

IBM TotalStorage<sup>®</sup> DS6000



# Recovering



IBM TotalStorage<sup>®</sup> DS6000



# Recovering

**Note:**

Before using this information and the product it supports, read the information in "Notices" on page 17.

**Twenty-fourth Edition (January 2007)**

**© Copyright International Business Machines Corporation 2004, 2007. All rights reserved.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

---

# Contents

## Tables . . . . . v

## Notices and publication information . . vii

Safety notices . . . . . vii

Environmental notices. . . . . vii

    Product recycling and disposal. . . . . vii

    Battery return program . . . . . viii

How to send your comments . . . . . ix

## Chapter 1. Recovering . . . . . 1

## Chapter 2. Disaster recovery using Copy Services . . . . . 3

Performing failover and failback operations (without Global Mirror). . . . . 4

    Moving production to Site B after planned outages (failover). . . . . 4

    Moving production to Site B after unplanned outages (failover). . . . . 6

    Returning production to Site A after planned and unplanned outages (failback) . . . . . 7

Performing failover and failback operations (with Global Mirror). . . . . 9

    Global Mirror setup (DS Storage Manager) . . . . 9

Using Global Mirror for a planned failover and failback. . . . . 10

Using Global Mirror for an unplanned failover and failback . . . . . 12

Recovering with and without the fast reverse restore feature . . . . . 14

## Notices . . . . . 17

Accessibility . . . . . 18

Trademarks . . . . . 19

Terms and conditions . . . . . 20

Electronic emission notices . . . . . 20

    Federal Communications Commission (FCC) statement . . . . . 21

    Industry Canada compliance statement . . . . 21

    European community compliance statement . . 21

    Japanese Voluntary Control Council for Interference (VCCI) class A statement. . . . . 22

    Korean Ministry of Information and Communication (MIC) statement . . . . . 22

    Taiwan class A compliance statement. . . . . 22

## Index . . . . . 23



---

## Tables

1. Failover and failback implementation . . . . 5





---

## Notices and publication information

This section contains information about safety notices that are used in this guide, environmental notices for this product, publication information, and information about sending your comments to IBM.

---

### Safety notices

Complete this task to find information about safety notices.

To find the translated text for a danger or caution notice:

1. Look for the identification number at the end of each danger notice or each caution notice. In the following examples, the numbers **1000** and **1001** are the identification numbers.

#### **DANGER**

**A danger notice indicates the presence of a hazard that has the potential of causing death or serious personal injury.**

**1000**

#### **CAUTION:**

**A caution notice indicates the presence of a hazard that has the potential of causing moderate or minor personal injury.**

**1001**

2. Find the number that matches in the *IBM System Storage Solutions Safety Notices for IBM Versatile Storage Server and IBM System Storage Enterprise Storage Server, GC26-7229*.

---


### Environmental notices

This section identifies the environmental guidelines that pertain to this product.

#### **Product recycling and disposal**

This unit contains recyclable materials.

This unit must be recycled or discarded according to applicable local and national regulations. IBM® encourages owners of information technology (IT) equipment to responsibly recycle their equipment when it is no longer needed. IBM offers a variety of product return programs and services in several countries to assist equipment owners in recycling their IT products. Information on IBM product recycling offerings can be found on IBM's Internet site at <http://www.ibm.com/ibm/environment/products/prp.shtml>.



**Notice:** This mark applies only to countries within the European Union (EU) and Norway.

Appliances are labeled in accordance with European Directive 2002/96/EC concerning waste electrical and electronic equipment (WEEE). The Directive determines the framework for the return and recycling of used appliances as applicable throughout the European Union. This label is applied to various products to indicate that the product is not to be thrown away, but rather reclaimed upon end of life per this Directive.


In accordance with the European WEEE Directive, electrical and electronic equipment (EEE) is to be collected separately and to be reused, recycled, or recovered at end of life. Users of EEE with the WEEE marking per Annex IV of the WEEE Directive, as shown above, must not dispose of end of life EEE as unsorted municipal waste, but use the collection framework available to customers for the return, recycling and recovery of WEEE. Customer participation is important to minimize any potential effects of EEE on the environment and human health due to the potential presence of hazardous substances in EEE. For proper collection and treatment, contact your local IBM representative.

## Battery return program

This product may contain sealed lead acid, nickel cadmium, nickel metal hydride, lithium, or lithium ion battery. Consult your user manual or service manual for specific battery information. The battery must be recycled or disposed of properly. Recycling facilities may not be available in your area. For information on disposal of batteries outside the United States, go to <http://www.ibm.com/ibm/environment/products/batteryrecycle.shtml> or contact your local waste disposal facility.

