IBM

IBM DS8870 Multiple Target Peer-to-Peer Remote Copy

- Enhanced resiliency functions
- Architecture and design
- Implementation and use

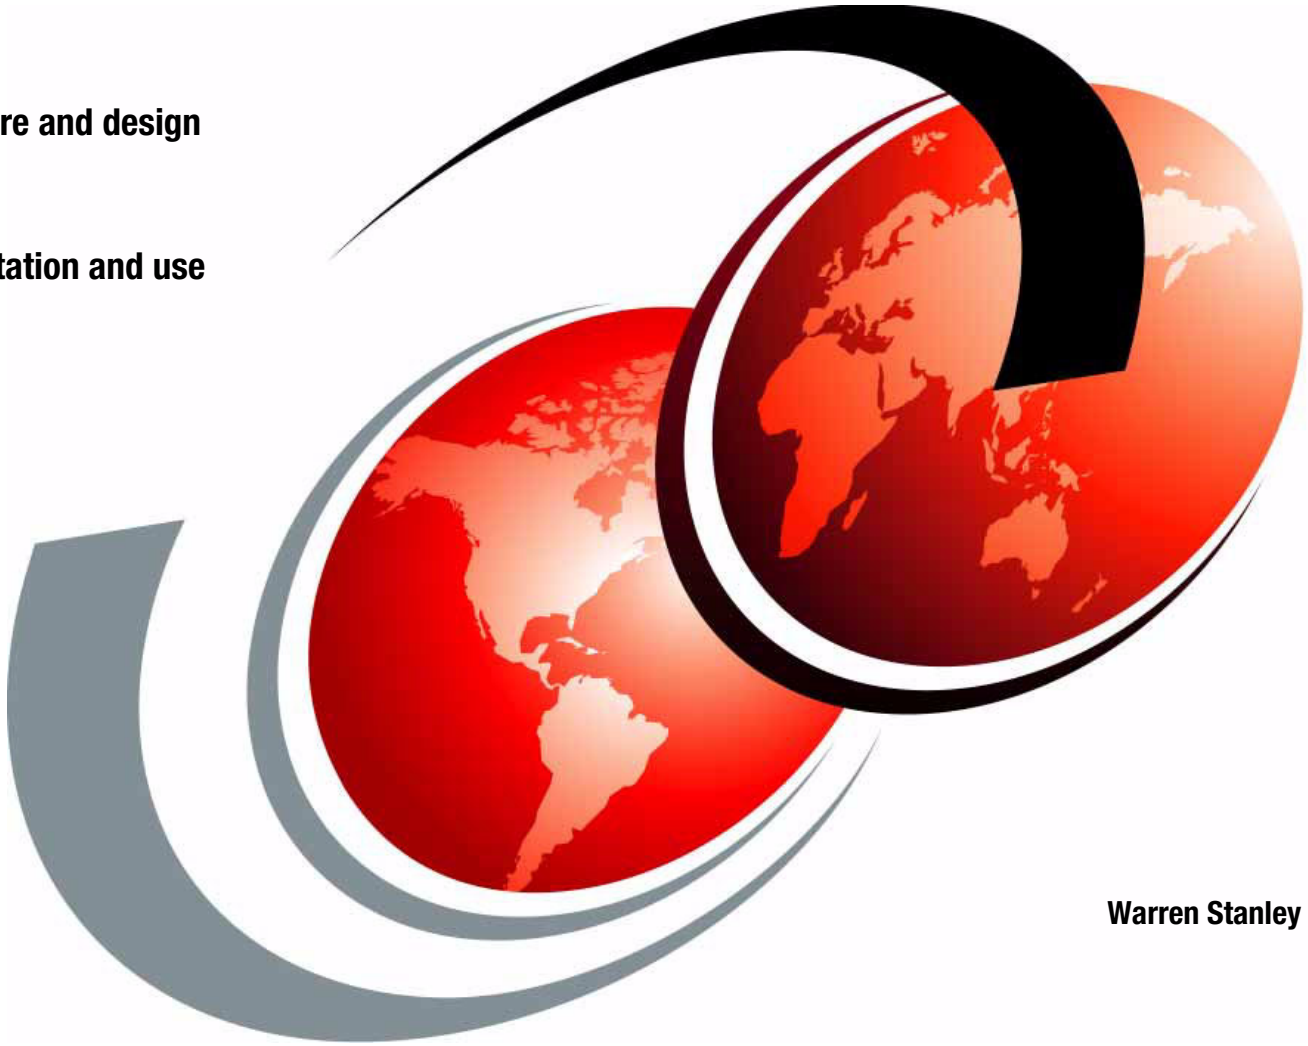

Warren Stanley

Redpaper

ibm.com/redbooks

International Technical Support Organization

IBM DS8870 MultipleTarget PPRC

February 2015

REDP-5151-00

Note: Before using this information and the product it supports, read the information in “Notices” on page vii.

First Edition (February 2015)

This edition applies to IBM System Storage® DS8000 series with DS8000 LMC 7.7.40.xx.xx (bundle version 87.40.xxx.xx). This code level is supported only with the IBM DS8870 hardware.

