

Symmetrix Connect User Guide

EDM Interface to Oracle (UNIX)

EDM Interface to Filesystems

EDM Release 4.6.0

P/N 300-113-005-04

Rev D



EMC Corporation, 171 South St., Hopkinton, MA 01748-9103

Corporate Headquarters: (508) 435-1000 (800) 424-EMC2

Fax: (508) 435-5374 Service: (800) SVC-4EMC

This software/documentation contains proprietary information of EMC Corporation; it is provided under a license agreement that contains restrictions on use and disclosure and is also protected by copyright law. Reverse engineering of the software is prohibited.

The information in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. EMC Corporation does not warrant that this document is error-free.

Restricted Rights Legend

Use, duplication, or disclosure by the Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of DFARS 252.227-7013, Rights in Technical Data and Computer Software (October 1988).

If this software/documentation is delivered to a U.S. Government Agency not within the Department of Defense, then it is delivered with "Restricted Rights," as defined in FAR 52.227-14, Rights in Data - General, including Alternate III (June 1987).

EMC², EMC, MOSAIC:2000, and Symmetrix are registered trademarks and EMC Enterprise Storage, EMC Enterprise Storage Network, EMC Enterprise Storage Specialist, EMC Storage Logic, The EMC Effect, The EMC Effect Alliance, The Enterprise Storage Company, Celerra, Connectrix, CopyPoint, EDM, E-Infostructure, FarPoint, InfoMover, PowerPath, SDMS, SRDF, SymmAPI, TimeFinder, Universal Data Tone, and Volume Logix are trademarks of EMC Corporation. Other trademarks are the property of their respective owners.

Contents

Preface	Who Should Use This Manual	<i>vii</i>
	How to Get the Information You Need	<i>viii</i>
	Books In PDF Format On Documentation CD	<i>viii</i>
	Online Help and Online Books	<i>viii</i>
	Hardcopy Documentation	<i>viii</i>
	Notes and Cautions	<i>x</i>
	Technical Support	<i>xi</i>
	Reader's Comments	<i>xi</i>

Chapter 1 Conceptual Overview

The EDM—Symmetrix Connection	<i>1-2</i>
Database and Filesystem Backup Interfaces	<i>1-4</i>
Mirroring for Backups	<i>1-5</i>
Mirroring for Restores	<i>1-6</i>
Symmetrix Device and Mirror Configurations	<i>1-6</i>
Backup Interface to Oracle (UNIX)	<i>1-8</i>
Oracle Online and Offline Backups	<i>1-8</i>
Online Backup Actions and Mirror Handling	<i>1-9</i>
Offline Backup Actions and Mirror Handling	<i>1-12</i>
Oracle Data Objects for Backup	<i>1-12</i>

Discovery of Data to Back Up	1-13
Data Layout and Backup Granularity	1-14
Device Layout Considerations	1-18
Movement of Backup Data	1-20
Restores for the Interface to Oracle	1-22
Restore Limitations for Data in Filesystems	1-23
Restore Destination Options and Limitations	1-24
Restore Actions, Including Mirror Handling	1-25
Backup Interface to UNIX Filesystems	1-28
UNIX Filesystem Backup Actions	1-28
UNIX Filesystem Backup and Restore Actions	1-29
Considerations for Backup Configuration	1-30

Chapter 2 Configuring Backups

Symmetrix Connect Prerequisites	2-2
Configuring Backups for Oracle	2-3
Configuring Backups for UNIX Filesystems	2-14
Creating a .sets File for Multi-BCVs	2-19
Basic Procedure to Create a .sets File	2-20
Format Variations for the .sets File	2-21
The Importance of the .sets File	2-22
Writing Scripts for Oracle Offline Backups	2-24
General Guidelines for Scripts	2-25
Sample Oracle Pre-Offline Scripts	2-26
Sample Oracle Post-Offline Scripts	2-28
Writing a Post-Mirror-Split Script	2-30
Writing Scripts for the Unix Filesystem Interface	2-31

Chapter 3 Running and Monitoring Backups

- Symmetrix Connect Backups 3-2
- Backing Up with the Backup Activity Wizard 3-3
- Backing Up with the eb_dc_backup Command 3-5
- Checking Log Files on Client and Server 3-6
- Stopping a Backup 3-6
- Monitoring Tape Operations 3-6

Chapter 4 Running Restores

- Symmetrix Connect Restores 4-2
- Restoring Old .dcrpt and .ddtab Files 4-2
- Restoring Oracle to an Alternate Destination 4-3
- Restoring UNIX Filesystems to an Alternate Destination 4-8
- Restoring with the Restore Window 4-12
- Restoring with the eb_dc_restore Command 4-17
 - Identifying Objects to Restore 4-17
 - Restore Command Options 4-18
- Restores to Alternate Host for Compaq and Sequent 4-20
- Restoring AdvFS Filesystems 4-20

Appendix A Sample Configuration: Oracle No-Mirror

Appendix B Sample Configuration: Oracle Mirrored

Appendix C Sample Configuration: UNIX Filesystem Interface

Appendix D Backup Configuration and Processing Concepts

Backup Configuration Results *D-2*

Symmetrix Connect Work Items *D-2*

Symmetrix Connect Kicker Work Item *D-3*

Work Groups for EDM Symmetrix Connect *D-4*

Templates for EDM Symmetrix Connect *D-5*

Network Work Item With Interface to Oracle or UNIX
Filesystem Interface *D-5*

Crontab Entry to Schedule Backups *D-6*

Backup Processing *D-7*

The .dcrpt and .ddtab Files *D-9*

Appendix E Setting Up a Tuning File

Tuning File Overview *E-2*

Creating a Tuning File *E-2*

Raw Filesystem Backups *E-4*

File Backups *E-5*

Appendix F Symmetrix Connect Prerequisites

EDM Prerequisites *F-2*

Backup Client Prerequisites *F-3*

Glossary

Index

Preface

This *EDM Symmetrix Connect — Interface to Oracle, UNIX Filesystems* guide documents the Unix client software that interfaces with EMC Data Manager (EDM) and Oracle's relational database management systems.

It provides information for configuring, backing up, and restoring data using EDM Symmetrix Connect interfaces to Oracle and UNIX filesystems.

Note: For other EDM Symmetrix Connect interfaces for UNIX, see *Oracle Backup Client* and *EMC Backint for SAP R/3 System*.

Note: The following guides provide information on Symmetrix Connect with the EDM-specific interfaces to NT filesystems and databases: *Windows NT Backup Client*, *Windows NT Oracle Backup Client*, *Windows NT SQL Server Backup Client*, and *Windows NT Exchange Backup Client*.

Who Should Use This Manual

This manual is intended for system administrators and Oracle database administrators who are responsible for either installing, configuring, administering, and operating EDM Symmetrix Connect or administering the Oracle database system or databases.

You will need to become familiar with Oracle database and Unix system administration. You should also understand your network environment, your Symmetrix environment, and the requirements of the various groups that you serve.

How to Get the Information You Need

Information about the EMC Data Manager is available through the graphical user interface (GUI) online help and online books, as well as hardcopy documentation, as described in the following sections.

Books In PDF Format On Documentation CD

The EDM Documentation CD contains the latest EDM documentation in Adobe PDF format. Use Adobe Acrobat Reader software (freely available from Adobe's web site) to display the documentation and do full-text searches across the document set with the indexes provided.

Online Help and Online Books

The EDM graphical user interface Help facility provides detailed information and procedures for managing the EDM server, its clients, and library units.

To access online help, log in to the EDM. Enter **edm** to display the EDM Main window, then click on the Help button. The Help facility includes context-sensitive help for input fields, options, and buttons as well as step-by-step instructions for tasks you perform to manage and monitor the EDM server.

Also available are online versions of selected manuals. To access these online books, select Help in the EDM Main window menu bar, then Online Books.

Hardcopy Documentation

The following documentation is available for use with EMC Data Manager (EDM).

- *EDM Software Reference* (P/N 300-113-001)
- *EDM Software Release Notes* (P/N 300-113-004)
- *EDM Storage Devices* (P/N 300-113-002)
- *EDM Server Error Messages* (P/N 300-113-014)

Symmetrix Connect Documentation

The *EDM Oracle Backup Client* guide describes the use of EDM Symmetrix Connect RMAN Proxy Copy with the EDM Oracle Backup Client for UNIX.

Note: However, the *EDM Symmetrix Connect User Guide* provides all the information on using EDM Symmetrix Connect with the EDM-specific interface to Oracle for UNIX and with the UNIX filesystem interface.

The *EMC Backint for SAP R/3 System* guide provides information on using EDM Symmetrix Connect with EMC Backint (for Oracle) on UNIX and Windows NT.

The *Windows NT Oracle Backup Client* guide describes the use of EDM Symmetrix Connect with the EDM-specific interface to Oracle for Windows NT.

The following guides describe the use of Symmetrix Connect with the EDM-specific interfaces to NT filesystems and other databases: *Windows NT Backup Client*, *Windows NT SQL Server Backup Client*, and *Windows NT Exchange Backup Client*.

Symmetrix Path Documentation

The *EDM Symmetrix Path User Guide* (P/N 300-113-007) provides information for configuring, backing up, and restoring data using EDM Symmetrix Path.

Clients and Databases

The following EMC Data Manager supplemental guides provide installation and configuration instructions for setting up the EDM backup clients. Release notes are also available.

- *NetWare Backup Client* (P/N 300-114-001)
- *OS/2 Backup Client* (P/N 300-118-001)
- *Windows NT Backup Client* (P/N 300-119-001)
- *OpenVMS Backup Client* (P/N 300-122-001)
- *Oracle Backup Client* (P/N 300-115-001)
- *Sybase Backup Client* (P/N 300-116-001)
- *Informix Backup Client* (P/N 300-117-001)
- *EMC Backint for SAP R/3 System* (P/N 300-120-001)
- *Windows NT SQL Server Backup Client* (P/N 300-121-001)
- *Windows NT Oracle Backup Client* (P/N 300-115-003)
- *Windows NT Exchange Backup Client* (P/N 300-119-003)
- *Windows NT Lotus Notes Backup Client* (P/N 300-119-005)

Notes and Cautions

Notes provide clarification or additional important information.

Note: It calls your attention to an operating procedure, practice, condition, or similar situation which is important to highlight.

Cautions are used to indicate the presence of a hazard.

CAUTION: It calls your attention to an operating procedure, practice, condition, or similar situation. Failure to observe a caution can result in minor personal injury, damage to a program, device, system, or loss of data.

Technical Support

If you need technical assistance with the EDM system, contact the EMC support hot line at:

1 800 SVC-4EMC (1 800 782-4362)

If you are located outside the United States or Canada, contact the nearest EMC office for assistance.

Reader's Comments

We welcome your comments on this documentation.

Send e-mail to:

`doc_comments@mil.emc.com`

Please include the part number and title of the manual.

1 Conceptual Overview

EMC Data Manager (EDM) Symmetrix Connect provides high-speed backup of Symmetrix-based data to tape. With Symmetrix Connect, EDM performs backups directly through its own Ultra-SCSI or Fibre Channel Symmetrix connections. When used with Symmetrix business continuity products, Symmetrix Connect can provide non-disruptive backups of large databases or data stored in filesystems.

This chapter describes the following topics:

- The EDM—Symmetrix Connection
- Symmetrix Device and Mirror Configurations
- Backup Interface to Oracle (UNIX)
- Discovery of Data to Back Up
- Movement of Backup Data
- Restores for the Interface to Oracle
- Backup Interface to UNIX Filesystems
- Considerations for Backup Configuration

Note: Along the way, this chapter poses some questions to help you actively think about how Symmetrix Connect can be best configured for your site.

The EDM—Symmetrix Connection


A database server whose data is stored on a Symmetrix can be backed up and restored through the EDM—Symmetrix connection. The EDM takes control of the entire movement of backup data; the data is not transferred through the database server or the network.

Thus, Symmetrix Connect provides *serverless backup*, removing the processing load from the database server, and *LAN-less backup*, bypassing network bottlenecks.¹

Note: In contrast to Symmetrix Connect, Symmetrix *Path* is a LAN-less backup solution that is not serverless.

Figure 1-1 illustrates a Symmetrix Connect site configuration. It shows two database servers and the EDM backup server, all of which are attached to the Symmetrix.

Figure 1-1



1. For Oracle, control file and archived redo logs are transferred through the database server and over the network.

The data resides on *Symmetrix devices*¹, which are logical disk devices composed of portions of hard drives. When Symmetrix devices are made visible to a host, they are said to be *exported* to the host. Because the Symmetrix is a multi-host-attachable disk array, the EDM can be configured to access the *same* Symmetrix devices as the database servers. (See Figure 1-2, left.)


To take a backup from the database's production devices in this so-called *no-mirror* configuration, Symmetrix Connect carefully coordinates the access to the devices. It takes steps to ensure the integrity of the data and then it enables the EDM to move the backup data directly from disk to tape.

If your site uses a Symmetrix business continuity product (EMC TimeFinder or SRDF), the EDM can access *mirror copies* of the production devices for the backup, instead of the production devices themselves. This minimizes the impact of the backup on your regular operations. (See Figure 1-2, right.)

Immediately before a backup, Symmetrix Connect *splits the mirrors* from the production devices and returns the database to normal operation on the production devices. Then the mirrors are backed up.

Figure 1-2

Devices Exported to the Database Server and to the EDM



1. Symmetrix documentation also calls them “hyper-volumes,” “meta-volumes,” or “logical volumes.” In this book, the term *logical volumes* is reserved for software-generated logical volumes, which are layered on the Symmetrix devices.

Database and Filesystem Backup Interfaces

EDM Symmetrix Connect can back up data from databases and from filesystems. This book documents its use for *Oracle databases* and *UNIX filesystems* with two *backup interfaces*:

- Backup Interface to Oracle (UNIX) (see page 1-8)
- Backup Interface to UNIX Filesystems (see page 1-28)

The Oracle interface is Symmetrix Connect's own interface to the Oracle database system for backups. It can back up Oracle datafiles on the Symmetrix devices whether or not the data is organized in filesystems.

Note: Symmetrix Connect also supports Oracle UNIX backups from two standard backup utilities. See the *EDM Oracle Backup Client* guide (for RMAN Proxy Copy) and *EMC Backint* (for SAP R/3 backups).

The UNIX filesystem interface enables the backup of data from other types of database systems or applications, as long as that data is organized within UNIX filesystems.

Note: Symmetrix Connect also supports backups of Windows NT filesystems and databases. See "Symmetrix Connect Documentation" on page ix.

Each backup interface provides Symmetrix Connect with the logical view of the data needed to: specify which data to back up (and restore), ensure the data's consistency before backup or mirror split operations, and return the database to normal operation.

For example, when discovering the data to back up, the interface to Oracle connects to the Oracle database, obtains a list of the datafiles, the archived redo log directory, and the control file. To prepare the database for backup or mirror split, the Oracle interface shuts down the database for offline backups, or puts it "online backup mode" for online backups. Finally, to return the database to normal operation, the interface to Oracle either starts up the database or takes it out of online backup mode.

Mirroring for Backups

Using Symmetrix device mirroring for backups offers this key advantage: you can back up a static mirror-copy of the data rather than the production data itself.

For example, Oracle databases need special handling to prepare them for backup; either shutdown or online backup mode. But with mirroring, they need this special handling for only the brief time that it takes to split the mirrors, not for the duration of the backup. During the backup, the database server can function normally using the production disks for the database.

Is database uptime a critical issue at your site?

Would you want to be able to take consistent database backups while the database is up?


Offline mirrored backups start up the database before backing up the mirrors.

Would you like the ability to do major updates to the database during online backups?

Online mirrored backups take the database out of Oracle's online backup mode before backing up the mirrors. With the database in normal mode, you can perform major updates to the database without the overhead of online backup mode.

Figure 1-3

Mirrored Backups Access Mirror Devices



Mirroring for Restores

Using Symmetrix device mirroring for restores can enable innovative restore strategies. After a backup, the EDM can leave the mirrors split until it is time for the next backup. In some cases, the mirror can be used as a rapidly restorable copy of the backup. If your production database were to become corrupted, you could rapidly resynchronize the mirror to the production devices and roll forward using the redo logs.

If your Symmetrix has been configured for *multiple sets* of mirrored devices, Symmetrix Connect can be configured to use a different set of mirrors for each subsequent backup, rotating through the mirror sets according to your pre-set policy. This capability also allows for creative backup and restore schemes.

Symmetrix Device and Mirror Configurations


A variety of Symmetrix device and mirror configurations are supported for backup:

- *No-mirror*. You can back up and restore directly from the production database devices. (See  in Figure 1-4.)
- *Local BCVs* (Business Continuance Volumes).  These are copies of production devices on the same Symmetrix, provided by EMC TimeFinder. (Production devices are called “standard” devices in TimeFinder environments.)
- *SRDF*.  The Symmetrix Remote Data Facility (SRDF) product provides copies of production devices on a separate Symmetrix. The remote “target” or “R2” device mirrors correspond with the production devices. (Production devices are called “source” or “R1” devices in SRDF environments.)
- *Remote BCVs*.  These are possible with SRDF and TimeFinder, with TimeFinder on the target Symmetrix.

Note: All remote BCV mirror operations, for both backups and restores, affect the TimeFinder link (between the R2 and remote BCV mirror) *only*. To ensure continuous SRDF operation, the link between R1 and R2 devices is not affected.

Figure 1-4

Symmetrix Mirror Configurations




Are local BCVs being used for backups at your site?

Are SRDF targets or remote BCVs being backed up?

With the *multi-BCV* feature, you can employ multiple sets of BCVs in your backup scheme. TimeFinder's Multi-BCV function allows up to eight separate BCV relationships to be paired with a single standard device. When TimeFinder is used with SRDF, BCVs on remote target Symmetrix can be paired with SRDF "target" or "R2" devices. (See Figure 1-5.)

Are multi-BCVs being used at your site? If so, how, if at all, do you want to employ them for backups?

Figure 1-5

Multi-BCV Example**Backup Interface to Oracle (UNIX)**

The backup interface to Oracle is essentially a specialized module of client software that enables the EDM to specify which Oracle data to back up and how. It enables the EDM to initiate and control online and offline backups of Oracle databases, archived redo logs, and control files.

The interface to Oracle supports these backup capabilities:

- Online and offline backups
- Database or tablespace backups

Oracle Online and Offline Backups

Oracle 7 refers to backups as being online (or hot) or offline (or cold), while Oracle 8 refers to backups as being open or closed. This book generally uses the terms online and offline.

The backup interface to Oracle makes the distinction between online and offline backups, adjusting its processing accordingly. For online backups, the affected tablespaces are put into Oracle's online backup mode. The online backup mode ensures that any inconsistencies that occur during an online backup can be resolved by the Oracle software during recovery.

The online backup mode saves full images of all changed database blocks to the redo logs (instead of logging just the changes within the blocks). This dramatically increases the amount of data that is logged. Therefore, as soon as it can, the interface to Oracle takes the tablespaces out of the mode, archives the redo logs, and backs up the archived redo logs.

For offline backups, the interface to Oracle shuts down the database instance and, as soon as it can, it starts it up again. However, the interface to Oracle allows you to write your own shutdown and startup scripts, giving you control of how the database will be handled. If you provide a *pre-offline* script, Symmetrix Connect checks that the database instance is shut down before performing the backup. If you provide a *post-offline* script, the script defines how the instance is to be started.

Online Backup Actions and Mirror Handling

For online backups in mirrored configurations, the tablespaces are put in online backup mode for *only the time it takes to split the mirrors*. As soon as the mirrors are split, the interface to Oracle takes the tablespaces out of online backup mode and backs up the mirror. This minimizes the processing overhead from full image logging.

In no-mirror configurations, the tablespaces are placed in online backup mode for the duration of the backup. EDM Symmetrix Connect backs up the database from the standard devices and when the backup is complete, the EDM takes the tablespaces out of online backup mode.