In the United States, IBM has established a return process for reuse, recycling, or proper disposal of used IBM sealed lead acid, nickel cadmium, nickel metal hydride, and other battery packs from IBM Equipment. For information on proper disposal of these batteries, contact IBM at 1-800-426-4333. Please have the IBM part number listed on the battery available prior to your call.

In the Netherlands the following applies:



For Taiwan:



Please recycle batteries.

廢電池請回收

---

## How to send your comments

Your feedback is important to help us provide the highest quality information. If you have any comments about this information or any other DS6000™ series documentation, you can submit them in the following ways:

- e-mail

Submit your comments electronically to the following e-mail address:

starpubs@us.ibm.com

Be sure to include the name and order number of the book and, if applicable, the specific location of the text you are commenting on, such as a page number or table number.

- Mail

Fill out the Readers' Comments form (RCF) at the back of this book. Return it by mail or give it to an IBM representative. If the RCF has been removed, you can address your comments to:

International Business Machines Corporation  
RCF Processing Department  
Department 61C  
9032 South Rita Road  
TUCSON AZ 85775-4401

---

## Chapter 1. Recovering

The topics in this section provide recovery information related to your DS6000. Topics covered include concepts and procedures related to copy services and clusters.



---

## Chapter 2. Disaster recovery using Copy Services

One of the main reasons for using Copy Services functions is to prepare for a possible disaster by backing up, copying, and mirroring your data both at the local (production) and remote sites.

Having a disaster recovery plan can ensure that critical data is recoverable at the time of a disaster. Because most disasters are unplanned, your disaster recovery plan must provide a way that allows you to recover your applications quickly, and more importantly, to access your data. Consistent data to the same point-in-time across all storage units is vital before you can recover your data at a backup (normally your remote) site.

Most users use a combination of remote mirror and copy and point-in-time copy (FlashCopy) features to form a comprehensive enterprise solution for disaster recovery. In an event of a planned event or unplanned disaster, you can use failover and failback modes as part of your recovery solution. Failover and failback modes help to reduce the time that is required to synchronize remote mirror and copy volumes after you switch between the local (or production) and the remote sites during planned and unplanned outages. Although failover transmits no data, it changes the status of a device, and the status of the secondary volume changes to a suspended primary volume. The Failback command transmits data and can go in either direction depending on which device the Failback command is issued to.

Recovery procedures that include failover and failback modes use remote mirror and copy functions, such as Metro Mirror, Global Copy, Global Mirror, Metro/Global Mirror, and FlashCopy.

**Note:** See the *IBM System Storage DS6000 Command-Line Interface User's Guide* for specific disaster recovery tasks.

Data consistency can be achieved using the following methods:

### **Manually using external software (without Global Mirror)**

If you use Metro Mirror, Global Copy, and FlashCopy functions to create a consistent and restartable copy at your recovery site, you must do a manual and periodic suspend operation at your local site. This means using *freeze and run* commands together with external automated software and then using the FlashCopy function to make a consistent copy of your target volume for backup or recovery purposes. (Automation software is not provided with the storage unit; it must be supplied by the user.)

**Note:** Freezing of the data is done at the same point-in-time across all links and all storage units.

### **Automatically (with Global Mirror and FlashCopy)**

If you use a two-site Global Mirror configuration, the process to create a consistent and restartable copy at your remote site is done using an automated process, with minimal or no interruption to your applications. Global Mirror operations automate the process of continually forming consistency groups. It combines Global Copy and FlashCopy operations to provide consistent data at the remote site. A master storage unit (along with subordinate storage units) internally manages data consistency using consistency groups within a Global Mirror configuration. Consistency

groups can be created many times per hour to increase the currency of data that is captured in the consistency groups at the remote site.

**Note:** A consistency group is a collection of volumes (grouped in a session) across multiple storage units that are managed together in a session during the creation of consistent copies of data. The formation of these consistency groups is coordinated by the master storage unit, which sends commands over remote mirror and copy links to its subordinate storage units.

In a two-site Global Mirror configuration, if you have a disaster at your local site and have to start production at your remote site, you can use the consistent point-in-time data from the consistency group at your remote site to recover when the local site is operational.

---

## Performing failover and failback operations (without Global Mirror)

In the event of a planned outage, such as a scheduled maintenance or an unplanned outage, such as a hardware failure that disables the operation of your production site, you can perform a failover operation to your recovery site to continue operations. After your production site is operational, you can perform a subsequent failback operation to move production back to its original location.

### Moving production to Site B after planned outages (failover)

When you schedule a planned outage at your production site (Site A), you can switch production to your recovery site (Site B), allowing the processing of data to resume at Site B. This process is known as a failover recovery.

The storage units at both Site A and Site B must be functional and accessible.

In a disaster recovery environment, when two storage units are set up in two geographically distinct locations, the storage unit at the production or local site is referred to as Site A and the storage unit at the remote or recovery site as Site B.

