the following:
 - HP Operations Manager software software UNIX integration
 - HP Operations Manager software software Windows integration
 - HP Service Information Portal and HP Reporter software
 - HP Performance Manager software and HP Performance Agent integration
 - HP Media Operations

Online help

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

You can access the online help from the top-level directory on the installation DVD without installing Data Protector:

- **Windows:** Unzip `DP_help.zip` and open `DP_help.chm`.
- **UNIX:** Unpack the zipped tar file `DP_help.tar.gz`, and access the online help system through `DP_help.htm`.

Documentation map

Abbreviations

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

Abbreviation	Guide
CLI	Command line interface reference
Concepts	Concepts guide
DR	Disaster recovery guide
GS	Getting started guide
Help	Online Help
IG-IBM	Integration guide—IBM applications
IG-MS	Integration guide—Microsoft applications
IG-O/S	Integration guide—Oracle, SAP R/3, and SAP DB/MaxDB
IG-OMU	Integration guide—HP Operations Manager software software, UNIX
IG-OMW	Integration guide—HP Operations Manager software software, Windows
IG-PM/PA	Integration guide—Performance Manager and Performance Agent software
IG-Report	Integration guide—HP Reporter software
IG-SIP	Integration guide—HP Service Information Portal
IG-Var	Integration guide—Sybase, Network Node Manager, NDMP and VMware
Install	Installation and licensing guide
MO GS	Media Operations getting started guide

Abbreviation	Guide
MO RN	Media Operations product announcements, software notes, and references
MO UG	Media Operations user guide
MPE/iX	MPE/iX system user guide
PA	Product announcements, software notes, and references
Trouble	Troubleshooting guide
ZDB Admin	ZDB administrator's guide
ZDB Concpt	ZDB concepts guide
ZDB IG	ZDB integration guide

Map

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

Integrations

Look in these guides for details of the following integrations:

Integration	Guide
HP Operations Manager software software	IG-OMU, IG-OMW
HP Performance Manager software	IG-PM/PA
HP Performance Agent software	IG-PM/PA
HP Reporter Light	IG-OMW
HP Reporter software	IG-R
HP Service Information Portal	IG-SIP
HP StorageWorks Disk Array XP	all ZDB
HP StorageWorks Enterprise Virtual Array (EVA)	all ZDB
HP StorageWorks Virtual Array (VA)	all ZDB
IBM DB2 UDB	IG-IBM
Informix	IG-IBM
Lotus Notes/Domino	IG-IBM
Media Operations	MO User
MPE/iX System	MPE/iX
Microsoft Exchange Server	IG-MS, ZDB IG
Microsoft Exchange Single Mailbox	IG-MS
Microsoft SQL Server	IG-MS, ZDB IG
Microsoft Volume Shadow Copy Service (VSS)	IG-MS, ZDB IG
NDMP Server	IG-Var
Network Node Manager (NNM)	IG-Var

Integration	Guide
Oracle	IG-O/S
Oracle ZDB	ZDB IG
SAP DB	IG-O/S
SAP R/3	IG-O/S, ZDB IG
Sybase	IG-Var
Symmetrix (EMC)	all ZDB
VMware	IG-Var

Document conventions and symbols

Table 2 Document conventions

Convention	Element
Blue text: Table 2 on page 22	Cross-reference links and e-mail addresses
Blue, underlined text: http://www.hp.com	website addresses
<i>Italic text</i>	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 text</i>	<ul style="list-style-type: none"> • Code variables • Command variables
text	Emphasized monospace text

 **CAUTION:**

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

 **IMPORTANT:**

Provides clarifying information or specific instructions.

 **NOTE:**

Provides additional information.

 **TIP:**

Provides helpful hints and shortcuts.

Data Protector graphical user interface

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

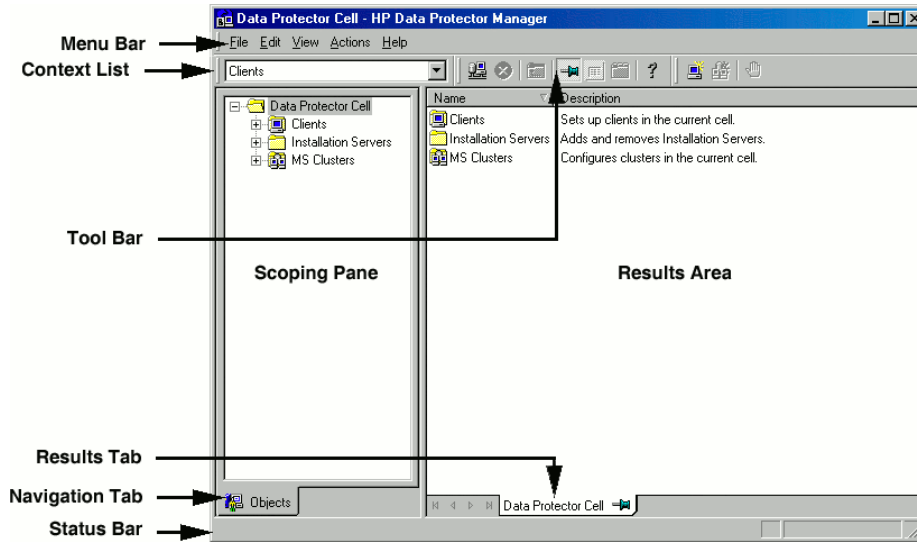


Figure 1 Data Protector graphical user interface

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://www.hp.com/service_locator
- <http://www.hp.com/support/manuals>
- <http://www.hp.com/support/downloads>

Documentation feedback

HP welcomes your feedback.

To make comments and suggestions about product documentation, please send a message to AppRM.DocFeedback@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
- Short restore times (minutes instead of hours)

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, arrays can be divided into many virtual disks. These can easily be copied within an array, perhaps many times dependent on array technology and 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 manual 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/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 more fully in [Chapter 2](#) on page 33):

- **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/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”](#) on page 35.

- **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”](#) on page 37.

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 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”](#) on page 69

1. Data in the replica is streamed to tape according to the tape backup type you have selected (Full, Incr, Incr1-9).
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” on page 70
The replica is kept on the array and used as the backup.
Data can be restored using instant recovery (see “Instant recovery” on page 31), which recovers the complete replica.
- **ZDB to disk+tape** — see “ZDB to disk+tape” on page 70
 1. Data in the replica is streamed to tape according to the tape backup type you have selected (Full, Incr, Incr1-9).
 2. The replica is kept on the 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” on page 31) of the complete replica

Support on arrays

Table 3 ZDB types v. arrays

	Split mirror		Snapshot	
ZDB techniques	XP	EMC	EVA	VA
ZDB to Tape local	Yes	Yes	Yes	Yes
ZDB to Tape remote	Yes	Yes	No	No
ZDB to Tape remote + local	Yes	Yes	Yes	No
ZDB to Disk local	Yes	No	Yes	Yes
ZDB to Disk+Tape local	Yes	No	Yes	Yes

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

- “[Local replication](#)” on page 35
- “[Remote replication](#)” on page 43
- “[Remote plus local replication](#)” on page 44

Instant recovery and restoring data from ZDB

Instant recovery

Instant recovery requires a replica to exist on the same disk array on which the data is to be restored. Application and backup systems are disabled and the contents of the replica are restored directly to their original locations (except in EVA, where the replica becomes the application data). Because the restore is performed internally within the array, it runs at very high speed.

Once the restore is completed, the sections of the database/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 more details, see “[Instant recovery](#)” on page 81.

Other methods of restore from ZDB

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

For details, see the *HP Data Protector concepts guide*.

However, with arrays that support split mirror replication, 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](#)” on page 84.

**NOTE:**

Replicas can be used for purposes other than ZDB+IR, such as data mining. Although Data Protector can create and administer replicas for such purposes, replicas intended for instant recovery should only be used for instant recovery. Otherwise, accurate restore of data cannot be guaranteed.

Restore possibilities for ZDB types

Table 4 ZDB types v. restore

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

From the operating system point of view, the contents of a replica of a particular source volume are the same, whatever 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 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 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 one or more source volumes.

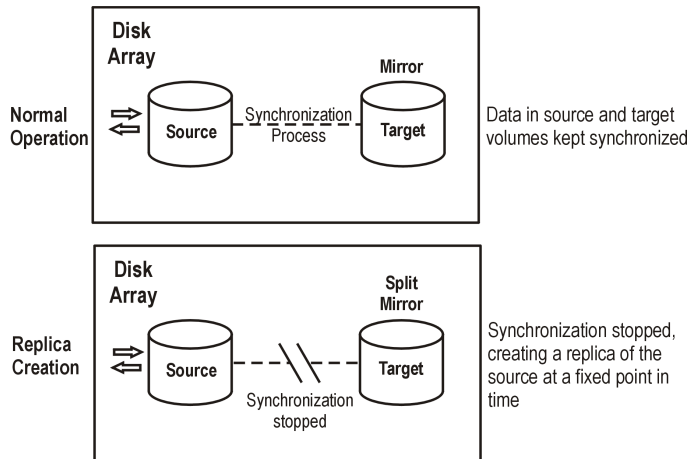


Figure 5 Split mirror replication

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

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 mirror is 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 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, 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 T_0 , storage capacity equal to that taken up by the source volumes concerned is allocated on the 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 T_0 exists in the target volumes and it is ready for use.
2. After snapshot creation, if T_0 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 T_0 still exists.

Characteristics of standard snapshots

- The 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 volumes 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 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:

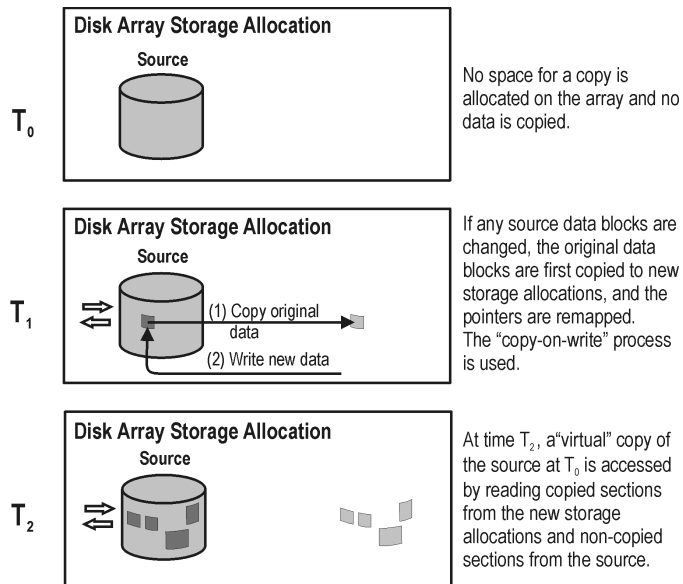


Figure 7 Creating a vsnap

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, if 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 standard snapshots, a vsnap is not an independent duplicate of the original data.
- Vsnaps require independent disk capacity management to guarantee enough space for replica growth. If space on the array runs out, snapshot updates will fail, and it could affect general array operation.
- They are space-efficient. The vsnap only uses the space it needs.
- They are 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 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:

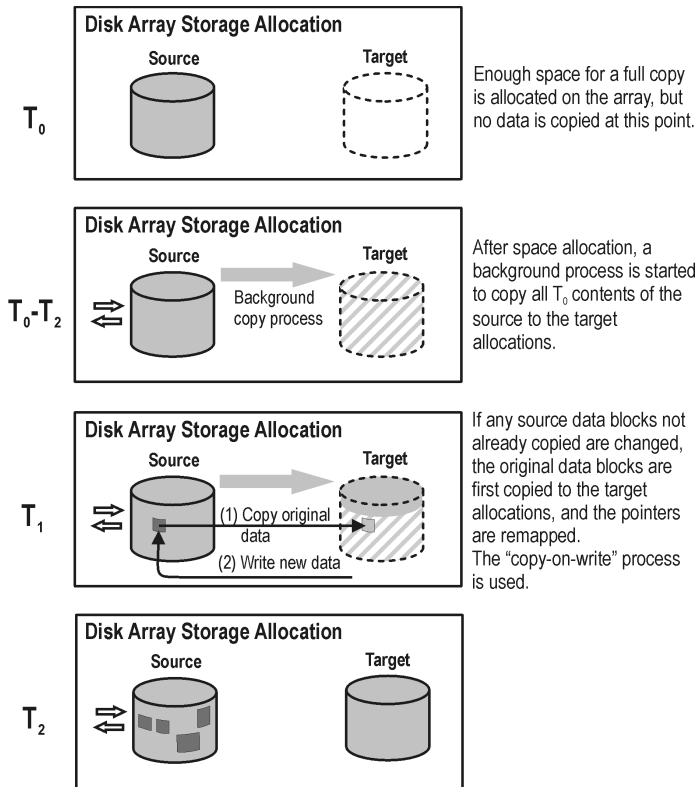


Figure 8 Creating a snapclone

1. A standard snapshot is created, including allocation of enough space for a full copy.
2. A background process starts to copy all unchanged data from source storage blocks to target storage blocks.
3. 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.

4. 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 volumes, 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 volumes are lost or corrupted, the contents of the target volumes 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.
- 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 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. On some arrays, instant recovery for the data backed up in LVM mirroring configurations is not supported.

Remote replication



In **remote replication**, data is replicated on a separate remote array. 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 or the entire computing center.
- 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 array. Once established, the target volumes are kept synchronized with the source volumes on the local 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 array then contains a fixed copy, or independent replica, of the source volumes on the local 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”](#) on page 44.

Remote plus local replication



Remote plus local replication uses both remote and local replication; replicas are created on a remote 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 EVA, 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 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 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 array (the replica of the replica) is then a fixed copy, or independent replica, of the source volumes on the local array.

Snapshot replication

This replication type is supported on EVA.

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 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 array.

Snapshot replicas are created at a particular instant and are immediately available for use. For more information, see “[Snapshot replication](#)” on page 37.



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+IR purposes:




Figure 9 Data Protector cell set up for ZDB+IR

To be able to use ZDB+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 tape library or other tape device is optional for ZDB+IR applications.

Cell components

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




Figure 10 Location of software components for ZDB+IR

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 array on which the application database/filesystem is installed. Each supported type of 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.

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 tape. It also performs various checks and administration functions.

The backup system must have a relevant Data Protector ZDB Agent installed. In some cases, it may also require an Application Integration Agent.

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 array that supports ZDB+IR within Data Protector:

- XPDB for HP StorageWorks Disk Array XP
- VADB for HP StorageWorks Virtual Array
- SMISDB for HP StorageWorks Enterprise Virtual Array

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 array. Generally speaking, each section contains the following types of information:

- Information on replicas kept on disk arrays, including:
 -
- Additional configuration information such as recorded XP command devices (on XP) and defining disk group pair relationships (on EVA).

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 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+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 are defined:




Figure 11 Data Protector GUI

CLI

You can use the CLI to perform most ZDB+IR operations available in the GUI, but some administration 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

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

Replication technologies	Supported disk arrays	Abbreviations
Split mirror	EMC Symmetrix Disk Array	EMC
	HP StorageWorks Disk Array XP	XP
Snapshot	HP StorageWorks Enterprise Virtual Array	EVA
	HP StorageWorks Virtual Array	VA

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

HP StorageWorks Disk Array XP

The following configurations are possible using the Data Protector XP 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 is streamed to tape from the replica after the pair has been split, so that during the backup, the application system remains online and available for use.

Local replication

For local replication, the **HP StorageWorks Business Copy (BC) XP configuration** is used. This enables you to create three **first-level mirrors** for instant recovery purposes, in other words, a replica rotation set of up to three replicas.




Figure 12 Example BC XP configuration

For further examples of XP configurations, see “Supported HP StorageWorks Disk Array XP configurations” on page 94

Local replication integrating with LVM mirroring

The Data Protector XP 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).

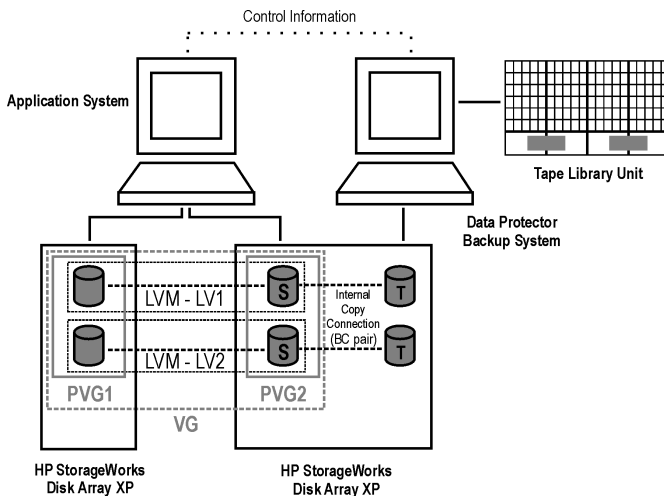


Figure 13 Example LVM mirroring configuration XP

Remote replication

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

The following two types of interfaces are supported for CA 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.