Contents

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 grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of g, 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 PUBLICATION “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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites 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 measurements 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.

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.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

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

DS8000®
FlashCopy®
GDPS®
HyperSwap®
IBM®
Redbooks®
Redpaper™
Redbooks (logo)  ®
System Storage®
System z®
Tivoli®
z/OS®
z/VM®

The following terms are trademarks of other companies:

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

Preface

Multiple Target Peer-to-Peer Remote Copy (MT-PPRC) was introduced with IBM® System Storage® DS8870 Version 7 Release 4 (License Machine Code 7.7.40.xx.xx).

MT-PPRC provides the capability to have two PPRC relationships on a single primary volume. With this enhancement, there can be another target available to provide additional data protection to act as a backup if a disaster occurs.

This IBM Redpaper™ publication provides practical information about the characteristics and functions of the IBM DS8000® Multiple Target Peer-to-Peer Remote Copy (MT-PPRC). The paper includes scenarios that illustrate the implementation and use of different topologies, as well as enhanced migration scenarios.

Authors

This paper was produced by a team of specialists from around the world working at the IBM International Technical Support Organization, San Jose Center.

Warren Stanley is a Senior Technical Staff Member and Master Inventor in the IBM DS8000 Copy Services microcode development organization. He is the microcode architect for Multiple Target PPRC, in addition to being a Copy Services architect and developer for PPRC replication and IBM z/OS® Global Mirror (XRC). Warren holds a degree in Electrical and Computer Engineering from California State Polytechnic University in Pomona, California, and he joined IBM in 1982.

Other contributors

Bert Dufrasne, Alexander Warmuth, Susan Candelaria, Craig Gordon, Hank Sautter, and Brian Sherman
IBM

Now you can become a published author, too

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time. Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us.

We want our papers to be as helpful as possible. Send us your comments about this paper or other IBM Redbooks® publications in one of the following ways:

- Use the online **Contact us** review Redbooks form::

  ibm.com/redbooks

- Send your comments by email:

  redbooks@us.ibm.com

- Mail your comments:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

1

Introducing Multiple Target Peer-to-Peer Remote Copy

This chapter provides an introduction to Multiple Target Peer-to-Peer Remote Copy (Multiple Target PPRC). It covers these topics:

- Multiple Target PPRC overview
- Use cases
- Review of Remote Mirror and Copy
- Multiple Target PPRC enhancement
- Licensing
- Terminology
- Table 1-1 on page 21

1.1 Multiple Target PPRC overview

IBM Multiple Target Peer-to-Peer Remote Copy (Multiple Target PPRC) enhances a multisite disaster recovery environment by providing the capability to have two PPRC relationships on a single primary volume. This adds data protection because there is an additional remote site.

In other words, with Multiple Target PPRC, the same primary volume can now have more than one target, which enables data to be mirrored from a single primary site to two different target sites. Before the availability of Multiple Target PPRC, it was possible for a primary volume to mirror its data to only one secondary volume.

Multiple Target PPRC provides the following enhancements:

- Mirrors data from a single local primary site to two remote secondary sites
- Increases capability and flexibility for disaster recovery solutions by using:
  - Synchronous replication
  - Asynchronous replication
  - Combination of synchronous replication and asynchronous replication configurations
- Improves a cascaded Metro/Global Mirror (MGM) configuration and simplifies recovery procedures

Note: Multiple Target PPRC is available in both open systems and IBM System z® environments.

Figure 1-1 shows a general Multiple Target PPRC topology where a single primary site is replicated to two secondary sites. Host I/O is directed to the primary site labeled H1, and Multiple Target PPRC mirrors the updated data to the two secondary sites, which are labeled H2 and H3.

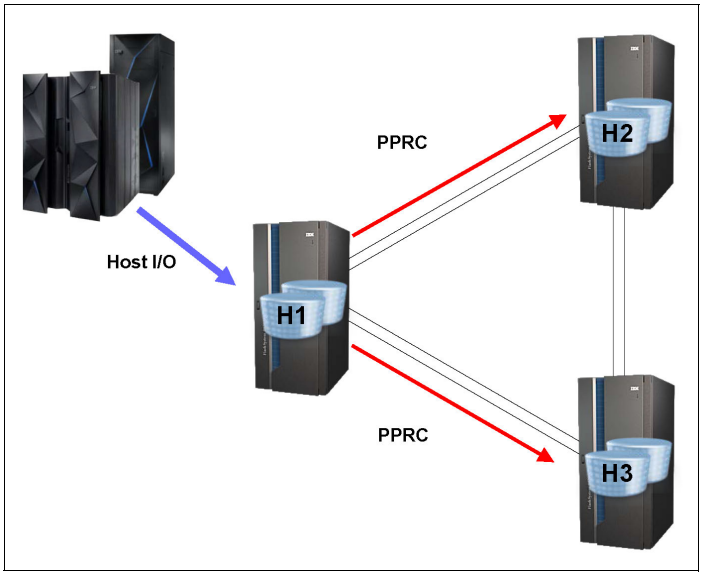

Figure 1-1 Multiple Target PPRC

Note: A primary volume can now have any combination of two Metro Mirror, Global Copy or Global Mirror relationships, with a restriction that a primary volume can belong to only one Global Mirror session at a time.