With online backups in a non-mirrored configuration, the extended time that the tablespaces are in online backup mode can be a problem, because if major database updating is performed, the overhead of extensive logging can be significant.

Note: For Compaq Tru64 and Sequent, if the datafiles reside in filesystems, online no-mirror backups are not supported. See “Backing Up Oracle Filesystems as Devices (Raw Filesystems)” on page 1-15.

Figure 1-6 illustrates the processing of online Oracle backups for mirrored and no-mirror configurations.





Figure 1-7 illustrates the processing of offline Oracle backups for mirrored and no-mirror configurations.

Figure 1-7

Offline Backups with the Interface to Oracle



Offline Backup Actions and Mirror Handling

For offline backups in mirrored configurations, the database is shut down for *only the time it takes to split the mirrors*. As soon as the mirrors are split, the interface to Oracle *starts up* the database *before* backing up the mirror. Therefore, mirrored configurations minimize the time the database is actually offline to a fraction of the duration of the backup.

For offline backups in no-mirror configurations, the Oracle database is shut down for the duration of the backup. EDM Symmetrix Connect backs up the database from the production devices and when the backup is complete, the EDM returns the database online. Because the database is down for an extended time, this might not be practical for your site.

Oracle Data Objects for Backup

When configuring a backup through Symmetrix Connect's backup interface to Oracle, you can specify a whole database for backup or selected *tablespaces*.

A *whole database* backup backs up all the tablespaces, and if tablespaces are added or deleted to the database, the whole database backup adjusts appropriately.

A *tablespace backup* backs up all of the individual datafiles that constitute the tablespace, and can be used for partial database backups. You cannot specify backups for selected datafiles; the finest level of granularity that you can specify is the tablespace.

For both offline and online backups, the *control file* must be backed up. For online backups, the *archived redo logs* directory also must be backed up. Symmetrix Connect's backup interface to Oracle interface takes care of this (backing up only the archived redo logs that have changed since the last backup).

The control file and archived redo logs are backed up through the database server and sent over the network (or over EDM Symmetrix Path, depending on the EDM backup configuration). Similarly, restores of the control file and logs take place over the network.


Discovery of Data to Back Up

Before each backup, Symmetrix Connect goes to the Symmetrix devices and discovers the Oracle data to back up. It identifies EDM-visible host pathnames for all requested datafiles, the archived log directory, and the control file.

The granularity of the backup depends on both the Oracle data layout on devices and the device layouts. Envision a pyramid of separate but related functional layers relating Symmetrix storage of data.

Figure 1-8

Storage Layers for Oracle Data



We have discussed the layers check-marked in Figure 1-8.

The following sections cover the host system layers:

- Data Layout and Backup Granularity
- Device Layout Considerations

Data Layout and Backup Granularity

The Oracle datafiles must reside on Symmetrix devices, but they can be placed either directly on the devices — as viewed by the host operating system or *Logical Volume Manager (LVM)* — or they can be placed within the operating system's filesystems.

- If each datafile is assigned directly to its own host device — whether they be native devices, native disk partitions, or LVM logical volumes — the datafiles can be backed up as those devices:
 - For example, a datafile could be backed up as a native Solaris device, such as `/dev/rdisk/c0t0d0s2`, or an AIX device, such as `/dev/rhdisk56`.
 - For those operating systems that use disk partitions or slices, a disk partition can be designated as a device (e.g., for Solaris, `/dev/rdisk/c0t0d0s3`).
 - For those operating systems that do not use partitions (and optionally, for those that do), a logical volume created by an LVM can be designated as a device. (e.g., for HPLVM, `/dev/vg1/rlvol10`).
- If Oracle datafiles are placed as regular files within filesystems (e.g., `/oradata1/prods01/datadbs01.dbf`):
 - Each filesystem can be backed up as a device, whether they be whole disks, disk partitions, or logical volumes.
 - For Solaris, HP, and AIX, each datafile can be backed up as a file.

Backing Up Partitions or Logical Volumes as Devices

If your Oracle datafiles are built directly on devices, EDM Symmetrix Connect backs up the *devices as devices*. (This is also known as backing up “raw partitions.”)

If each datafile is assigned its own device, restore granularity is at the datafile level. The backups are well suited to how the data is placed on the device; you are backing up precisely the data you specify and no more.

The devices are backed up as devices whether they are devices, partitions, or logical volumes.

Note: The logical volumes can be *striped* or *concatenated*. See “Device Layout Considerations” on page 1-18.

Are your datafiles built on devices (or partitions, or logical volumes)?

Backing Up Oracle Filesystems as Devices (Raw Filesystems)

If your Oracle datafiles are built in filesystems, EDM Symmetrix Connect can back up the *filesystems as devices*. For example, /oracle/product/8.1.6/datadbs01.dbf (a datafile in a filesystem) could be backed up in its entirety as the device /dev/rdisk/c0t0d0s3. (Since the device contains a filesystem, this is referred to as backing up “raw filesystems.”)

If your Oracle datafiles reside in Compaq Tru64 or Sequent filesystems, you have no choice. EDM Symmetrix Connect must back up each filesystem as a device.

Are your Oracle datafiles in Compaq Tru64¹ or Sequent filesystems?

If your datafiles reside in Solaris, HP, or AIX filesystems, backing up filesystems-as-devices is optional.

Backing up filesystems-as-devices has the advantage of high backup performance, but it comes at a cost: the lack of file-by-file granularity. Instead of one datafile per device, there could be hundreds of datafiles on the device. And the lack of file-by-file granularity results in a serious limitation: restores of a filesystem-as-a-device can only reconstruct the entire filesystem, overwriting all other files already present.


1. For AdvFS (Advanced File System) with Compaq Tru64, Symmetrix Connect provides only limited backup support. The *domain* is the smallest unit of backup. Each domain can contain only one fileset and each domain must be backed up as a device.

Figure 1-9

Single Datafile on Raw Partition
or on Logical Volume
`/dev/rdsk/c0t0d0s3`



Backed Up as Raw Device



`/dev/rdsk/c0t0d0s3`


Restores the entire raw partition,
which happens to be just the
single datafile

Oracle Data Layout and Datafile Restore Granularity

Datafiles in Filesystem
`/oradata1/prods01/datadbs01.dbf`
`/oradata1/prods01/datadbs02.dbf`
`/oradata1/prods01/datadbs03.dbf`




Either Backed Up within
Filesystem-as-Device



`/dev/rdsk/c0t0d0s3`

Restores the entire raw partition
which includes many more
datafiles than the one desired

Or Backed Up as a File



`/oradata1/prods01/datadbs01.dbf`
`/oradata1/prods01/datadbs02.dbf`
`/oradata1/prods01/datadbs03.dbf`

Restores just the intended datafile
or datafiles

Backing Up Datafiles as Regular Files

If your datafiles reside in Solaris, HP, or AIX filesystems, you can back up the individual *datafiles as files*. This gives you the advantage of file-by-file granularity. You can restore individual files without overwriting other files in the filesystem.¹

Backups and restores can run slower with the datafiles-as-files approach, but actual performance depends upon the operating system platform, how the data is dispersed on the disks, and the number of files in question for a particular backup or restore.

1. For non-Oracle files, using the Backup Interface to UNIX Filesystems, you do not have this option for any operating system. See "Backing Up UNIX Filesystems as Devices" on page 1-29.

Therefore, when choosing between filesystems-as-devices backups and datafiles-as-files backups, you have to balance the trade-offs: possible backup performance improvements *vs.* improved restore granularity (with restrictions, described next).

Note: You cannot configure filesystems-as-devices and datafiles-as-files backups together in a single Symmetrix Connect backup.

Note: Restores always match the backup method.

Are you backing up using the filesystem-as-device method or the datafiles-as-files method?

Datafiles-as-Files Backup Limitations

Although the datafiles-as-files functionality does provide file-level granularity, it does not fully support standard backup functionality in the following respects:

- Directories are not backed up; you do not get a backup of the directory hierarchies leading to the client files.
- Extended file attributes, such as ACLs, are not backed up.
- Sparse files are not supported. (Oracle does not create or use sparse datafiles. But, as a safeguard, Symmetrix Connect examines for gaps in files, and if detected, fails the backup or restore).


The first limitation causes an important restore limitation: Symmetrix Connect cannot restore directories. You must manually recreate the directory paths to all restore targets, at the start of the restore session. See “Restoring Oracle to an Alternate Destination” on page 4-3.

Also, the client filesystems must be mounted at the start of the restore (to obtain extent maps of the files).

Device Layout Considerations

Device layout — how the data is actually stored on the disks — affects just how granular the data captured for backup can be. Entire disks, or disk partitions, or logical volumes (created by an LVM) can be designated as a device to the operating system.

Figure 1-10



Note: Symmetrix Connect does not support LVM-mirrored volumes.

Logical volumes that are *concatenated* are written to and backed up in the same concatenated order. Therefore, a restore destination can be made to any number and size of device portions, as long as the overall size of the logical volume is equal to or greater than the size of the backup. (See Figure 1-11.)

Logical volumes that are *striped* are written to differently than concatenated volumes, but for Compaq Tru64 and Sequent, backups are taken in the same order as they are for concatenated. Therefore, the data is *not* captured in “data-logical order.” This causes stringent restrictions on restores for these platforms.


In contrast, on Solaris, HP, and AIX platforms, backups of striped logical volumes capture the data in “data-logical order.” Therefore, restores are possible without regard to the physical layouts in effect at backup or restore time. For these platforms, data in any striped or concatenated configuration can be restored to new volume layouts if desired, as long as the overall size of the logical volume is equal to or greater than the size of the backup.

Figure 1-11


Concatenated and Striped Logical Volumes

Concatenated Logical Volume

LVM fills each portion of each device consecutively




Symmetrix Connect backs up sequentially




Restores can be to any number and size of device portions (but overall size must be same or larger)

**Striped Logical Volume**


LVM fills each portion of each device alternately




Compaq Tru64, Sequent:
Symm Connect backs up sequentially




Solaris, HP, and AIX:
Symm Connect backs up logically



RESTRICTION: Restores can *only* be to **exact** number and size of stripes



Restores can be to *any* number and size of stripes



Movement of Backup Data

To back up a large database, you need to be able to back up multiple streams of data concurrently. For backups, you need to read multiple streams from disk and write them to tape. For restores, you need to read multiple streams from tape and write them to disk.

The EDM's backup configuration specifies the streams for moving the backed up datafiles and logs between disk and tape for both backup and restore operations.

Each Symmetrix Connect backup is identified by a work item base name, which creates a set of work items. The work items define the front-end backup work (reading the data).

There are these three types of work items:


- Symmetrix Connect work items, which back up the datafiles over the direct EDM—Symmetrix connection. Each one specifies one stream to be read directly from disk for backups (and to be restored directly for restores).
- A network work item, which backs up the archived redo logs and the control file over the network. The network work item can be used for restoring datafiles as well as the archived redo logs and the control file.
- A kicker work item, which initiates the work item processing for the entire set of work items with the same base name.

The back-end of the backup process (writing to tape) is identified by a unique trail base name. The trail base name creates a group of specifications called trails, which define the streams of data to be written to tape.

It is possible for you to configure more than one work item per trail to obtain better backup performance (see Figure 1-12). When there is more than one work item per trail, the backup data is interleaved on the tape. For maximum restore performance, you would configure one work item per trail.

Figure 1-12

Symmetrix Connect Backup Data Streams



Note: See Appendix D, “Backup Configuration and Processing Concepts” for more information on work items and trails. Also see Appendix E, “Setting Up a Tuning File” for information on performance tuning.

Restores for the Interface to Oracle

The backup interface to Oracle enables the EDM to initiate and control restores of Oracle databases, archived redo logs, and control files.

Symmetrix Connect can restore the data from tape to the Symmetrix device over the direct EDM—Symmetrix connection. There are some exceptions — for some alternate destination restores — where the restore data movement of the datafile backups (which were backed up over the direct connection) must go over the network instead.

The backup interface to Oracle supports the following restore features:

- Specifying and running restores
- Selecting data to restore (database, selected tablespaces, and individual datafiles of the database)
- In-place *and* alternate destination restores that take place over the EDM—Symmetrix connections
- Alternate restore data movement, restoring datafiles that were backed up Symmetrix Connect over the network instead
- Alternate mirror restores, including restores to the production devices after establishing and splitting mirrors to obtain a snapshot of the current production devices

Note: These features are subject to various limitations as described in the following sections.

Restore Limitations for Data in Filesystems

EDM Symmetrix Connect does not fully support filesystem restore functionality in the following respects:

- The filesystem-as-device method has these limitations:
 - More files will be restored than requested, unless you request the whole filesystem, because the whole filesystem is restored.
 - The database server's filesystems must be mounted at the start of the restore.
 - The restore unmounts the database server's filesystems during the data movement.
- The datafiles-as-files method has these restrictions:
 - Directories are not restored. You must manually recreate the directory paths to all physical restore targets prior to the restore session.
 - The server's filesystems must be mounted at the start of the restore to obtain extent maps of the files.
 - The restore creates the database server's files prior to actually restoring the file contents. It needs to obtain extent maps for files. Therefore, the restore pre-allocates the files, in their proper sizes, *on the production devices*, even when restoring to a mirror).

Note: This might require extra time and consume the database server's system resources.

- Split mirror restores must start by establishing the mirrors before splitting them; this is needed for the file pre-allocation.
- The restore unmounts the data server's filesystems during the data movement to ensure the files cannot be accessed and to invalidate any in-memory image of the file data on the client during the restore.

Restore Destination Options and Limitations

Symmetrix Connect has the following restore destination options for restores that can take place over the direct connection.

- Alternate directory; you can specify an alternate directory at restore time. There are these exceptions:
 - A restore destination that is a partition or logical volume must exist prior to the restore and must be of a size that is equal to or greater than the backup size.
 - Alternate directory restores can *only* go *over the network* (or Symmetrix Path) for Compaq Tru64 and Sequent instead of over the direct connection.
 - Alternate directory restores cannot be performed at all for LVM-striped logical volumes for Compaq Tru64 and Sequent (and for pre-4.6.0 LVM-striped logical volume backups of Sun, HP, and AIX), because these restores cannot go over the network.
 - At the start of the restore session, you must manually create a link from the backed-up pathname to the new pathname.

Note: See “Restoring Oracle to an Alternate Destination” on page 4-3.

- Alternate client of the same platform type; you can specify the alternate client at restore time. There are these restrictions:
 - A restore destination that is a partition or logical volume must exist prior to the restore and must be of a size that is equal to or greater than the backup size.
 - A restore destination for LVM-striped backups of Compaq Tru64 and Sequent (and also for old backups taken prior to EDM 4.5.0, on Solaris, HP and AIX) must be exactly the same size, with exactly the same size and number of stripes.

Restore Actions, Including Mirror Handling

Restores can be done to tablespaces that are offline while other tablespaces remain online, as long as no other portion of the database is located on any of the Symmetrix devices being written to.


Note: The restore process does not shut down the database or take tablespaces offline.

If the site is using a Symmetrix business continuity product, such as EMC TimeFinder or SRDF, the EDM can restore to *mirror copies* of the production devices (instead of, or in addition to, the production devices themselves).

By default, the restore process first establishes and then splits the mirrors. It then restores from backup media over the Symmetrix connection to the mirror. The default post-restore mirror policy is to leave mirrors down. This gives you control as to when you want mirrors “restored,” meaning that the data on the mirror overwrites the standard devices.

Figure 1-13

Oracle Restore to Local BCV



Note: When restoring backups that were taken from mirrored configurations, if the datafiles are built on filesystems, EDM Symmetrix Connect runs a filesystem check (fsck) to check and recover the filesystems, which might not have been unmounted when the backup was taken.

Alternate Mirror Restores

Restores of EDM Symmetrix Connect backups can go to an alternate mirror set, rather than the one that was backed up. The restore takes place over the EDM—Symmetrix connection.

Note: To restore to an alternate mirror, the alternate mirror devices must be exported to (visible to) the EDM.

You specify an alternate mirror at restore time, as well as pre-restore and post-restore handling of mirrors.

For example, if the Remote BCVs are backed up, you could choose to restore to SRDF target (R2) volumes. Or, you could restore to the standard volumes. (This is referred to as restoring to “no-mirror.”)

Note: Restores to Remote BCVs must be started with the SRDF mirror up for Solaris, HP, and AIX hosts using file level granularity. An error message will be issued and the restore will fail if this is not the case.

Restoring to No-Mirror with Snapshot in Mirrored Configurations

If you are restoring datafiles-as-files, the pre-allocation activity creates the files on the *production devices (even if you restore to a mirror)*. Therefore, a snapshot capability is available to save the existing state of the production devices.


Note: The “no-mirror” restore capability requires that the production devices be exported to (visible to) the EDM.

For datafiles-as-files restores, the “no-mirror” restore with snapshot capability is your best choice for the mirror type. Since the datafiles-as-files restore must *alter the production devices* with the pre-allocation, no matter which mirror you restore to, it is efficient to do the *restore* to the production devices.

The snapshot function copies the production devices to the mirror type of your choice. It synchronizes the production devices to mirrors before the no-mirror restore takes place. This gives you the ability to roll-back from the restore should the restore fail or if the successfully restored data is, for whatever reason, not what you want.

Figure 1-14

Oracle No-Mirror Restore With Snapshot (Example of Local BCV)



Note: If you are doing a no-mirror restore and choose a snapshot to a Remote BCV, the snapshot will be taken from *the SRDF* devices, *not* the *production* devices. If you do the restore with the SRDF link down, the snapshot will not contain a copy of production data. Consider starting up the SRDF link before doing the restore or consider choosing an SRDF snapshot instead of the Remote BCV snapshot.

Backup Interface to UNIX Filesystems

Symmetrix Connect is best-suited to backing up large amounts of data, collectively and rapidly. Accordingly, the interface to UNIX filesystems enables Symmetrix Connect backup of filesystems *as devices* only. Backing up filesystems-as-devices has the advantage of high backup performance, but lacks of file-by-file granularity.

You can use this interface to back up databases and software applications that store their data in the supported UNIX filesystems. The interface to UNIX filesystems enables you to provide your own database or application shutdown and startup scripts to complete such an implementation.

A database or application that would be a good candidate for Symmetrix Connect backups with this interface would:

- Use large amounts of Symmetrix storage.
- Use small amounts of non-Symmetrix storage.
- Have some means of freezing the disk image, such as taking the application offline, to ensure the backup data will be internally consistent.
- Be able to list all the filenames that contain the Symmetrix-based data and the non-Symmetrix data.

UNIX Filesystem Backup Actions

Without the use of any scripts, the filesystem remains mounted for backup in a mirrored configuration. The filesystem I/O buffer is flushed immediately before the mirrors are split to pick up changes that haven't been synchronized to disk.

In non-mirrored configurations, the filesystem is unmounted by the EDM Symmetrix Connect software for the duration of the backup.

Backing Up UNIX Filesystems as Devices

For the interface to UNIX filesystems, EDM Symmetrix Connect must back up each *filesystem as a device*. For example, /home (a filesystem) is backed up in its entirety as the device /dev/rdisk/c0t0d0s3.

For the filesystem interface, for non-Oracle files, this is the *only* backup method used for all supported operating systems, Solaris, HP, and AIX, as well as Compaq Tru64¹ and Sequent.

UNIX Filesystem Backup and Restore Actions

When restoring backups from mirrored configurations, EDM Symmetrix Connect unmounts the targeted filesystems, runs a filesystem check (fsck) to check and recover the filesystems (after the data is on the filesystems), and remounts the filesystems.

Restore Limitations for UNIX Filesystem Interface

Restores of a filesystem-as-a-device can only reconstruct the entire filesystem, overwriting all other files already present. You must manually recreate the mount point at the start of the restore session. See “Restoring UNIX Filesystems to an Alternate Destination” on page 4-8.

See also “Restore Destination Options and Limitations” on page 1-24.

1. For AdvFS, the *domain*— not the filesystem — is the smallest unit of backup. See “Backing Up Oracle Filesystems as Devices (Raw Filesystems)” on page 1-15.