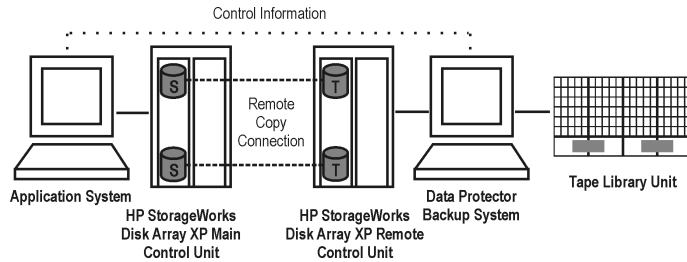


Figure 14 Example CA XP configuration

Remote plus local replication

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

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.

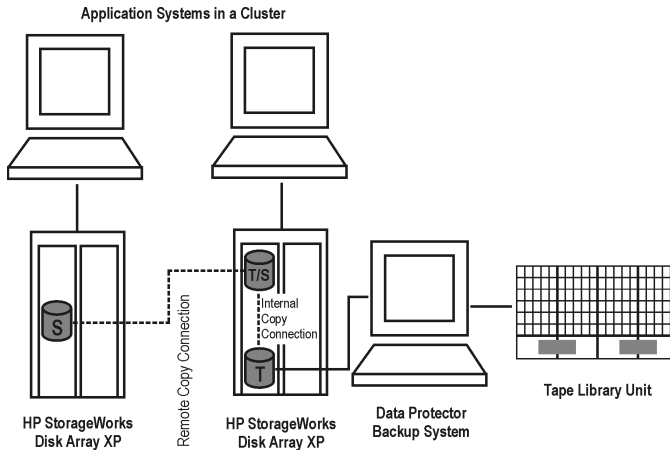


Figure 15 CA+BC Configuration in a cluster

For more information about cluster configurations, see the *HP Data Protector zero downtime backup administrator's guide*.

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](#)" on page 105.

Local replication

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




Figure 16 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.

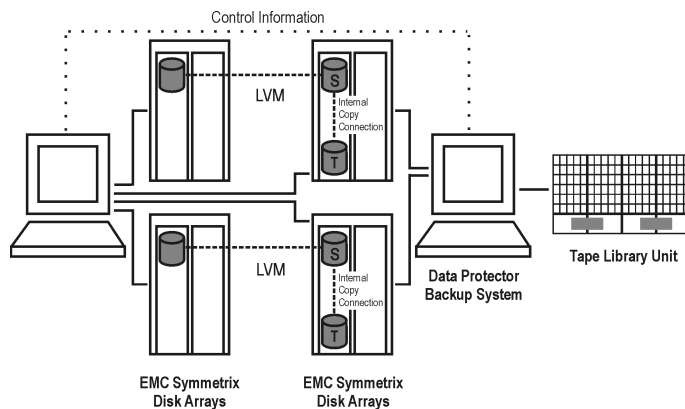


Figure 17 Example LVM mirroring configuration EMC

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 machine.

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 18 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 machine, and then creation of local replicas of those replicas on the remote machine. 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.

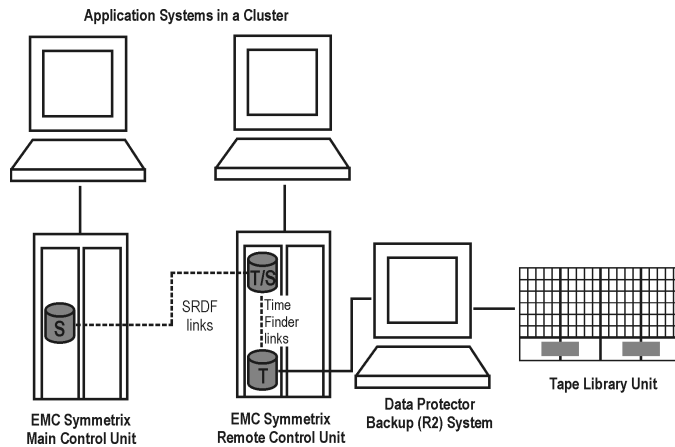


Figure 19 Example SRDF+TimeFinder configuration in a cluster

For more information about cluster configurations, see the *HP Data Protector zero downtime backup administrator's guide*.

HP StorageWorks Virtual Array

The Data Protector VA integration supports the creation of standard snapshots (pre-allocated snapshots).

The following configurations are possible using this integration:

- Local replication
- Local replication integrating with LVM mirroring

For further examples of VA configurations, see [Appendix A](#) on page 93.

VA storage presentation

HP StorageWorks Virtual Array is an array of physical disks configured in such a way that it appears as one or two large blocks of data storage that can be divided into a number of smaller logical storage blocks or LUNs. Data written to a LUN is distributed across the physical disks, providing data redundancy and hence improved data protection.

Local replication

For local replication, the **HP StorageWorks Business Copy (BC) VA configuration** is used. This enables you to create replicas to be used for instant recovery. With this

configuration, you can use large replica sets, the number of members being limited primarily by the available space on the array.

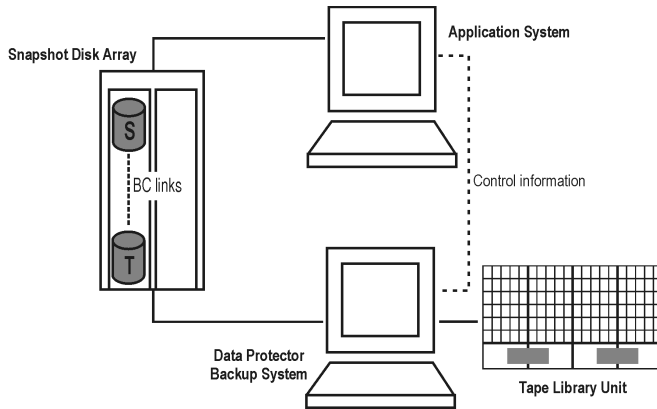


Figure 20 Example BC snapshot configuration

Local replication integrating with LVM mirroring

The Data Protector VA integration supports LVM mirroring in configurations where VA source volumes are LVM mirrored from one (or more) VA(s) to another VA(s). The LVM-mirrored source volumes and their LVM mirrors belong to the same logical volume. The application system has to be connected to the disk arrays containing logical units belonging to the LVM-mirrored logical volumes.

At least two HP StorageWorks Virtual Arrays, located in physically separate sites, are needed for such a configuration.



NOTE:

Instant recovery is not supported when LVM mirroring is used.

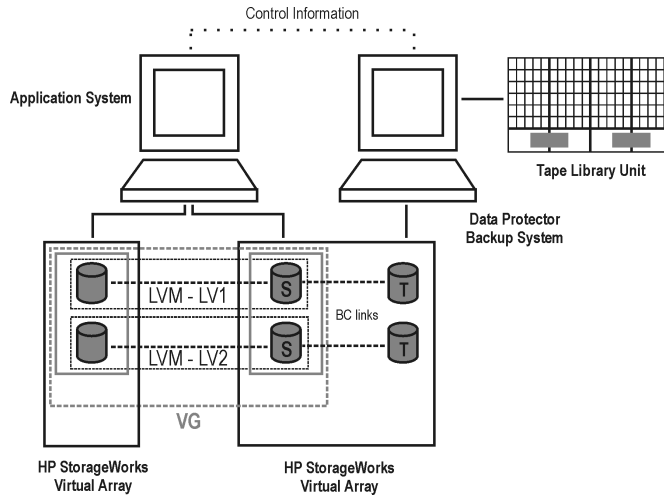


Figure 21 Example LVM mirroring configuration VA

HP StorageWorks Enterprise Virtual Array

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

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

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

For further examples of EVA configurations, see [Appendix A](#) on page 93.

EVA storage presentation

EVA 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 StorageWorks Business Copy (BC) EVA configuration** is used. This enables you to create replicas to be used for instant recovery purposes (note that only snapclones can be used for instant recovery). With this, large replica sets can be used, the number of members being limited primarily by the available space on the array. The maximum number of standard snapshots and vsnaps is limited by the EVA storage system.

Local replication integrating with LVM mirroring

The Data Protector EVA integration supports LVM mirroring in configurations where volume groups are LVM-mirrored from one (or more) EVA(s) to another EVA(s). 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.

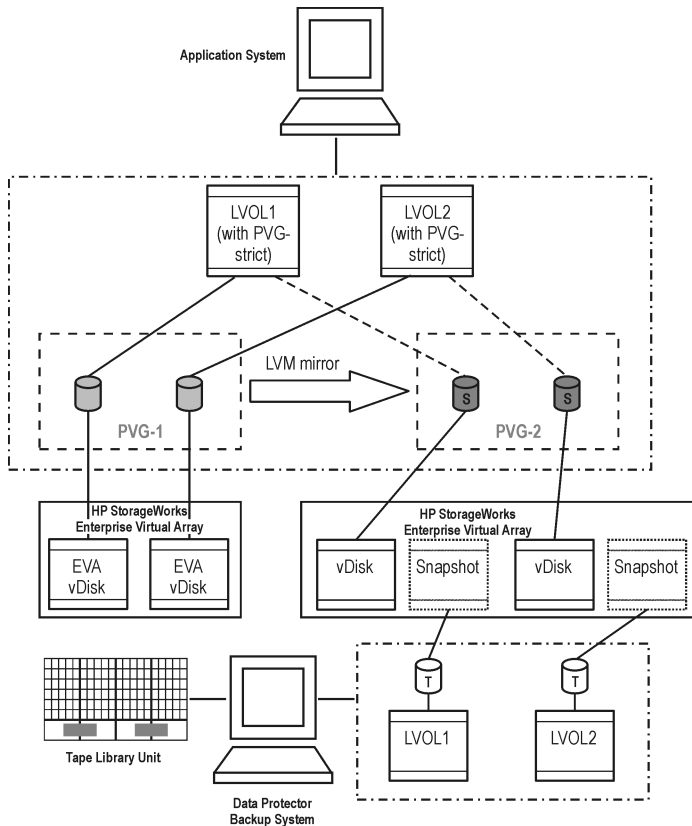


Figure 22 Example LVM mirroring configuration EVA

Remote plus local replication

For remote plus local replication, a combination of BC EVA and HP StorageWorks **Continuous Access (CA)** 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 23 Example CA+BC EVA configuration

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 offline backup

Microsoft Exchange Server and Microsoft SQL Server are also supported through the VSS integration. See *HP Data Protector zero downtime backup integration guide*.

For information on online and offline backup, see [“Freezing operation of the application or database”](#) on page 74.

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. See the Data Protector support matrices for current details.

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 in to backup mode (online backup) 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 logs should not be backed up as part of a replica.

You can back up the logs by scheduling a separate conventional Data Protector backup to tape, but not in parallel with the ZDB session. The logs are backed up after the ZDB session. Alternatively, use the session `post-exec` option to automatically start a standard Data Protector backup to tape for the archived transaction logs.

Restore

See the Data Protector support matrices for details of restore methods available with supported database applications.

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.

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*.

4 The replica life-cycle

Overview

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

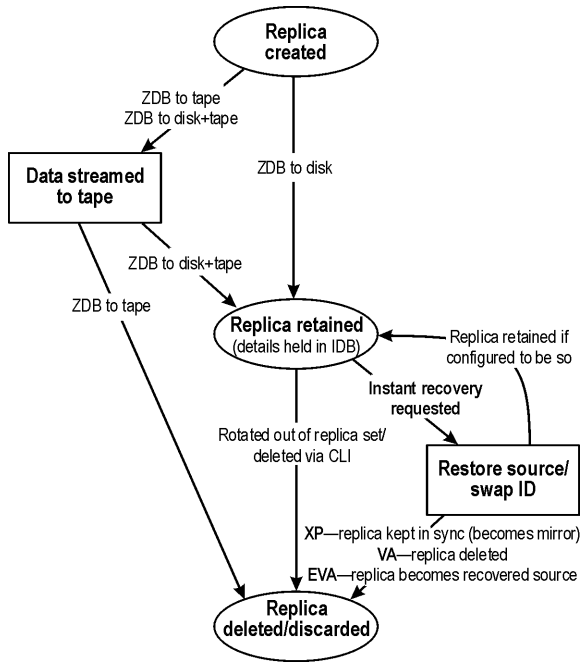


Figure 24 Replica life-cycle




Figure 25 Replica life-cycle

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 array, which can then be presented to a host system.

In all cases, only complete source volumes on the 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](#)” on page 68) to be created
- The type of 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 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.

In Data Protector, the members of a set can undergo **replica set rotation**, either interactively or at times specified in the scheduler.

Replica set rotation

When you create a backup specification for ZDB and instant recovery purposes, you specify the number of replicas in the replica set. Each time the backup is run, a new replica in the replica set is created until the specified number is reached. The next replica to be created then replaces the oldest replica in the set. With some replica

types, this is achieved by directly overwriting the existing replica, in other cases, the oldest replica must be deleted before the new replica is created.

The maximum number of replicas that you can define in a set depends on the type of array being used.

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

For ZDB to disk and ZDB to disk+tape, the replica continues to exist until it is rotated out of the replica set, you delete it through the CLI, or it is used for instant recovery (although you can configure it to be retained afterwards for XP and VA arrays).

ZDB to tape

With ZDB to tape, a replica is normally only kept on an 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 array. You can, however, opt to keep the replica on the array to reserve space on the array for future ZDB-to-tape sessions using the same backup specification. This way, you guarantee there is enough space on the 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 array, freeing the space.	Instant recovery is not possible.
Extensive array support.	

ZDB to disk

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

One or more replicas can be kept on an 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 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/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 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 the disk-based replica created in ZDB to disk and ZDB to disk+tape, 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](#)” on page 81.

What happens to the replica afterwards depends on the array and the configuration:

- **XP:** The replica becomes a mirror in synchronization with the recovered source. You can however configure the instant recovery so that the replica is retained afterwards.
- **VA:** The replica is deleted after recovery, unless you configure the instant recovery so that it is retained.
- **EVA:** The replica becomes the recovered source, and so ceases to exist as a replica.

Deleting replicas

Replicas can be deleted automatically or manually:

- *Automatically:*
 - When a replica becomes the oldest member of a replica rotation set, it is automatically overwritten (or deleted) when a new replica is created in the set. You can however protect replicas on XP and VA using an exclusion list. For more information, see the *HP Data Protector zero downtime backup administrator's guide*.
 - If a replica is used for ZDB to tape, it is automatically deleted after the backup session, unless you specify in the backup specification that it should be retained.
 - Effectively, a replica is deleted after instant recovery, unless configured otherwise. In fact, with XP array, the replica becomes a mirror in synchronization with the recovered source. With EVA, the replica becomes the recovered source and so ceases to exist as a replica.
- *Manually:* When replicas are no longer required within Data Protector, you can delete them from the array using the Data Protector CLI.

5 ZDB session process

ZDB process overview

With conventional Data Protector backup to tape, application operation is affected for the whole of the backup session, until the streaming of data to tape is complete. However, with ZDB, application operation is only affected during the creation of a replica.

The principle steps in a ZDB process are:

1. Locate the data objects for backup. See [“Locating data objects”](#) on page 73.
2. Freeze operation of the application database. See [“Freezing operation of the application or database”](#) on page 74
3. Create a replica containing the specified data objects. See [“Creating a replica”](#) on page 75.
4. If backup to tape is required, stream the replica to tape. See [“Streaming the replica to tape”](#) on page 75.
5. If the ability to restore from disk is required, record information about the session. See [“Recording session information”](#) on page 77.

Locating data objects

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 host 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 database integrations.

For more details, see the *HP Data Protector concepts guide*.

3. The application host is prepared, bringing data into a consistent state. For online backup, the database is quiesced. For offline backup, the database is brought offline. If the ZDB option **Dismount application system** is set ON, 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 host is resumed. Any dismounted filesystem is remounted.
In the case of an offline backup, the database can be put 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 host 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 array then creates the replica on the array.

The method of replication depends on the type of disk array being used, whether local or remote replication is required, and so on. For information on how split mirror and snapshot replication is performed, see [Chapter 2](#) on page 33.

Streaming the replica to tape

1. In ZDB to tape and ZDB to disk+tape, the replica is streamed to tape.
2. The backup host 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.

Direct backup

With certain versions of HP-UX and on XP, it is possible to use Data Protector direct backup functionality to move data directly to a backup device within a SAN environment. For information, see the *HP Data Protector concepts guide* and the online Help index "direct backup environment".

The replica after creation

- With *ZDB to disk* and *ZDB to disk+tape*, the replica remains on the array for instant recovery purposes. If it is part of a replica set, it remains until it is the oldest in the set (excluding any specified in the exclude list for XP or VA). It will

then be replaced by the next ZDB-to-disk or ZDB-to-disk+tape session performed using the same backup specification.

- After a *ZDB to tape* session, when the data has been backed up to tape, by default the replica is automatically deleted. You can opt to keep the replica on the 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*.

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

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

- XPDB for XP
- VADB for VA
- SMISDB for EVA

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

Restore process overview

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 following are available:

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” on page 31.

Because of the different types of replicas involved and various array limitations, the detailed restore process is different for each 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 replications:

- 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 online Help index: “standard restore procedure”.

Split mirror restore

NOTE:

With the speed of today's SAN-attached, very fast tape drives, it is usually quicker to restore directly to the application host than use split mirror restore.

Availability

In local replications on split mirror arrays:

- 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](#)” on page 84.

