HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide

Manual Edition: October 2004

Manufacturing Part Number: B6960-90112

Release A.05.50

Legal Notices

Contents

Contents

Contents

Printing History

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

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

Table 1 **Edition History**

Part Number	Manual Edition	Product
B6960-90112	October 2004	Data Protector Release A.05.50

Conventions

The following typographical conventions are used in this manual.

Table 2

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

Data Protector provides a cross-platform (Windows and UNIX) graphical user interface.

Figure 1 Data Protector Graphical User Interface

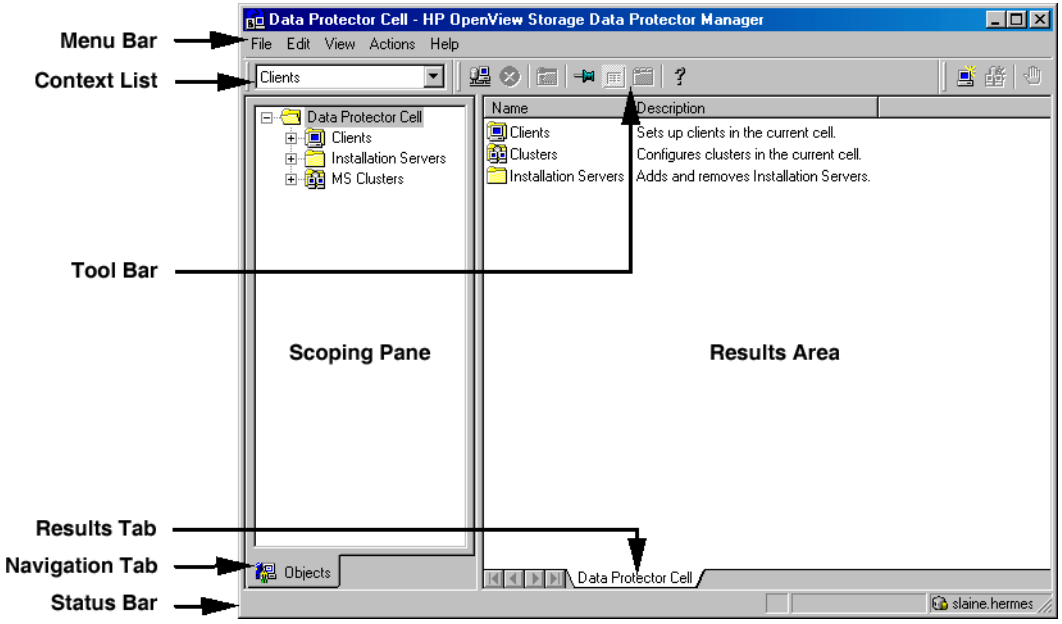

Contact Information

General Information

General information about Data Protector can be found at

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

Technical Support

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

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

http://www.hp.com/support

Information about the latest Data Protector patches can be found at

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

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

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

Documentation Feedback

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

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

Training Information

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

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

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

Data Protector Documentation

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

Manuals

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

HP OpenView Storage Data Protector Concepts Guide

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

HP OpenView Storage Data Protector Administrator's Guide

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

HP OpenView Storage Data Protector Installation and Licensing Guide

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

HP OpenView Storage Data Protector Integration Guide

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

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

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

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

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

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

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

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

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

HP OpenView Storage Data Protector Integration Guide for HP OpenView

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

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

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

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

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

HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide

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

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

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

HP OpenView Storage Data Protector Zero Downtime Backup Integration Guide

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

HP OpenView Storage Data Protector MPE/iX System User Guide

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

HP OpenView Storage Data Protector Media Operations User's Guide

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

HP OpenView Storage Data Protector Software Release Notes

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

Online Help

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

In This Book

The *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide* describes zero downtime backup and instant recovery concepts and how these are used within Data Protector.

Audience

This manual 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 OpenView Storage Data Protector Concepts Guide* and the task-oriented *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide* and *HP OpenView Storage Data Protector Zero Downtime Backup Integration Guide*.

Organization

The manual is organized as follows:

1 Overview

In This Chapter

This chapter provides an introduction to zero downtime backup and instant recovery. After reading it you will know:

- What zero downtime backup and instant recovery are.
- What a replica is.
- The available replica types and their relative merits.
- How replicas are used in zero downtime backup and instant recovery.

Introduction

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 (Redundant Array of Independent Disks) technology. Such disk arrays are capable of holding large application databases, containing vast amounts of data.

As the arrays have increased in sophistication, so have the tools for administering the information stored on them. Of these, storage virtualization techniques are amongst the most important. The whole of an array, with no software running on it is, essentially, just a very large block of storage. However, by using storage virtualization techniques, this can be divided into many virtual disks or file systems in the same way as a single disk can be.

Virtual disks, file systems, etc., can easily be copied within an array, perhaps many times if the array is big enough. This offers array users the opportunity to perform operations on copies of their data, as if they were the originals, without any risk to their original data.

In the realms of backup technology, the ability to create copies of application data has been exploited to produce backup solutions for applications in high availability and mission critical areas.

When using conventional tape backup and restore techniques, handling the enormous amounts of data involved can be very time consuming, which can be unacceptable, particularly in the modern internet environment where information is expected to be available 24 hours a day.

The need for new methods to overcome these problems has lead to the development of the zero downtime backup and instant recovery techniques described in this manual.

The basic principles behind the two techniques are very simple:

- Create, at high speed on the array, a copy of the data to be backed up and then perform backup operations on the copy, rather than on the original data.
- Restore a backup copy of data, held on the array, to its original location on the array to facilitate high speed recovery.

However, the rapid increase in the number of disk arrays available, with their various different copying methods and the, all too often (almost obligatory) new terminology introduced with each, have served to make the field appear very complicated.

The purpose of this manual is to attempt to reduce the apparent complexity of the subject and present the general principles behind the techniques in an understandable way.

Where possible, terms that are compliant with standard Storage Network Industry Association (SNIA) terminology have been used.

An Introduction to Zero Downtime Backup and Instant Recovery

The use of conventional backup to tape techniques for large database applications, can be very problematic: to produce a backup of a consistent database, the database concerned either has to be taken offline or, if the application allows it, has to be put into a “hot-backup mode” while data in it is streamed to tape.

The first option can cause major disruption to the application’s operation. The second can produce a lot of large transaction log files, and can put extra load on the application system.

However, if an application database is installed on one of a disk array that can be integrated with Data Protector, it is possible to reduce the time for which the database is taken offline or placed in hot-backup mode. This can be done by first creating a point-in-time copy of the database on the array and then streaming data from the copy to tape. The copy concerned can be created very quickly and, afterwards, the database can be returned to normal operation before any streaming of data to tape is started.

As an extension of this approach, it is not actually necessary to stream data from the copy to tape to produce a backup. The copy itself can be kept on the array as the backup and a restore can be performed directly from it.

These are the basic principles of zero downtime backup and instant recovery.

There are various types of copy that can produced: depending on the type of array they may be either exact duplicates, or virtual copies. Wherever possible in this guide, we will use the generic term **replica** to simplify the explanations, unless a specific type of copy is being described. The creation of the replicas concerned, we will refer to as **replication**.

What Are Zero Downtime Backup and Instant Recovery?

These are two of the most important backup and restore techniques available for applications using disk arrays. They are especially important for high availability applications.

Zero Downtime Backup

Zero downtime backup is the term used, within Data Protector, to describe a backup approach in which replication techniques are used to minimize the impact of backup operations on an application database or file system: a replica of the data to be backed up is created first and all subsequent backup operations are performed on the replicated data rather than the original data.

Figure 1-1 **Zero Downtime Backup and Instant Recovery Concept**

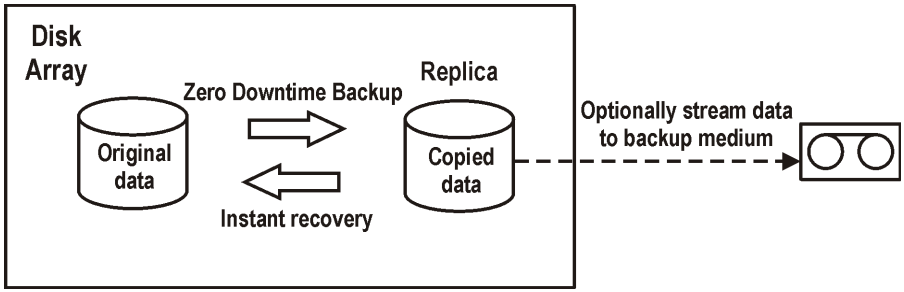

Instant Recovery

Instant recovery is a process in which a replica of data objects, which has been kept on the array, is used to restore the data objects to their states at the time at which the replica was created, avoiding the need to perform a restore from tape. Depending on the application/database concerned, this may be all that is required, or other steps, such as the application of transaction log files, may be required for full recovery.

What Is a Replica?

In recent years, a common approach to producing replicas on large arrays has been to maintain an exact copy (a so-called mirror) of application data that has to be operated upon. With mirror technology, during normal operation, each time the application data is updated, an identical update is made to the mirrored data and the two are kept synchronized. In this way two identical sets of data are constantly maintained.

When an administrative task (such as performing a backup) has to be performed on the data, the synchronization can be stopped (the mirror can be split) and the mirrored data can be used for the task, leaving the application to continue virtually unaffected, using the original data. If

necessary, after the work on the mirrored data is complete, the synchronization of the two sets of data can be resumed until mirrored data is required for another administrative task.

As array technology has advanced, a number of new techniques for producing copies of application data have been developed. Instead of maintaining a constant mirror of application data, these techniques create a dynamic copy (or snapshot) of application data required for administration purposes. Snapshots are generally virtual copies and contain no real data when first created: only pointers to the original data are created at first and data is copied later. This is why more dynamic copy creation is possible. More details on snapshot creation are given later.

A range of snapshot techniques is now available on several different types of disk array, each with its own particular merits from the storage and data processing points of view.

Replica Definition

A **replica**, therefore, can be regarded as a copy of application/filesystem data at a particular point in time. Depending on the hardware/software with which it is created, it may be an exact duplicate, or a virtual copy of the data being backed up.

Replica Types

Currently, there are currently two basic families of replicas:

- **Split-mirror**

  This type of replica is produced using mirroring technology, provided by array hardware such as the HP StorageWorks Disk Array XP and the EMC Symmetrix Disk Array.

  Mirroring technology allows a duplicate of filesystem/application data to be created and maintained during normal application use. This duplicate is called a “mirror” of the source (or original) data. During normal application usage, the mirrored data can be kept synchronized with the source data, that is, any data updates to the source data are also applied to the duplicate.

  If a permanent replica of the data at a fixed point in time is required, the synchronization is stopped (the mirror is split), leaving an independent *split mirror* replica of the source data.

For more detailed information, see “Local Split Mirror Replication” on page 19.

- **Snapshot**

  This type of replica is produced using snapshot technology, provided by array hardware such as the HP StorageWorks Enterprise Virtual Array (VA) or the HP StorageWorks Enterprise Virtual Array (EVA).

  A snapshot replica can also be regarded as a copy, or image of filesystem/ application data. However, depending upon the type of array, it may not be a data duplicate, as a split mirror replica is, but rather a virtual copy, with pointers to the original data rather than to copied data in separate storage locations.

  From an application point of view, a snapshot replica can be considered as being created at one particular point in time, but generally, background replication processes continue for some time afterwards. However, pre-configuration, such as setting up a synchronized mirror, is not required.

  For more detailed information, see “Local Snapshot Replication” on page 21.

How Are Replicas Used in Zero Downtime Backup and Instant Recovery?

Zero Downtime Backup

As said earlier, the aim of zero downtime backup is to reduce the effect of backup on application/filesystem performance.

This can be achieved in many ways and, for applications that are tightly integrated with Data Protector, various techniques are available.

Online Backup

For database applications, such as Oracle, it is possible for Data Protector to ask the application to place the database into hot-backup mode prior to performing a backup. In this mode, the database is placed into a state in which the database files can be copied and increased information is written to the transaction logs (required to make the database consistent afterwards). This allows the database to be operated upon, without stopping the application.

While the database is in this mode, it is possible to create a replica of the sections of the database to be backed up. After the replica has been created, the database can be switched back to normal operational mode. This is a form of "online" backup. However, with ZDB, the hot-backup database mode is only required while the replica is created, not while subsequent backup operations, such as streaming to tape, are performed.

Offline Backup

Alternatively, database operation can be stopped, without the use of transaction logs, while the replica is created. This is a form of "offline" backup, which stops the application, but again, with ZDB, only for the period during which the replica is created.

In both cases, the effect of the backup process on the application is limited to the period during which the replica is created, much less than with standard tape backup techniques. 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.

Once the replica of the data to be backed up has been produced, firstly, it must be mounted to a backup system (to take full advantage of the ZDB system, it should be a separate computer system) connected to the array on which the replica has been created. Then, it can be used in various ways during the rest of the backup process:

- Data can be streamed from the replica to a tape backup device. In this case, the total backup process is called **ZDB to tape**. Data is normally restored using standard restore from tape, so no information about the replica is stored and it can be discarded after the backup process is complete.
- The replica can be kept on the array and used as the backup. In this case, the total backup process is called **ZDB to disk**. Data backed up using this method can be restored directly from the replica using instant recovery functionality. If replicas are created for this purpose, important array related information about the replica must be recorded, to allow the data to be restored.

  With ZDB to disk, one or more replicas of the data backed up can be kept on an array. In fact, a time-based series can be set up, with each replica corresponding to a particular point in time.
- Data can be streamed from the replica to a tape backup device *and*, afterwards, the replica can be kept on the array and used as a backup. In this case, the total backup process is called **ZDB to**

disk+tape. This provides extra flexibility: data can be restored both using standard Data Protector restore from tape (allowing restore of individual backup objects) and directly from the replica using instant recovery functionality (allowing recovery of the complete replica).

The three methods make up the Data Protector ZDB family.

Instant Recovery

Under normal circumstances, backups are performed regularly and restores are performed infrequently. So, in many cases, users may be content to perform ZDB to tape and accept the time required to perform a restore from tape, if necessary.

However, particularly with high availability systems, the ability to perform a high-speed restore as well as high-speed backups may be essential. In such cases, it is better to perform ZDB to disk or ZDB to disk+tape, so that high speed restore is possible, using instant recovery functionality.

When restoring data with instant recovery, the application and backup systems are disabled and the contents of a replica are restored directly to their original locations. Because the restore is performed internally within the array, it is very high speed. Once the restore has been completed, the sections of the database/file system concerned have been returned to their states at the time the replica was created and the application system can be re-enabled.

Database Recovery

Note that to fully recover a database and make it consistent, it may also be necessary to apply any archived transaction logs, backed up separately, following a restore using instant recovery functionality.

Instant recovery is performed without the need to first restore data from tape.

NOTE

To be available for restore using instant recovery, replicas must have been produced using ZDB to disk or ZDB to disk+tape. Otherwise, information about the replicas, required for the restore, is not recorded within the Data Protector IDB.

Other Methods of Restore from ZDB Sessions

Data backed up to tape using ZDB to tape or ZDB to disk+tape is generally restored using standard Data Protector restore from tape techniques.

Split Mirror Restore

However, there is a special case in which 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 available with arrays that support split mirror replication and is known as split mirror restore.

This is not actually known as instant recovery, but the restore of the replica contents to their original locations is a similar process and it is only necessary to suspend application operation during this stage, minimizing the impact on the application concerned.

Alternative Replica Usage

Replicas are not only used for backup and recovery purposes. There are many applications, such as data mining for which they are very useful. Data Protector ZDB+IR functionality is sometimes used to create replicas for such applications. However, replicas produced for ZDB+IR purposes should not be used for any other purpose, otherwise the validity of the data concerned for restore purposes cannot be guaranteed.

2 Replication Techniques

In This Chapter

This chapter provides an introduction to the replication techniques available for zero downtime backup and instant recovery within Data Protector. After reading it you will know:

- The basic principle behind disk virtualization on disk arrays.
- The types of replication techniques available with the supported disk arrays.
- The array processes used for replication on the same array (local replication).
- The array processes used for replication on a different array (remote replication).

Array Basics

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

Data Protector supports replication on the following disk arrays:

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

These allow a range of replication techniques within the split mirror and snapshot families to be supported. For further details, refer to Chapter 3, “Data Protector ZDB+IR Operational Overview,” on page 35.

All of these arrays support disk virtualization techniques, which allow the creation of virtual disks, logical volumes, and so on.

Figure 2-1 Disk Virtualization

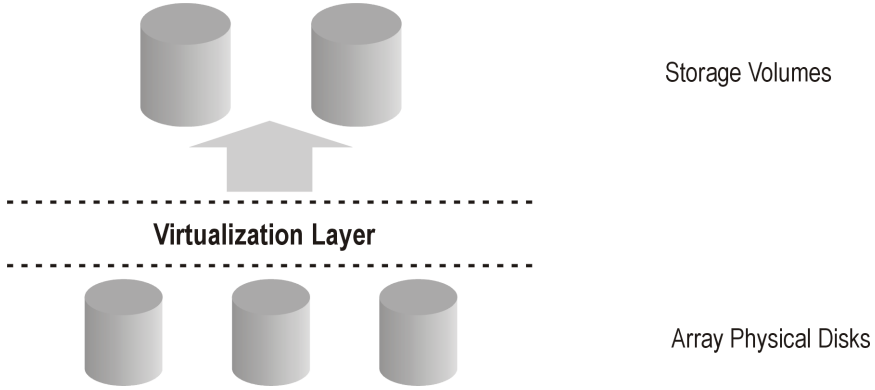

An array of physical disks is configured in such a way that it appears as one large block of data storage. This can then be divided into a number of virtual storage blocks, which are presented to the host/operating system.

These blocks can go under a variety of names, but basically the techniques for their production are similar and, for simplicity, here we will consider them all as **storage volumes**.

RAID Technology

If the disk array concerned uses **RAID technology**, a second level of virtualization is applied to the available storage by the RAID system, to provide data redundancy and improved data protection.

Figure 2-2 Disk Virtualization with RAID

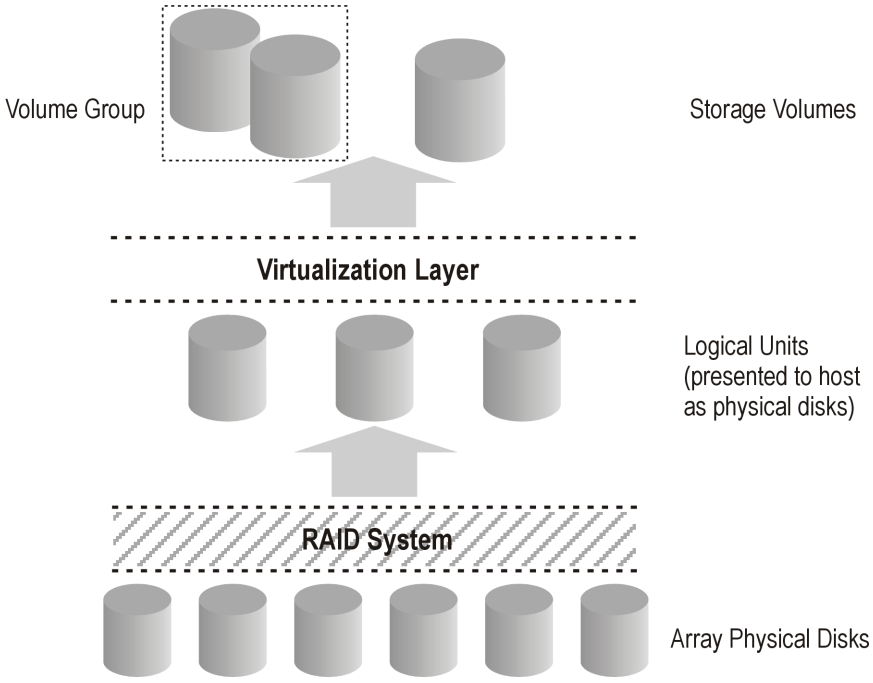

Various RAID levels are available, providing different levels of data redundancy, speed, and access time. In some cases it is possible to adjust the balance between these attributes according to the amount of free storage available.

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.

NOTE

Data Protector can operate on the logical unit level if used on Windows systems, and on both the logical unit and volume groups level if used on HP-UX or Solaris systems (volume groups level is possible if the Logical Volume Manager is used).

Local Replication Basics

Of the array integrations supported by Data Protector, those for the HP StorageWorks Disk Array XP and EMC Symmetrix Disk Array support split mirror techniques and those for the HP StorageWorks Virtual Array and HP StorageWorks Enterprise Virtual Array support snapshot replication techniques. In both cases, the basic ideas of replication are the same, in that copies, or images, of the volumes containing the specified source data, or data objects, are produced. These copies are produced in other logical volumes on the same array, which can then be presented to a host system.

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

Source Volumes/ Target Volumes

The volume(s) containing the source data or data objects to be replicated are known as the **source volumes**. These are replicated to an equivalent number of **target volumes**, which contain the replicated data. When the replication process is complete, the data in the target volumes produced constitute the replica.

From the operating system point of view, the contents of a replica of a particular set of source data objects is the same, irrespective of the method 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.

There can also be certain hardware or software limitations dependent on the array type.

All these aspects should be taken into account when deciding which replication technique to use for a particular purpose. To make this task easier, the principles of operation for each technique are described in the following sections.

Local Split Mirror Replication

In disk array terms, a **mirror** is a copy of one or more source volumes.

Figure 2-3 Split Mirror Replication

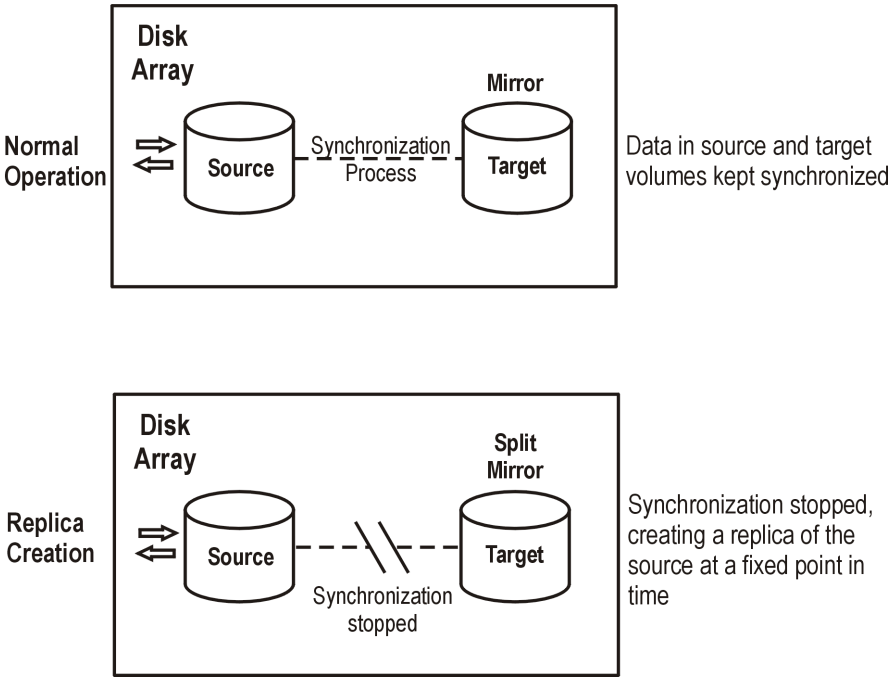

When a mirror is first created, the data contained in it is identical to that in the source volumes, i.e., it is a data duplicate of the source volumes. Once established, the mirror volumes (target volumes) can be kept synchronized with the source volumes: any updates to the contents of the source volumes are applied equally to the contents of the target volumes keeping the contents identical.

The source and target volumes are often referred to as mirrored volumes, mirrored disks or as a mirrored pair.

When a replica is required that is fixed at a particular point in time, the synchronization between the mirrored volumes is stopped (This is known as splitting the mirrors) leaving a fixed copy, or independent split mirror replica of the source volumes. This process is very fast and has minimal impact on the application system.

The split mirror replica produced is:

- A complete duplicate (or clone) of the source volumes.

  With a clone:

  — From a host/operating system point of view, the contents of the target volumes are identical in every way to the contents of the source volumes at the point that the replica was created.

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

- Completely independent of the original. Because of the complete physical copy, if the contents of the source volumes are lost or corrupt, the contents of the target volumes are not affected.

Local Snapshot Replication

In modern disk array terminology, the term “snapshot” has come to cover a family of replicas that are created at the point in time that they are required.

Snapshot replicas are considered to be created at one particular instant and are immediately available for use. However, background copying processes may continue for some time afterwards. This will be explained more fully below.

The common factor that differentiates snapshot replicas from split mirror replicas is that no synchronized copies are required before creating the fixed point-in-time copies. This is because, at the time a snapshot is created, no data is copied. Only a copy of the pointer table to the original data is created. So the replica of the source volumes presented to a host at this point is, effectively, a “virtual” copy. It should be stressed, however, that from a host/operating system’s point of view, the replica always contains a full copy of the source volumes at the time it was created.

The processes involved in snapshot creation, therefore, are very different to those involved in split mirror creation. To explain how the different types of snapshots are created in a way that will allow you to judge their relative merits, the operations involved are described below at the physical storage (physical disk or logical unit) level, not at the storage volume level. This should be born in mind while reading this section.

The supported array integrations allow you to create the following family of snapshots:

- Pre-allocated (or fully-allocated) snapshot.
- Virtually capacity-free (or demand-allocated) snapshot (VSNAP).
- Snapclone.

Pre-Allocated Snapshot

Also known as “fully-allocated snapshot”, “standard snapshot” or frequently just as “snapshot”, this form of replica is created/maintained as shown in the following figure and described more fully below:

Figure 2-4 Pre-Allocated Snapshot Creation

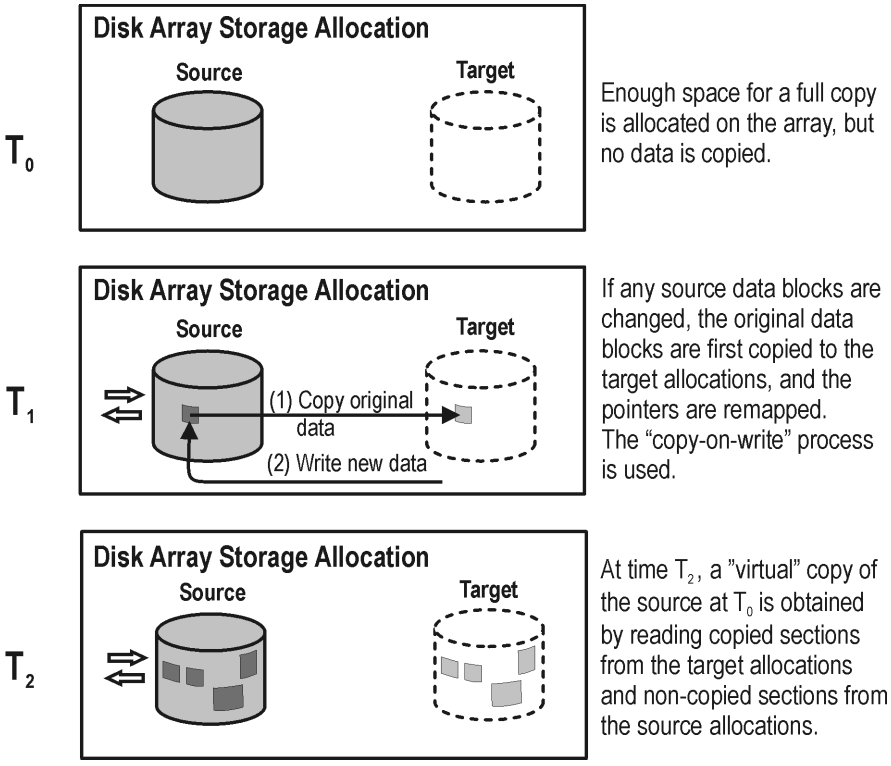

1. At time T_0, storage capacity equal to that taken up by the source volumes concerned is allocated on the array for the target volumes.

   No data is copied from the source storage blocks to the target storage blocks. Instead, pointers are set to point to the storage blocks holding the original data. At this point, the copy is completely virtual, but, from a host’s perspective, 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 is due to be updated, it is first copied to target storage blocks and the pointers are remapped to these blocks before the update is performed. There is therefore, still a complete record of the original data on the array.

   This procedure is also known as “copy on write”.

3. The snapshot is now partly real (the section for which source data has been copied) and partly virtual. When the replica is accessed at any point in time T_2, 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.

The pre-allocated snapshot produced:

- Does not contain a complete *independent* duplicate of the original data (even though, at a later stage, every single storage block in the source volume has been changed since replication).
- Has adequate space guaranteed for the snapshot, even if all the data in the source volumes is changed.
- Is space-inefficient: enough space is always reserved for all the data to be changed, but normally only part of this space is used. While the snapshot exists, the rest of the reserved space cannot be used for any other purpose.

Virtually Capacity-Free Snapshot (VSNAP)

With this type of snapshot, also known as “demand-allocated snapshot”, no storage capacity is reserved at the start. Otherwise, the creation/maintenance process is very similar to that for the pre-allocated snapshot and is shown in the following figure and described more fully below:

Figure 2-5 VSNAP Creation

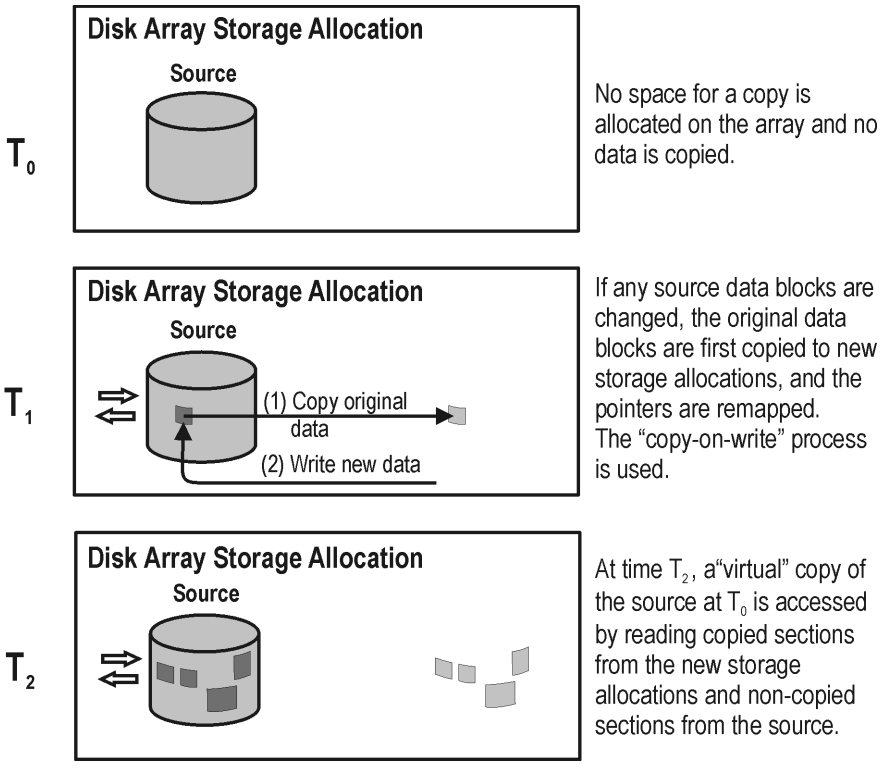

1. At time T_0, no storage capacity is reserved on the array for the target volumes.

   No data is copied from the source storage blocks to new (target) storage blocks. Instead, pointers are mapped to the storage blocks holding the original data and the copy is completely virtual. But, from

a host's perspective, a complete replica of the source volumes at time T_0 exists in the target volumes and it is ready for use. It takes up no storage space other than that required for the pointers.

2. After snapshot creation, if T_0 source data is due to be updated, it is first copied to new storage blocks and the pointers remapped to these blocks before the update is performed. There is therefore, still a complete record of the original data on the array.

   This is again the "copy on write" procedure, as used with the pre-allocated snapshot.

   In this case, the storage space taken up is only that required for the changed data. If all the original data were updated, the space required would be the same as for the pre-allocated snapshot.

3. When the replica is accessed at any point in time, T_2, The copy is partly real and partly virtual. Any previously copied data is read from the new 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.

The VSNAP produced:

- Does not contain a complete *independent* duplicate of the original data (unless, at a later stage, all the original data in the source volumes has been changed).
- Is space-efficient.
- Requires independent disk capacity management to guarantee enough space for replica growth. Otherwise, if space on the array runs out, not only could snapshot update fail, but it could affect general array operation.

Snapclone

This form of snapshot starts as a pre-allocated snapshot and ends up as a clone, or duplicate, similar to a split mirror replica. This is accomplished as shown in the following figure and described more fully below:

Figure 2-6 Snapclone Creation

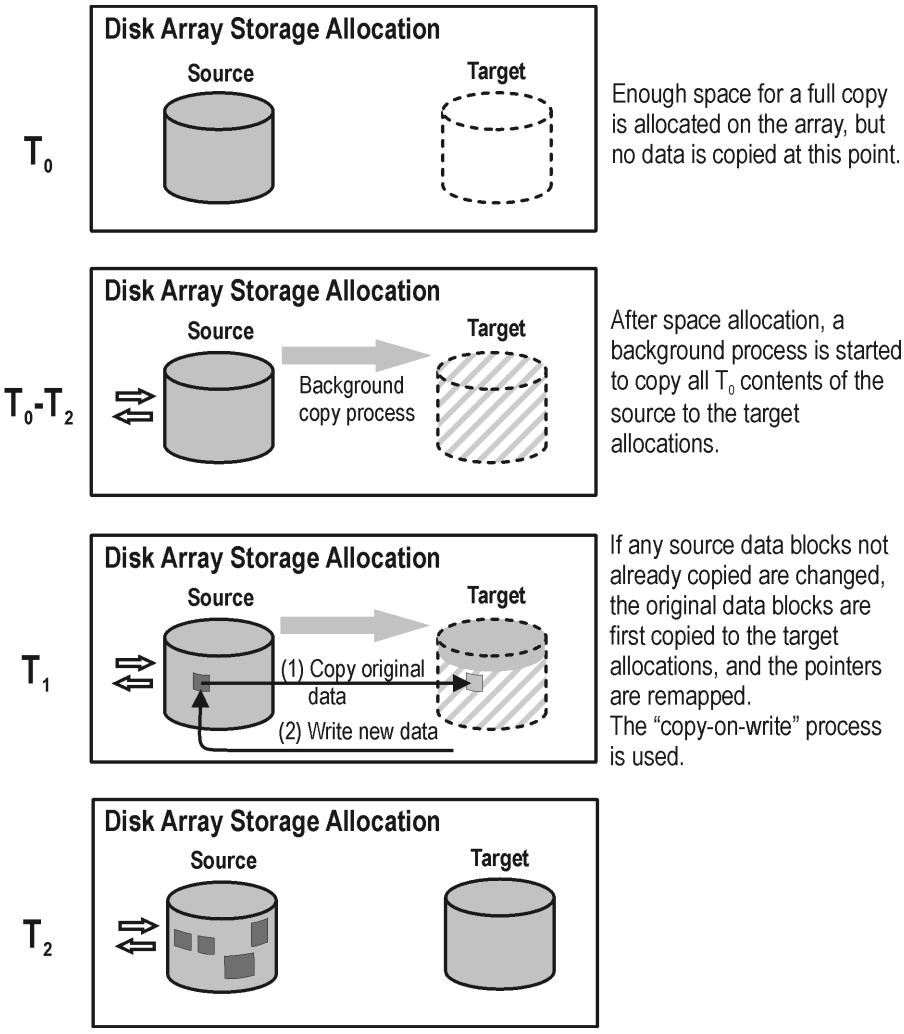

1. A pre-allocated snapshot is created. For details, refer to

"Pre-Allocated Snapshot" on page 22, steps 1 and 2. Once it is created, it is fully usable as a normal pre-allocated snapshot, updated using the same "copy on write" procedure.

2. A background process is started to copy all the data from the source storage blocks to the target storage blocks. During this process, if the snapshot is required for use, the copy is partly virtual and partly real: any data that has not been copied is accessed from the source storage blocks and any that has been copied is accessed from the target storage blocks.
3. When all data has been copied to the target storage locations, the background process is stopped and a standalone clone, or duplicate of the source at point in time T_0 remains.

The snapclone produced (after the background copying process has completed) is:

- A complete duplicate (or clone) of the source volumes.

  With a clone:

  — From a host/operating system point of view, the contents of the target volumes are identical in every way to the contents of the source volumes at the time the replica was created.

