

# **EMC Best Practices: Symmetrix Connect and File Level Granularity**

February 27, 2001

No part of the publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, with the prior written consent of EMC Corporation. The information contained in this document is subject to change without notice. EMC Corporation assumes no responsibility for any errors that may appear. All computer software programs, including but not limited to microcode, described in this document are furnished under a license, and may be used or copied only in accordance with the terms of such license. EMC Corporation either owns or has the right to license the computer software programs described in this document. EMC Corporation retains all rights, title and interest in the computer software programs.

THE INFORMATION CONTAINED IN THIS DOCUMENT IS PROVIDED “AS IS”. EMC Corporation makes no warranties, express or implied by operation of law or otherwise, relating to this document or to the solutions, products or the computer software programs described herein. EMC CORPORATION DISCLAIMS ALL WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. In no event will EMC Corporation be liable for (a) incidental, indirect, special or consequential damages or (b) any damages whatsoever resulting from the loss of use, data or profits, arising out of this document, even if advised or the possibility of such damages.

EMC<sup>2</sup>, EMC, and Symmetrix are registered trademarks and EDM, SRDF, TimeFinder, and where information lives are trademarks of EMC Corporation.

All other brand names are trademarks or registered trademarks of their respective owners.

Copyright © 2001 EMC Corporation. All rights reserved

C872

Printed 8/31/2001

## Table of Contents

<b>Executive Summary</b> .....	<b>4</b>
<b>Introduction</b> .....	<b>5</b>
<b>Symmetrix Connect Technical Explanation</b> .....	<b>5</b>
New Symmetrix Connect with Symmetrix Mapping .....	6
Device, Raw Filesystem, and File-level Backups .....	6
Symmetrix Mirroring Basics .....	7
Stripe Sets and New Disk Geometry Alternatives .....	11
Pre-allocation of Disk Space during Restores of File-level Backups.....	13
Side Effects of Mirror Restores .....	13
<b>Symmetrix Connect Restore Options</b> .....	<b>14</b>
No Mirror Restore.....	15
Restore to a Mirrored Disk .....	15
No Mirror Restore with Protective Mirrored Snapshot .....	16
Out of Place Restore.....	17
Pros and Cons of Each Restore Option.....	18
<b>Backup and Restore Scenarios</b> .....	<b>19</b>
Backup Scenarios and How they Relate to EDM Mirroring Policy .....	19
Restore Scenarios and How they Relate to EDM Mirroring Policy.....	22
<b>Maximizing Symmetrix Connect Utility and Performance</b> .....	<b>23</b>
Notes on Stripe Width and Number of Workitems for Striped Logical Volumes.....	23
Using Symbolic Links or Mountpoints to Better Manage your Restored Data.....	24
<b>Additional References</b> .....	<b>26</b>
<b>Appendix A: Information About Raw Filesystem Backup and Restore</b> .....	<b>27</b>
Side Effects of Raw Filesystem Restores .....	27
Raw Filesystems (files as device) vs. File-Level (files as files) Backup .....	28

## Executive Summary

The EMC Data Manager (EDM™) is a centralized, open systems backup solution that solves a key problem facing IS managers daily: How can I safely and effectively backup and catalog the increasing volume of enterprise data, while reducing the operations window and staffing needs required to accomplish this. EDM has been designed as a highly scalable platform that fits into the open systems environment. This system can meet the needs of today's most demanding information infrastructures. A properly configured EDM can successfully protect any computing environment, from client/server to high-speed Internet Service Providers (ISPs).

In today's marketplace, it may be important to locate and restore a single file from the filesystem, but still backup the system quickly and efficiently. Release 4.6 of the EDM includes Symmetrix® Connect backup with Symmetrix Mapping. This functionality gives EDM users file-level granularity via a fast Symmetrix Connect backup. This paper explains the technical detail of how Symmetrix Connect works and the issues to remember when performing file system backups and restores.

## Introduction

This document begins with a thorough technical explanation on Symmetrix Connect, helping the reader understand important technical details. These details are necessary background for subsequent sections of the document.

The next section outlines restore options available to EDM users. It is important to consider a restore strategy before that need to restore arises because restore options define the backup method used. This section helps users to understand the options available to you and how to implement each option during backup.

The third section describes the mirror options available with Symmetrix Connect and explains the mirror management policies associated with different backup and restore scenarios.

The last section, Maximizing Symmetrix Connect Utility and Performance, describes the relevant technical factors that can help users make important decisions about how to configure Symmetrix Connect to meet the needs of the organization. Symmetrix Connect is a valuable tool for protecting key information. This document can help users understand how to use the version 4.6 features of the tool most effectively.

## Symmetrix Connect Technical Explanation

This section describes the technical framework of a Symmetrix Connect backup and restore. This information can help the reader understand how Symmetrix Connect Backups occur. Topics in this section include:

- New Symmetrix Connect with Symmetrix Mapping
- Device, raw filesystem, and file-level backups explained
- Basic information about Symmetrix mirroring
- Stripe sets and new disk geometry alternatives
- Pre-allocation of disk space during file-level restores
- Side effects of mirror restores
- Side effects of raw filesystem restores

Each of these topics is discussed in detail in the following sections.

## ***New Symmetrix Connect with Symmetrix Mapping***

This paper discusses enhancements that are available with the new Symmetrix Connect with Symmetrix Mapping. The enhancements provide the following new benefits to Symmetrix Connect users:

- Backup and restore can be done on individual files in filesystems, removing any earlier restrictions that allowed only device-level backups.
- Data from all backup types are stored in logical order so there are virtually no restrictions on how data may be cross-restored to different file types and logical volume layouts (including different striping geometries).