Instant recovery

With instant recovery, lost or corrupt data is replaced with known good data, which was previously replicated to other volumes on an 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, data in the source volumes is replaced by that in the target volumes internally within the array, involving no other backup medium or device. This makes the restore very fast.

Even though individual backup objects may be specified in the backup specification, it is generally not possible to restore just those individual backup objects; only a complete session can be selected for restore and, hence, only the complete replica can be restored. Thus, not only the originally selected backup objects are restored, but the complete contents of all the volume groups that contained them; their contents will all be returned to the state they were in when the replica was created.

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

Instant recovery process

The following is an example of instant recovery:




Figure 26 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 provided primarily for data security purposes.

Depending on a disk array type, these options enable you to:

- **LVM on HP-UX only:** Check that 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.

- **XP and VA only:** Keep the replica on the array after instant recovery for potential problems with any recovery step after the restore.
 - **EVA only:** Remove the presentation of replicas to any hosts.
3. Optionally, perform a preview of the instant recovery session to provide an extra level of security.
 4. Start the instant recovery.

Data Protector then:

1. Starts processes on the application and backup systems.
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 restore are met (including any instant recovery options specified).
4. Prepares the application system by deactivating any volume groups on HP-UX and dismounts any filesystems associated with the replica.
5. Restores the replica to the original source volumes as follows:
 - On XP, this is done by synchronizing the source volumes with those of the selected split mirror replica.
 - On VA, all other replicas in the replica set created by the associated backup specification are first deleted and their entries removed from the ZDB database.
 - On EVA, the selected snapclone replica is substituted with the original source volumes, which are then deleted, together with their ZDB database entries. 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.
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 BC XP or BC EVA 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, it is necessary to perform additional steps. For information, see the *HP Data Protector zero downtime backup administrator's guide*.

Split mirror restore

NOTE:

With the speed of today's SAN-attached tape drives, it is usually quicker to restore directly to the application host than use split mirror restore.

In split mirror restore, backup objects are first moved from tape to a split mirror replica (either existing or newly created for the purpose) on the backup system. The replica is then re-synchronized with its original source on 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 XP, using the Business Copy (BC) XP configuration.
- On EMC, using the Symmetrix TimeFinder, SRDF, or combined (SRDF+TimeFinder) configurations.

Split mirror process

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




Figure 27 Split mirror restore example

1. Select a mirror to use for the restore. If it is still being synchronized with the source volumes, split the link.
2. If the link is already split, optionally resynchronize the mirror with the source volumes and split the link again, to produce an up-to-date split mirror replica.
3. Restore the required objects from tape to the split mirror replica via the backup system.
4. Synchronize the source volumes with the split mirror replica, effectively replacing the source volumes with the replica.

After synchronization, the contents of the selected mirror 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 mirror was split.

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 array.
- Back up log files regularly.

To control use of disk array space, use a time-based replica set rotation. The number of replicas in the set will depend on the available disk array space and the time range required.

Split mirror disk arrays

The HP StorageWorks Disk Array XP and EMC Symmetrix 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.

See the *HP Data Protector zero downtime backup administrator's guide* for example backup policies.

The following are general recommendations and limitations on the split mirror array's performance.

Snapshot disk arrays—VA and EVA

If you use the Data Protector VA or EVA integrations, consider the following when planning backup strategy:

- Type of snapshot (standard, vsnap or snapclone)—discussed below
- Snapshot policy (strict or loose)—see the *HP Data Protector zero downtime backup administrator's guide*
- Disk array specific considerations—see “[Other array-specific considerations](#)” on page 89
- Instant recovery—see the *HP Data Protector zero downtime backup administrator's guide*

Snapshot types

See “[Snapshot replication](#)” on page 37 for general considerations.

Standard snapshots

VA	Only standard snapshots are supported, so you have to use them for all ZDB types. Pre-allocation of storage prevents running out of space. <i>Important:</i> Since source data will usually be shared by replica and source, certainly when the snapshot is first created, loss or corruption of the source volume will probably make the replica unusable.
EVA	Standard snapshots are intended to be short-lived. You can only use them for ZDB-to-tape sessions because instant recovery is not supported from them.

Vsnap snapshots

Vsnaps are only supported on EVA.

Vsnaps can only be used for ZDB-to-tape sessions because instant recovery is not supported from them.

Snapclones

Snapclones are only supported on EVA.

Snapclones are the only type of snapshot that can be used for instant recovery.

Other array-specific considerations

Replica creation on EVA

A second snapclone for a source volume can only be created if the first has finished. If not, Data Protector automatically retries the operation a configurable number of times at configurable intervals. See the *HP Data Protector zero downtime backup administrator's guide* for details.

Replica sets on EVA

- A replica cannot be reused in the following cases:
 - One of the volumes in the snapclone has a snapshot attached to it.
 - One of the target volumes to be reused is presented to a system.

In such situations, the EVA agent aborts the session.

(With EVA, “reuse” means that the snapshot in the replica set is deleted 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 a replica to be reused is in use and therefore locked by another session, Data Protector reuses the second oldest replica.

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 online 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 split mirror disk arrays), the lock occurs after the mirror copy is prepared (after the split of links between the mirrored disks), but before the mirrored data is moved from a tape device to the mirror copy.

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 2 or 3 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 from the tape.

A Supported configurations

Introduction

This appendix gives you information on the configurations supported on different disk arrays.

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

Table 7 Disk arrays integrating with Data Protector

Replication technology	Supported disk arrays	Abbreviation
Split mirror	HP StorageWorks Disk Array XP	XP
	EMC Symmetrix Disk Array	EMC
Snapshot	HP StorageWorks Enterprise Virtual Array	EVA
	HP StorageWorks Virtual Array	VA

The configurations described are supported by Hewlett-Packard. For an updated list of supported configurations, see the latest support matrices available at <http://www.hp.com/support/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.

For all supported configurations, for ZDB, a backup specification can only cover one single application system and one backup system. You can, however, have multiple 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](#)” on page 76.

Single-host configurations, where a single system is used as both application and backup system, are not recommended because of performance issues. Only disk image and filesystem backups are possible using the single-host configuration. For a list of the supported single-host configurations, see the disk arrays support matrices in the *HP Data Protector product announcements, software notes, and references*.

With all configurations, application 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 StorageWorks Disk Array XP configurations

Local replication configurations

Figures [Figure 28](#) on page 95 through [Figure 30](#) on page 96 are examples of supported local replication configurations on XP.




Figure 28 BC XP configuration 1




Figure 29 BC XP configuration 2

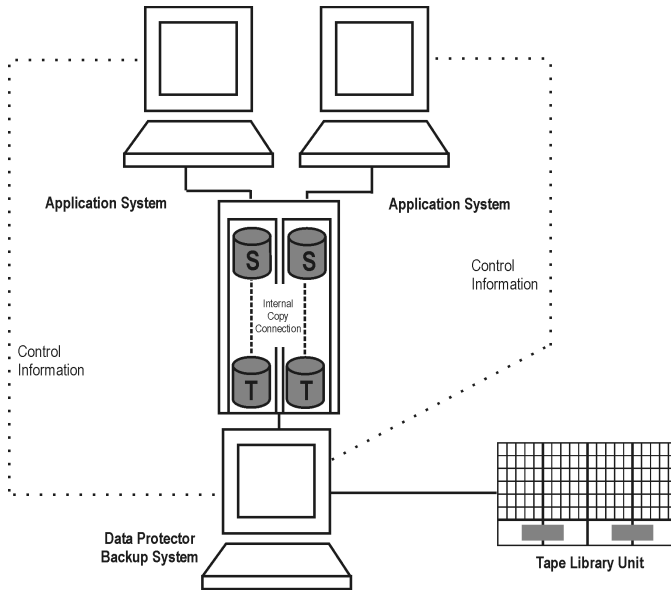


Figure 30 BC XP configuration 3

Single-Host (BC1) configuration

The following figure shows a single-host configuration, also called **BC1 configuration**:




Figure 31 BC1 XP configuration

Cascading configurations

With the Data Protector XP integration, you can create two additional copies for each first-level mirror by using a **cascading configuration**. However, only first-level mirrors (to a maximum of three) are supported for instant recovery purposes. This means you can configure the integration to use cascading, but the additional second-level mirrors (up to six) can only be used for purposes other than ZDB+IR.

The following figure is an example of the cascading configuration, where MU:0, MU:1 and MU:2 are first-level mirrors, and the six mirrors underneath are the second-level mirrors:

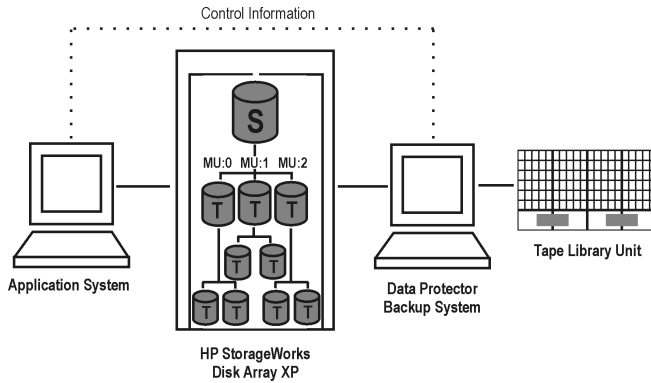


Figure 32 Cascading configuration

Local replication configurations with HP-UX LVM mirroring

Figures [Figure 33](#) on page 98 through [Figure 37](#) on page 100 are examples of supported LVM mirroring configurations on XP:

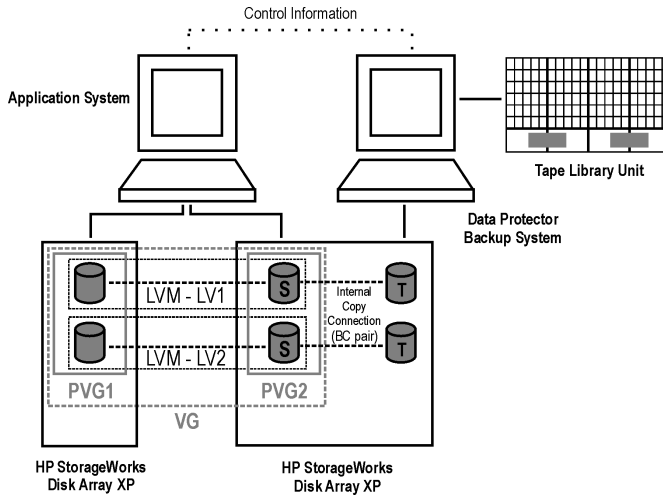


Figure 33 LVM mirroring configuration 1

 **NOTE:**

The Data Protector design requires that either all devices in the PV1 columns have an XP mirror or all of the devices in the PV2 columns have an XP mirror. Mirrors of a logical volume must be fully hosted by physical volumes belonging to the same physical volume group. The physical extents of a mirror must be fully allocated from the physical volumes from the same array.

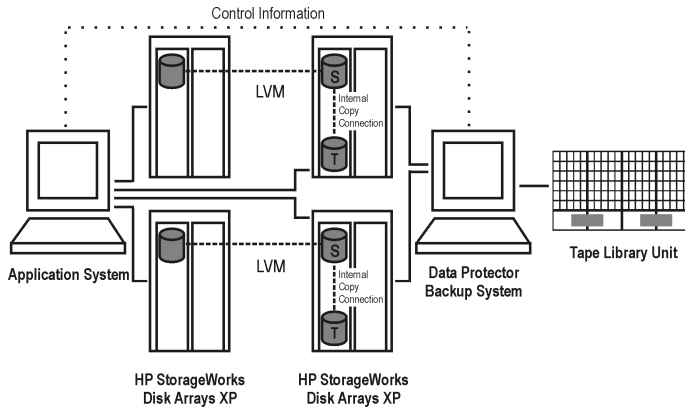


Figure 34 LVM mirroring configuration 2

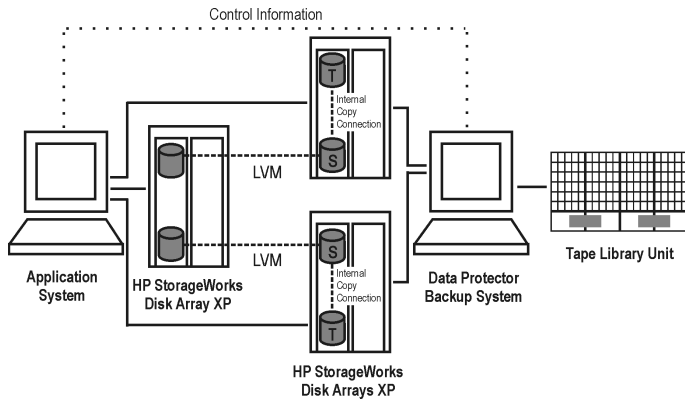


Figure 35 LVM mirroring configuration 3




Figure 36 LVM mirroring configuration 4

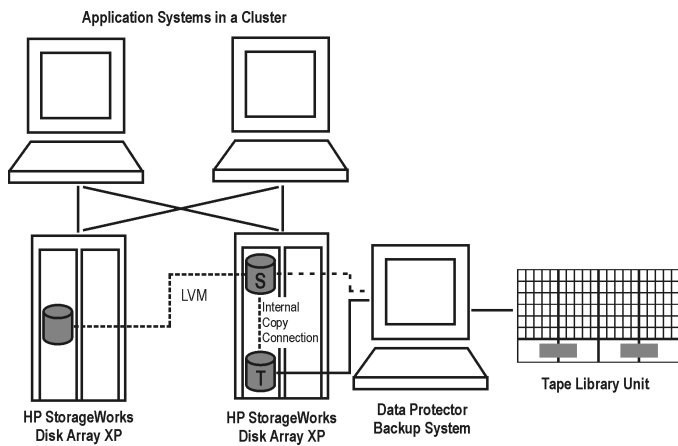


Figure 37 LVM mirroring configuration in a cluster

Remote replication configurations

A single backup system and a single XP array can be used to back up multiple main disk arrays. See [Figure 41](#) on page 102. 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.

Figures [Figure 14](#) on page 54 through [Figure 41](#) on page 102 are examples of supported remote replication configurations on XP:

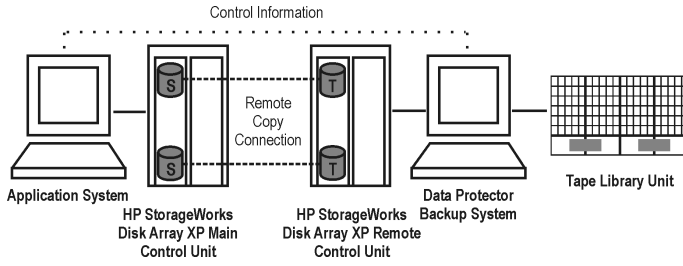


Figure 38 CA XP configuration 1

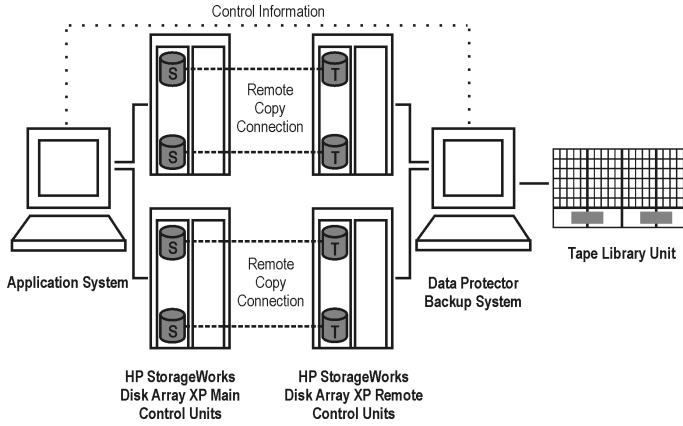


Figure 39 CA XP configuration 2

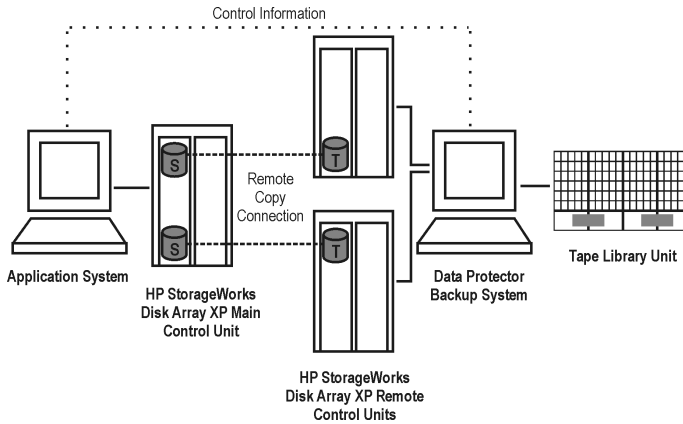


Figure 40 CA XP configuration 3




Figure 41 CA 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 about this, 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.