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

- Completely independent of the original. Because of the complete physical copy, if the contents of the source volumes are lost or corrupt, the contents of the target volumes are not affected.

However, while the background data copying process is running, it can have an impact on application performance, due to the competition for resources. This copying process could take a significant amount of time if snapclones of large databases are being produced.

General Snapshot Definition

From the above descriptions of the supported snapshot family, you will see that a general definition for a snapshot could be considered as follows.

A copy, either real or virtual, of application or filesystem data contained in storage volumes. It is regarded as being created at one particular point in time and does not have to exist beforehand (unlike with split

mirror creation). However, copying processes may continue for a period after initial creation. During this time, the copy can be a mixture of real and virtualized disk blocks.

Remote Replication Basics

In addition to local replication, it is also possible to perform remote replication. In this case, replicas are created on a separate array, not on the array containing the source volumes.

For this method, mirroring techniques are generally used.

Remote Split Mirror Replication

With this technique, mirrored volumes (or a mirrored pair) are set up in a similar way to those used in local split mirror replication, but the source and target volumes are generally on separate disk arrays. This is because the technique was primarily developed for disaster recovery applications on high availability systems.

Once established, the mirror (target) volumes on the remote array are kept synchronized with the source volumes on the local array, that is, any updates to contents of the source volumes are applied equally to the contents of the target volumes, maintaining them as duplicates, as with local split mirror replication. However, this time, the synchronization may have to take place over several kilometers, if the arrays are installed at separate sites, and application performance can be adversely affected. This is because normally, for data security reasons, the link to the remote system is synchronous.

Using this technique, a mirror duplicate of volumes containing application data can be constantly maintained on the remote (target) disk array.

When a replica is required that is fixed at a particular point in time, the synchronization between the mirrored volumes is stopped (splitting the mirrors) leaving, on the remote array, a fixed copy, or independent split mirror replica, of the source volumes.

Data Protector supports the use of mirroring techniques for remote replication with the HP StorageWorks Disk Array XP and EMC Symmetrix Disk Array.

A split mirror replica produced on the remote array in this way is:

- A complete duplicate of the source volumes on the source array.
- Completely independent of the original. If the contents of the source volumes on the local array are lost or corrupt, the contents of the split mirror target volumes on the remote array are not affected.

For further information on split mirror creation, refer to “Local Split Mirror Replication” on page 19.

NOTE Continuous remote synchronization over a long link could potentially have a performance impact on the application system.

Remote Plus Local Replication Basics

Remote plus local replication is composed of both techniques: duplicate volumes are created on a remote array using remote replication, and then used as the source for a local replication. This is done in exactly the same way as described in “Local Replication Basics” on page 18, but this time, the target volumes of the remote replication become the source volumes for the local replication.

If you want to regularly create independent, fixed point-in-time replicas on a remote system, you will need to use this combined remote and local replication. This might be desirable for disaster recovery purposes.

For this method, mirroring techniques are generally used.

Remote Plus Local Split Mirror Replication

Remote Replication

With this technique, for the remote replication, mirrored volumes (or a mirrored pair) are set up with the source and target volumes on separate disk arrays as with remote split mirror replication.

Once established, the mirror (target) volumes on the remote array are kept synchronized with the source volumes on the local array, that is, any updates to contents of the source volumes are applied equally to the contents of the target volumes, maintaining them as duplicates. The link between the arrays is synchronous.

Using this technique, a mirror duplicate of the source volumes on the local array can be constantly maintained on the remote (target) disk array.

Local Replication

In this case, the target volumes of the remote replication stage become the source volumes for the local replication stage (on the remote array).

The local replication stage is then performed in the same way as described in “Local Split Mirror Replication” on page 19.

When a replica is required that is fixed at a particular point in time, the synchronization between the remotely mirrored volumes is not altered, but the synchronization between the locally mirrored volumes is stopped (splitting the mirrors) leaving, on the remote array, a fixed copy, or independent split mirror replica, of the source volumes on the local array.

Data Protector supports the use of mirroring techniques for remote plus local replication with the HP StorageWorks Disk Array XP and EMC Symmetrix Disk Array.

A split mirror replica produced on the remote array in this way:

- Is a complete duplicate of the data in the source volumes on the source array.
- Is completely independent of the original. If the contents of the source volumes on the local array are lost or corrupt, the contents of the split mirror target volumes on the remote array are not affected.
- Has no impact on the link between the local and remote arrays, which can be maintained continuously.

For further information on split mirror replicas, refer to “Local Split Mirror Replication” on page 19.

3 Data Protector ZDB+IR Operational Overview

In This Chapter

This chapter explains Data Protector operation from the perspective of zero downtime backup and instant recovery. After reading it you will know:

- How a cell is set up for ZDB+IR purposes.
- How the required Data Protector cell components interact.
- The array integrations supported by Data Protector and the techniques each array uses for replication and recovery.
- Application integrations that support ZDB+IR and types of replication available with these integrations.

Data Protector Cell Concept

Data Protector uses the concept of the managed cell. How a cell is set up for ZDB+IR purposes is briefly described below.

Figure 3-1 Data Protector Cell Schematic for ZDB+IR

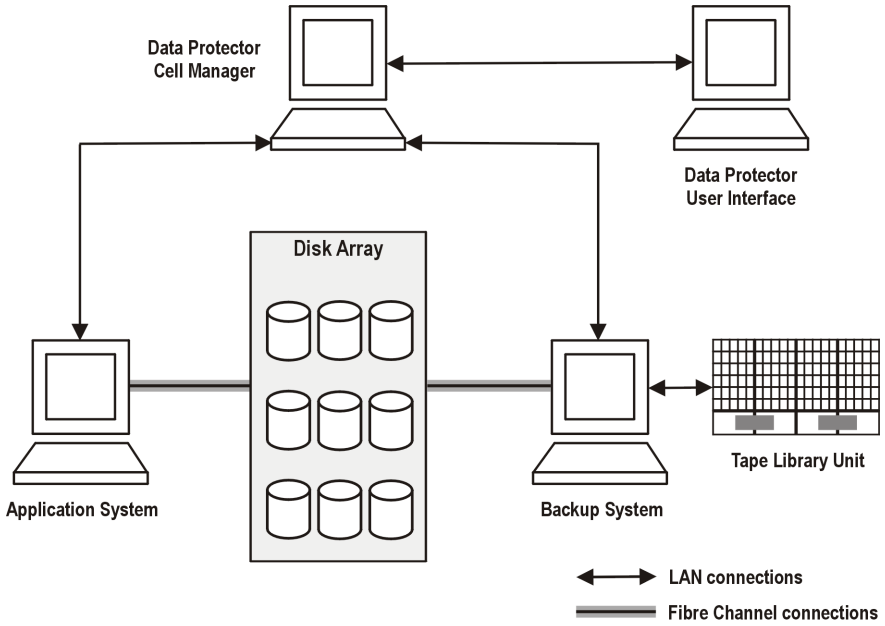

Within a **Data Protector cell** a single, central control server, the **Cell Manager**, controls:

- Data Protector **clients** installed on applications systems. For simplicity, at this stage, only one application client system is shown.
- A separate backup system with, for tape backup, a tape device attached.

To be able to use ZDB+IR techniques, the application database or file system data to be backed up must be on a disk array, to which the application and backup systems are both directly attached. The tape device becomes optional for ZDB+IR applications.

In addition, there is a user interface, consisting of a **Graphical User Interface (GUI)** and a **Command-Line Interface (CLI)** that is used to control the Data Protector operations. The user interface does not necessarily have to be on a system with any other Data Protector cell components installed: it can be installed on any system(s) connected to the Cell Manager via a LAN connection.

Data Protector Cell Components

A brief description of the Data Protector cell components and how they interact is required here, to be able explain later how the replication techniques required for ZDB+IR are applied in Data Protector.

For the Data Protector cell shown in Figure 3-1 on page 37 to function, the correct operational software components must be installed on each of the hardware elements.

The required components for a typical application integration are shown in Figure 3-2 below and their functions are outlined in the following section.

Figure 3-2 Cell Operational Components for ZDB+IR

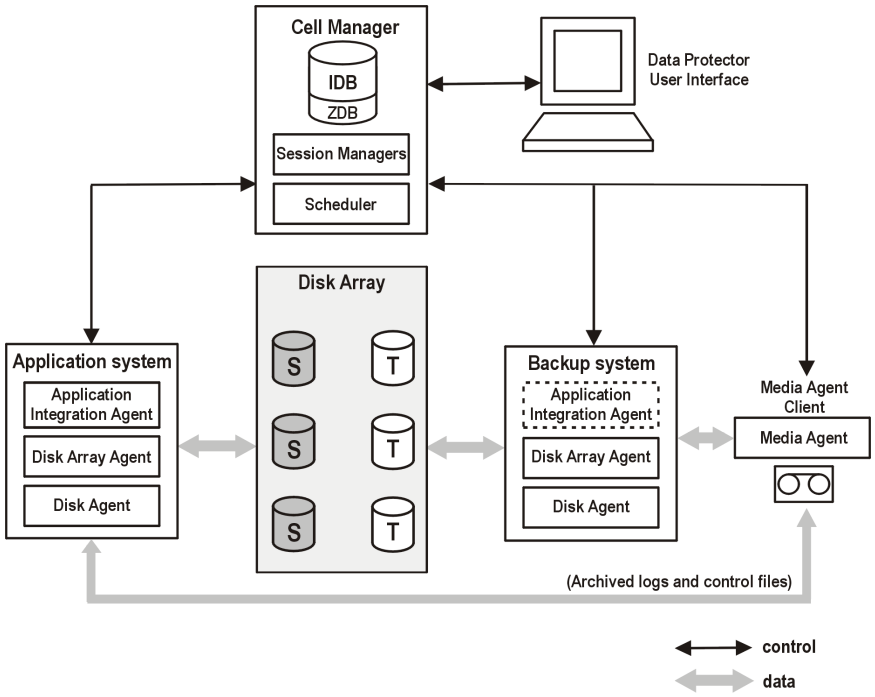

Cell Manager

The Data Protector Cell Manager runs the core Data Protector software and controls all the cell operations required for backup and restore. It performs the following functions:

- Manages Data Protector cell operations from a central point.
- Contains the Data Protector **Internal Database (IDB)**, which holds Data Protector information such as:
  - **backup specifications**.
  - the schedule for running **backup sessions**.
  - the results of backup sessions.

  The IDB contains an extension - the ZDB database section. This has a separate sub-section for each disk array integration installed, holding all the information specific for any backup session involving ZDB+IR and associated with that array. For more information about ZDB database, refer to “Additional Information” on page B-1.

- Contains the **scheduler**, which is responsible for starting non-interactive backup sessions according to the backup specifications.
- Controls **session managers** that are responsible for such things as:
  - — controlling and terminating running backup sessions.
  - — writing backup session information to the Internal Database (IDB).
  - — controlling and terminating running restore sessions.

Application Systems

Each application system for which replicas are to be created must have a Data Protector client installed. In addition, it must have:

- A **Disk Array Agent (DAA)**. This is used to control the interaction between the Data Protector Cell Manager and the array on which the application database/file system is installed. Each type of array supported has its own dedicated agent.
- An **Application Integration Agent (AIA)**. This acts as the interface between the Data Protector Cell Manager and the application and is required for Data Protector to be able perform functions such as controlling the state of the database during the backup and restore sessions for database applications.

Backup System

This system must have a Data Protector client installed, together with the relevant Disk Array Agent. In some cases, it may also require an Application Integration Agent. It is the system to which a replica is presented after it has been 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.

Various checks and administration functions are also performed by this system.

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

Systems with a User Interface

Data Protector can be controlled from any system on the network on which the Data Protector GUI or CLI is installed. This means that, even if the main cell components are located in a computer room, Data Protector operation can be managed from an office desktop system. Refer to “User Interfaces” on page 42 for more details.

ZDB+IR operations can be controlled using the GUI and the CLI.

Component Interaction

The Cell Manager sends instructions to the Application Integration Agents and Disk Array Agents on the application and replica administration systems. Here the instructions are interpreted and passed on to the array. All actual replication and manipulation is performed within the array itself.

User Interfaces

Both the Data Protector graphical user interface (GUI) and command-line interface (CLI) can be used to perform the operations required for zero downtime backup and instant recovery.

GUI

GUI allows you to administer your ZDB environment from a single system. Using the GUI, you can perform various backup tasks, monitor your active operations and use Data Protector reporting and notification capabilities. Backup tasks include creating a backup specification, defining the backup options and the schedule, and starting a ZDB session. In the `Instant Recovery` context, you can browse for the sessions marked for instant recovery, define the necessary options and start an instant recovery session. Note that instant recovery is possible for the data objects stored on a disk array (backed up in ZDB-to-disk or ZDB-to-disk+tape sessions). The data stored on a backup medium (backed up in ZDB-to-tape or ZDB-to-disk+tape sessions) is to be restored using the standard Data Protector restore from tape procedure. These sessions are displayed in the `Restore` context. Thus, the Data Protector GUI provides the easy differentiation between the sessions marked and not marked for instant recovery.

Figure 3-3 on page 43 is an example of the GUI window, where the backup options for a ZDB session are defined.

Figure 3-3 Data Protector GUI

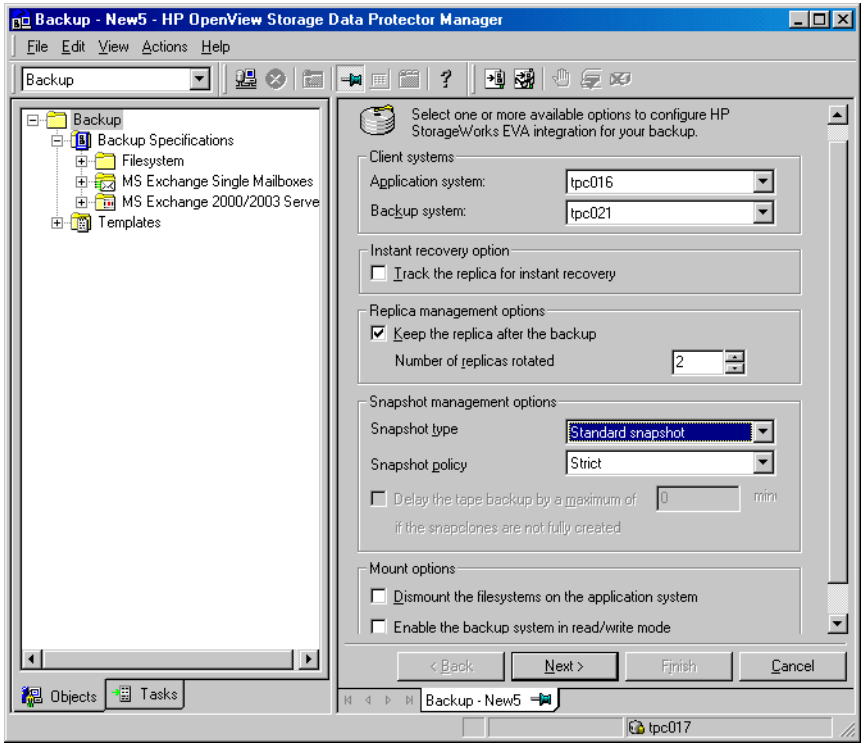

CLI

Most of the ZDB+IR operations available in GUI can be performed using the CLI. However, some administration tasks can be done using the CLI only. These include querying, synchronizing and purging the ZDB database, checking its consistency, and some other. In addition, using the CLI, you can manually delete a replica or a replica set when it is no longer needed, together with the information on this replica/replica set stored in the ZDB database. An example below shows you how to delete a single replica created on the HP StorageWorks Enterprise Virtual Array:

```
omnidbeva -delete -session <session_ID>
```

Array Integrations Available with Data Protector

Array integrations with Data Protector are available for the following disk arrays:

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

All of these integrations are capable of creating replicas and, in most cases, replica sets that can be used for the purposes described in Chapter 1, “Overview,” on page 1. Each, however, uses different techniques for replication and recovery and has its own advantages and disadvantages for the various applications of the techniques.

HP StorageWorks Disk Array XP

The Data Protector HP StorageWorks Disk Array XP integration supports the creation of replicas using the split mirror technique. This means that the target volumes (T) produced are exact duplicates of the source volumes (S).

With HP StorageWorks Disk Array XP, you can use local and remote replication techniques, or combine the two for the best level of data protection. The basic concepts, as well as advantages and disadvantages of each technique are described below in this section.

NOTE

Replicas created using local replication can be restored using either instant recovery or split mirror restore functionality. If you use remote or remote plus local replication techniques, you restore your data using the standard Data Protector restore from tape procedure described in the *HP OpenView Storage Data Protector Administrator’s Guide*.

For the overview of all supported Disk Array XP configurations, refer to “Supported HP StorageWorks Disk Array XP Configurations” on page A-3.

Local Replication

For local replication, the HP StorageWorks Business Copy (BC) XP configuration is used, which allows the creation of three **first-level mirrors** (replicas) to be used for instant recovery (refer to “Supported HP StorageWorks Disk Array XP Configurations” on page A-3 for more information). It means that a replica rotation set created for instant recovery purposes consists of a maximum of three replicas. Once established, BC operations continue unattended, providing the data replication within the same array.

Local replication gives you much flexibility in choosing your backup strategy. With this, you can perform any of the three ZDB types supported with Data Protector - ZDB to disk, ZDB to tape, or ZDB to disk+tape. Note that instant recovery is only supported with ZDB to disk or ZDB to disk+tape. If you do ZDB to tape, you can either restore your data following the Data Protector standard restore from tape procedure, or use the Data Protector split mirror restore technique.

For more information about ZDB types, refer to Chapter 1, “Overview,” on page 1 of this guide.

NOTE

For ZDB to tape and ZDB to disk+tape, a separate backup system is connected to the disk array with the target volumes, while the source volumes are connected to the application system. The backup is performed from the replica after the pair has been split, meaning that the application system, during the backup, remains online and available for use.

Figure 3-4 **Example of BC XP Configuration**

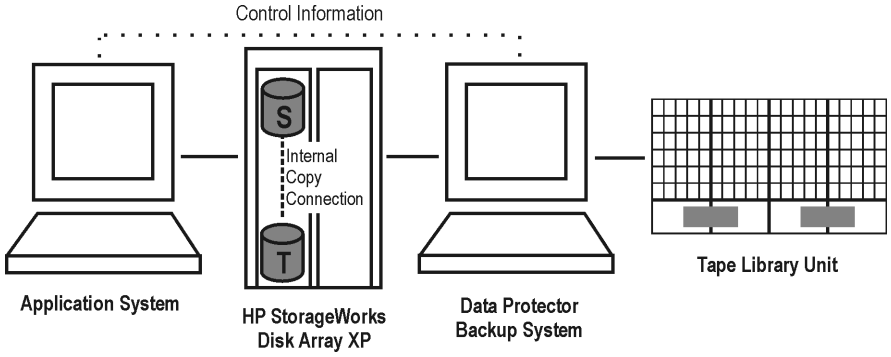

Split Mirror Restore

To restore filesystem objects, disk images, MS Exchange 2000 and SAP R/3 integration data, backed up in ZDB-to-tape sessions, you can either use the Data Protector standard restore from tape procedure or the Data Protector split mirror restore technique. With the split mirror restore, the data is first moved from the medium to a replica, which is then synchronized with the source volume. For more information about split mirror restore, refer to Chapter 6, “Restore Techniques from ZDB-to-Tape Sessions,” on page 83 of this guide.

NOTE

To restore Oracle8/9 and MS SQL integration data from a ZDB-to-tape session, a standard restore from tape procedure is used.

Remote Replication

For remote replication, the HP StorageWorks Continuous Access (CA) XP configuration is used, which allows the creation of remote split mirror replicas on a remote machine, up to 27 miles (43 km) away. Once established, CA operations continue unattended, providing continuous, real-time remote data replication.

CA operations are non-disruptive and allow the source volumes of each CA volume pair to remain online for all client systems for both read and write operations.

A separate backup system needs to be connected to the disk array with the target volumes, while the source volumes are connected to the application system. Streaming of data to tape is performed from the replica after the pair has been split, meaning that the application system, during the backup, remains online and available for use.

The diagram below gives an example of the CA XP configuration.

Figure 3-5 Example of CA XP Configuration

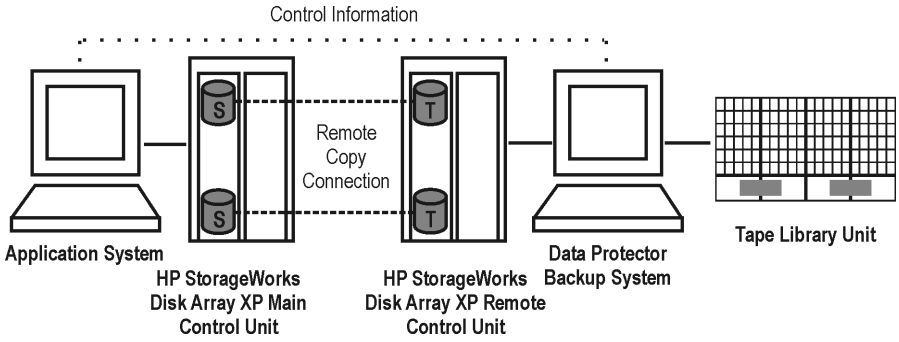

Interface Types

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

- Extended Serial Adapter (ESCON) for distances up to 43 km.
- Fibre Channel (FC) for distances up to 2 km.

The Fibre Channel distance can be extended by using FC switches that have single-mode fibre multiplexors built in.

Remote Plus Local Replication

For remote plus local replication, the combination of HP StorageWorks Continuous Access (CA) XP and HP StorageWorks Business Copy (BC) XP configurations is used. This allows 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 HP StorageWorks Disk Array XPs, located in physically separate sites, are needed for such a configuration.

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

When a replica is required, the integration splits the BC pair. In order to ensure data consistency, the CA pair status is checked before the BC pair split is executed. In a synchronous CA configuration, this ensures that all data from the Main Control Unit is in the Remote Control Unit.

Cluster Configurations

This configuration is supported in a cluster using MC/ServiceGuard or Microsoft Cluster Server. An example of the cluster configuration is presented below.

Figure 3-6 Combined CA+BC Configuration in a Cluster

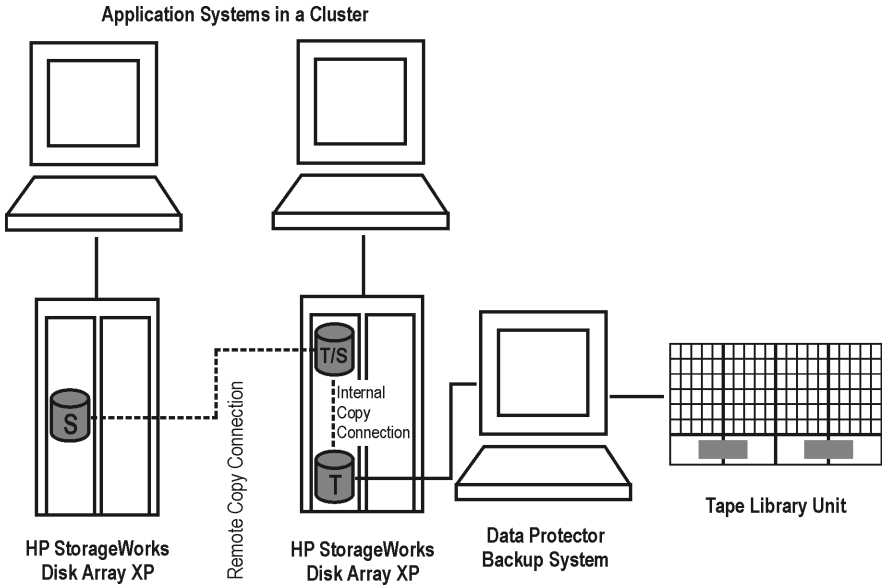

For more information about cluster configurations, refer to Appendix A of this guide and to the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide*.

HP-UX LVM Mirroring

The Data Protector HP StorageWorks Disk Array XP integration supports HP-UX Logical Volume Manager (LVM) mirroring in a configuration where HP StorageWorks Disk Array XP LDEVs are LVM-mirrored from one or more HP StorageWorks Disk Array XP unit(s) to one or more other HP StorageWorks Disk Array XP unit(s). HP-UX LVM mirroring is done on the level of the logical volume. LVM-mirrored

LDEVs and their LVM mirrors belong to the same logical volume. It is recommended that LDEVs in every HP StorageWorks Disk Array XP belong to a different physical volume group, so that extending logical volumes will not give unpredictable results such as mirroring a logical volume onto the same disk. The LDEVs in the array(s) that is/are connected to the backup system need to have their BC pairs assigned. The application system has to be connected to those HP StorageWorks Disk Array XP units that contain LDEVs belonging to LVM-mirrored logical volumes.

In other words, LVM mirroring in such a configuration can be regarded as a substitute for the CA remote copy connection in a combined CA+BC configuration, which thus, from the Data Protector HP StorageWorks Disk Array XP integration point of view, becomes a BC configuration.

Replicas created using LVM mirroring can be used for ZDB+IR. For the detailed procedure on how to perform instant recovery from such replicas, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide*.

An example of the configuration that uses LVM mirroring is given below.

Figure 3-7 LVM Mirroring Configuration on XP

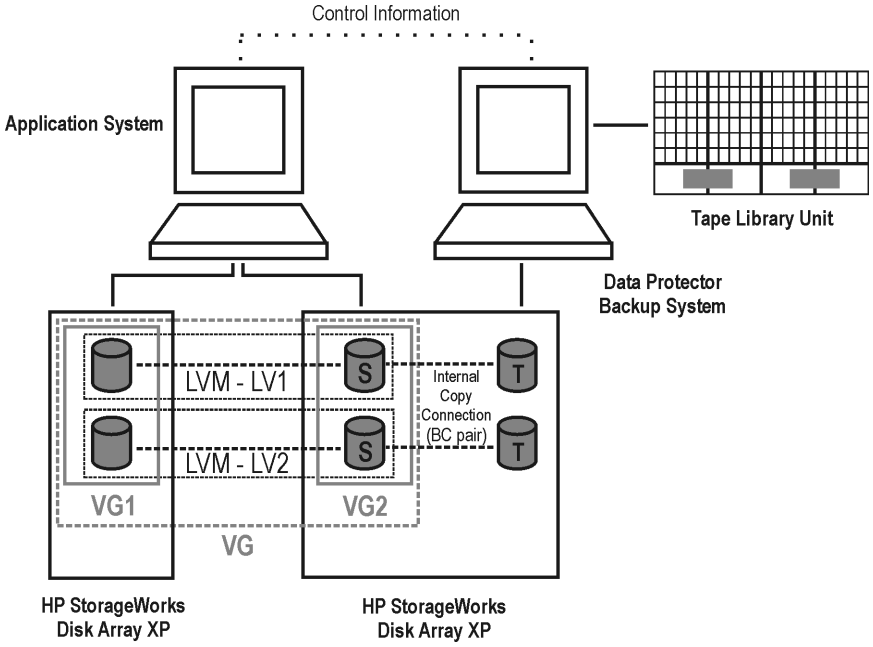

For an overview of all the supported LVM configurations, refer to “Supported HP-UX LVM Mirroring Configurations” on page A-14.

EMC Symmetrix Disk Array

The Data Protector EMC Symmetrix integration supports the creation of single replicas using the split mirror technique. It means that the target volumes (T) produced are exact duplicates of the source volumes (S).

With EMC Symmetrix disk array, you can use local and remote replication techniques, or combine the two for the best level of data protection. The basic concepts, as well as advantages and disadvantages of each technique are described below in this section.

For the overview of the supported EMC Symmetrix configurations, refer to “Supported EMC Symmetrix Configurations” on page A-18.

NOTE Instant recovery is not supported with the Data Protector EMC Symmetrix integration. It means that the only ZDB type available with the integration is ZDB to tape.

Local Replication

For local replication, the EMC Symmetrix TimeFinder configuration is used, which allows the creation of single mirrors (replicas) for ZDB and split mirror restore purposes. Once established, TimeFinder operations run unattended, providing the data replication within the same array.

For a backup, a separate backup (R2) system needs to be connected to the disk array with the target volumes, while the source volumes are connected to the application (R1) system. Streaming of data to tape is done from the replica after the pair has been split, meaning that the application (R1) system, during the backup, remains online and available for use.

For retrieving the data, you can use either the split mirror restore technique, or a standard Data Protector restore procedure. You can use split mirror restore for restoring filesystem objects, disk images, and the SAP R/3 integration data. With this technique, the data is first moved from the medium to the replica, which is then synchronized to its original source on the application (R1) system. For more information about split mirror restore, refer to Chapter 6, “Restore Techniques from ZDB-to-Tape Sessions,” on page 83 of this guide. For the instructions on how to perform a standard Data Protector restore from tape, refer to the *HP OpenView Storage Data Protector Administrator’s Guide*.

An example of the EMC Symmetrix TimeFinder configuration is given below.

Figure 3-8 Example of TimeFinder Configuration

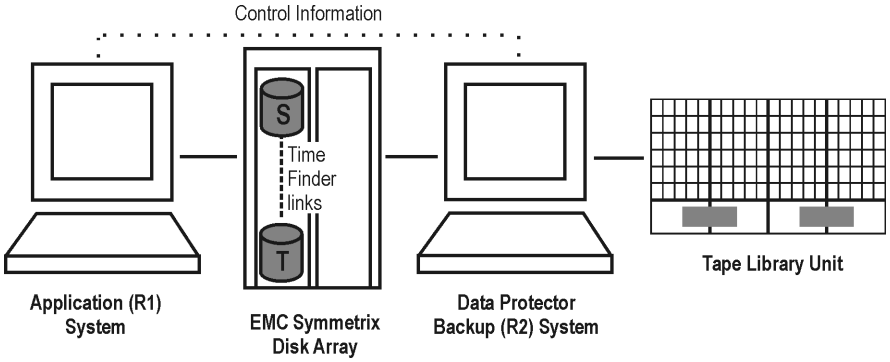

Remote Replication

For remote replication, the EMC Symmetrix Remote Data Facility (SRDF) configuration is used, which allows the creation of a single split mirror replica on a remote machine. At least two EMC Symmetrix Disk Arrays, located in physically separate sites, are used for such a configuration. Once established, SRDF operations continue unattended, providing continuous real-time remote data replication.

Figure 3-9 on page 52 shows an example of the EMC Symmetrix SRDF configuration.

Figure 3-9 SRDF Configuration

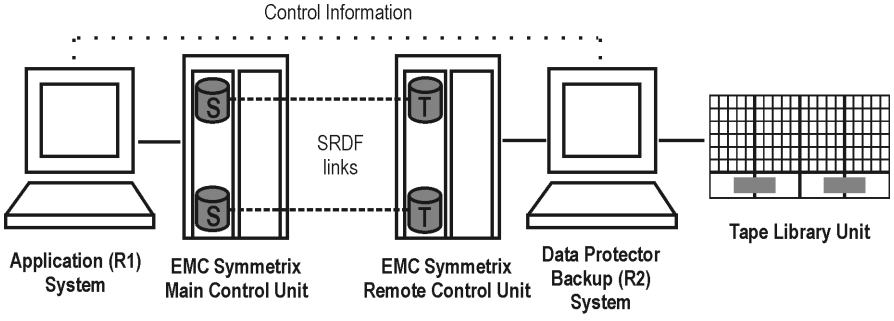

Remote Plus Local Replication

For remote plus local replication, the combination of EMC Symmetrix Remote Data Facility (SRDF) and EMC Symmetrix TimeFinder configurations is used. This allows the creation of a (secondary) split mirror replica on a remote machine, and then creation of local replicas of that secondary replica on that remote machine. At least two EMC Symmetrix Disk Arrays, located in physically separate sites, are needed for such a configuration.

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.

When a replica is required, the integration splits the TimeFinder pair. In order to assure 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.

Cluster Configurations

In order to automate the failover, an application like MC/ServiceGuard can be used.

This configuration is supported in a cluster using MC/ServiceGuard. Figure 3-10 on page 54 presents an example of the cluster configuration.

Figure 3-10 Combined SRDF+TimeFinder Configuration in a Cluster

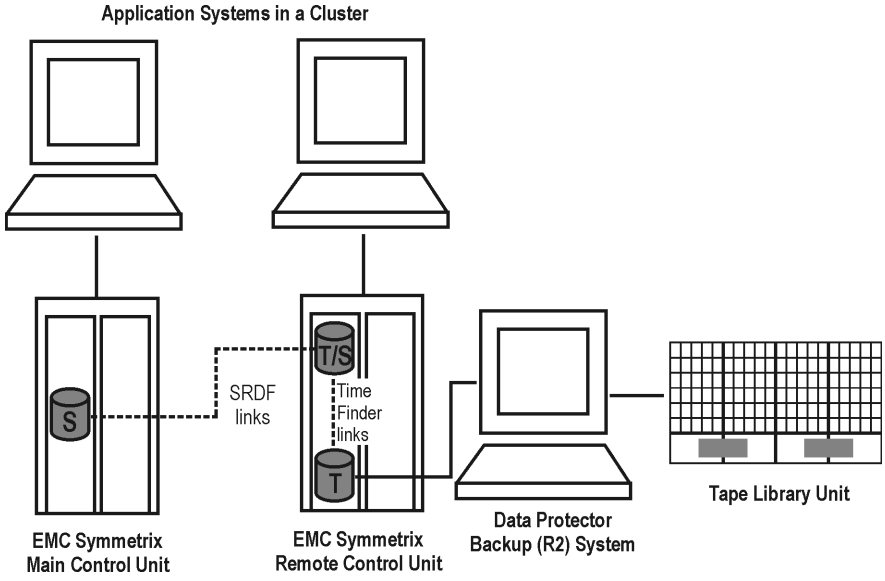

For more information about cluster configurations, refer to Appendix A of this guide and to the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide*.

HP StorageWorks Virtual Array

The Data Protector HP StorageWorks Virtual Array integration supports the creation of pre-allocated replicas and replica sets using snapshot techniques. This means that the target volumes (T) produced are logical copies, or images, of the source volumes (S).

With HP StorageWorks Virtual Array, you can use local replication technique, or a combined configuration of remote and local replication using the **HP-UX Logical Volume Manager (LVM) mirroring**. The basic concepts, as well as advantages and disadvantages of each technique are described below in this section.

For the overview of the supported VA configurations, refer to "Supported Snapshot Configurations" on page A-28.

VA Storage Presentation

HP StorageWorks Virtual Array is an array of physical disks configured in such a way that it appears as one large block of data storage that can be divided into a number of smaller logical storage blocks or **logical units** (**LUN**s). The data written to a LUN can be 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, which allows the data replication within the same array. With this, large replica sets can be used, the number of members being limited primarily by the available space on the array. Once established, BC operations continue unattended, providing local data replication.

Local replication gives you much flexibility in choosing your backup strategy. With this, you can perform any of the three ZDB types supported with Data Protector - ZDB to disk, ZDB to tape, or ZDB to disk+tape. Note that instant recovery is only supported with ZDB to disk or ZDB to disk+tape. If you do ZDB to tape, you restore your data using the Data Protector standard restore from tape procedure.

NOTE The maximum number of snapshot replicas that can be produced on one array is 1024.

For an example of the BC snapshot configuration, see Figure 3-11 on page 56.

Figure 3-11 BC Snapshot Configuration

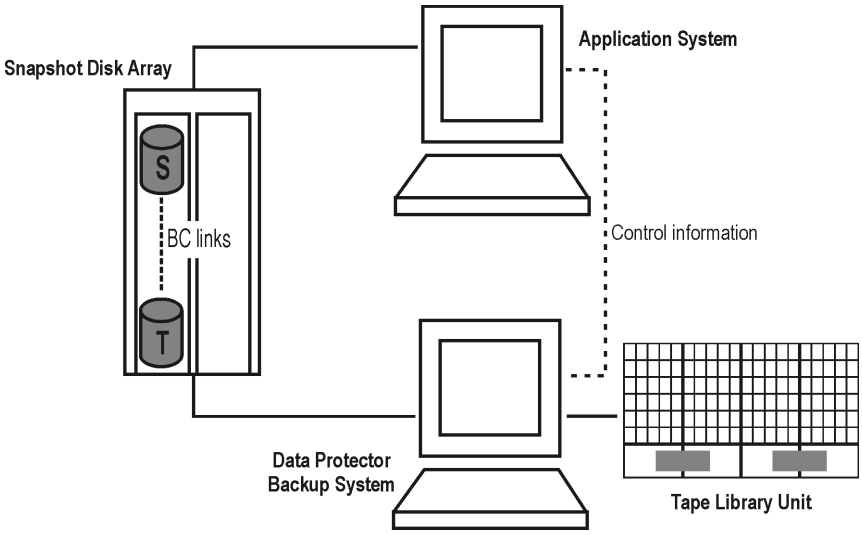

Remote Plus Local Replication

Remote plus local replication is possible using the HP-UX LVM mirroring, which effectively provides a software technique for remote replication. With this, you create replicas on a remote machine and, after that, local replicas of those remote replicas can be created using the snapshot technique.