Considerations for Backup Configuration

The following key points to consider prior to configuring Symmetrix Connect backups.

- ❑ Client platform type: Identify your type of operating system, filesystem, Logical Volume Manager, and device layout. Learn more about how each is supported, what it really means for your backups and restores, and how it affects the choices you can make.
- ❑ Client and data identifying information: Know the hostname for the client. For Oracle backups, know the name of the database instance. You must also know the TNS alias of the Oracle instance, and the username and password pair for connecting to the instance. For non-Oracle filesystem backups, know the filesystems to backup.
- ❑ Database backup methods: How is your database is laid out on disk? If your data is in filesystems, have you considered the trade-offs of raw filesystem *vs.* files-in-files backups? How do you want backups to proceed? Offline or online? Whole or partial? If partial, which tablespaces to backup.
- ❑ Restores: Envision the kind of restores that are most important for you to be able to perform. Logical repairs or disaster recovery? Use of mirrors for restore?
- ❑ Mirroring configuration: Understand how your Symmetrix was configured for mirrors. Understand which mirroring configuration to specify at backup configuration time and if you can optionally choose a different mirror type. Do you have multiple BCVs associated with any standard device? (If you do, you need to create a .sets file as described in “Creating a .sets File for Multi-BCVs” on page 2-19. Optionally do this before performing the backup configuration in the Backup Configuration Wizard.)

- ❑ Mirroring handling options: Understand the expected state of the mirrors before backup and the state to leave the mirrors in after the backup. (See step 19 on page 2-7.) If you have multiple BCVs and a .sets file, think through how your mirror options will work in conjunction with your multiple BCV policy.
- ❑ Understand which mirroring options to specify at restore time and which different mirror types you might be able to choose. Understand the implications of restore mirror synchronization.
- ❑ Backup streams and media considerations: The number of drives for the backup. The the ratio of work items per trails. You might also want to think of meaningful names for the work item base name, template names, and trail names.
- ❑ Other EDM backup configuration features: Are you going to use media duplication? What is backup rotation period? Do you want to start new tapes for every backup? How long do you want to retain backups? Do you want EDM to automatically schedule the backups?

2 Configuring Backups

This chapter describes how to install and configure the Oracle database server for EDM Symmetrix Connect's own interface to Oracle.

Initial installation is performed by EMC service personnel.

This chapter describes the following topics:

- Symmetrix Connect Prerequisites
- Configuring Backups for Oracle
- Configuring Backups for UNIX Filesystems
- Creating a .sets File for Multi-BCVs
- Writing Scripts for Oracle Offline Backups
- Writing a Post-Mirror-Split Script
- Writing Scripts for the Unix Filesystem Interface

Symmetrix Connect Prerequisites

Prior to performing the Symmetrix Connect configuration procedures in this chapter, the following overall activities should have been performed:

- ❑ Oracle host system installation including the operating system, any volume manager software, and the Oracle database product.

Note: The Oracle database server's **filesystems must be mounted** for the configuration.

Note: The **database instance must be up and running** for the configuration.

- ❑ EDM hardware installation performed by EMC Customer Engineer (CE).
- ❑ Symmetrix hardware and software installation performed by EMC Customer Engineer (CE), including cabling between Symmetrix and Oracle host, and between Symmetrix and EDM.
- ❑ EDM software installation performed by EMC Regional Technical Support representative (RTS), including EDM preparation for EDM Symmetrix Connect.

For more detailed prerequisites, see Appendix F.

Configuring Backups for Oracle

Use the Backup Configuration Wizard to configure EDM Symmetrix Connect backups for Oracle. The wizard guides you through the steps.

Note: Alternatively, you can use the **eb_config** command, a command-line version of the wizard.


To configure the Oracle Symmetrix Connect backups:

1. Log into the EDM as root.
2. Set your display environment variable and enter the **edm** command to start the EDM interface:


```
edm# setenv DISPLAY displayname:0.0
edm# edm &
```

Note: See the *EMC Data Manager Software Reference* manual for more information about starting the GUI.

3. From the Backup menu, choose Configure -> Wizard... to open the EDM Backup Configuration Wizard.
4. In the “Select the clients...” panel, select the UNIX database server (“summer,” in this example) to configure for backup. Click Next.



5. A panel asks if you are creating a new configuration or modifying, duplicating, or deleting an existing one. Select New and click Next. (The rest of the procedure shown here is for a new configuration.)
6. In the “Select the objects to configure” panel, scroll down the list of filesystems and databases to find the Oracle database (“CLE,” in this example). Click in the checkbox to select it. Click Next.



7. In the “Select the application” panel, select “EDM Symmetrix Connect” to choose the Symmetrix Connect interface to Oracle.

8. If the client has been enabled for Symmetrix Path, a panel asks if you want to use it. Choosing Yes backs up the archived redo logs and control file over Symmetrix Path. But, since Symmetrix Path is not designed for small files, choose No to back them up over the network.
9. Select a backup mode, either Online or Offline. See “Oracle Online and Offline Backups” on page 1-8.
10. If you specified Offline, you now are asked if you want to initiate backups of the archived redo log directory, even though this is an offline backup. Choose No, since this is an offline backup, or choose Yes to back up the archived redo logs to tape anyway.
11. You are next asked “Do you want to back up database datafiles located within filesystems as raw devices?”

Before answering the prompt, ask yourself these questions:

- a. Are the datafiles located within filesystems?

If not, choose No and click Next. (The datafiles will be backed up as devices anyway, but this question is about backing up *filesystems* as devices.) Go to step 12.

- b. Is the platform Compaq Tru64 or Sequent?

If yes, choose Yes and click Next. (You must choose Yes in order to back up the database.) Click Yes in the pop-up warning. Go to step 12.

- c. Do you want to collectively back up all the datafiles in each filesystem as a single device?

If yes, choose Yes and click Next. Click Yes in the pop-up warning. Go to step 12.

If your answer is no, it is because you want to back up the individual datafiles *as files*. Choose No and click Next.

Note: See “Data Layout and Backup Granularity” on page 1-14.

12. Enter the Oracle information needed for connecting to the database.

Provide the following Oracle database information:



TNS Alias:

Username:

Password:

The TNS alias identifies the database SID. The username and password are for the Oracle database account.

- a. Enter the TNS alias for the Oracle instance.
- b. Enter the Username (typically “system”).
- c. Enter the password.

Click Next.

13. Select the whole database or individual tablespaces for backup. The panel lists one item: the entire database, which you can:
 - Select by clicking in the checkbox, or
 - Expand by clicking on the “+” sign.


Select the tablespace(s) to configure:

☒

CLE

Note: Selecting the entire database ensures that any tablespaces added to the database after this configuration session will be backed up.

If you do expand the database to list all the tablespaces, you can increase the size of the wizard panel so that you do not have to scroll to see all the tablespaces.



CAUTION: Only Database Administrators or other trained personnel should set up and change individual tablespace selections.

After selecting the database or tablespaces, click Next.

14. Specify the UNIX account name for this Oracle database's DBA and click Next.
15. Choose the Symmetrix mirroring configuration that has been set up. See "Symmetrix Device and Mirror Configurations" on page 1-6. Click Next.

Note: If you select "No mirror," the next panel is at step 23 on page 2-8.

16. If you selected Local BCV or Remote BCV, and there are multiple-BCVs for any standard device involved in the backup, select "Multiple BCVs present." Otherwise, select "One BCV present for every device." Click Next.

17. If multiple BCVs are present, select to either rotate through sets of multiple BCVs or to just use one specified set of BCVs. Choose either: “Round Robin rotation of BCV sets” or “Use One BCV set” and click Next.

Note: See “Creating a .sets File for Multi-BCVs” on page 2-19.

18. If you are specifying a particular set, enter the BCV Set Name now. For example, enter “set1”. Click Next.

Note: If you have not created the .sets file, just enter set1. When you create the .sets file, define the first BCV set to contain the BCVs you want for the first set. You do not have to label the column (“set1” is the system default) but if you do label it, label it correctly.

19. You are now prompted to “Select the actions to take on the Symmetrix mirrors” *before* and *after* the backup. Click Next to take the defaults.

“Before Backup” specifies what Symmetrix Connect should do with mirrors *before* backup starts:

- **Bring mirrors down**
Expect to find the mirrors established and split the mirrors.
- **Verify that mirrors are down**
Expect to find the mirrors split and require the operator to verify that mirrors are split.
If used with “Leave mirrors up” after backup, the operator must manually split the mirrors before each backup.
- **Bring mirrors down if needed**
Expect to find the mirrors either split or established. If established, split the mirrors.
- **Bring mirrors down after establishing**
Expect to find the mirrors either split or established. If split, re-establish and then split the mirrors. If established, split the mirrors.

“After Backup” specifies what Symmetrix Connect should do with mirrors *after the backup* is completed:

- **Leave mirrors up**
Reestablish the mirrors.
 - **Leave mirrors down**
Do not reestablish the mirrors. *If used with “Bring mirrors down” before backup, the operator must manually establish the mirrors before the next backup.*
 - **Leave mirrors as found**
Reestablish the mirrors if they had been established before the backup started.
20. If you do not select Leave mirrors down, you are now prompted “Do you want the EDM to wait for tracks to resynchronize before displaying a message that the backup has finished?” Select from the dropdown list either: “Wait for resync” or “Don’t wait for resync.” Click Next.
 21. You are asked “Do you want to run a script on the client after mirrors are split, but before the backup begins?” This is optional.
 - Do *not* select this for the *pre- and post- offline scripts* described in “Writing Scripts for Oracle Offline Backups” on page 2-24. Click Next.
 - Select this only for the post-mirror-split script described in “Writing a Post-Mirror-Split Script” on page 2-30. Click Next.
 22. If you are using a post mirror-split script:
 - a. Enter the *full pathname* for the script. You can either browse existing files or enter the name. (Notice that the default is only a directory, not a file.)
 - b. Enter the Oracle username (UNIX account) to run it. Click Next.
 23. If there is more than one media type available on the EDM, you are now prompted to “Select a media type,” for example, DLT tape. Click Next.


24. You are now prompted to “Select the number of trails and streams for this configuration.” This is number of Symmetrix Connect trails and the number of Symmetrix Connect streams per trail. (See Figure 2-1 for more information.)

- a. In “Trails:”, enter the number of trails to use concurrently for the EDM Symmetrix Connect portion of this backup. (This is how many separate sets of media to use when backing up the database. It is also the number of drives to be written to simultaneously.)

For optimal backup performance, the number of trails should equal the number of available drives of the media type you are selecting. It cannot exceed the Server parameter for maximum concurrent trails for this media type, which is the default value. (In the Backup Configuration window, select the Server tab. Under Media Type, see the parameter labeled “Current Max.”)

- b. In “Number of streams per trail:”, enter the maximum number of concurrent data streams per trail for the EDM to read directly from the Symmetrix volumes. For faster backups, you could raise it higher than 1, (depending on disk layout and the use of a tuning file). For optimal restore performance, it should equal 1.

Figure 2-1



25. Enter the base names for the Symmetrix Connect work item and trail.

Provide base names for the work item and trail.



Work Item Base Name:

Trail Base Name:

These names describe the data that you want the EDM to back up. The work item base name must be unique for each configuration.

For the Symmetrix Connect's own Oracle interface:

The default work item base name is:

dc_oracle_mirroring_clientname.

The default trail base name is dc.

Note: At the end of the configuration session, by default, the templates and trails for Symmetrix Connect will be named using the base trail name in these formats:
template_basetrailname_trailnumber
trail_basetrailname_trailnumber

26. You are asked if you want to run a configuration check without running a backup. This will check whether the backup configuration is valid by performing a test backup (by running **eb_dc_backup -check**).

WARNING: Be aware that this validation process will split and re-establish the mirrors as configured. This could alter the data on the mirrors if they are already split. You do not get a choice to skip validation of the mirror handling parameters in the wizard, but you do get that choice with the eb_dc_backup_ -check command.

CAUTION: The configuration check can take hours to complete on very large databases; it is started before you complete the remaining wizard steps and you cannot finish until the check finishes.


Note: Skip the check if you have multi-BCVs and you have yet to create the .sets file (see “Creating a .sets File for Multi-BCVs” on page 2-19).

Choose No to skip this check. You are warned to run an **eb_dc_backup -check** *work_item_base_name* before running a backup.

Choose Yes to run the configuration check:


- a. You are prompted “Do you want verbose output?” Click Yes or No and click Next.
- b. If you decided to run the configuration check, you are prompted to click Run to begin the check.
- c. Once you click Run, you have the option to cancel while the check occurs.
- d. When the check is finished, you have the option to save the output.

27. Choose to enable media duplication or not.
28. Choose the backup rotation period. This value determines only the media rotation period. This is the duration after which, by default, EDM starts writing the backups to a new volume rather than appending the backups to the volume containing the previous backups.
29. Choose the backup retention period.
30. You are asked “Do you want the EDM to initiate the back up?” Choose Yes and set the time if you want the EDM to regularly run (from **cron** on a daily basis) the backup job (**ebbackup**) for this backup configuration. Choose No if you are using another scheduling tool or if you want to manually edit the crontab file that controls this function.
31. Choose Yes if you want the EDM to reschedule and reprocess any backups that fail in the current backup run (after all normally scheduled backups have completed).
32. Choose if you want to append new data to the current volume until the first day of a new rotation or if you want to use a new tape volume for each new backup (which enables you to eject the backup media for the previous backup).
33. Enter the name to use to identify this configuration in case you want to modify it in this backup configuration wizard (it is not used for anything else). For the Symmetrix Connect Oracle interface, the default configuration name is *dc_oracle_mirroring_clientname*.



34. You might receive a warning if the configuration will cause changes you need to know about. Click OK.

35. The last panel summarizes your configuration. Click Save to complete the configuration.



As a result of the Backup Configuration Wizard session:

- Symmetrix Connect configuration information is stored in a EDM Symmetrix Connect-specific ddtab (*work_item_base_name.cfg*); there is one such file for each work item base name.
- Backup configuration information is written to EDM's backup configuration (eb.cfg) file. See “Backup Configuration Results” on page D-2 for more information.

The backup configuration for the Oracle database is now complete.

Modifying a Configuration

Use the Backup Configuration wizard to modify an existing configuration. After you select the client to back up, select Modify Configuration, and then select a work item base name.


Configuring Backups for UNIX Filesystems

Use the Backup Configuration Wizard to configure EDM Symmetrix Connect backups for UNIX filesystems.


Note: Alternatively, you can use the **eb_config** command, a command-line version of the wizard.

To configure the UNIX filesystem Symmetrix Connect backups:

1. From the Backup menu, choose Configure -> Wizard... to open the EDM Backup Configuration Wizard.
2. In the “Select the clients...” panel, select the UNIX host (“summer,” in this example) to configure for backup. Click Next.



3. Next, select New to create a new configuration. Click Next.
4. In the “Select the objects to configure” panel, click the checkboxes for the filesystems to back up (“/u01” and “/u02,” in this example). Click Next.



5. In the “Select the application” panel, select “Symmetrix Connect backup” to choose the Symmetrix Connect interface to UNIX filesystems.
6. Choose the Symmetrix mirroring configuration that has been set up. See “Symmetrix Device and Mirror Configurations” on page 1-6. Click Next.

Note: If you select “No mirror,” the next panel is at step 14 on page 2-16.

7. If you selected Local BCV or Remote BCV, and there are multiple-BCVs for any standard device involved in the backup, select “Multiple BCVs present.” Otherwise, select “One BCV present for every device.” Click Next.

8. If multiple BCVs are present, select to either rotate through sets of multiple BCVs or to just use one specified set of BCVs. Click Next.

Note: See “Creating a .sets File for Multi-BCVs” on page 2-19.

9. If you are specifying a particular set, enter the BCV Set Name now. For example, enter “set1”. Click Next.

Note: If you have not created the .sets file, just enter set1. When you create the .sets file, define the first BCV set to contain the BCVs you want for the first set.

10. You are now prompted to “Select the actions to take on the Symmetrix mirrors” *before* and *after* the backup. Click Next to take the defaults.

“Before Backup” specifies what Symmetrix Connect should do with mirrors *before* backup starts:

- **Bring mirrors down**
Expect to find the mirrors established and split the mirrors.
- **Verify that mirrors are down**
Expect to find the mirrors split and require the operator to verify that mirrors are split.
If used with “Leave mirrors up” after backup, the operator must manually split the mirrors before each backup.
- **Bring mirrors down if needed**
Expect to find the mirrors either split or established. If established, split the mirrors.
- **Bring mirrors down after establishing**
Expect to find the mirrors either split or established. If split, re-establish and then split the mirrors. If established, split the mirrors.

“After Backup” specifies what Symmetrix Connect should do with mirrors *after the backup* is completed:

- **Leave mirrors up**
Reestablish the mirrors.
 - **Leave mirrors down**
Do not reestablish the mirrors. *If used with “Bring mirrors down” before backup, the operator must manually establish the mirrors before the next backup.*
 - **Leave mirrors as found**
Reestablish the mirrors if they had been established before the backup started.
11. If you did not select Leave mirrors down, choose if the EDM should wait for tracks to resynchronize before displaying a message that the backup has finished. Click Next.
 12. You are asked “Do you want to run a script on the client after mirrors are split, but before the backup begins?” This is optional.
 - Do *not* select this for the *pre- and post- offline scripts* described in “Writing Scripts for the Unix Filesystem Interface” on page 2-31. Click Next.
 - Select this only for the script described in “Writing a Post-Mirror-Split Script” on page 2-30. Click Next.
 13. If you are using a post mirror-split script:
 - a. Enter the *full pathname* for the script. You can either browse existing files or enter the name. (Notice that the default is only a directory, not a file.)
 - b. Enter the username (UNIX account) to run it. The default is root.Click Next.
 14. If there is more than one media type available on the EDM, you are now prompted to “Select a media type,” for example, DLT tape. Click Next.

15. You are now prompted to “Select the number of trails and streams for this configuration.” In “Trails:”, enter the number of trails to use concurrently for this EDM Symmetrix Connect backup. In “Number of streams per trail:”, enter the maximum number of concurrent data streams per trail for the EDM to read directly from the Symmetrix volumes.
16. Enter the base names for the Symmetrix Connect work item and trail. For the Symmetrix Connect’s UNIX filesystem interface:

The default work item base name is:
dc_filesystem_mirroring_clientname.

The default trail base name is *dc*.
17. You are asked if you want to run a configuration check without running a backup. This will check whether the backup configuration is valid by performing a test backup (by running **eb_dc_backup -check**).

WARNING: Be aware that this validation process will split and re-establish the mirrors as configured. This could alter the data on the mirrors if they are already split. You do not get a choice to skip validation of the mirror handling parameters in the wizard, but you do get that choice with the **eb_dc_backup -check** command.

CAUTION: The configuration check can take hours to complete; it is started before you complete the remaining wizard steps and you cannot finish until the check finishes.

Note: Skip the check if you have multi-BCVs and you have yet to create the .sets file (see “Creating a .sets File for Multi-BCVs” on page 2-19).

Choose No to skip this check. You are warned to run an **eb_dc_backup -check** *work_item_base_name* before running a backup.

Choose Yes to run the configuration check:

- a. You are prompted “Do you want verbose output?” Click Yes or No and click Next.
 - b. If you decided to run the configuration check, you are prompted to click Run to begin the check.
 - c. Once you click Run, you have the option to cancel while the check occurs.
 - d. When the check is finished, you have the option to save the output.
18. Choose if you want to enable media duplication.
 19. Choose the backup rotation period.
 20. Choose the backup retention periods.
 21. Choose Yes and set the time if you want the EDM to regularly run the backup job. Choose No if you are using another scheduling tool or if you want to manually edit the crontab file.
 22. Choose Yes if you want the EDM to reschedule and reprocess any backups that fail in the current backup run.
 23. Choose if you want to append new data to the current tape volume until the first day of a new rotation or if you want to use a new volume for each new backup (which enables you to eject the backup media for the previous backup).
 24. Enter the name to use to identify this configuration for whenever you want to modify it in this backup configuration wizard (it is not used for anything else). For the Symmetrix Connect UNIX filesystem interface, the default configuration name is `dc_filesystem_mirroring_clientname`.
 25. You might receive a warning if the configuration will cause changes you need to know about. Click OK.
 26. The last panel summarizes your configuration. Click Save to complete the configuration.