At the time of this writing, there are three platforms that can take advantage of this new Symmetrix Connect functionality. These platforms are:

- Sun Solaris
- HP-UX
- IBM AIX

Other platforms will be added to this list in subsequent releases.

## ***Device, Raw Filesystem, and File-level Backups***

With the advent of EMC Data Manager version 4.6 and Symmetrix Connect with Symmetrix Mapping, EDM has much tighter integration with the Symmetrix. Thanks to this tighter integration, Symmetrix Connect can now be configured to back up data using one of the following three methods:

- **Device Backups** — When there is no filesystem present the user must back up the data by device. Any restore from this backup must restore the entire device.
- **Raw Filesystem Backups** — When a filesystem exists, the user has two options for backing up the devices. If the user chooses the option to “back up database data files located within filesystems as raw devices” in the GUI (or sets `rawfs` parameter to `on` in the CLI), then a raw filesystem backup occurs. This method is sometimes referred to as an “all-or-nothing filesystem backup” or “files as devices” because any restore from this backup restores the entire device, even though the administrator can “browse” for a certain file in the restore GUI.
- **File-Level backups** — If the user does not choose the option to “back up database data files located within filesystems as raw devices” in the GUI (or sets `rawfs` parameter to `off` in the CLI), then a file-level backup occurs. This method is sometimes referred to as “files as files” because any restore from this backup has file-level granularity.

Each of these methods offers the right mix of functionality in different situations. The tradeoffs between granularity and performance for these three backup methods are summarized in Table 1.

**Table 1. Symmetrix Connect Tradeoffs.**

Method	Granularity	Performance
File-level	File-level	Good
Raw Filesystem	Device-level	Better
Device	Device-level	Best

There are many factors that affect performance, including striping, average file size, head movement, amount of fragmentation on the device, hardware, etc. The chart above is designed to give a general comparison of the effect each backup method has on performance if all other factors remain unchanged.

## ***Symmetrix Mirroring Basics***

Symmetrix Connect provides automatic mirror management during backup and restore; therefore, the following sections describe mirror terminology, benefits of mirroring, Symmetrix mirroring options available, and mirror management strategies.

### **Key Mirror-Related Terminology**

When a mirror is *established* the Symmetrix is automatically maintaining the mirror image. If the primary disk changes, the same changes are immediately applied to the mirrored disk.

When a mirror is *split* the Symmetrix isolates the mirrored version of the disk and no further changes are applied to the mirrored volume. After a split is complete, the primary disk can continue to change, but the mirror maintains the point-in-time data that existed at the time of the split.

Mirrors can be *synchronized* in either direction. For example, if you wish to apply all the changes that occurred on the primary disk that occurred after a split to the mirror, you can synchronize from the primary to the mirror and the changes will be applied to the mirrored disk. This brings the mirrored disk current with the primary. If you synchronize in the other direction you can make the primary disk match the mirror. This is often the final step during a restore.

## Mirror Benefits

Mirror management is important because of the benefits mirrors provide. Mirrors can create and maintain an exact copy of a disk volume on a second volume. Mirror functionality provides the following key benefits:

- Security from data loss via easily recoverable short-term backups
- Disaster recovery
- Greater data availability and availability options
- Point-in-time data for reporting or other management purposes

## Mirroring Options Available on the Symmetrix

There are three types of mirrors that may be used in a Symmetrix configuration. These types are Local Business Continuance Volumes (LBCVs) created with the EMC TimeFinder™ product, Symmetrix Remote Data Facility (SRDF™), and Remote Business Continuance Volumes (RBCVs), also created using TimeFinder. Each of these options is defined further below:

- **LBCV** — As the name implies, these mirrored disk volumes are located within the same Symmetrix system where the standard volumes reside. LBCVs are commonly used for short-term backup, availability, and storing point-in-time data (see Figure 1).
- **SRDF** — These mirrored disk volumes are usually located in a physically different location from the standard volumes. This can be in a different part of the same building or in a different building. SRDF™ provides disaster recovery solutions by placing the data in physically separate locations (see Figure 1).
- **RBCV** — These RBCVs are located on the same Symmetrix system that contains the SRDF volumes. This RBCV mirrors the SRDF volumes for an extra measure of protection for the remote data (see Figure 1).

The EDM can create backups from any mirrored volumes that are visible to it. For a mirrored volume to be visible to the EDM, the following conditions must be met:

- The EDM must have a direct SCSI or Fibre connection to the Symmetrix where the volume resides
- The Symmetrix must have exported the mirrored volume(s) to the port where the EDM has been connected.



For a graphical representation of all of these mirror options, see the illustration in Figure 1.