LVM mirroring can be used for ZDB-to-tape sessions only. It means that you cannot perform instant recovery to restore your data; instead, standard Data Protector restore from tape procedure is used.

At least two HP StorageWorks Virtual Arrays, located in physically separate sites, are needed for such a configuration. The VA source volumes are LVM mirrored from one or more local VA disk array(s) to one or more remote VA disk array(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.

Refer to HP-UX documentation for more information about LVM mirroring.

Figure 3-12 HP-UX LVM Mirroring on VA

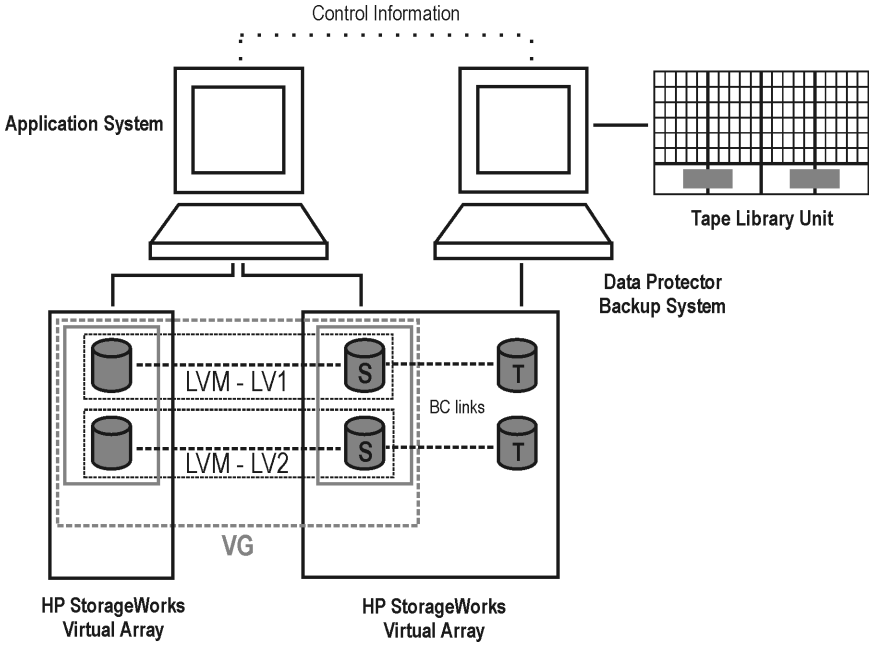

HP StorageWorks Enterprise Virtual Array

The Data Protector HP StorageWorks Enterprise Virtual Array integration supports the creation of pre-allocated, VSNAP, and snapclone replicas and replica sets on the local array using the snapshot technique. This means that the target volumes (T) produced are logical copies, or images, of the source (S) volumes.

With the Data Protector HP StorageWorks Enterprise Virtual Array integration, only local replication is supported. Examples of the configurations are shown in “Supported Snapshot Configurations” on page A-28 of this guide.

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 have control over the exact allocation of virtual disks on physical disks, but you can influence it by choosing different

protection characteristics. For that, the Redundant Array of Independent Disks (RAID) technology is used, which provides various levels of data redundancy, speed, and access time. For more information about the RAID technology, refer to the EVA-related documentation.

Local Replication on EVA

For local replication, the HP StorageWorks Business Copy (BC) EVA configuration is used, which allows data replication within the same array. With this, large replica sets can be used, the number of members being limited primarily by the available space on the array. Once established, BC operations continue unattended, providing local data replication.

Local replication gives you much flexibility in choosing your backup strategy. With this, you can perform any of the three ZDB types supported with Data Protector - ZDB to disk, ZDB to tape, or ZDB to disk+tape. Note that instant recovery is only supported with ZDB to disk or ZDB to disk+tape. If you do ZDB to tape, you restore your data using the Data Protector standard restore from tape procedure.

NOTE The maximum number of VSNAPs and pre-allocated snapshots for a replica rotation set is limited to 7.

For an example of the BC snapshot configuration, refer to Figure 3-11 on page 56.

Application Integrations

Data Protector provides application integrations that support ZDB+IR with the following database applications:

- Oracle8/9
- SAP R/3
- MS SQL Server
- MS Exchange Server 2000

Replication

With Oracle8/9 and SAP R/3 integrations, it is possible to perform the following types of replication:

- Online replication:

  During the creation of a replica, the database on the application system is placed into hot-backup mode. While in this mode, all transactions are cached and the replica of the database can be produced without application operations having to be suspended.

  When the replication process has completed, the cached information is applied to the database and the application is returned to normal operation. Using this method of replication reduces the impact on the application to a minimum, making it suitable for uninterrupted operations.

- Offline replication:

  During replication, the application database is shut down. The time to create the replica is short, but during that time the application is offline, making this method less suitable for high availability applications. With the SAP R/3 integration, instant recovery from offline replication is not possible.

NOTE

With the MS SQL Server integration, you can perform online replication only. With the MS Exchange Server 2000 integration, only offline replication is available, since a filesystem replica is created.

Transaction Logs

When backing up database applications, it is also required to back up separately any archived database transaction logs. To enable this, the application *must not* be configured to use circular logging.

As stated above, only offline replication is supported with the MS Exchange 2000 integration. However, to be able to perform roll-forward recovery, you need to back up transaction log files. Otherwise, only recovery to the point in time at which the replica has been created is possible.

Instant Recovery

With all of the supported application integrations, except for the SAP R/3 integration data backed up in an offline replication session, it is possible to perform instant recovery.

NOTE

With the application integrations, supported by Data Protector, you can recover the data to the point in time, at which a replica was created, using instant recovery. However, to fully recover a database, the transaction logs must be applied afterwards. Using these logs, it is also possible to roll forward from the point at which the replica was created.

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

4 Replica Operations Concepts

In This Chapter

This chapter explains the operations performed on replicas in Data Protector zero downtime backup and instant recovery. After reading it you will know:

- How replicas are created within Data Protector.
- How replicas are manipulated during backup and restore operations.
- What happens to replicas after backup and restore operations.

Replica Operations Concepts

Replication

Replication is controlled by the Backup Session Manager on the Cell Manager. It starts the backup session when instructed to by the Data Protector scheduler, or a user interactively.

ZDB Specification

A Data Protector backup specification defines all the items that Data Protector requires to perform a backup session. A ZDB specification contains all the information required to run a ZDB session, in particular:

- The type of application/file system data to be backed up.
- The source data that are to be to backed up.
- The type of replica (or replica set)* to be created.
- The type of array on which the data resides.
- The application system to be used.
- The backup system to be used.
- Replica management options.
- Replica mounting options.

* Replica set is explained later in this section.

For applications that are not fully integrated with Data Protector, it is also possible to set options to stop the application before replication and restart the application after replication.

After the ZDB specification has been created, it is stored on the Cell Manager in the IDB and can be reviewed or updated at any time.

Backup Session

The Backup Session Manager reads the ZDB specification 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 Application Integration Agent puts the application database/file system into the required state prior to replication: this could be with all database/file system updates stopped for an “offline” replication, or with all database updates re-routed to log files in the case of an “online” replication.

With the database/file system 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 them and, on the backup system, the agent allocates the volumes required for the replica.

The replica is then created on the array, the steps required being performed by the array firmware. For further information on how replicas are created on arrays, see Chapter 2, “Replication Techniques,” on page 13.

Once a replica has been created, you can use it in various ways:

- Keep it on the array for instant recovery purposes.
  If this is done, the array specific details (logical volume information, etc.) for the associated backup session are also recorded in the ZDB database. This information is required for instant recovery purposes.
  This backup process is known as **ZDB to disk**.
- Stream, from the replica, the data objects specified in the backup specification to tape and, afterwards, delete the replica.
  In this case, no instant recovery information is recorded in the ZDB database.
  This backup process is known as **ZDB to tape**.
- Stream, from the replica, the data objects specified in the backup specification to tape but, afterwards, keep the replica on the array for instant recovery purposes.
  In this case, the array specific details (logical volume information, etc.) for the associated backup session are also recorded in the ZDB database for instant recovery purposes.
  This backup process is known as **ZDB to disk+tape**.

Note that, if you want to perform ZDB to disk or ZDB to disk+tape, you must select the relevant option to track the replica storage for instant recovery in the backup specification, otherwise you will only be able to

perform ZDB to tape. You will then be able to choose between ZDB to disk or ZDB to disk+tape when you manually start or schedule your backup, but it will not be possible to select ZDB to tape.

After replication, the Application Integration Agent allows the application to resume normal operation.

If ZDB to tape has been selected, the selected backup objects are then streamed to tape.

After successful backup completion, details of the backup session are saved to the IDB as normal.

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 creating a backup of the log to a file. This can be streamed to tape, if required.

It is generally not recommended to include the log file in the volumes to be replicated.

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

Replica Set Creation

In a Data Protector backup specification, it is possible to either create a single replica with no relationship to any other, or to create a replica as part of a **replica set**.

A replica set consists of a group of replicas with a defined relationship to the source database/file system objects, created using the same backup specification.

In Data Protector, the members of a set can be used in rotation, either interactively or at times specified in the scheduler (refer to “Replica Set Rotation” on page 67).

Replica sets are normally used when creating replicas for instant recovery purposes, where the replicas are kept on the array for the potential recovery of data to its state at specific points in time.

Scheduling Replication

If a specified replication session is to be run automatically, the details of the required times for the specification are entered into the Data Protector **scheduler** and are stored in the IDB. You can either schedule a single session at a specific time, or regular sessions, repeated over periods of days, weeks or months.

When a ZDB specification is displayed in the Data Protector GUI, the schedule information related to it is shown attached to it.

Replica Manipulation

Once you have created your replicas or replica sets, you can handle them in various ways for backup and recovery purposes.

Recovery to a Point in Time

Data Protector allows you to perform recovery of data objects to their states at a particular point in time.

The main step in the process to recover to a point in time is instant recovery.

Instant recovery is a process that you can use to replace lost or corrupted data (or rather, the whole volumes containing them) with known good data, previously replicated to other volumes on an array. The remainder of the process is dependent on the application being recovered.

In the case of a file system that has been replicated, this instant recovery step is all that is required to return the data to its state at the point in time at which the replica was created.

In the case of a database application, after performing an instant recovery procedure, it may be necessary to perform other operations, such as applying archived log files, to fully recovery the database. In this way, it is possible to recover a database to the point in time at which a replica was created, or, to a later point, if log files for that time exist (commonly known as roll forward). For more information, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Integration Guide*.

For maximum flexibility in recovery to a point in time, replicas should be regularly created and kept on the array and log file backups should be regularly performed, thus increasing the time range for which recovery can be performed. However, this eventually becomes very costly on disk array space, so normally a time-based replica rotation set is used, the number of replicas in the set depending on the available disk array space and the time range required.

Replica Set Rotation

If you create a ZDB specification to create a replica for instant recovery purposes, you can specify that a member of a replica rotation set should be used. After saving the specification, the time interval between replica creations within the rotation set is specified in the Data Protector scheduler.

The number of members in the set is specified in the ZDB specification.

Initially, no members of the set exist (in general*) and new replicas are created, at the specified times, and added to the set until the complete set exists on the array.

Once the set is full, the next replica to be created replaces the oldest existing replica in the set. With some array types, this is achieved by directly overwriting the existing replica, in other cases it can only be achieved by first deleting the existing replica and then creating the new one. The differences are because of hardware limitations.

* When a new set is first created, in general, volume allocations for the replicas are done automatically, but with an HP StorageWorks Disk Array XP, it is necessary to pre-configure them.

NOTE

It is only possible to create replica sets for instant recovery if full replica details are saved to the ZDB database. For this, you must specify that you want to track the replicas for instant recovery purposes in the ZDB specification.

ZDB to Tape

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

Since the replica is not to be kept in the array for instant recovery, instant recovery specific information is not written to the ZDB database during creation.

After creation, the replica is mounted on the backup system and the (replicated) backup objects specified in the backup specification are streamed to tape (or other backup medium) in the same way as if they were source backup objects.

After the streaming process, the replica is no longer required for backup purposes, so it is normally deleted from the array. However, if you want to, you can choose to keep them on the array, singly or as part of a replica set, but they will not be available for instant recovery.

ZDB to Disk+Tape

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

Backup objects specified in the backup specification are streamed from the replica to the backup medium. However, afterwards, the replica is not deleted, but kept on the array for instant recovery, so instant recovery information is written to the ZDB database during creation.

Split Mirror Restore

If a ZDB to tape or ZDB to disk+tape has been used to produce a backup to tape from a split mirror replica, it is possible to perform a split mirror restore. With this technique, a replica is modified before performing a restore. Backup objects are restored from tape to a split mirror replica (either previously existing or newly created specifically for the purpose) replacing some or all of the replica contents. The source volumes are then synchronized with the replica, effectively replacing the existing contents with those of the replica.

Other Uses for Replicas

There are many uses to which replicas can be put besides ZDB+IR, for example data mining. Data Protector can be used to create and administer replicas for such purposes, but it should be emphasized that replicas intended for instant recovery should never be used for any other purpose. If they are, accurate restore of data cannot be guaranteed.

Replica Deletion

When replicas are no longer required within Data Protector, if they are not members of a replica rotation set, they can be deleted from the array using the CLI.

If they are members of a replica rotation set, they are overwritten (or deleted so that a new one can be added) automatically when they, in turn, become the oldest in the set. Replicas can, however, be protected using an exclusion list on VA and EVA. Refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide* for information on exclusion lists.

If a replica is used for ZDB to tape, it is normally deleted automatically after completion of the backup session, unless you specifically ask in the backup specification to keep on the array.

5 ZDB to Tape

In This Chapter

This chapter explains how the various replication techniques are used within Data Protector for zero downtime backup to tape. After reading it you will know about:

- The principal steps involved in the ZDB-to-tape process.
- The types of replication configuration that can be used for ZDB to tape (local, remote, or remote plus local).
- The replication techniques usable for each replication configuration.
- Which array configurations can be used with Data Protector for the techniques involved.

ZDB-to-Tape Process Overview

With ZDB to tape, data replication is used to greatly reduce the impact on an application during backup operations, compared with conventional tape backup.

The principle steps in the ZDB-to-tape process are as follows:

1. Place the application or database into a stable state.
2. Create a replica containing the specified data objects.
3. Return the application/database to normal operation.
4. Move data from the copies of the specified data objects in the replica to the backup medium.
5. Perform post-backup processing on the replica.

Placing the Application or Database into a Stable State

With conventional Data Protector backup to tape, application operation is affected for the whole of the backup session, i.e., until the streaming of data to the backup medium is complete. However, with ZDB to tape, application operation is only affected during the time that a replica of the data to be backed up is created.

Online/Offline Backup

While a replica is being created, the operation of the section of the file system or database concerned must be frozen. This can be done in two ways:

- The simplest method is by performing an offline backup. In this case, the application or database is taken offline, i.e., all file I/O is stopped. With database applications, the database is usually placed into a consistent state, for instance by applying any previously unapplied transaction logs after stopping the I/O.
- Alternatively you can perform an online backup. This is normally only possible with database applications which can be placed into **hot-backup mode**. In this mode, all database I/O is diverted to log files for the duration of, in the case of ZDB, the replication part of the backup session. In this way a backup can be performed without the application having to be taken offline.

NOTE

When performing an online backup of a database application, it is normally also necessary to back up the archived transaction logs, to be able to perform a complete database recovery.

The steps concerned in these operations can be controlled automatically when backing up applications with application integrations supplied with Data Protector, but it is also possible to set up similar behavior when backing up other applications or file systems: pre- and post-exec options allow you to specify scripts to run before and after replication.

NOTE

Note that 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:

- 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.
- With ZDB to tape, a database is only put into hot-backup mode or taken offline for the time that a replica is created. Normal database operation can be resumed for the rest of the backup session while the data is being streamed to the backup medium.

Creating a Replica Containing the Specified Data Objects

The methods of replication available for ZDB to tape are dependent on a combination of things such as the type of disk array being used, whether local or remote replication is required, etc.

But, whichever method is used, there are some aspects of replica handling that are specific to ZDB to tape.

If you are only performing ZDB to tape (i.e., you do not want to mark replicas for instant recovery purposes, which enters information in the ZDB database) it is still possible to specify an option to keep a replica on the array after a backup and re-use it for further ZDB-to-tape sessions that use the same backup specification. In this way, you can guarantee that there is always enough space on the array for your backup, so that your backups will not fail.

In general, all replica types are suitable for ZDB to tape, though there may be some special recommendations in some cases.

ZDB to Tape Using Local Replication

Local replication techniques for ZDB to tape offer a very comprehensive level of support for backup purposes:

- They are available for all the array types supported by Data Protector.
- They are available for all application integrations supported by Data Protector.
- Replication and synchronization processes are all performed on the local array. This means that:
  - The processes are fast.
  - The disruption to the application or file system involved is minimized.
  - They are relatively easy to set up compared with other available ZDB techniques, such as remote plus local replication.

- They produce sessions from which individual backup objects can be restored.
- They offer a good level of data protection.
- They are suitable both for backup purposes and disaster recovery purposes, but, in the event of a disaster, restore of a complete session for a large database from tape would be very long for a high availability system.
- Split mirror restore (available for disk image, filesystem and filesystem based application backups) can be used effectively to perform a low impact restore for a system that is partially corrupted, but still operational.

Split Mirror Techniques

ZDB to tape using local replication with split mirror techniques is available on the following arrays:

- HP StorageWorks Disk Array XP using the Business Copy (BC) XP configuration.
- EMC Symmetrix Disk Array using the EMC Symmetrix TimeFinder configuration.

For details, refer to “Array Integrations Available with Data Protector” on page 44.

Snapshot Techniques

ZDB to tape using local replication with snapshot techniques is available on the following arrays:

- HP StorageWorks Virtual Array using the Business Copy (BC) VA configuration.
- HP StorageWorks Enterprise Virtual Array using the Business Copy (BC) EVA configuration.

For details, refer to “Array Integrations Available with Data Protector” on page 44.

ZDB to Tape Using Remote Replication

Remote replication techniques for ZDB to tape offer support for backup and disaster recovery purposes:

- They are available for arrays supported by Data Protector that use split mirror techniques.
- They are available for all application integrations supported by Data Protector.
- They can be used to create replicas on a separate remote array. Refer, for instance, to “HP StorageWorks Disk Array XP” on page 44 for information on possible connection distances between arrays.
- The backup has no impact on the application system, because it is performed on a remote array after the link between the source volumes on the local array and the target volumes on the remote system has been broken.

However, this configuration is only supported with Data Protector for synchronous connections, which can have a large impact on application performance when used over long distances. Most users would 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. Compare “ZDB to Tape Using Remote plus Local Replication” on page 77.

Split Mirror Techniques

ZDB to tape using remote replication with split mirror techniques is available on the following arrays:

- HP StorageWorks Disk Array XP, using the Continuous Access XP (CA XP) configuration.
- EMC Symmetrix Disk Array (EMC), using the Symmetrix Remote Data Facility (SRDF) configuration.

For details, refer to “Array Integrations Available with Data Protector” on page 44.

NOTE ZDB to tape using remote replication is not available with snapshot arrays.

ZDB to Tape Using Remote plus Local Replication

Remote plus local replication techniques for ZDB to tape offer support for backup of systems set up for disaster recovery purposes.

- They are available, in two forms:
  - A hardware solution is supported on split mirror arrays.
  - A software solution is supported with HP-UX on selected split mirror and snapshot arrays.
- They are available for all application integrations supported by Data Protector.
- They use at least two supporting arrays located in physically separate sites.
- They are typically used if the remote site functions as a disaster recovery site and a split of the remote pairs is not possible. In order to automate failover, an application such as MC/ServiceGuard or Microsoft Cluster Server can be used.

  Unlike ZDB to tape using remote replication, they allow backup operations, without impact on disaster recovery capability or any extra impact on application system performance.
- They produce sessions from which individual backup objects can be restored.

Hardware Solution

The hardware solution using remote plus local replication techniques is available on the following split mirror arrays:

- HP StorageWorks Disk Array XP using the Continuous Access (CA) XP plus Business Copy (BC) XP configuration.
- EMC Symmetrix Disk Array using the Symmetrix Remote Data Facility (SRDF) plus TimeFinder configuration.

For details, refer to “Array Integrations Available with Data Protector” on page 44.

NOTE ZDB to tape using remote plus local replication (hardware solution) is not available on snapshot arrays.

Software Solution

The software solution is only available with HP-UX. For the remote part of the process, it uses LVM mirroring, which can be configured to provide functionality similar to that of Continuous Access (CA) XP in a remote plus local replication environment on split mirror and snapshot arrays as follows:

- HP StorageWorks Disk Array XP using LVM mirroring plus the Business Copy (BC) XP configuration.
- HP StorageWorks Virtual Array using LVM mirroring plus the Business Copy (BC) VA configuration.

For details about LVM mirroring on XP, refer to “HP-UX LVM Mirroring” on page 48. For information about LVM mirroring on VA, refer to “Remote Plus Local Replication” on page 56.

Returning the Application/Database to Normal Operation

After the required replica for the backup has been produced, operation of the section of the file system or database concerned can be returned to normal.

In the case of an offline backup, the application or database can be put back online and normal operation started again.

In the case of an online backup, any archived transaction logs that you may need to apply to a restored database should also be backed up after the application/database is returned to normal operation.

Moving Data from the Replica to the Tape (Backup Medium)

Before data can be moved to tape from the copied data objects in a replica, the replica must first be mounted on the backup system.

Mount Point Creation

Once replication is complete, Data Protector creates mount points on the backup system and mounts filesystems in the replica to them. The required process for performing this depends on an application, disk image or file system backup is being performed. For more information, refer to “Additional Information” on page B-1.

System Locks

At this point, Data Protector also applies two forms of lock on the system, in case of possible concurrent processes:

- **Device lock.** For ZDB to tape, the tape device is not locked from the beginning of the backup session as with conventional backup to tape, but after the replica is created and before the start of data movement to tape. The lock is kept in place until data movement to tape has completed.
- **Logical resource lock** (also called disk lock). Whenever Data Protector accesses storage volumes for a backup (or restore) session, it prevents access to these volumes by any other Data Protector session until the current session’s requirement is complete, i.e., in this case, until all the required data has been moved to tape.

Standard Data Movement to Tape

Once the required locks are in place, the specified data objects can be moved to tape. This is normally done using standard Data Protector Media Agent functionality.

Data Protector writes the information to the tape as though the data objects were taken 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. Restore of data objects from ZDB-to-tape sessions can therefore be performed directly to the application system, using conventional restore procedures.

Direct Backup

With certain versions of HP-UX, it is possible to use Data Protector direct backup functionality to move data directly to a backup device within a SAN environment. For further information, refer to the *HP OpenView Storage Data Protector Concepts Guide* and the *HP OpenView Storage Data Protector Administrator's Guide*.

Performing Post-Backup Processing on the Replica

In a ZDB-to-tape session, after the data movement to tape has completed, the default behavior is for the replica to be deleted automatically. However, there are two other options available:

- `Keep the replica on the array`. In this case, the replica will not be deleted. No instant recovery information is written to the ZDB database, but the replica will be available for use other than instant recovery.
- `Use an existing replica`. If the option `Keep the replica on the array` has been selected during a previous backup, it is possible to re-use the replica for another ZDB to tape that uses the same backup specification. However, it is important that there is an existing replica on the array, or a backup with this option specified will fail.

6 Restore Techniques from ZDB-to-Tape Sessions

In This Chapter

This chapter gives an overview of the available techniques for restore from zero downtime backup-to-tape sessions:

- Standard Data Protector restore from tape.
- Split mirror restore.

Restore Process Overview

It is possible to perform a restore from any ZDB-to-tape sessions, using standard Data Protector restore from tape techniques, regardless of the types of array and application integrations are being used. In this case, it is possible to restore individual data objects directly to the application system. For further information, refer to the *HP OpenView Storage Data Protector Administrator's Guide*.

In addition, in some cases, if ZDB-to-tape sessions were produced using split mirror techniques, a split mirror restore technique is also available.

Split Mirror Restore

A split mirror restore is performed by restoring data from tape to the backup system and then re-synching this data with that available to the application system. It can be used to restore complete sessions or individual backup objects.

It can be used to restore data from file system or disk image ZDB-to-tape sessions produced under the following conditions:

- On an HP StorageWorks Disk Array XP using the Business Copy (BC) XP configuration.
- On an EMC Symmetrix Disk Array using the Symmetrix TimeFinder, SRDF, or combined (SRDF+TimeFinder) configurations.
- On an HP StorageWorks Disk Array XP using LVM mirroring, together with the Business Copy (BC) XP configuration (HP-UX only).

General Split Mirror Restore

An example of a split mirror restore on an HP StorageWorks Disk Array XP is shown below.

IMPORTANT Because of the different types of replicas involved and various array limitations, the detailed restore process is different for each array type. For further information refer to the “Limitations and Considerations” for each of the arrays in the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator’s Guide*.

Figure 6-1 Split Mirror Restore Example

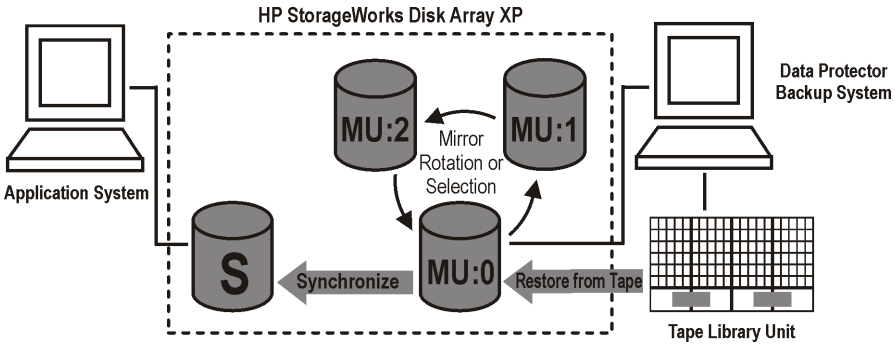

1. Select a mirror to use for the restore. If it is still being synchronized with the source volumes, split the link.
2. If the link is already split, optionally resynchronize the mirror with the source volumes and split the link again, to produce an up-to-date split mirror replica.
3. Restore the required 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.

The result is that, after synchronization, the contents of the selected mirror replace those of the source volumes as follows:

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

For further information on split mirror restore, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide*.

7 ZDB to Disk

In This Chapter

This chapter explains how the various replication techniques are used within Data Protector for zero downtime backup to disk. After reading it you will know about:

- The principal steps involved in the ZDB-to-disk process.
- The type of replication configuration that can be used for ZDB to disk (local).
- The replication techniques usable for the supported replication configuration.
- Which array configurations can be used with Data Protector for the techniques involved.

ZDB-to-Disk Process Overview

With ZDB to disk, data replication is used to greatly reduce the impact on an application during backup operations, compared with conventional tape backup. After replication, data is not streamed to tape, but the replicas are kept on the array and can be used for restore using instant recovery techniques.

The principle steps in the ZDB-to-disk process are as follows:

1. Place the application or database into a defined state.
2. Create a replica containing the specified data objects.
3. Return the application/database to normal operation.
4. Perform post-backup processing on the replica.
5. Record session and instant recovery information in the IDB.

Placing the Application or Database into a Defined State

With conventional Data Protector backup to tape, application operation is affected for the whole of the backup session, until the streaming of data to a backup medium is complete. However, with ZDB to disk, as with ZDB to tape, application operation is only affected during the time that a replica of the data to be backed up is created.

The main differences with ZDB to disk are that, after the replica has been produced:

- No data is streamed from the replica to tape. Instead, the replica is kept on the array for future use.
- Information about the replica is written to the ZDB database.

However, the preparation for replication is the same as that for ZDB to tape.

Online/Offline Backup

While a replica is being created, the operation of the section of the file system or database concerned must be frozen. This can be done in two ways:

- The simplest method is by performing an offline backup. In this case, the application or database is taken offline, i.e., all file I/O is stopped. With database applications, the database is usually placed into a consistent state, for instance by applying any previously unapplied transaction logs after stopping the I/O.
- Alternatively you can perform an online backup. This is normally only possible with database applications, which can be placed into **hot-backup mode**. In this mode, all database I/O is diverted to the transaction log files for the duration of, in the case of ZDB, the replication part of the backup session. In this way a backup can be performed without the application having to be taken offline.

The steps concerned in these operations can be controlled automatically when backing up applications with Data Protector integrations, but it is also possible to set up similar behavior when backing up other applications or file systems: pre- and post-exec options allow you to specify scripts to run before and after replication.

NOTE

When performing an online backup of a database application, it is normally also necessary to back up the archived transaction logs, to be able to perform a complete database recovery. The logs must not be backed up as part of the ZDB-to-disk replica. They can be backed up by scheduling a separate conventional Data Protector backup to tape. This does not have to be done in parallel with the ZDB-to-disk session. It can be done after the ZDB-to-disk session has finished.

The post-exec option can be used within a ZDB-to-disk session to automatically start a standard Data Protector backup to tape for the archived transaction logs.

Creating a Replica Containing the Specified Data Objects

The methods of replication available for ZDB to disk are dependent on a combination of things such as the type of disk array being used, whether local or remote replication is required, etc.

But, whichever method is used, there are some aspects of replica handling that are specific to ZDB to disk.

The main requirement when creating a backup specification for ZDB to disk is that an option `Track the replica for instant recovery` is specified. This has three major effects:

- It ensures that array specific information about the replica(s) produced is stored in the ZDB database. This information is mapped to the session information stored in the IDB and is essential if you want to be able to use the session for instant recovery purposes. If this option is not selected, no instant recovery specific information about the replica is stored in the ZDB database and replicas produced using the backup specification cannot be used for instant recovery.
- It allows a choice to be made between ZDB to disk and ZDB to disk+tape, when scheduling or starting a ZDB session interactively. No ZDB to tape option is available.
- In the case of the supported snapshot arrays, it allows a replica set to be created that can be used for replica set rotation. This is not supported for the EMC Symmetrix Disk Array.

  The number of replicas in the set to be rotated can be specified, from 1 upwards. The maximum possible number in a replica set is dependent on the type of array being used. See “Array Integrations Available with Data Protector” on page 44 for more information.

In general, all replica types are suitable for ZDB to disk, though the snapshot sub-types supported vary according to the array. For further information on what is supported for the various disk arrays, refer to “Array Integrations Available with Data Protector” on page 44.

ZDB to Disk Using Local Replication

Only local replication techniques are supported for ZDB to disk. They offer a very comprehensive level of support for backup purposes:

- They are available for all the array types supported by Data Protector.
- They are available for all application integrations supported by Data Protector.
- Replication and synchronization processes are all performed on the local array. This means that:
  - The processes are fast.
  - They are relatively easy to set up compared with other available ZDB techniques, such as remote plus local replication.
  - The disruption to the application or file system involved is minimized.
- Replica set rotation functionality is available.
- They produce sessions from which complete replicas can be restored using instant recovery functionality. This functionality allows the states of all the volumes containing the backup objects specified in a ZDB specification to be returned to those at a specific point in time at high speed.
- They are suitable for backup purposes where there is no requirement to restore individual backup objects.

Split Mirror Techniques

ZDB to disk using local replication with split mirror techniques is available on the following arrays:

- HP StorageWorks Disk Array XP using the Business Copy (BC) XP configuration.
- EMC Symmetrix Disk Array using the EMC Symmetrix TimeFinder configuration.

For details, refer to “Array Integrations Available with Data Protector” on page 44.

Snapshot Techniques

ZDB to disk using local replication with snapshot techniques is available on the following arrays:

- HP StorageWorks Virtual Array using the Business Copy (BC) VA configuration.
- HP StorageWorks Enterprise Virtual Array using the Business Copy (BC) configuration.

NOTE Of the snapshot types available on this array, only snapclone is supported for instant recovery purposes.

For details, refer to “Array Integrations Available with Data Protector” on page 44.

Returning the Application/Database to Normal Operation

After the required replica for the backup has been produced, the operation of the section of the file system or database concerned can be returned to normal.

In the case of an offline backup, the application or database can be put back online and normal operation started again.

In the case of an online backup, any archived transaction logs that you may need to apply to a restored database should also be backed up after the application/database is returned to normal operation. This can be done by scheduling a standard Data Protector backup to tape.

Transaction logs up must not be backed up as part of the ZDB-to-disk replica.

Performing Post-Backup Processing on the Replica

With a ZDB-to-disk session, after the replica has been produced and the array specific information about it recorded in the ZDB database, it is kept on the array for instant recovery purposes, either alone, or as part of a replica set.

If replica set rotation is being used for the backup specification concerned, the replica remains on the array until:

- It is the oldest in the set.
- If older replicas exist, but they are on the exclude list for HP StorageWorks Disk Array XP or HP StorageWorks Virtual Array.

Then, when the next ZDB-to-disk session is performed using the same backup specification, the existing replica is replaced (is either overwritten, or is deleted and a new one added to the set in its place, depending on the array type). The existing entry in the ZDB database is then deleted and a new one added.

The same thing happens if the replica set size is 1, i.e., there is effectively no replica rotation.

Recording Session and Instant Recovery Information in the IDB

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

However, to be able to perform a ZDB to disk, the option `Track the replica for instant recovery` must be selected in the backup options. This means that, after a replica has been produced, array-specific information for instant recovery purposes is also written to the ZDB database. After a ZDB-to-disk session, the associated restore objects and restore sessions can be viewed in the `Instant Recovery` context.

For more information about ZDB database, refer to “ZDB Database” on page B-7.

8 Restore Techniques from ZDB-to-Disk Sessions

In This Chapter

This chapter gives an overview of the available techniques for restore from zero downtime backup-to-disk sessions. After reading it you will understand the basics of:

- The principal steps involved in restore from ZDB-to-disk session.
- The general instant recovery process.
- Database recovery.
- Instant recovery in special situations.

Restore Process Overview

Instant Recovery

When ZDB-to-disk sessions are performed, the replicas produced can only be restored using instant recovery techniques.

During a ZDB-to-disk session, information is saved in the IDB and the related array specific information required for instant recovery in the ZDB database.

The replicas produced during these sessions cannot be displayed or selected directly in the Data Protector GUI, but the sessions available for restore using instant recovery can. This can be done using a special `Instant Recovery` context available in the GUI. Alternatively, the Data Protector CLI can be used.

During a restore using instant recovery, the data in the original source volumes is replaced by that in the replica target volumes internally within the array, involving no other backup medium or device. This makes the restore very fast.

It is important to note, with instant recovery, that even though several individual backup objects may have been selected in the original ZDB specification, it is 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. The implications of this are that 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 their states at the time the replica was created.

Additional Steps

If restoring a file system or a disk image, the instant recovery session may be all that is required. If recovering a database application, additional steps, such as the subsequent restore and application of transaction logs, may be required to make the database fully operational again. This assumes that backups of the transaction logs exist for points in time *after* the creation time of the replica restored using instant recovery. This normally involves the use of another backup medium or device. For information on database recovery, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Integration Guide*.

General Instant Recovery Process

IMPORTANT

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

An example of restore using instant recovery is shown below.

Figure 8-1 Instant Recovery Example

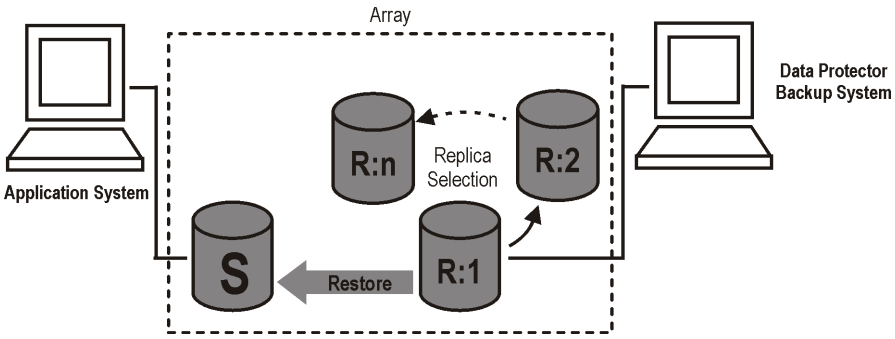

1. Decide which replica you want to restore and select the ZDB session that was used to create it.

   In a dedicated `Instant Recovery` context in the Data Protector GUI, you can display ZDB sessions for restore in two ways:

   - As `Restore Sessions`. All restore sessions listed by date and time.
   - As `Restore Objects`. Restore sessions are listed in separate groups according to the type of ZDB session performed, i.e., `Filesystem`, `Disk Image`, `SAP R/3`, `MS SQL Server`, etc.