For this scenario, assume that all I/O to Site A has ceased because of a planned outage, such as a scheduled maintenance. The failover operation is issued to the storage unit that will become the primary. That is, production is moved to Site B during this outage, which makes the target volumes at Site B convert to source volumes and causes them to enter a suspended state. Your original source volumes at Site A remain in the state they were in at the time of the site switch. Table 1 on page 5 provides an example of the implementation of failover and failback operations.

**Note:** The failover recovery operation does not reverse the direction of a remote mirror and copy pair. It changes a target volume into a suspended source volume, while leaving the source volume in its current state.

The following assumptions are made for this scenario:

- Applications continue to update the source volumes that are located at Site A.
- Paths are established from Site A to Site B.
- Volume pairs are in duplex state.

The following steps summarize the actions that you must take to move production to Site B after you initiate a planned outage at Site A.



1. Quiesce applications to cease all write I/O from updating the source volumes when the planned outage window is reached. Quiescing your applications might occur as part of a planned outage, but the delay in processing caused by the quiesce action should be brief.

**Note:** On some host systems, such as AIX®, Windows®, and Linux®, before you perform FlashCopy operations, you must quiesce the applications that access FlashCopy source volumes. The source volumes must be unmounted (depending on the host operating system) during FlashCopy operations. This ensures that there is no data in the buffers that might be flushed to the target volumes and potentially corrupt them.

2. Perform a failover recovery operation to Site B. After the failover operation has processed successfully, the volumes at Site B transition from target to source volumes.
3. Create paths in the opposite direction from Site B to Site A depending on your path design and when the source storage unit becomes available. You need the paths in the opposite direction because you want to transfer the updates back to Site A.
4. Rescan (this is dependant on your operating system) your fibre-channel devices. The rescanning removes device objects for the site A volumes and recognizes the new source volumes.
5. Mount your target volumes (now the new source volumes) on the target storage unit at Site B.
6. Start all applications. After the applications start, all write I/O operations to the source volumes are tracked. Depending on your plans regarding Site A, the volume pairs can remain suspended (if you want to do offline maintenance).
7. Initiate a failback recovery operation when your scheduled maintenance is complete. The failback recovery operation initiates the transfer of data back to Site A. This process resynchronizes the volumes at Site A with the volumes at Site B.

**Note:** Failback recovery operations are usually used after a failover recovery has been issued to restart mirroring either in the reverse direction (remote site to local site) or original direction (local site to remote site).

Table 1 provides an example of the implementation of failover and failback operations:

*Table 1. Failover and failback implementation*

Step	Operation	MC connectivity required to	Format of source volume and target volume	Format of source and target volume pair	Result: Site A	Result: Site B
1						

Table 1. Failover and failback implementation (continued)

Step	Operation	MC connectivity required to	Format of source volume and target volume	Format of source and target volume pair	Result: Site A	Result: Site B
Disaster at Site A	Failover	Site B	Volume B, Volume A	Volume B1 : Volume A1	Volume A1 -> Volume B1 (Suspended) The volume pair might display as full or pending duplex state if host write operations have stopped.	Volume B1 -> Volume A1 (Suspended)
2 (Site A volumes must be in a suspended state)						
Return production to Site A	Failback	Site A	Volume A, Volume B	Volume A1 : Volume B1	Volume A1 -> Volume B1	Volume A1 -> Volume B1
3a (Site B volumes must be in a suspended state)						
Return to production (Site B) <b>Note:</b> If Site A is still not operational; production can continue at Site B.	Failback	Site B	Volume B, Volume A	Volume B1 : Volume A1	Volume B1 -> Volume A1	Volume B1 -> Volume A1
3b (prepare to return to production (Site A) from production (Site B))	Failover	Site A	Volume A, Volume B	Volume A1: Volume B1	Volume A1 -> Volume B1	Volume B1 -> Volume A1 (Suspended state; the volume pair might display full or pending state if host write operations have stopped.)
3c (Site A volumes must be in a suspended state)						
Return to production - Site A	Failback	Site A	Volume A, Volume B	Volume A1: Volume B1	Volume A1 -> Volume B1	Volume A1 -> Volume B1

## Moving production to Site B after unplanned outages (failover)

For this scenario, assume that an unexpected failure occurs at Site A. The failure at Site A causes the volumes to be suspended or causes a mix of suspended and full duplex volume pairs because the input might have been written to those volumes when the failure occurred.

In a disaster recovery environment, when two storage units are set up in two geographically distinct locations, the storage unit at the production or local site is referred to as Site A, and the storage unit at the remote or recovery site is referred to as Site B.

The failover operation is performed on the storage unit that will become the primary. Production is moved to Site B during this outage, which makes the target volumes at Site B convert to source volumes. The volumes are designated as in a suspended state. Your original source volumes at Site A remain in the state that they were in at the time of the site switch. When Site A is available again, application I/O is switched back from Site B to Site A.