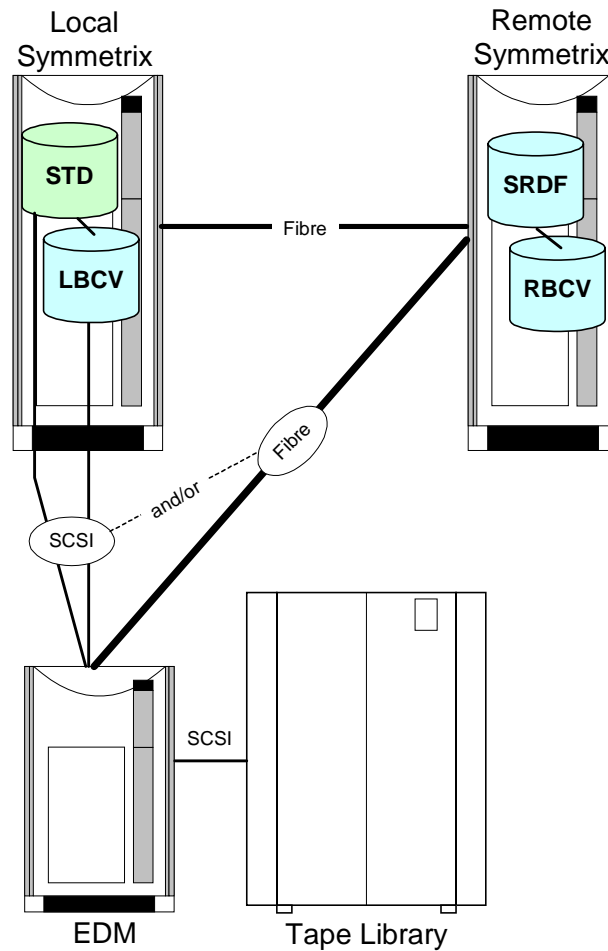


Figure 1. Mirror Types.

## Mirror Management Options

The EDM offers several options for controlling mirror behavior before and after a backup or restore. The option that is right for any particular situation is based on several factors. This section describes each mirror management option and explains what that option does.

### Pre-backup Mirror Policy

During pre-backup processing, mirror management can be configured to do any of the following:

- **Bring Mirrors Down** — This option expects to find the mirrors established and it will split the mirrors automatically before the backup. If the mirrors are down already, the backup will fail and report an error. The error is designed to prevent the system from backing up mirrors that are in an unexpected state.
- **Verify Mirrors are Down** — This option expects to find the mirrors split and it will leave them split and perform the backup. If the mirrors are established at the time of backup, the

backup will fail and report an error. The error is designed to ensure that the backup is taken for the specific point in time that the mirrored data represents.

- **Bring Mirrors Down if Needed** — This option checks whether the mirrors are established or split and it will split the mirrors if they are established. If you select this option, the backup will not fail regardless of the state of the mirrors.
- **Bring Mirrors Down after Establishing** — This option checks the mirrors and if they are not established, the EDM first establishes the mirror to ensure that it is an exact copy of data on the primary volume. Then the EDM splits the mirrors to perform the backup.

### Post-backup Mirror Policy

During post-backup processing, mirror management can be configured to do any of the following:


- **Bring Mirrors Up** — After the restore is complete, the EDM automatically resynchronizes the mirror to the primary disk.
- **Leave Mirrors Down** — After the restore is complete, the EDM leaves the mirrors split from the primary disk.
- **Leave Mirrors as Found** — After the restore is complete, the EDM resynchronizes the mirrors to the primary disk if they were established to begin with. If not, the EDM leaves the mirrors split.

See the section entitled *Backup and Restore Scenarios* for some guidance on how to set these options.

These options are also available at restore time and the restore time options are described in the section entitled *Restore Scenarios and How they Relate to EDM Mirroring Policy* later in this document.

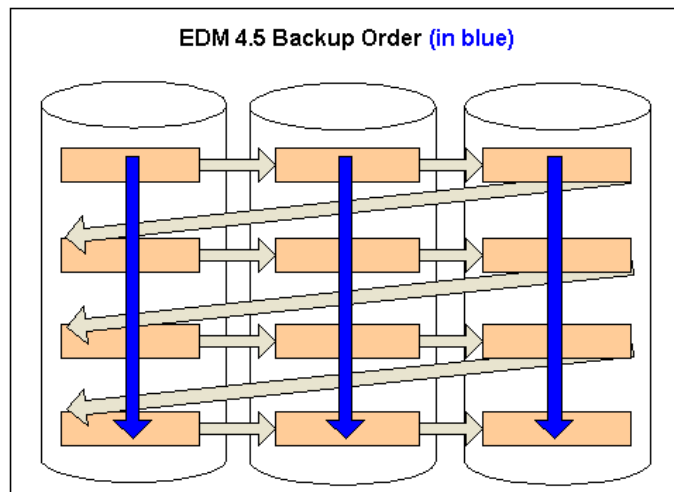
## Stripe Sets and New Disk Geometry Alternatives

Another benefit of the new Symmetrix Connect model with Symmetrix Mapping lies in the EDM's capacity (in version 4.6) to move striped data to tape in a logical order. As the name implies, striped data is stored on disk in stripes that traverse physical disk boundaries (see Figure 2).



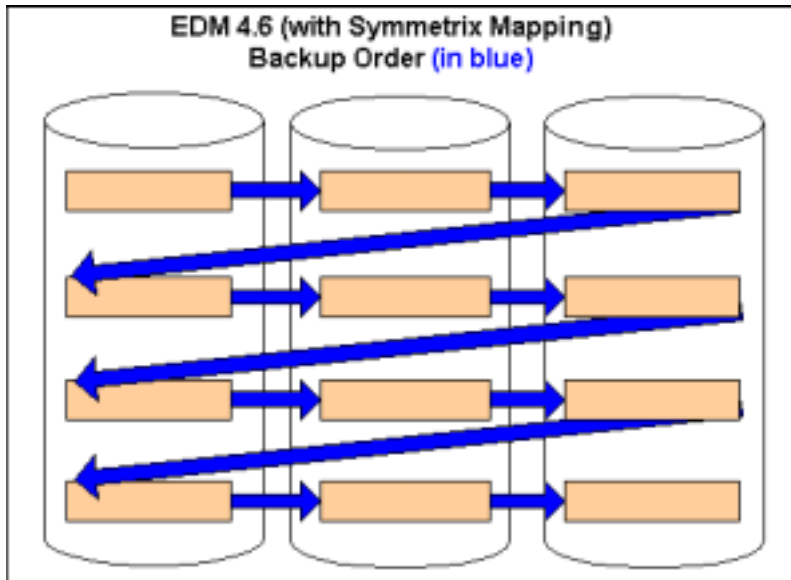
**Figure 2. Stripe Sets (Arrows show logical data organization).**

In EDM version 4.5 and earlier (and in version 4.6 on clients that do not take advantage of the new Symmetrix Mapping), backing up the physical structure backs up the data out of order because each separate disk gets backed up independently (see Figure 3).