The backup configuration for the UNIX filesystems is now complete.

Creating a .sets File for Multi-BCVs

With the Multi-BCV feature, up to eight BCVs can be configured for each standard device. If multiple BCVs are being used for the devices on which your database resides, you need to configure Symmetrix Connect; you need to specify which of the BCVs to use for backups and restores.

To specify the BCVs, you must create a new text file for each backup configuration. The “.sets” file specifies which BCVs (per standard device) should be used to take a backup and for restores.

This prevents backups and restores from failing or from using a BCV device that was not intended for backup or restore. For more about what occurs without this file, see “The Importance of the .sets File” on page 2-22.

The text file must specify a series of standard-to-BCV matchups, with multiple BCVs allowed for each standard. For each standard, the file specifies either the one BCV — or the rotation of BCVs — to use for each subsequent backup.

For example:

STD	set1	set2	set3	set4	set5	set6	set7	set8
#--	----	----	----	----	----	----	----	----
001	011	021	031	041	051	061	071	081

The pathname for the file is:

`/usr/epoch/EB/config/work_item_base_name.sets`

For restores, you also need to configure the devices that would be used in either a “no-mirror restore with snapshot” or an alternate-mirror restore, too, not just the BCVs used for backups.

When TimeFinder is used with SRDF, BCVs can reside on both the source Symmetrix system and a remote target Symmetrix. Up to eight additional BCVs can be paired with each SRDF “target” or “R2” device.

Basic Procedure to Create a .sets File

Before you can start, you need to obtain the device IDs for the standard devices and their BCVs. Consult your Symmetrix documentation for instructions.

Note: This basic procedure assumes a single Symmetrix. If you have a multiple Symmetrix configuration, review this procedure first and then refer to “Format Variations for the .sets File” on page 2-21.

To create a basic .sets file:

1. For each backup configuration (each different *work_item_base_name*), create a file named *work_item_base_name.sets* in */usr/epoch/config/* .
2. In the first column of the file, list all the device IDs for at least those standard devices that you expect might have data involved with this backup (plus any additional standard devices as well). For example,

```
001
002
003
004
```

3. In the second column (separated by one or more spaces and/or tabs), list the corresponding BCV device IDs for each STD. For example,

```
001      011
002      012
003      013
004      014
```

4. Optionally add additional columns for additional BCV sets to be used for backups. You need to do this if you want to implement a rotation policy for the multi-BCV backups. For example,

```
001      011      021      031      041
002      012      022      032      042
003      013      023      033      043
004      014      024      034      044
```

5. Optionally add headings containing the set names to the columns. If you do add headings, you must name the first column “STD.” If you do not add headings, the system defaults are the names shown in the following example:

```
STD      set1  set2  set3  set4  set5  set6  set7  set8
```

6. Save the file.

Format Variations for the .sets File

For multiple Symmetrix configurations, there are various format variations for the .sets file. These variations enable you to specify Symmetrix serial numbers along with the standard device ids.

Note: Before you can start, you need to obtain the Symmetrix serial numbers and the device IDs. Consult your Symmetrix documentation.

The full format for the file is:

```
[<Symm Serial#>:]<STD Device ID>  [<Symm Serial#>:]<1st BCV Device ID>  [[Symm Ser#>:]<nth BCV Device ID>]
```

Note: Only BCVs are allowed. You cannot include SRDF “R2” target devices.

This breaks down into the following variations (for illustration purposes, only one BCV set is shown):

- Multiple Symmetrix, Local BCV Format

```
<Symm Serial#>:<STD Device ID>      <BCV Device ID>
```

- Single Symmetrix for Primary Data, Remote BCV Format

```
<STD Device ID>  <Symm Serial#>:<BCV Device ID>
```

- Multiple Symmetrix, Remote BCV Format

```
<Symm Serial#>:<STD Device ID>  <Symm Serial#>:<BCV Device ID>
```

The following sections explain each of these variations and provides examples.

Specifying Multiple Symmetrix, Local BCV

For backups from Local BCVs on multiple Symmetrix units, follow the basic procedure, except: In the first column, include the Symmetrix Serial Number prior to any STD device IDs (separated by a colon) to limit the use of this device ID to that particular Symmetrix. For example,

```
000182600431:001      011
000182600431:002      012
000182600433:001      011
000182600433:002      012
```

**Specifying Single Symmetrix for
Primary Data, Remote BCV**

For backups from Remote BCVs corresponding to database devices on a single Symmetrix, follow the basic procedure, except: In the second column, include the Symmetrix Serial Number before the corresponding BCV device IDs (separated by a colon). For example,

005	000182600432:011
006	000182600432:012
007	000182600432:013
008	000182600432:014

**Specifying Multiple Symmetrix,
Remote BCV**

For backups from Remote BCVs corresponding to database devices on more than one Symmetrix, follow the basic procedure, except: In the first column, include the Symmetrix Serial Number prior to any STD device IDs (separated by a colon) to limit the use of this device ID to that particular Symmetrix. And in the second column, include the Symmetrix Serial Number before the corresponding BCV device IDs (separated by a colon). For example,

000182600431:001	000182600432:011
000182600431:002	000182600432:012
000182600433:001	000182600434:011
000182600433:002	000182600434:012

**The Importance of
the .sets File**

This file prevents EDM from failing or from using a BCV device that was not intended for backup or restore.

Without this file, EDM Symmetrix Connect uses the established mirror BCV or, if mirrors are not established, the most recently established BCV for backups and restores. This might cause unintended results.

But backups fail if the established or most recently established BCV is not visible to the EDM. (Under the same circumstances, restores to the same-mirror-as-backup will fail.) If this occurs, create the .sets file and within it, specify the BCV devices to use.

Configuration sessions fail in the Backup Configuration Wizard if there is no .sets file and you specify that multiple BCVs are present for Local BCV and Remote BCV backup mirror configurations. If this occurs, create the .sets file and within it, specify the BCV devices to use, and rerun the configuration.

CAUTION: Backup configuration sessions do not prompt you about whether you intend to perform restores in a Multi-BCV environment. Therefore, you can successfully configure no-mirror and SRDF backups but encounter errors from a missing .sets file if you try to perform no-mirror restore-with-snapshot or alternate-mirror restores in a multi-BCV environment without a .sets file.

As a result, you might encounter these specific restore problems:

- A no-mirror restore with a snapshot fails if there are multiple BCVs for the standard device but there is no .sets file. If this occurs, create the .sets file and within it, specify the BCV devices to use for the snapshot.
- An alternate-mirror restore to Local BCV fails if there are multiple BCVs for the standard device but there is no .sets file. If this occurs, create the .sets file and within it, specify the Remote BCV devices to use.
- An alternate-mirror restore to Remote BCV fails if there are multiple BCVs for the SRDF device but there is no .sets file. If this occurs, create the .sets file and within it, specify the Remote BCV devices to use.

Writing Scripts for Oracle Offline Backups

EDM Symmetrix Connect's Oracle interface takes certain default actions before and after an Oracle offline backup, or in a mirrored configuration, before and after the splitting of mirrors for an offline backup. You can customize the before and after actions with your own, user-supplied scripts

The interface to Oracle shuts down the database instance with the following Oracle Server Manager commands: **shutdown immediate**, **startup restrict**, and **shutdown normal**.

When the mirrors are split (in mirrored configurations) or when the backup is complete (in non-mirrored configurations), the interface to Oracle starts up the database instance.

You can customize the before and after actions with your own, user-supplied scripts. But if the scripts exist, they must include the database shutdown and startup explicitly, because EDM Symmetrix Connect will no longer perform these actions by default.

If you choose to use the *pre-offline* script, Symmetrix Connect checks that the database instance is shut down before performing the offline backup. If the script exists, but the database instance has not been shut down, the backup will fail.

If you choose to use the *post-offline* script, Symmetrix Connect does not perform any checking.

To use a pre- and post-offline script, you need to do the following:

1. Put the scripts in the `$ORACLE_HOME/epoch/admin` directory on the database server.
2. Name the pre- and post-offline scripts as follows:
 - **offline_pre_sid.sh**
 - **offline_post_sid.sh**

Note: The EDM will not perform any default pre-offline actions if it finds a file named `offline_pre_sid.sh` in `$ORACLE_HOME/epoch/admin` directory on the database server. The EDM will not perform any default post-offline actions if it finds a file named `offline_post_sid.sh` in the directory.

3. Make sure the script is owned by the oracle user and is executable by the oracle user and group.
4. Do not assume any Unix environment variables will be inherited by the script. Explicitly set: `PATH`, `ORACLE_SID`, `ORACLE_HOME`, `LD_LIBRARY_PATH`, and any other environment variables you might need.
5. Use the following guidelines.

General Guidelines for Scripts

General guidelines for Symmetrix Connect scripts:

- Each script should be a shell script, whether Bourne shell (sh), C shell (csh), or Korn shell (ksh). All scripts require their respective headers. For example, for sh, the script needs the header `#!/bin/sh` to work correctly.
- Do not send large volumes of output to either stdout or stderr. If you do, these will be truncated. If more than a few lines are needed, set up and use your own log file.
- Return a 0 (zero) status to continue with the backup; return a non-zero status to fail the backup.

Sample Oracle Pre-Offline Scripts

The following sample is a basic pre-offline script (with embedded Server Manager commands) to bring the database instance offline. Add the additional actions desired for your site.

Figure 2-2

Basic Pre-Offline Script to Shut Down Database Instance

```
#!/bin/sh
ORACLE_SID=jared
export ORACLE_SID
ORACLE_HOME=/jared
export ORACLE_HOME
PATH=$ORACLE_HOME/bin:$PATH
export PATH
svrmgrl << EOF
connect internal
shutdown immediate;
startup restrict;
shutdown normal;
EOF
rc=$?
exit $rc
```

You will probably want to fail the backup because of errors occurring while shutting down the database instance. Therefore the following, more robust, script checks for errors in running Server Manager and in shutting down the instance. This script:

- Calls a separate sql command file to: connect, turn on spooling, shut down the instance, and turn off spooling.
- Checks the exit code for the sql command file, so if there is a failure in running Server Manager, the backup can be cancelled.
- Searches the spooled log file for error strings. If there are any, you could generate your own SNMP event.

Figure 2-3 Pre-Offline Script, With Error Checks

```
#!/bin/sh
ORACLE_SID=CAT815
export ORACLE_SID
ORACLE_HOME=/data/oracle8.1.5/app/oracle/product/8.1.5
export ORACLE_HOME
PATH=$ORACLE_HOME/bin:/usr/bin:$PATH
export PATH
rm -f /tmp/offline.log
svrmgrl command=@/data/oracle8.1.5/app/oracle/product/8.1.5/scripts/offline.sql

rc=$?
if [ $rc -ne 0 ]
then
    ##### Your choice if you want to abort the backup because of a SVRMGRl failure #####
    exit $rc
fi
tmp='egrep '^ORA-[0-9]*:' /tmp/offline.log'
if [ ! -z "$tmp" ]
then
    ##### Generate Your SNMP Event #####
fi
rm -f /tmp/offline.log
exit 0
```

offline.sql

```
connect internal
spool /tmp/offline.log
shutdown immediate;
startup restrict;
shutdown normal;
spool off
exit
```

Sample Oracle Post-Offline Scripts

While you probably do not want to fail the backup because of errors occurring while starting up the database instance, you might want to generate an SNMP event to notify you that this occurred.

The following script, therefore, checks for errors in running Server Manager and in opening the database. This script:

- Calls a separate sql command file to: connect, turn on spooling, start up the instance, and turn off spooling.
- Checks the exit code for the sql command file. If there is a failure in running Server Manager, you could choose to have the backup cancelled.
- Searches the spooled log file for error strings. If there are any, you could generate your own SNMP event.

Figure 2-4 **Post-Offline Script, With Error Checks**

```

#!/bin/sh
ORACLE_SID=CAT815
export ORACLE_SID
ORACLE_HOME=/data/Oracle8.1.5/app/oracle/product/8.1.5
export ORACLE_HOME
PATH=$ORACLE_HOME/bin:/usr/bin:$PATH
export PATH
rm -f /tmp/online.log
svrmgrl command=@/data/Oracle8.1.5/app/oracle/product/8.1.5/scripts/online.sql

rc=$?
if [ $rc -ne 0 ]
then
    ##### Your choice if you want to abort the backup because of a SVRMGRl failure #####
    exit $rc
fi
tmp='egrep '^ORA-[0-9]*:' /tmp/online.log'
if [ ! -z "$tmp" ]
then
    ##### Generate Your SNMP Event #####
fi
rm -f /tmp/online.log
exit 0

```

alter_online.sql

```

connect internal
spool /tmp/online.log
startup;
spool off
exit

```

Writing a Post-Mirror-Split Script

A post-mirror-split script is similar to the post-offline script for Oracle and the “post” script for UNIX filesystems, but there are some differences.

- It differs from the post-offline script for Oracle as follows:
 - You can name the script and its directory location as you wish (as opposed to using the specific pathname `$ORACLE_HOME/epoch/admin/offline_post_sid.sh`).
 - If you specify one, it will run for *online* backups as well as for offline backups (instead of just offline backups).
 - If used in conjunction with the “post-offline” script, this script is executed immediately after the “post-offline” script.
 - The script can be run as root (instead of the Oracle user).

Note: The script must be run as the Oracle user if it is to start up the Oracle instance.

- It differs from the “post” script for UNIX filesystems in that:
 - You can name the script and its directory location as you wish (as opposed to using the specific pathname `filesystem_root/.emc_edm/EMC_OFFLINE_POST`).
 - If used in conjunction with the “post” script, this script is executed immediately after the “post” script.

To use a post-mirror-split script, you need to do the following:

1. Name the script and its location while configuring the backup in the wizard.
2. In the same wizard panel, name the UNIX user account to run the script. The script can be run as root.
3. Write the script. Follow the guidelines in “General Guidelines for Scripts” on page 2-25.

Writing Scripts for the Unix Filesystem Interface

The interface to UNIX filesystems takes certain default actions before and after backups (in non-mirrored configurations) or the splitting of mirrors. You can add custom before and after actions by specifying them in “pre” and “post” scripts.

With interface to UNIX filesystems, your custom pre and post scripts enable you to take actions to ensure data consistency for applications that store data in many interrelated files or in a few very large files. (See “Backup Interface to UNIX Filesystems” on page 1-28 for more background information.)

Your scripts could implement offline backups by shutting down and starting up your application.

For mirrored configurations only, the scripts could possibly put your database or application into (and out of) an online backup mode, if such a mode is available.

Note: For the UNIX filesystem interface of Symmetrix Connect, there is no “online” or “offline” backup mode that you can pick during the configuration.

With *or without* the use of scripts, the interface *always* performs the following actions. In mirrored configurations, the filesystem interface always flushes the filesystem I/O buffer immediately before splitting mirrors to pick up changes that haven’t been synchronized to disk. In non-mirrored configurations, the filesystem interface always unmounts the filesystem for the duration of the backup.

Unlike the interface to Oracle, there are no default actions omitted when you use the scripts. (Also, unlike the interface to Oracle, there is no checking for a state after which the backup fails.)

You can add custom before and after actions with your own, user-supplied scripts.

Note: If a script acts on multiple filesystems, you only need to place the script in any one of the affected filesystems.

To use the pre and post scripts for the interface to UNIX filesystems, you need to do the following:

1. Name the pre and post scripts as follows:
 - *filesystem_root/.emc_edm/EMC_OFFLINE_PRE*
 - *filesystem_root/.emc_edm/EMC_OFFLINE_POST*

Note: The scripts must be named exactly as shown.

For example, if the mounted filesystem is /home/doc, the scripts should be located in the directory:
/home/doc/.emc_edm/

2. Make sure the script is owned by the root user or the application user and is executable by the root or application user and group.
3. Do not assume any Unix environment variables will be inherited by the script. Explicitly set PATH and any other environment variables you might need.
4. Follow the guidelines in “General Guidelines for Scripts” on page 2-25.

3 Running and Monitoring Backups

For the Oracle interface and the UNIX Filesystem interface, you initiate Symmetrix Connect backups from the EDM.

This chapter covers the following topics:

- Symmetrix Connect Backups
- Backing Up with the Backup Activity Wizard
- Backing Up with the `eb_dc_backup` Command
- Checking Log Files on Client and Server
- Stopping a Backup
- Monitoring Tape Operations

Symmetrix Connect Backups

You can start Symmetrix Connect backups, for either the Oracle interface or the UNIX filesystem interface, as follows:

- automatically from the cron scheduler on the EDM (see “Crontab Entry to Schedule Backups” on page D-6)
- interactively with the Backup Activity Wizard
- with the **eb_dc_backup** command on the EDM, either interactively or in scripts


Symmetrix Connect backups using the Oracle interface cannot be initiated from the Oracle database server.

Note: Before you run the Oracle backups, the database instance must be up and running.


Backing Up with the Backup Activity Wizard

To manually initiate a Symmetrix Connect backup from the EDM with the Backup Activity Wizard:

1. From the EDM main window's Backup menu, select Activity Wizard.
2. Read the opening panel and click Next. The next panel asks which backup operation you want to perform.




3. Select Start Backups and click Next. The next panel prompts you to select the backup objects to be started.
4. Scroll until you find the name of **the EDM**:
 - *Not the database server* (Oracle interface)
 - *Not the filesystem host* (UNIX filesystem interface)
5. Expand the EDM's sublist by clicking on the "+".



6. Select the “kicker” work item that initiates the backups for the configuration — *work_item_base_name_k* — by checking the box. (See “Symmetrix Connect Kicker Work Item” on page D-3.) Click Next.



7. For backup level, always select “Full” for Symmetrix Connect with either the interface to Oracle or the UNIX Filesystem interface.
8. Specify whether to append to existing media or use new media.
9. Verify that your choices are correct. Click Finish to initiate the backup or Prev if you want to go back and change something.



Backing Up with the eb_dc_backup Command

To manually initiate a backup from the EDM with the eb_dc_backup command:

1. Log on to the EDM as root.
2. At a command prompt, enter the **eb_dc_backup** command.

The syntax for backing up a specific work item is:

```
eb_dc_backup [-mirror_set multi-BCV_set_name] [-new] [-retry] workitem_basename
```

-mirror_set *multi-BCV_set_name*

Use this Multi-BCV mirror set for the backup, overriding either the fixed-policy set name or the set that would be due according to the round-robin policy.

-new

Write to a new media set.

-retry

Reschedule and reprocess failed work-item in the current backup. (Symmetrix Connect does one retry.)

3. Use the Backup Report window, or use the terminal window where you ran the eb_dc_backup command, to check on the status of the backup.

Checking Log Files on Client and Server

Examine the logged status messages from the backup processing in the following files on the EDM:

- /usr/epoch/EB/log/*basetrailname*.log
is the Symmetrix Connect log for the backup processing of the specific *basetrailname*. This is a cumulative log, so no client data is overwritten.
- /usr/epoch/EB/log/backups.log
is the EDM backup log for the processing of all backups

Note: For information on restore processing, look in:
/usr/epoch/EB/log/recoveries.log

Stopping a Backup

To stop a backup in progress, use the Backup Activity Wizard:

1. From the EDM main window's Backup menu, select Activity Wizard.
 2. Read the opening panel and click Next.
 3. Choose to either stop individual active backups or to stop all backup processing and click Next.
 4. If you are stopping all backup processing on the EDM, click Finish.
 5. To stop an individual backup, select the "kicker" work item that initiated the backups for the configuration — *work_item_base_name_k*. Click Next.
 6. Verify that your choice is correct and click Finish.
-

Monitoring Tape Operations

To monitor tape operations, use the EDM graphical user interface's Library Unit Manager window as your central console. You are prompted whenever operator intervention is required (for example, to supply a fresh batch of backup tapes).