The following steps summarize the actions you must take to move production to Site B as the result of an unplanned outage and then return production to Site A after it recovers.

1. Perform a failover recovery operation to Site B. After the failover operation has processed successfully, the volumes at Site B transition from target to source volumes.
2. Mount your target volumes on your server at Site B.
3. Start your applications on your server at Site B.
4. After Site A recovers, proceed with the following steps, which are the first steps toward the recovery of the volumes at Site A.
  - a. Create paths between LSSs at Site B to Site A to allow the volumes at Site A to be synchronized with the Site B volumes.
  - b. Delete any remote mirror and copy volume relationships that still exist from the source volumes.
  - c. Wait until the volumes are in full duplex state, and then schedule a time to perform a failback recovery operation using the volumes at Site A. This process resynchronizes the volumes at Site A with the volumes at Site B.

**Note:** Failback recovery operations are usually used after a failover recovery has been performed to restart mirroring either in the reverse direction (remote site to local site) or in the original direction (local site to remote site).

## Returning production to Site A after planned and unplanned outages (failback)

Returning production to its original implementation is called a failback recovery. After restoring operations at Site A, you can schedule a failback operation to synchronize data and to enable production to resume at your original site, Site A.

Before you run a failback operation, you must create paths from Site B to Site A between the specific LSSs.

For this scenario, assume that Site A is operational and that connectivity from Site B to Site A is available. Use this procedure to restart your production environment using the volumes from Site B. See Table 1 on page 5 for an example of the implementation of failover and failback operations.

**Note:** The process to move production back to your local site (Site A) for a planned outage is identical to the one that is used for an unplanned outage. Therefore, this procedure outlines the steps for both outages.

The failback operation resynchronizes the volumes in the following manner depending on the volume state:

- If a volume at Site A is in simplex state, all of the data for that volume is sent from Site B to Site A.
- If a volume at Site A is in full-duplex or suspended state and without changed tracks, only the modified data on the volume at Site B is sent to the volume at Site A.
- If a volume at Site A is in a suspended state but has tracks that have been modified, the volume at Site B will discover which tracks were modified at any site and send both the tracks that were changed on Site A and the tracks that were marked at Site B from Site A to Site B.

The following assumptions are made for this scenario:

- Paths from Site B to Site A are created.
- Remote mirror and copy volume pairs are created. Site B volume is the source volume of the failback operation. This volume was initially the target volume of the relationship.

**Note:** The failback recovery operation can be issued against any remote mirror and copy volume that is in a primary suspended state. The operation copies required data from the source volume to the target volume in order to resume mirroring. Failback recovery operations are usually used after a failover recovery has been issued to restart mirroring either in the reverse direction (remote site to local site) or original direction (local site to remote site). However, this process also works if the target volume is in simplex state.

Perform the following steps using the DS Storage Manager. You can also use the DS CLI to perform Copy Services functions.

1. Perform a failback recovery operation using volumes at Site B. This process copies all changed tracks from the target volumes back to the source volumes and copies over any tracks that were modified on the original source volumes.
2. Before returning to normal operation, quiesce your applications (still updating volumes at Site B) to cease all write I/O from updating the source volumes.

**Note:** On some host systems, such as AIX, Windows, and Linux, before performing FlashCopy operations, you must quiesce your applications that access FlashCopy source volumes. The source volumes must then be unmounted during the FlashCopy establishment. This is to ensure that there is no data in the buffers that could be flushed to the target volumes and potentially corrupt them. Depending on the host operating system, it might be necessary to unmount the source volumes.

3. From Site A, perform a failover recovery operation using the source volumes. This process converts the full-duplex target volumes at the Site A to suspended source volumes. The volumes at Site A start the change recording process while in failover mode.
4. Depending on your operating system, it might be necessary to rescan fibre-channel devices and mount the new source volumes at Site A.
5. From Site A, perform another failback recovery operation using the source volumes. This process resynchronizes the volumes at Site A with volumes at Site B.

- Note:** Failback recovery operations are usually used after a failover recovery has been issued to restart mirroring either in the reverse direction (remote site to local site) or original direction (local site to remote site).
6. Mount your volumes at Site A and start your applications on your primary server.

---

## Performing failover and failback operations (with Global Mirror)

Global Mirror is a long-distance, two-site, remote copy solution that uses asynchronous technology. Global Mirror is based on existing Global Copy and FlashCopy functions.

The following tasks are illustrated:

- Setting up your environment to use Global Mirror
- Using Global Mirror for a planned failover
- Using Global Mirror for an unplanned failover
- Using recovery procedures with and without the fast reverse feature

### Global Mirror setup (DS Storage Manager)

Use this DS Storage Manager process to set up your environment to use Global Mirror. Global Mirror processing is used for general transactions as well as with data recovery transactions. When Global Mirror is used for data recovery it asynchronously copies data from a host to a recovery site and maintains data on a storage unit at the recovery site.