**Figure 3. EDM version 4.5 backup order (backs up the data out of order).**

This means that with EDM version 4.5 and earlier, the customer must restore the data to a machine with the same stripe geometry as the original. For example, if the striped data was configured on a three-disk stripe set, it could not be restored to a four-disk stripe set. Also, the size of the stripe blocks must be the same when restoring backups from EDM 4.5 or earlier and data cannot be backed up with Symmetrix Connect then restored using network restore.

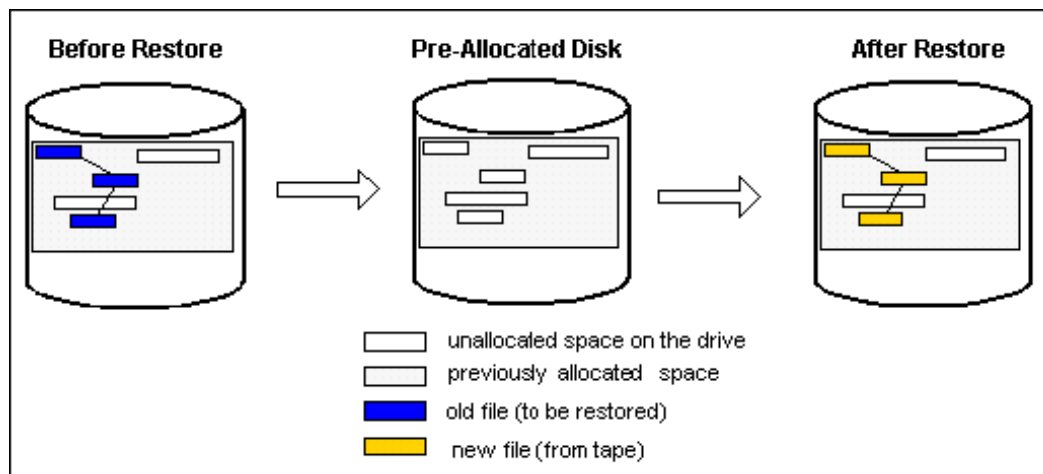


**Figure 4. EDM version 4.6 (with Symmetrix Mapping) backup order (backs up the data in logical order).**

With EDM 4.6, data is backed up in a logical order, exactly as it was added to the disk (see Figure 4). Data backed up in this way can be restored to a striped drive set with a different geometry or even to a concatenated (non-striped) drive configuration and backups can be restored out of place using network restore functionality. This functionality broadens the customer's options when restoring data from a stripe set backup.

## Pre-allocation of Disk Space during Restores of File-level Backups

When a customer sets the `rawfs` parameter to `off` at backup time, the EDM version 4.6 performs a file-level backup. When the customer restores that data, the EDM pre-allocates disk space on the system to which the files are to be restored. This is required to ensure that the precise amount of space has been reserved on disk for each file that will be restored through the Symmetrix Connect data path. Figure 5 shows space allocated for the selected file on the left. When the restore occurs, the EDM first pre-allocates the space occupied by the old version of the file and sets that empty space aside for the data that will be written to it from tape.



**Figure 5. Side Effects of file Pre-allocation.**

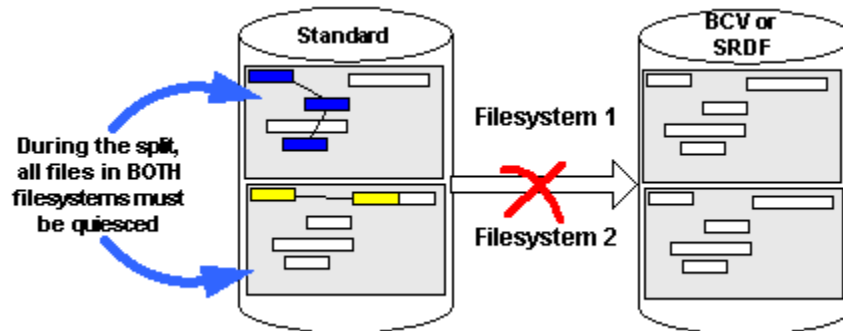
*Note: Regardless of whether the system is performing a no mirror restore to the standard volume or a mirrored restore to a BCV or SRDF mirror, pre-allocation must first occur on the standard volume. The destination filesystem must be mounted to the client, and if mirrors are involved, the standard drive must be synchronized to the mirrored volume(s).*

Pre-allocation prepares the disk to accept the restored files, but pre-allocation can also introduce some potential concerns if the standard volume is active during the restore. These issues are described in the next section.

## Side Effects of Mirror Restores

Mirror restores have certain side effects that may occur when the data on the standard volume changes during a restore to a mirror, such as a BCV or an SRDF. This section describes the issue and provides strategies for avoiding the issue during restore.