   Note that, within a session, it is only possible to display:

   - Complete volumes or mount points for filesystem ZDB sessions.
   - "Database" for application integrations.

IMPORTANT Any filesystems that were not part of the ZDB-to-disk specification, but reside on the same *volume group* as the backed up objects are not displayed in the GUI, but these will also be restored as a result of the instant recovery process.

2. Select from a variety of instant recovery options. These are provided primarily for data security purposes.

   These allow you to:

   - Check that the configurations of the volume groups involved in the instant recovery have not changed since the replica to be restored was created.

     On the HP StorageWorks Virtual Array, this check also verifies that the results of a CRC check performed on the data in the replica to be restored match those produced when the replica was created.

   - Keep the replica on the array after it has been restored. It is advisable to do this, so that the replica is still available if there are any problems with any recovery step after the restore.

     This is not available with the HP StorageWorks Enterprise Virtual Array, because of the way the restore process is performed on that array.

   - On the HP StorageWorks Enterprise Virtual Array, remove the presentation to any hosts of any volumes making up a snapclone replica that is to be restored. If any such volumes are left presented to any hosts, a restore cannot be performed using the snapclone concerned.

3. Optionally, perform a preview of the restore of the instant recovery session to provide an extra level of security.
4. Start the restore.

Data Protector then:

- Connects itself to the application and backup systems.
- Extracts the session information from the IDB and the array-specific information associated with the session from the ZDB database.

- Performs the necessary checks to verify that all the required conditions for a successful restore are met (including any instant recovery options specified).
- Deactivates any volume groups and dismounts any filesystems associated with the replica.
- Restores the replica to the original source volumes.
  - On the HP StorageWorks Disk Array XP, this is done by synchronizing the source volumes with those of the selected split mirror replica.
  - On the HP StorageWorks Virtual Array, all other replicas in the replica set created by the associated ZDB specification are first deleted and their entries removed from the ZDB database.
  - On the HP StorageWorks Enterprise Virtual Array, the selected snapclone replica is substituted for the original source volumes which are then deleted, together with their ZDB database entries. Any host presentations that were made to the original source volumes are then made to the restored snapclone volumes which then effectively become the new source volumes. As far as Data Protector is concerned, the snapclone replica is deleted from the associated replica set and the array.
- Re-enables any volume groups that it disabled and re-mounts any filesystems that it dismounted.

IMPORTANT

If you are restoring a disk image using instant recovery, it is important to realize that, for such a backup, no configuration information for the disk will be available. Therefore to avoid the possibility of data corruption:

- If the backup objects involve volume groups, these must be manually deactivated before the instant recovery process and re-activated after it.
- If the backup objects involve logical volumes as raw disk, these must be manually dismounted before the instant recovery session and re-mounted afterwards, if necessary.
- If the backup objects involve filesystem objects, the filesystem must be manually dismounted before the instant recovery process and remounted afterwards.

The result of the restore is that the contents of the source volumes are returned to their states at the time the replica was created.

After Instant Recovery

Following the restore using instant recovery, if you are restoring a database application, you must normally perform further database or application specific steps, such as restoring and applying transaction logs to make the database consistent and fully operational.

For further information on restore using instant recovery, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide*.

Instant Recovery and LVM Mirroring

Restore using instant recovery is supported for ZDB sessions produced on HP-UX systems with an LVM mirroring plus BC XP configuration on the HP StorageWorks Disk Array XP only. However, this requires specialist steps to be performed, in addition to those given above. For further information, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide*.

Instant Recovery in a Cluster

Restore using instant recovery is supported for an application or a filesystem running in an MC/ServiceGuard or Microsoft Cluster Server on the application system. However, it is necessary to perform some extra specialist steps. For further information, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide*.

9 ZDB to Disk+Tape

In This Chapter

This chapter explains how the various replication techniques are used within Data Protector for zero downtime backup to disk+tape. After reading it you will know about:

- The principal steps involved in the ZDB-to-disk+tape process.
- The type of replication configuration that can be used for ZDB to disk+tape (local).
- The replication techniques usable for the supported replication configuration.
- Which array configurations can be used with Data Protector for the techniques involved.

ZDB-to-Disk+Tape Process Overview

With ZDB to disk+tape, data replication is used to greatly reduce the impact on an application during backup operations, compared with conventional tape backup.

The principle steps in the ZDB-to-disk+tape process are as follows:

1. Place the application or database into a defined state.
2. Create a replica containing the specified data objects.
3. Return the application/database to normal operation.
4. Move data from the copies of the specified data objects in the replica to the backup medium.
5. Perform post-backup processing on the replica.
6. Record session and instant recovery information in the IDB.

Placing the Application or Database into a Defined State

With conventional Data Protector backup to tape, application operation is affected for the whole of the backup session, until the streaming of data to a backup medium is complete. However, with ZDB to disk+tape, as with ZDB to tape, application operation is only affected during the time that a replica of the data to be backed up is created.

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

With ZDB to disk+tape, after the replica has been produced:

- Information about the replica is written to the ZDB database, for instant recovery purposes.
- Data is streamed from the replica to tape, so that it is possible to perform a conventional Data Protector restore.
- After the streaming to tape has completed, the replica is kept on the array for future use, so that it is possible to restore using instant recovery functionality.

Online/Offline Backup

While a replica is being created, the operation of the section of the file system or database concerned must be frozen. This can be done in two ways:

- The simplest method is by performing an offline backup. In this case, the application or database is taken offline, i.e., all file I/O is stopped. With database applications, the database is usually placed into a consistent state, for instance by applying any previously unapplied redo logs after stopping the I/O.
- Alternatively you can perform an online backup. This is normally only possible with database applications which can be placed into **hot-backup mode**. In this mode, all database I/O is diverted to log files for the duration of, in the case of ZDB, the replication part of the backup session. In this way a backup can be performed without the application having to be taken offline.

NOTE When performing an online backup of a database application, it is normally also necessary to back up the archived redo logs, to be able to perform a complete database recovery.

The steps concerned in these operations can be controlled automatically when backing up applications with application integrations supplied with Data Protector, but it is also possible to set up similar behavior when backing up other applications or file systems: user exits are supplied, which allow you to specify scripts to run before and after replication.

NOTE Note that 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:

- 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.
- With ZDB to disk+tape, a database is only put into hot-backup mode or taken offline for the time that a replica is created. Normal database operation can be resumed for the rest of the backup session while data is being streamed to the backup medium.

Creating a Replica Containing the Specified Data Objects

ZDB to disk+tape is not merely a combination of ZDB to disk and ZDB to tape, as might be expected, because the replication methods supported for these two backup types do not match.

From the functionality point of view, ZDB to disk+tape is in fact ZDB to disk with the added capability to stream data from the replica to tape, or other backup medium, after replication. So, the replication method/array support is the same as for ZDB to disk, which is much more limited than that for ZDB to tape.

It is possible specify ZDB-to-disk+tape sessions in the same schedule as ZDB-to-disk sessions, using the same backup specification. This allows you to 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 allows greater flexibility for restore. Note that the same replica set will be used for both types of session.

When creating a backup specification for ZDB to disk+tape, the option `Track the replica for instant recovery` must be selected as with ZDB to disk. Thus:

- Array-specific information about the replica(s) produced is stored in the ZDB database and mapped to the session information stored in the IDB. This allows the replica produced in the session to be used for instant recovery purposes. If this option is not selected, no information about the replica is stored in the ZDB database and replicas produced using the backup specification cannot be used for instant recovery.
- All of supported arrays, except the EMC Symmetrix Disk Array, allow creation of replica sets that can be used for replica set rotation.

  The number of replicas in the set can be specified, from 1 upwards. The maximum number is dependent on the type of array being used. See “Array Integrations Available with Data Protector” on page 44 for more information.

In general, all replica types are suitable for ZDB to disk+tape, though the snapshot sub-types supported vary according to the array. See “Array Integrations Available with Data Protector” on page 44 for more information.

ZDB to Disk+Tape Using Local Replication

Local replication techniques for ZDB to disk+tape offer a very comprehensive level of support for backup purposes:

- They are available for all the array types supported by Data Protector.
- They are available for all application integrations supported by Data Protector.
- Replication and synchronization processes are all performed on the local array. This means that:
  - The processes are fast.
  - They are relatively easy to set up compared with other available ZDB techniques, such as remote plus local replication.
  - The disruption to the application or file system involved is minimized.
- Replica set rotation functionality is available, even for tape.
- Data backed up in a single session can be restored in different ways as follows:
  - A complete replica can be restored, at high speed, using instant recovery functionality. In this case, all the volumes containing the data objects specified in a ZDB specification are returned to their states at a specific point in time.
  - Individual backup objects can be restored to the same point in time, using normal Data Protector restore from tape.
  - If using a split mirror disk array, using split mirror restore, individual backup objects can first be restored from tape to update a replica, which can then be restored.

Split Mirror Techniques

ZDB to disk+tape using local replication with split mirror techniques is available on the following arrays:

- HP StorageWorks Disk Array XP using the Business Copy (BC) XP configuration.
- EMC Symmetrix Disk Array using the EMC Symmetrix TimeFinder configuration.

For details, refer to “Array Integrations Available with Data Protector” on page 44.

Snapshot Techniques

ZDB to disk+tape using local replication with snapshot techniques is available on the following arrays:

- HP StorageWorks Virtual Array using the Business Copy (BC) VA configuration.
- HP StorageWorks Enterprise Virtual Array using the Business Copy (BC) EVA configuration.

NOTE Of the snapshot types available with the EVA, only snapclone is supported for instant recovery purposes.

For details, refer to “Array Integrations Available with Data Protector” on page 44.

Returning the Application/Database to Normal Operation

After the required replica for the backup has been produced, the operation of the section of the file system or database concerned can be returned to normal.

In the case of an offline backup, the application or database can be put back online and normal operation started again.

In the case of an online backup, any archived transaction logs that you may need to apply to a restored database should also be backed up after the application/database is returned to normal operation.

Moving Data from the Replica to the Tape (Backup Medium)

Before data can be moved to tape from the copied data objects in a replica, the replica must first be mounted on the backup system.

Mount Point Creation

Once replication is complete, Data Protector creates mount points on the backup system and mounts filesystems in the replica to them. The required process for performing this depends on an application, disk image or file system backup being performed. For more information, refer to “Additional Information” on page B-1.

System Locks

At this point, Data Protector also applies two forms of lock on the system, in case of possible concurrent processes:

- **Device lock**. For ZDB to disk+tape, the tape device is not locked from the beginning of the backup session as with conventional backup to tape, but after the replica is created and before the start of data movement to tape. The lock is kept in place until data movement to tape has completed.
- **Logical resource lock** (also called disk lock). Whenever Data Protector accesses logical volumes for a backup (or restore) session, it prevents access to these volumes by any other Data Protector session until the current session’s requirement is complete, i.e., in this case, until all the required data has been moved to tape.

Standard Data Movement to Tape

Once the required locks are in place, the specified data objects can be moved to tape. This is normally done using standard Data Protector media agent functionality.

Data Protector writes the information to the tape as though the data objects were taken 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. Restore of data objects backed up to tape during ZDB-to-disk+tape sessions can therefore be performed directly to the application system, using conventional restore procedures.

Direct Backup

With certain versions of HP-UX, it is possible to use Data Protector direct backup functionality to move data directly to a backup device within a SAN environment. For further information, refer to the *HP OpenView Storage Data Protector Concepts Guide* and the *HP OpenView Storage Data Protector Administrator's Guide*.

Performing Post-Backup Processing on the Replica

With a ZDB-to-disk+tape session, after the replica has been produced and the array specific information about it recorded in the ZDB database, it is kept on the array for instant recovery purposes, either alone, or as part of a replica set. This is the same behavior as ZDB to disk, but markedly different to ZDB to tape.

If replica set rotation is being used for the backup specification concerned, the replica remains on the array until:

- It is the oldest in the set.
- If older replicas exist, but they are on the exclude list for HP StorageWorks Disk Array XP or HP StorageWorks Virtual Array.

Then, when the next ZDB-to-disk+tape session is performed using the same backup specification, the existing replica is replaced (is either overwritten, or is deleted and a new one added to the set in its place, depending on the array type). The existing entry in the ZDB database is then deleted and a new one added.

The same thing happens if the replica set size is 1, i.e., there is effectively no replica rotation.

Recording Session and Instant Recovery Information in the IDB

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

However, as with ZDB to disk, to be able to perform a ZDB to disk+tape, the option `Track the replica for instant recovery` must be selected in the backup options. This means that, after a replica has been produced, array specific information for instant recovery purposes is also written to the ZDB database.

After a ZDB-to-disk+tape session, the associated restore objects and restore sessions can be viewed in two places in the GUI:

- In the `Restore` context, allowing restore of data objects from tape.
- In the `Instant Recovery` context, allowing restore from replicas.

For more information about ZDB database, refer to “ZDB Database” on page B-7.

10 Restore Techniques from ZDB-to-Disk+Tape Sessions

In This Chapter

This chapter gives an overview of the available techniques for restore from zero downtime backup-to-disk+tape sessions:

- Restore directly from tape to the application system.
- Instant recovery.
- Split mirror restore.

Restore Process Overview

Restore Techniques

ZDB-to-disk+tape sessions provide more restore flexibility than either ZDB-to-disk or ZDB-to-tape as the replicas produced are kept on the array and are also streamed to tape.

The restore methods available depend upon the type of the array being used.

Snapshot Arrays

The HP StorageWorks Virtual Array and the HP StorageWorks Enterprise Virtual Array both support:

- Restore using instant recovery, allowing the restore of the total contents of a replica, with minimum impact on the application system concerned.
- Restore directly from tape to the application system, allowing the restore of any of the individual backup objects available from the tape session.

NOTE

What is available for restore depends on what was actually streamed to tape. This, in turn, is dependent upon how the ZDB-to-disk+tape specification was set up. If the complete contents of the source volumes were specified by selecting at the highest level, all objects will have been streamed to tape. If not, only the selected backup objects will have been streamed to tape, even though the whole of the source volumes will have been replicated.

Split Mirror Arrays

The HP StorageWorks Disk Array XP supports:

- Restore using instant recovery, allowing the restore of the total contents of a replica, with minimum impact on the application system concerned.

- Split mirror restore, allowing (potentially) restore of anything from an individual backup object to the whole contents of the replica, with minimum impact on the application system.

NOTE What is available for restore depends on what was actually streamed to tape. This, in turn, is dependent upon how the ZDB-to-disk+tape specification was set up. If the complete contents of the source volumes were specified by selecting at the highest level, all objects will have been streamed to tape. If not, only the selected backup objects will have been streamed to tape, even though the whole of the source volumes will have been replicated.

- Restore directly from tape to the application system, allowing the restore of any of the individual backup objects available from the tape session.

NOTE With the Data Protector EMC Symmetrix Disk Array integration, ZDB-to-disk+tape sessions cannot be performed.

For further information on restore using instant recovery, refer to Chapter 8, “Restore Techniques from ZDB-to-Disk Sessions,” on page 101.

For further information on split mirror restore, refer to Chapter 6, “Restore Techniques from ZDB-to-Tape Sessions,” on page 83.

For further information on restoring directly from tape the application system, refer to the *HP OpenView Storage Data Protector Concepts Guide* and the *HP OpenView Storage Data Protector Administrator’s Guide*.

11 Important Considerations

In This Chapter

This chapter discusses the following important considerations you need to take into account when planning your ZDB strategy:

- Disk array specific considerations.
- ZDB performance tuning.
- Instant recovery strategy.
- Security aspects.

It also gives you recommendations that will help you design your backup solution and improve your ZDB performance.

Optimizing ZDB Performance

In business-critical environments, it is a key requirement to minimize the time needed for data recovery in case of a corrupt database or a disk crash. ZDB solutions enable you to create a copy of data almost instantaneously; however, there are some important things to consider when planning ZDB performance.

General ZDB strategy planning is a process that includes the following steps:

1. Defining the requirements and constraints for backups, for example, how often your data needs to be backed up or whether you need additional copies of the backed up data on additional media sets.
2. Understanding the factors that affect disk array performance.
3. Preparing the backup strategy that supports your backup concept and how it is implemented.

This section provides you with some important information and considerations that help you plan your backup solution and improve your ZDB performance. The recommendations are given for every type of disk arrays supported by Data Protector.

Split Mirror Disk Arrays Considerations

The Data Protector HP StorageWorks Disk Array XP and EMC Symmetrix integrations provide a set of options enabling you to define your backup policy (moving the mirror copy of the original data to media, leaving the disks split or resynchronizing them, preparing the next disk for the next backup, and so on). Refer to *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide* for the examples of the backup policies. Below are some general recommendations on the split mirror arrays performance.

The interaction between Data Protector and applications during a ZDB session is complex and there are some important considerations to be taken into account.

Parallel Session Considerations

When running parallel ZDB sessions that back up one or more of the backup objects residing on the same source volumes, there must be no overlap during the phase of the backup process before the start of data streaming to tape. As a consequence of this, ZDB-to-disk sessions that back up the same objects must not be run in parallel.

Backup Specification Considerations

- With filesystem and disk image backups, parallel ZDB-to-tape or ZDB-to-disk+tape sessions can be run using the same backup specification, but they must not overlap during the phase of the backup process before the start of data streaming to tape.
- With application integrations, parallel ZDB sessions must *not* be run using the same backup specification.
- Different backup specifications that back up one or more of the backup objects residing on the same source volumes can be used to run parallel ZDB-to-tape or ZDB-to-disk+tape backup sessions, but there must be no overlap during the phase of the backup process before the start of data streaming to tape.
- ZDB and instant recovery sessions that involve the backup objects residing on the same source volumes cannot be run concurrently. In such cases, the sessions must be run sequentially.

If you consider running parallel backup sessions, it is advisable to measure the time taken to reach the start of the data copy to tape in each case and add on a large buffer to safeguard against overlap when scheduling your backup sessions (according to the conditions above).

Snapshot Disk Arrays Considerations

If you use the Data Protector HP StorageWorks Virtual Array or HP StorageWorks Enterprise Virtual Array integration, you need to consider the following factors when planning your backup strategy:

- Type of snapshot. Refer to “Snapshot Types” on page 132.
- Snapshot policy. Refer to “Snapshot Policy” on page 135.
- Some snapshot disk arrays specific considerations. Refer to “Other Considerations” on page 136.

Snapshot Types

Depending on the type of snapshot replica you have selected to create, you need to be aware of some factors that may have an impact on your ZDB performance. Refer to the below sections for more details.

Pre-Allocated (Standard) Snapshot

When this type of snapshot is used, the same amount of disk space inside the array as taken by the source volumes is pre-allocated for the target volumes.

For more information about pre-allocated snapshots, refer to Chapter 2, “Replication Techniques,” on page 13.

Performance

When a snapshot is accessed by a backup system, it is accessed by reading the disk blocks from the source volumes and the replica. Therefore, both the application and the backup systems disk resources are used, which can result in the application performance degradation.

Recommended Usage for VA

On VA, pre-allocated snapshots are the only supported type of snapshot, so they have to be used for all ZDB types. Pre-allocation of storage protects such snapshots from running out of storage on a disk array. It is important to know that if the data in the source volume is lost or corrupted, the data in the associated replica is useless and cannot be restored using instant recovery.

Recommended Usage for EVA

On EVA, pre-allocated snapshots are intended to be short-lived. They can be used for ZDB-to-tape sessions only, because instant recovery from such snapshots is not supported.

Virtually Capacity-Free Snapshot (VSNAP)

This type of snapshot is supported on EVA only. Its creation and maintenance is very similar to that for the pre-allocated snapshot; however, no storage pre-allocation on the array is required for the target volumes. This makes VSNAPs very space-efficient.

For more information about VSNAPs, refer to Chapter 2, “Replication Techniques,” on page 13.

Performance

When a VSNAP is accessed by a backup system, it is accessed by reading the disk blocks from the source volumes and the replica. Therefore, both the application and the backup systems disk resources are used, which can result in the application performance degradation.

Recommended Usage

VSNAPs are intended to be short-lived. Since the storage requirement for VSNAPs is dynamic, such snapshots may run out of storage on a disk array, should there be 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.

VSNAPs can be used for ZDB-to-tape sessions only, because instant recovery from such snapshots is not supported.

Snapclone

This type of snapshot is supported on EVA only. Snapclone is a full copy of one or more source volumes, completely independent of the original. The first part of its creation is similar to the creation of a pre-allocated snapshot, which is then followed by the cloning process. During this process, all data from the source volume is copied.

For more information about snapclones, refer to Chapter 2, “Replication Techniques,” on page 13.

Performance

If a snapclone is accessed by a system before the cloning process is finished, the disk blocks that are not copied yet are read from the source volume rather than the snapclone. In the case of ZDB to tape or ZDB to disk+tape, the data is read by using both the application and the backup systems disk resources, which can result in application performance degradation. Therefore, Data Protector by default delays copying the

snapclone data to tape by at most 90 minutes if the cloning process is still in progress. You can change the default behavior in the Data Protector GUI when configuring a backup specification.

Recommended Usage

Snapclones are intended to be long-lived. They are the only type of replica that can be used for instant recovery.

When all the data is copied, the snapclone is no longer dependent on its original virtual disk and behaves as a regular standalone virtual disk. If the data in the source volumes is lost or corrupted, the data in the replica can be used for instant recovery.

Allocating Snapclones

When configuring the EVA virtual disks, you define in which disk group source volumes should reside. By default, the same disk group as used for the source volumes is used for their snapclone allocation. Thus, the source volumes and the snapclone reside in the same disk group and therefore on the same set of physical disks.

However, you can allocate snapclones to a different disk group than the one used for the source volumes by editing the **EVA disk group pairs configuration file**. For the detailed instructions on how to change the snapclone allocation, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide*.

The benefits of using different disk groups for the source and the target volumes are the following:

- Read and write operations on replicas have less impact on the application performance, since different physical disks are used for the source volumes and the replica.
- It is less likely that a hardware failure on the source volumes will also destroy the replica.
- You can force replica allocation on low-performance physical disks while keeping fast physical disks available for the application data.

TIP

It is recommended you use low-performance disks for ZDB to tape only. During the instant recovery process on EVA, snapclones become source data after the instant recovery, so the physical location of your source data changes. If you use low-performance disks for ZDB to disk or ZDB to

disk+tape, your application starts to run on these low-performance disks after the instant recovery, which can cause the application performance degradation.

When defining disk group pairs, note the following:

- A disk group for a certain EVA can be a member of only one disk group pair specified in the disk group configuration file.
- If a disk group specified in the disk group configuration file does not exist or the syntax of the configuration file is inaccurate, snapclones are created in the same disk group where their source volumes reside.

Snapshot Policy

The Data Protector EVA integration introduces the concept of snapshot policy. With this policy, only one type of snapshot replica can be created for the same source volumes. For example, if you want to create a VSNAP for the source volumes for which a pre-allocated snapshot already exists, you need to delete the pre-allocated snapshot to be able to create a VSNAP successfully.

NOTE The above-stated limitation is not applicable if you want to create either a pre-allocated snapshot or a VSNAP for the source volumes, for which snapclones already exist.

When creating a Data Protector EVA backup specification, you can choose between strict and loose policy. Note that if a replica was created using the loose policy, instant recovery from such a replica is not possible.

- **Strict policy**

  If the strict policy is selected, Data Protector creates snapshots of the type you have selected (standard, VSNAP, or snapclone). However, if some of the source volumes used in the backup session already have existing snapshots of a different type, the backup session fails.

- **Loose policy**

  If the loose policy is selected, Data Protector may create snapshots of a different type than selected by the user. For example, if you want to create snapclones, but VSNAPs or pre-allocated snapshots of these source volumes already exist, either VSNAPs or pre-allocated snapshots are created instead.

For more information about the snapshot policy, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide*.

Other Considerations

The Data Protector ZDB solutions are implemented using the snapshot technology on various disk arrays. Each of these disk arrays has certain limitations and considerations you need to take into account when planning your backup strategy. Refer to below sections for more information.

VA Considerations

- If you use the same source volume in more than one backup specification, you will not be able to perform instant recovery from the replicas created in such environment. Therefore, to enable instant recovery, you need to manually delete all the replicas containing this source volume (the replicas created using the backup specification, different to the one you used for the replica from which you are restoring).

  Note that instant recovery will also fail if the replicas, not created by Data Protector but containing the same source volume as used in the replica created by Data Protector, exist on the array. In this case, you need to delete such replicas before you start instant recovery.

- Before an instant recovery session, ensure that no process accesses any of the volumes to which the data will be restored. Data Protector is not able to block access to VA by other applications. Therefore, such access during the restore could result in data corruption.

EVA Considerations

- **Replica creation**

  — If the cloning process of the source volume is in progress, another snapclone for the same volume cannot be created until the cloning process finishes. In this case, Data Protector retries the operation several times and, if the cloning process is still in progress, the session fails. By default, there are 10 retries with 10 seconds waiting between each retry. These default values can be changed by setting the `EVA_EMAPI_MAX_RETRY` and `EVA_EMAPI_RETRY_DELAY omnirc` variables. For more information about these variables, refer to *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide*.

- **Replica Set**

  — A replica cannot be reused in the following cases:

  One of the volumes in the snapclone has a snapshot attached to it, or:

  One of the target volumes to be reused is presented to a system.

  In such situations, the EVA agent aborts the session.

  — If a replica to be reused is in use and therefore locked by another session, Data Protector reuses the second oldest replica.

- **Backup**

  — In a situation when you have already reached the specified number of replicas rotated, but you then set a smaller number in the backup specification, the oldest replicas in the replica set are deleted during the next backup session. For example, you already have 7 replicas created with the backup specification in which the rotation number is set to 7. If you now change the backup specification and set this number to 5, Data Protector will delete three oldest replicas from the replica rotation set before creating a new replica.

  — If you modify a backup specification with the backup option `Keep the replica after the backup` selected, so that you deselect this option, all the replicas already created with this backup

specification are deleted during the next backup session. Thus, when the backup session is finished, there are no replicas belonging to this backup specification on a disk array.

- **Instant Recovery**

  — Before an instant recovery session, Data Protector checks that no snapclone from the selected replica is presented to a host. If a snapclone is presented, instant recovery fails. However, if you select the `Force the removal of all replica presentations` option, Data Protector removes host presentations before the actual instant recovery, and the session does not fail.

CAUTION If there are systems using the disks that have been removed, they will not be able to access these disks any longer. This may result in the system crash.

The source volumes that are to be replaced after instant recovery *should not be presented* to any non-application system (except clustered nodes of the application system). If a target virtual disk is presented to a system and you perform instant recovery, Data Protector does an unpresent operation and, after instant recovery, a present operation of the disk. Therefore, after instant recovery, these systems may not function properly.

— To perform a ZDB session marked for instant recovery, snapclones must be selected as the snapshot type. Since you can also define in which disk group the snapclones will be created, it is important to consider the following:

After instant recovery, the replica becomes the source, therefore the physical location of the source volumes changes. Thus, if a snapclone to be restored belongs to a different disk group than its source volumes, the disk group of the snapclone becomes the disk group of the source volumes.

For example, if a disk group used for a snapclone replica consists of disks with lower performance than the source volumes, application performance degrades after instant recovery.

Therefore, for ZDB-to-disk and ZDB-to-disk+tape sessions, it is recommended to create a snapclone in the same disk group as the source volumes (the default behavior). Otherwise, make sure that

the disk group to be used for the snapclone has the same performance capabilities as the disk group used for its source volumes.

Running Concurrent Sessions for One Application System

When running concurrent ZDB sessions that back up one or more of the backup objects residing on the same source volumes, there must be no overlap during the phase of the backup process before the start of data streaming to tape. As a consequence of this, ZDB-to-disk sessions that back up the same objects must not be run in parallel.

Additional considerations with regard to using the same backup specification are as follows:

- With filesystem and disk image backups, concurrent sessions can be run using the same backup specification, but they must not overlap during the phase of the backup process before the start of the data copy to tape media.
- With application integrations, concurrent sessions must not be run using the same backup specification.

ZDB and instant recovery sessions that involve the backup objects residing on the same source volumes cannot be run concurrently. In such cases, the sessions must be run sequentially.

Planning Security

When you plan your backup environment, you need to consider security aspect. A well designed and implemented security plan prevents the unauthorized access and corruption of data. General security considerations related to managing Data Protector cells, user groups, user rights, and other are described in the *HP OpenView Storage Data Protector Concepts Guide*. This section provides you with the security information related to the ZDB environment.

Backup Device and Disk Locking Concept

Backup Device Locking

Regular (non-ZDB) Data Protector backup and restore sessions lock a tape device used in the session at the beginning of a backup or restore session and unlock it at the end of the session. The Data Protector tape device locking is described in detail in the *HP OpenView Storage Data Protector Administrator's Guide*.

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. A device is released when streaming of data to a tape device is finished.
- 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 an instant recovery session, tape devices are not used, so there is no tape device locking with these two operations.

Disk Locking

In order to prevent a ZDB or an 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 fails to lock storage volumes needed for the required operation (if they are already locked by another process).

LUN Security on VA

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

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

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

A Supported Configurations

In This Appendix

This appendix gives you the information on the configurations supported with different array types. It is organized as follows:

The configurations described are supported by Hewlett-Packard. For an updated list of supported configurations, please consult http://www.openview.hp.com/products/datapro/spec_0001.html.

In the event that 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.

Note that for all supported configuration, described in this appendix, you can have only one backup system per application, database, filesystem or disk image backup. You cannot back up the same application, database, filesystem or disk image to more than one backup system simultaneously.

Supported HP StorageWorks Disk Array XP Configurations

The following configurations are possible using Data Protector HP StorageWorks Disk Array XP integration:

- Local replication configurations. Refer to “Local Replication Configurations” on page A-4.
- Remote replication configurations. Refer to “Remote Replication Configurations” on page A-7.
- Remote plus local replication configuration. Refer to “Remote Plus Local Replication Configurations” on page 10.
- HP-UX LVM Mirroring Configurations. Refer to “Supported HP-UX LVM Mirroring Configurations” on page 14.

Replicas are created using the split mirror technique. This means that the target volumes (T) produced are exact duplicates of the source volumes (S).

NOTE

For ZDB to tape and ZDB to disk+tape, a separate backup system is connected to the disk array with the target volumes, while the source volumes are connected to the application system. Streaming of data to tape is done from the replica after the pair has been split, meaning that the application system, during the backup, remains online and available for use.

In all the example configurations presented in this section, you can have more than one application system on the application side of the configuration. For more information on configurations with multiple application systems refer to “Backup System Mount Point Creation” on page B-3.

Note that, with all types of configurations, it is also possible to have the application and the backup data spread across multiple disk arrays of the same type.

Local Replication Configurations

Local replication within the HP StorageWorks Disk Array XP is possible if the Business Copy (BC) XP feature is used. This allows you to create three first-level mirrors (replicas) that can be used for ZDB+IR purposes. For the detailed information on how replicas can be used, refer to Chapter 4, “Replica Operations Concepts,” on page 61. The current section provides you with the examples of the configurations supported for the Disk Array XP local replication.

Figures A-1 through A-3 demonstrate the supported local replication configurations. Each configuration has a specific behavioral pattern imposing specific requirements on the control functions in order to guarantee backup and recovery functionality.

Figure A-1 Supported BC XP Configuration 1

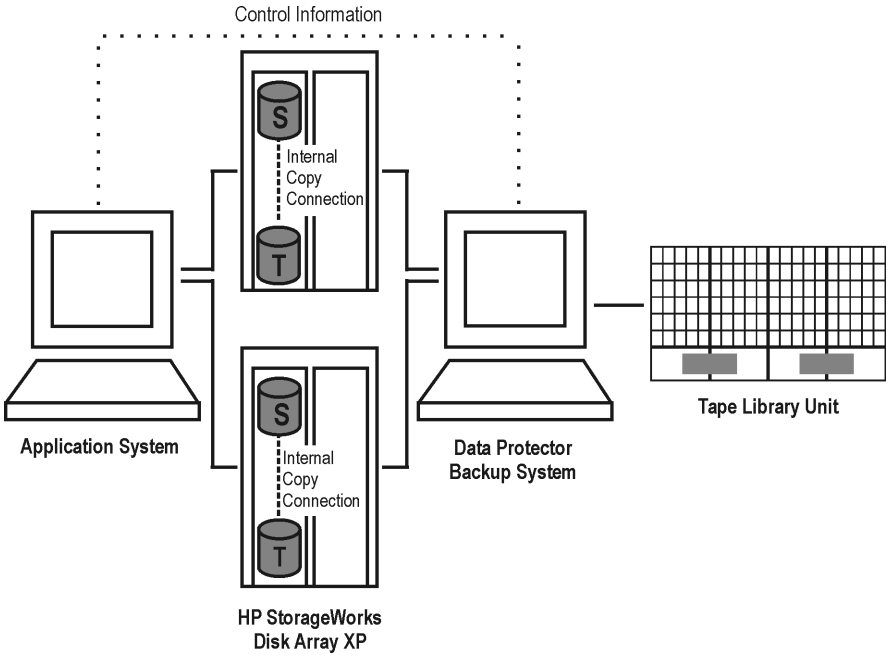

Figure A-2 Supported BC XP Configuration 2

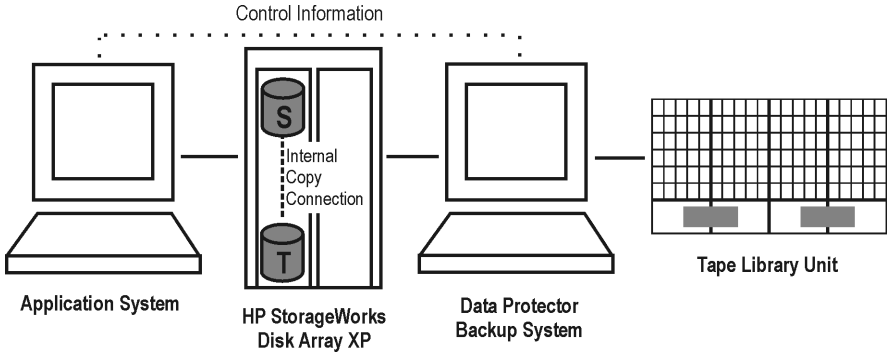

Figure A-3 Supported BC XP Configuration 3

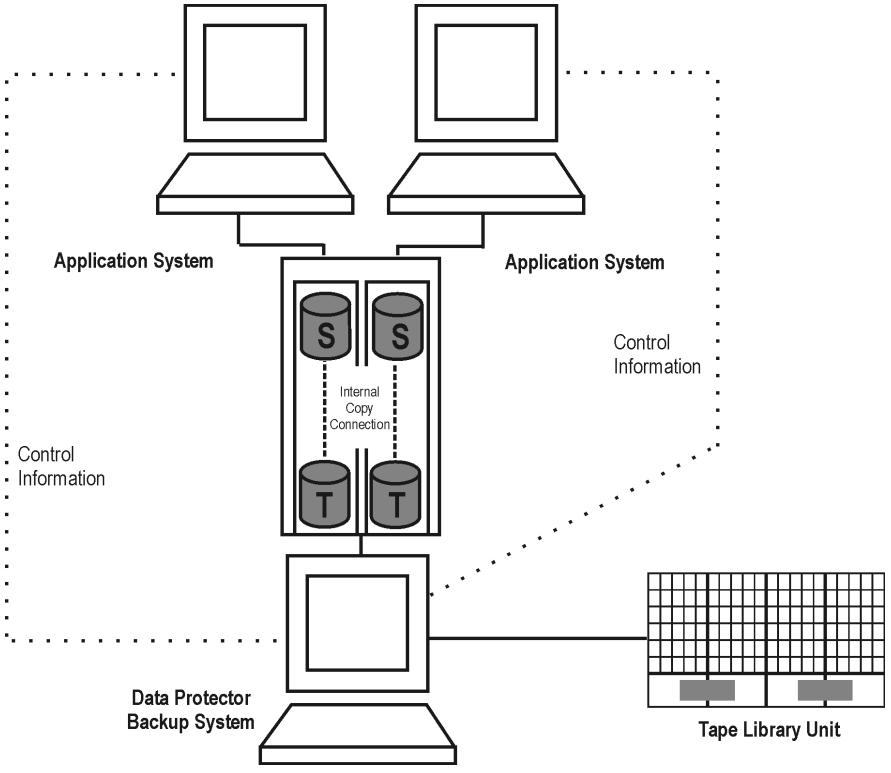

BC1 Configuration

The BC1 configuration, where a single system is used as both the application and the backup system, is not recommended because of the performance issues. Only disk image and filesystem backups are possible using the BC1 configuration. To check on which operating systems the BC1 configuration is supported, check the support matrices in the *HP OpenView Storage Data Protector Software Release Notes*. The example of the BC1 configuration is shown below.