When both relationships are synchronous Metro Mirror, the transfers of updated data to the remote sites are performed in parallel to minimize the impact on host response time in comparison to one Metro Mirror relationship.

1.2 Use cases

Many companies and businesses require their applications to be continuously available and cannot tolerate any service interruption. A loss of a disaster recovery site is often considered to be a severe impact on their business. If their local production site fails, swapping to a Metro Mirror target allows applications to continue running. However, without another target to act as a backup for disaster protection, business applications are left unprotected from a subsequent failure. Multiple Target PPRC addresses these needs by allowing an additional replication target.

Multiple Target PPRC also complements the use of cascaded three-site mirroring topologies, such as MGM, by adding flexibility and simplifying certain recovery scenarios.

1.3 Review of Remote Mirror and Copy

The use of Multiple Target PPRC requires an understanding of the replication technologies of the IBM System Storage DS8000 series. The sections that follow give brief descriptions of these solutions. See “Related publications” on page 123 for sources for more detailed information.

1.3.1 Metro Mirror

Metro Mirror is a synchronous mirror of logical volumes between two DS8000 storage systems. With a synchronous mirror, each write to the local storage system is written to both the local and remote sites before the I/O is considered to be complete. This provides the potential for a zero Recovery Point Objective (RPO).

The host write response time is extended by the amount of time that is required to transfer the data to the remote location, which limits the distance at which a synchronous mirror is a practical solution. Metro Mirror can be used at distances of up to 300 km (or longer by contacting IBM), provided that application performance requirements can be met. Metro Mirror has been implemented in production environments at over 800 km. However, with functionality such as IBM HyperSwap®, there are additional availability benefits when Metro Mirror is deployed at shorter distances. HyperSwap is typically used at distances of less than 50 km, but distances up to approximately 100 km are sometimes feasible.

1.3.2 Global Copy

Global Copy is an asynchronous copy of data, which makes it suitable for essentially unlimited distances. With an asynchronous copy, a write to the local storage system is considered complete after being completed at the local site. The copy of the data to the remote location occurs later, asynchronously, to the host I/O. Because the host write operation does not wait for the data to be sent to the remote location, there is little impact on the write response time.

With Global Copy, data is not necessarily copied to the remote site in the same order that it was written to the local storage system. Therefore, the data at the target site is not logically consistent. Additional steps must be taken to obtain a usable copy of data at the target site, such as quiescing writes for a period of time and waiting for all of the data to flow to the remote site, at which time the data will become consistent.

A common use of Global Copy is for migration purposes over long distances. Another common and important use is as the data transfer mechanism for Global Mirror, which does ensure data consistency at the remote location.

1.3.3 Global Mirror

Global Mirror is an asynchronous mirror of data with logical data consistency. Global Mirror integrates Global Copy and IBM FlashCopy® technologies into a long-distance disaster recovery solution.

Global Copy is used as the transfer mechanism to copy data to the remote site. Global Mirror periodically creates a consistency group across all participating storage systems by using a combination of bitmap files and a small cache sidefile (for colliding writes) to track changed data for each consistency group. Data written before a specific time is copied to the remote site, and data written after that time is held at the local site. When all data for the consistency group has been copied to the remote site, a FlashCopy of the remote site's data is made. This process then repeats and another consistency group is formed. In a well-configured system, consistency groups can be formed every few seconds, resulting in an RPO of a few seconds or less.

1.3.4 z/OS Global Mirror

Like Global Mirror, z/OS Global Mirror (z/GM) is asynchronous mirroring with consistency. However, unlike Global Mirror, z/GM uses an IBM DFSMS/MVS component called the System Data Mover (SDM) to transfer the data and create consistency groups at the remote location. As its name implies, z/GM is available only for z/OS environments.

As writes occur to the local primary volumes, the DS8000 time stamps the data and buffers that data in a cache sidefile on the primary storage system. The SDM manages the process of reading the updated data, forming consistency groups, and writing the data to the remote volumes.

Figure 1-2 on page 17 shows a z/GM topology where the SDM running on a z/OS system at the remote site is reading updated data from the primary H1 volumes, forming consistency groups, and writing the data to the remote H2 volumes.

z/OS Global Mirror is capable of mirroring data for large z/OS environments while maintaining a low RPO.