If you are restoring certain files to a volume on a mirrored disk, while the standard volume is in use, there is a risk that other files unrelated to the restore may reside on the standard and may be in use during the restore. There is a risk to any files that are changed while the mirror is split to perform the restore, because the mirror will ultimately be resynchronized over the standard copy and all changes to the other files will be lost. Therefore, while mirrors are split, all files on the standard disks must be quiesced to prevent them from changing.



**Figure 6. Side Effects of Mirrored Restores.**

There are several ways to avoid these side effects. These methods are listed below.

- Restore to the original or standard volume (see No Mirror Restore)
- Restore to the original or standard volume after taking a protective mirrored snapshot to a separate mirror (see No Mirror Restore with Protective Mirrored Snapshot)
- Keep everything on the disk quiesced while the restore is in progress
- Perform an out-of-place restore to a separate set of disks on the same or a new client

These options are explored in detail in the section entitled *Symmetrix Connect Restore Options*.

## Symmetrix Connect Restore Options

This section describes the technical details of Symmetrix Connect restore options. This section addresses the following restore types:

- No mirror restore (offline restore to the original standard volume)
- Restore to a mirrored disk
- No Mirror Restore with Protective Mirrored Snapshot (restore to the original standard volume after a mirror snapshot)
- Out of place restore

See the following sections for more information about each of these restore types.

## No Mirror Restore

The No Mirror restore is the most straightforward restore method. With this method, any target data files or databases must be offline during the entire restore process and Symmetrix Connect will unmount any filesystem that is being restored. The data is transferred directly to the standard volume from tape. The disadvantage to this method is that any problem that the restore might introduce cannot be repaired without restoring from a different backup. See Figure 7 for more information about how this restore scenario is configured.




Figure 7. No Mirror Restore

## Restore to a Mirrored Disk

Restore to a mirrored disk is similar to the restore scenario described above. Again, any target data files, tablespaces, or databases must be offline during the entire restore process and Symmetrix Connect will unmount any filesystem that is being restored. The difference here is that the restore is written to an unsynchronized BCV (or SRDF mirrored volume) and therefore the customer has the option not to resynchronize the restored data to the standard volume.

See Figure 8 for more information about how this restore scenario is configured.





Figure 8. Offline Restore to a Mirrored Disk

## **No Mirror Restore with Protective Mirrored Snapshot**

A no mirror restore with protective mirrored snapshot is a variation of the no mirror configuration. This configuration saves a copy of the data to a mirrored volume prior to the restore. In the event that the administrator wishes to retrieve the pre-restore image, the mirrored volume can be resynchronized to retrieve the original state of the standard volume before the restore.

In this scenario, standard volume is synchronized to a BCV or an SRDF mirror, the mirror is split, and then the data is restored to the standard volume. With this scenario, it is not necessary to synchronize data back to the standard volume. However, if a problem occurs due to the restore, the mirrored volume can be resynchronized and the original data can be retrieved. See Figure 9 for more information about how this restore scenario is configured.

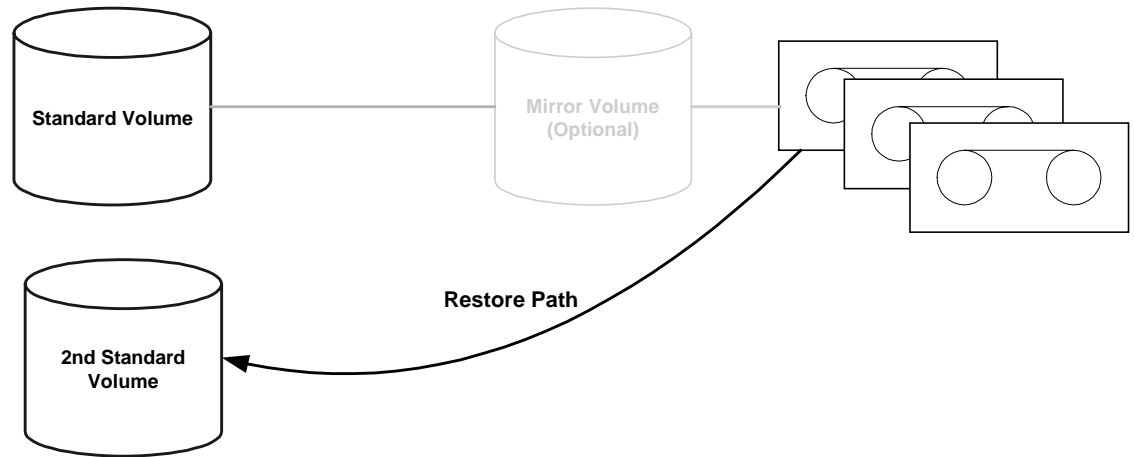


**Figure 9. No Mirror Restore with Protective Mirrored Snapshot.**



## ***Out of Place Restore***

An out of place restore scenario describes how a backup can be restored to a completely separate standard volume on the same or a different machine. This scenario is used for special tasks like restoring Exchange mailboxes, duplicating systems for deployment in remote areas, etc. The actual restore configuration is similar to a no mirror restore because data is restored to an additional standard volume or a mirrored volume that has been remounted to the original host to mimic a standard volume.