Figure A-4 BC1 XP Configuration

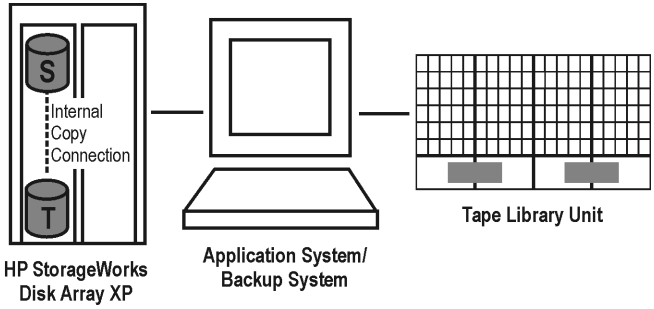

Cascading Configurations

With the Data Protector HP StorageWorks Disk Array XP integration, it is possible to create additional two copies for each of the first-level mirrors if the **cascading** configuration is used. However, only the three first-level mirrors are supported for ZDB+IR purposes. This means you can configure the integration to use cascading, but you need to be aware that the additional second-level mirrors (up to six) will have to be used for the purposes, other than ZDB+IR.

An example of the cascading configuration is presented in Figure A-5 on page A-7, where MU:0, MU:1 and MU:2 are the first-level mirrors, and the six mirrors underneath are the second-level mirrors.

Figure A-5 Cascading Configuration

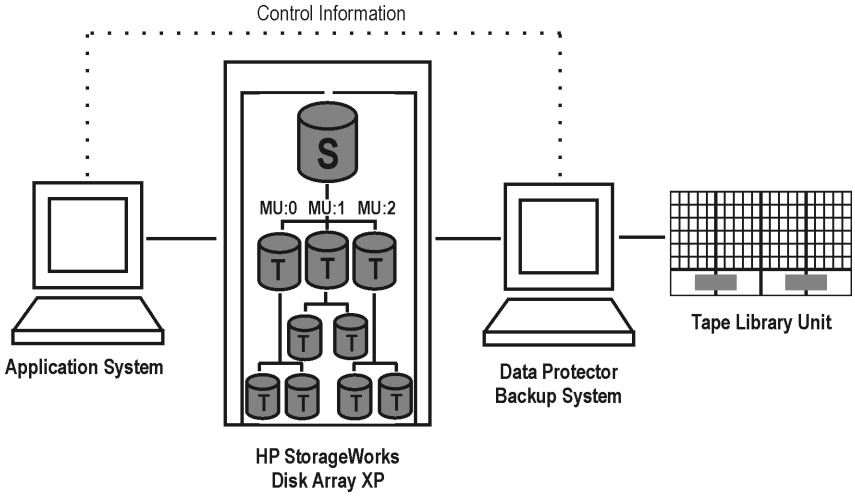

Remote Replication Configurations

Remote replication within the HP StorageWorks Disk Array XP is possible if the Continuous Access (CA) XP feature is used. This allows you to create mirrors (replicas), which can be used for ZDB+IR purposes, on a remote machine, up to 27 miles (43 km) away. For the detailed information on how replicas can be used, refer to Chapter 4, “Replica Operations Concepts,” on page 61. The current section provides you with the examples of the configurations supported for the Disk Array XP remote replication.

A single backup system and a single HP StorageWorks E Disk Array XP can be used to back up multiple main disk arrays. See Figure A-9 on page 9. With this approach you can build a central backup site.

Figures A-6 through A-9 demonstrate the supported remote replication configurations. Each configuration has a specific behavioral pattern imposing specific requirements on the control functions in order to guarantee backup and recovery functionality.

Figure A-6 **Supported CA XP Configuration 1**

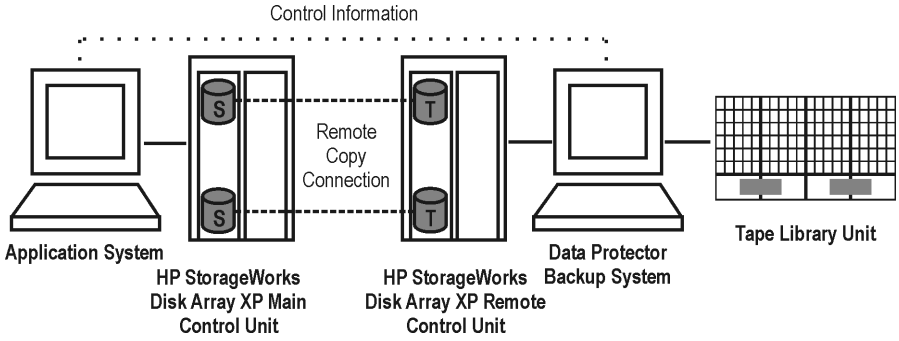

Figure A-7 **Supported CA XP Configuration 20**

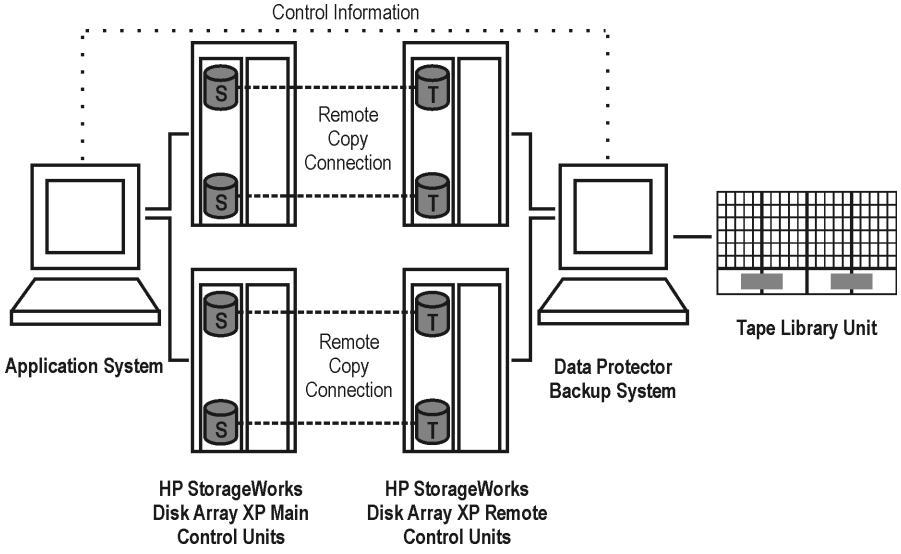

Figure A-8 Supported CA XP Configuration 3

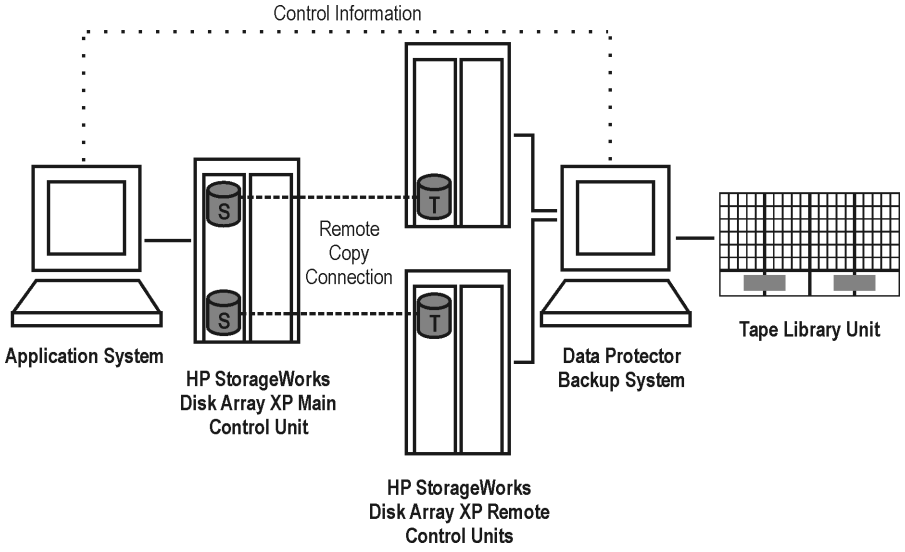

Figure A-9 Supported CA XP Configuration 4

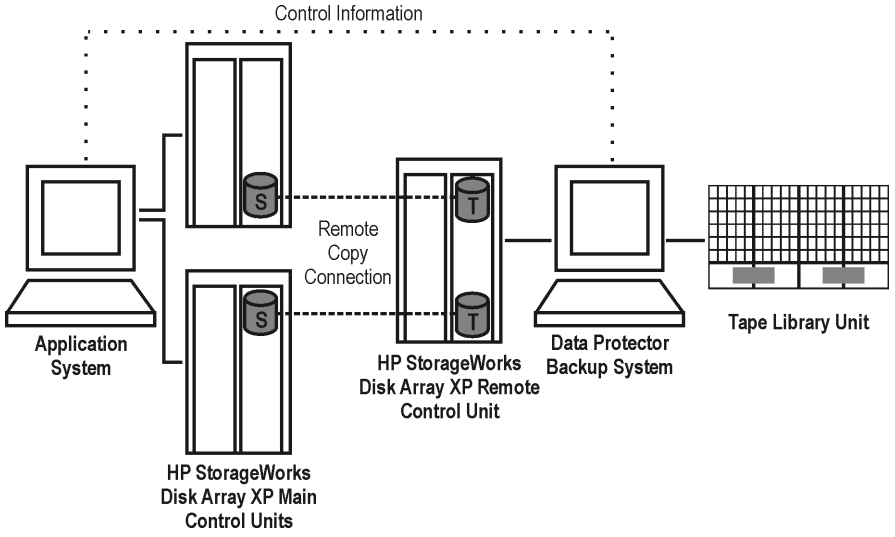

Remote Plus Local Replication Configurations

Remote plus local replication within the HP StorageWorks Disk Array XP is possible if the combination of Continuous Access (CA) XP and Business Copy (BC) XP configurations is used. This allows the creation of (secondary) split mirror replicas on a remote machine, and then creation of local replicas or replica sets of that secondary replica on that remote machine. The replicas created can be used for ZDB+IR purposes. For the detailed information on how replicas can be used, refer to Chapter 4, "Replica Operations Concepts," on page 61. The current section provides you with the examples of the configurations supported for the Disk Array XP remote plus local replication.

NOTE At least two HP StorageWorks Disk Array XPs, located in physically separate sites, are needed for such remote plus local replication.

When a replica is required, the integration splits the BC pair. In order to ensure data consistency, the CA pair status is checked before the BC pair split is executed. In a synchronous CA configuration, this ensures that all data from the Main Control Unit is in the Remote Control Unit.

Cluster Configurations

If the application system is running in a cluster, the backup system must be outside this cluster (it may run in a different cluster, or may not be part of a cluster at all). The reason for this limitation is that during the backup, the filesystem/database structure (filesystem and volume group/disk group) is active on the backup system and would prevent an activation during the failover process.

For an application in a cluster, a floating IP address can be used rather than a static one. This allows a successful start of a backup even after a local failover.

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

IMPORTANT

If a failover to the remote site happens, the backup configuration changes from the previous combined CA+BC to a BC only configuration. This means that the next backup can no longer start automatically, so the backup specification must be updated to reflect the configuration change.

Limitations

- On Windows systems, the target volumes (T) of the CA pair, which are also the source volumes (S) of the BC pair, should not be connected to the backup system together with the target volumes of the BC pair. This is a Windows limitation: two disks with the same signature cannot be connected to the same system. Since the source and the target volumes of the BC pair are a mirror image of one another, they have the same disk signature.
- The asynchronous CA configuration as a part of the combined CA + BC configuration is not supported.

HP-UX Systems

On HP-UX, it is recommended that only the BC target volume be 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, refer to Chapter 3, “HP StorageWorks Disk Array XP” in the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator’s Guide*.

Figures A-26 through A-29 demonstrate the supported remote plus local replication configurations. Each configuration has a specific behavioral pattern imposing specific requirements on the control functions in order to guarantee backup and recovery functionality.

Figure A-10 Supported CA+BC XP Configuration 1

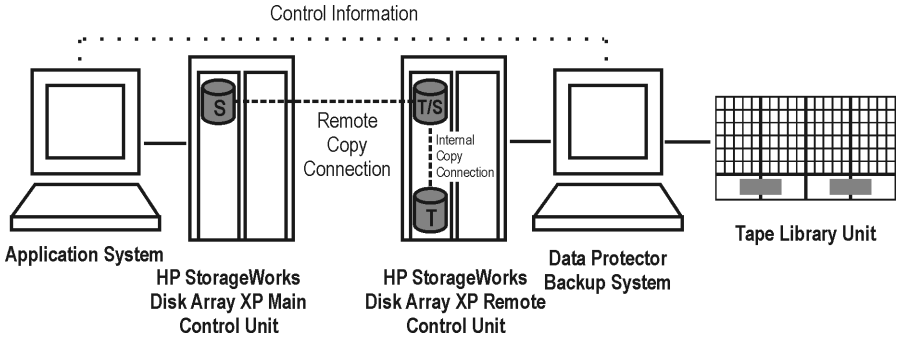

Figure A-11 Supported CA+BC XP Configuration 2

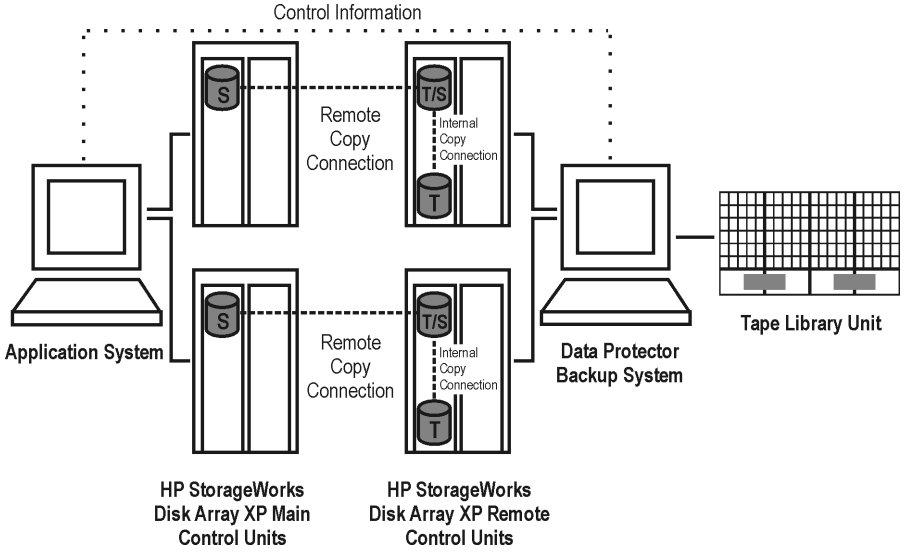

Figure A-12 Supported CA+BC XP Configuration 3

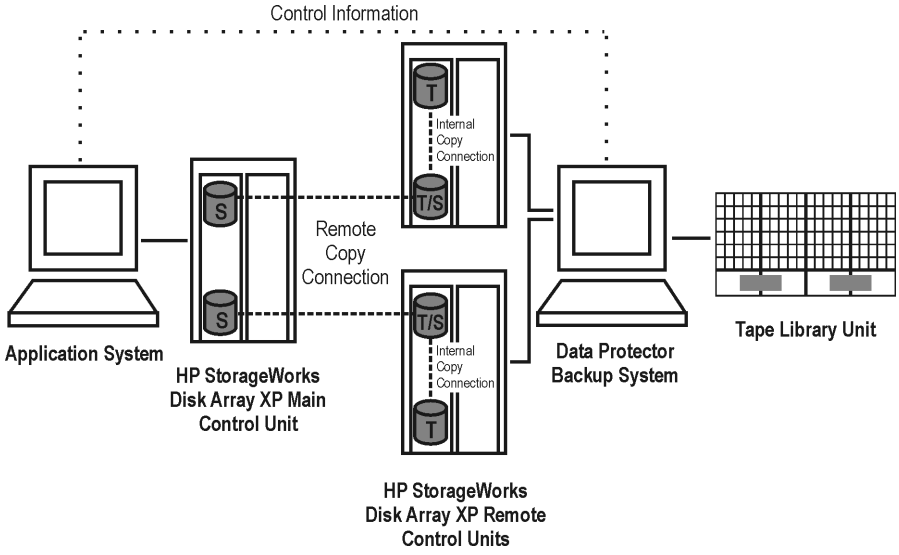

Figure A-13 Supported CA+BC XP Configuration 4

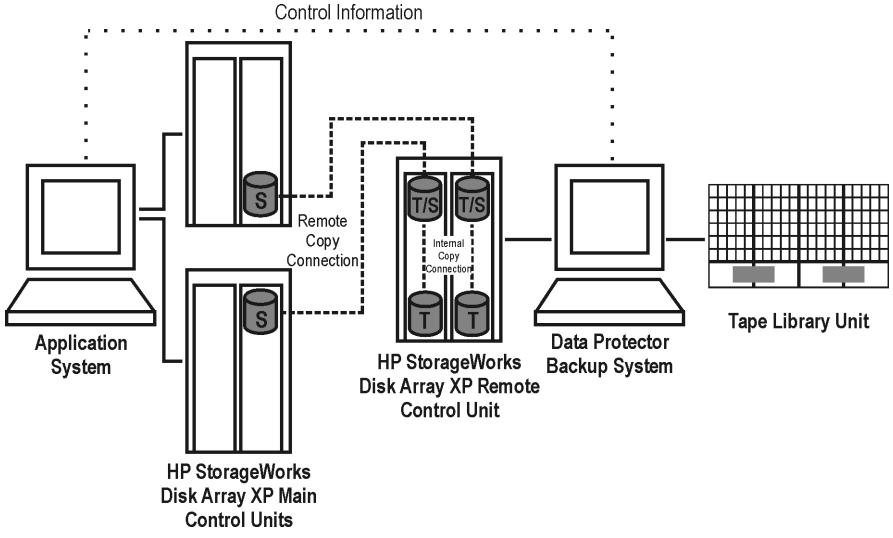

Supported HP-UX LVM Mirroring Configurations

This integration supports HP-UX Logical Volume Manager Mirroring (LVM Mirroring) in a configuration where HP StorageWorks Disk Array XP LDEVs are LVM mirrored from one or more HP StorageWorks Disk Array XP unit(s) to one or more other HP StorageWorks Disk Array XP unit(s). In other words, LVM mirroring can be regarded as a substitute for the CA remote copy connection in a combined CA+BC configuration, which, from the Data Protector HP StorageWorks Disk Array XP integration point of view, becomes a BC configuration.

It is recommended that LDEVs in every HP StorageWorks Disk Array XP unit belong to a different physical volume group, so that extending logical volumes will not give unpredictable results such as mirroring a logical volume onto the same disk. The LDEVs in the unit(s) that is/are connected to the backup system need to have their BC pairs assigned. The application system has to be connected to those HP StorageWorks Disk Array XP units that contain LDEVs belonging to LVM mirrored logical volumes.

Figures A-14 through A-18 demonstrate the supported LVM Mirroring configurations. Each configuration has a specific behavioral pattern imposing specific requirements on the control functions in order to guarantee backup and recovery functionality.

NOTE

If the LVM Mirroring configuration is used, a warning message is issued in the Data Protector monitor during the backup or restore process, since the volume group LDEVs in the physical volume group on the application system do not have their BC pairs assigned. This warning message should be ignored.

Figure A-14 Supported LVM Mirroring Configuration 1

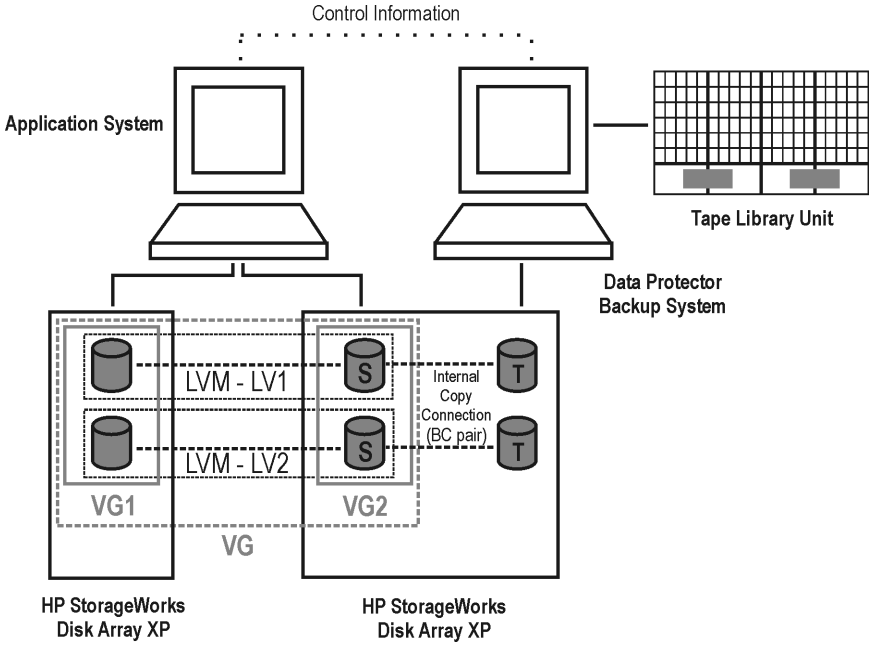

Figure A-15 Supported LVM Mirroring Configuration 2

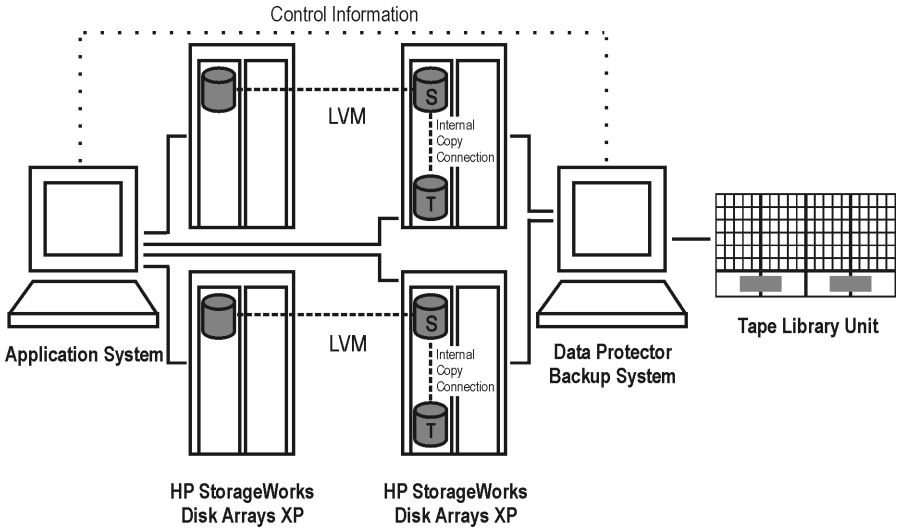

Figure A-16 Supported LVM Mirroring Configuration 3

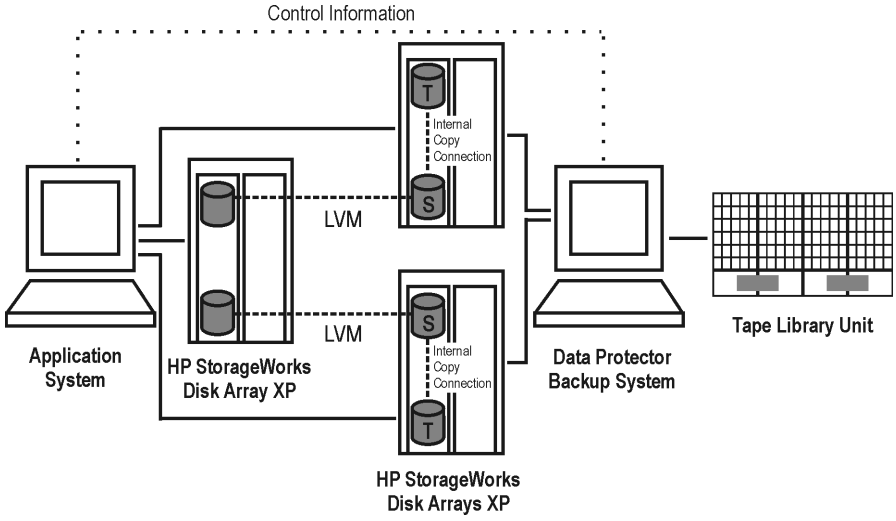

Figure A-17 Supported LVM Mirroring Configuration 4

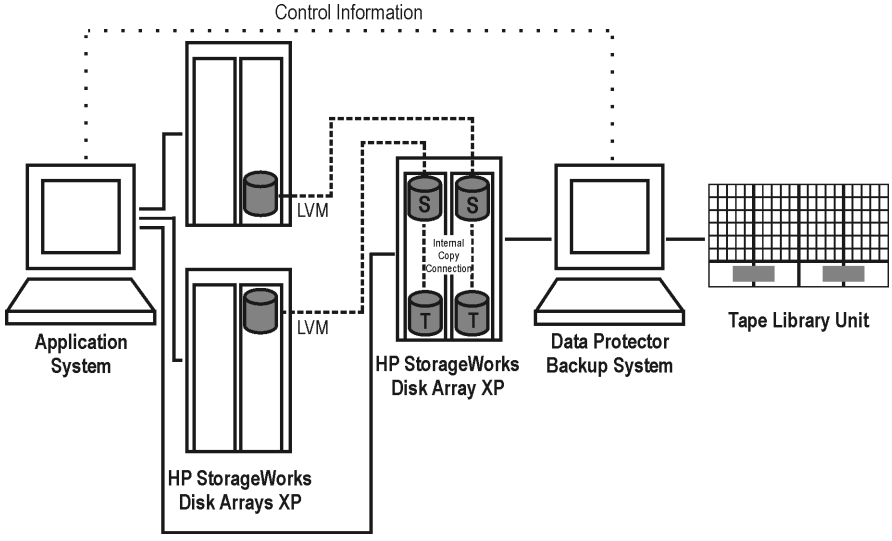

Figure A-18 Supported LVM Mirroring Configuration in a Cluster

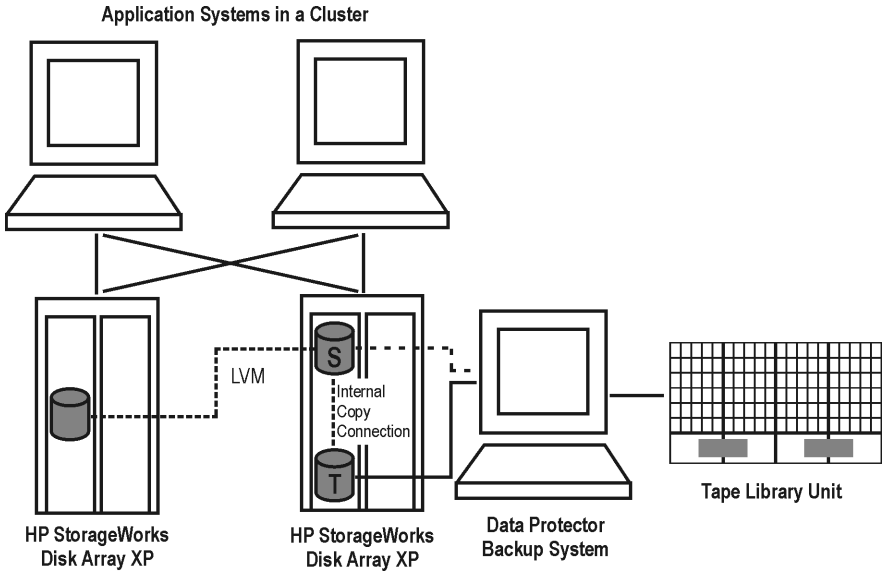

Supported EMC Symmetrix Configurations

The following configurations are possible using Data Protector HP StorageWorks Disk Array XP integration:

- Local replication configurations. Refer to “Local Replication Configurations” on page A-19.
- Remote replication configurations. Refer to “Remote Replication Configurations” on page A-21.
- Remote plus local replication configuration. Refer to “Remote Plus Local Replication Configurations” on page 23.

Replicas are created using the split mirror technique. This means that the target volumes (T) produced are exact duplicates of the source volumes (S).

NOTE

For a backup, a separate backup (R2) system needs to be connected to the disk array with the target volumes, while the source volumes are connected to the application (R1) system. Streaming of data to tape is done from the replica after the pair has been split, meaning that the application (R1) system, during the backup, remains online and available for use.

Limitation

Instant recovery is not supported with the Data Protector EMC Symmetrix integration.

In all the example configurations presented in this section, you can have more than one application system (R1) on the application side of the configuration. For more information on configurations with multiple application systems, refer to “Backup System Mount Point Creation” on page B-3.

Note that, with all three types of configurations, it is also possible to have the application and the backup data spread across multiple EMC Symmetrix disk arrays on the application and backup side of the configuration.

Local Replication Configurations

Local replication within the EMC Symmetrix Disk Array is possible if the TimeFinder configuration is used. This allows you to create single mirrors (replicas) that can be used for ZDB and split mirror restore purposes. For the detailed information on how replicas can be used, refer to Chapter 4, “Replica Operations Concepts,” on page 61. The current section provides you with the examples of the configurations supported for the EMC Symmetrix local replication.

Figures A-19 through A-21 demonstrate the supported local replication configurations. Each configuration has a specific behavioral pattern imposing specific requirements on the control functions in order to guarantee backup and restore functionality.

Figure A-19 Supported TimeFinder Configuration 1

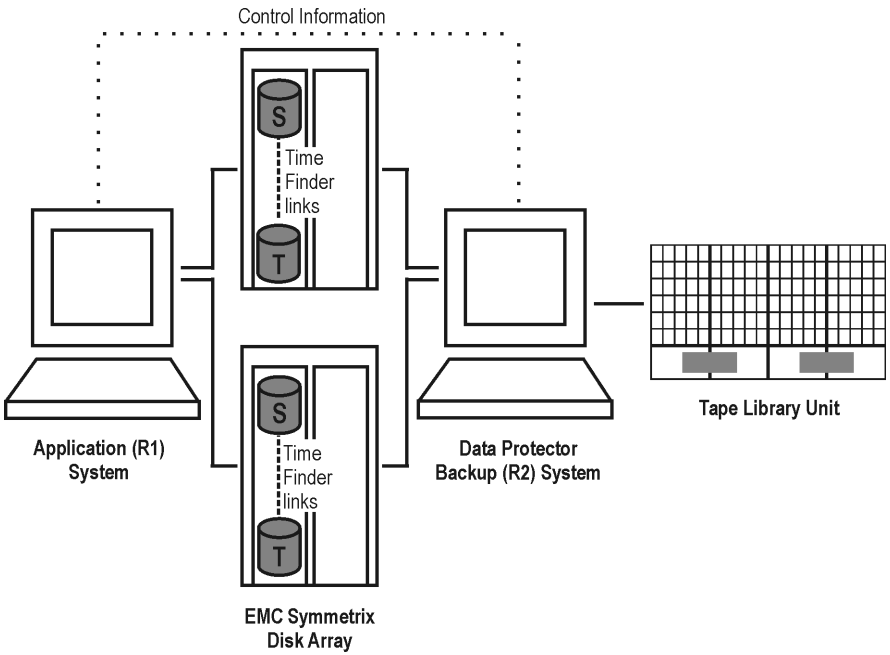

Figure A-20 Supported TimeFinder Configuration 2

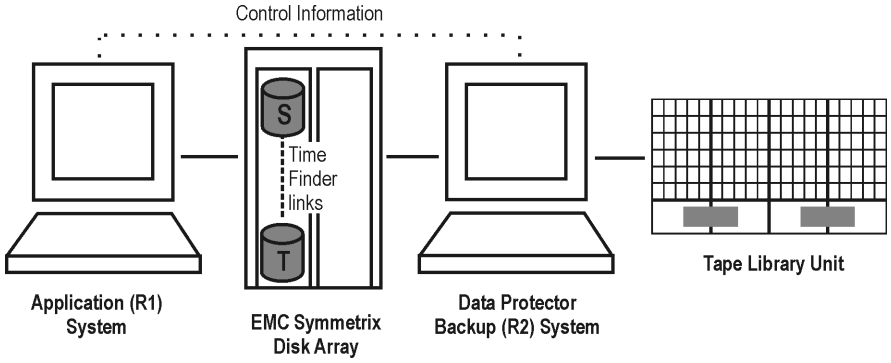

Figure A-21 Supported TimeFinder Configuration 3

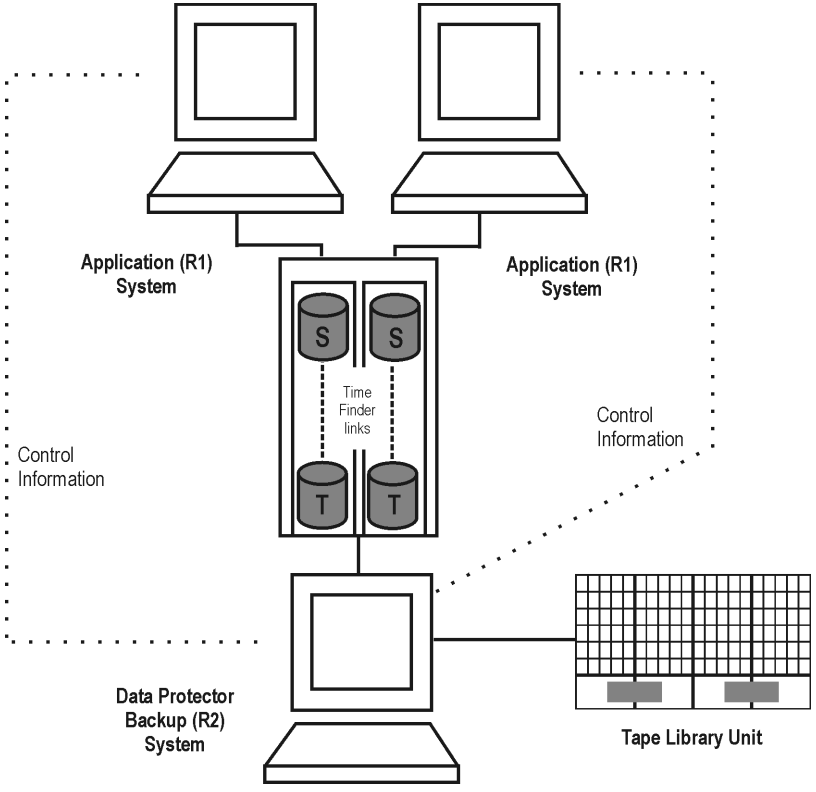

TimeFinder1 Configuration

TimeFinder1 configuration, where a single system is used as both the application and the backup system, is not recommended because of the performance issues. Only disk image and filesystem backups are possible using the TimeFinder1 configuration. For a list of the operating systems on which this configuration is supported, check the support matrices in the *HP OpenView Storage Data Protector Software Release Notes*.

Remote Replication Configurations

Remote replication within the EMC Symmetrix Disk Array is possible if the EMC Symmetrix Remote Data Facility (SRDF) configuration is used. This allows you to create single split mirror replica on a remote machine. For the detailed information on how replicas can be used, refer to Chapter 4, “Replica Operations Concepts,” on page 61. The current section provides you with the examples of the configurations supported for the EMC Symmetrix remote replication.

Limitation

A cluster configuration is not supported in this environment.

Figures A-22 through A-25 demonstrate the supported remote replication configurations. Each configuration has a specific behavioral pattern imposing specific requirements on the control functions in order to guarantee backup and recovery functionality.