Note: For more information, see online help on Library Unit Manager Overview and on Library Unit Manager Tasks.

4 Running Restores

For the Oracle interface and the UNIX Filesystem interface, you initiate Symmetrix Connect restores from the EDM.

This chapter covers the following topics:

- Symmetrix Connect Restores
- Restoring Old .dcrpt and .ddtab Files
- Restoring Oracle to an Alternate Destination
- Restoring UNIX Filesystems to an Alternate Destination
- Restoring with the Restore Window
- Restoring with the eb_dc_restore Command
- Restores to Alternate Host for Compaq and Sequent
- Restoring AdvFS Filesystems

Symmetrix Connect Restores

You can start restores of Symmetrix Connect backups, for either the Oracle interface or the UNIX filesystem interface, with either the EDM's Restore window or with the **eb_dc_restore** command on the EDM.

Note: When restoring backups taken more than 30 backups ago, you must first perform a network restore (of the work item base name's .dcrpt and .ddtab files) using the **ebrestore** command on the EDM (not the Restore window or eb_dc_restore). See "Restoring Old .dcrpt and .ddtab Files" next.

Oracle archived redo logs and control files are actually restored over the network or Symmetrix Path, but they can be restored with the EDM Restore window or the eb_dc_restore command.

Symmetrix Connect restores of Oracle interface backups cannot be initiated from the Oracle database server.

At the beginning of a restore to a Symmetrix mirror, EDM unmounts filesystems to prevent corruption of filesystems that could occur after eb_dc_restore, when the Symmetrix mirror is restored to the standard device.

Note: See "Restores to Alternate Host for Compaq and Sequent" on page 4-20.

Note: See "Restoring AdvFS Filesystems" on page 4-20.

Restoring Old .dcrpt and .ddtab Files

If you are restoring a backup older than the most recent thirty backups, you must perform a network restore of the work item's .dcrpt and .ddtab files for that particular backup.

These files are always backed up to tape, but the files for the particular backup date must reside on the EDM (on disk) for the restore to work.

Note: Symmetrix Connect restores rely on EDM-archived .dcrpt and ddtab files to link Symmetrix Connect information to EDM backup catalog information. See "The .dcrpt and .ddtab Files" on page D-9.

To perform a network restore of the work item's .dcrpt and ddtab files:

1. Run **ebrestore** on the EDM.
2. Select the EDM as the source host to restore from.
3. Select the EDM as the destination host to restore to.
4. Choose the appropriate kicker work item that you want to restore from *work_item_base_name_k*.
5. Look in the /usr/epoch/EB/Symmetrix_Connect/Archives/*work_item_base_name* directory.
6. Select the *date-time_stamp.dcrpt* and *date-time_stamp.ddtab* files for the date and time of the backup. For example:
 2000.06.15.14.06.19.dcrpt
 2000.06.15.14.06.19.ddtab
7. Restore in place.

Restoring Oracle to an Alternate Destination

Out-of-place restores follow the same pathnames as in-place restores *except* that the location you specify in Directory is *pre-pended* to the in-place pathname. You must create symbolic links from the old pathnames to the new pathnames.

In-place restores occur as follows:


- *Datafiles* that were backed up *as files* are restored to their *original file pathnames*. The pathnames may contain links.
 - *Filesystems* backed up *as devices* (raw filesystems) are restored to their original *mount point*, unless the “Data Path” is “Network/ Symmetrix Path”, in which case they are restored to the original *device*. The original mount point may contain links.
- Note:** Filesystems backed up on Compaq Tru64 and Sequent are always restored to the *device*.
- Devices are restored to their original *device* pathnames, for example /dev/vol01/tool01 is restored to /dev/vol01/tool01. You cannot have links in the /dev paths.

Figure 4-1

Restoring Oracle To a Device In-Place and Out-of-Place

If a Single Datafile on Raw Partition
or on Logical Volume was...

/dev/rdisk/c0t0d0s3



Or if Datafiles in Filesystem were...

/oradata1/prods01/datadbs01.dbf
/oradata1/prods01/datadbs02.dbf
/oradata1/prods01/datadbs03.dbf



...and if restore **Data Path** is "Network/Symmetrix Path"
(or for any restore of Compaq Tru64 or Sequent)

Then, the Restore is To a Device

You can Restore to a Device, either:

In-Place, which is to the original device pathname

/dev/rdisk/c1t1d1s3




Or, out-of-Place (alternate directory), which is to a new directory path

/tmp

pre-pended to the original device pathname

/dev/rdisk/c1t1d1s3

where the original device name is a symbolic link to the new device



To restore a device or filesystem-as-device backup *to a device* out-of-place:

1. Open a shell window.
2. Go to the location you will specify in the Directory field of the Destination Options dialog box (step 8 on page 4-13), for example:

```
# cd /tmp
```

3. From that point, create any directory structures from the pathname of the backup as needed so that the entire pathname of the original backup exists. For example, if the backup was of /dev/rdsk/c1t1d1s3:

```
# mkdir dev
# mkdir dev/rdsk
```

4. Go to the *directory at the end of that path*, in the example:


```
# cd dev/rdsk
# pwd
/tmp/dev/rdsk
```

5. Create a symbolic link from the backed-up device to the new device, in the example:

```
# ln -s /dev/rdsk/c2t2d2s4 c1t1d1s3
# ls -l
lrwxrwxrwx 1 root root 55 Sep 7 20:42
c1t1d1s3 -> /dev/rdsk/c2t2d2s4
```

Note: If the backup was done using the raw filesystem functionality, the new location can be a sub-directory, (e.g., /tmp/dev/rdsk) but the link must point to the raw device for the filesystem. The entire filesystem will be overlaid onto the new device.

Figure 4-2 Restoring Oracle To a Mount Point or Directory Out-of-Place



To restore a **filesystem-as-device** (raw filesystem) to a *mount point* (available on Solaris, AIX, and HP only):

1. In a shell window, go to the location you will specify in the Directory field of the Destination Options dialog box (step 8 on page 4-13), for example:

```
# cd /tmp
```

2. Create a symbolic link from the backup mount point to the new mount point. For example,

```
# pwd
/tmp
# ln -s /oradata2 oradata1
# ls -l
lrwxrwxrwx 1 root root 7 Sep 7 20:16
oradata1 -> /oradata2
```

To restore **datafiles-as-files** to a new *directory* (available on Solaris, AIX, and HP only):

1. In a shell window, go to the location you will specify in the Directory field, for example:

```
# cd /tmp
```

2. From that point, create any needed directory structures that precede the link. For example, if the backup was of /oradata1/prods01/ datadbs01.dbf:

```
# mkdir oradata1
```

3. Create a symbolic link from the backup to the new location. For example,

```
# pwd
/tmp
# ln -s /oradata2/parts01 oradata1/prods01
# ls -l
lrwxrwxrwx 1 root root 7 Sep 7 20:16
oradata1/prods01 -> /oradata2/parts01
```

4. In the location pointed to by the link (parts01 in the Figure 4-2 example), create any additional directory structures needed to hold your files.

Restoring UNIX Filesystems to an Alternate Destination

Out-of-place restores go to the same pathnames as in-place restores *except* that the location you specify in Directory is *pre-pended* to the in-place pathname. Therefore you are responsible for creating symbolic links from the old pathnames to the new pathnames.

In-place restores of UNIX filesystems, which are backed up as devices, are restored to their original *mount point*, unless the “Data Path” is “Network/ Symmetrix Path”, in which case they are restored to the original *device*.

Note: Filesystems backed up on Compaq Tru64 and Sequent are always restored to the *device*.

To restore *to a device* out-of-place:

1. Open a shell window.
2. Go to the location you will specify in the Directory field of the Destination Options dialog box (step 8 on page 4-13), for example:

```
# cd /tmp
```

3. From that point, create any directory structures from the pathname of the backup as needed so that the entire pathname of the original backup exists. For example, if the backup was of /dev/rdisk/c1t1d1s3:

```
# mkdir dev
# mkdir dev/rdisk
```

4. Go to the *directory at the end of that path*, in the example:

```
# cd dev/rdisk
```

5. Create a symbolic link from the backed-up device to the new device, in the example:

```
# pwd
/tmp/dev/rdisk
# ln -s /dev/rdisk/c2t2d2s4 c1t1d1s3
# ls -l
lrwxrwxrwx 1 root root    55 Sep 7 20:42
c1t1d1s3 -> /dev/rdisk/c2t2d2s4
```

Note: The new location can be a sub-directory, (e.g., /tmp/dev/rdsk) but the link must point to the device for the filesystem. The entire filesystem will be overlaid onto the new device.

Figure 4-3

Restoring UNIX Filesystem To a Device In-Place and Out-of-Place

The Datafiles in Filesystem were...

/app1/users01/data01.dat
/app1/users01/data02.dat
/app1/users01/data03.dat



...Backed Up in
Filesystem-as-Device



/dev/rdsk/c0t0d0s3

If the restore **Data Path** is “Network/Symmetrix Path”
(or any restore for Compaq Tru64 or Sequent)

Then the Restore is To a Device

You can Restore to a Device, either:

In-Place, which is to the original device pathname

/dev/rdsk/c1t1d1s3



Or, out-of-Place, which is to a new directory path

/tmp

pre-pended to the original device pathname

/dev/rdsk/c1t1d1s3

where the original device name is a symbolic link to the new device




Figure 4-4

Restoring UNIX Filesystem To a Mount Point

The Datafiles in Filesystem were...

/app1/users01/data01.dat
/app1/users01/data02.dat
/app1/users01/data03.dat



...Backed Up in
Filesystem-as-Device



/dev/rdisk/c0t0d0s3

If the restore **Data Path** is “Symmetrix Connect”
(and if host is not Compaq Tru64 or Sequent)

Then the Restore is To a Filesystem Mount Point

You can Restore to a Mount Point, either:

In-Place, which is to the original mount point

/app1




Or, out-of-Place, which is to a new directory path
/tmp

pre-pended to the original mount point

/app1

where the original mount point

is a symbolic link to the new mount point



To restore to *mount point* (available on Solaris, AIX, and HP only):

1. Open a shell window.
2. Go to the location you will specify in the Directory field of the Destination Options dialog box (step 8 on page 4-13), for example:


```
# cd /tmp
```

3. Create a symbolic link from the backup mount point name to the new mount point. For backed-up filesystems, you can create a symbolic link to the mount point, that is, create the link at the beginning of the original backup path. For example,

```
# pwd
/tmp
# ln -s /app2 app1
# ls -l
lrwxrwxrwx 1 root  root    7 Sep 7 20:16
oradata1 -> /oradata2
```

Restoring with the Restore Window

- To restore a Symmetrix Connect backup of Oracle using the EDM Restore window:
1. Take the Oracle database or tablespace offline.
 2. On the EDM, bring up the Restore window, either from the EDM main window (Backup > Restore) or by entering the **edmrestore** command on the EDM.



3. In the browser panel (upper left), expand the database server and the work item you want to restore:

4. To restore the full database (and its archived redo logs and its control file), expand the work item in the browser and select the database.

To restore selected tablespaces, selected datafiles, the archived redo logs, or the control file, expand the database object in the browser and select the desired object in the list on the right.

5. Examine the Mark Summary Tab and verify that you have selected the object(s) you want to restore.
6. Examine the Media Tab for a list of backup volumes needed. If any are offline, the Library Unit Manager will open the Media Request window when the media is needed.
7. Press the Start Restore button.
8. In the Destination Options window, the Overwrite Policy is initially set to Always, the Data Path to Symmetrix Connect, and the Restore Location to In Place.

EDM(tippu): Destination Options

Overwrite Policy

- ☒ Always
- ☐ Older Files Only
- ☐ Never

Data Path

- ☐ Network / Symmetrix Path
- ☒ Symmetrix Connect

Restore Location

- ☒ In Place
- ☐ Specified Location:
 - Host Name:
 - Directory:

- a. For alternate destination restores:


Enter in Directory the path that will be pre-pended to the original backup pathname.

Create directories and the link to the original backup pathname as directed in “Restoring Oracle to an Alternate Destination” on page 4-3.

- b. Change the options in Overwrite and Data Path categories as needed.


Note: Depending on options that you select in one category (for example, data path) other options might become available in another (for example, in overwrite policy). Also note the dependency on Data Path for how filesystems-as-devices are restored, both in-place and out-of-place (see “Restoring Oracle to an Alternate Destination” on page 4-3.)

9. Next, select the mirroring type.



Note: See “Alternate Mirror Restores” on page 1-26 and “Restoring to No-Mirror with Snapshot in Mirrored Configurations” on page 1-26 for descriptions of these features. See also “Creating a .sets File for Multi-BCVs” on page 2-19.

10. If you selected no mirror, select the snapshot policy.



11. If there are multi-BCVs (for mirror types or snapshots to Local BCV and Remote BCV), you are prompted for which BCV set to use.

Note: If the restore BCV set is not defined in the .sets file, you are prompted for a correct one or if you want to cancel the restore. (Also, if you receive errors relating to a missing .sets file for Multi-BCVs, see “The Importance of the .sets File” on page 2-22.)

12. For all mirror types except No Mirror, select the pre-restore mirror policy (what you want to do with mirrors before restoring data):


- **Bring Mirrors Down**
Expect to find the mirrors established and split them.
- **Bring Mirrors Down if Needed**
Expect to find the mirrors either split or established. If established, split the mirrors.
- **Verify Mirrors are Down**
Expect to find the mirrors split and require the operator to verify that mirrors are split.

13. Unless you selected “No-mirror” and “No snapshot copy”, select the post-restore mirror policy (what you want to do with mirrors after restoring data):

- **Leave Mirrors Down**
Do not reestablish the mirrors.
- **Leave Mirrors Up**
Reestablish the mirrors.
- **Leave mirrors as found**
Reestablish the mirrors if they had been established before the restore.

14. If you did not select Leave mirrors down, choose the “Resync policy” (whether or not the EDM wait for tracks to resynchronize before displaying a message that the restore has finished).

15. If you selected the Data Path “Symmetrix Connect”, select whether you want to be prompted before writing the restore to disk.
16. Verify that you want to restore.



If the restore fails, the status bar turns red and a Restore Failed! popup window appears with the work item name and the reason for the failure.

If you want to cancel the restore, click the Cancel button on the Restore Progress window and reply Yes in the Cancel Received pop-up.

17. If you selected to be prompted before writing the restore to disk, a Commit Selection window asks if you want to commit or discard the restore.
18. Close the EDM Restore window before your next backups are scheduled to run.

Restoring with the eb_dc_restore Command

As an alternative to the EDM Restore graphical interface, you can use the eb_dc_restore command to restore a Symmetrix Connect backup of Oracle.

To use eb_dc_restore command on the EDM server to restore an Oracle backup, you need to:

1. Take the Oracle database or tablespace offline.
2. Identify the objects to restore, using one of the eb_dc_restore options that list such objects.
3. Run eb_dc_restore with the appropriate command option(s).

Identifying Objects to Restore

Use one of the following commands to identify backups available to be restored:

- eb_dc_restore -listnames *pattern*
Lists work items whose base names match the pattern.
- eb_dc_restore -listdates *work_item_base_name*
Lists backups for that work item base name.
- eb_dc_restore -listfiles *work_item_base_name* [*date*]
Lists filesystems backed up for that work item base name.

The following example gives sample output for eb_dc_restore -listfiles, showing restorable objects for the most recent backup for the work item base name "rbcv."

Figure 4-5

Sample Output from eb_dc_restore -listfiles

```
[roadster] # eb_dc_restore -listfiles rbcv
Listing files from the backup dated 07/18/2000_11:03:17
-----file-----size-----status-----
Tablespace "RBS"
/u02/oradata/iota0/rbs01.dbf                23.9 MB      Good
Tablespace "MBCV_TS"
/rbcv/rbcv_ts.dbf                          504.0 KB     Good
Control files
/u01/app/oracle/product/8.1.5/dbs/cntrliota0.dbf.000718110055
                                                                Good
```

Restore Command Options

The `eb_dc_restore -restore` command restores entire databases, individual tablespaces, archived redo logs, and filesystems. This section describes sample commands for restoring.

Restoring the Entire Database

To restore the entire work item, use the `-ALL` option and specify the work item base name.

```
edm# eb_dc_restore -restore rbcv -ALL
```

Restoring a Tablespace

To restore a tablespace, use the `-TS` option and specify the tablespace name. The following command is an example based on the sample in Figure 4-5 on page 4-17.

```
edm# eb_dc_restore -restore rbcv -TS PROJ
```

Tablespace names that include lower case characters or spaces must be quoted, and the quotes must be escaped. In the following example, the tablespace name is “proj” (all lowercase), so the argument to the `-TS` option is specified:

```
edm# eb_dc_restore -restore rbcv -TS \"proj\"
```

Restoring Data Files

To restore the data files for a database, use the `-DF` option.

```
edm# eb_dc_restore -restore rbcv -DF /u02/oradata/iota0/rbs01.dbf
```

Restoring Control Files

To restore the control file, use the `-CF` option.

```
edm# eb_dc_restore -restore rbcv -CF
```

Restoring Archived Redo Logs

To restore archived redo logs, use the `-AR` option.

```
edm# eb_dc_restore -restore rbcv -AR
```

Restoring to a Different Directory

To restore to a different directory over the network, specify the new location with the option **-destination** *directory*.

Note: See “Restoring Oracle to an Alternate Destination” on page 4-3 and for required manual steps to create the destination directory. See also “Restoring UNIX Filesystems to an Alternate Destination” on page 4-8.

Restoring to an Alternate Mirror Type

To restore to an alternate mirror type, specify the mirror type with the option **-mirror** *mirror_type*.

Note: See “Alternate Mirror Restores” on page 1-26 for a description of this feature.

Restoring to No Mirror With Snapshot

To restore to “no-mirror with snapshot,” specify the snapshot-mirror type with the options

-mirror no_mirror -snapshot *snapshot_mirror_type*.

Note: See “Restoring to No-Mirror with Snapshot in Mirrored Configurations” on page 1-26 for a description of this feature.

Setting Mirror Policy for When Restore is Done

To set the mirror policy for when the restore is done, specify the policy by using one of the options from this list:

-est_then_split, **-leave_mirrors_up**, **-leave_mirrors_up**, or **-leave_mirrors_as_found**.

Restoring With Multi-BCVs

To restore backups to a mirror in a multiple-BCV backup environment, specify the particular BCV set to receive the restored data with **-mirror_set** *multi-BCV_set_name*.

Note: See “Creating a .sets File for Multi-BCVs” on page 2-19.

Restores to Alternate Host for Compaq and Sequent

For Compaq Tru64 or Sequent restores to a different host name (cross-restore), the destination system must match the system that was backed up with respect to striping and volume sizes.

If striping was used, the size and number of disks the new volume is striped across must be the same.

If striping was not used, the sizes of the new volumes must be the same or larger.

Restoring AdvFS Filesystems

Symmetrix Connect allows you to specify datafiles, tablespaces, or databases for restore. However, in the case of AdvFS filesystems on Compaq Tru64, the entire AdvFS domain corresponding to the datafile/tablespace/database is restored, not any single entity within it.

So consider the following guidelines when you attempt to restore any AdvFS filesystem:

- The restore proceeds only if the configuration of the target disks is identical to that of the backed up disks. During the restore of AdvFS filesystems, the number and sequence of disks in the domain must equal the number and sequence of disks that were backed up. The name of each disk used for the restore must match the name of the disk that was backed up. Finally, the size of each target disk should be equal to or greater than the corresponding size of the backed up disk.
- During a restore of AdvFS filesystems, the filesystem for the current fileset in the domain is unmounted and the backed-up image of the domain is restored, resulting in replacement of the current domain image.