Figures [Figure 42](#) on page 103 through [Figure 45](#) on page 104 are examples of supported remote plus local replication configurations on XP:




Figure 42 CA+BC XP configuration 1

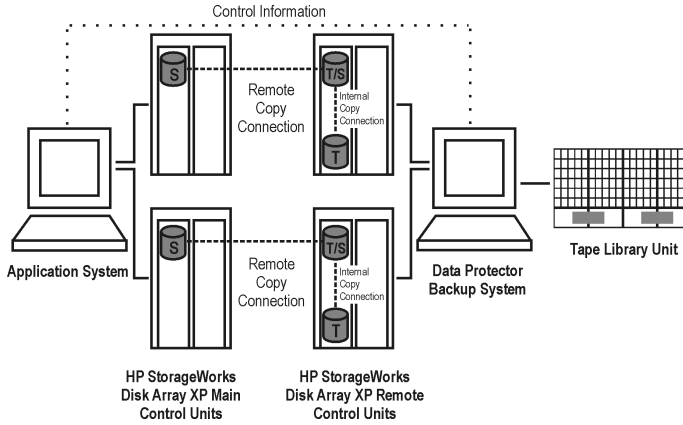


Figure 43 CA+BC XP configuration 2




Figure 44 CA+BC XP configuration 3




Figure 45 CA+BC XP configuration 4

Cluster configurations

The following figure below is an example of a CA+BC configuration in a cluster:

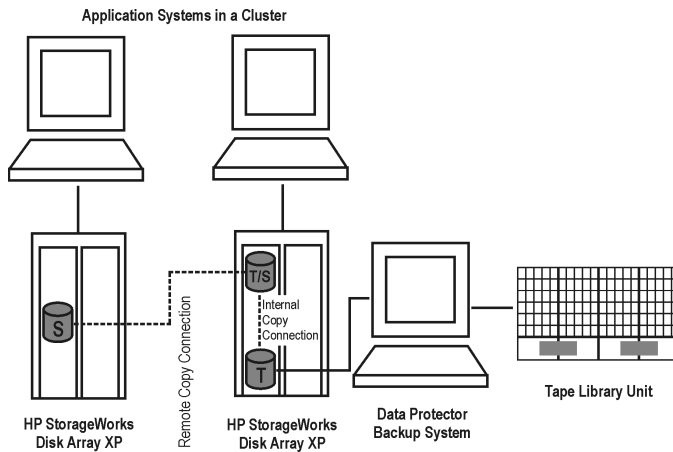


Figure 46 CA+BC configuration in a cluster

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.

Figures [Figure 47](#) on page 106 through [Figure 49](#) on page 107 are examples of supported local replication configurations on EMC:

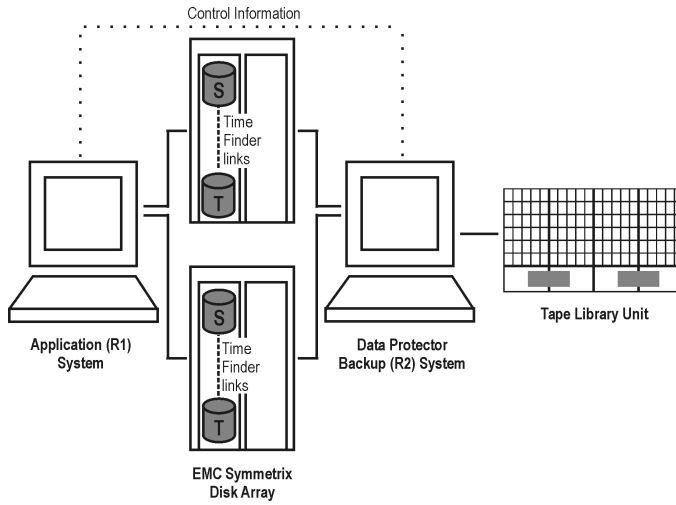


Figure 47 TimeFinder configuration 1




Figure 48 TimeFinder configuration 2

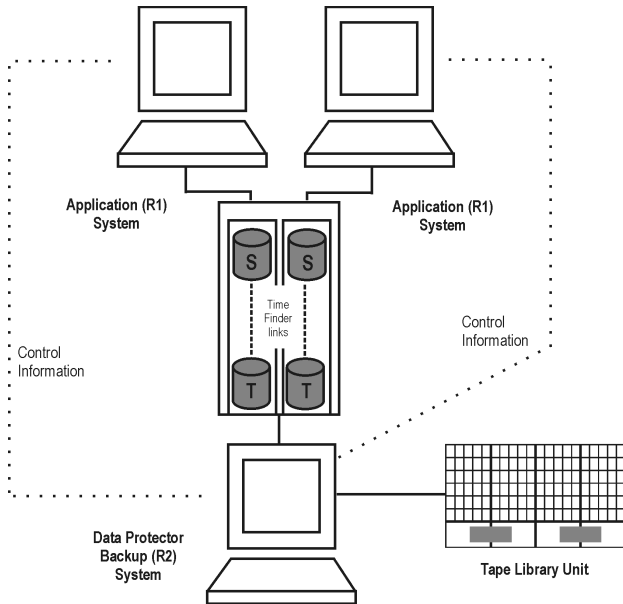


Figure 49 TimeFinder configuration 3

Local replication configurations with HP-UX LVM mirroring

Figures [Figure 50](#) on page 108 through [Figure 54](#) on page 110 are examples of supported LVM mirroring configurations on EMC:

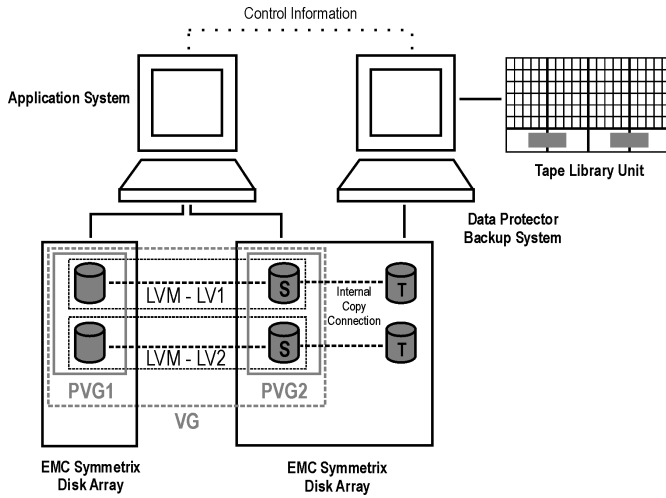


Figure 50 LVM mirroring configuration 1




Figure 51 LVM mirroring configuration 2

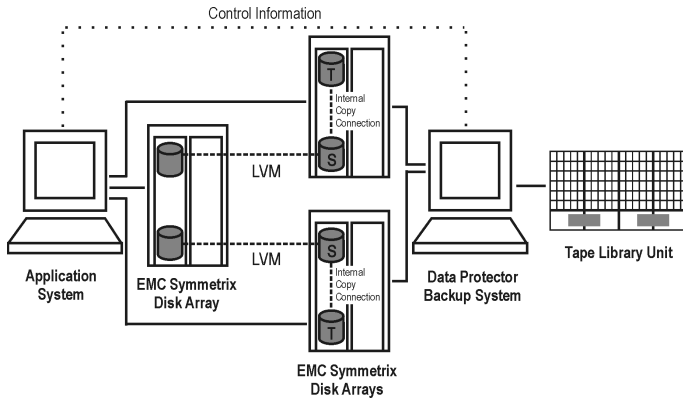


Figure 52 LVM mirroring configuration 3




Figure 53 LVM mirroring configuration 4

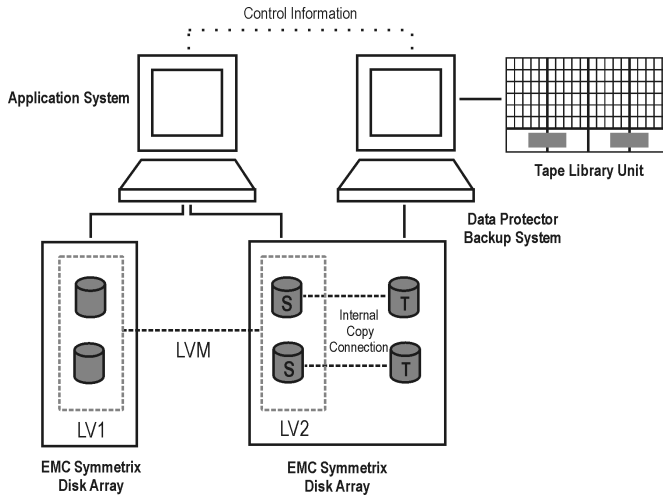


Figure 54 LVM mirroring configuration 5

Remote replication configurations

Figures [Figure 18](#) on page 57 through [Figure 58](#) on page 112 are examples of supported remote replication configurations on EMC:




Figure 55 SRDF configuration 1

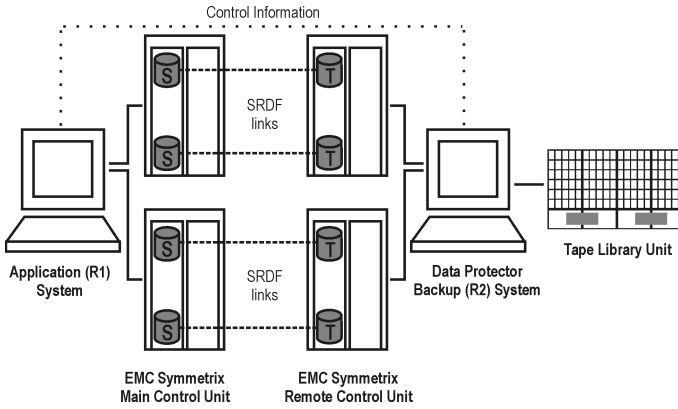


Figure 56 SRDF configuration 2




Figure 57 SRDF configuration 3




Figure 58 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 about this, see the *HP Data Protector zero downtime backup administrator's guide*.

Figures [Figure 59](#) on page 113 through [Figure 19](#) on page 58 are examples of supported remote plus local replication configurations on EMC:




Figure 59 SRDF+TimeFinder configuration 1

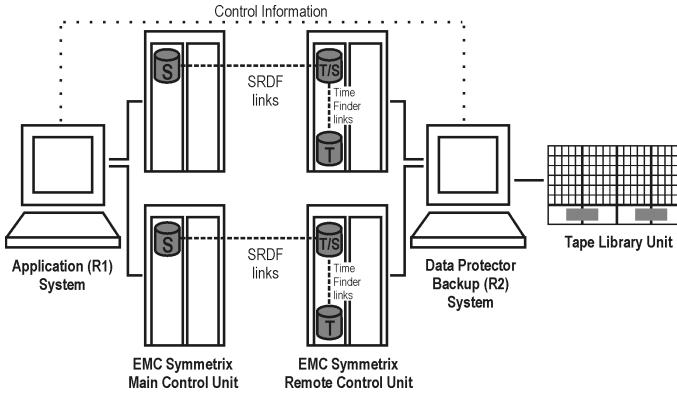


Figure 60 SRDF+TimeFinder configuration 2




Figure 61 SRDF+TimeFinder configuration 3

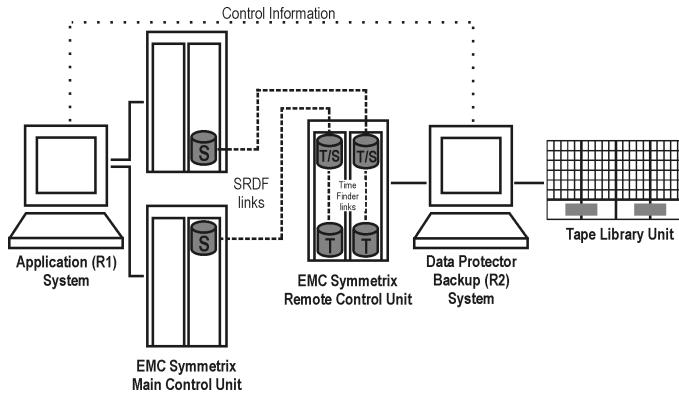


Figure 62 SRDF+TimeFinder configuration 4

Cluster configurations

The following figure is an example of an SRDF+TimeFinder configuration in a cluster:

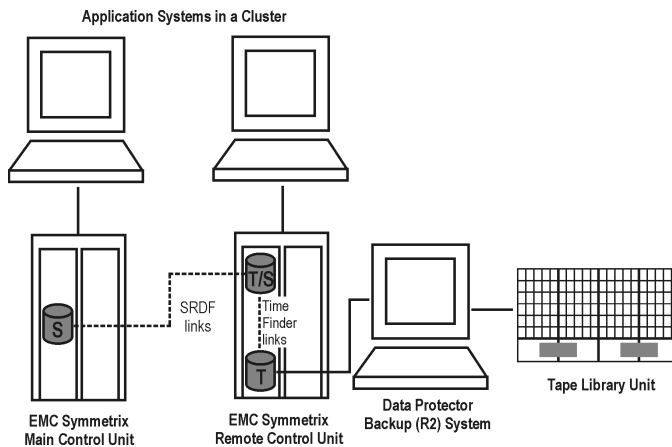


Figure 63 SRDF+TimeFinder configuration in a cluster

For more information about cluster configurations, see the *HP Data Protector zero downtime backup administrator's guide*.

Supported snapshot configurations

Local replication configurations on VA and EVA

For local replication, BC VA or BC EVA configuration is used.

A separate backup system needs to be connected to a disk array. After the replicas are created, Data Protector scans for new disks on the backup system, creates device files (on UNIX), 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.

Figures [Figure 64](#) on page 116 through [Figure 66](#) on page 117 are examples of supported local replication configurations:

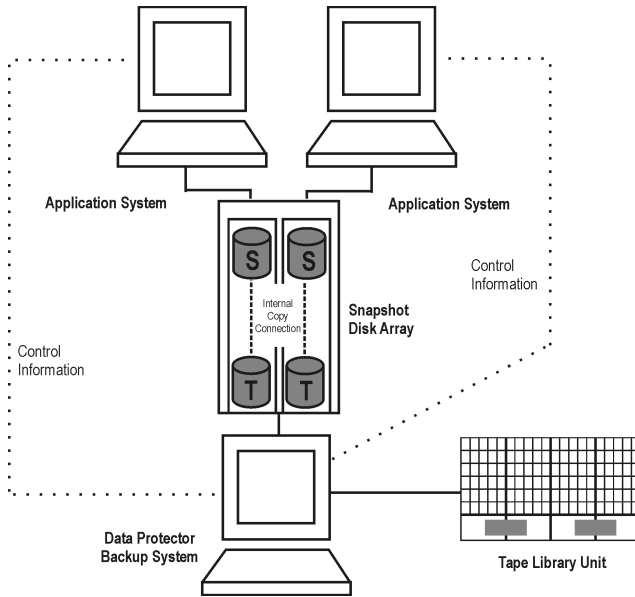


Figure 64 BC snapshot configuration 1




Figure 65 BC snapshot configuration 2

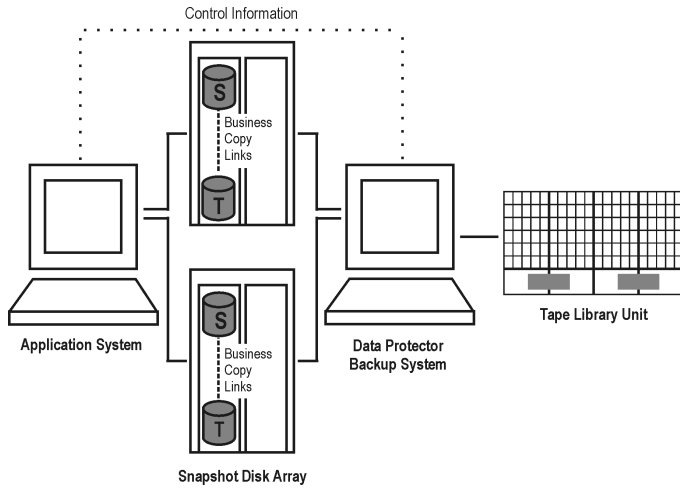


Figure 66 BC snapshot configuration 3

Local replication configurations with HP-UX LVM mirroring on EVA

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.

Figures [Figure 67](#) on page 118 and [Figure 69](#) on page 120 are examples of supported LVM mirroring configurations on EVA:

Figure 67 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 SMI-S Agent always attempts to select a secondary mirror, it will choose PVG-2 for the BC pair replication.

Figure 68 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 69 LVM mirroring configuration 3

Some of the secondary mirror members are hosted by the primary mirror 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*.

Local replication configurations with HP-UX LVM mirroring on VA

Figures [Figure 21](#) on page 60 through [Figure 74](#) on page 123 are examples of supported LVM mirroring configurations on VA:




Figure 70 Supported LVM mirroring configuration 1




Figure 71 Supported LVM mirroring configuration 2




Figure 72 Supported LVM mirroring configuration 3




Figure 73 Supported LVM mirroring configuration 4

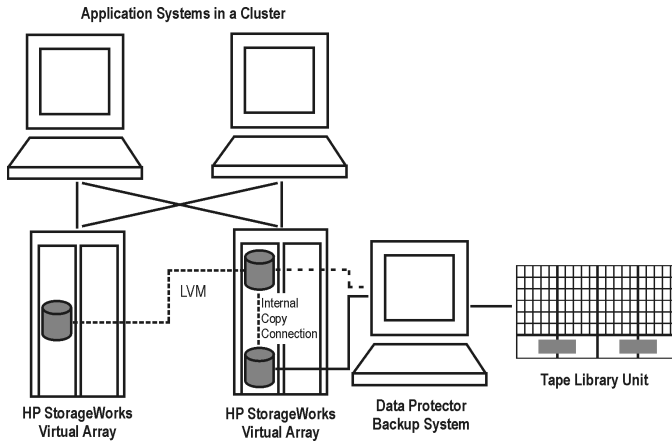


Figure 74 LVM mirroring configuration in a cluster

Remote plus local replication configurations on EVA

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

Figures [Figure 75](#) on page 123 through [Figure 77](#) on page 124 are examples of supported remote plus local configurations on EVA:




Figure 75 CA+BC EVA configuration 1

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




Figure 76 CA+BC EVA configuration 2

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




Figure 77 CA+BC EVA configuration 3

This configuration represents a failover scenario with maintained replica location.