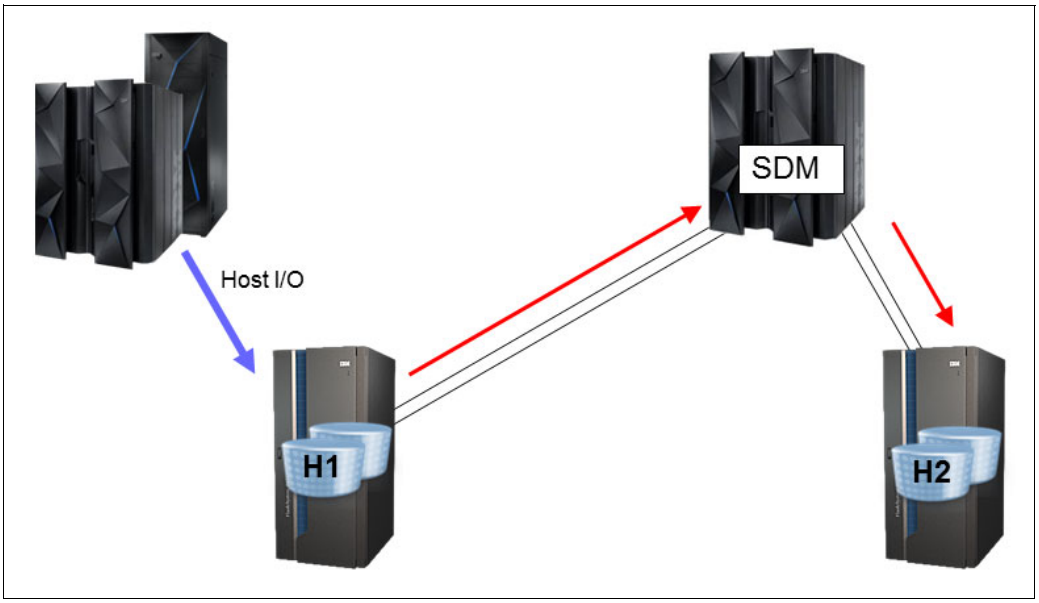

Figure 1-2 z/OS Global Mirror

1.3.5 Metro Global Mirror and Metro z/Global Mirror

Metro Global Mirror (MGM) is a three-site topology containing both Metro Mirror and Global Mirror. This configuration gives both a synchronous mirror for high availability with RPO=0 and a long-distance disaster recovery solution.

Figure 1-3 shows a cascaded MGM topology where there is Metro Mirror from H1 to H2 and Global Mirror from H2 to H3. The J3 volumes are the FlashCopy journal target volumes which contain the data for the last consistency point created by Global Mirror. As described in section 2.3, “Metro Mirror and Global Mirror” on page 25, with Multiple Target PPRC it is now possible to also have an MGM topology where the Global Mirror relationship is from H1 to H3 rather than from H2 to H3.

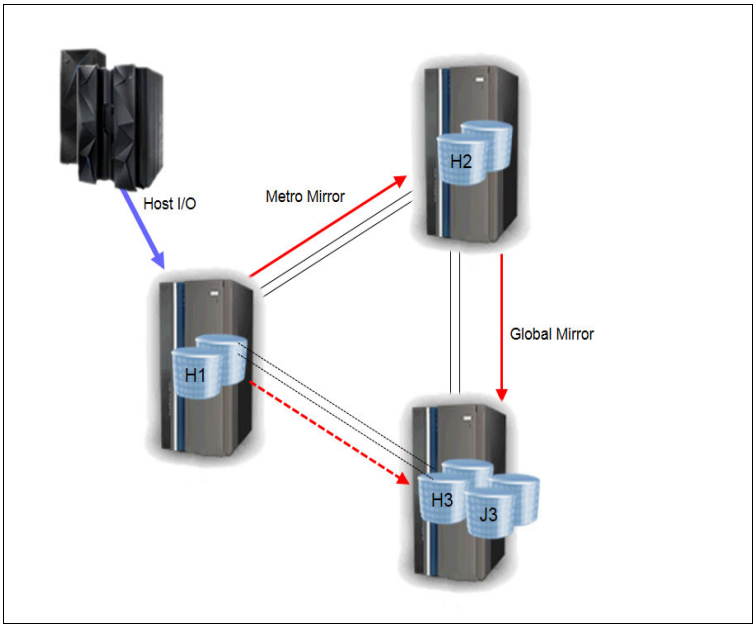

Figure 1-3 Metro Global Mirror

Similar to MGM, a Metro z/Global Mirror (MzGM) topology is a three-site solution with Metro Mirror and zGM as shown in Figure 1-4. This provides the high availability of synchronous Metro Mirror from H1 to H3 with the disaster recovery capability of z/OS Global Mirror from H1 to H2. There can also be connectivity from H3 to the SDM to provide an Incremental Resynchronization capability from H3 to H2 in the event of an outage at the H1 site.