- The restore process does not create AdvFS domains, so that in disaster situations where a host or disks are lost, the file domains to be restored to must first be recreated so that they exist, have the same name, fileset name, and configuration as they did at backup.
- During a restore, the Symmetrix Connect software temporarily disables the AdvFS domain being restored. AdvFS commands on this domain might report false errors until the restore is completed.

A Sample Configuration: Oracle No-Mirror

This appendix contains the Backup Configuration wizard's prompts for configuring Symmetrix Connect Oracle, no-mirror backups.

See these other appendices for other scenarios:

- Appendix B, Configuring Symmetrix Connect Oracle, Mirrored
- Appendix C, Configuring Symmetrix Connect UNIX Filesystem Backups

Table A-1 Configuring Symmetrix Connect Oracle, No Mirror	
Prompt	Description/Actions
Select the clients that you want to configure for backup:	Choose the network hostname of the client hosting the database to be configured for Symmetrix Connect backups with the Oracle interface.
What type of configuration do you want? ◆ New Configuration ◆ Modify Configuration ◆ Duplicate Configuration ◆ Delete Configuration	Choose: New Configuration (It's the <u>DEFAULT</u>) Delete Configuration erases an entry in the configuration file having a specified configuration name.
Select the object(s) to configure for backup.	Scroll down the list to find the Oracle database and click in the checkbox to select it.
Select the application to use for this configuration. —————Backup Application————— ◆ RMAN/EBU ◆ EMC Backint Backup Interface for SAP R/3 System ◆ EDM Symmetrix Connect	Choose: EDM Symmetrix Connect to choose the Symmetrix Connect interface to Oracle.
Do you want to use Symmetrix Path? ◆ Yes ◆ No	If the client has been enabled for Symmetrix Path, a panel will ask if you want to use Symmetrix Path. In this case, where you have already selected Symmetrix Connect backup, the prompt about Symmetrix Path is only for the archive logs and control files, not for the datafiles.
Do you want an online or offline backup? ◆ Online ◆ Offline	For this scenario, select Offline.

Table A-1

Configuring Symmetrix Connect Oracle, No Mirror

Prompt	Description/Actions
Do you want to back up the archived redo log directory? ◆ Yes ◆ No Note: Presented for offline backups only, because online backups always back up the archived redo log directory.	Do you want EDM to initiate backup of the archived redo logs along with the database, even though this is an offline backup? Choose No, since this is an offline backup.
Do you want to back up database datafiles located within filesystems as raw devices? ◆ Yes ◆ No	Before answering the prompt, ask yourself these questions: Are the datafiles located within filesystems? If not, choose No in the panel and click Next. (The datafiles will be backed up as devices anyway, but this question is about backing up <i>filesystems</i> as devices.) Is the platform Compaq Tru64 or Sequent? If yes, choose Yes in the panel and click Next. (You must choose Yes in order to back up the database.) Click Yes in the pop-up warning. Do you want to collectively back up all the datafiles in each filesystem as a single device? If yes, choose Yes in the panel and click Next. Click Yes in the pop-up warning. If your answer is no, it is because you want to back up the individual data files <i>as files</i> . Choose No in the panel and click Next. Note: See “Data Layout and Backup Granularity” on page 1-14.
Provide the following Oracle database information: TNS alias for the database Username Password	• Enter the TNS alias for the Oracle instance. • Enter the Username (typically “system”). • Enter the password.

Table A-1 Configuring Symmetrix Connect Oracle, No Mirror	
Prompt	Description/Actions
Select the tablespace(s) to configure:	The panel lists one item: the entire Oracle database. Either click in the checkbox to select it or double-click on the database icon to view components of the database and check them individually. In this case, click on the whole database.
Specify the UNIX account name for this Oracle database. UNIX Account: <u>oracle</u>	Enter the Unix username for the client's Oracle user.
Select the mirror configuration you want to use. ————— Mirror Configuration ————— ◆ Local BCV ◆ SRDF ◆ No mirror ◆ Remote BCV	Select which one of the following Symmetrix Configurations has already been set up: - Local BCV (TimeFinder) mirroring <u>DEFAULT</u> - SRDF mirroring. - No mirroring. - Remote BCV (TimeFinder mirroring of SRDF mirror). For this scenario, select No mirror.
Select the media type you want for this backup. Choose one of: _____ Note: Omitted if there is one media type available.	Choose the media type of library unit to be used for backup from the available types in the dropdown list.

Table A-1

Configuring Symmetrix Connect Oracle, No Mirror

Prompt	Description/Actions
Select the number of trails and streams for this configuration. Trails: Streams per trail:	- Enter the number of trails to use concurrently for the EDM Symmetrix Connect portion of this backup. Note: (Depending on how many network streams you selected and how many streams per trail you specify below, the first EDM Symmetrix Connect trails might or might not be also used for network backups.) Note: For optimal backup performance, the number of trails should equal the number of available drives of the media type you are selecting. It cannot exceed the server parameter for maximum concurrent trails for this media type. (The default is the value for the <code>server</code> parameter for maximum concurrent trails for this media type.) - Enter the Number of streams per trail. This is the maximum number of concurrent data streams per trail for database backup using EDM Symmetrix Connect. Note: For optimal <i>backup</i> performance, it could be more than 1, depending on disk layout and the use of a tuning file. The maximum is the value for the Media tab parameter for Limit Maximum Concurrent Backups for this media trail. For optimal <i>restore</i> performance, it should equal 1.

Configuring Symmetrix Connect Oracle, No Mirror

EDM Symmetrix Connect — Interface to Oracle, UNIX Filesystems

Table A-1

Configuring Symmetrix Connect Oracle, No Mirror

Prompt	Description/Actions
Do you want to enable media duplication for this media set? ◆ Yes ◆ No	Choose whether or not to use media duplication. (No is <u>DEFAULT</u>) Note: For this scenario, take the default.
How often do you want to backups to rotate? Rotate Every <u>14</u> days	This value determines the media rotation period. This is the duration after which, by default, EDM starts writing the backups to a new volume rather than appending the backups to the volume containing the previous backups. Note: For Symmetrix Connect, EDM does not perform automatic scheduling of full vs. incremental backups.
How long do you want to retain your backups? ◆ Forever ◆ For <u>1</u> years	Choose the period of time to retain your backups, their catalogs, and duplicates. <u>DEFAULT</u> is 1 year
Do you want EDM to initiate the backup? ◆ Yes <u>18:00</u> ◆ No	Choose Yes and set the time if you want the EDM to regularly run (from cron on a daily basis) the backup job (ebbbackup) for this backup configuration. Choose No if you are using another scheduling tool or if you want to customize manually the crontab file that controls this function.
Do you want to automatically retry failed backups? ◆ Yes ◆ No Note: Omitted if EDM not initiating backup regularly from cron .	Choose No. You would choose Yes if you want the EDM to reschedule and reprocess any backups that fail in the current backup run. The retries occur after all normally scheduled backups have completed.

Table A-1 Configuring Symmetrix Connect Oracle, No Mirror	
Prompt	Description/Actions
Do you want backup to append to the current volume or use a new one? ◆ Append to Current Volume ◆ Use a New Volume	Choose if you want to append new data to the current volume until the first day of a new rotation or if you want to use a new volume for each new backup (which enables you to eject the backup media for the previous backup).
What do you want to name the new configuration? Name:	Enter the name to use to identify this configuration, if you later want to modify, duplicate, or delete it using the Backup Configuration Wizard. The default name is the work item base name that you provided in the panel “Provide base names for the work item and trail.”
Summary of your configuration: Configuration options:	The last panel summarizes your configuration. Click Save to complete the configuration.

B Sample Configuration: Oracle Mirrored

This appendix contains the Backup Configuration wizard's prompts configuring Symmetrix Connect Oracle, mirrored (Local BCV) backups.

See these other appendices for other scenarios:

- Appendix A, Configuring Symmetrix Connect Oracle, No Mirror
- Appendix C, Configuring Symmetrix Connect UNIX Filesystem Backups

Table B-1 Configuring Symmetrix Connect Oracle, Mirrored	
Prompt	Description/Actions
Select the clients that you want to configure for backup:	Choose the network hostname of the client hosting the database to be configured for Symmetrix Connect backups with the Oracle interface.
What type of configuration do you want? ◆ New Configuration ◆ Modify Configuration ◆ Duplicate Configuration ◆ Delete Configuration	Choose: New Configuration (It's the <u>DEFAULT</u>) Delete Configuration erases an entry in the configuration file having a specified configuration name.
Select the object(s) to configure for backup.	Scroll down the list to find the Oracle database and click in the checkbox to select it.
Select the application to use for this configuration. —————Backup Application————— ◆ RMAN/EBU ◆ EMC Backint Backup Interface for SAP R/3 System ◆ EDM Symmetrix Connect	Choose: EDM Symmetrix Connect to choose the Symmetrix Connect interface to Oracle.
Do you want to use Symmetrix Path? ◆ Yes ◆ No	If the client has been enabled for Symmetrix Path, a panel will ask if you want to use Symmetrix Path. In this case, where you have already selected Symmetrix Connect backup, the prompt about Symmetrix Path is only for the archive logs and control files, not for the datafiles.
Do you want an online or offline backup? ◆ Online ◆ Offline	For this scenario, select Online.

Table B-1

Configuring Symmetrix Connect Oracle, Mirrored

Prompt	Description/Actions
Do you want to back up database datafiles located within filesystems as raw devices? ◆ Yes ◆ No	Before answering the prompt, ask yourself these questions: Are the datafiles located within filesystems? If not, choose No in the panel and click Next. (The datafiles will be backed up as devices anyway, but this question is about backing up <i>filesystems</i> as devices.) Is the platform Compaq Tru64 or Sequent? If yes, choose Yes in the panel and click Next. (You must choose Yes in order to back up the database.) Click Yes in the pop-up warning. Do you want to collectively back up all the datafiles in each filesystem as a single device? If yes, choose Yes in the panel and click Next. Click Yes in the pop-up warning. If your answer is no, it is because you want to back up the individual data files <i>as files</i>. Choose No in the panel and click Next. Note: See “Data Layout and Backup Granularity” on page 1-14.
Provide the following Oracle database information: TNS alias for the database Username Password	• Enter the TNS alias for the Oracle instance. • Enter the Username (typically “system”). • Enter the password.
Select the tablespace(s) to configure:	The panel lists one item: the entire Oracle database. Either click in the checkbox to select it or double-click on the database icon to view components of the database and check them individually. For this scenario, open the individual tablespaces and select one or more of them.

Table B-1 Configuring Symmetrix Connect Oracle, Mirrored	
Prompt	Description/Actions
Specify the UNIX account name for this Oracle database. UNIX Account: <u>oracle</u>	Enter the Unix username for the client's Oracle user.
Select the mirror configuration you want to use. ————— Mirror Configuration ————— ◆ Local BCV ◆ SRDF ◆ No mirror ◆ Remote BCV	Select which one of the following Symmetrix Configurations has already been set up: • Local BCV (TimeFinder) mirroring <u>DEFAULT</u> • SRDF mirroring. • No mirroring. • Remote BCV (TimeFinder mirroring of SRDF mirror).
Are multiple BCVs present for any standard device involved with the backup? ————— BCV Choices ————— ◆ Multiple BCVs present ◆ One BCV present for every device Note: Omitted for “no-mirror” and “SRDF”.	Specify if there are multiple-BCVs for any standard device involved in the backup; if yes, select the first choice, if no, select the second: • Multiple BCVs present • One BCV present for every device <u>DEFAULT</u> See “Creating a .sets File for Multi-BCVs” on page 2-19.
Select the mirror set policy you wish to use? ————— Mirror Set ————— ◆ Round Robin rotation of BCV sets ◆ Use One BCV set Note: Only received if you set “Multiple BCVs present”. Note: Omitted for “no-mirror”.	If multiple BCVs are present, select to either rotate through sets of multiple BCVs or to just use one specified set of BCVs: • Round Robin rotation of BCV sets • Use One BCV set <u>DEFAULT</u>

Table B-1

Configuring Symmetrix Connect Oracle, Mirrored

Prompt	Description/Actions
Name of the BCV mirror set you want to use? BCV Set Name <u>set 1</u> Note: Only received if you set “Use One BCV set”. Note: Omitted for “no-mirror”.	If you are specifying a particular set, enter the BCV Set Name now. For example, enter “set1”.
Select the actions to take on the Symmetrix mirrors: Before Backup ◆ Bring mirrors down ◆ Verify that mirrors are down ◆ Bring mirrors down if needed ◆ Bring mirrors down after establishing Note: Omitted for “no-mirror”.	What should software do with mirrors before backup starts: • Bring mirrors down Split mirrors. (The software expects mirrors to be currently established.) • Verify that mirrors are down Require user to verify that mirrors are split. (Mirrors expected to be already split.) <i>If used with Leave Mirrors Up after backup, user must manually split mirrors before each backup.</i> • Bring mirrors down if needed Split mirrors if mirrors are established and run backup. (All mirrors may be either established or already split.) • Bring mirrors down after establishing <u>DEFAULT</u> Automatically re-establish mirrors if already split, then split them and run backup. If mirrors found established, split mirrors and then back up.
After Backup ◆ Leave mirrors up ◆ Leave mirrors down ◆ Leave mirrors as found Note: Omitted for “no-mirror”.	What should software do with mirrors after backup completes: • Leave mirrors up Reestablish mirrors. <u>DEFAULT</u> • Leave mirrors down <u>DEFAULT</u> Do not reestablish mirrors. <i>If used with bring_mirrors_down before backup, user must manually establish mirrors before next backup.</i> • Leave mirrors as found Reestablish mirrors if they had been established before backup started.

Table B-1 Configuring Symmetrix Connect Oracle, Mirrored	
Prompt	Description/Actions
Do you want the EDM to wait for tracks to resynchronize before displaying a message that the backup has finished? Wait for Mirror Sync: Note: Omitted for Leave mirrors down (and for “no-mirror”).	Select from the dropdown list either: Wait for resync Don't wait for resync
Do you want to run a script on the client after mirrors are split, and before the backup begins? ☐ Run Script After Mirrors Are Split Note: Omitted for “no-mirror”.	Choose whether to run a post-mirror-split script.
Specify the script to run after mirrors are split. Script: <u>/usr/epoch/</u> Browse: Username: <u>root</u> Note: Omitted if not running script after mirror split (and for “no-mirror”).	Pathname and username for Symmetrix post-mirror-split script: • Enter the <i>full pathname</i> for the script. You can either browse existing files or enter the name. Note: The default is only a directory, not a file. • Enter the Oracle username (UNIX account) to run it.
Select the media type you want for this backup. Choose one of: _____ Note: Omitted if there is one media type available.	Choose the media type of library unit to be used for backup from the available types in the dropdown list.

Table B-1

Configuring Symmetrix Connect Oracle, Mirrored

Prompt	Description/Actions
Select the number of trails and streams for this configuration. Trails: Streams per trail:	- Enter the number of trails to use concurrently for the EDM Symmetrix Connect portion of this backup. Note: (Depending on how many network streams you selected and how many streams per trail you specify below, the first EDM Symmetrix Connect trails might or might not be also used for network backups.) Note: For optimal backup performance, the number of trails should equal the number of available drives of the media type you are selecting. It cannot exceed the server parameter for maximum concurrent trails for this media type. (The default is the value for the <code>server</code> parameter for maximum concurrent trails for this media type.) - Enter the Number of streams per trail. This is the maximum number of concurrent data streams per trail for database backup using EDM Symmetrix Connect. Note: For optimal <i>backup</i> performance, it could be more than 1, depending on disk layout and the use of a tuning file. The maximum is the value for the Media tab parameter for Limit Maximum Concurrent Backups for this media trail. For optimal <i>restore</i> performance, it should equal 1.

Table B-1 Configuring Symmetrix Connect Oracle, Mirrored	
Prompt	Description/Actions
Provide base names for the work item and trail Work Item Base Name: Trail Base Name:	Enter the base name to be used in naming the trail. - For Symmetrix Connect's interface to Oracle, the default work item base name is <code>dc_oracle_mirroring_clientname</code> - The default trail base name is <code>dc</code>. At the end of the configuration, by default, the templates and trails for Symmetrix Connect will be named using the base trail name these formats: <code>template_basetrailname_trailnumber</code> <code>trail_basetrailname_trailnumber</code>
Do you want to run a configuration check? ◆ Yes ◆ No	You are asked if you want to run a configuration check without running a backup. WARNING: Be aware that this validation process will split and re-establish the mirrors as configured. This could alter the data on the mirrors if they are already split. The check could take hours to complete. Note: Skip the check if you have multi-BCVs and you have yet to set up the .sets file (see “Creating a .sets File for Multi-BCVs” on page 2-19). Choose No to skip this check. You are warned to run an eb_dc_backup -check work_item_base_name before running a backup. Choose Yes to run the configuration check.

Table B-1

Configuring Symmetrix Connect Oracle, Mirrored

Prompt	Description/Actions
Do you want verbose output for the configuration check? ◆ Yes ◆ No Note: Omitted if no configuration check.	If you choose Yes to run the configuration check, Click Yes or No for verbose output.
Select “Run” to begin the configuration check. Note: Omitted if no configuration check.	If you decided to run the configuration check, you are prompted to click Run to begin the check. Once you click Run, you have the option to cancel while the check occurs. When the check is finished, you have the option to save the output.
Do you want to enable media duplication for this media set? ◆ Yes ◆ No	Choose whether or not to use media duplication. (No is <u>DEFAULT</u>) Note: For this scenario, take the default.
How often do you want to backups to rotate? Rotate Every <u>14</u> days	This value determines the media rotation period. This is the duration after which, by default, EDM starts writing the backups to a new volume rather than appending the backups to the volume containing the previous backups. Note: For Symmetrix Connect, EDM does not perform automatic scheduling of full vs. incremental backups.
How long do you want to retain your backups? ◆ Forever ◆ For <u>1</u> years	Choose the period of time to retain your backups, their catalogs, and duplicates. <u>DEFAULT</u> is 1 year

Table B-1

Configuring Symmetrix Connect Oracle, Mirrored

Prompt	Description/Actions
Do you want EDM to initiate the backup? ◆ Yes <u>18:00</u> ◆ No	Choose Yes and set the time if you want the EDM to regularly run (from cron on a daily basis) the backup job (ebbackup) for this backup configuration. Choose No if you are using another scheduling tool or if you want to customize manually the crontab file that controls this function.
Do you want to automatically retry failed backups? ◆ Yes ◆ No Note: Omitted if EDM not initiating backup regularly from cron .	Choose No. You would choose Yes if you want the EDM to reschedule and reprocess any backups that fail in the current backup run. The retries occur after all normally scheduled backups have completed.
Do you want backup to append to the current volume or use a new one? ◆ Append to Current Volume ◆ Use a New Volume	Choose if you want to append new data to the current volume until the first day of a new rotation or if you want to use a new volume for each new backup (which enables you to eject the backup media for the previous backup).
What do you want to name the new configuration? Name:	Enter the name to use to identify this configuration, if you later want to modify, duplicate, or delete it using the Backup Configuration Wizard. The default name is the work item base name that you provided in the panel “Provide base names for the work item and trail.”
Summary of your configuration: Configuration options:	The last panel summarizes your configuration. Click Save to complete the configuration.