You can configure the following Global Mirror environment, which uses two sites (local and remote) and two or more storage units.

#### Local site

Contains A volumes (the source volume), which are copied to the recovery site using Global Copy

#### Recovery (or remote) site

Contains B volumes (the target volume and FlashCopy source volume) and C volumes (the FlashCopy target volume)

**Note:** A storage unit at the local site is designated as the Global Mirror master and all other local (or production) storage units are designated as subordinate storage units. The master storage unit sends commands to its subordinate storage units. These subordinates work together to create a consistency group and to communicate the FlashCopy commands to the recovery (or remote) site. All status is relayed back to the Global Mirror master.

You can perform this task using either the DS CLI or the DS Storage Manager.

#### DS Storage Manager setup steps

Perform the following steps to set up your environment for Global Mirror processing and to create a consistent backup copy of your data.

1. **Ensure that the storage units that you are using are configured, assigned, and operating in a normal state.** See "Storage Units — Main Page" for details.
2. **Ensure that fibre-channel paths are established.**
  - All Global Copy source and target pairs
  - Master and subordinate storage units

**Note:** To establish a volume pair, there must be a path between the LSSs in which the volumes reside. See "Creating remote mirror and copy paths" for the task procedure.

3. **Create Global Copy pairs.** Create the volumes from the local storage units to the recovery storage units using the path that you established. See "Creating Metro Mirror volume pairs" for the task procedure.

**Note:** When you create volume pairs using Global Copy, data on the volumes is not time consistent and generally not useful for the recovery site unless all of the data from the local site has been copied up to a specific point in time. To be useful for disaster recovery, the target volumes must periodically be made consistent and a copy of the data made. Be aware that the Global Copy source volumes are not active in the Global Mirror session until they have completed the first pass of their initial copy.

4. **Create FlashCopy relationships.** Create the relationships at the recovery site between the Global Copy secondary volumes and the FlashCopy target volumes. Ensure that you resynchronize all volumes that will be in the FlashCopy relationships. See "Creating a FlashCopy volume pair" for the task procedure. When creating the FlashCopy relationship, select the following options:

**Enable Change Recording**

Select this option to activate change recording on the volume pair participating in a FlashCopy relationship.

**Note:** The Persistent FlashCopy option is the default with this option.

**Inhibit writes to target volume**

Select this option to ensure that updates cannot be made to the target volume. This ensures data consistency on the target volume. If you select the Inhibit writes to target option, the change recording feature is not active on the target volume, only on the source volume.

**Note:** Do not select the **Initiate background copy** option so that data will only be copied from the source volume to the target volume if a track on the source volume is modified.

5. **Create your Global Mirror session.** See "Creating a new Global Mirror session" for the task procedure.
6. **Add volumes to your Global Mirror session.** See "Adding volumes to a Global Mirror session" for the task procedure.

After the setup is done, the following process describes how data consistency is created between the two sites:

1. Consistency groups of volumes are created at the local site.
2. Incremental copies of data are sent to the recovery site.
3. FlashCopy operation is performed at the recovery site.
4. Global Copy operations are resumed (only out-of-sync data is copied).
5. Steps are repeated, which are determined by the defined time periods.

## Using Global Mirror for a planned failover and failback

Use this process to manage a planned failover involving two sites, which are referred to as local (or production) and remote (or recovery) sites, using Global Mirror.

During a planned outage, all storage units, servers, and networks are functional. Last and current states of all components are well-defined.

Examples for planned outages are system maintenance, disaster recovery tests, and training. The objective of the failover and failback procedures is to continue with consistent and current data (without any loss of transactions).

The procedure that is presented in this section makes the following assumptions:

- All storage units (local and remote) are functional and accessible.
- Applications are updating the primary volumes that are located at the local site.
- Paths are created from the local to the remote site.
- You have already created an environment to manage a planned outage.

Perform the following steps using the DS Storage Manager to begin the planned failover. You can also use the DS CLI.

**Note:** For details on individual tasks, see the related topics section.

1. Quiesce host I/O that is updating the primary volumes that is located in the local site.
2. Check the status of Global Mirror to ensure that Global Mirror is in a good state and that a consistency group has formed successfully at the desired interval.
3. Wait until two consistency groups form successfully.
4. Pause Global Mirror and suspend the Global Copy pairs (A->B).
5. Create paths from the B volumes to the A volumes.
6. Perform recovery failover on the B volumes (B->A).
7. Perform the fast reverse restore process from the C volumes to the B volumes, selecting the **Initiate background copy** option.