Figure A-22 Supported SRDF Configuration 1

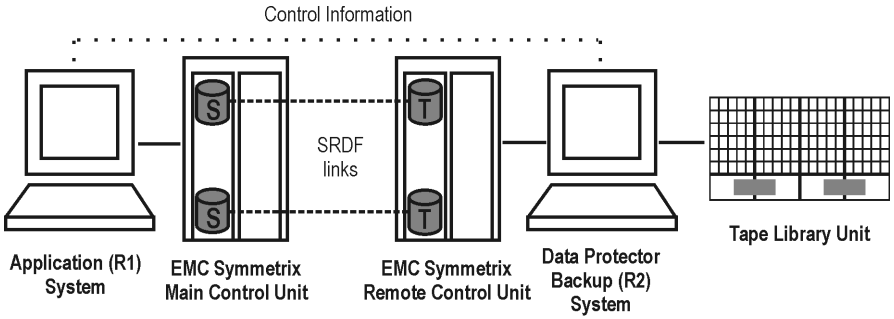

Figure A-23 Supported SRDF Configuration 2

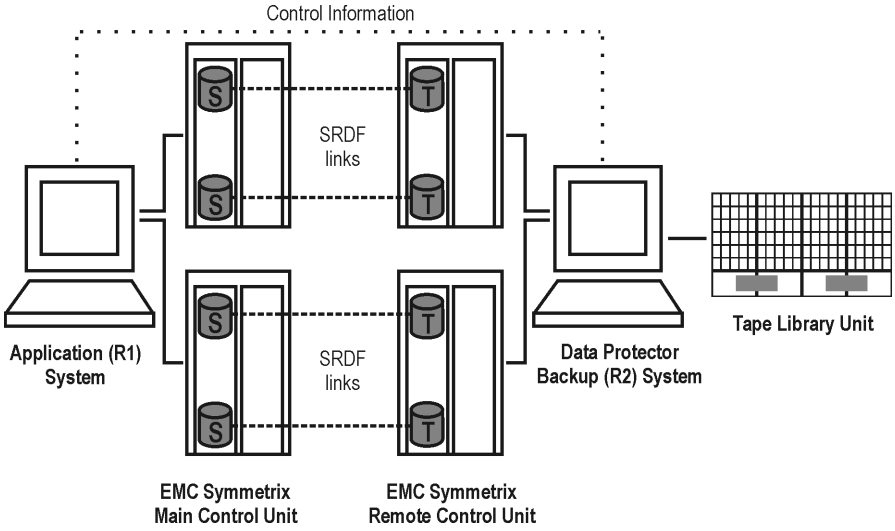

Figure A-24 Supported SRDF Configuration 3

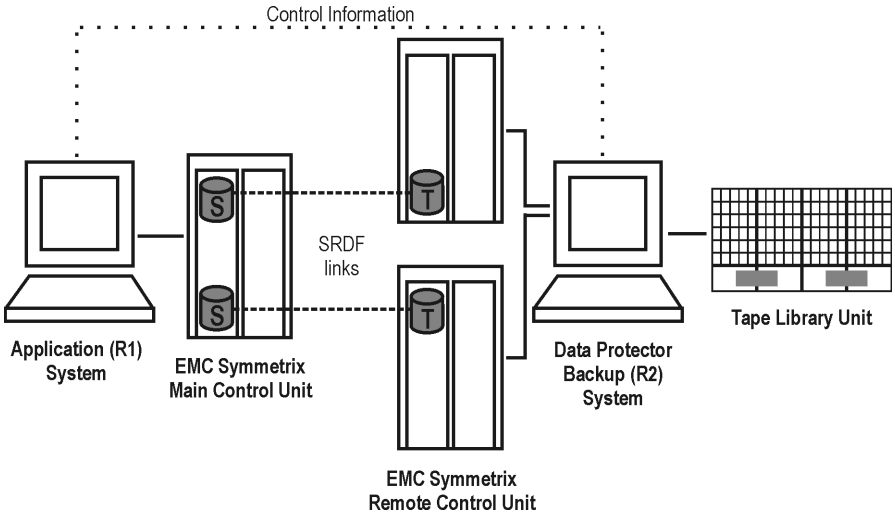

Figure A-25 Supported SRDF Configuration 4

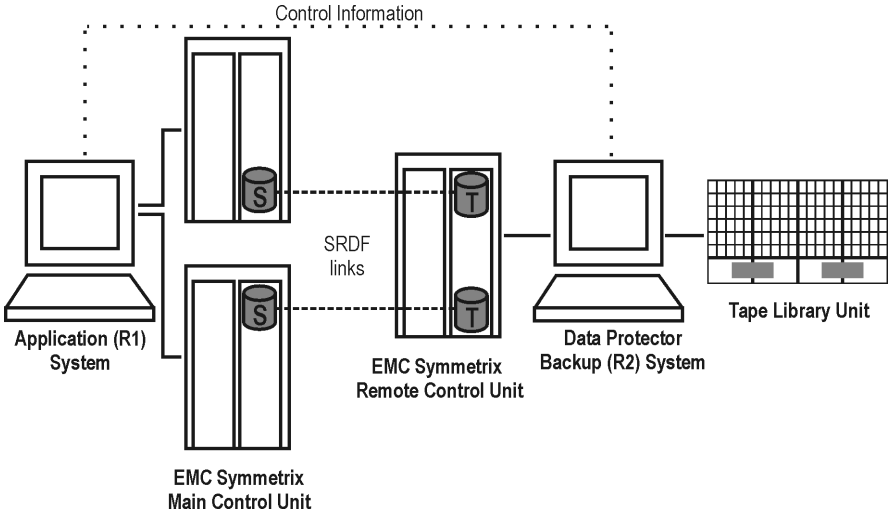

Remote Plus Local Replication Configurations

Remote plus local replication within the EMC Symmetrix Disk Array is possible the combination of EMC Symmetrix Remote Data Facility (SRDF) and EMC Symmetrix TimeFinder configurations is used. This allows the creation of a (secondary) split mirror replica on a remote machine, and then creation of local replicas of that secondary replica on that remote machine. The replicas created can be used for ZDB and split mirror restore purposes. For the description of the split mirror restore, see Chapter 6, “Restore Techniques from ZDB-to-Tape Sessions,” on page 83. The detailed information on how replicas can be used is described in Chapter 4, “Replica Operations Concepts,” on page 61. The current section provides you with the examples of the configurations supported for the EMC Symmetrix remote plus local replication.

NOTE At least two EMC Symmetrix Disk Arrays, located in physically separate sites, are needed for remote plus local replication.

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.

Cluster Configurations

If the application (R1) system is running in a cluster, the backup (R2) system must be outside this cluster (it may run in a different cluster, or may not be part of a cluster at all). The reason for this limitation is that during the backup, the filesystem/database structure (filesystem and volume group/disk group) is active on the backup (R2) system and would prevent an activation during the failover process.

For an application in a cluster, a floating IP address can be used rather than a static one. This allows a successful start of a backup even after a local failover.

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

IMPORTANT

If a failover to the remote site happens, the backup configuration changes from the previous combined SRDF+TimeFinder to a TimeFinder only configuration. This means that the next backup can no longer start automatically, so the backup specification must be updated to reflect the configuration change.

It is recommended that only the TimeFinder target volume be 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, refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide*.

Figures A-26 through A-30 demonstrate the supported remote plus local replication configurations. Each configuration has a specific behavioral pattern imposing specific requirements on the control functions in order to guarantee backup and recovery functionality.

Figure A-26 Supported SRDF+TimeFinder Configuration 1

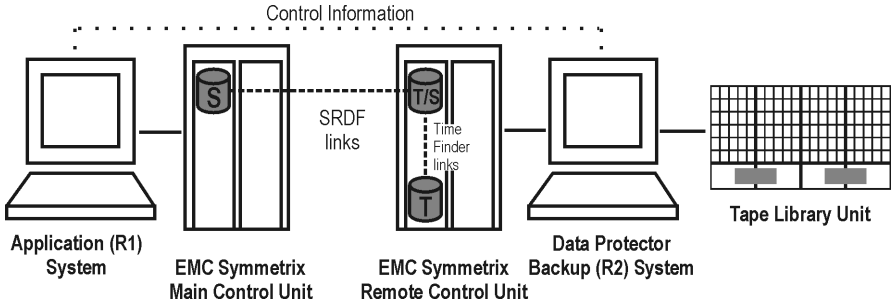

Figure A-27 Supported SRDF+TimeFinder Configuration 2

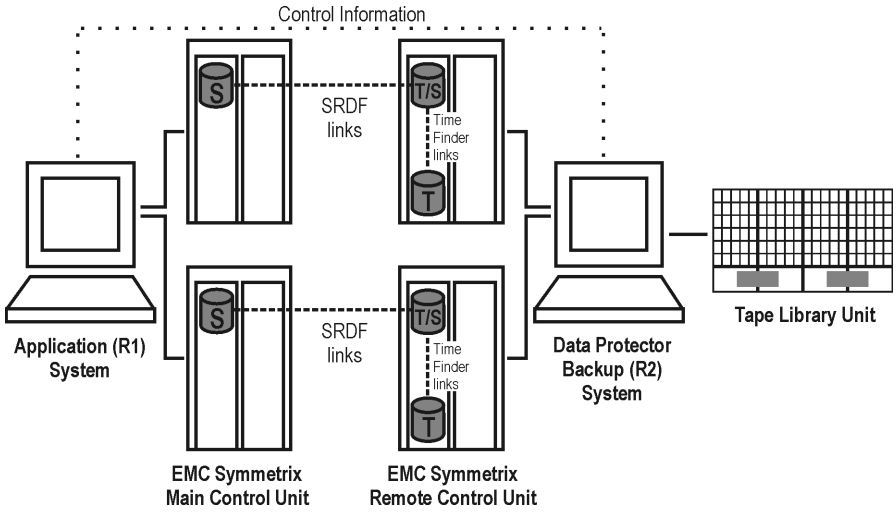

Figure A-28 Supported SRDF+TimeFinder Configuration 3

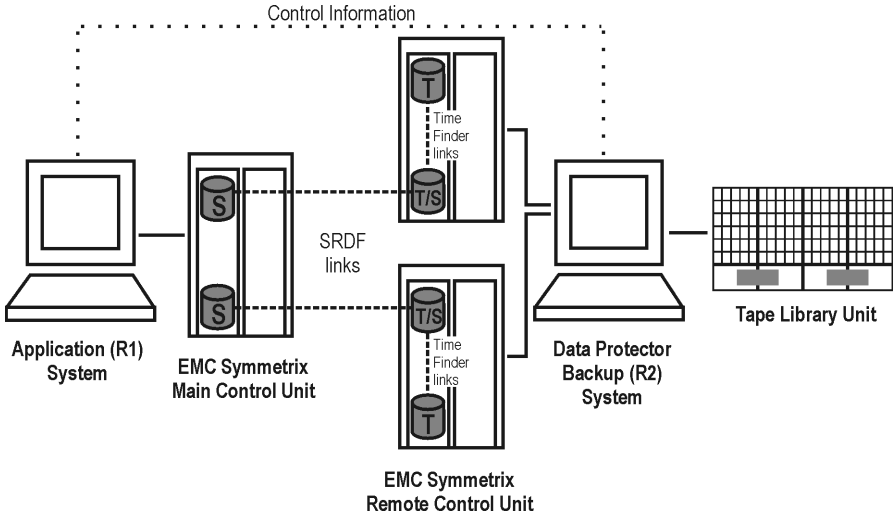

Figure A-29 Supported SRDF+TimeFinder Configuration 4

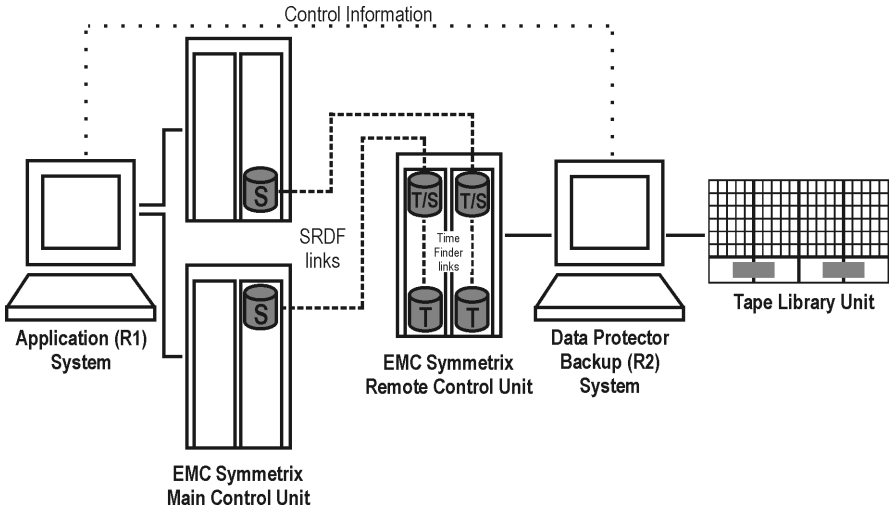

Figure A-30 Supported SRDF+TimeFinder Configuration in a Cluster

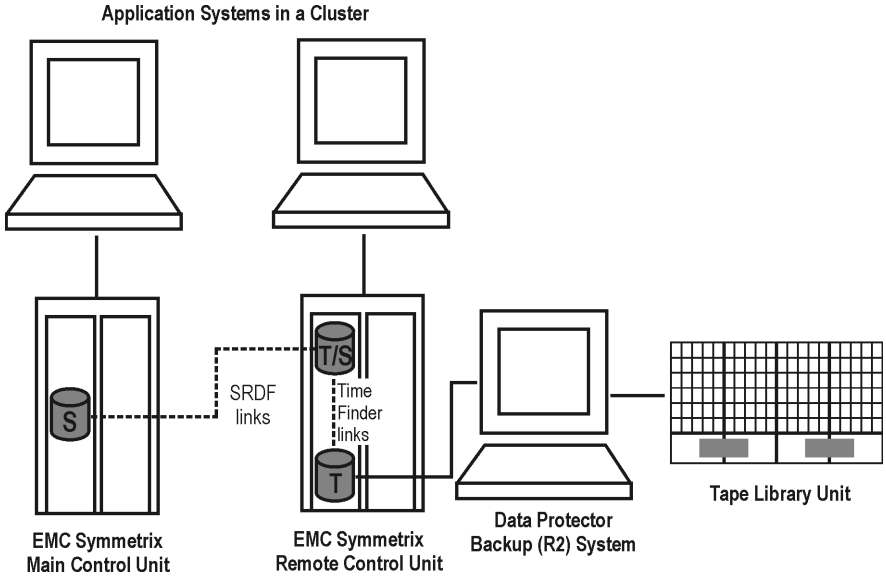

Supported Snapshot Configurations

HP StorageWorks Virtual Array and HP StorageWorks Enterprise Virtual Array configurations are referred to as snapshot configurations, as both arrays provide creation of replicas using snapshot technique. This means that the target volumes (T) produced are logical copies, or images, of the source volumes (S).

The following configurations are possible using Data Protector snapshot integrations:

- Local replication. Refer to “Local Replication Configurations” on page 28.
- Remote plus local replication (HP StorageWorks Virtual Array only). Refer to “Remote Plus Local Replication on HP StorageWorks Virtual Array” on page 31.

NOTE

For ZDB to tape and ZDB to disk+tape, a separate backup system is normally 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 systems), and performs all other necessary steps to mount the filesystems on the backup system so that it can access the replicated data. Streaming of data to tape is performed from the replica, while the application system continues with operations.

In all the example configurations presented in this section, it is possible to have more than one application system on the application side of the configuration. For more information on how the mount points are created on the backup system in configurations with multiple application systems, refer to “Backup System Mount Point Creation” on page B-3.

Note that, with all three types configurations, you can have the application and the backup data spread across multiple disk arrays.

Local Replication Configurations

Local replication within HP StorageWorks Virtual Array and HP StorageWorks Enterprise Virtual Array is possible if the Business Copy (BC) VA or Business Copy (BC) EVA configuration is used. This allows

the data replication within the same array. With this, large replica sets can be used, the number of members being limited primarily by the available space on the array.

Figures A-31 through A-33 demonstrate the supported local replication configurations. Each configuration has a specific behavioral pattern imposing specific requirements on the control functions in order to guarantee backup and recovery functionality.

Figure A-31 Supported BC Snapshot Configuration 1

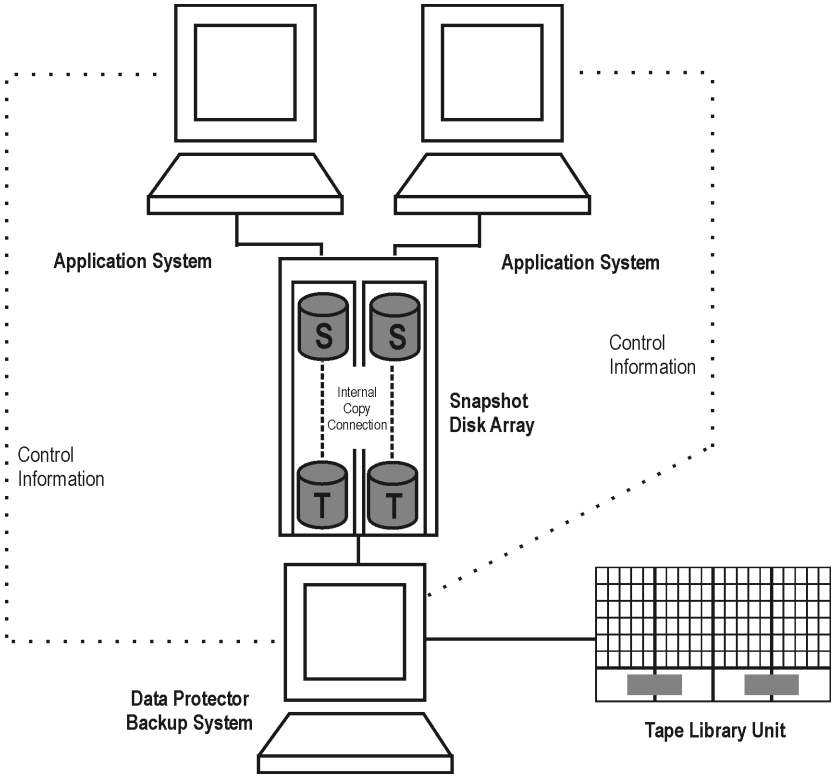

Figure A-32 Supported BC Snapshot Configuration 2

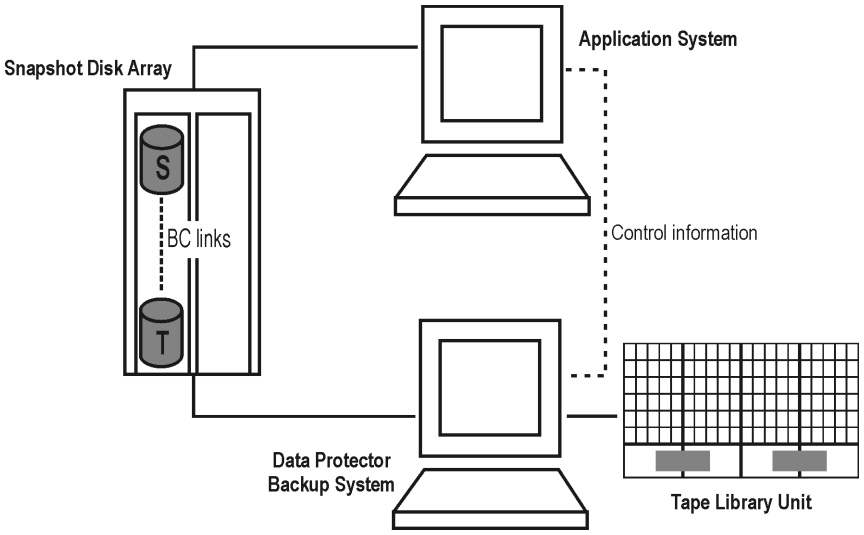

Figure A-33 Supported BC Snapshot Configuration 3

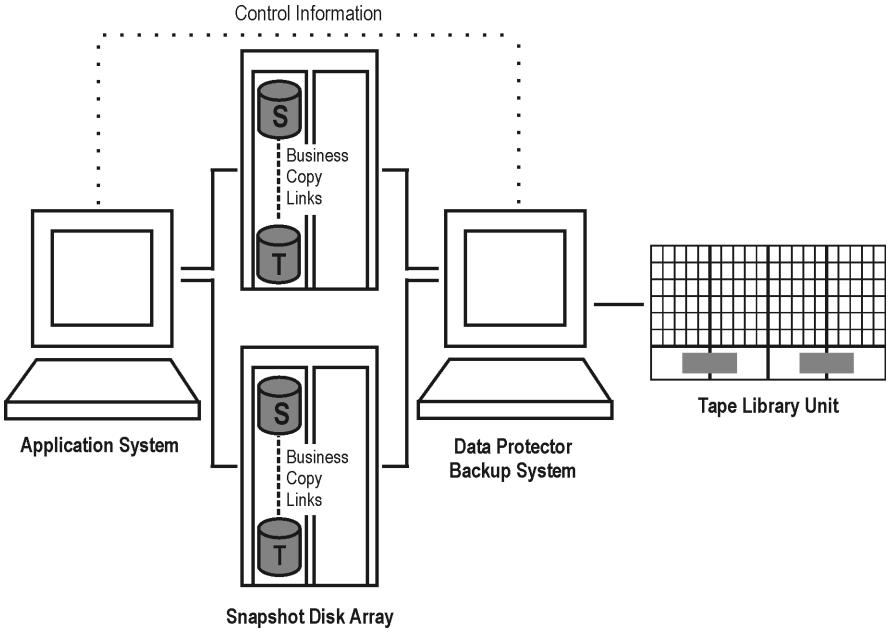

Single Host Configuration

Single host configuration, where a single system is used as both the application and the backup system, is not recommended because of performance issues. Only disk image and filesystem backups are possible using the single host configuration. Clusters are not supported in such a configuration.

Remote Plus Local Replication on HP StorageWorks Virtual Array

Remote plus local replication within HP StorageWorks Virtual Array is possible if the HP-UX Logical Volume Manager Mirroring (LVM Mirroring) is used. This allows the creation of replicas on a remote machine; after that, local replicas of that remote replica can be created using the snapshot technique.

At least two HP StorageWorks Virtual Arrays, located in physically separate sites, are needed for such a configuration. The VA source volumes are LVM-mirrored from one or more local VA disk array(s) to one or more remote VA disk array(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.

Figures A-34 through A-38 demonstrate the supported LVM mirroring configurations. Each configuration has its own specific behavioral pattern, which imposes specific requirements on the control functions, in order to guarantee backup and restore functionality.

IMPORTANT

Instant recovery is not supported when LVM mirroring is used.

Figure A-34 Supported LVM Mirroring Configuration 1

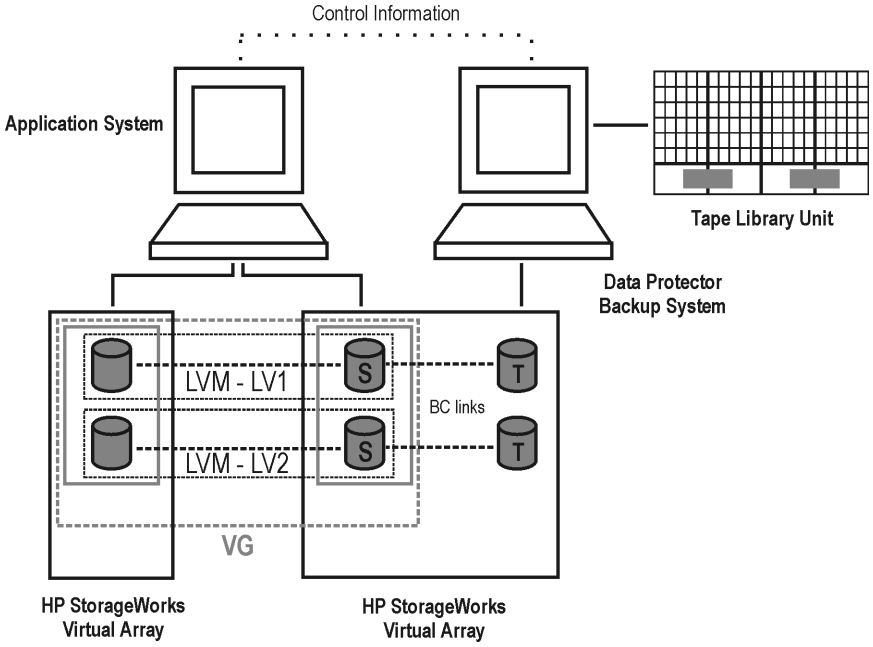

Figure A-35 Supported LVM Mirroring Configuration 2

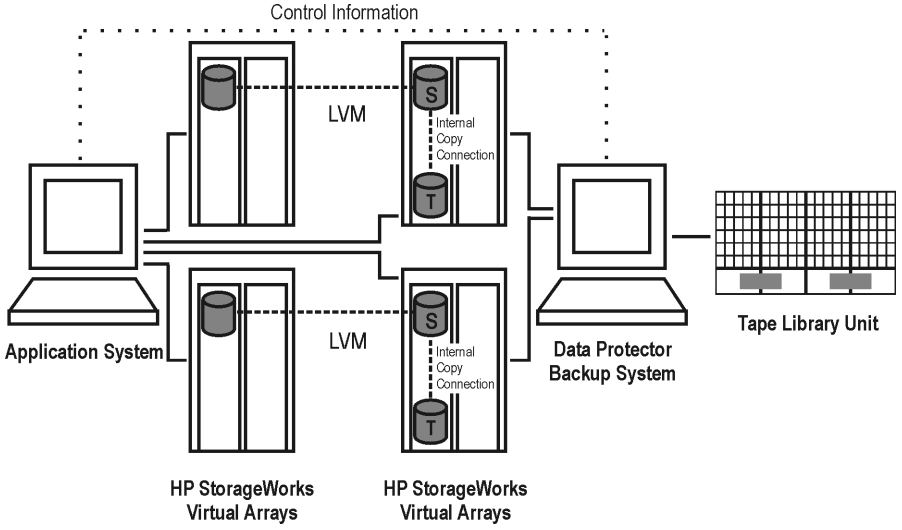

Figure A-36 Supported LVM Mirroring Configuration 3

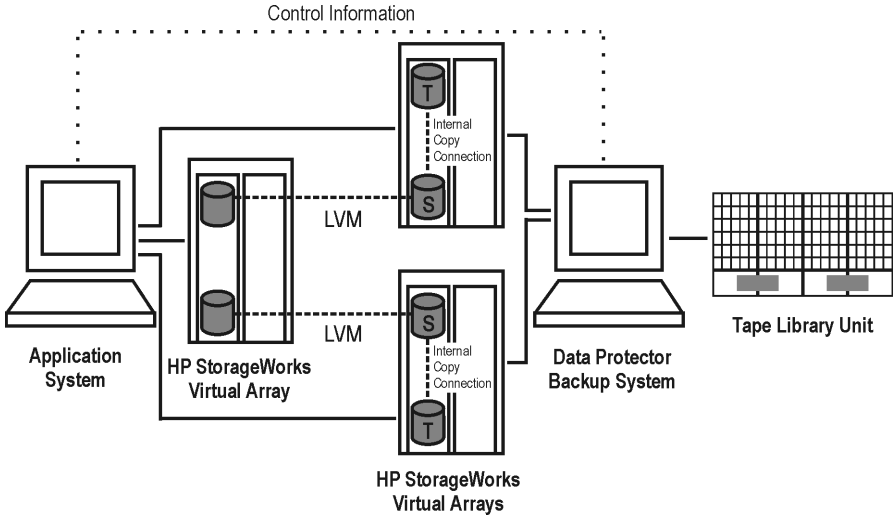

Figure A-37 Supported LVM Mirroring Configuration 4

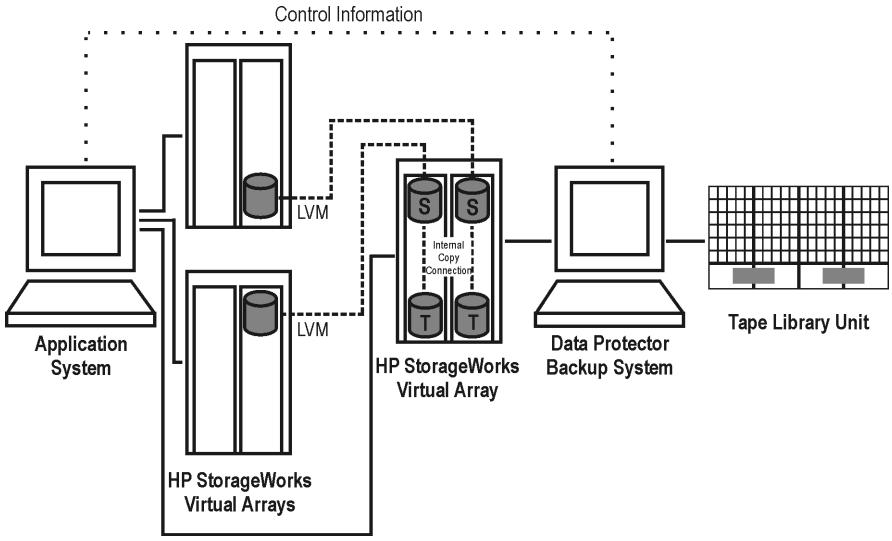

Figure A-38 LVM Mirroring Configuration in a Cluster

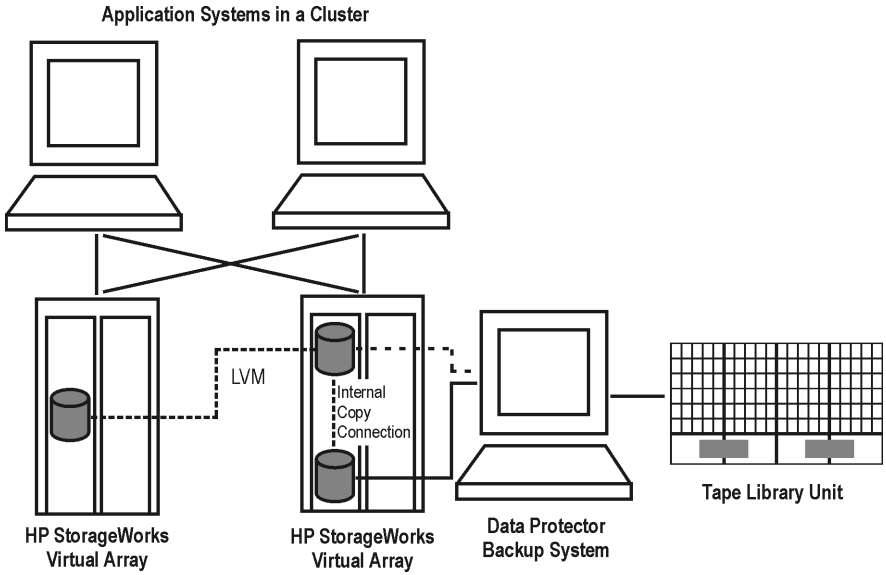

B Additional Information

In This Appendix

This appendix explains how the backup system mount point creation is done. It also gives you the information on the ZDB database.

The appendix is organized into the following sections:

Backup System Mount Point Creation

Data Protector disk array integrations support configurations where multiple application systems are connected to a disk array and one system (the backup system) is responsible for backing up these applications. Local, remote, or remote plus local replication configuration (if supported on a particular array) can be used for ZDB in such a configuration. For more information on supported configurations, refer to Appendix A.

Each application system uses its own original storage, from which replicas are created; in case of ZDB to tape and ZDB to disk+tape, filesystems are mounted on the backup system.

Filesystem and MS Exchange 2000 Backup

To perform a concurrent backup of multiple application systems, the mount points assigned to the filesystems in the original storage *do not need to be* different for each application system. The backup of the MS Exchange 2000 application is performed as *filesystem* backup. With filesystem backup, Data Protector, during a ZDB session, creates or reuses unique mount points on the backup system. Data Protector then mounts filesystems to these mount points. For more information on the backup process for a particular array type, refer to *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide*.

Figure B-1 **Backup System Mount Point Creation: Filesystem and MS Exchange 2000 Backup**

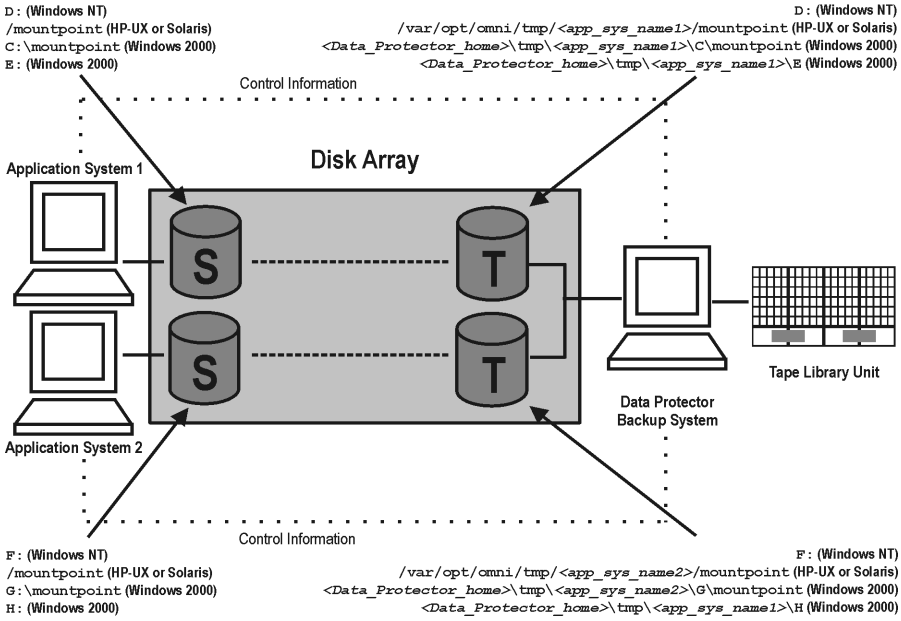

NOTE

The above example depicts the default Data Protector behavior. You can change the backup system mount point pathname creation by setting the `ZDB_PRESERVE_MOUNTPOINTS`, `ZDB_MOUNT_PATH` and `ZDB_MULTI_MOUNT` variables in the `.omnirc` file. Refer to *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide* for more information on these variables.

Application and Disk Image Backup

The information in this section applies only for the backup of the following:

- Disk images
- Oracle 8i/9i
- SAP R/3

- MS SQL Server 2000

For a list of applications, supported for a particular type of a disk array, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

Applications on Filesystems

To perform a concurrent backup of multiple application systems, the mount points or drive letters assigned to the original storage *must be* different for each application system. Data Protector, during a ZDB session, creates mount points or drive letters with the same names as on the application system. Data Protector then mounts filesystems in a replica to these mount points.

If the mount points or drive letters are the same for different application systems, concurrent backup of such systems is not possible; backup of objects that belong to these mount points or drive letters must be run sequentially.

Applications on Disk Images + Disk Image Backup

If your application uses raw disk images as the data source, or if you are performing a disk image backup without an application, the following applies: Data Protector, during a ZDB session, finds and uses raw device files (UNIX systems) or physical drive numbers (Windows systems) for the replica created from the original storage raw device files (UNIX systems) or physical drive numbers (Windows systems) on the backup system. Therefore, make sure the device file names and physical drive numbers are the same on the application and the backup systems.

Note that due to the limitation described above, snapshot integrations are not suitable for such backups (with snapshot integrations, Data Protector cannot guarantee that after presentation to the backup system replicas are assigned the same raw device files or physical drive numbers as on the application system).

NOTE

On HP StorageWorks Disk Array XP, if the BC first level mirrors are configured, the integration always mounts the selected first level mirror to the same mount point.