**Figure 10. Out of Place Restore.**

## Pros and Cons of Each Restore Option

The following table summarizes the pros and cons of each of the restore options described above:

Restore Option	Pros	Cons
No Mirror Restore	<ul style="list-style-type: none"> <li>No waiting for mirror sync</li> <li>No mirror sync side effects</li> <li>Fewer pre-allocation side effects</li> </ul>	<ul style="list-style-type: none"> <li>Cannot retrieve original (pre-restore) data from a mirror</li> <li>EDM may not have access to standard disks, if not this option is not available</li> </ul>
No Mirror Restore with Protective Mirror Snapshot	<ul style="list-style-type: none"> <li>No pre-allocation side effects</li> <li>Can retrieve original (pre-restore) data from a mirror if desired</li> </ul>	<ul style="list-style-type: none"> <li>Possible mirror sync side effects</li> <li>EDM may not have access to standard disks, if not this option is not available</li> </ul>
Restore to a Mirrored Disk	<ul style="list-style-type: none"> <li>Can avoid resync if an error occurs during restore</li> <li>When the EDM cannot access the standard volume, this is the only option</li> <li>Separate copy of the data exists until resync for device or raw file system backups (File-level backups require an out of place restore to get a separate copy of the data.)</li> </ul>	<ul style="list-style-type: none"> <li>Pre-allocation side effects can occur with restore of a file-level backup</li> <li>Possible mirror sync side effects</li> </ul>

## **Backup and Restore Scenarios**

This section describes various backup and restore scenarios and the considerations required with each.

### ***Backup Scenarios and How they Relate to EDM Mirroring Policy***

This section defines how backup mirror policy can be used to make your operation more efficient and reduce the number of steps performed manually by an administrator.

#### **Scenario 1: Backing up a Mirror used for Short Term Backup**

The most popular use of mirror technology is to protect your enterprise from data loss by keeping an easily recoverable short-term backup of the data on a BCV. These short-term BCV backups are often used to run online backups to tape using the EDM. In this scenario, the maximum protection from loss can be achieved by keeping the mirror constantly established and synchronized. To back up to tape, the EDM must split the BCV to perform the backup.

#### **Scenario 2: Backing up a Mirror used for Disaster Recovery**

Mirroring for disaster recovery is very similar to mirroring for short-term backups described earlier. The only difference is that the administrator should use the SRDF volume for disaster recovery. SRDF allows you to mirror data to another Symmetrix that can be located in a physically separate building. If a localized disaster affects the primary machine, the SRDF machine can be used as a substitute.

The mirroring policy choices are the same as those outlined for mirroring short-term backups.

#### **Scenario 3: Backing up a Mirror used for Reporting and/or Data Management**

Sometimes companies use BCV or SRDF mirrors to capture a stable copy of data as it existed at a given point in time. For example, a company may split a BCV or SRDF at 11:59 PM. In this scenario, the BCV or SRDF may be split from the primary volume at backup time and it should remain split.

See Table 2 for suggested backup policy associated with each scenario.

**Table 2. Backup Policy Suggestions for Various Backup Scenarios Outlined Above.**

Scenario	Suggested Pre-backup Policy	Suggested Post-backup Policy	Notes
Scenario 1: Backing Up a Mirror Used for Short-term Backup	<p align="center"><b>Bring Mirrors Down</b></p> <p align="center"><b>OR</b></p>		<p>Ensures that the mirrors have remained established and are established at the time of backup.</p> <p>With this option selected, if the mirrors went down due to some unscheduled event, the backup will fail.</p>
	<p align="center"><b>Bring Mirrors Down after Establishing</b></p>		<p>Ensures that the BCV is an exact copy of the primary disk at the time of the backup.</p> <p>If you choose this option, the backup will not fail regardless of the state of the mirrors at backup time.</p>
		<p align="center"><b>Bring Mirrors Up</b></p>	<p>Ensures that the mirrors get reestablished after each backup.</p> <p>An established mirror provides up-to-the-minute backups that protect the data should some unforeseen event destroy or corrupt data on the primary volume.</p>

Scenario	Suggested Pre-backup Policy	Suggested Post-backup Policy	Notes
Scenario 2: Backing up a Mirror Used for Disaster Recovery	<p align="center"><b>Bring Mirrors Down</b></p> <p align="center"><b>OR</b></p>		<p>Ensures that the mirrors have remained established and are established at the time of backup.</p> <p>With this option selected, if the mirrors went down due to some unscheduled event, the backup will fail.</p>
	<p align="center"><b>Bring Mirrors Down after Establishing</b></p>		<p>Ensures that the BCV is an exact copy of the primary disk at the time of the backup.</p> <p>If you choose this option, the backup will not fail regardless of the state of the mirrors at backup time.</p>
		<p align="center"><b>Bring Mirrors Up</b></p>	<p>Ensures that the mirrors get reestablished after each backup.</p> <p>An established mirror provides up to the minute backups that protect the data should some unforeseen event destroy or corrupt data on the primary volume.</p>
Scenario 3: Backing Up a Mirror Used for Reporting and Data Management	<p align="center"><b>Verify Mirrors are Down</b></p> <p align="center"><b>OR</b></p>		<p>Ensures that the backup will not take place unless the mirrors are split.</p> <p>With this option selected, if the mirrors were established at backup time, the backup will fail.</p>
	<p align="center"><b>Bring Mirrors Down if Needed</b></p>		<p>Will split the mirrors if they are established.</p> <p>If you choose this option, the backup will not fail regardless of the state of the mirrors at backup time.</p>
		<p align="center"><b>Leave Mirrors Down</b></p> <p align="center"><b>OR</b></p>	<p>Ensures that mirrors will be left down in every case after the backup.</p>
		<p align="center"><b>Leave Mirrors as Found</b></p>	<p>Restores the mirrors to the state in which they were found.</p>

## ***Restore Scenarios and How they Relate to EDM Mirroring Policy***

This section defines various restore scenarios and the affect each scenario has on mirror policy (if any). Restores are performed for many reasons to various destinations. Many factors can affect the mirroring policy that an administrator uses during restore. All scenarios shown here relate specifically to file-level backups.