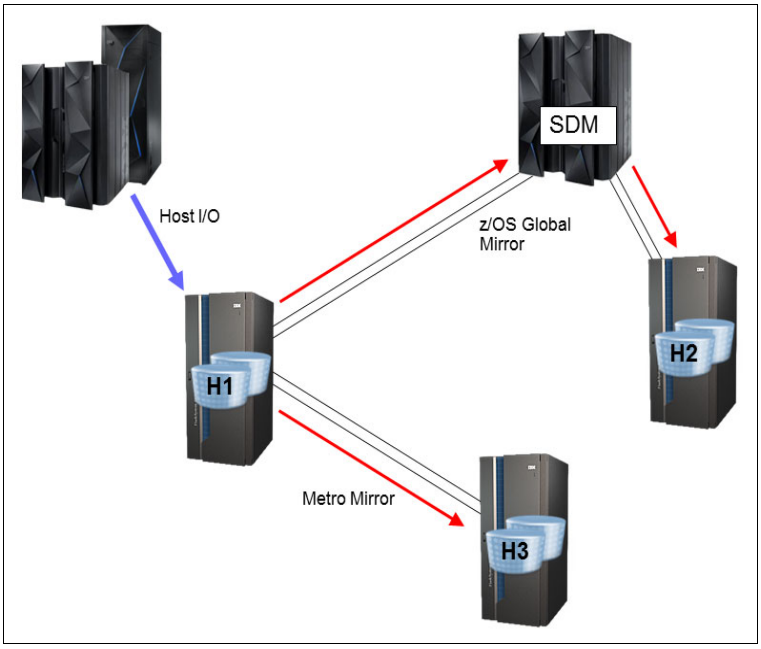

Figure 1-4 Metro z/OS Global Mirror

1.4 Multiple Target PPRC enhancement

Multiple Target PPRC extends the basic Metro Mirror, Global Copy and Global Mirror solutions. With two PPRC relationships on a volume, additional topologies and flexibility are possible.

In the event of a primary site failure and recovery at the first recovery site, an additional recovery site is still available to provide protection from a second site failure.

Current cascaded solutions such as MGM and Mz/GM provide three-site solutions and remain fully supported. Multiple Target PPRC adds the capability for three-site synchronous and three-site asynchronous topologies.

1.4.1 Requirements and restrictions

All DS8870 storage systems in a Multiple Target PPRC environment are required to have an LMC 7.7.40.xx.xx (bundle version 87.40.xxx.xx) or later to support Multiple Target PPRC. This requirement applies to all secondary DS8870 storage systems in addition to the primary DS8870.

Important: A Multiple Target PPRC environment requires that all DS8870 storage systems have Multiple Target PPRC support.

Take Note: Managing a Multiple Target PPRC environment is not supported through the DS GUI. You must use the DS CLI or other management software such as GDPS or Copy Services Manager, as mentioned in 1.5, “Software considerations”.

1.5 Software considerations

A Multiple Target PPRC environment includes, but is not limited to, the following software requirements:

- Activate Multiple Target PPRC support on the DS8870 storage system by using indicator 0745 and feature code 7025. DFSMS software checks the indicator (0745) that is set by the DEVSUPxx PARMLIB member.
- Management software requires that all systems be upgraded to support the function. Otherwise, the function is in single-target mode. After systems are upgraded, update your DEVSUPxx member of PARMLIB to indicate that Multiple Target PPRC mode is supported by using a new keyword: `PPRCMT`.
- Support for IBM System z is provided through SPEs z/OS (including DFSMS), ICKDSF, IBM z/VM® virtualization platform, and management software (IBM GDPS® and IBM Copy Services Manager, formerly known as Tivoli® Storage Productivity Center for Replication). Software support is provided in z/OS Version 2.1 and later through the SDM component, ICKDSF Release 17, and z/VM Release 6.3.
- Updated queries support multiple secondaries for a Metro Mirror primary volume. Be aware that if ANTRQST PQUERY is used, it is required to convert to a new FORMAT mapping macro. A keyword in the DEVSUPxx member of PARMLIB is provided to enable the new function. If it is not set, it retains existing recovery states for cascading PPRC environments. Also, TSO, ICKDSF, and DS command-line interface (CLI) queries will also show multiple secondary volumes, and any automation based on these needs to be checked and verified.

1.6 Licensing

To use the Multiple Target PPRC function, licensed function indicator 0745 and feature code 7025 are required and must be applied to all DS8870 storage systems that are part of the Multiple Target PPRC configuration. This includes the primary and all secondary storage systems. In addition, Metro Mirror, Global Mirror, Metro/Global Mirror, z/OS Global Mirror, or FlashCopy licenses are required for all DS8870s that use those functions within the Multiple Target PPRC configuration.

1.7 Terminology

The term *site* is used to designate the data that is to be replicated. A site is a specific set of volumes. This can be a group of volumes within a single storage system or several volumes across multiple storage systems that are being replicated.

These sites are labeled as H1, H2, H3, and so on. To simplify the procedures, the descriptions often refer to the sites as though they are a single volume, but it is important to remember that they are actually groups of volumes.

For example, Figure 1-5 on page 20 shows a sample configuration where the different sites are labeled.