**Notes:**

- a. When you initiate the fast reverse restore process, Volume C becomes unusable.
  - b. There must be *no* I/O allowed to the B or C volumes during the fast reverse restore process.
8. Wait for the background copy to complete before continuing to the next step. The C to B FlashCopy relationship ends when the background copy completes.
  9. Initiate the FlashCopy from Volume B back to Volume C. Ensure that you also select the **Enable Change Recording** and **Inhibit writes to target volume** options. This creates a backup copy of the consistency group before applications begin to update the B volumes.
  10. Start the host I/O at the remote site on the B volumes. You will remain in production on the remote site in this configuration until you are ready to return production to the local site.
  11. When you are ready to return production to the local site, perform recovery failback (B->A) to resynchronize the A volumes. The application at the remote site remains active.
  12. After the resynchronize process has completed its first pass, quiesce the applications at your remote site so that the resynchronization can complete.



13. When the resynchronization has completed (no out-of-sync tracks), perform recovery failover and failback with Global Copy on Volume A to re-create the Volume A to Volume B Metro Mirror relationship.
14. Start the host I/O at the local site on the A volumes.
15. Resume the Global Mirror process.

## Using Global Mirror for an unplanned failover and failback

Use this process to manage an unplanned failover and failback that involves two sites which are referred to as the local (or production) and remote (or recovery) sites.

Global Mirror provides two-site extended distance remote copy disaster recovery. When a disaster occurs at the local site, you must initiate the failover and failback recovery of consistent data on the remote site. Host activity can resume on the local site when the host recovers but not before a consistent set of data is copied to all primary volumes on the local site.

With Global Mirror, the data that the host writes to the storage unit at the local site is asynchronously shadowed to the storage unit at the remote site. A consistent copy of the data is then automatically maintained on the storage unit at the remote site.

Use of Global Mirror does not guarantee against data loss. During a disaster, data can only be restored to the last known consistent increment that was created. This means that data that is written to the primary site and is waiting to be transferred to the secondary site is lost whenever the two storage units can no longer communicate.

The following considerations apply when you use the Global Mirror recovery process:

- The Global Mirror master might still be running at the local site, especially if the disaster at the local site is a rolling disaster in which not all components fail simultaneously.
- The consistent copy at the remote site is *not* the secondary volume, but it is the FlashCopy target whose source is the secondary volume.
- Formation of a consistency group might have been in progress at the time of the failure.
- You can speed up recovery processing by choosing the **Fast Reverse** restore process that is explained later in this section.

Perform the following steps to use Global Mirror for an unplanned failover using the DS Storage Manager. You can also use the DS CLI for these steps.

**Note:** For details on individual tasks, see the related topics section.

1. Enter the Global Mirror session at the local site.

**Note:** Wait until the master storage unit completes the termination processing or enters the fatal state before continuing with the next step. Of course, this might not be possible if the local site has completely failed. In that situation, proceed to the next step without waiting.

2. Issue a recovery failover request on the Global Copy volumes pair to force a stop of the Volume A to Volume B extended distance relationship and create a Volume B to Volume A Global Copy relationship.



**Note:** All B volumes must successfully process the recovery failover request before you can move to the next step.

3. Look at the session properties for Volume B and Volume C to ascertain the state of the consistency group between the B volume and C volume. You are looking primarily at the FlashCopy relationships and your analysis determines your next step in the recovery process. Act on your analysis as follows:
  - a. FlashCopy relationships are nonrestorable and all the sequence numbers are equal. No action to the consistency group is necessary.
  - b. FlashCopy relationships are restorable and all the sequence numbers are equal. Issue the FlashCopy **Discard changes** command to all the FlashCopy relationships in the consistency group.
  - c. All the FlashCopy sequence numbers are equal and at least one of the FlashCopy relationships is nonrestorable. Issue a FlashCopy **Commit changes** command to all the FlashCopy relationships in the consistency group that are restorable.
  - d. You have a mixed list of FlashCopy relationships; some are restorable and some are nonrestorable. The sequence numbers of the relationships that are restorable all have the same sequence number. The sequence numbers of the relationships that are nonrestorable are all equal, but they have a different number from the sequence number of those that are restorable. Issue a FlashCopy **Commit changes** command to all the FlashCopy relationships in the consistency group that are restorable.
  - e. You have a mixed list of FlashCopy relationships; some are restorable and some are nonrestorable. The sequence numbers are not the same within each type of relationship. The recovery plan cannot continue. The Global Mirror process has been corrupted. If the Global Mirror process has been corrupted you must recover your data using your last good backup.

**Note:** When the state of all the FlashCopy relationships are known, you might want to initiate a tape backup of Volume C.

4. Issue the fast reverse restore process from the C volumes to the B volumes, selecting the **Initiate background copy** option.