Figure B-2 **Backup System Mount Point Creation: Application or Disk Image Backup**

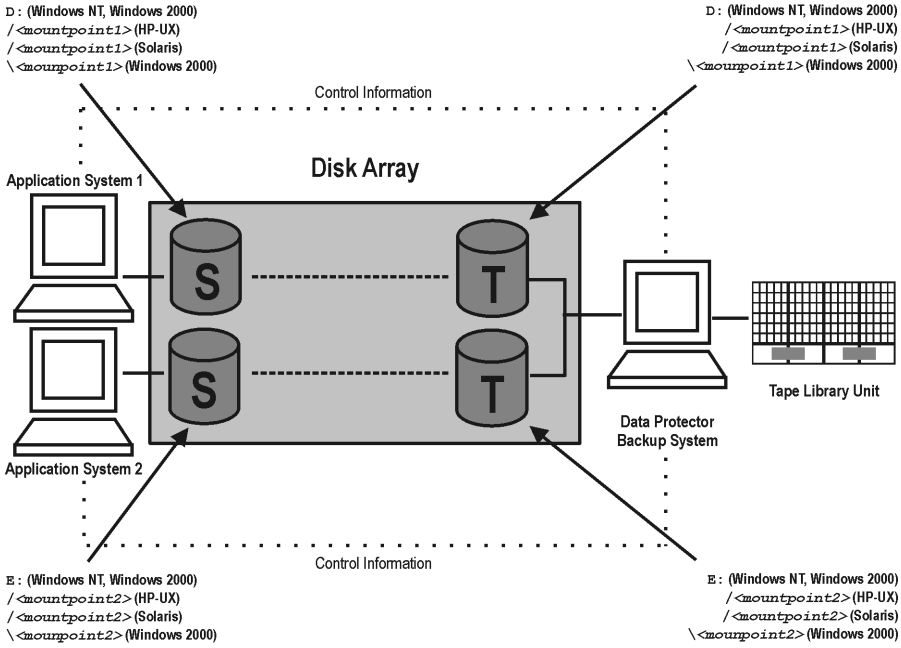

ZDB Database

The ZDB database is an extension to the Data Protector internal database (IDB) on the Cell Manager. It is used to hold array specific information about a replica that is required, in addition to the normal backup session information held in the IDB, 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.
- EVADB and SMISDB for HP StorageWorks Enterprise Virtual Array.

The exact information stored in the ZDB varies a little for each array-related section, because of differences in the way the arrays operate. Generally speaking, each section contains the following type of information:

- Information on the replicas that are kept on disk array(s). This includes:
  - The backup session ID.
  - Information on when the backup session was performed.
  - Name of the backup specification used in the backup session.
  - Name, ID, and WWN of the target volume created in the backup session.
  - Name and ID of the EVA storage system on which the target volume resides.
  - On EVA, the information on the target volume type (VSNAP, pre-allocated snapshot, or snapclone).
  - ID of a source volume used in the backup session.
  - Information on whether the target volume can be used for instant recovery (IR flag).
  - Information on whether the target volume should be deleted (purge flag).

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

- Some of the disk array security information.
- CRC check information calculated during the ZDB-to-disk session. This is applicable to the Data Protector HP StorageWorks Disk Array XP and HP StorageWorks Virtual Array integrations.
- On HP StorageWorks Enterprise Virtual Array, the information on disk group pairs and some information on the EVA hardware configuration.
- On HP StorageWorks Disk Array XP, the information about XP command devices.

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

The ZDB database stores the information only about those 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 the Data Protector internal database (IDB).

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

Glossary

access rights
See **user rights**.

ACSLS *(StorageTek specific term)*
The Automated Cartridge System Library Server (ACSLS) software that manages the Automated Cartridge System (ACS).

Active Directory *(Windows specific term)*
The directory service in a Windows network. It contains information about resources on the network and makes them accessible to users and applications. The directory services provide a consistent way to name, describe, locate, access and manage resources regardless of the physical system they reside on.

AML *(EMASS/GRAU specific term)*
Automated Mixed-Media library.

application agent
A component needed on a client to back up or restore online database integrations.
See also **Disk Agent**.

application system *(ZDB specific term)*
A system the application or database runs on. The application or database data is located on source volumes.
See also **backup system** and **source volume**.

archived redo log *(Oracle specific term)*
Also called offline redo log. If the Oracle database operates in the ARCHIVELOG mode, as each online redo log is filled, it is copied to one (or more) archived log destination(s). This copy is the Archived Redo Log. The presence or absence of an Archived Redo Log is determined by the mode that the database is using:

- ARCHIVELOG - The filled online redo log files are archived before they are reused. The database can be recovered from an instance and disk failure. The “hot” backup can be performed only when the database is running in this mode.

- NOARCHIVELOG - The filled online redo log files are not archived.

See also **online redo log.**

archive logging *(Lotus Domino Server specific term)*
Lotus Domino Server database mode where transaction log files are overwritten only after they have been backed up.

ASR Set
A collection of files stored on several diskettes, required for proper reconfiguration of the replacement disk

(disk partitioning and logical volume configuration) and automatic recovery of the original system configuration and user data that was backed up during the full client backup.
These files are stored as an ASR archive file on the Cell Manager (in *<Data_Protector_home>*\Config\Server\dr\asr on a Windows Cell Manager or in /etc/opt/omni/server/dr/asr/ on a UNIX Cell Manager) as well as on the backup medium. The ASR archive file is extracted to three diskettes for 32-bit Windows systems or four diskettes for 64-bit Windows systems after a disaster occurs. You need these diskettes to perform ASR.

autochanger
See **library**

autoloader
See **library**

BACKINT *(SAP R/3 specific term)*
SAP R/3 backup programs can call the Data Protector backint interface program via an open interface, which enables them to communicate with Data Protector software. For backup and restore, SAP R/3 programs issue orders for the Data Protector backint interface.

backup API
The Oracle interface between the Oracle backup/restore utility and the backup/restore media management layer. The interface defines a set of routines to allow the reading and writing of data to the backup media, the creation, searching and removing the backup files.

backup chain
This relates to a situation where full and incremental backups are performed. Based on the level of the incremental backups used (Incr, Incr 1, Incr 2, and so on), simple or rather complex dependencies of incrementals to previous incrementals can exist. The backup chain are all backups, starting from the full backup plus all the dependent incrementals up to the desired point in time.

backup device
A device configured for use with Data Protector, which can write data to and read data from storage media. This can be, for example, a standalone DDS/DAT drive or a library.

backup generation
One backup generation includes one full backup and all incremental backups until the next full backup.

backup ID
An identifier of an integration object that equals the session ID of the backup of this object. The backup ID is preserved when an object is copied, exported, or imported.

backup object
A backup unit that contains all items backed up from one disk volume (logical disk or mount point). The backed up items can be any number of files, directories, or the entire disk or mount point. Additionally, a backup object can be a database entity or a disk image (rawdisk).

A backup object is defined by:

- Client name: hostname of the Data Protector client where the backup object resides.
- Mount point: the access point in a directory structure (drive on Windows and mount point on UNIX) on the client where the backup object is located.
- Description: uniquely defines backup objects with identical client name and mount point.
- Type: backup object type (for example filesystem or Oracle).

backup owner
Each backup object in the IDB has an owner. The default owner of a backup is the user who starts the backup session.

backup session
A process that creates a copy of data on storage media. The activities are specified in a backup specification or an interactive session. All clients configured in one backup specification are backed up together in one backup session using the same backup type (full or incremental). The result of a backup session is a set of media, which was written to, also called the backup or media set.
See also **incremental backup** and **full backup.**

backup set
A complete set of integration objects associated with a backup.

backup set *(Oracle specific term)*
A logical grouping of backed up files that are created using the RMAN backup command. A backup set is a complete set of files associated with a backup. The files can be multiplexed to increase performance. A backup set contains either datafiles or archived logs, but not both together.

backup specification
A list of objects to be backed up, together with a set of devices or drives to be used, backup options for all objects in the specification, days and time that you want backups to be performed. The objects are entire disks/volumes or parts of them such as files, directories, or even the Windows

Registry for example. File selection lists such as include-lists and exclude-lists can be specified.

backup system *(ZDB specific term)*
A system connected to target volumes of one or multiple application systems. The backup system is typically connected to a backup device to perform the backup of the data in a replica.
See also **application system**, **target volume**, and **replica**.

backup types
See **incremental backup**, **differential backup**, **transaction backup**, **full backup** and **delta backup**.

backup view
Data Protector provides different views for backup specifications: By Type - according to the type of data available for backups/templates. Default view. By Group - according to the group to which backup specifications/templates belong. By Name - according to the name of backup specifications/templates. By Manager - if you are running MoM, you can also set the Backup view according to the Cell Manager to which backup specifications/templates belong.

BC *(EMC Symmetrix specific term)*
Business Continuance are processes that allow customers to access and manage instant copies of EMC Symmetrix standard devices.
See also **BCV**.

BC *(HP StorageWorks Disk Array XP specific term)*
The Business Copy XP allows to maintain internal copies of HP StorageWorks Disk Array XP LDEVs for purposes such as data backup or data duplication. The copies (secondary volumes or S-VOLs) can be separated from the primary volumes (P-VOLs) and connected to a different system for various purposes, such as backup and development. For backup purposes, P-VOLs should be connected to the application system, and one of the S-VOL mirror sets to the backup system.
See also **HP StorageWorks Disk Array XP LDEV**, **CA**, **Main Control Unit**, **application system**, and **backup system**.

BC Process *(EMC Symmetrix specific term)*
A protected storage environment solution that has defined specially configured EMC Symmetrix devices as mirrors or Business Continuance Volumes to protect data on EMC Symmetrix standard devices.
See also **BCV**.

BC VA *(HP StorageWorks Virtual Array specific term)*
Business Copy VA allows you to

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

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

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

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

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

BRBACKUP *(SAP R/3 specific term)*
An SAP R/3 backup tool that allows an online or offline backup of the control file, of individual data files, or of all tablespaces and, if necessary, of the online redo log files.
See also **SAPDBA**, **BRARCHIVE** and **BRRESTORE**.

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

- Database data files, control files, and online redo log files saved with BRBACKUP

- Redo log files archived with BRARCHIVE
- Non-database files saved with BRBACKUP

You can specify files, tablespaces, complete backups, log sequence numbers of redo log files, or the session ID of the backup.
See also **SAPDBA**, **BRBACKUP** and **BRARCHIVE**.

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

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

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

catalog protection
Defines how long information about backed up data (such as file names and file versions) is kept in the IDB.
See also **data protection**.

CDB
The Catalog Database is a part of the IDB that contains information about backups, object copies, restores, media management sessions, and backed up data. Depending on the selected logging level, it also contains file names and file versions. This part of the database is always local to the cell.
See also **MMDB**.

CDF file *(UNIX specific term)*
A Context Dependent File is a file consisting of several files grouped under the same pathname. The system ordinarily selects one of the files using the context of the process. This mechanism allows machine dependent executables, system data, and device files to work correctly from all hosts in a cluster while using the same pathname.

cell
A set of systems that are under the control of a Cell Manager. The cell typically represents the systems on a site or an organizational entity, which are connected to the same LAN. Central control is available to administer the backup and restore policies and tasks.

Cell Manager
The main system in the cell where the essential Data Protector software is installed and from which all backup and restore activities are managed. The GUI used for management tasks can be located on a different system. Each cell has one Cell Manager system.

centralized licensing
Data Protector allows you to configure centralized licensing for the whole enterprise environment consisting of several cells. All Data Protector licenses are installed and kept on the Enterprise Cell Manager system. You can then allocate licenses to specific cells to suit your needs.
See also **MoM**.

Centralized Media Management Database (CMMDB)
See **CMMDB**.

channel *(Oracle specific term)*
An Oracle Recovery Manager resource allocation. Every allocated channel starts a new Oracle process, which performs backup, restore, and recovery actions. The type of channel allocated determines the type of media used:

- type “disk”
- type ‘SBT_TAPE’

If the specified channel is type ‘SBT_TAPE’ and Oracle is integrated with Data Protector, the server process will attempt to read backups from or write data files to Data Protector.

circular logging *(Microsoft Exchange Server and Lotus Domino Server specific term)*
Circular logging is a Microsoft Exchange Server database and Lotus Domino Server database mode, in which the transaction log file contents are periodically overwritten after the corresponding data is committed to the database. Circular logging reduces disk storage space requirements.

client backup
A backup of all filesystems mounted on a client. Filesystems mounted on the client after the backup specification was created are not automatically detected.

client backup with disk discovery
A backup of all filesystems mounted on a client. When the backup starts, Data Protector discovers the disks on the clients. Client backup with disk

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

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

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

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

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

COM+ Registration Database
(Windows specific term)
The COM+ Registration Database and the Windows Registry store COM+ application attributes, class attributes, and computer-level attributes. This guarantees consistency among these attributes and provides common operation on top of these attributes.

command-line interface
A set of DOS and UNIX like commands that you can use in shell scripts to perform Data Protector configuration, backup, restore, and management tasks.

Command View (CV) EVA *(HP StorageWorks EVA specific term)*
The user interface that allows you to configure, manage, and monitor your HP StorageWorks EVA storage system. It is used to perform various storage management tasks, for example, creating virtual disk families, managing storage system hardware, creating snapclones and snapshots of virtual disks. The Command View EVA software runs on the HP OpenView

Storage Management Appliance, and is accessed by a Web browser.
See also **HP StorageWorks EVA Agent (legacy)** and **HP StorageWorks EVA SMI-S Agent**.

concurrency
See **Disk Agent concurrency**.

control file *(Oracle and SAP R/3 specific term)*
An Oracle data file that contains entries specifying the physical structure of the database. It provides database consistency information used for recovery.

CRS
The Cell Request Server process (service) runs on the Data Protector Cell Manager. It starts and controls the backup and restore sessions. The service is started as soon as Data Protector is installed on the Cell Manager.
CRS runs under the account root on UNIX systems, and under any Windows account. By default, it runs under the account of the user, specified at installation time.

CSM
The Data Protector Copy Session Manager process controls the object copy session and runs on the Cell Manager system.

data file *(Oracle and SAP R/3* specific term)
A physical file created by Oracle that contains data structures such as tables and indexes. A data file can only belong to one Oracle database.

data protection
Defines how long the backed up data on media remains protected, that is, Data Protector will not overwrite it. When the protection expires, Data Protector will be able to reuse the media in one of the next backup sessions.
See also **catalog protection**.

Data Protector Event Log
A central repository of all Data Protector related notifications. By default, all notifications are sent to the Event Log. The Event Log is accessible only to Data Protector users in the Admin group and to Data Protector users who are granted the Reporting and notifications user rights. You can view or delete all events in the Event Log.

Data Protector user account
You can use Data Protector only if you have a Data Protector user account, which restricts unauthorized access to Data Protector and to backed up data. Data Protector administrators create this account specifying a user logon name, the systems from which the user can log on, and a Data Protector user group

membership. This is checked whenever the user starts the Data Protector user interface or performs specific tasks.

data stream
Sequence of data transferred over the communication channel.

database library
A Data Protector set of routines that enables data transfer between Data Protector and a server of an online database integration, for example, the Oracle Server.

database parallelism
More than one database is backed up at a time if the number of available devices allows you to perform backups in parallel.

database server
A computer with a large database stored on it, such as the SAP R/3 or Microsoft SQL database. A server has a database that can be accessed by clients.

Dbobject *(Informix specific term)*
An Informix physical database object. It can be a blobspace, dbspace, or logical-log file.

DC directory
The Detail Catalog (DC) directory consists of DC binary files, which store information about file versions. It represents the DCBF part of the IDB, which occupies approximately 80% of the IDB. The default DC directory is called the dcbf directory and is located in the <*Data_Protector_home*>\db40 directory on a Windows Cell Manager and in the /var/opt/omni/server/db40 directory on a UNIX Cell Manager. You can create more DC directories and locate them as appropriate to you. Up to 10 DC directories are supported per cell. The default maximum size of a DC directory is 4 GB.

DCBF
The Detail Catalog Binary Files (DCBF) part of the IDB stores information about file versions and attributes. It occupies approximately 80% of the IDB. One DC binary file is created for each Data Protector medium used for backup.

delta backup
A delta backup is a backup containing all the changes made to the database from the last backup of any type.
See also **backup types**

device
A physical unit which contains either just a drive or a more complex unit such as a library.

device chain
A device chain consists of several standalone devices configured for sequential use. When a medium in one

device gets full, the backup automatically continues on a medium in the next device in the device chain.

device group *(EMC Symmetrix specific term*)
A logical unit representing several EMC Symmetrix devices. A device cannot belong to more than a single device group. All devices in a device group must be on the same EMC Symmetrix unit. You can use a device group to identify and work with a subset of the available EMC Symmetrix devices.

device streaming
A device is streaming if it can feed enough data to the medium to keep it moving forward continuously. Otherwise, the tape has to be stopped, the device waits for more data, reverses the tape a little and resumes to write to the tape, and so on. In other words, if the data rate written to the tape is less or equal the data rate which can be delivered to the device by the computer system, then the device is streaming. Streaming significantly improves the performance of the device and use of space.

DHCP server
A system running the Dynamic Host Configuration Protocol (DHCP) providing dynamic configuration of IP addresses and related information.

differential backup
An incremental backup (incr) based on any previous Data Protector backup (full or any incremental), which must still be protected.
See **incremental backup**.

differential backup *(MS SQL specific term)*
A database backup that records only the data changes made to the database after the last full database backup.
See also **backup types**.

differential database backup
A differential database backup records only those data changes made to the database after the last full database backup.

direct backup
A SAN-based backup solution in which data movement directly from disk to tape (or to other secondary storage) is facilitated by the SCSI Extended Copy (Xcopy) command. Direct backup lessens the backup I/O load on systems in a SAN environment. The data movement is facilitated directly from disk to tape (or to other secondary storage) by the SCSI Extended Copy (XCopy) command. The command is provided by any element of the infrastructure including bridges, switches, tape libraries, and disk subsystems.
See also **XCopy engine**.

directory junction *(Windows specific term)*
Directory junctions use the reparse point concept of Windows. An NTFS 5 directory junction allows you to redirect a directory/file request to another location.

Directory Store (DS) *(Microsoft Exchange specific term)*
A part of the Microsoft Exchange Server directory. The Microsoft Exchange Server directory contains objects used by Microsoft Exchange applications in order to find and access services, mailboxes, recipients, public folders, and other addressable objects within the messaging system.
See also **Information Store (MDB)**.

disaster recovery
A process to restore a client's main system disk to a state close to the time when a (full) backup was performed.

Disk Agent
A component needed on a client to back it up and restore it. The Disk Agent controls reading from and writing to a disk. During a backup session, the Disk Agent reads data from a disk and sends it to the Media Agent, which then moves it to the device. During a restore session the Disk Agent receives data from the Media Agent and writes it to the disk.

Disk Agent concurrency
The number of Disk Agents that are allowed to send data to one Media Agent concurrently.

disk discovery
The detection of disks during client backup with disk discovery. During this backup, Data Protector discovers (detects) the disks that are present on the client — even though they might not have been present on the system when the backup was configured — and backs them up. This is particularly useful in dynamic environments, where configurations change rapidly. After the disks are expanded, each inherits all options from its master client object. Even if pre- and post-exec commands are specified once, they are started many times, once per each object.

disk group *(Veritas Volume Manager specific term)*
The basic unit of data storage in VxVM system. A disk group can consist of one or more physical volumes. There can be more than one disk group on the system.

disk image (rawdisk) backup
A high-speed backup where Data Protector backs up files as bitmap images. A disk image (rawdisk) backup does not track the files and directory structure stored on the disk, but stores a disk image structure on byte level. You

can perform a disk image backup of either specific disk sections or a complete disk.

disk quota
A concept to manage disk space consumption for all or a subset of users on a computer system. This concept is used by several operating system platforms.

disk staging
The process of backing up data in several phases to improve the performance of backups and restores, reduce costs of storing the backed up data, and increase the data availability and accessibility for restore. The backup stages consist of backing up data to one media type first (for example disk) and later copying it to a different media type (for example tape).

Distributed File System (DFS)
A service that connects file shares into a single namespace. The file shares can reside on the same or on different computers. DFS provides client access to the resources in a location-transparent manner.

DMZ
The Demilitarized Zone (DMZ) is a network inserted as a "neutral zone" between a company's private network (intranet) and the outside public network (Internet). It prevents outside users from getting direct access to company servers in the intranet.

DNS server
In the DNS client-server model, this is the server containing information about a portion of the DNS database that makes computer names available to client resolvers querying for name resolution across the Internet.

domain controller
A server in a network that is responsible for user security and verifying passwords within a group of other servers.

DR image
Data required for temporary disaster recovery operating system (DR OS) installation and configuration.

DR OS
A disaster recovery operating system is an operating system environment in which disaster recovery runs. It provides Data Protector a basic runtime environment (disk, network, tape, and filesystem access). The OS has to be installed and configured before the Data Protector disaster recovery can be performed. DR OS not only hosts the Data Protector disaster recovery process but is also a part of the restored system because it replaces its own configuration data with the original configuration data.

drive
A physical unit that receives data from a computer system and can write it onto a magnetic medium (typically a tape drive). It can also read the data from the medium and send it to the computer system.

drive index
A number that identifies the mechanical position of a drive inside a library device. This number is used by the robotic control to access a drive.

dynamic client
See **client backup with disk discovery**.

EMC Symmetrix Agent (SYMA)

(EMC Symmetrix specific term)
See **Symmetrix Agent (SYMA)**

emergency boot file *(Informix specific term)*
An Informix configuration file that resides in the <*INFORMIXDIR*>\etc directory (on HP-UX) or <*INFORMIXDIR*>/etc directory (on Windows) and is called ixbar.<*server_id*>, where <*INFORMIXDIR*> is the OnLine Server home directory and <*server_id*> is the value of the SERVERNUM configuration parameter. Each line of the emergency boot file corresponds to one backup object.

Enterprise Backup Environment
Several cells can be grouped together and managed from a central cell. The enterprise backup environment includes all clients located in several Data Protector cells which are managed and administered from a central cell using the Manager-of-Managers concept.
See also **MoM**.

Event Logs
Files in which Windows logs all events, such as the starting or stopping of services and the logging on and off of users. Data Protector can back up Windows Event Logs as part of the Windows configuration backup.

exchanger
Also referred to as SCSI Exchanger.
See also **library**.

exporting media
A process that removes all data about backup sessions, such as systems, objects, and file names, which reside on the media from the IDB. Information about the media and their relation to a pool is also removed from the IDB. The data on the media remains unchanged.
See also **importing media.**

Extensible Storage Engine (ESE)

(Microsoft Exchange Server 2000/2003 specific term)
A database technology used as a storage system for information exchange in Microsoft Exchange Server 2000/2003.

failover
Transferring of the most important cluster data, called group (on Windows) or package (on Unix) from one cluster node to another. A failover can occur mostly because of software or hardware failures or maintenance on the primary node.

FC bridge
See **Fibre Channel bridge**

Fibre Channel
An ANSI standard for high-speed computer interconnection. Using either optical or copper cables, it allows the high speed bidirectional transmission of large data files and can be deployed between sites kilometers apart.
Fibre Channel connects nodes using three physical topologies: point-to-point, loop, and switched.

Fibre Channel bridge
A Fibre Channel bridge or multiplexer provides the ability to migrate existing parallel SCSI devices, like RAID arrays, solid state disks (SSD), and tape libraries to a Fibre Channel environment. On one side of the bridge or multiplexer there is a Fibre Channel interface while on the other side there are parallel SCSI ports. The bridge or multiplexer enables SCSI packets to be moved between the Fibre Channel and parallel SCSI devices.

file depot
A file containing the data from a backup to a file library device.

file jukebox device
A device residing on disk consisting of multiple slots used to store file media.

file library device
A device which resides on a disk emulating a library with several media, hence containing multiple files, referred to as file depots.

File Replication Service (FRS)
A Windows service that replicates the domain controller store logon scripts and group policies. FRS also enables replication of Distributed File System (DFS) shares between systems and allows any server to perform replication activity.

file version
The same file can be backed up multiple times in case of full backups and incremental backups (if the file changed). If the log level ALL is selected for backup, Data Protector

retains one entry in the IDB for the filename itself and one for each version (date/time) of the file.

filesystem
The organization of files on a hard disk. A filesystem is backed up so that the file attributes and the file contents are stored on the backup media.

first level mirror (*HP StorageWorks Disk Array XP specific term*)
HP StorageWorks Disk Array XP allows up to three mirror copies of a Primary Volume and each of these copies can have additional two copies. The three mirror copies are called first level mirrors.
See also **Primary Volume**, and **MU numbers**.

fnames.dat
The fnames.dat files of the IDB contain information on the names of the backed up files. Typically, these files occupy about 20% of the IDB, if filenames are stored.

formatting
A process that erases any data contained on a medium and prepares it for use with Data Protector. Information about media (media ID, description, and location) is saved in the IDB as well as on the respective media (media header). Data Protector media with protected data are not formatted until the protection expires or the media are unprotected/recycled.

free pool
An auxiliary source of media for use by media pools when they run out of media. The media pools must be configured to use free pools.

full backup
A backup in which all selected objects are backed up, whether or not they have been recently modified.
See also **backup types**.

full database backup
A backup of all data in a database, not only the data that has been changed after the last (full or incremental) backup of the database. A full database backup does not depend on any other backup.

full mailbox backup
A full mailbox backup is a backup of the entire mailbox content.

global options file
A file that allows you to customize Data Protector. It explains the global options, which cover various aspects of Data Protector, typically time-outs and limits, and affect the entire Data Protector cell. The file is located in the /etc/opt/omni/server/options directory on HP-UX and Solaris systems and in the

<Data_Protector_home>\Config\Server\Options directory on Windows systems.

group *(Microsoft Cluster Server specific term*)
A collection of resources (for example disk volumes, application services, IP names and addresses) that are needed to run a specific cluster-aware applications.

GUI
A cross-platform (HP-UX, Solaris, and Windows) graphical user interface, provided by Data Protector for easy access to all configuration, administration, and operation tasks.

hard recovery *(Microsoft Exchange Server specific term)*
A Microsoft Exchange Server database recovery that is performed after a restore by the database engine, using transaction log files.

heartbeat
A cluster data set with a time stamp carrying information about the operational status of a particular cluster node. This data set or packet is distributed among all cluster nodes.

Hierarchical Storage Management (HSM)
A method for optimizing the use of expensive hard disk storage by migrating less frequently used data to less expensive optical platters. When needed, the data is migrated back to hard disk storage. This balances the need for fast retrieval from hard disk with the lower cost of optical platters.

Holidays file
A file that contains information about holidays. You can set different holidays by editing the Holidays file: /etc/opt/omni/server/Holidays on the UNIX Cell Manager and *<Data_Protector_home>*\Config\Server\holidays on the Windows Cell Manager.

host backup
See **client backup with disk discovery**.

hosting system
A working Data Protector client used for Disk Delivery Disaster Recovery with a Data Protector Disk Agent installed.

HP ITO
See **OVO**.

HP OpC
See **OVO**.

HP OpenView SMART Plug-In (SPI)
A fully integrated, out-of-the-box solution which "plugs into" HP OpenView Operations, extending the managed domain. Through the Data Protector integration, which is implemented as an HP OpenView

SMART Plug-In, a user can have an arbitrary number of Data Protector Cell Managers monitored as an extension to HP OpenView Operations (OVO).

HP OVO
See **OVO**.

HP StorageWorks Disk Array XP LDEV
A logical partition of a physical disk within an HP StorageWorks Disk Array XP. LDEVs are entities that can be replicated in the Continuous Access XP (CA) and Business Copy XP (BC) configurations, or can be used as standalone entities.
See also **BC** *(HP StorageWorks Disk Array XP specific term)*, **CA** *(HP StorageWorks Disk Array XP specific term)*, and **replica**.

HP StorageWorks EVA Agent (legacy)
A Data Protector software module that executes all tasks required for the HP StorageWorks Enterprise Virtual Array integration operating on HP StorageWorks EVA with Command View (CV) EVA software v3.1 or lower, and the EVA VCS firmware v3.01x or lower.
See also **Command View (CV) EVA** and **HP StorageWorks EVA SMI-S Agent**.

HP StorageWorks EVA SMI-S Agent
A Data Protector software module that executes all tasks required for the HP StorageWorks Enterprise Virtual Array integration operating on HP StorageWorks EVA with Command View (CV) EVA software starting with v3.2. With the EVA SMI-S Agent, the control over the array is established through HP StorageWorks SMI-S EVA provider, which directs communication between incoming requests and CV EVA.
See also **Command View (CV) EVA**, **HP StorageWorks SMI-S EVA provider**, and **HP StorageWorks EVA Agent (legacy)**.

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

S Agent and **Command View (CV) EVA**.

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

HP VPO
See **OVO**.

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

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

importing media
A process that re-reads all data about backup sessions which are on the medium back into the IDB. This then allows for fast and convenient access to the data on the media.
See also **exporting media.**

incremental backup
A backup that selects only files that have changed since a previous backup. Several levels of incremental backup are available, allowing selective backup of only files that have changed since the last incremental backup.
See also **backup types**.

incremental backup *(Microsoft Exchange Server specific term)*
A backup of the Microsoft Exchange Server data that has changed since the last full or incremental backup. With the incremental backup, only the transaction log files are backed up.
See also **backup types**.

incremental mailbox backup
An incremental mailbox backup backs up all the changes made to the mailbox after the last backup of any type.

incremental1 mailbox backup
An incremental1 mailbox backup backs up all the changes made to the mailbox after the last full backup.

incremental (re)-establish *(EMC Symmetrix specific term)*
A BCV or SRDF control operation. In BCV control operations, an incremental establish causes the BCV device to be synchronized incrementally and to function as an EMC Symmetrix mirrored medium. The EMC Symmetrix devices must have been previously paired.
In SRDF control operations, an incremental establish causes the target (R2) device to be synchronized incrementally and to function as an EMC Symmetrix mirrored medium. The EMC Symmetrix devices must have been previously paired.

incremental restore *(EMC Symmetrix specific term)*
A BCV or SRDF control operation.
In BCV control operations, an incremental restore reassigns a BCV device as the next available mirror of the standard device in the pair. However, the standard devices are updated with only the data that was written to the BCV device during the time of the original pair split, and the data that was written to the standard device during the split is overwritten with data from the BCV mirror. In SRDF control operations, an incremental restore reassigns a target (R2) device as the next available mirror of the source (R1) device in the pair. However, the source (R1) devices are updated with only the data that was written to the target (R2) device during the time of the original pair split, and the data that was written to the source (R1) device during the split is overwritten with data from the target (R2) mirror.

Inet
A process that runs on each UNIX system or service that runs on each Windows system in the Data Protector cell. It is responsible for communication between systems in the cell and for starting other processes needed for backup and restore. The Inet service is started as soon as Data Protector is installed on a system. The Inet process is started by the inetd daemon.

Information Store *(Microsoft Exchange Server 2000/2003 specific term)*
The Microsoft Exchange Server 2000/ 2003 service that is responsible for storage management. Information Store in Microsoft Exchange Server 2000/ 2003 manages two kinds of stores: mailbox stores and public folder stores. A mailbox store consists of mailboxes that belong to individual users. A public folder store contains public folders and messages that are shared among several users.
See also **Key Management Service** and **Site Replication Service**.

Information Store *(Microsoft Exchange Server 5.5 specific term)*
This is the default message store provider for the Microsoft Exchange Server 5.5. Information Store consists of the following stores:

- public information store
- private information store
- personal folder store
- offline information store.

The public information store contains public folders and messages that can be shared among multiple users and applications. A single public store is shared by all users within an Exchange Server 5.5 organization, even if multiple Exchange Servers are used. The private information store consists of mail boxes that can belong to users or to applications. The mail boxes reside on the server running the Exchange Server 5.5.
See also **Directory Store (DS)**.

initializing
See **formatting**.

Installation Server
A computer system that holds a repository of the Data Protector software packages for a specific architecture. The Installation Server is used for remote installation of Data Protector clients. In mixed environments at least two Installation Servers are needed: one for UNIX systems and one for Windows systems.

instant recovery (*ZDB specific term*)
A process in which a replica, produced by a ZDB-to-disk or a ZDB-to-disk+tape sessions, is used to restore the contents of the source volumes to their states at the time at which the replica was created, avoiding the need to perform a restore from tape. Depending on the application/database concerned, this may be all that is required, or other steps, such as the application of transaction log files, may be required for full recovery.
See also **replica**, **zero downtime backup (ZDB)**, **ZDB to disk**, and **ZDB to disk+tape**.

integrated security *(MS SQL specific term)*
Integrated security allows the Microsoft SQL Server to use Windows authentication mechanisms to validate Microsoft SQL Server logins for all connections. Using integrated security means that users have one password for both Windows and Microsoft SQL Server. Integrated security should be used in environments where all clients support trusted connections. Connections validated by Windows Server and accepted by Microsoft SQL

Server are referred to as trusted connections. Only trusted connections are allowed.

integration object
A backup object of a Data Protector integration, such as Oracle or SAP DB.

Internet Information Server (IIS)

(Windows specific term)
Microsoft Internet Information Server is a network file and application server that supports multiple protocols. Primarily, IIS transmits information in Hypertext Markup Language (HTML) pages by using the Hypertext Transport Protocol (HTTP).

IP address
Internet Protocol address is a numeric address of a system used to uniquely identify the system on the network. The IP address consists of four groups of numbers separated by periods (full stops).

ISQL *(Sybase specific term)*
A Sybase utility used to perform system administration tasks on Sybase SQL Server.

ITO
See **OVO**.

jukebox
See **library**.

jukebox device
A device consisting of multiple slots used to store either optical or file media. When being used to store file media, the jukebox device is known as the “file jukebox device”.

Key Management Service *(Microsoft Exchange Server 2000/2003 specific term)*
The Microsoft Exchange Server 2000/ 2003 service that provides encryption functionality for enhanced security.
See also **Information Store** and **Site Replication Service**.

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

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

lights-out operation or **unattended operation**
A backup or restore operation that takes

place outside of normal business hours without an operator. This implies that no operator personnel is present to work with the backup application or service mount requests, for example.

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

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

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

lock name
You can configure the same physical device several times with different characteristics, by using different device names.

The lock name is a user specified string that is used for locking all such device configurations to prevent collision if several such devices (device names) are used concurrently. Use an identical lock name for all device definitions which use the same physical device.

log_full shell script *(Informix UNIX specific term)*
A script provided by ON-Bar that you can use to start backing up logical-log files when OnLine Server issues a log-full event alarm. The Informix ALARMPROGRAM configuration parameter defaults to the <*INFORMIXDIR*>/etc/log_full.sh, where <*INFORMIXDIR*> is the OnLine Server home directory. If you do not want logical logs to be backed up continuously, set the ALARMPROGRAM configuration parameter to <*INFORMIXDIR*>/etc/no_log.sh.

logging level
The logging level determines the amount of details on files and directories written to the IDB during backup or object copying. You can always restore your data, regardless of the logging level used during backup. Data Protector provides four logging levels: Log All, Log Directories, Log Files, and No Log. The different logging level settings influence the IDB growth, backup speed, and the convenience of browsing data for restore.

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

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

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

- <user_name> is the name by which a user is known to Oracle Server and to other users. Every user name is associated with a password and both have to be entered to connect to an Oracle Target Database. This user must have been granted Oracle SYSDBA or SYSOPER rights.

- <password> is a string used for data security and known only to its owner. Passwords are entered to connect to an operating system or software application. The password has to be the same as the password specified in the Oracle password file (orapwd), which is used for authentication of users performing database administration.

- <service> is the name used to identify an SQL*Net server process for the target database.

login information to the Recovery Catalog Database *(Oracle specific term)*
The format of the login information to the Recovery (Oracle) Catalog Database is <user_name>/

<password>@<service>, where the description of the user name, password, and service name is the same as in the Oracle SQL*Net V2 login information to the Oracle target database. In this case, <service> is the name of the service to the Recovery Catalog Database, not the Oracle target database.

Note that the Oracle user specified here has to be the owner of the Oracle Recovery (Oracle) Catalog.

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

LVM
A Logical Volume Manager is a subsystem for structuring and mapping physical disk space to logical volumes on UNIX systems. An LVM system consists of several volume groups, where each volume group has several volumes.

Magic Packet
See **Wake ONLAN**.

mailbox (*Microsoft Exchange Server specific term)*
The location to which e-mail is delivered, which is set up by the administrator for each user. If a set of personal folders is designated as the e-mail delivery location, e-mail is routed from the mailbox to this location.

Mailbox Store (*Microsoft Exchange Server 2000/2003 specific term)*
A part of the Information Store that maintains information about user mailboxes. A mailbox store consists of a binary rich-text .edb file and a streaming native internet content .stm file.

Main Control Unit (MCU) *(HP StorageWorks Disk Array XP specific term)*
An HP StorageWorks XP disk array that contains the primary volumes for the Continuous Access configuration and acts as a master device.
See also **BC** *(HP StorageWorks Disk Array XP specific term)*, **CA** *(HP StorageWorks Disk Array XP specific term)*, and **HP StorageWorks Disk Array XP LDEV**.

Manager-of-Managers (MoM)
See **Enterprise Cell Manager**.

Media Agent
A process that controls reading from and writing to a device, which reads from or writes to a medium (typically a tape). During a backup session, a Media Agent receives data from the Disk Agent and sends it to the device for writing it to the medium. During a restore session, a Media Agent locates data on the backup

medium and sends it to the Disk Agent. The Disk Agent then writes the data to the disk. A Media Agent also manages the robotics control of a library.