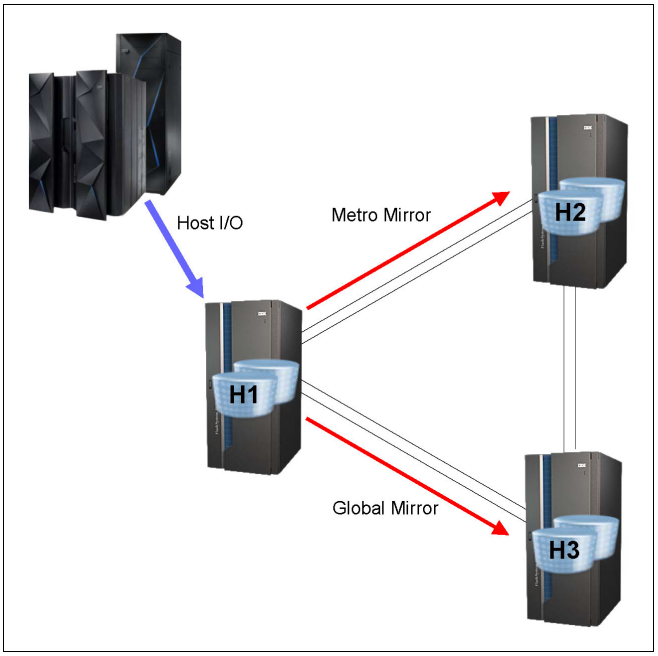

Figure 1-5 Sample configuration for describing terminology

The local site is labeled H1. Site H1 can be a single volume, a subset of volumes within a storage system, all volumes within a storage system, or a group of volumes across multiple storage systems. All volumes referred to by “site H1” are related and are used in a similar manner.

Site H2 in Figure 1-5 is a different set of volumes, which can be within a single storage system or spread across multiple storage systems. This H2 site is often at a different physical location from H1, but it is not required to be. In some situations, such as initial testing when limited hardware is available, it is even possible that H1 and H2 are different sets of volumes within a single storage system.

Using this terminology, it is possible to refer to operations such as “establish H1:H2” as a type of shorthand. Because H1 and H2 refer to groups of related volumes, this means to establish a pair from each volume at site H1 to its corresponding volume at H2.

PPRC paths are created at the LSS or LCU level rather than the volume level, so an operation such as “establish paths H1:H2” means to establish paths from each LSS or LCU at site H1 to its corresponding LSS or LCU at site H2. (The number of LSSes or LCUs supported at the secondary is now up to 16 with DS8870 LMC 7.7.40.xx.xx.)

The PPRC failover command deserves some special consideration in the terminology used. Consider a configuration of Metro Mirror from H1 to H2 as shown in Figure 1-6 on page 21:

- Site H1 contains PPRC primary volumes that are being mirrored to PPRC secondary volumes at H2.
- A PPRC failover command is used to change the H2 volumes to become primary volumes whose corresponding secondary volumes are then H1. This is described by the term “failover H2:H1,” which indicates that the command is issued to the H2 and H1are the secondary volumes specified in the command.
- Global Mirror journal volumes are designated as J1, J2, and so on.

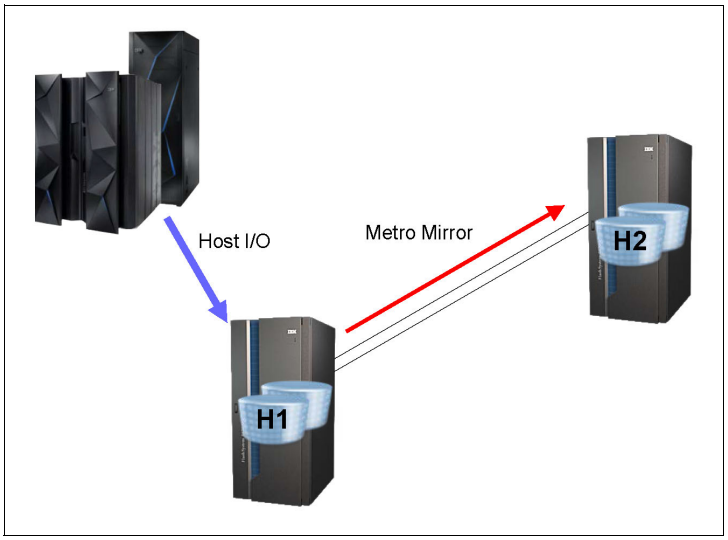

Figure 1-6 Sample Metro Mirror

1.8 Commands

There are several different interfaces available to manage a PPRC environment. The examples in this paper use the DS CLI (the DS GUI cannot be used for managing MT-PPRC). The following table lists equivalent TSO and ICKDSF commands for some operations.

Table 1-1 PPRC commands

Description	DS CLI	TSO	ICKDSF
Create paths	mkpprcpath	CESTPATH	PPRCOPY ESTPATH
Remove paths	rmpprcpath	CDELPATH	PPRCOPY
Create pair	mkpprc	CESTPAIR	PPRCOPY ESTPAIR
Remove pair	rmpprc	CDELPAIR	PPRCOPY DELPAIR
Suspend pair	pausepprc	CSUSPEND	PPRCOPY SUSPEND
Resume pair	resumepprc	CESTPAIR MODE(RESYNC)	PPRCOPY ESTPAIR MODE(RESYNC)
Failover	failoverpprc	CESTPAIR ACTION(FAILOVER)	PPRCOPY ESTPAIR FAILOVER
Failback	failbackpprc	CESTPAIR ACTION(FAILBACK)	PPRCOPY ESTPAIR FAILBACK
Freeze	freezepprc	CGROUP FREEZE	PPRCOPY FREEZE
Unfreeze	unfreezepprc	CGROUP RUN	PPRCOPY RUN
Set PPRC Characteristics	chpprc	PSETCHAR	PPRCOPY SETCHARACTERISTICS
Create FlashCopy	mkflash	FCESTABL	FLASHCPY ESTABLISH

2

Topologies

This chapter describes the different topologies that are possible using Multiple Target PPRC.