C Sample Configuration: UNIX Filesystem Interface

This appendix contains the Backup Configuration wizard's prompts for the following scenario:

Configuring Symmetrix Connect UNIX Filesystem Backups

See these other appendices for other scenarios:

- Appendix A, Configuring Symmetrix Connect Oracle, No Mirror
- Appendix B, Configuring Symmetrix Connect Oracle, Mirrored

Table C-1 Configuring Symmetrix Connect UNIX Filesystem Backups	
Prompt	Description/Actions
Select the clients that you want to configure for backup.	Choose the network hostname of the client hosting the filesystems to be configured for Symmetrix Connect backup with the UNIX filesystem interface.
What type of configuration do you want? ◆ New Configuration ◆ Modify Configuration ◆ Duplicate Configuration ◆ Delete Configuration	Choose: New Configuration (It's the <u>DEFAULT</u>) Delete Configuration erases an entry in the configuration file having a specified configuration name.
Select the object(s) to configure:	Click on the Select Unconfigured button to select all the filesystems.
Select the application to use for this configuration. —————Backup Application————— ◆ Symmetrix Connect backup ◆ Network or Symmetrix path backup	Choose: Symmetrix Connect backup to choose the Symmetrix Connect interface to UNIX filesystems.
Select the mirror configuration you want to use. —————Mirror Configuration————— ◆ Local BCV ◆ SRDF ◆ No mirror ◆ Remote BCV	Select which one of the following Symmetrix Configurations has already been set up: • Local BCV (TimeFinder) mirroring <u>DEFAULT</u> • SRDF mirroring. • No mirroring. • Remote BCV (TimeFinder mirroring of SRDF mirror).

Table C-1

Configuring Symmetrix Connect UNIX Filesystem Backups

Prompt	Description/Actions
Are multiple BCVs present for any standard device involved with the backup? BCV Choices ◆ Multiple BCVs present ◆ One BCV present for every device Note: Omitted for “no-mirror” and “SRDF”.	Specify if there are multiple-BCVs for any standard device involved in the backup; if yes, select the first choice, if no, select the second: - Multiple BCVs present - One BCV present for every device <u>DEFAULT</u> See “Creating a .sets File for Multi-BCVs” on page 2-19.
Select the mirror set policy you wish to use? Mirror Set ◆ Round Robin rotation of BCV sets ◆ Use One BCV set Note: Only received if you set “Multiple BCVs present”. Note: Omitted for “no-mirror”.	If multiple BCVs are present, select to either rotate through sets of multiple BCVs or to just use one specified set of BCVs: - Round Robin rotation of BCV sets - Use One BCV set <u>DEFAULT</u>
Name of the BCV mirror set you want to use? BCV Set Name <u>set1</u> Note: Only received if you set “Use One BCV set”. Note: Omitted for “no-mirror”.	If you are specifying a particular set, enter the BCV Set Name now. For example, enter “set1”.

Configuring Symmetrix Connect UNIX Filesystem Backups

EDM Symmetrix Connect — Interface to Oracle, UNIX Filesystems

Table C-1

Configuring Symmetrix Connect UNIX Filesystem Backups

Prompt	Description/Actions
Do you want to run a script on the client after mirrors are split, and before the backup begins? ☐ Run Script After Mirrors Are Split Note: Omitted for “no-mirror”.	Choose whether to run a post-mirror-split script.
Specify the script to run after mirrors are split. Script: <u>/usr/epoch/</u> Browse: Username: <u>root</u> Note: Omitted if not running script after mirror split (and for “no-mirror”).	Pathname and username for Symmetrix post-mirror-split script: - Enter the <i>full pathname</i> for the script. You can either browse existing files or enter the name. Note: The default is only a directory, not a file. - Enter the username (UNIX account) to run it. The default is root
Select the media type you want for this backup. Choose one of: _____ Note: Omitted if there is one media type available.	Choose the media type of library unit to be used for backup from the available types in the dropdown list.

Configuring Symmetrix Connect UNIX Filesystem Backups

EDM Symmetrix Connect — Interface to Oracle, UNIX Filesystems

Table C-1

Configuring Symmetrix Connect UNIX Filesystem Backups

Prompt	Description/Actions
Provide base names for the work item and trail Work Item Base Name: Trail Base Name:	Enter the base name to be used in naming the trail. - For Symmetrix Connect's interface to UNIX Filesystems, the default work item base name is <i>dc_filesystem_mirroring_clientname</i> - The default trail base name is <i>dc</i>. At the end of the configuration, by default, the templates and trails for Symmetrix Connect will be named using the base trail name these formats: <i>template_basetrailname_trailnumber</i> <i>trail_basetrailname_trailnumber</i>
Do you want to run a configuration check? ◆ Yes ◆ No	You are asked if you want to run a configuration check without running a backup. WARNING: Be aware that this validation process will split and re-establish the mirrors as configured. This could alter the data on the mirrors if they are already split. The check could take hours to complete. Note: Skip the check if you have multi-BCVs and you have yet to set up the .sets file (see "Creating a .sets File for Multi-BCVs" on page 2-19). Choose No to skip this check. You are warned to run an eb_dc_backup -check work_item_base_name before running a backup. Choose Yes to run the configuration check.

Table C-1 Configuring Symmetrix Connect UNIX Filesystem Backups	
Prompt	Description/Actions
Do you want verbose output for the configuration check? ◆ Yes ◆ No	If you choose Yes to run the configuration check, Click Yes or No for verbose output.
Select “Run” to begin the configuration check.	If you decided to run the configuration check, you are prompted to click Run to begin the check. Once you click Run, you have the option to cancel while the check occurs. When the check is finished, you have the option to save the output.
Do you want to enable media duplication for this media set? ◆ Yes ◆ No	Choose whether or not to use media duplication. (No is <u>DEFAULT</u>) Note: For this scenario, take the default.
How often do you want to backups to rotate? Rotate Every <u>14</u> days	This value determines the media rotation period. This is the duration after which, by default, EDM starts writing the backups to a new volume rather than appending the backups to the volume containing the previous backups. Note: For Symmetrix Connect, EDM does not perform automatic scheduling of full vs. incremental backups.
How long do you want to retain your backups? ◆ Forever ◆ For <u>1</u> years	Choose the period of time to retain your backups, their catalogs, and duplicates. <u>DEFAULT</u> is 1 year

Table C-1

Configuring Symmetrix Connect UNIX Filesystem Backups

Prompt	Description/Actions
Do you want EDM to initiate the backup? ◆ Yes <u>18:00</u> ◆ No	Choose Yes and set the time if you want the EDM to regularly run (from cron on a daily basis) the backup job (ebbackup) for this backup configuration. Choose No if you are using another scheduling tool or if you want to customize manually the crontab file that controls this function.
Do you want to automatically retry failed backups? ◆ Yes ◆ No Note: Omitted if EDM not initiating backup regularly from cron .	Choose No. You would choose Yes if you want the EDM to reschedule and reprocess any backups that fail in the current backup run. The retries occur after all normally scheduled backups have completed.
Do you want backup to append to the current volume or use a new one? ◆ Append to Current Volume ◆ Use a New Volume	Choose if you want to append new data to the current volume until the first day of a new rotation or if you want to use a new volume for each new backup (which enables you to eject the backup media for the previous backup).
What do you want to name the new configuration? Name:	Enter the name to use to identify this configuration, if you later want to modify, duplicate, or delete it using the Backup Configuration Wizard. The default name is the work item base name that you provided in the panel "Provide base names for the work item and trail."
Summary of your configuration: Configuration options:	The last panel summarizes your configuration. Click Save to complete the configuration.

D Backup Configuration and Processing Concepts

Read this chapter to better understand the backup configuration and concepts of EDM Symmetrix Connect. The backup processing phases are also described.

This chapter covers the following topics:

- Backup Configuration Results
- Crontab Entry to Schedule Backups
- Backup Processing
- The .dcrpt and .ddtab Files

Backup Configuration Results

As a result of the Backup Configuration Wizard (or **eb_config** command) session for Symmetrix Connect, information specific to Symmetrix Connect is written to *work_item_base_name.cfg*; there is one such file for each work item base name.

In addition, backup configuration information is written to EDM's backup configuration (*eb.cfg*) file and can be viewed in the EDM GUI in the Backup Configuration window:

- Symmetrix Connect work items (see page D-2)
- A Symmetrix Connect kicker work item (see page D-3)
- Symmetrix Connect work groups (see page D-4)
- A template, trailset, and trail (see page D-5)
- A network work item (see page D-5)

Also, a command to automatically schedule backups is written to the EDM's crontab. See "Crontab Entry to Schedule Backups" on page D-6.

Symmetrix Connect Work Items

Each time you run the EDM Backup Configuration Wizard (or **eb_config** command) for Symmetrix Connect, it creates as many EDM Symmetrix Connect work items (type:dc_worker) as the number of trails multiplied by the number of Symmetrix Connect streams per trail. (This equals the maximum concurrent streams from the database (-**n concurrent_work_items**). The names of these work items are in the format:

workitembasename_streamnumber

Figure D-1

EDM Symmetrix Connect Work Items for Streams 2,6, 10

```
"dc_oracle_local_bcv_summer_2"
"dc_oracle_local_bcv_summer_6"
"dc_oracle_local_bcv_summer_10"
```

Symmetrix Connect Kicker Work Item

For Symmetrix Connect, the EDM Backup Configuration Wizard (or **eb_config** command) creates one kicker work item for the client.

The names of these work items are in the format:

workitembasename_k

This kicker type (*workitembasename_k*) is a Symmetrix Connect kicker used to schedule backups for the application interface to Oracle and the application interface to UNIX filesystems.

The Symmetrix Connect kicker work item is put in the first EDM Symmetrix Connect work group.

This is the work item you specify to start the Symmetrix Connect backup from the Backup Activity wizard or from **ebbbackup -i** in crontab. See “Crontab Entry to Schedule Backups” on page D-6 and “Backing Up with the Backup Activity Wizard” on page 3-3.

Work Groups for EDM Symmetrix Connect

The EDM Backup Configuration Wizard (or **eb_config** command) creates as many EDM Symmetrix Connect work groups as the maximum concurrent trails (**-t concurrent_trails**). The names of the EDM Symmetrix Connect work groups are in the format:

wg_workitembasename_trailnumber

The EDM Backup Configuration Wizard (or **eb_config** command) distributes the EDM Symmetrix Connect work items among the EDM Symmetrix Connect work groups evenly, assigning the stream numbers as shown in Figure D-2. (Only one work group shown).

Figure D-2

EDM Symmetrix Connect Work Group for Streams 2,6, 10

```
"wg_dc_oracle_local_bcv_summer_2"  
"dc_oracle_local_bcv_summer_2"  
"dc_oracle_local_bcv_summer_6"  
"dc_oracle_local_bcv_summer_10"
```

Templates for EDM Symmetrix Connect

The EDM Symmetrix Connect work groups are assigned to schedule templates created to correspond to each trail. The names of the EDM Symmetrix Connect templates are in the format:

template_basetrailname_trailnumber

The names of the EDM Symmetrix Connect trails are in the format:

trail_basetrailname_trailnumber

Figure D-3

Template for EDM Symmetrix Work Group for Trail 2

backup template: "template_dc_2"

primary trailset: "trail_dc_2"

work group list:

"wg_dc_oracle_local_bcv_summer_2";

Network Work Item With Interface to Oracle or UNIX Filesystem Interface

Network work items are used for backing up control files and archived redo logs.

A network work item is set up automatically when you run the EDM Backup Configuration Wizard (or **eb_config** command) for a Symmetrix Connect configuration with the Oracle interface or the UNIX filesystem interface. The default name of this work item is in the format:

work_item_base_name_0

Crontab Entry to Schedule Backups

The EDM Backup Configuration Wizard and the Backup Configuration window maintain crontab entries for backups they configure. See step 30 on page 2-12 (or step 21 on page 2-18).

The command crontab uses for starting the backup is **ebbackup**. A basic backup job is run by **ebbackup** *template*. A specific Symmetrix Connect backup for the Oracle interface or the UNIX filesystem interface can be run by **ebbackup -i** *work_item_base_name_k* *template*

See “Symmetrix Connect Kicker Work Item” on page D-3.

The following is a sample of the area reserved in the crontab for EDM backups.

```
## Begin EDM GUI maintained backups #####
## This section of the root crontab is maintained using the EDM GUI. #
## Do not manually edit these entries - do not modify this table. ####
0 18 * * * /usr/epoch/EB/bin/ebbackup template_dc_1 > /dev/null 2>&1 # 'template_dc_1' 0 0
## End EDM GUI maintained backups #####
```

Backup Processing

The kicker work item has a backup initialization command that runs the script **eb_dc_backup** for the *work item base name*.

Note: The work item base name corresponds to a file named *work_item_base_name.cfg*, which provides a formatted list of the configuration parameters specified during the EDM Symmetrix Connect Configuration for that work item.

A Symmetrix Connect backup for the Oracle interface (run by the **eb_dc_backup** command) progresses through the following five phases:

1. Discovery
2. Preparation
3. Execution
4. Cleanup
5. Backup Completion

Phase 1 - Discovery

During the discovery phase, the Symmetrix Connect software learns what needs to be backed up. Symmetrix Connect with the interface to Oracle connects to the Oracle database, obtains a list of datafiles and the archived log directory, and creates a coherent copy of the control file. Symmetrix Connect then determines which Symmetrix devices are in use for the database, for SRDF, and for TimeFinder. Then Symmetrix Connect verifies that all requested data files have Symmetrix disk ids and EDM-side pathnames, confirming that the files reside on Symmetrix devices that are either visible to the EDM or mirrored to Symmetrix devices that are visible to the EDM.

Phase 2 - Preparation

During the preparation phase, the interface to Oracle shuts down the database or puts tablespaces into online backup mode. Then, for mirrored configurations only, Symmetrix Connect splits the Symmetrix mirrors and the Oracle interface brings the database back online (for offline backups), or takes the affected tablespaces out of online backup mode (for online backups).

Phase 3 - Execution

During the execution phase, the EDM schedules and runs the set of Symmetrix Connect work items, copying the data files from the Symmetrix disks to tape. Then the EDM schedules and runs the network work item for backup, which copies the archived redo logs and the backup control file over the network. The archived redo logs are mandatory for online backups, optional for offline backups. Each work item creates a separate saveset and catalog, listing all the files and partitions that were backed up by that work item.

Phase 4 - Cleanup

During the cleanup phase: For no-mirror configurations only, Symmetrix Connect remounts any “raw filesystem” filesystems and the interface to Oracle starts up the database or takes it out of online backup mode. For mirrored configurations only, the software resynchronizes Symmetrix mirroring, if so configured.

Phase 5 - Backup Completion

During the backup completion phase, Symmetrix Connect verifies the cataloging of the backup files, creates the final *work_item_base_name.dcrpt* file with the backup status and saveset ID information for individual work items, archives the final .dcrpt and .ddtab files to disk and backs them up to tape, and provides notification of the status of the backup. (See “The .dcrpt and .ddtab Files” on page D-9.)

The .dcrpt and .ddtab Files

Each time **eb_dc_backup** runs, it creates a new file named *work_item_base_name.dcrpt*, with information about the date's backup. Each line in the file describes one object that was backed up and includes: the name of the Symmetrix Connect work item that did the back up; the names of the instance, the tablespace, and the pathname; and the type of file (datafile, control, or archivelog directory). The most recent *work_item_base_name.dcrpt* file is located in the */usr/epoch/EB/server_name* directory. The final version of the .dcrpt file is backed up by the kicker work item.

Similarly, each time **eb_dc_backup** runs, it creates a new *work_item_base_name.ddtab* file. It contains information about that date's backup, such as the database, filesystem, LVM, and Symmetrix configuration for data files backed up.

The thirty latest versions of .dcrpt and .ddtab files are stored on disk are stored in the */usr/epoch/EB/Symmetrix_Connect/Archives/work_item_base_name* directory, and are named as *backuptimestamp.dcrpt* where "backuptimestamp" has a YYYY.MM.DD.hh.mm.ss format.

Symmetrix Connect restores rely on EDM-archived .dcrpt and ddtab files to link Symmetrix Connect information to EDM backup catalog information.

Before you can issue a **eb_dc_restore** command on the kicker work item whose related dcrpt and ddtab files are more than thirty versions old, you need to restore the files from tape. See "Restoring Old .dcrpt and .ddtab Files" on page 4-2.

E Setting Up a Tuning File

By default, Symmetrix Connect balances the backup data load across the number of work items that have been configured for a given work item base name. But, if you are familiar with the physical disk locations of the Oracle datafiles, you might achieve even better performance by creating a tuning file.

Tuning File Overview

A tuning file is a means by which your data load can be more evenly distributed across physical hardware. The use of a tuning file might be able to improve the backup rate for a particular configuration.

The tuning file specifies the order of backup (which file follows which file) and the stream to which each file should be backed up. Data files per work item in the *work_item_base_name.ddtab* that are not specified in the configuration tuning file will still be backed up, but the order and stream are not predictable. If you create a tuning file, the Symmetrix Connect software will consult it, and generate file lists and partition lists for the work items that implement your specified order.

Creating a Tuning File

Locate the tuning file at
/usr/epoch/EB/config/work_item_base_name.tune.

Create a line oriented text file. Each line has three columns:


- Column 1: A pattern that matches one or more of the datafiles to backup. Get this value from a previously executed *work_item_base_name.dcrpt* file. See “The .dcrpt and .ddtab Files” on page D-9.
- Column 2: A numeric value that specifies the work item number to which this file should be sent
- Column 3: A numeric value that specifies the precedence of this file in the work item. Lower numbers have the high precedence.

Smaller precedence numbers get backed up before larger numbers within the same work item.

Refer to the following abbreviated example of a tuning file in Figure E-1.

Figure E-1

work_item_base_name.tune File (Example)



In this example, the work groups 1 through 4, each composed of multiple work items, may be backed up in parallel or not, depending on eb.cfg file parameters and available drives with media.

A file that does not match any of the patterns in the tuning file will be backed up to by a stream that has the smallest amount of allocated data, as though it had a precedence value of 1 (default).

Note: The only special characters (wild cards) that are supported in a tuning file are “*” and “?” (and they must be placed at the end of the file pattern).

Note: Warning messages associated with tuning files are described in the *EDM Server Error Messages* manual.

Raw Filesystem Backups

When you create a tuning file for a raw filesystem backup configuration (filesystems backed up as devices), the file entry in the tuning file should be the device on which the filesystem (that is to be backed up) resides, not the files that constitute the filesystem. The device that is to be backed up would be the *client_file* field in the *work_item_base_name* .ddtab corresponding to the raw filesystem backup configuration.

For example, the following files are part of a filesystem that resides on the device /dev/rivers/rcoclv1.

```
/oraclefs2/data1_concat.dbf
/oraclefs2/data_FS1.dbf
/oraclefs2/data2_RAW.dbf
```

When creating a tuning file for a raw filesystem backup, /dev/rivers/rcoclv1 should be specified just once in it, not the above files. So to tune the backup, the entry in the tuning file would be:

```
# This entry says to put the device below in Work
Item 2 with a priority of 1
```

```
/dev/rivers/rcoclv1 2 1
```

File Backups

When you create a tuning file for backing up filesystem files as files, the file entries in the tuning file should be the file pathnames as Oracle knows it. Do not resolve the pathnames to the underlying device pathname.

For example, from SVRMGR you enter “select * from v\$datafile;” and a file is listed as:

`/oracle/app/oracle/product/8.0.5/oradata6/GT4gb_1/GT4gb.data1`

Enter the following in the tuning file:

`/oracle/app/oracle/product/8.0.5/oradata6/*`

F Symmetrix Connect Prerequisites

This appendix lists the specific prerequisites that should be satisfied for the EDM and the Oracle host. These prerequisites have to be satisfied prior to this book's procedures.

This appendix describes the following topics:

- EDM Prerequisites
- Backup Client Prerequisites

EDM Prerequisites

Prerequisites for the EDM prior to this book's procedures are:

- ❑ Solaris 2.6 operating system already installed or upgraded.
- ❑ Solutions Enabler package installed — the same versions as on the client; (includes SymmCLI software and SymmAPI library).
- ❑ Appropriate Symmetrix logical disks are exported to EDM:
 - For non-mirrored configurations, the devices that either contain database's data files or symbolic links to them.
 - For SRDF, the target (R2) devices.
 - For TimeFinder, BCV mirror devices.
 - For Remote BCV, the remote BCV devices.
- ❑ One SYMAPI gatekeeper device for Symmetrix Connect in addition to the gatekeepers for other SYMAPI applications (one gatekeeper per application).