MAPI *(Microsoft Exchange specific term)*
The MAPI (Messaging Application Programming Interface) is the programming interface that lets applications and messaging clients interact with messaging and information systems.

media allocation policy
Determines in which sequence media are used for backup. The Strict allocation policy directs Data Protector to prompt for a specific medium. The Loose policy directs Data Protector to prompt for any suitable medium. The Formatted First policy directs Data Protector to give preference to unknown media, even if unprotected media are available in the library.

media condition
The quality of a medium as derived from the media condition factors. Heavy usage and age result in an increased number of read and write errors with tape media. Media need to be replaced when they are marked as POOR.

media condition factors
The user-assigned age threshold and overwrite threshold used to determine the state of a medium.

media ID
A unique identifier assigned to a medium by Data Protector.

media label
A user-defined identifier used to describe a medium.

media location
A user-defined physical location of a medium, such as "building 4" or "off-site storage".

media management session
A session performing some action on a medium, such as initializing, scanning the content, verifying data on a medium, or copying a medium.

media pool
A set of media of the same type (such as DDS) used and tracked as a group. Media are formatted and assigned to a media pool.

media set
The result of a backup session is data backed up on a group of media called media set. Depending on the media usage policy, several sessions can share the same media.

media type
The physical type of media, such as DDS or DLT.

media usage policy
The media usage policy controls how new backups are added to the already used media. It can be Appendable, Non-Appendable, or Appendable for incrementals only.

merging
This defines one mode to resolve file conflicts during restore. If the file to be restored already exists at the destination, the one with the more recent modification date is kept. Files not present on the disk are always restored. *See also* **overwrite**.

MFS
The Migrating File System enables a standard JFS filesystem with migration capabilities (on HP-UX 11.00). The MFS is accessed via a standard filesystem interface (DMAPI), it is mounted to a directory the same way as any HP-UX filesystem. In an MFS, only the superblock, the inode and the 'extended attribute' information remain permanently on the hard disk and are never migrated.
See also **VBFS**.

Microsoft Exchange Server
A “client-server” messaging and a workgroup system that offers a transparent connection to many different communication systems. It provides users with an electronic mail system, individual and group scheduling, online forms, and workflow automation tools. It provides a developer with a platform on which to build custom information-sharing and messaging-service applications.

Microsoft Management Console (MMC) *(Windows specific term)*
An administration model for Windows-based environments. It provides a simple, consistent, and integrated administration user interface allowing management of many applications through the same GUI, provided that the applications adhere to the MMC model.

Microsoft SQL Server 7.0/2000
A database management system designed to meet the requirements of distributed "client-server" computing.

Microsoft Volume Shadow Copy service (VSS)
A software service that provides a unified communication interface to coordinate backup and restore of a VSS-aware application regardless of its specific features. This service collaborates with the backup application, writers, shadow copy providers, and the operating system kernel to implement the management of volume shadow copies and shadow copy

sets.
See also **shadow copy, shadow copy provider, writer**.

mirror *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)*
See **target volume**.

mirror rotation *(HP StorageWorks Disk Array XP specific term)*
See **replica set rotation**.

MMD
The Media Management Daemon process (service) runs on the Data Protector Cell Manager and controls media management and device operations. The process is started when Data Protector is installed on the Cell Manager.

MMDB
The Media Management Database (MMDB) is a part of the IDB that contains information about media, media pools, devices, libraries, library drives, and slots configured in the cell, as well as the Data Protector media used for backup. In an enterprise backup environment, this part of the database can be common to all cells.
See also **CMMDB**, **CDB**.

MoM
Several cells can be grouped together and managed from a central cell. The management system of the central cell is the Manager-of-Managers (MoM). The MoM allows you to configure and manage multiple cells from a central point.

mount request
A screen prompt that tells you to insert a specific medium into a device. Once you respond to the mount request by providing the required medium and confirm the mount request, the session continues.

mount point
The access point in a directory structure for a disk or logical volume, for example /opt or d:. On UNIX, the mount points are displayed using the bdf or df command.

MSM
The Data Protector Media Session Manager, which runs on the Cell Manager and controls media sessions, such as copying media.

MU number (*HP StorageWorks Disk Array XP specific term*)
A Mirror Unit number is an integer number (0, 1 or 2), used to indicate a first level mirror.
See also **first level mirror**.

multi-drive server
A license that allows you to run an unlimited number of Media Agents on a

single system. This license, which is bound to the IP address of the Cell Manager, is no longer available.

obdrindex.dat
An IDB file with information about IDB backups, media, and devices used for the backup. This information can significantly simplify IDB recovery. It is recommended to relocate the file, together with IDB transaction logs, on a separate physical disk from other IDB directories, and, additionally, to make a copy of the file and locate it where you want.

OBDR capable device
A device that can emulate a CD-ROM drive loaded with a bootable disk and can thus be used as a backup or boot device for disaster recovery purposes.

object
See **backup object**

object copy
A copy of a specific object version that is created during an object copy session or a backup session with object mirroring.

object copy session
A process that creates an additional copy of the backed up data on a different media set. During an object copy session, the selected backed up objects are copied from the source to the target media.

object copying
The process of copying selected object versions to a specific media set. You can select object versions from one or several backup sessions to be copied.

Object ID *(Windows specific term)*
The object IDs (OIDs) enable access to NTFS 5 files no matter where in the system the files reside. Data Protector treats the OIDs as alternate streams of the files.

object mirror
A copy of a backup object created using object mirroring. Object mirrors are often referred to as object copies.

object mirroring
The process of writing the same data to several media sets during a backup session. Data Protector enables you to mirror all or some backup objects to one or more media sets.

offline backup
A backup during which an application database cannot be used by the application.

- For simple backup methods (non ZDB), the database is generally put into a quiescent state that allows use

by the backup system, but not the application, for the whole backup period (~minutes/hours). For instance, for backup to tape, until streaming of data to the tape is finished.

- For ZDB methods, the database is also put into the quiescent state, but for the period of the data replication process only (~seconds). Normal database operation can then be resumed for the rest of the backup process.

See also **zero downtime backup (ZDB)** and **online backup**.

offline recovery
Offline recovery is performed if the Cell Manager is not accessible, for example, due to network problems. Only standalone and SCSI library devices can be used for offline recovery. Recovery of the Cell Manager is always offline.

offline redo log
See **archived redo log**

OmniStorage
Software providing transparent migration of less frequently used data to the optical library while keeping more frequently used data on the hard disk. HP OmniStorage runs on HP-UX systems.

On-Bar *(Informix specific term)*
A backup and restore system for OnLine Server. ON-Bar enables you to create a copy of your OnLine Server data and later restore the data. The ON-Bar backup and restore system involves the following components:

- onbar utility
- Data Protector, as the backup solution
- XBSA interface
- ON-Bar catalog tables, which are used to back up dbobjects and track instances of dbobjects through multiple backups.

onbar utility *(Informix specific term)*
The Informix utility that communicates backup and restore requests to OnLine Server. The utility uses XBSA to exchange control data and back up and restore data with Data Protector.

ONCONFIG *(Informix specific term)*
An environment variable that specifies the name of the active ONCONFIG configuration file. If the ONCONFIG environment variable is not present, OnLine uses the configuration values from the file <*INFORMIXDIR*>/etc/onconfig (on HP-UX) or <*INFORMIXDIR*>\etc\onconfig (on Windows).

online backup
A backup performed while a database application remains available for use. The database is placed into a special backup mode of operation for the time period that the backup application requires access to the original data objects. During this period, the database is fully operational, but there may be a small performance impact and log files may grow very quickly.

- For simple backup methods (non ZDB), backup mode is required for the whole backup period (~minutes/ hours). For instance, for backup to tape, until streaming of data to tape is finished.

- For ZDB methods, backup mode is required for the short period of the data replication process only (~seconds). Normal database operation can then be resumed for the rest of the backup process.

In some cases, transaction logs may also have to be backed up to allow a consistent database to be restored.
See also **zero downtime backup (ZDB)** and **offline backup**.

online redo log *(Oracle specific term)*
Redo logs that have not been archived, but are either available to the instance for recording database activity or are filled and waiting to be archived or reused.
See also **archived redo log**.

OnLine Server *(Informix specific term)*
Refers to INFORMIX-OnLine Dynamic Server.

OpC
See **OVO**.

Oracle instance *(Oracle specific term)*
Each installation of an Oracle database on one or more systems. One computer system can have several instances of a database running.

ORACLE_SID *(Oracle specific term)*
A unique name for an Oracle Server instance. To switch among Oracle Servers, specify the desired <*ORACLE_SID*>. The <*ORACLE_SID*> is included in the CONNECT DATA parts of the connect descriptor in a TNSNAMES.ORA file and in the definition of the TNS listener in the LISTENER.ORA file.

original system
The system configuration backed up by Data Protector before a computer disaster hits the system.

overwrite
An option that defines one mode to resolve file conflicts during restore. All

files are restored from a backup even if they are older than existing files.
See also **merging**.

OVO
HP OpenView Operations for Unix provides powerful capabilities for operations management of a large number of systems and applications on in a network. Data Protector provides an integration into this management product. This integration is implemented as a SMART Plug-In for OVO management servers on HP-UX and Solaris. Earlier versions of OVO were called IT/Operation, Operations Center and Vantage Point Operations.
See also **merging**.

ownership
The ownership of a backup determines who can restore from the backup. The user who starts an interactive backup is the session owner. If a user starts an existing backup specification without modifying it, the session is not considered interactive. In that case, if the backup owner has been defined in the backup specification, they remain the session owner. Otherwise, the session owner becomes the user who started the backup in question. For the scheduled backups, by default, the session owner is for the UNIX Cell Manager: root.sys@<Cell Manager>, and for the Windows Cell Manager, the user that was specified during the installation of the Cell Manager. It is possible to modify the ownership, so that the specific user becomes the session owner.

P1S file
P1S file contains information on how to format and partition all disks installed in the system during Enhanced Automated Disaster Recovery (EADR). It is created during full backup and is saved on backup medium and on Cell Manager into <*Data_Protector_home*>\Config\Server\dr\p1s directory on a Windows Cell Manager or in /etc/opt/omni/server/dr/p1s directory on a UNIX Cell Manager with the filename recovery.p1s.

package *(MC/ServiceGuard and Veritas Cluster specific term)*
A collection of resources (for example volume groups, application services, IP names and addresses) that are needed to run a specific cluster-aware application.

pair status *(HP StorageWorks Disk Array XP specific term*)
A mirrored pair of disks can have various status values depending on the action performed on it. The three most important status values are:

- COPY - The mirrored pair is currently resynchronizing. Data is transferred from one disk to the other. The disks do not contain the same data.

- PAIR - The mirrored pair is completely synchronized and both disks (the primary volume and the mirrored volume) contain identical data.

- SUSPENDED - The link between the mirrored disks is suspended. That means that both disks are accessed and updated independently. However, the mirror relationship is still maintained and the pair can be resynchronized without transferring the complete disk.

parallel restore
Restoring backed up data to multiple disks at the same time (that is, in parallel) by running multiple Disk Agents, that receive data from one Media Agent. For the parallel restore to work, select data that is located on different disks or logical volumes and during backup, the data from the different objects must have been sent to the same device using a concurrency of 2 or more. During a parallel restore, the data for multiple objects selected for restore is read from media at the same time, thereby improving performance.

parallelism
The concept of reading multiple data streams from an online database.

physical device
A physical unit that contains either a drive or a more complex unit such as a library.

post-exec
A backup option that executes a command or script after the backup of an object or after the entire session completes. Post-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows and as shell scripts on UNIX.
See also **pre-exec**.

pre- and post-exec commands
Pre- and post-exec commands are used to perform additional action before and after a backup or restore session. They are not supplied by Data Protector. You need to create your own commands. They can be written as executables or batch files on Windows and as shell scripts on UNIX.

prealloc list
A subset of media in a media pool that specifies the order in which media are used for backup.

pre-exec
A backup option that executes a command or script before the backup of an object or before the entire session is started. Pre-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows and as shell scripts on UNIX.
See also **post-exec**.

Primary Volume (P-VOL) *(HP StorageWorks Disk Array XP specific term)*
Standard HP StorageWorks Disk Array XP LDEVs that act as a primary volume for the CA and BC configurations. The P-VOL is located in the MCU.
See also **Secondary Volume (S-VOL)**.

Private Information Store *(Microsoft Exchange Server 5.5 specific term)*
A part of the Information Store that maintains information in user mailboxes. A mailbox store consists of a binary rich-text .edb file.

protection
See **data protection** and also **catalog protection**.

public folder store (*Microsoft Exchange Server 2000/2003 specific term)*
The part of the Information Store that maintains information in public folders. A public folder store consists of a binary rich-text .edb file and a streaming native internet content .stm file.

public/private backed up data
When configuring a backup, you can select whether the backed up data will be:

- public, that is visible (and accessible for restore) to all Data Protector users
- private, that is, visible (and accessible for restore) only to the owner of the backup and administrators

RAID
Redundant Array of Inexpensive Disks.

RAID Manager Library *(HP StorageWorks Disk Array XP specific term)*
The RAID Manager Library is used internally by Data Protector on Solaris systems to allow access to HP StorageWorks Disk Array XP configuration, status, and performance data and to key HP StorageWorks Disk Array XP features through the use of function calls translated into a sequence of low level SCSI commands.

RAID Manager XP *(HP StorageWorks Disk Array XP specific term)*
The RAID Manager XP application

provides an extensive list of commands to report and control the status of the CA and BC applications. The commands communicate through a RAID Manager instance with the HP StorageWorks Disk Array XP Disk Control Unit. This instance translates the commands into a sequence of low level SCSI commands.

rawdisk backup
See **disk image backup**.

RCU *(HP StorageWorks specific term)*
The Remote Control Unit acts as a slave of an MCU in a CA configuration. In bidirectional configurations, the RCU can act as an MCU.

RDBMS
Relational Database Management System.

RDF1/RDF2 *(EMC Symmetrix specific term)*
A type of SRDF device group. Only RDF devices can be assigned to an RDF group. An RDF1 group type contains source (R1) devices and an RDF2 group type contains target (R2) devices.

RDS
The Raima Database Server process (service) runs on the Data Protector Cell Manager and manages the IDB. The process is started when Data Protector is installed on the Cell Manager.

Recovery Catalog *(Oracle specific term)*
A set of Oracle tables and views that are used by Recovery Manager to store information about Oracle databases. This information is used by Recovery Manager to manage the backup, restore, and recovery of Oracle databases. The recovery catalog contains information about:

- The physical schema of the Oracle target database
- Data file and archived log backup sets
- Data file copies
- Archived Redo Logs
- Stored scripts.

Recovery Catalog Database *(Oracle specific term)*
An Oracle database that contains a recovery catalog schema. You should not store the recovery catalog in your target database.

RecoveryInfo
When backing up Windows configuration files, Data Protector collects the information about the current system configuration (information on disk layout, volume,

and network configuration). This information is needed for disaster recovery.

Recovery Manager (RMAN) *(Oracle specific term)*
An Oracle command-line interface that directs an Oracle Server process to back up, restore, or recover the database it is connected to. RMAN uses either the recovery catalog or the control file to store information about backups. This information can be used later in restore sessions.

recycle
A process that removes the data protection from all backed up data on a medium, allowing Data Protector to overwrite it during one of the next backups. Data that belongs to the same session(s) but resides on other media is also unprotected. Recycling does not actually alter the data on the medium.

redo log *(Oracle specific term)*
Every Oracle database has a set of two or more redo log files. The set of redo log files for the database is known as the database's redo log. Oracle uses the redo log to record all changes made to data.

Remote Control Unit *(HP StorageWorks Disk Array XP specific term)*
The Remote Control Unit (RCU) acts as a slave of an MCU in a CA configuration. In bidirectional configurations, the RCU can act as an MCU.

Removable Storage Management Database *(Windows specific term)*
A Windows service used for managing removable media (such as tapes and disks) and storage devices (libraries). Removable Storage allows applications to access and share the same media resources.

reparse point *(Windows specific term)*
A system-controlled attribute that can be associated with any directory or file. The value of a reparse attribute can have user-controlled data. The format of the data is understood by the application that stored the data and a filesystem filter that was installed to interpret the data and process such files. Whenever the filesystem encounters a file with a reparse point, it attempts to find the filesystem filter associated with the data format.

replica *(ZDB specific term)*
An image, at a particular point in time, of the data in source volumes that contain user-specified backup objects. Depending on the hardware/software with which it is created, the image may be an independent exact duplicate (clone) of the storage blocks at the physical disk level (for example, a split mirror), or a virtual copy (for example, a

snapshot). From a host's perspective, on a basic UNIX or Windows system, the complete physical disk containing a backup object is replicated. However, if a volume manager is used on UNIX, the whole volume/disk group containing a backup object is replicated.
See also **snapshot**, **snapshot creation**, **split mirror**, and **split mirror creation**.

replica set *(ZDB specific term)*
A group of replicas, all created using the same backup specification.
See also **replica** and **replica set rotation**.

replica set rotation *(ZDB specific term)*
The use of a replica set for regular backup production: Each time the same backup specification requiring the use of a replica set is run, a new replica is created and added to the set, until the maximum number for the set is reached. After that, the oldest replica in the set is replaced, maintaining the maximum number of replicas in the set.
See also **replica** and **replica set**.

restore session
A process that copies data from backup media to a client.

RMAN *(Oracle specific term)*
See **Recovery Manager**.

RSM
The Data Protector Restore Session Manager controls the restore session. This process always runs on the Cell Manager system.

RSM *(Windows specific term)*
Removable Storage Manager (RSM) includes a media management service that facilitates communication among applications, robotic changers, and media libraries. It enables multiple applications to share local robotic media libraries and tape or disk drives and to manage removable media.

SAPDBA *(SAP R/3 specific term)*
An SAP R/3 user interface that integrates the BRBACKUP, BRARCHIVE, and BRRESTORE tools.

scan
A function that identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library).

scanning
A function which identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library). It is useful to perform a scan and check the actual media in the

device if someone has manually manipulated media without using Data Protector to eject or enter, for example.

Scheduler
A function that controls when and how often automatic backups occur. By setting up a schedule, you automate the start of backups.

Secondary Volume (S-VOL) *(HP StorageWorks Disk Array XP specific term*)
Secondary Volumes, or S-VOLs, are XP LDEVs that act as a secondary CA or BC mirror of another LDEV (a P-VOL). In the case of CA, S-VOLs can be used as failover devices in a MetroCluster configuration. The S-VOLs are assigned separate SCSI addresses, different from the addresses used by the P-VOLs.
See also **Primary Volume (P-VOL).**

session
See **backup session**, **media management session,** and **restore session**.

session ID
An identifier of a backup, restore, object copy, or media management session, consisting of the date when the session ran and a unique number.

session key
This environment variable for the Pre- and Post-exec script is a Data Protector unique identification of any session, including preview sessions. The session key is not recorded in the database, and it is used for specifying options for the omnimnt, omnistat and omniabort CLI commands.

shadow copy *(MS VSS specific term)*
A volume that represents a duplicate of the original volume at a certain point in time. The data is then backed up from the shadow copy and not from the original volume. The original volume continues to changes as the backup process continues, but the shadow copy of the volume remains constant.
See also **Microsoft Volume Shadow Copy service**.

shadow copy provider *(MS VSS specific term)*
An entity that performs the work on creating and representing the volume shadow copies. Providers own the shadow copy data and expose the shadow copies. Providers can be software (for example, system providers) or hardware (local disks, disk arrays).
See also **shadow copy.**

shadow copy set *(MS VSS specific term)*
A collection of shadow copies created at the same point in time.
See also **shadow copy**.

shared disks
A Windows disk on another system that has been made available to other users on the network. Systems with shared disks can be backed up without a Data Protector Disk Agent installed.

SIBF
The Serverless Integrations Binary Files (SIBF) is a part of the IDB that stores raw NDMP meta data. This data is necessary to perform restore of NDMP objects.

Site Replication Service *(Microsoft Exchange Server 2000/2003 specific term)*
The Microsoft Exchange Server 2000/ 2003 service that permits compatibility with Microsoft Exchange Server 5.5 by emulating the Exchange Server 5.5 directory service.
See also **Information Store** and **Key Management Service**.

slot
A mechanical position in a library. Each slot can hold a medium, such as a DLT tape. Data Protector references each slot with a number. To read a medium, a robotic mechanism moves the medium from a slot into the drive.

SMB
See **split mirror backup**.

SMBF
The Session Messages Binary Files (SMBF) part of the IDB stores session messages generated during backup, object copy, restore, and media management sessions. One binary file is created per session. The files are grouped by year and month.

snapshot *(HP StorageWorks VA and HP StorageWorks EVA specific term)*
A form of replica produced using snapshot creation techniques. A range of snapshot types is available, with different characteristics, depending on the arrays/techniques used. Such replicas are dynamic and may be either virtual copies, still reliant upon the contents of the source volumes, or independent exact duplicates (clones), depending on the snapshot type and the time since creation.
See also **replica** and **snapshot creation**.

snapshot backup *(HP StorageWorks VA and HP StorageWorks EVA specific term)*
See **ZDB to tape**, **ZDB to disk**, and **ZDB to disk+tape**.

snapshot creation *(HP StorageWorks VA and HP StorageWorks EVA specific term)*
A replica creation technique, in which copies of source volumes are created using storage virtualization techniques. The replicas are considered to be created

at one particular point-in-time, without pre-configuration, and are immediately available for use. However background copying processes normally continue after creation.
See also **snapshot**.

source (R1) device *(EMC Symmetrix specific term)*
An EMC Symmetrix device that participates in SRDF operations with a target (R2) device. All writes to this device are mirrored to a target (R2) device in a remote EMC Symmetrix unit. An R1 device must be assigned to an RDF1 group type.
See also **target (R2) device**.

source volume *(ZDB specific term)*
A storage volume containing data to be replicated.

sparse file A file that contains data with portions of empty blocks. Examples are: -A matrix in which some or much of the data contains zeros -files from image applications -high-speed databases If sparse file processing is not enabled during restore, it might be impossible to restore this file.

split mirror *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)*
A replica created using split mirror techniques. Such a replica provides an independent, exact duplicate, or clone, of the contents of the source volumes.
See also **replica** and **split mirror creation**.

split mirror backup *(EMC Symmetrix specific term)*
See **ZDB to tape**.

split mirror backup *(HP StorageWorks Disk Array XP specific term)*
See **ZDB to tape**, **ZDB to disk**, and **ZDB to disk+tape**.

split mirror creation *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)*
A replica creation technique, in which a pre-configured set of target volumes (a mirror) is kept synchronized with a set of source volumes until the time at which a replica of the contents of the source volumes is required. Then, the synchronization is stopped (the mirror is split) and a split mirror replica of the source volumes at the time of the split remains in the target volumes.
See also **split mirror**.

split mirror restore *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)*
A process in which data backed up in a ZDB-to-tape or a ZDB-to-disk+tape session is restored from tape media to a split mirror replica, which is then synchronized to the source volumes. Individual backup objects or complete

sessions can be restored using this method.
See also **ZDB to tape**, **ZDB to disk+tape**, and **replica**.

sqlhosts file *(Informix specific term)*
An Informix connectivity-information file that contains the names of each of the database servers and any aliases to which the clients on a host computer can connect.

SRD file
The Data Protector System Recovery Data (SRD) file contains system information required for installing and configuring the operating system in case of a disaster. The SRD file is an ASCII file, generated when a CONFIGURATION backup is performed on a Windows client and stored on the Cell Manager.

SRDF *(EMC Symmetrix specific term)*
The EMC Symmetrix Remote Data Facility is a business continuation process that enables effective, real-time data replication of SLDs between dislocated processing environments. These environments could be situated within the same root computer environment or separated by long distances.

SSE Agent *(HP StorageWorks Disk Array XP specific term)*
A Data Protector software module that executes all tasks required for a split mirror backup integration. It communicates with the HP StorageWorks Disk Array XP storing system using the RAID Manager XP utility (HP-UX and Windows systems) or RAID Manager Library (Solaris systems).

sst.conf file
The file /usr/kernel/drv/sst.conf is required on each Data Protector Sun Solaris client to which a multi-drive library device is connected. It must contain an entry for the SCSI address of the robotic mechanism of each library device connected to the client.

st.conf file
The file /kernel/drv/st.conf is required on each Data Protector Solaris client with a backup device connected. It must contain device information and a SCSI address for each backup drive connected to that client. A single SCSI entry is required for a single-drive device and multiple SCSI entries are required for a multi-drive library device.

stackers
Devices with multiple slots for media storage usually with only one drive. A stacker selects media from the stack sequentially. In contrast, a library can randomly select media from its repository.

standalone file device
A file device is a file in a specified directory to which you back up data.

standard security *(MS SQL specific term)*
Standard security uses the login validation process of the Microsoft SQL Server for all connections. Standard security is useful in network environments with a variety of clients, some of which may not support trusted connections. It also provides backward compatibility for older versions of the Microsoft SQL Server.
See also **integrated security**.

Storage Group
(Microsoft Exchange Server 2000/2003 specific term)
A collection of databases (stores) that share a common set of transaction log files. Exchange manages each storage group with a separate server process.

StorageTek ACS library
(StorageTek specific term)
Automated Cartridge System is a library system (also known as Silo) consisting of one Library Management Unit (LMU) and one to 24 Library Storage Modules (LSM) connected to the unit.

storage volume *(ZDB specific term)*
A storage volume represents an object that may be presented to an operating system or some other entity (for example, a virtualization mechanism) upon which volume management systems, file systems, or other objects may exist. The volume management systems, file systems are built on this storage. Typically, these can be created or exist within a storage system such as a disk array.

switchover
See **failover**

Sybase Backup Server API *(Sybase specific term)*
An industry-standard interface developed for the exchange of backup and recovery information between a Sybase SQL Server and a backup solution like Data Protector.

Sybase SQL Server *(Sybase specific term)*
The server in the Sybase "client-server" architecture. Sybase SQL Server manages multiple databases and multiple users, keeps track of the actual location of data on disks, maintains mapping of logical data description to physical data storage, and maintains data and procedure caches in memory.

Symmetrix Agent (SYMA) *(EMC Symmetrix specific term)*
The Data Protector software module that prepares the EMC Symmetrix environment for backup and restore operations.

System Backup to Tape *(Oracle specific term)*
An Oracle interface that handles the actions required to load, label, and unload correct backup devices when Oracle issues a backup or restore request.

system databases *(Sybase specific term)*
The four system databases on a newly installed Sybase SQL Server are the:

- master database (master)
- temporary database (tempdb)
- system procedure database (sybsystemprocs)
- model database (model).

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

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

System State *(Windows specific term)*
The System State data comprises the Registry, COM+ Class Registration database, system startup files, and the Certificate Services database (if the server is a certificate server). If the server is a domain controller, Active Directory directory services and the Sysvol directory are also contained in the System State data. If the server is running the Cluster service, the System State data also includes resource registry checkpoints and the quorum resource recovery log, which contains the most recent cluster database information.

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

SysVol (*Windows specific term)*
A shared directory that stores the server copy of the domain's public files, which are replicated among all domain controllers in the domain.

tablespace
A part of a database structure. Each database is logically divided into one or more tablespaces. Each tablespace has data files or raw volumes exclusively associated with it.

tapeless backup *(ZDB specific term)*
See **ZDB to disk**.

target database *(Oracle specific term)*
In RMAN, the target database is the database that you are backing up or restoring.

target (R2) device *(EMC Symmetrix specific term)*
An EMC Symmetrix device that participates in SRDF operations with a source (R1) device. It resides in the remote EMC Symmetrix unit. It is paired with a source (R1) device in the local EMC Symmetrix unit and receives all write data from its mirrored pair. This device is not accessed by user applications during normal I/O operations. An R2 device must be assigned to an RDF2 group type.
See also **source (R1) device**

target system *(Disaster Recovery specific term)*
A system after a computer disaster has occurred. The target system is typically in a non-bootable state and the goal of disaster recovery is to restore this system to the original system configuration. The difference between a crashed system and a target system is that a target system has all faulty hardware replaced.

target volume *(ZDB specific term)*
A storage volume to which data is replicated.

Terminal Services *(Windows specific term)*
Windows Terminal Services provide a multi-session environment that allows clients to access a virtual Windows desktop session and Windows-based programs running on the server.

thread *(MS SQL Server 7.0/2000 specific term)*
An executable entity that belongs to only one process. It comprises a program counter, a user-mode stack, a kernel-mode stack, and a set of register values. Several threads can run at a time within one process.

TimeFinder *(EMC Symmetrix specific term)*
A business continuation process that creates an instant copy of single or multiple Symmetrix Logical Devices (SLDs). The instant copy is created on specially preconfigured SLDs called BCVs and is accessible via a separate device address to the system(s).

TLU
Tape Library Unit.

TNSNAMES.ORA *(Oracle and SAP R/3 specific term)*
A network configuration file that

contains connect descriptors mapped to service names. The file may be maintained centrally or locally, for use by all or individual clients.

transaction
A mechanism for ensuring that a set of actions is treated as a single unit of work. Databases use transactions to keep track of database changes.

transaction backup
Transaction backups generally use fewer resources than database backups, so they can be created more frequently than database backups. By applying transaction backups, you can recover the database to a specific point in time prior to when a problem occurred.

transaction backup *(Sybase and SQL specific term*)
A backup of the transaction log providing a record of changes made since the last full or transaction backup.

transaction log backup
Transaction log backups generally use fewer resources than database backups so they can be created more frequently than database backups. By applying transaction log backups, you can recover the database to a specific point in time.

transaction log files
Files that record transactions of the database modifications, and provide fault tolerance in case of a database disaster.

transaction logs (*Data Protector specific term)*
Keeps track of IDB changes. The archiving of transaction logs should be enabled to prevent you from losing the transaction log files that are created after the last IDB backup and are necessary for IDB recovery.

transaction log table *(Sybase specific term)*
A system table in which all changes to the database are automatically recorded.

transportable snapshot *(MS VSS specific term)*
A shadow copy that is created on the application system and can be presented to the backup system which performs the backup.
See also **Microsoft Volume Shadow Copy service (VSS)**.

TSANDS.CFG file *(Novell NetWare specific term)*
A file that allows you to specify the names of containers where you want backups to begin. It is text file located in the SYS:SYSTEM\TSA directory on the server where TSANDS.NLM is loaded.

unattended operation
See **lights-out operation**.

user account
You can use Data Protector only if you have a Data Protector user account, which restricts unauthorized access to Data Protector and to backed up data. Data Protector administrators create this account specifying a user logon name, the systems from which the user can log on, and a Data Protector user group membership. This is checked whenever the user starts the Data Protector user interface or performs specific tasks.

user disk quotas
NTFS quota-management support enables an enhanced tracking mechanism and control over disk space usage on shared storage volumes. Data Protector backs up user disk quotas on the whole system and for all configured users at a time.

user group
Each Data Protector user is member of a User Group. Each User Group has a set of user rights that are given to every user in that User Group. The number of User Groups with their associated user rights can be defined as desired. Data Protector provides three default user groups: admin, operator, and user.

user profile *(Windows specific term)*
Configuration information retained on a user basis. This information includes desktop settings, screen colors, network connections, and so on. When the user logs on, the user profile is loaded and the Windows environment is set accordingly.

user rights
User rights or access rights are the permissions needed to perform specific Data Protector tasks. Configuring a backup, starting a backup session, or starting a restore session are typical user rights. Users have the access rights of the user group to which they belong.

vaulting media
The process of storing media to a safe and remote place. The media are brought back to the data center when they are needed for restore or are ready for reuse in subsequent backups. The vaulting procedure depends on your company's backup strategy and policies for data protection/reliability.

VBFS *(OmniStorage specific term)*
A Very Big File System is an extension of the standard HP-UX file system on HP-UX 9.x. It is mounted to a directory the same way as any HP-UX file system. In a VBFS, only the superblock, the inode and the 'extended attribute'

information remain permanently on the hard disk and are never migrated.
See also **MFS**.

verify
A function that lets you check whether the Data Protector data on a specified medium is readable. Additionally, consistency within each block can be checked if the backup was performed with the cyclic redundancy check (CRC) option ON.

Virtual Controller Software (VCS)

(HP StorageWorks EVA specific term)
The firmware that manages all aspects of storage system operation, including communication with Command View EVA through the HSV controllers.
See also **Command View (CV) EVA**.

Virtual Device Interface *(MS SQL Server 7.0/2000 specific term)*
This is a SQL Server 7.0/2000 programming interface that allows fast backup and restore of large databases.

virtual disk *(HP StorageWorks EVA specific term)*
A unit of storage allocated from an HP StorageWorks Enterprise Virtual Array storage pool. Virtual disks are the entities that are replicated using the HP StorageWorks Enterprise Virtual Array snapshot functionality.
See also **source volume** and **target volume**.

virtual server
A virtual machine in a cluster environment defined in a domain by a network IP name and address. Its address is cached by the cluster software and mapped to the cluster node that is currently running the virtual server resources. This way all requests for a particular virtual server are cached by a specific cluster node.

volser *(ADIC and STK specific term)*
A VOLume SERial number is a label on the medium to identify the physical tape used in very large libraries. A volser is a naming convention specific to ADIC/GRAU and StorageTek devices.

volume group
A unit of data storage in an LVM system. A volume group can consist of one or more physical volumes. There can be more than one volume group on the system.

volume mountpoint (*Windows specific term*)
An empty directory on a volume that can be used to mount another volume. The volume mount point acts as a gateway to the target volume. Provided that the volume is mounted, users and applications can refer to the data on the

mounted volume by the full (merged) filesystem path as if both volumes are one integral part.

Volume Shadow Copy service
See **Microsoft Volume Shadow Copy service**.

VPO
See **OVO**.

VSS
See **Microsoft Volume Shadow Copy service**.

VxFS
Veritas Journal Filesystem.

VxVM (Veritas Volume Manager)
A Veritas Volume Manager is a system for managing disk space on Solaris platforms. A VxVM system consists of an arbitrary group of one or more physical volumes organized into logical disk groups.

Wake ONLAN
Remote power-up support for systems running in power-save mode from some other system on the same LAN.

Web reporting
The Data Protector functionality that allows you to view reports on backup status and Data Protector configuration using the Web interface.

wildcard character
A keyboard character that can be used to represent one or many characters. The asterisk (*), for example, typically represents one or more characters, and the question mark (?) typically represents a single character. Wildcard characters are often used in operating systems as a means of specifying more than one file by name.

Windows CONFIGURATION backup
Data Protector allows you to back up Windows CONFIGURATION, including Windows Registry, user profiles, Event Logs, and WINS and DHCP server data (if configured on a system) in one step.

Windows Registry
A centralized database used by Windows to store configuration information for the operating system and the installed applications.

WINS server A system running Windows Internet Name Service software that resolves Windows networking computer names to IP addresses. Data Protector can back up WINS server data as part of the Windows configuration.

writer
(MS VSS specific term)
A process that initiates change of data on the original volume. Writers are typically applications or system services that write persistent information on a volume. Writers also participate in the shadow copy synchronization process by assuring data consistency.

XBSA interface *(Informix specific term)*
The onbar utility and Data Protector communicate with each other through the X/Open Backup Specification Services Programmer's Interface (XBSA).

XCopy engine *(direct backup specific term)*
A SCSI-3 copy command that allows you to copy data from a storage device having a SCSI source address to a backup device having a SCSI destination address, thus enabling direct backup. The data flows from a source device (either block or streaming, that is, disk or tape) to the destination device (either block or streaming) through XCopy. This releases the controlling server of reading the data from the storage device into memory and then writing the information to the destination device.
See also **direct backup**.

ZDB
See **zero downtime backup (ZDB)**.

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

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

ZDB to disk+tape *(ZDB specific term)*
A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time, in the same way as with ZDB to disk. However, data in the replica is also streamed to a backup medium, as with ZDB to tape. If this backup method is used, data backed up in the same session can be restored

using the instant recovery process, the standard Data Protector restore from tape, or on split mirror arrays, split mirror restore.
See also **zero downtime backup (ZDB)**, **ZDB to disk**, **ZDB to tape**, **instant recovery**, **replica**, and **replica set rotation**.

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

zero downtime backup (ZDB)
A backup approach in which data replication techniques, provided by a disk array, are used to minimize the impact of backup operations on an application system. A replica of the data to be backed up is created first. All subsequent backup operations are performed on the replicated data rather than the original data, while the application system can be returned to normal operation.
See also **ZDB to disk**, **ZDB to tape**, **ZDB to disk+tape**, and **instant recovery**.

Z