Glossary

access rights	See user rights .
ACSL	<i>(StorageTek specific term)</i> The Automated Cartridge System Library Server (ACSL) software that manages the Automated Cartridge System (ACS).
Active Directory	<i>(Windows specific term)</i> The directory service in a Windows network. It contains information about resources on the network and makes them accessible to users and applications. The directory services provide a consistent way to name, describe, locate, access and manage resources regardless of the physical system they reside on.
AML	<i>(EMASS/GRAU specific term)</i> Automated Mixed-Media library.
application agent	A component needed on a client to back up or restore online database integrations. See also Disk Agent .
application system	<i>(ZDB specific term)</i> A system the application or database runs on. The application or database data is located on source volumes. See also backup system and source volume .
archived redo log	<i>(Oracle specific term)</i> 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](#).

archive logging	<i>(Lotus Domino Server specific term)</i> Lotus Domino Server database mode where transaction log files are overwritten only after they have been backed up.
ASR Set	A collection of files stored on several diskettes, required for proper reconfiguration of the replacement disk (disk partitioning and logical volume configuration) and automatic recovery of the original system configuration and user data that was backed up during the full client backup. These files are stored as an ASR archive file on the Cell Manager (in <code>Data_Protector_home\Config\Server\dr\asr</code> on a Windows Cell Manager or in <code>/etc/opt/omni/server/dr/asr/</code> on a UNIX Cell Manager) as well as on the backup medium. The ASR archive file is extracted to three diskettes for 32-bit Windows systems or four diskettes for 64-bit Windows systems after a disaster occurs. You need these diskettes to perform ASR.
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	<i>(Oracle specific term)</i> Automatic Storage Management is an Oracle 10g integrated filesystem and volume manager that manages Oracle database files. It eliminates complexity associated with managing data and disk and provides striping and mirroring capabilities to optimize performance.
automigration	<i>(VLS specific term)</i> The functionality that allows data backups to be first made to the VLS' virtual tapes and then migrated to physical tapes (one virtual tape emulating one physical tape) without using an intermediate backup application.

See also [Virtual Library System \(VLS\)](#) and [virtual tape](#).

BACKINT	<p>(SAP R/3 specific term) SAP R/3 backup programs can call the Data Protector backint interface program via an open interface, which enables them to communicate with Data Protector software. For backup and restore, SAP R/3 programs issue orders for the Data Protector backint interface.</p>
backup API	<p>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.</p>
backup chain	<p>See restore chain.</p>
backup device	<p>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.</p>
backup generation	<p>One backup generation includes one full backup and all incremental backups until the next full backup.</p>
backup ID	<p>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.</p>
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 (rawdisk). 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 and mount point on UNIX). For integration objects — backup stream identification, indicating the backed up database/application items.• Description: For filesystem objects — uniquely defines objects with identical client name and mount point. For integration

	<p>objects — displays the integration type (for example, SAP or Lotus).</p> <ul style="list-style-type: none"> • 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 (full or incremental). The result of a backup session is a set of media, which was written to, also called the backup or media set.</p> <p>See also backup specification, incremental backup, and full backup.</p>
backup set	A complete set of integration objects associated with a backup.
backup set	<p><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.</p>
backup specification	<p>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.</p>
backup system	<p><i>(ZDB specific term)</i> A system connected to target volumes of one or multiple application systems. The backup system is typically connected to a backup device to perform the backup of the data in a replica.</p> <p>See also application system, target volume, and replica.</p>
backup types	See incremental backup , differential backup , transaction backup , full backup , and delta backup .

backup to IAP	A Data Protector based backup to the HP Integrated Archiving Platform (IAP) appliance. It takes advantage of the IAP capability to eliminate redundancies in the stored data at a block (or chunk) level, by creating a unique content address for each data chunk. Only changed chunks are transmitted over the network and added to the store.
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	<i>(HP StorageWorks Disk Array XP specific term)</i> The Business Copy XP allows to maintain internal copies of HP StorageWorks Disk Array XP LDEVs for purposes such as data backup or data duplication. The copies (secondary volumes or S-VOLs) can be separated from the primary volumes (P-VOLs) and connected to a different system for various purposes, such as backup and development. For backup purposes, P-VOLs should be connected to the application system, and one of the S-VOL mirror sets should be connected to the backup system. See also HP StorageWorks Disk Array XP LDEV, CA, Main Control Unit, application system, and backup system .
BC EVA	<i>(HP StorageWorks EVA specific term)</i> Business Copy EVA is a local replication software solution enabling you to create point-in-time copies (replicas) of the source volumes using the snapshot and clone capabilities of the EVA firmware. See also replica , source volume , snapshot , and CA+BC EVA .
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](#).

BC VA *(HP StorageWorks Virtual Array specific term)* Business Copy VA allows you to maintain internal copies of HP StorageWorks Virtual Array LUNs for data backup or data duplication within the same virtual array. The copies (child or Business Copy LUNs) can be used for various purposes, such as backup, data analysis or development. When used for backup purposes, the original (parent) LUNs are connected to the application system and the Business Copy (child) LUNs are connected to the backup system. See also [HP StorageWorks Virtual Array LUN](#), [application system](#), and [backup system](#).

BCV *(EMC Symmetrix specific term)* Business Continuance Volumes, or BCV devices, are dedicated SLDs that are pre-configured in the ICDA on which the business continuation operation runs. BCV devices are assigned separate SCSI addresses, differing from the addresses used by the SLDs they mirror. The BCV devices are used as splittable mirrors of the primary EMC Symmetrix SLDs that need to be protected. See also [BC](#) and [BC Process](#).

Boolean operators The Boolean operators for the full text search functionality of the online Help system are AND, OR, NOT, and NEAR. Used when searching, they enable you to define your query precisely by creating a relationship between search terms. If no operator is specified in a multi-word search, AND is used by default. For example, the query manual disaster recovery is equivalent to manual AND disaster AND recovery.

boot volume/disk/partition A volume/disk/partition with files required for the initial step of the boot process. Microsoft terminology defines the boot volume/disk/partition as a volume/disk/partition containing the operating system files.

BRARCHIVE *(SAP R/3 specific term)* An SAP R/3 backup tool that allows you to archive redo log files. BRARCHIVE also saves all the logs and profiles of the archiving process. See also [BRBACKUP](#), and [BRRESTORE](#).

BRBACKUP *(SAP R/3 specific term)* An SAP R/3 backup tool that allows an online or offline backup of the control file, of individual data

files, or of all tablespaces and, if necessary, of the online redo log files.

See also [BRARCHIVE](#), and [BRRESTORE](#).

BRRESTORE

(SAP R/3 specific term) An SAP R/3 tool that can be used to restore files of the following type:

- Database data files, control files, and online redo log files saved with BRBACKUP
- Redo log files archived with BRARCHIVE
- Non-database files saved with BRBACKUP

You can specify files, tablespaces, complete backups, log sequence numbers of redo log files, or the session ID of the backup.

See also [BRBACKUP](#), and [BRARCHIVE](#).

BSM

The Data Protector Backup Session Manager controls the backup session. This process always runs on the Cell Manager system.

CA

(HP StorageWorks Disk Array XP specific term) Continuous Access XP allows you to create and maintain remote copies of HP StorageWorks Disk Array XP LDEVs for purposes such as data duplication, backup, and disaster recovery. CA operations involve the main (primary) disk arrays and the remote (secondary) disk arrays. The main disk arrays contain the CA primary volumes (P-VOLs), which contain the original data and are connected to the application system. The remote disk arrays contain the CA secondary volumes (S-VOLs) connected to the backup system.

See also [BC](#) *(HP StorageWorks Disk Array XP specific term)*, [Main Control Unit](#) and [HP StorageWorks Disk Array XP LDEV](#).

CA+BC EVA

(HP StorageWorks EVA specific term) The combination of Continuous Access (CA) EVA and Business Copy (BC) EVA enables you to create and maintain copies (replicas) of the source volumes on a remote EVA, and then use these copies as the source for local replication on this remote array.

See also [BC EVA](#), [replica](#), and [source volume](#).

CAP

(StorageTek specific term) Cartridge Access Port is a port built into the door panel of a library. The purpose is to enter or eject media.

catalog protection	Defines how long information about backed up data (such as file names and file versions) is kept in the IDB. See also data protection .
CDB	The Catalog Database is a part of the IDB that contains information about backups, object copies, restores, media management sessions,, and backed up data. Depending on the selected logging level, it also contains file names and file versions. This part of the database is always local to the cell. See also MMDB .
CDF file	<i>(UNIX specific term)</i> A Context Dependent File is a file consisting of several files grouped under the same pathname. The system ordinarily selects one of the files using the context of the process. This mechanism allows machine dependent executables, system data, and device files to work correctly from all hosts in a cluster while using the same pathname.
cell	A set of systems that are under the control of a Cell Manager. The cell typically represents the systems on a site or an organizational entity, which are connected to the same LAN. Central control is available to administer the backup and restore policies and tasks.
Cell Manager	The main system in the cell where the essential Data Protector software is installed and from which all backup and restore activities are managed. The GUI used for management tasks can be located on a different system. Each cell has one Cell Manager system.
centralized licensing	Data Protector allows you to configure centralized licensing for the whole enterprise environment consisting of several cells. All Data Protector licenses are installed and kept on the Enterprise Cell Manager system. You can then allocate licenses to specific cells to suit your needs. See also MoM .
Centralized Media Management Database (CMMDB)	See CMMDB .

Change Journal	<i>(Windows specific term)</i> 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	<i>(Windows specific term)</i> A module that can be queried to determine which objects on a filesystem have been created, modified, or deleted.
channel	<p><i>(Oracle specific term)</i> 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>
chunking	<p><i>(IAP specific term)</i> The process of dividing data into blocks (chunks), where each chunk gets a unique content address. This address is then used to determine whether a particular chunk is already backed up to the IAP appliance. If the duplicate data is identified (two addresses are identical, that is the address is the same as for another data chunk already stored into IAP), it is not backed up. This way, the data redundancy is eliminated and the optimal data storage is achieved.</p> <p>See also backup to IAP.</p>
circular logging	<i>(Microsoft Exchange Server and Lotus Domino Server specific term)</i> Circular logging is a Microsoft Exchange Server database and Lotus Domino Server database mode, in which the transaction log file contents are periodically overwritten after the corresponding data is committed to the database. Circular logging reduces disk storage space requirements.
client backup	A backup of all filesystems mounted on a client. Filesystems mounted on the client after the backup specification was created are not automatically detected.
client backup with disk discovery	A backup of all filesystems mounted on a client. When the backup starts, Data Protector discovers the disks on the clients. Client backup with disk discovery simplifies backup configuration

and improves backup coverage of systems that often mount or dismount disks.

client or **client system** Any system configured with any Data Protector functionality and configured in a cell.

cluster-aware application It is an application that supports the cluster Application Programming Interface. Each cluster-aware application declares its own critical resources (disk volumes (on Microsoft Cluster Server), volume groups (on MC/ServiceGuard), application services, IP names and addresses, and so on).

Cluster Continuous Replication (*Microsoft Exchange Server specific term*) Cluster continuous replication (CCR) is a high availability solution that uses cluster management and failover options to create and maintain an exact copy (CCR copy) of a storage group. A storage group is replicated to a separate server. CCR removes any single point of failure in your Exchange back-end servers. You can perform backups using VSS on your passive Exchange Server node where a CCR copy is located and thus reducing the load on the active node.

A CCR copy is used for disaster recovery since you can switch to the CCR copy in a few seconds. A replicated storage group is represented as a new instance of Exchange writer called Exchange Replication Service and can be backed up (using VSS) like an ordinary storage group.

See also [Exchange Replication Service](#) and [Local Continuous Replication](#).

CMD Script for Informix Server (*Informix Server specific term*) A Windows CMD script that is created in INFORMIXDIR when an Informix Server database is configured. The CMD script is a set of system commands that export environment variables for Informix Server.

CMMDB The Data Protector Centralized Media Management Database (CMMDB) is the result of merging MMDBs from several cells in the MoM environment. It allows you to share high-end devices and media across multiple cells in a MoM environment. One cell can control the robotics, serving the devices that are connected to systems in other cells. The CMMDB must reside on the MoM Manager. A reliable network connection between the MoM cell and the other Data Protector cells is highly recommended
See also [MoM](#).

COM+ Registration Database	<i>(Windows specific term)</i> The COM+ Registration Database and the Windows Registry store COM+ application attributes, class attributes, and computer-level attributes. This guarantees consistency among these attributes and provides common operation on top of these attributes.
command-line interface (CLI)	A set of DOS and UNIX like commands that you can use in shell scripts to perform Data Protector configuration, backup, restore, and management tasks.
Command View (CV) EVA	<i>(HP StorageWorks EVA specific term)</i> The user interface that enables you to configure, manage, and monitor your HP StorageWorks EVA storage system. It is used to perform various storage management tasks, for example, creating virtual disk families, managing storage system hardware, and creating snapclones and snapshots of virtual disks. The Command View EVA software runs on the HP OpenView Storage Management Appliance, and is accessed by a Web browser. See also HP StorageWorks EVA SMI-S Agent and HP StorageWorks SMI-S EVA provider .
Command View VLS	<i>(VLS specific term)</i> A web browser-based GUI that is used to configure, manage, and monitor the VLS through a LAN. See also Virtual Library System (VLS) .
concurrency	See Disk Agent concurrency .
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	<i>(HP StorageWorks EVA specific term)</i> A pair that consists of the source volumes on a local EVA and their replica on a remote EVA. See also source volume , replica , and CA+BC EVA
CRS	The Cell Request Server process (service), which runs on the Data Protector Cell Manager, and starts and controls the backup and restore sessions. The service is started as soon as Data Protector is installed on the Cell Manager. On Windows systems, the CRS runs under the account of the user specified at installation time. On UNIX systems, it runs under the account <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	<i>(Oracle and SAP R/3 specific term)</i> 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. <i>See also catalog protection.</i>
data stream	Sequence of data transferred over the communication channel.
database library	A Data Protector set of routines that enables data transfer between Data Protector and a server of an online database integration, for example, 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.
Data Replication (DR) group	<i>(HP StorageWorks EVA specific term)</i> A logical grouping of EVA virtual disks. It can contain up to eight copy sets provided they have common characteristics and share a common CA EVA log. <i>See also copy set.</i>
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.
Dobject	<i>(Informix Server specific term)</i> An Informix Server physical database object. It can be a blob space, db space, or logical log file.
DC directory	The Detail Catalog (DC) directory contains DC binary files, which store information about file versions. It represents the DCBF part of the IDB, which occupies approximately 80% of the IDB. The default DC directory is called the <code>dcbf</code> directory and is located in the <code>Data_Protector_home\db40</code> directory on a Windows Cell Manager and in the <code>/var/opt/omni/server/db40</code> directory on a UNIX Cell

Manager. You can create more DC directories and use a custom location. Up to 10 DC directories are supported per cell. The default maximum size of a DC directory is 16 GB.