**Note:**

- a. When you initiate the fast reverse restore process, Volume C becomes unusable.
  - b. There must be *no* I/O allowed to the B or C volumes during the fast reverse restore process.
  - c. If you do not want to use the fast reverse restore process, use the Recovering from a disaster without using fast reverse restore procedure instead of this step.
5. Wait for the background copy to complete before continuing to the next step. The C to B FlashCopy relationship ends when the background copy completes.
6. Initiate the FlashCopy from Volume B back to Volume C. Ensure that you also select the **Enable Change Recording** and **Inhibit writes to target volume** options. This creates a backup copy of the consistency group before applications begin to update the B volumes.
7. Start the host I/O at the remote site on the B volumes. You will remain in production on the remote site in this configuration until you are ready to return production to the local site.

8. When you are ready to return production to the local site, perform the recovery failback (B -> A) to resynchronize the A volumes. The application at the remote site remains active.
9. After the resynchronize process has completed its first pass, quiesce the applications at your remote site so that the resynchronization can complete.
10. When the resynchronization has completed (no out-of-sync tracks), perform recovery failover and failback with Global Copy on Volume A to re-create the Volume A to Volume B Metro Mirror relationship.
11. Start the host I/O at the local site on the A volumes.
12. Resume the Global Mirror process.

## Recovering with and without the fast reverse restore feature

The fast reverse restore feature, when used with Global Mirror functions, allows a FlashCopy relationship to be reversed without waiting for the background copy of the original FlashCopy to complete.

You can reverse a FlashCopy relationship that has change recording enabled and restore the tracks from the target volume to the source volume. However, the background copy process of the original FlashCopy must complete before you can reverse the order of the FlashCopy relationship to its original source and target relationship.

With the fast reverse restore feature, you can reverse a FlashCopy relationship that has change recording enabled without waiting for the background copy of the original FlashCopy to complete. It is not required that you use the fast reverse restore feature during disaster recovery operations; however, the time to recover could be significantly longer if you do not.

### Recovering from a disaster using the fast reverse process

Use this process for the Global Mirror fast reverse restore feature, which allows a FlashCopy relationship to be reversed without waiting for the background copy of the original FlashCopy to complete.

You can use the fast reverse restore feature of FlashCopy to speed up the recovery process. Using the DS Storage Manager, you make the following selections:

1. Create a FlashCopy target on the existing Metro Mirror source by choosing the FlashCopy target and a source Metro Mirror .
2. Select the **Fast Reverse** restore option.

FlashCopy copies the partial data that had been copied to Volume C before the disaster occurred to Volume B. This creates a consistent copy on Volume B. Processing also provides a background copy of the consistent group for Volume B.

**Note:** When you initiate the fast reverse restore process, Volume C becomes unusable.

### Recovering from a disaster without the fast reverse process

Use this procedure in place of the fast reverse process. The fast reverse process allows a FlashCopy relationship to be reversed without waiting for the background copy of the original FlashCopy to complete.

It is not necessary to use the fast reverse feature of FlashCopy for your recovery. However, the wait for recovery is significantly longer if you do not.

1. Initiate a background copy on the Volume B to Volume C FlashCopy relationship.
2. Wait for the background copy to complete. Observe the out-of-sync tracks.
3. Select the **Reverse relationship** and **Initiate background copy** options on the Volume B to Volume C FlashCopy relationship. This reverses the FlashCopy relationship such that Volume C is the source and Volume B is the target.



---

## Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:**

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATIONS "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been

estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

---

## Accessibility

Accessibility features provide users who have disabilities with the ability to successfully access information and use technology.

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully.

### Features

These are the major accessibility features in the IBM System Storage<sup>™</sup> DS6000 information:

- You can use screen-reader software and a digital speech synthesizer to hear what is displayed on the screen. IBM Home Page Reader version 3.0 has been tested.
- You can operate features using the keyboard instead of the mouse.

### Navigating by keyboard

You can use keys or key combinations to perform operations and initiate menu actions that can also be done through mouse actions. You can navigate the IBM System Storage DS6000 information from the keyboard by using the shortcut keys for your browser or Home Page Reader. See your browser Help for a list of shortcut keys that it supports. See the following Web site for a list of shortcut keys supported by Home Page Reader: [http://www-306.ibm.com/able/solution\\_offerings/keyshort.html](http://www-306.ibm.com/able/solution_offerings/keyshort.html)

### Accessing the publications

You can find HTML versions of the IBM System Storage DS6000 information at the following Web site: <http://www.ehone.ibm.com/public/applications/publications/cgibin/pbi.cgi>

You can access the information using IBM Home Page Reader 3.0.