Multiple Target PPRC allows for many different combinations of Metro Mirror, Global Copy, and Global Mirror. It is also possible to use Multiple Target PPRC in combination with IBM z/OS Global Mirror.

The descriptions cover these topologies:

- Two Metro Mirror relationships
- Metro Mirror and Global Mirror
- Metro Mirror and Global Copy
- Global Copy plus Global Mirror
- Metro Global Mirror plus Metro Mirror
- z/OS Global Mirror and two Metro Mirror pairs

2.1 Multiple Target PPRC topologies

Figure 2-1 shows a typical diagram that is used for describing the different topologies and configurations throughout this paper.

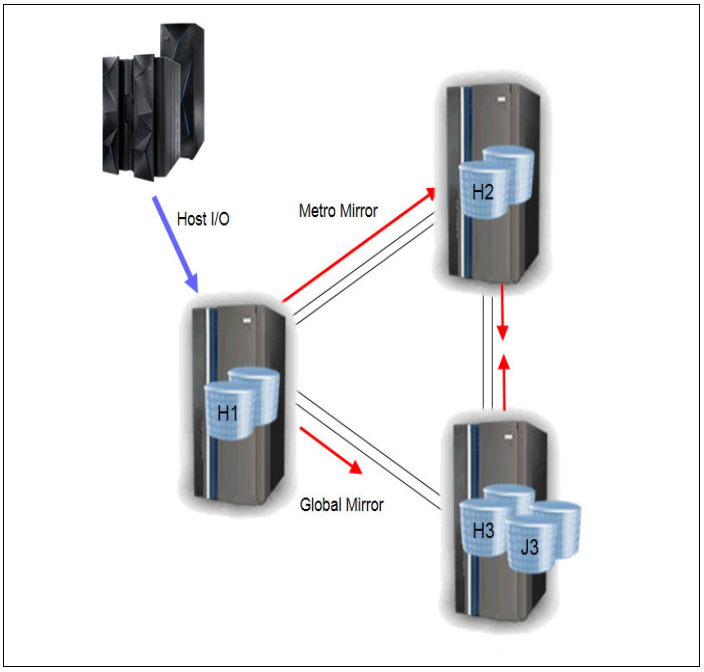

Figure 2-1 Sample figure showing Multiple Target PPRC topology

In Figure 2-1, the different locations are identified by the labels H1, H2, and H3. Each of these refers to a site that typically has many volumes. Each site can also contain multiple DS8870 storage systems. The Global Mirror journal volumes on the Global Mirror recovery sites are indicated by J3 in this example.

The double thin lines between each of the sites represent PPRC paths. PPRC paths are normally established in both directions so that they are available when needed. The path direction is not usually indicated in the figures in the interest of keeping the diagrams simple.

The Host I/O is shown with an arrow going to the primary production site.

The type of mirror or copy (Metro Mirror, Global Mirror, or Global Copy) is shown for each relationship. Active relationships are shown by an arrow from the primary to the secondary. Suspended relationships are indicated by a short arrow from the primary site.

Figure 2-1 represents a Multiple Target PPRC configuration with Metro Mirror H1:H2 and Global Mirror H1:H3. The short arrow for the Global Mirror H1:H3 indicates that the relationship is suspended. The Global Mirror journal volumes are indicated by J3. There are PPRC paths established among all three sites.

Figure 2-1 also shows suspended pairs between H2 and H3. These are internally created pairs, as described in 3.2, “Multiple Target Incremental Resynchronization” on page 33.

2.2 Two Metro Mirror relationships

In a two Metro Mirror pair topology, data is synchronously mirrored to both secondary sites at the same time.

One possible configuration is to have H1 and H2 at the same location, perhaps even within the same data center. H3 can be at a different location, although still within Metro Mirror distance. In this configuration, H2 provides a high-availability copy in case of a failure of H1. In addition, the Incremental Resynchronization capability described in 3.2, “Multiple Target Incremental Resynchronization” on page 33, provides the capability of establishing Metro Mirror H2:H3 so that even after a failure there is a capability to quickly resume the protection of another Metro Mirror copy.

H3 provides a remote copy of data in case of a more widespread failure that affects both H1 and H2.

Figure 2-2 shows a topology with two Metro Mirror pairs.