### **Scenario 1: Restore of a File-level Backup (Files as Files) to the Standard Volume**

Target files are pre-allocated on the standard devices so that the precise amount of disk space is reserved for each file prior to putting back their data. Note that this will overwrite any pre-existing files in the target locations (see scenario 3 for “out-of-place” techniques). After the new files location have been discovered the data is restored via the Symmetrix Connect data path. With this in-place, no-mirror scenario, if there are problems with the restore there is no way of rolling back to the pre-restore versions of the file. See scenario 2 for a way to prevent this drawback.

### **Scenario 2: Restore of a File-level Backup to the Standard Volume using a No Mirror Restore with Protective Mirrored Snapshot**

A no mirror restore with protective mirrored snapshot is a variation of the no mirror configuration. With this option, the administrator first saves a copy of the data to a mirrored volume prior to the restore, and then the steps in scenario 1 are performed. In the event that the administrator wishes to retrieve the pre-restore image, the mirrored volume can be resynchronized to retrieve the original state of the standard volume before the restore. This is a very safe restore method.

This method acts exactly like a “no-mirror” restore so there are no pre- and post-mirror activities.

### **Scenario 3: Restore of a File-level Backup Out of Place**

Out of place restores are another very safe way to restore information from file-level backups. This method restores to a standard volume mounted to the same or a secondary client similar to scenario 1, but the administrator takes steps to keep the restored data separate from the backup location.

One variation of this scenario uses a BCV, splits that BCV from its primary volume, and remounts it to the same client or a completely separate client as a new standard volume. (These are manual steps performed before the restore can take place.) Once the data is restored to this volume, two techniques can be used to make the data appear in the appropriate location on the system. The system can use links in the pathnames to make the data appear in a different location or this new volume can be mounted under the original mountpoint of the standard it is replacing.

This method acts exactly like a “no-mirror” restore so there are no pre- and post-mirror activities.

### **Scenario 4: Restore of a File-level Backup (Files as Files) to a BCV**

In this scenario, target files are pre-allocated on the standard devices while the mirrors are established in order to reserve space in the filesystem, but the data is actually restored to the BCV. This scenario may be necessary if the EDM cannot access the standard devices. Note that the pre-allocation step will overwrite the standard copies of the target files unless they are linked to different filesystems (refer to the out-of-place techniques in scenario 3). After the data is restored to the BCV, the BCV can be restored to the standard.

If you do perform a standard BCV restore, the applicable pre- and post-restore mirror policies are:

**Pre-restore Mirror Policy** — You should choose **Bring Mirrors Down after Establishing**. This policy makes sure that the most up-to-date data is stored on the BCV before the restore takes place. It is also important to stop all activity on the drive even though it remains mounted because activity on the drive may invoke the affected entities issue defined earlier in this document.

**Post-restore Mirror Policy** — **Leave Mirrors Up** should be the post-restore mirror policy.

*Note: If your policy is to export only mirrored volumes to the EDM, then you must choose this restore method.*

## Scenario 5: Restore Symmetrix Connect Files via the Network or Symmetrix Path

For small restores, using network or Symmetrix Path restores can be a very convenient way to restore files that were backed up with Symmetrix Connect. Files restored in this scenario will always be written to standard devices, but there is no space pre-allocation overhead, and the files may be linked to alternate locations.

This method acts exactly like a “no-mirror” restore so there are no pre- and post-mirror activities.

## Maximizing Symmetrix Connect Utility and Performance

Symmetrix Connect is a high-performance centralized backup system scaled for use with very large database files. It is critical to appropriately consider and balance benefits of certain restore options with the performance implications of each option. Different restore options are better suited to different backup needs. This section can help users identify the best restore option for a specific need. The following considerations will be explored:

- Notes on Stripe Width and Number of Workitems for Striped Logical Volumes
- Using symbolic links or mountpoints to better manage your restored data

### ***Notes on Stripe Width and Number of Workitems for Striped Logical Volumes***

When selecting stripe widths and number of workitems for stripe sets, there are some rules of thumb derived by EMC’s performance group that can help you achieve the best possible performance. Those rules are as follows:

- Performance tests show that the I/O operations associated with stripe widths of between 32 KB and 1024 KB are the most efficient sizes. Stripe widths of 16 KB are about 25 percent less efficient than 32 KB stripe widths. Stripe widths that are smaller than 32 KB definitely adversely affect performance.
- Performance tests also indicate that for stripe sets with a large number of spindles (for 24, 36, and 48 spindles) the ratio of drives to workitems is about 3 to 1.

Please remember that these performance related guidelines can be affected by many factors and that your results may differ.

## ***Using Symbolic Links or Mountpoints to Better Manage your Restored Data***

Symbolic links are a way to place data at a mountpoint on one volume and reference it via a link from another mount point which can be on a completely different volume. Of course an entire volume, restored with an out-of-place restore, can be mounted to any appropriate mountpoint.