Note: Symmetrix Connect uses the gatekeeper on the EDM to discover devices and to perform mirror operations.

- ❑ The **symcfg discover** command has been run on the EDM:
khayyam# **/usr/symcli/bin/symcfg discover**

Note: Run the **symcfg discover** command whenever changes are made to the Symmetrix, such as when new devices are exported to the database system. Run **symcfg help** for more information.

- ❑ EDM 4.6.0 software installed.
- ❑ All EDM 4.6.0 software patches installed.
- ❑ Configured with **edminstall** (which includes **lmconfig** and **eb_server_config**).

Backup Client Prerequisites

Prerequisites for the Oracle host (the backup client host) prior to this book's procedures are:

- ❑ Supported operating system installed (see *EMC Data Manager Release Notes*).
- ❑ Solutions Enabler package installed — the same versions as on the EDM; (includes SymmCLI software and SymmAPI library).
- ❑ If the Oracle host is an HP system using the older HP-PB (NIO) fast-wide differential SCSI board and it is to be used in a TimeFinder configuration (either Local BCV or Remote BCV), pseudo-devices were created.
- ❑ Appropriate Symmetrix logical disks exported to client host:
 - For non-mirrored configurations, the devices that either contain database's data files or symbolic links to them.
 - For SRDF, the target (R2) devices.
 - For TimeFinder, BCV mirror devices.
 - For Remote BCV, the remote BCV devices.
- ❑ One SYMAPI gatekeeper device for Symmetrix Connect in addition to the gatekeepers for other SYMAPI applications (one gatekeeper per application).

Note: Symmetrix Connect uses the gatekeeper on the client to discover devices.

- ❑ The **symcfg discover** command has been run on the client:

```
khayyam# /usr/symcli/bin/symcfg discover
```

Note: Run the **symcfg discover** command whenever new STD devices are exported to the database system. Run **symcfg help** for more information.

- ❑ The Oracle product installed, including one or more Oracle databases created (and running in archivelog mode, with auto-archiving started).

Glossary

alternate mirror	A Symmetrix device mirror configuration type, including no-mirror, other than the one that was backed up.
autoconfiguration	A feature of EDM backup software where the software automatically creates work items to backup the known filesystems on a client. Often performed in conjunction with client installation.
backint	SAP AG defined interface to utilities sapdba , brbackup , brarchive , and brrestore and to external backup and media management products such as EMC Data Manager.
backup catalog	Identifies a backup at the file level by recording the names and attributes of each file on the client system at the time of the backup. Backup catalogs also keep track of the location of backed up data for each file that was selected for backup.
backup template	Specifies all the information about how to perform a backup. It includes the workgroup(s) to backup, the rotation period and the backup shift lengths, the trailsets on which to store the backup data, and the backup schedule that dictates a backup level for each day.

backup trailset	Defines the media trails to which the backup data is sent and the type of media to use. It also defines how long to save the backup data and its associated elements.
BCV device	Business Continuance Volume; a special kind of logical disk, designated for use as a TimeFinder or remote BCV target disk.
BCV pair	Consists of a Symmetrix standard device and a BCV device attached together.
brarchive	SAP utility used to back up the Oracle archived redo log files. The utility can be accessed from the sapdba menu or be used from the command-line.
brbackup	SAP utility used to back up the Oracle database files, including the control files online redo log files, configuration files, and catalog files. The utility can be accessed from the sapdba menu or be used from the command-line.
brrestore	SAP utility used to restore databases, including all files (even archived redo logs). The utility can be accessed from the sapdba menu or be used from the command-line.
datafile	A database's datafiles contain all the database data. A datafile is the smallest identifiable backup object. Restoring it may require restoring a larger object containing it. The data of logical database structures, such as tables and indexes, is physically stored in the datafiles allocated for a database. A datafile, located in a raw device or filesystem, can be associated with only one database, and once created, a datafile cannot change in size. In Oracle one or more datafiles form a logical unit of database storage called a tablespace.

dcrpt report	Lists details of an actual backup down to the data file, separated by <i>work_item_base_name</i> and timestamp.
device name	A name that the operating system associates with a physical object, where a physical object is either a partition or a disk drive. For example, c0t0d0 is a device name of a disk and c0t0d0s2 is a device name of a partition.
ddtab	The discovery data table, constructed during the discovery phase, contains the data required to reconstruct the client environment on the EDM server unit in order to backup the target Symmetrix. The details file, in ddtab format, records the configuration parameters specified by the user during the configuration process. The ddtab files are used for communicating between storage management processing layers.
disk	A direct access storage medium with magnetically-recorded data.
directory	A logical subdivision of a computer's disk space, used to organize a group of files.
domain	A domain is a pool of storage with multiple devices, in Compaq Tru64 AdvFS.
EDM	EMC Data Manager. A computer system composed of a central processing unit, magnetic disks, high speed network interfaces, robotic library or libraries, and backup software. EDM is dedicated to managing and storing large amounts of data. Its primary use is for backup and recovery operations.
EMC TimeFinder	An EMC software product that allows customers to use special devices that contain copies of Symmetrix standard devices while the standard devices are online for regular Symmetrix operation.

fileset	A fileset is the Compaq Tru64 AdvFS equivalent to UNIX filesystems, except that it can reside on the multiple devices of the domain.
filesystem	A collection of files and directories on a disk or tape in standard UNIX file system format. Either a) A file tree on a specific disk partition or other mount point, or b) The entire set of all files.
gatekeeper	Commands executed by SymmAPI are transferred to the Symmetrix via a “gatekeeper” device. The gatekeeper device is one of the Symmetrix volumes accessible by the database host or EDM using SymmAPI. Symmetrix Connect uses the gatekeeper on the EDM to discover devices and to perform mirror operations. Symmetrix Connect uses the gatekeeper on the client to discover devices.
interface	In this book, specifically used to characterize Symmetrix Connect modules that facilitate connection and operation between the core Symmetrix Connect software and the database or filesystem.
logical volume	A virtual disk device, presented by a logical volume manager, that appears to applications, databases, and file systems like a physical disk partition, but does not have the physical limitation of a physical disk partition. A volume created using a Logical Volume Manager (LVM) like VERITAS.
logical volume group	A concept within logical volume managers. A collection of logical disks that are grouped together (or share a common configuration). Several logical volumes may be presented from one group of logical disks; no logical volume can span different logical volume groups.

logical volume manager	A software product that has the capabilities of presenting a part of one or more logical disk devices as a logical volume. The logical volume manager maps I/O requests to the logical volume to I/O requests to the underlying logical disk drives transparently to the application.
magnetic disk	A non-volatile computer storage device which uses one or more magnetic-coated disk platters that can be recorded on over and over again.
mount	An operation where a logical volume or raw partition containing a filesystem is associated with a mount point, enabling access to the files within a filesystem.
network work item	For Symmetrix Connect, a “listener” database work item that backs up all client database control files and archived log directory over the network to the EDM server unit. (For network or Symmetrix Path backups, listener work items handle all database backup data.)
Oracle sid	Database instances are identified by a sid, system identification, number which does not have to be the same as the database name
partition (raw partition)	A contiguous portion of a logical drive, used to store related information, such as a filesystem.
production device	In a non-mirrored backup, this is the Symmetrix disk that was backed up or the source volume. When restoring to no-mirrors, you are restoring to the production devices.
raw device	Storage devices that are not buffered through an operating system cache.

redo log	A sequential log of transactions that are to be reapplied to the database, if they did not get written to disk. The log always consists of at least two files: one is optionally being spooled to disk while the other is being written. When the one currently being written fills up, the next one is reused.
remote BCV	Refers to the BCV mirror copy of the standard device on an SRDF target Symmetrix.
resynchronize mirrors	Returning Symmetrix devices to a mirrored state. The differences between the source and target drive are removed, and then mirroring is enabled so that the one drive continues to be a copy of the other. Re-establishing refers to standard device overwriting the BCV. Restoring refers to the BCV overwriting the standard.
root disk	The disk drive where the operating system resides.
rotation period	Specifies the number of backups after which a new set of media is started for the next backup.
sapdba	SAP menu-based utility started from the command-line used to access the brbackup , brarchive , and brrestore utilities.
saveset records	Contains information about an entire backup: the start time and the trails that the backup program uses to write the backup data.
server	A computer system that runs the database management system or other application using the data you are backing up with Symmetrix Connect.

source	In an SRDF Symmetrix mirrored disk environment, the logical disk that can be read from and written to by a host. Updates to the source Symmetrix are automatically made to the corresponding target.
splitting a mirror	The act of (temporarily) disabling the mirroring that was active between two Symmetrix devices. At the point when the mirror is split, the two logical disks are exact copies of each other. Subsequently, these two copies will drift apart as updates are made to one or the other independently.
SRDF	Symmetrix Remote Data Facility is a capability of Symmetrix which allows two or more Symmetrix to be connected to each other, allowing logical disks from one Symmetrix to be mirrored at high speed onto logical disks in another Symmetrix.
standard device	A Symmetrix device, designated for use by a client. Standard devices may be mirrored by Symmetrix.
STG	Symmetrix Transport Group; a collection of Symmetrix devices, formatted and configured for use together as transport media for EDM Symmetrix Path. Each STG is shared among connected hosts, such as the EDM and the database server, thus constituting a transport path enabling data transfer.
STP	Symmetrix Transport Protocol; a software implementation of rules defining communication over the Symmetrix. Allows communication between hosts, using Symmetrix as the transport media.
Symmetrix device	A virtual disk drive within a Symmetrix unit as a single SCSI target and LUN disk. Same as a Symmetrix “logical volume” or hypervolume.

SymmAPI	Symmetrix Application Program Interface that consists of a library of commands which provide a way to monitor and control Symmetrix operations. For example, a local BCV backup uses the SymmAPI to split the mirrors.
tablespaces	An Oracle database is divided into logical storage units. A tablespace is used to group related logical structures together. For example, tablespaces commonly group all of an application's objects to simplify certain administrative operations.
target	In a mirrored disk SRDF Symmetrix environment, a logical disk that is a copy of the source, which cannot be written to directly by a host.
templates	A set of backup scheduling policies that apply to all of the work groups that belong to the template. A schedule template has a primary media set and may also have an alternate media set.
trail	A media use policy, including media type and label class. Simultaneous backups that use the same trail will share the same tape device.
trailset	A set of media use policies. A trail set specifies the trails to use for various types of backup (level 0, level 9, baseline, etc.) and the rotation time for each.
unmount	An operation which disables access to a filesystem, and splits the association with the mount point. Most operating systems do not allow a filesystem to be unmounted when there are processes that are using files or directories within the filesystem.
work group	A group of work items. To function, a work group must belong to one backup schedule template.

- work item** The smallest unit of a backup job, identifying what is to be backed up, the client on which the backup is to run, and certain backup configuration parameters. Work items belong to work groups. A stream of data reading from disk and writing to tape.
- work item base name** The name of a configured EDM Symmetrix Connect backup. All work items and work groups associated with a backup have names that contain the same *work_item_base_name*. The backup must have been configured by the **eb_config** command or the EDM GUI's Backup Configuration Symmetrix Connect wizard.

Index

A

- AdvFS filesystems
 - backup support 1-15
 - restore support 4-20
 - restoring 4-20
- alternate client restore 1-24
- alternate destination restore 1-24
- applications
 - prompt for 2-14
 - types for Symmetrix Connect 1-4
- archived redo logs
 - backup 1-12, D-8
 - prompt 2-4

B

- Backint 1-4
- backup completion phase D-8
- backup configuration file D-2
- backup failure 2-22, 3-6
- backup processing phases D-7
- BCVs, multiple 1-7
 - eb_dc_backup -mirror_set 3-5
 - error 2-22
 - restore prompt 4-15
 - .sets file 2-19
 - wizard prompts 2-6
- brbackup 1-4

C

- CD containing documentation viii
- cleanup phase D-8
- client prerequisites F-3
- client, configuring
 - Oracle 2-3
 - UNIX filesystem 2-14
- closed backups 1-8
- configuration check
 - failure 2-23
 - prompt 2-11, 2-17
- control file
 - backup 1-12, D-8

D

- database objects 1-12
- database server requirements F-3
- datafiles
 - backing up as device 1-14
 - backing up as files 1-12
 - restoring as device 4-5
 - restoring as files 4-7
 - tuning file E-5
- DBA 2-6
- dcrpt and ddtab
 - description D-9
 - restoring 4-2

defaults

- configuration name
 - Oracle interface 2-12
 - UNIX filesystem 2-18

pre-backup

- Oracle offline 2-24

pre-backup actions

- Oracle offline 1-12
- Oracle online 1-9
- UNIX filesystems 1-28

trail base name 2-10, 2-17

work item base name

- Oracle interface 2-10
- UNIX filesystem 2-17

device backups

filesystems

- concept 1-15
- restoring 4-5
- tuning file E-4

discovery

- backup discovery
 - concepts 1-13
 - phase D-7
- symcfg discover F-2

documentation viii

domain, AdvFS 1-15, 4-20

E

eb_dc_backup 3-5

- check 2-17

eb_dc_restore

- listfiles 4-17
- mirror set 4-19
- restore 4-18

eb_restore 4-3

eb.cfg file D-2

edm 2-3

edmrestore 4-12

execution phase D-8

export

- command in scripts 2-26
- requirements for F-2
- Symmetrix concept 1-3

F

failures

- backups 3-6
- configuration check 2-23
- restore 4-16
- svrmgrl 2-26

file-level granularity 1-17

filesets, AdvFS 1-15, 4-20

filesystems

- AdvFS 1-15, 1-29
 - restore limitations 4-20
- AIX 1-16, 1-29
- Compaq Tru64 1-15, 1-29
 - restore limitations 4-20
- HP-UX 1-16, 1-29
 - restore limitations 1-23, 4-20
- Sequent 1-15, 1-29
 - restore limitations 4-20
- Solaris 1-16, 1-29
- Windows NT 1-4

G

gatekeeper device

- for database server F-3
- for EDM F-2

Glossary 1

granularity

- device-level 1-18
- file-level backups 1-17
- filesystem-as-device 1-15, 1-23

H

hardcopy documentation viii

hot backups 1-8

hyper-volumes 1-3

I

- installation prerequisites F-3
- instance 2-5
- interface to Oracle 1-8
 - configuring 2-3
 - restoring
 - out-of-place 4-3
 - window session 4-12
 - scripts for 2-24
 - starting backups 3-3, 4-5
 - stopping backups 3-6
- interface to UNIX filesystems 1-28
 - configuring 2-14
 - restoring
 - out-of-place 4-8
 - window session 4-12
 - scripts for 2-31
 - starting backups 3-3, 4-5
 - stopping backups 3-6
- interface types supported 1-4

K

- kicker work item 1-20, 3-4, 4-3, D-3

L

- limitations
 - datafile backups 1-17
 - datafile restores 1-23
 - directory backups 1-17
 - filesystem backups 1-15
 - filesystem restores 1-23
- local BCV
 - backups, concept 1-6
 - backups, prompt 2-6
 - restores, concept 1-26
 - restores, prompt 4-14
- logical volume manager (LVM) 1-14, 1-24
- logical volumes 1-14, 1-18, 1-19, 1-24, 4-20
 - Symmetrix 1-3

M

- media rotation 2-12, 2-18, A-7, B-9, C-8
- meta-volumes 1-3
- mirror policies
 - eb_dc_restore 4-19
 - restore prompt 4-15
 - wizard prompt 2-7
- mirror set
 - eb_dc_backup 3-5
 - eb_dc_restore 4-19
 - restore prompt 4-15
 - wizard prompt 2-7
- mirrored backups
 - scripts 2-16, 2-30
- mount point, restore to 4-7
- multi-BCVs 1-7
 - eb_dc_backup -mirror_set 3-5
 - error 2-22
 - restore prompt 4-15
 - .sets file 2-19
 - wizard prompts 2-6

N

- new devices F-2
- no-mirror
 - backups, concept 1-6
 - backups, prompt 2-6
 - restores, concept 1-26
 - restores, prompt 4-14

O

- offline backups 1-8
 - scripts 2-26, 2-29
- online backup mode 1-4, 1-8
- online documentation viii
- online help viii
- Oracle instance 2-5

- Oracle interface 1-4, 1-8, 2-5
 - configuring 2-3
 - restoring
 - out-of-place 4-3
 - window session 4-12
 - scripts for 2-24
 - starting backups 3-3, 4-5
 - stopping backups 3-6
- Oracle pathnames
 - tuning file E-5
- Oracle, other Symm Connect
 - RMAN Proxy Copy 1-4
 - Windows NT 1-4
- ORACLE_SID 2-26
- overwrite standard 1-25

P

- partial database 1-12
- partitions 1-14
- pathnames
 - as Oracle knows E-5
- PDF format document viii
- post-restore 1-25
- pre-and post-backup scripts
 - Oracle interface 2-24
 - UNIX filesystems 2-31
- preparation phase D-7
- prerequisites
 - database server F-3
 - EDM F-2
 - general 2-2
 - SymmAPI, SymmCLI F-2, F-3
- Proxy Copy 1-4

R

- R1, R2 devices 1-6
- raw filesystem 1-15
 - See also* filesystem-as-device
- raw partitions 1-14

- remote BCV
 - backups, concept 1-6
 - backups, prompt 2-6
 - restores, concept 1-26
 - restores, prompt 4-14
 - SRDF link 1-27
- Remote Data Facility 1-6
- requirements
 - database server F-3
 - EDM F-2
- restores
 - 30-day plus old backups 4-2
 - AdvFS filesystems 4-20
 - alternate location 1-24
 - Compaq Tru64 4-20
 - failure, multi-BCV 2-23
 - limitations
 - destination 1-24
 - filesystems 1-23
 - mirrors 1-25
 - mount point 4-7
 - network 1-24
 - no-mirror 1-26
 - out-of-place 1-24
 - Sequent 4-20
 - snapshot 1-26
- rotation period 2-12, 2-18, A-7, B-9, C-8

S

- SAP 1-4
- script
 - post-mirror-split
 - description 2-30
 - prompt 2-16
 - post-offline 2-29
 - pre-offline 2-26
 - UNIX filesystem 2-31
- search capability viii
- sets file 2-19, 3-5
- SID 2-5, 2-26

- snapshot 1-26
- source SRDF device 1-6
- SRDF
 - backups, concept 1-6
 - backups, prompt 2-6
 - link, with remote BCV 1-27
 - restores, concept 1-26
 - restores, prompt 4-14
- standard TimeFinder devices 1-6
- stopping backups 3-6
- striped LVM logical volumes 1-24
- svrmgrl
 - default shutdown commands 2-24
 - failure 2-26
 - sample post-offline 2-29
 - sample pre-offline 2-26
 - select * from v\$datafile E-5
- SymmAPI, SymmCLI F-2, F-3
- Symmetrix Connect work item D-2
- Symmetrix device
 - id verification D-7
- Symmetrix devices 1-3
- Symmetrix Path, use of 1-12, 2-4

T

- tablespace 1-12
- target
 - SRDF device 1-6
- TimeFinder 1-6
- TNS alias 2-5
- tuning file E-1
 - datafiles as files E-5
 - filesystem-as-device E-4

U

- UNIX filesystem interface 1-4, 1-28
 - configuring 2-14
 - restoring
 - out-of-place 4-8
 - window session 4-12

- scripts 2-31
- starting backups 3-3, 4-5
- stopping backups 3-6
- username prompt
 - Oracle
 - DBA UNIX account 2-6
 - post-mirror-split 2-8
 - TNS alias 2-5
 - UNIX filesystem
 - post-mirror-split 2-16

W

- Windows NT 1-4
- wizard failure
 - configuration check 2-11, 2-17
 - multi-BCV 2-23
- work item
 - base name A-6, B-8, C-7
 - kicker D-3
 - network with Symmetrix Connect D-5
 - Symmetrix Connect D-2
 - work item D-3
 - type dc_kicker D-3
 - type dc_worker D-2