DCBF	The Detail Catalog Binary Files (DCBF) part of the IDB stores information about file versions and attributes. It occupies approximately 80% of the IDB. One DC binary file is created for each Data Protector medium used for backup. Its maximum size is limited by the file system settings.
delta backup	A delta backup is a backup containing all the changes made to the database from the last backup of any type. See also backup types .
device	A physical unit which contains either just a drive or a more complex unit such as a library.
device chain	A device chain consists of several standalone devices configured for sequential use. When a medium in one device gets full, the backup automatically continues on a medium in the next device in the device chain.
device group	<i>(EMC Symmetrix specific term)</i> 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.
direct backup	A SAN-based backup solution in which data movement directly from disk to tape (or to other secondary storage) is facilitated by the SCSI Extended Copy (Xcopy) command. Direct backup lessens the backup I/O load on systems in a SAN environment. The data movement is facilitated directly from disk to tape (or to other secondary storage) by the SCSI Extended Copy (XCOPY) command. The command is provided by any element of the infrastructure including bridges, switches, tape libraries, and disk subsystems. See also XCOPY engine .
directory junction	<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.
Disk Agent	A component needed on a client to back it up and restore it. The Disk Agent controls reading from and writing to a disk. During a backup session, the Disk Agent reads data from a disk and sends it to the Media Agent, which then moves it to the device. During a restore session the Disk Agent receives data from the Media Agent and writes it to the disk.
Disk Agent concurrency	The number of Disk Agents that are allowed to send data to one Media Agent concurrently.
disk discovery	The detection of disks during client backup with disk discovery. During this backup, Data Protector discovers (detects) the disks that are present on the client — even though they might not have

been present on the system when the backup was configured — and backs them up. This is particularly useful in dynamic environments, where configurations change rapidly. After the disks are expanded, each inherits all options from its master client object. Even if pre- and post-exec commands are specified once, they are started many times, once per each object.

disk group	<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 (rawdisk) backup	A high-speed backup where Data Protector backs up files as bitmap images. A disk image (rawdisk) backup does not track the files and directory structure stored on the disk, but stores a disk image structure on byte level. You can perform a disk image backup of either specific disk sections or a complete disk.
disk quota	A concept to manage disk space consumption for all or a subset of users on a computer system. This concept is used by several operating system platforms.
disk staging	The process of backing up data in several phases to improve the performance of backups and restores, reduce costs of storing the backed up data, and increase the data availability and accessibility for restore. The backup stages consist of backing up data to one media type first (for example disk) and later copying it to a different media type (for example tape).
distributed file media format	A media format, available with the file library, which supports a space efficient type of synthetic backup called virtual full backup. Using this format is a prerequisite for virtual full backup. See also virtual full backup .
Distributed File System (DFS)	A service that connects file shares into a single namespace. The file shares can reside on the same or on different computers. DFS provides client access to the resources in a location-transparent manner.
DMZ	The Demilitarized Zone (DMZ) is a network inserted as a "neutral zone" between a company's private network (intranet) and the outside public network (Internet). It prevents outside users from getting direct access to company servers in the intranet.

DNS server	In the DNS client-server model, this is the server containing information about a portion of the DNS database that makes computer names available to client resolvers querying for name resolution across the Internet.
domain controller	A server in a network that is responsible for user security and verifying passwords within a group of other servers.
DR image	Data required for temporary disaster recovery operating system (DR OS) installation and configuration.
DR OS	A disaster recovery operating system is an operating system environment in which disaster recovery runs. It provides Data Protector a basic runtime environment (disk, network, tape, and filesystem access). The OS has to be installed and configured before the Data Protector disaster recovery can be performed. DR OS not only hosts the Data Protector disaster recovery process but is also a part of the restored system because it replaces its own configuration data with the original configuration data.
drive	A physical unit that receives data from a computer system and can write it onto a magnetic medium (typically a tape drive). It can also read the data from the medium and send it to the computer system.
drive index	A number that identifies the mechanical position of a drive inside a library device. This number is used by the robotic control to access a drive.
dynamic client	See client backup with disk discovery .
EMC Symmetrix Agent (SYMA) (EMC Symmetrix specific term)	See Symmetrix Agent (SYMA) .
emergency boot file	<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) or <code>INFORMIXDIR\etc</code> (on UNIX). <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.

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 Event Log is accessible only to Data Protector users in the <code>Admin</code> group and to Data Protector users who are granted the <code>Reporting and notifications</code> user rights. You can view or delete all events in the Event Log.
Event Logs	Files in which Windows logs all events, such as the starting or stopping of services and the logging on and off of users. Data Protector can back up Windows Event Logs as part of the Windows configuration backup.
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. <i>See also Cluster Continuous Replication and Local Continuous Replication.</i>
exchanger	Also referred to as SCSI Exchanger. <i>See also library.</i>
exporting media	A process that removes all data about backup sessions, such as systems, objects, and file names, which reside on the media from the IDB. Information about the media and their relation to a pool is also removed from the IDB. The data on the media remains unchanged. <i>See also importing media.</i>

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.
failover	Transferring of the most important cluster data, called group (on Windows) or package (on UNIX) from one cluster node to another. A failover can occur mostly because of software or hardware failures or maintenance on the primary node.
failover	<i>(HP StorageWorks EVA specific term)</i> An operation that reverses the roles of source and destination in CA+BC EVA configurations. See also CA+BC 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 StorageWorks Disk Array XP specific term)</i> HP StorageWorks Disk Array XP allows up to three mirror copies of a primary volume and each of these copies can have additional two copies. The three mirror copies are called first-level mirrors. See also primary volume and MU number .
flash recovery area	<i>(Oracle specific term)</i> Flash recovery area is an Oracle 10g managed directory, filesystem, or Automatic Storage Management disk group that serves as a centralized storage area for files related to backup and recovery (recovery files). See also recovery files .
fnames.dat	The <code>fnames.dat</code> files of the IDB contain information on the names of the backed up files. Typically, these files occupy about 20% of the IDB, if filenames are stored.
formatting	A process that erases any data contained on a medium and prepares it for use with Data Protector. Information about media (medium ID, description, and location) is saved in the IDB as well as on the respective media (media header). Data Protector media with protected data are not formatted until the protection expires or the media are unprotected/recycled.
free pool	An auxiliary source of media for use by media pools when they run out of media. The media pools must be configured to use free pools.
full backup	A backup in which all selected objects are backed up, whether or not they have been recently modified. See also backup types .

full database backup	A backup of all data in a database, not only the data that has been changed after the last (full or incremental) backup of the database. A full database backup does not depend on any other backup.
full mailbox backup	A full mailbox backup is a backup of the entire mailbox content.
full ZDB	A ZDB to tape or ZDB to disk+tape session in which all selected objects are streamed to tape, even if there are no changes from the previous backup. See also incremental ZDB .
global options file	A file that allows you to customize Data Protector. It explains the global options, which cover various aspects of Data Protector, typically time-outs and limits, and affect the entire Data Protector cell. The file is located in the <code>/etc/opt/omni/server/options</code> directory on HP-UX and Solaris systems and in the <code>Data_Protector_home\Config\Server\Options</code> directory on Windows systems.
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 cross-platform (HP-UX, Solaris, Linux, and Windows) graphical user interface, provided by Data Protector for easy access to all configuration, administration, and operation tasks. Besides the original Data Protector GUI, Data Protector also provides a Java-based graphical user interface with the same look and feel. As Java can run on numerous platforms, the Data Protector Java GUI is supported on a larger number of platforms than the original Data Protector GUI.
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: <code>/etc/opt/omni/server/Holidays</code> on the UNIX Cell Manager and <code>Data_Protector_home\Config\Server\holidays</code> on the Windows Cell Manager.
host backup	See client backup with disk discovery .
hosting system	A working Data Protector client used for Disk Delivery Disaster Recovery with a Data Protector Disk Agent installed.
HP ITO	See OM .
HP OpC	See OM .
HP OpenView SMART Plug-In (SPI)	A fully integrated, out-of-the-box solution which "plugs into" HP Operations Manager software, extending the managed domain. Through the Data Protector integration, which is implemented as an HP OpenView SMART Plug-In, a user can have an arbitrary number of Data Protector Cell Managers monitored as an extension to HP Operations Manager software (OM).
HP OM	See OM
HP StorageWorks Disk Array XP LDEV	A logical partition of a physical disk within an HP StorageWorks Disk Array XP. LDEVs are entities that can be replicated in the Continuous Access XP (CA) and Business Copy XP (BC) configurations, or can be used as standalone entities. See also BC , CA (HP StorageWorks Disk Array XP specific term), and replica .
HP StorageWorks EVA SMI-S Agent	A Data Protector software module that executes all tasks required for the HP StorageWorks Enterprise Virtual Array integration. With the EVA SMI-S Agent, the control over the array is established through HP StorageWorks SMI-S EVA provider, which directs communication between incoming requests and CV EVA.

See also [Command View \(CV\) EVA](#) and [HP StorageWorks SMI-S EVA provider](#).

HP StorageWorks SMI-S EVA provider

An interface used for controlling HP StorageWorks Enterprise Virtual Array. SMI-S EVA provider runs as a separate service on the HP OpenView Storage Management Appliance system and acts as a gateway between incoming requests and Command View EVA. With the Data Protector HP StorageWorks EVA integration, SMI-S EVA provider accepts standardized requests from the EVA SMI-S Agent, communicates with Command View EVA for information or method invocation, and returns standardized responses.

See also [HP StorageWorks EVA SMI-S Agent](#) and [Command View \(CV\) EVA](#).

HP StorageWorks Virtual Array LUN

A logical partition of a physical disk within an HP StorageWorks Virtual Array. LUNs are entities that can be replicated in the HP StorageWorks Business Copy VA configuration, or can be used as standalone entities.

See also [BC VA](#) and [replica](#).

HP VPO

See [OM](#).

ICDA

(EMC Symmetrix specific term) EMC's Symmetrix Integrated Cached Disk Arrays (ICDA) is a disk array device that combines a set of physical disks, a number of FWD SCSI channels, an internal cache memory, and control and diagnostic software commonly referred to as the microcode.

IDB

The Data Protector Internal Database is an embedded database located on the Cell Manager that keeps information regarding which data is backed up, on which media it is backed up, how backup and restore sessions are run, and which devices and libraries are configured.

IDB recovery file

An IDB file (obrindex.dat) with information about IDB backups, media, and devices used for the backup. This information can significantly simplify IDB recovery. It is recommended to relocate the file, together with IDB transaction logs, to a separate physical disk from other IDB directories, and, additionally, to make an additional copy of the file.

importing media	A process that re-reads all data about backup sessions which are on the medium back into the IDB. This then allows for fast and convenient access to the data on the media. <i>See also exporting media.</i>
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 backup types.</i>
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 backup types.</i>
incremental mailbox backup	An incremental mailbox backup backs up all the changes made to the mailbox after the last backup of any type.
incremental1 mailbox backup	An incremental1 mailbox backup backs up all the changes made to the mailbox after the last full backup.
incremental (re)-establish	<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 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.
See also [full ZDB](#).
- Inet** A process that runs on each UNIX system or service that runs on each Windows system in the Data Protector cell. It is responsible for communication between systems in the cell and for starting other processes needed for backup and restore. The Inet service is started as soon as Data Protector is installed on a system. The Inet process is started by the inetd daemon.
- Information Store** (*Microsoft Exchange Server specific term*) The Microsoft Exchange Server service that is responsible for storage management. Information Store in Microsoft Exchange Server manages two kinds of stores: mailbox stores and public folder stores. A mailbox store consists of mailboxes that belong to individual users. A public folder store contains public folders and messages that are shared among several users.
See also [Key Management Service](#) and [Site Replication Service](#).
- Informix Server** (*Informix Server specific term*) Refers to Informix Dynamic Server.
- initializing** See [formatting](#).
- Installation Server** A computer system that holds a repository of the Data Protector software packages for a specific architecture. The Installation Server is used for remote installation of Data Protector clients. In mixed environments at least two Installation Servers are needed: one for UNIX systems and one for Windows systems.
- instant recovery** (*ZDB specific term*) A process in which a replica, produced by a ZDB-to-disk or a ZDB-to-disk+tape session, is used to restore the contents of the source volumes to their states at the time at which the replica was created, avoiding the need to perform a restore from tape. Depending on the application or database concerned, this may be all that is required, or other steps, such as the application of transaction log files, may be required for full recovery.
See also [replica](#), [zero downtime backup \(ZDB\)](#), [ZDB to disk](#), and [ZDB to disk+tape](#).

integration object	A backup object of a Data Protector integration, such as Oracle or SAP DB.
Internet Information Services (IIS)	<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).
IP address	An Internet Protocol address is a numeric address of a system used to uniquely identify the system on the network. The IP address consists of four groups of numbers separated by periods (full stops).
ISQL	<i>(Sybase specific term)</i> A Sybase utility used to perform system administration tasks on Sybase SQL Server.
ITO	See OM .
Java GUI Client	The Java GUI Client is a component of the Java GUI that contains only user interface related functionalities and requires connection to the Java GUI Server to function.
Java GUI Server	The Java GUI Server is a component of the Java GUI that is installed on the Data Protector Cell Manager system. The Java GUI Server receives requests from the Java GUI Client, processes them and then sends the responses back to the Java GUI Client. The communication is done through Hypertext Transfer Protocol (HTTP) on port 5556.
jukebox	See library .
jukebox device	A device consisting of multiple slots used to store either optical or file media. When being used to store file media, the jukebox device is known as the “file jukebox device”.
Key Management Service	<i>(Microsoft Exchange Server specific term)</i> The Microsoft Exchange Server service that provides encryption functionality for enhanced security. See also Information Store and Site Replication Service .
keychain	A tool that eliminates the supply of a passphrase manually when decrypting the private key. It needs to be installed and

configured on the Installation Server if you perform remote installation using secure shell.

LBO *(EMC Symmetrix specific term)* A Logical Backup Object (LBO) is an object of data storage/retrieval in the EMC Symmetrix environment. It is stored/retrieved by EMC Symmetrix as one entity and can only be restored as a whole.

library Also called autochanger, jukebox, autoloader, or exchanger. A library contains media in repository slots. Each slot holds one medium (for example, DDS/DAT). Media are moved between slots and drives by a robotic mechanism, allowing random access to media. The library can contain multiple drives.

lights-out operation or **unattended operation** A backup or restore operation that takes place outside of normal business hours without an operator. This implies that no operator personnel is present to work with the backup application or service mount requests, for example.

LISTENER.ORA *(Oracle specific term)* An Oracle configuration file that describes one or more Transparent Network Substrate (TNS) listeners on a server.

load balancing By default, Data Protector automatically balances the usage of devices selected for backup, so that they are used evenly. Load balancing optimizes the device usage by balancing the number of objects written to each device. Since load balancing is done automatically during backup time, you do not need to manage how the data is actually backed up. You just specify the devices to be used. If you do not want to use load balancing, you can select which device will be used for each object in the backup specification. Data Protector will access the devices in the specified order.

local and remote recovery Remote recovery is performed if all Media Agent hosts specified in the SRD file are accessible. If any of them fails, the disaster recovery process fails over to the local mode. This means that the target system is searched for locally attached devices. If only one device is found, it is automatically used. Otherwise, Data Protector prompts you to select the device, which will be used for restore.

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. An LCR copy is used for disaster recovery because you can switch to the LCR copy in a few seconds. If an LCR copy is used for backup and if it is located on a different disk than the original data, then the I/O load on a production database is minimal. A replicated storage group is represented as a new instance of Exchange writer called Exchange Replication Service and can be backed up (using VSS) as a normal storage group. See also Cluster Continuous Replication and Exchange Replication Service.</p>
lock name	<p>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.</p>
log_full shell script	<p>(<i>Informix Server UNIX 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>.</p>
logging level	<p>The logging level determines the amount of details on files and directories written to the IDB during backup, object copying, or object consolidation. You can always restore your data, regardless of the logging level used during backup. Data Protector provides four logging levels: Log All, Log Directories, Log Files, and No Log. The different logging level settings</p>

influence the IDB growth, backup speed, and the convenience of browsing data for restore.

logical-log files

This applies to online database backup. Logical-log files are files in which modified data is first stored before being flushed to disk. In the event of a failure, these logical-log files are used to roll forward all transactions that have been committed as well as roll back any transactions that have not been committed.

login ID

(Microsoft SQL Server specific term) The name a user uses to log on to Microsoft SQL Server. A login ID is valid if Microsoft SQL Server has an entry for that user in the system table syslogin.

login information to the Oracle Target Database

(Oracle and SAP R/3 specific term) The format of the login information is *user_name/password@service*, where:

- *user_name* is the name by which a user is known to Oracle Server and to other users. Every user name is associated with a password and both have to be entered to connect to an Oracle Target Database. This user must have Oracle SYSDBA or SYSOPER rights.
- *password* must be the same as the password specified in the Oracle password file (*orapwd*), which is used for authentication of users performing database administration.
- *service* is the name used to identify an SQL*Net server process for the target database.

login information to the Recovery Catalog Database

(Oracle specific term) The format of the login information to the Recovery (Oracle) Catalog Database is *user_name/password@service*, where the description of the user name, password, and service name is the same as in the Oracle SQL*Net V2 login information to the Oracle target database. In this case, *service* is the name of the service to the Recovery Catalog Database, not the Oracle target database. Note that the Oracle user specified here must be the owner of the Oracle Recovery Catalog.

Lotus C API

(Lotus Domino Server specific term) An interface for the exchange of backup and recovery information between Lotus Domino Server and a backup solution, like Data Protector.

LVM

A Logical Volume Manager is a subsystem for structuring and mapping physical disk space to logical volumes on UNIX

systems. An LVM system consists of several volume groups, where each volume group has several volumes.