To create a symbolic link (to a file or a directory) use the following command:

```
ln -s <full_path_link_points_to> <link_name>
```

For example...

```
ln -s /oracle/dat /data
```

(creates a link called /data to the directory /oracle/dat)

```
ln -s /tmp/myfile.dat datafile.dat
```

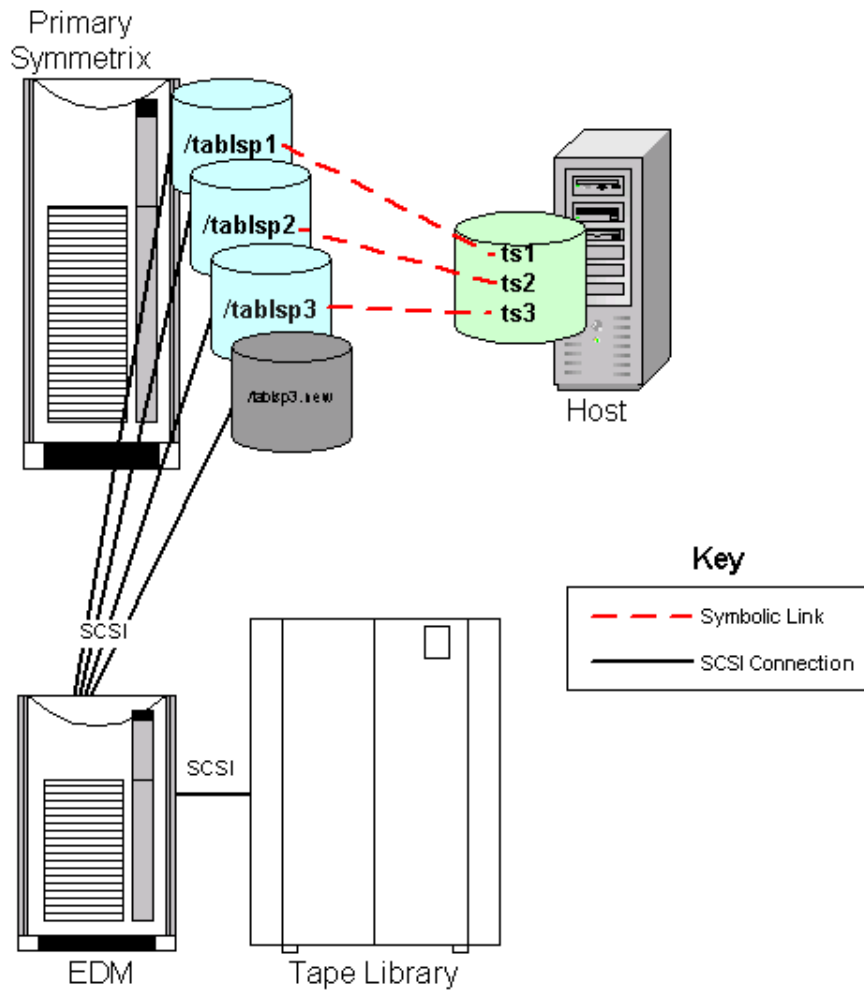
(creates a link called datafile.dat to the file  
myfile.dat in the /tmp directory)

Many administrators use these symbolic links to manage their data. Symbolic links can also make it easier to repair part of your database without some of the risks associated with a direct restore from tape to the standard volume. The following example shows how these links can make it possible to restore data with a minimum of risk.

Consider the illustration in Figure 11. The tablespaces for the Oracle database are all stored on separate volumes and linked to from a single mount point using symbolic link.

If one of these tablespaces becomes corrupt or unstable for any reason, an out of place file-level restore of just the files associated with that filespace can be restored to /tmp and then the symbolic link can be changed to point to /tmp. See Figure 13 for an example of what the configuration would look like after the restore.





**Figure 11. Using Symbolic Links.**




Figure 12. Configuration after a restore using symbolic links.

## Additional References

For more information about Symmetrix Connect, reference the following sources:

- *EMC Data Manager Symmetrix Connect User's Guide*
- *EMC Data Manager Oracle Backup Client 3.0.1 Release Notes*
- *EMC Data Manager Oracle Backup Client*
- *EMC Data Manager Software Installation*
- *EMC Data Manager Demonstrates Symmetrix Connect Functionality Over Nortel OPTera Network At Comdex Fall 2000*
- *Validation of Oracle RMAN Proxy Copy with EDM Symm Connect*

## **Appendix A: Information About Raw Filesystem Backup and Restore**

This section describes some pertinent information about the Raw Filesystem backup and restore. This method is also referred to as “files as devices” because although the raw filesystem backup contains the identity of all the files in the filesystem, the data can only be restored at device granularity.

### ***Side Effects of Raw Filesystem Restores***

A side effect of Raw Filesystem restores that can cause data loss occurs when an administrator tries to restore a single file from a raw filesystem backup. Suppose the administrator selects only a single file from a raw filesystem backup, the entire device will be restored, overwriting any data that has changes since the backup.

That is why the raw filesystem backup is sometimes referred to as a filesystem all-or-nothing backup. For file-level restore granularity, the system must be backed up using the file-level backup method described in this document.

## ***Raw Filesystems (files as device) vs. File-Level (files as files) Backup***

The pros and cons of each of these methods are described in the following table:

<b>Method</b>	<b>Pros</b>	<b>Cons</b>
<b>Raw Filesystem (Files as Devices) Backup and Restore</b>	<ul style="list-style-type: none"><li>• Filesystem backup without extra overhead in certain layouts.</li><li>• No file pre-allocation issues.</li></ul>	<ul style="list-style-type: none"><li>• Device-level granularity only.</li><li>• Each filesystem must be backed up using only one stream.</li></ul>
<b>File-level (Files as Files) Backup and Restore</b>	<ul style="list-style-type: none"><li>• File-level granularity.</li><li>• Administrators can backup and restore only those files that are needed.</li><li>• Decreased backup windows because only selected files are backed up.</li><li>• In some cases, other parts of the system can continue to operate while backup occurs.</li><li>• Multiple streams can be used to restore data.</li></ul>	<ul style="list-style-type: none"><li>• Requires pre-allocation of file space on the drive.</li></ul>