---

## Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

- AIX
- DB2
- DFSMS/MVS
- DFSMS/VM
- DS4000
- DS6000
- DS8000
- e (logo)
- Enterprise Storage Server
- ES/9000
- ESCON
- FICON
- FlashCopy
- Graphically Dispersed Parallel Sysplex
- HACMP
- i5/OS
- IBM
- IntelliStation
- MVS/ESA
- Netfinity
- NetVista
- Operating System/400
- OS/400
- RS/6000
- S/390
- Seascape
- SNAP/SHOT
- SP
- System/390
- System p5
- System Storage
- Versatile Storage Server
- Virtualization Engine
- VSE/ESA
- z/Architecture
- z/OS
- z/VM
- zSeries

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

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

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, and service names may be trademarks or service marks of others.

---

## Terms and conditions

Permissions for the use of these publications is granted subject to the following terms and conditions.

**Personal Use:** You may reproduce these Publications for your personal, non commercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these Publications, or any portion thereof, without the express consent of IBM.

**Commercial Use:** You may reproduce, distribute and display these Publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these Publications, or reproduce, distribute or display these Publications or any portion thereof outside your enterprise, without the express consent of IBM.

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the Publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the Publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED “AS-IS” AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

---

## Electronic emission notices

This section contains the electronic emission notices or statements for the United States and other countries.



## **Federal Communications Commission (FCC) statement**

This equipment has been tested and complies with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, might cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used to meet FCC emission limits. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors, or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the users authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device might not cause harmful interference, and (2) this device must accept any interference received, including interference that might cause undesired operation.

## **Industry Canada compliance statement**

This Class A digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe A est conform à la norme NMB-003 du Canada.

## **European community compliance statement**

This product is in conformity with the protection requirements of EC Council Directive 89/336/EEC on the approximation of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a nonrecommended modification of the product, including the fitting of non-IBM option cards.

### **Germany only**

**Zulassungsbescheinigung laut Gesetz ueber die elektromagnetische Vertraeglichkeit von Geraeten (EMVG) vom 30. August 1995.**

Dieses Geraet ist berechtigt, in Uebereinstimmung mit dem deutschen EMVG das EG-Konformitaetszeichen - CE - zu fuehren.

Der Aussteller der Konformitaetserklaeung ist die IBM Deutschland.

Informationen in Hinsicht EMVG Paragraph 3 Abs. (2) 2:

Das Geraet erfuehlt die Schutzanforderungen nach EN 50082-1 und EN 55022 Klasse A.

EN 55022 Klasse A Geraete beduerfen folgender Hinweise:

Nach dem EMVG:

"Geraete duerfen an Orten, fuer die sie nicht ausreichend entstoert sind, nur mit besonderer Genehmigung des Bundesministeriums fuer Post und Telekommunikation oder des Bundesamtes fuer Post und Telekommunikation

betrieben werden. Die Genehmigung wird erteilt, wenn keine elektromagnetischen Störungen zu erwarten sind." (Auszug aus dem EMVG, Paragraph 3, Abs.4)

Dieses Genehmigungsverfahren ist nach Paragraph 9 EMVG in Verbindung mit der entsprechenden Kostenverordnung (Amtsblatt 14/93) kostenpflichtig.

Nach der EN 55022:

"Dies ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funkstörungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Massnahmen durchzuführen und dafür aufzukommen."

Anmerkung:

Um die Einhaltung des EMVG sicherzustellen, sind die Geräte wie in den Handbüchern angegeben zu installieren und zu betreiben.

## **Japanese Voluntary Control Council for Interference (VCCI) class A statement**

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラス A 情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

## **Korean Ministry of Information and Communication (MIC) statement**

Please note that this device has been certified for business use with regard to electromagnetic interference. If you find this is not suitable for your use, you may exchange it for one of residential use.

## **Taiwan class A compliance statement**

**警告使用者:**

這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

VS07171L

---

# Index

## A

accessibility 18

## C

Copy Services  
disaster recovery 3

## D

disaster recovery  
Copy Services 3  
failback 4  
failover 4  
Global Mirror 9  
Global Mirror planned scenario 11  
Global Mirror scenario for  
unplanned 12  
two-site solution 9  
using the fast reverse restore  
feature 14  
without using the fast reverse  
feature 14

## F

failover  
site switch 4, 7  
failover and fallback 3  
fast reverse  
Global Mirror 14  
fast reverse restore  
advantages of using 14

## G

Global Mirror 3  
disaster recovery for unplanned 12  
disaster recovery without fast  
reverse 14  
fast reverse restore 14  
planned failover and fallback  
scenario 11  
setting up 9  
unplanned failover and fallback 12

## K

keyboards  
accessibility features 18

## L

legal  
terms and conditions 20

## P

performing  
site switch  
failover 4, 7

## R

recovering 9  
with and without the fast reverse  
restore feature 14

## S

setting up  
Global Mirror 9

## T

Trademarks 19

## U

using  
fast reverse restore 14







Printed in USA