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 <code>.edb</code> file and a streaming native internet content <code>.stm</code> file.
Main Control Unit (MCU)	<i>(HP StorageWorks Disk Array XP specific term)</i> An HP StorageWorks XP disk array that contains the primary volumes for the CA and BC configurations and acts as a master device. See also BC (HP StorageWorks Disk Array XP specific term), CA (HP StorageWorks Disk Array XP specific term), and HP StorageWorks Disk Array XP LDEV .
Manager-of-Managers (MoM)	See 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 session, a Media Agent locates data on the backup medium and sends it to the Disk Agent. The Disk Agent then writes the data to the disk. A Media Agent also manages the robotics control of a library.
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.
medium ID	A unique identifier assigned to a medium by Data Protector.
media label	A user-defined identifier used to describe a medium.
media location	A user-defined physical location of a medium, such as "building 4" or "off-site storage".
media management session	A session performing some action on a medium, such as initializing, scanning the content, verifying data on a medium, or copying a medium.
media pool	A set of media of the same type (such as DDS) used and tracked as a group. Media are formatted and assigned to a media pool.
media set	The result of a backup session is data backed up on a group of media called media set. Depending on the media usage policy, several sessions can share the same media.
media type	The physical type of media, such as DDS or DLT.
media usage policy	The media usage policy controls how new backups are added to the already used media. It can be <code>Appendable</code> , <code>Non-Appendable</code> , or <code>Appendable for incrementals only</code> .
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 overwrite.</i>

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. See also shadow copy , shadow copy provider , replica , and writer .
mirror (EMC Symmetrix and HP StorageWorks Disk Array XP specific term)	See target volume .
mirror rotation (HP StorageWorks Disk Array XP specific term)	See replica set rotation .
MMD	The Media Management Daemon process (service) runs on the Data Protector Cell Manager and controls media management and device operations. The process is started when Data Protector is installed on the Cell Manager.
MMDB	The Media Management Database (MMDB) is a part of the IDB that contains information about media, media pools, devices,

libraries, library drives, and slots configured in the cell, as well as the Data Protector media used for backup. In an enterprise backup environment, this part of the database can be common to all cells.

See also [CMMDB](#), [CDB](#).

- MoM** Several cells can be grouped together and managed from a central cell. The management system of the central cell is the Manager-of-Managers (MoM). The cells are called MoM clients. The MoM enables you to configure and manage multiple cells from a central point.
- mount request** A screen prompt that tells you to insert a specific medium into a device. Once you respond to the mount request by providing the required medium and confirm the mount request, the session continues.
- mount point** The access point in a directory structure for a disk or logical volume, for example `/opt` or `d:`. On UNIX, the mount points are displayed using the `bdf` or `df` command.
- MSM** The Data Protector Media Session Manager, which runs on the Cell Manager and controls media sessions, such as copying media.
- MU number** (*HP StorageWorks Disk Array XP specific term*) Mirror Unit number. An integer number (0, 1 or 2), used to indicate a first-level mirror.
See also [first-level mirror](#).
- multi-drive server** A license that allows you to run an unlimited number of Media Agents on a single system. This license, which is bound to the IP address of the Cell Manager, is no longer available.
- obdrindex.dat** See [IDB recovery file](#).
- OBDR capable device** A device that can emulate a CD-ROM drive loaded with a bootable disk and can thus be used as a backup or boot device for disaster recovery purposes.
- object** See [backup object](#).
- object 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.
offline backup	A backup during which an application database cannot be used by the application. <ul style="list-style-type: none">• For simple backup methods (non ZDB), the database is generally put into a quiescent state that allows use by the backup system, but not the application, for the whole backup period (several minutes or hours). For instance, for backup to tape, until streaming of data to the tape is finished.• For ZDB methods, the database is also put into the quiescent state, but for the period of the data replication process only (several seconds). Normal database operation can then be resumed for the rest of the backup process.

See also [zero downtime backup \(ZDB\)](#) and [online backup](#).

offline recovery	Offline recovery is performed if the Cell Manager is not accessible, for example, due to network problems. Only standalone and SCSI library devices can be used for offline recovery. Recovery of the Cell Manager is always offline.
offline redo log	See archived redo log .
On-Bar	<p><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:</p> <ul style="list-style-type: none"> • the <code>onbar</code> 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	<p><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 <code>onconfig</code> file in the directory <code>INFORMIXDIR\etc</code> (on Windows) or <code>INFORMIXDIR/etc/</code> (on UNIX).</p>
online backup	<p>A backup performed while a database application remains available for use. The database is placed into a special backup mode of operation for the time period that the backup application requires access to the original data objects. During this period, the database is fully operational, but there may be a small performance impact and log files may grow very quickly.</p> <ul style="list-style-type: none"> • For simple backup methods (non ZDB), backup mode is required for the whole backup period (several minutes or hours). For instance, for backup to tape, until streaming of data to tape is finished. • For ZDB methods, backup mode is required for the short period of the data replication process only (several seconds). Normal database operation can then be resumed for the rest of the backup process. <p>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.</p>

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 .
OpC	See OM .
OpenSSH	A set of network connectivity tools used to access remote machines securely, by using a variety of authentication and encryption methods. It needs to be installed and configured on the Installation Server and the client if you perform remote installation using secure shell.
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 <code>ORACLE_SID</code> . The <code>ORACLE_SID</code> is included in the <code>CONNECT DATA</code> parts of the connect descriptor in a <code>TNSNAMES.ORA</code> file and in the definition of the TNS listener in the <code>LISTENER.ORA</code> file.
original system	The system configuration backed up by Data Protector before a computer disaster hits the system.
overwrite	An option that defines one mode to resolve file conflicts during restore. All files are restored from a backup even if they are older than existing files. See also merging .

OM	<p>HP Operations Manager software for UNIX provides powerful capabilities for operations management of a large number of systems and applications on in a network. Data Protector provides an integration into this management product. This integration is implemented as a SMART Plug-In for OM management servers on HP-UX, Solaris, and Linux. Earlier versions of OM were called IT/Operation, Operations Center and Vantage Point Operations.</p> <p>See also merging.</p>
ownership	<p>The ownership of a backup determines who can restore from the backup. The user who starts an interactive backup is the session owner. If a user starts an existing backup specification without modifying it, the session is not considered interactive. In that case, if the backup owner has been defined in the backup specification, they remain the session owner. Otherwise, the session owner becomes the user who started the backup in question. For the scheduled backups, by default, the session owner is for the UNIX Cell Manager: <code>root.sys@Cell Manager</code>, and for the Windows Cell Manager, the user that was specified during the installation of the Cell Manager. It is possible to modify the ownership, so that the specific user becomes the session owner.</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 <code>Data_Protector_home\Config\Server\dr\pls</code> directory on a Windows Cell Manager or in <code>/etc/opt/omni/server/dr/pls</code> directory on a UNIX Cell Manager with the filename <code>recovery.pls</code>.</p>
package	<p><i>(MC/ServiceGuard and Veritas Cluster specific term)</i> 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><i>(HP StorageWorks Disk Array XP specific term)</i> A mirrored pair of disks can have various status values depending on the action performed on it. The three most important status values are:</p>

- COPY - The mirrored pair is currently re-synchronizing. Data is transferred from one disk to the other. The disks do not contain the same data.
- PAIR - The mirrored pair is completely synchronized and both disks (the primary volume and the mirrored volume) contain identical data.
- SUSPENDED - The link between the mirrored disks is suspended. That means that both disks are accessed and updated independently. However, the mirror relationship is still maintained and the pair can be re-synchronized without transferring the complete disk.

parallel restore	Restoring backed up data to multiple disks at the same time (that is, in parallel) by running multiple Disk Agents, that receive data from one Media Agent. For the parallel restore to work, select data that is located on different disks or logical volumes and during backup, the data from the different objects must have been sent to the same device using a concurrency of 2 or more. During a parallel restore, the data for multiple objects selected for restore is read from media at the same time, thereby improving performance.
parallelism	The concept of reading multiple data streams from an online database.
physical device	A physical unit that contains either a drive or a more complex unit such as a library.
post-exec	A backup option that executes a command or script after the backup of an object or after the entire session completes. Post-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows and as shell scripts on UNIX. See also pre-exec .
pre- and post-exec commands	Pre- and post-exec commands are used to perform additional action before and after a backup or restore session. They are not supplied by Data Protector. You need to create your own commands. They can be written as executables or batch files on Windows and as shell scripts on UNIX.
prealloc list	A subset of media in a media pool that specifies the order in which media are used for backup.

pre-exec	A backup option that executes a command or script before the backup of an object or before the entire session is started. Pre-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows and as shell scripts on UNIX. See also post-exec .
primary volume (P-VOL)	<i>(HP StorageWorks Disk Array XP specific term)</i> Standard HP StorageWorks Disk Array XP LDEVs that act as a primary volume for the CA and BC configurations. The P-VOL is located in the MCU. See also secondary volume (S-VOL) and Main Control Unit (MCU) .
protection	See 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 <code>.edb</code> file and a streaming native internet content <code>.stm</code> 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
RAID	Redundant Array of Inexpensive Disks.
RAID Manager Library	<i>(HP StorageWorks Disk Array XP specific term)</i> The RAID Manager Library is used internally by Data Protector on Solaris systems to allow access to HP StorageWorks Disk Array XP configuration, status, and performance data and to key HP StorageWorks Disk Array XP features through the use of function calls translated into a sequence of low level SCSI commands.
RAID Manager XP	<i>(HP StorageWorks Disk Array XP specific term)</i> The RAID Manager XP application provides an extensive list of commands to report and control the status of the CA and BC applications. The commands communicate through a RAID Manager instance with the HP StorageWorks Disk Array XP Disk Control Unit. This

instance translates the commands into a sequence of low level SCSI commands.

rawdisk backup	See disk image backup .
RCU	See 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.
RDS	The Raima Database Server process (service) runs on the Data Protector Cell Manager and manages the IDB. The process is started when Data Protector is installed on the Cell Manager.
Recovery Catalog	<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: <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	<i>(Oracle specific term)</i> Recovery files are Oracle 10g specific files that reside in the flash recovery area: the current control file, online redo logs, archived redo logs, flashback logs, control file autobackups, datafile copies, and backup pieces. See also flash recovery area .
RecoveryInfo	When backing up Windows configuration files, Data Protector collects the information about the current system configuration (information on disk layout, volume, and network configuration). This information is needed for disaster recovery.

Recovery Manager (RMAN)	<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.
recycle	A process that removes the data protection from all backed up data on a medium, allowing Data Protector to overwrite it during one of the next backups. Data that belongs to the same session(s) but resides on other media is also unprotected. Recycling does not actually alter the data on the medium.
redo log	<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 StorageWorks Disk Array XP specific term)</i> The Remote Control Unit (RCU) acts as a slave of an MCU in a CA configuration. In bidirectional configurations, the RCU can act as an MCU.
Removable Storage Management Database	<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 UNIX, the whole volume or disk group containing a backup object (logical volume) is replicated. If partitions are used on Windows, the whole physical volume containing the selected partition is replicated.

See also [snapshot](#), [snapshot creation](#), [split mirror](#), and [split mirror creation](#).

- replica set** *(ZDB specific term)* A group of replicas, all created using the same backup specification.
See also [replica](#) and [replica set rotation](#).
- replica set rotation** *(ZDB specific term)* The use of a replica set for regular backup production: Each time the same backup specification requiring the use of a replica set is run, a new replica is created and added to the set, until the maximum number for the set is reached. After that, the oldest replica in the set is replaced, maintaining the maximum number of replicas in the set.
See also [replica](#) and [replica set](#).
- restore chain** All backups that are necessary for a restore of a backup object to a certain point in time. A restore chain consists of a full backup of the object and any number of related incremental backups.
- restore session** A process that copies data from backup media to a client.
- resync mode** *(HP StorageWorks Disk Array XP VSS provider specific term)* One of two XP VSS hardware provider operation modes. When the XP 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 \(S-VOL\)](#), [MU number](#), and [replica set rotation](#).
- RMAN (Oracle specific term)** See [Recovery Manager](#).
- RSM** The Data Protector Restore Session Manager controls the restore session. This process always runs on the Cell Manager system.

RSM	<i>(Windows specific term)</i> 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.
scan	A function that identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library).
scanning	A function which identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library). It is useful to perform a scan and check the actual media in the device if someone has manually manipulated media without using Data Protector to eject or enter, for example.
Scheduler	A function that controls when and how often automatic backups occur. By setting up a schedule, you automate the start of backups.
secondary volume (S-VOL)	<i>(HP StorageWorks Disk Array XP specific term)</i> secondary volumes, or S-VOLs, are XP LDEVs that act as a secondary CA or BC mirror of another LDEV (a P-VOL). In the case of CA, S-VOLs can be used as failover devices in a MetroCluster configuration. The S-VOLs are assigned separate SCSI addresses, different from the addresses used by the P-VOLs. See also primary volume (P-VOL) 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, 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 <code>omnimnt</code> , <code>omnistat</code> , and <code>omniabort</code> commands.

shadow copy	<p><i>(Microsoft VSS specific term)</i> A volume that represents a duplicate of the original volume at a certain point in time. The data is then backed up from the shadow copy and not from the original volume. The original volume continues to change as the backup process continues, but the shadow copy of the volume remains constant.</p> <p>See also Microsoft Volume Shadow Copy Service and replica.</p>
shadow copy provider	<p><i>(Microsoft VSS specific term)</i> 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).</p> <p>See also shadow copy.</p>
shadow copy set	<p><i>(Microsoft VSS specific term)</i> A collection of shadow copies created at the same point in time.</p> <p>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>
SIBF	<p>The Serverless Integrations Binary Files (SIBF) is a part of the IDB that stores raw NDMP meta data. This data is necessary to perform restore of NDMP objects.</p>
single instancing	<p><i>(IAP specific term)</i> The process of recognizing redundancy of data, at both a whole object and a chunk level. It computes a strong hash for each data chunk and uses it as a unique content address needed to determine whether attempts to store duplicates are being made.</p> <p>See also backup to IAP.</p>
Site Replication Service	<p><i>(Microsoft Exchange Server specific term)</i> The Microsoft Exchange Server 2000/2003 service that permits compatibility with Microsoft Exchange Server 5.5 by emulating the Exchange Server 5.5 directory service.</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</p>

number. To read a medium, a robotic mechanism moves the medium from a slot into the drive.

- SMB** See [split mirror backup](#).
- smart copy** (*VLS specific term*) A copy of the backed up data created from the virtual tape to the physical tape library. The smart copy process allows Data Protector to distinguish between the source and the target medium thus enabling media management. See also [Virtual Library System \(VLS\)](#).
- smart copy pool** (*VLS specific term*) A pool that defines which destination library slots are available as smart copy targets for a specified source virtual library. See also [Virtual Library System \(VLS\)](#) and [smart copy](#).
- SMBF** The Session Messages Binary Files (SMBF) part of the IDB stores session messages generated during backup, restore, object copy, object consolidation, and media management sessions. One binary file is created per session. The files are grouped by year and month.
- snapshot** (*HP StorageWorks VA and HP StorageWorks EVA specific term*) A form of replica produced using snapshot creation techniques. A range of snapshot types is available, with different characteristics, depending on the arrays/techniques used. Such replicas are dynamic and may be either virtual copies, still reliant upon the contents of the source volumes, or independent exact duplicates (clones), depending on the snapshot type and the time since creation. See also [replica](#) and [snapshot creation](#).
- snapshot backup**
(*HP StorageWorks VA and HP StorageWorks EVA specific term*) See [ZDB to tape](#), [ZDB to disk](#), and [ZDB to disk+tape](#).
- snapshot creation** (*HP StorageWorks VA and HP StorageWorks EVA specific term*) A replica creation technique, in which copies of source volumes are created using storage virtualization techniques. The replicas are considered to be created at one particular point in time, without pre-configuration, and are immediately available for

use. However background copying processes normally continue after creation.

See also [snapshot](#).

- source (R1) device** *(EMC Symmetrix specific term)* An EMC Symmetrix device that participates in SRDF operations with a target (R2) device. All writes to this device are mirrored to a target (R2) device in a remote EMC Symmetrix unit. An R1 device must be assigned to an RDF1 group type.
See also [target \(R2\) device](#).
- source volume** *(ZDB specific term)* A storage volume containing data to be replicated.
- sparse file** A file that contains data with portions of empty blocks. Examples are: a matrix in which some or much of the data contains zeros, files from image applications, and high-speed databases. If sparse file processing is not enabled during restore, it might be impossible to restore this file.
- split mirror** *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)* A replica created using split mirror techniques. Such a replica provides an independent, exact duplicate, or clone of the contents of the source volumes.
See also [replica](#) and [split mirror creation](#).
- split mirror backup** *(EMC Symmetrix specific term)* See [ZDB to tape](#).
- split mirror backup** *(HP StorageWorks Disk Array XP specific term)* See [ZDB to tape](#), [ZDB to disk](#), and [ZDB to disk+tape](#).
- split mirror creation** *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)* A replica creation technique, in which a pre-configured set of target volumes (a mirror) is kept synchronized with a set of source volumes until the time at which a replica of the contents of the source volumes is required. Then, the synchronization is stopped (the mirror is split) and a split mirror replica of the source volumes at the time of the split remains in the target volumes.
See also [split mirror](#).

split mirror restore	<i>(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)</i> A process in which data backed up in a ZDB-to-tape or a ZDB-to-disk+tape session is restored from tape media to a split mirror replica, which is then synchronized to the source volumes. Individual backup objects or complete sessions can be restored using this method. See also ZDB to tape , ZDB to disk+tape , and replica .
sqlhosts file	<i>(Informix Server specific term)</i> An Informix Server connectivity information file (on UNIX) or registry (on Windows) that contains the names of each of the database servers and any aliases to which the clients on a host computer can connect.
SRD file	The Data Protector System Recovery Data (SRD) file contains system information required for installing and configuring the operating system in case of a disaster. The SRD file is an ASCII file, generated when a CONFIGURATION backup is performed on a Windows client and stored on the Cell Manager.
SRDF	<i>(EMC Symmetrix specific term)</i> 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	<i>(HP StorageWorks Disk Array XP specific term)</i> A Data Protector software module that executes all tasks required for a split mirror backup integration. It communicates with the HP StorageWorks Disk Array XP storing system using the RAID Manager XP utility (HP-UX and Windows systems) or RAID Manager Library (Solaris systems).
sst.conf file	The file <code>/usr/kernel/drv/sst.conf</code> is required on each Data Protector Sun Solaris client to which a multi-drive library device is connected. It must contain an entry for the SCSI address of the robotic mechanism of each library device connected to the client.
st.conf file	The file <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	<i>(Microsoft Exchange Server specific term)</i> 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.
StorageTek ACS library	<i>(StorageTek specific term)</i> Automated Cartridge System is a library system (also known as Silo) consisting of one Library Management Unit (LMU) and one to 24 Library Storage Modules (LSM) connected to the unit.
storage volume	<i>(ZDB specific term)</i> A storage volume represents an object that may be presented to an operating system or some other entity (for example, a virtualization mechanism) upon which volume management systems, file systems, or other objects may exist. The volume management systems, file systems are built on this storage. Typically, these can be created or exist within a storage system such as a disk array.
switchover	See failover .
Sybase Backup Server API	<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.
Symmetrix Agent (SYMA)	<i>(EMC Symmetrix specific term)</i> The Data Protector software module that prepares the EMC Symmetrix environment for backup and restore operations.

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	<i>(Sybase specific term)</i> The four system databases on a newly installed Sybase SQL Server are the: <ul style="list-style-type: none"> • master database (master) • temporary database (tempdb) • system procedure database (sybprocedure) • model database (model).
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.

tablespace	A part of a database structure. Each database is logically divided into one or more tablespaces. Each tablespace has data files or raw volumes exclusively associated with it.
tapeless backup (ZDB specific term)	See ZDB to disk .
target database	<i>(Oracle specific term)</i> In RMAN, the target database is the database that you are backing up or restoring.
target (R2) device	<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. See also source (R1) device .
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 crashed 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 logs	<i>(Data Protector specific term)</i> Keep track of IDB changes. The archiving of transaction logs should be enabled to prevent you from losing the transaction log files that are created after the last IDB backup and are necessary for IDB recovery.
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) .
TSANDS.CFG file	<i>(Novell NetWare specific term)</i> A file that allows you to specify the names of containers where you want backups to begin. It is a text file located in the <code>SYS:SYSTEM\TSA</code> directory on the server where <code>TSANDS.NLM</code> is loaded.

UIProxy	The Java GUI Server (<code>UIProxy</code> service) runs on the Data Protector Cell Manager. It is responsible for communication between the Java GUI Client and the Cell Manager, moreover, it performs business logic operations and sends only important information to the client. The service is started as soon as Data Protector is installed on the Cell Manager.
unattended operation	See lights-out operation .
user account (Data Protector user account)	You can use Data Protector only if you have a Data Protector user account, which restricts unauthorized access to Data Protector and to backed up data. Data Protector administrators create this account specifying a user logon name, the systems from which the user can log on, and a Data Protector user group membership. This is checked whenever the user starts the Data Protector user interface or performs specific tasks.
user 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.
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 StorageWorks EVA specific term)</i> The firmware that manages all aspects of storage system operation, including communication with Command View EVA through the HSV controllers. See also Command View (CV) EVA .
Virtual Device Interface	<i>(Microsoft SQL Server specific term)</i> This is a SQL Server programming interface that allows fast backup and restore of large databases.
virtual disk	<i>(HP StorageWorks EVA specific term)</i> A unit of storage allocated from an HP StorageWorks Enterprise Virtual Array storage pool. Virtual disks are the entities that are replicated using the HP StorageWorks Enterprise Virtual Array snapshot functionality. See also source volume and target volume .
virtual 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. See also Virtual Library System (VLS) and Virtual Tape Library .

Virtual Tape Library (VTL)	<i>(VLS specific term)</i> An emulated tape library that provides the functionality of traditional tape-based storage. See also Virtual Library System (VLS) .
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 mount point	<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 .
VPO	See OM .
VSS	See Microsoft Volume Shadow Copy Service .
VSS compliant mode	<i>(HP StorageWorks Disk Array XP VSS provider specific term)</i> One of two XP VSS hardware provider operation modes. When the XP 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 of 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.

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.
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).
XCopy engine	<i>(direct backup specific term)</i> A SCSI-3 copy command that allows you to copy data from a storage device having a SCSI source address to a backup device having a SCSI destination address, thus enabling direct backup. The data flows from a source device (either block or streaming, that is, disk or tape) to the destination device (either block or streaming) through XCopy. This releases

the controlling server of reading the data from the storage device into memory and then writing the information to the destination device.

See also [direct backup](#).

ZDB

See [zero downtime backup \(ZDB\)](#).

ZDB database

(ZDB specific term) A part of the IDB, storing ZDB related information such as source volumes, replicas and security information. The ZDB database is used for ZDB, instant recovery, and split mirror restore.

See also [zero downtime backup \(ZDB\)](#).

ZDB to disk

(ZDB specific term) A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time. Multiple replicas, produced using the same backup specification at different times, can be kept in a replica set. A replica from ZDB to disk can be restored using the instant recovery process.

See also [zero downtime backup \(ZDB\)](#), [ZDB to tape](#), [ZDB to disk+tape](#), [instant recovery](#), and [replica set rotation](#).

ZDB to disk+tape

(ZDB specific term) A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time, in the same way as with ZDB to disk. However, data in the replica is also streamed to a backup medium, as with ZDB to tape. If this backup method is used, data backed up in the same session can be restored using the instant recovery process, the standard Data Protector restore from tape, or on split mirror arrays, split mirror restore.

See also [zero downtime backup \(ZDB\)](#), [ZDB to disk](#), [ZDB to tape](#), [instant recovery](#), [replica](#), and [replica set rotation](#).

ZDB to tape

(ZDB specific term) A form of zero downtime backup where data in the replica produced is streamed to a backup medium, typically tape. Instant recovery is not possible from such a backup, so the replica need not be retained on the disk array after backup completion. The backed-up data can be restored using standard Data Protector restore from tape. On split mirror arrays, split mirror restore can also be used.

See also [zero downtime backup \(ZDB\)](#), [ZDB to disk](#), [instant recovery](#), [ZDB to disk+tape](#), and [replica](#).

zero downtime backup (ZDB)

A backup approach in which data replication techniques, provided by a disk array, are used to minimize the impact of backup operations on an application system. A replica of the data to be backed up is created first. All subsequent backup operations are performed on the replicated data rather than the original data, while the application system can be returned to normal operation.

See also [ZDB to disk](#), [ZDB to tape](#), [ZDB to disk+tape](#), and [instant recovery](#).

Index

A

- Application Integration Agent, 48
- application systems, 48
- audience, 15

B

- backup types, 29
 - incremental ZDB, 76
 - ZDB to disk, 70
 - ZDB to disk+tape, 30, 70
 - ZDB to tape, 69
- backup scenarios, 91
- backup specifications, 67
- backup systems, 48
- backup types, 30
 - ZDB to disk, 30
 - ZDB to tape, 29
- BC configuration
 - EVA, 115
- BC configuration
 - EVA, 61
 - VA, 58, 115
 - XP, 52, 95
- BC1 configuration
 - XP, 96
- Business Copy configuration
 - See BC

C

- CA configuration
 - XP, 53, 100

- CA+BC configuration
 - EVA, 62, 123
 - XP, 54, 102

- cascading configuration
 - XP, 96

- clusters

- instant recovery, 84
- CA+BC XP, 104
- LVM mirroring EMC, 110
- LVM mirroring VA, 123
- LVM mirroring XP, 100
- SRDF+TimeFinder EMC, 114

- concurrency handling
 - device locking, 90
 - disk locking, 90

- configurations

- BC, VA, 115
- BC, EVA, 115
- BC, XP, 95
- BC1, XP, 96
- CA+BC, XP, 102
- CA+BC, EVA, 123
- CA+BC, XP, 102
- CA, XP, 100
- cascading, XP, 96
- LVM mirroring, EMC, 107
- LVM mirroring, XP, 97
- LVM mirroring, EVA, 117
- LVM mirroring, VA, 120
- SRDF+TimeFinder, EMC, 112
- SRDF, EMC, 110
- TimeFinder, EMC, 105

- Continuous Access configuration
 - See CA

conventions
document, 22
creating replicas, 29, 67, 75

D

Data Facility configuration
See RDF

Data Protector cell, 47 - 51
ZDB database, 49
application systems, 48
backup systems, 48
components, 47

database applications, 63 - 64
MS Exchange Server, 63
MS SQL Server, 63
offline backup, 28
offline backup, 74
online backup, 28, 74
Oracle, 63
restore, 64
SAP R/3, 63
supported database applications, 63
transaction log backup, 64

deleting replicas, 72

device locking, 90

direct backup, XP, 76

disk arrays, supported configurations
XP, 52

disk locking, 90

Disk Array XP
See XP

Disk Array Agent, 48

disk arrays, introduction, 33 - 34
disk virtualization, 33
RAID technology, 33
storage volumes, 33

disk arrays, supported configurations
EMC, 55
VA, 115

disk arrays, supported ZDB techniques,
30, 93

disk arrays, supported configurations,
52 - 63
EMC, 105
EVA, 60, 115
VA, 58
XP, 94

disk virtualization, 27, 33

document
conventions, 22
related documentation, 15

documentation
HP website, 15
providing feedback, 25

E

EMC Symmetrix
See EMC

EMC, configurations
LVM mirroring, 107
SRDF, 56
TimeFinder, 56

EMC, backup
local replication, 56
local replication integrating with LVM
mirroring, 56
local replication, 105 - 107
local replication using LVM mirroring,
107 - 110
remote plus local replication, 57
remote replication, 56
remote plus local replication, 112 - 114
remote replication, 110 - 112

EMC, configurations
SRDF, 57, 110
SRDF+TimeFinder, 112
TimeFinder, 57, 105

EMC, restore
split mirror restore, 84

Enterprise Virtual Array
See EVA

EVA, configurations
BC, 61

- EVA, backup
 - local replication, 61
 - local replication integrating with LVM mirroring, 61
 - local replication, 115
 - planning ZDB strategy, 88
 - remote plus local replication, 62, 123
 - remote plus local replication using LVM mirroring, 117

- EVA, configurations
 - BC, 115
 - CA+BC, 62, 123
 - LVM mirroring, 117

- EVA, introduction, 60

- EVA, restore
 - instant recovery, 71, 81

F

- full ZDB, 91

H

- help

 - obtaining, 24

- hot-backup mode, 27

- hot-backup mode, 28, 74

- HP

 - technical support, 24

- HP StorageWorks Disk Array XP

 - See XP

- HP StorageWorks Enterprise Virtual Array

 - See EVA

- HP StorageWorks Virtual Array

 - See VA

I

- incremental ZDB, 91

- incremental ZDB, 76

- instant recovery

 - introduction, 31

- instant recovery, 71, 81 - 84

 - advantages, 27

 - clusters, 84

 - LVM mirroring, 83

 - overview, 79

 - process, 82

- IR

 - See instant recovery

L

- local replication, 35 - 42

 - integrating with LVM mirroring, 42

 - advantages, 35

 - disadvantages, 35

 - snapshot replication, 37

 - split mirror replication, 35

- local replication using LVM mirroring, 59

- locking

 - devices, 90

 - disks, 90

- Logical Volume Manager mirroring

 - See LVM mirroring

- LVM mirroring

 - instant recovery, 83

- LVM mirroring

 - EMC, 56, 107

 - EVA, 61, 117

 - local replication, 42

 - VA, 59, 120

 - XP, 53, 97

M

- mirrors, 35

- MS Exchange Server integration, 63

- MS SQL Server integration, 63

O

- offline backup, 28, 74

- online backup
 - hot-backup mode, 74
- online backup, 28, 74
 - hot-backup mode, 28
- Oracle integration, 63

P

- planning ZDB strategy, 87 - 92
 - backup scenarios, 91
 - concurrency handling, 90
 - flexibility in recovery, 87
 - introduction, 87
 - snapshot disk arrays, 88
 - split mirror disk arrays, 87
- pre-allocated snapshots
 - See standard snapshots

R

- RAID technology, 33
- related documentation, 15
- Remote Data Facility configuration
 - See SRDF
- remote plus local replication, 44 - 46
 - advantages, 45
 - disadvantages, 45
 - snapshot replication, 45
 - split mirror replication, 45
- remote replication, 43 - 44
 - advantages, 43
 - disadvantages, 43
 - split mirror replication, 44
- replica sets, 68
 - rotation, 68
- replicas
 - creating, 29, 67, 75
 - deleting, 72
 - introduction, 29
 - lifecycle, 65
 - streaming to tape, 75
 - using, 69, 76

- replication
 - local, 35 - 42
 - remote, 43 - 44
 - remote plus local, 44 - 46
 - scheduling, 69
 - techniques, 34
- restore from ZDB, 32, 79 - 85
 - instant recovery, 31
 - instant recovery, 81 - 84
 - split mirror restore, 31, 84 - 85
 - standard Data Protector restore, 31, 80
- roll forward, 64, 82

S

- SAP R/3 integration, 63
- scheduling replication, 69
- single-host configuration, 94, 96
- snapclones, 37, 41, 89
- snapshot replication
 - local, 37 - 42
 - planning, 88
 - remote plus local, 45 - 46
- snapshot types
 - snapclones, 37, 41
 - standard snapshots, 37, 38
 - vsnap, 37, 39
- source volumes, 29
- split mirror replication
 - planning, 87
- split mirror replication
 - local, 35 - 37
 - mirrors, 35
 - remote, 44 - 44
 - remote plus local, 45 - 45
- split mirror restore, 84 - 85
 - overview, 81
 - process, 84
- SRDF configuration
 - EMC, 56, 57, 110
- SRDF+TimeFinder configuration
 - EMC, 112

- standard Data Protector restore overview, 80
- standard snapshots, snapshot replication, 37
- standard snapshots, snapshot replication, 38
- storage volumes, 33
- Subscriber's Choice, HP, 25
- supported database applications, 63
- supported disk arrays, 30, 93
 - configurations, 93 - 124

T

- target volumes, 29
- technical support
 - HP, 24
- technical support
 - service locator website, 25
- TimeFinder configuration
 - EMC, 56, 57, 105
- transaction logs, 28, 31
- transaction logs, 28, 64

U

- user interfaces, 49
 - Data Protector CLI, 51
 - Data Protector GUI, 50

V

- VA, configurations
 - BC, 58
- VA, backup, 58 - 60
 - local replication, 58
 - local replication integrating with LVM mirroring, 59
 - local replication, 115
 - planning ZDB strategy, 88
 - remote plus local replication using LVM mirroring, 120

- VA, configurations
 - BC, 115
 - LVM mirroring, 120
- VA, introduction, 58
- VA, restore
 - instant recovery, 71, 81
- Virtual Array
 - See VA
- virtualization, 27, 33
- virtually capacity-free snapshots
 - See vsnaps;
 - vsnaps, 37, 39, 88

W

- websites
 - HP Subscriber's Choice for Business, 25
 - HP, 25
 - product manuals, 15

X

- XP, configurations
 - BC, 52
 - CA, 53
 - LVM mirroring, 97
- XP, backup, 52 - 55
 - local replication, 52
 - local replication integrating with LVM mirroring, 53
 - local replication, 94 - 97
 - local replication using LVM mirroring, 97 - 100
 - remote plus local replication, 54
 - remote replication, 53
 - remote plus local replication, 102 - 104
 - remote replication, 100 - 102

- XP, configurations
 - BC, 95
 - BC1, 96
 - CA, 100
 - CA+BC, 54, 102
 - cascading, 96
- XP, restore
 - instant recovery, 71
 - instant recovery, 81
 - split mirror restore, 84

Z

- ZDB Agent, 48
- ZDB database, 49, 77
- ZDB to disk, 70
- ZDB to disk+tape, 70
- ZDB to tape, 69
- ZDB, backup process
 - streaming replicas to tape, 75
- ZDB, backup types, 29
- ZDB, planning backup strategy, 87 - 92
- ZDB, backup process, 73 - 77
 - creating replicas, 75
 - freezing database applications, 74
 - locating data objects, 73
 - overview, 73
 - recording session information, 77
- ZDB, backup types, 30
 - incremental ZDB, 76
 - ZDB to disk, 70
 - ZDB to disk+tape, 70
 - ZDB to tape, 69
 - ZDB to disk+tape, 30
 - ZDB to disk, 30
 - ZDB to tape, 29

- ZDB, introduction, 27 - 32
 - advantages, 27
 - backup types, 29
 - concepts, 27
 - database application backup, 28
 - replicas, 28
 - replication, 28
 - snapshot backup, 29
 - source volumes, 29
 - split mirror backup, 29
 - target volumes, 29
- ZDB, planning backup strategy
 - backup scenarios, 91
 - concurrency handling, 90
 - flexibility in recovery, 87
 - introduction, 87
 - snapshot disk arrays, 88
 - split mirror disk arrays, 87
- zero downtime backup
 - ZDB;, 27