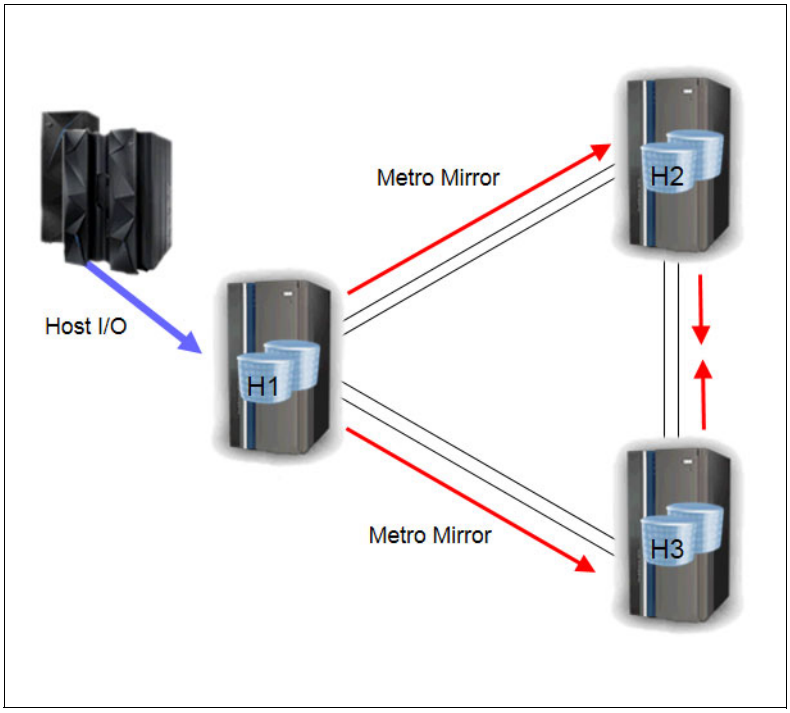

Figure 2-2 Two Metro Mirror pairs

The specifics of using this topology, including several scenarios for recovering after failures at the different sites, are described in detail in Chapter 4, “Implementing a two Metro Mirror topology” on page 45.

2.3 Metro Mirror and Global Mirror

A topology with Metro Mirror and Global Mirror provides a local high availability capability and a long-distance disaster recovery capability at the same time. In such a configuration, data is synchronously mirrored with Metro Mirror to one secondary site while being asynchronously mirrored by Global Mirror to a separate disaster recovery site.

In Figure 2-3, H1:H2 is a Metro Mirror relationship. This configuration provides a synchronous mirror for recovery in case of a failure at the H1 site. It can be used for HyperSwap in a HyperSwap environment.

H1:H3 is an asynchronous Global Mirror relationship. The Global Mirror FlashCopy journal volumes are indicated by J3. Because this is an asynchronous copy, H3 can be located in a different region at a long distance from the local primary site. In the event of a widespread disaster that affects both H1 and H2, production can be restarted at this remote location.

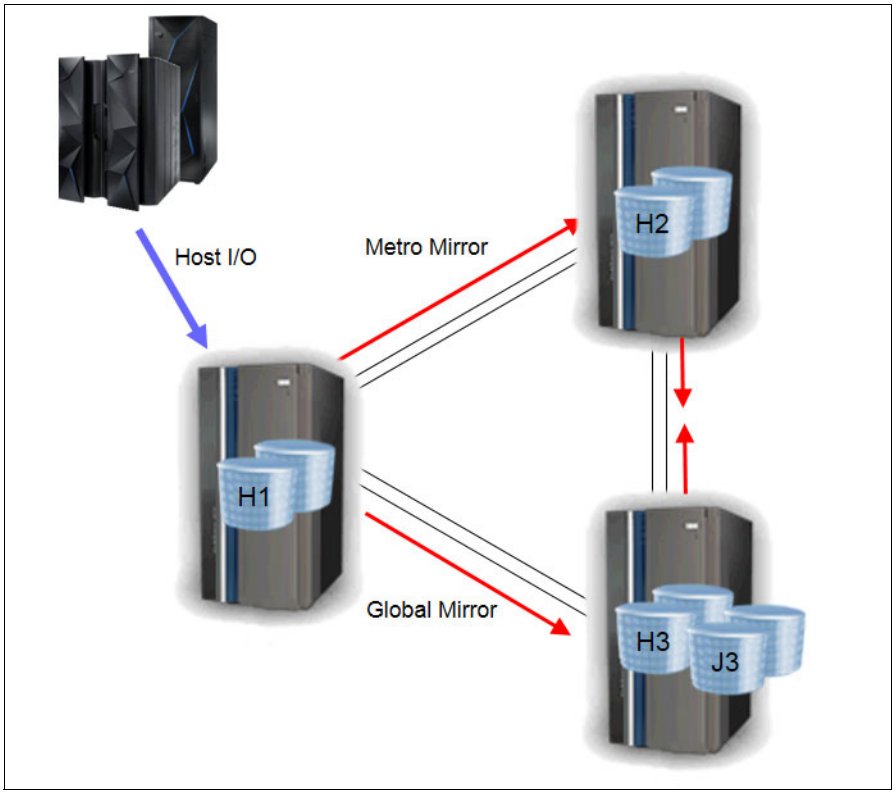

Figure 2-3 Metro Mirror and Global Mirror

The specifics of using this topology, including several scenarios for recovering after failures at the different sites, are described in detail in Chapter 5, “Implementing a Metro Mirror and Global Mirror topology” on page 63.

2.4 Metro Mirror and Global Copy

Like the topology with Metro Mirror and Global Mirror, a topology with Metro Mirror and Global Copy provides both a local synchronous copy and a long-distance asynchronous copy. Unlike Global Mirror, however, Global Copy does not create a dependent write-consistent copy for disaster recovery purposes. However, this topology has several uses:

- Intermediate configuration for two Metro Mirror pairs

  In a two Metro Mirror pair configuration, the Metro Mirror pairs can be initially established as Global Copy. After all of the pairs have completed copying the majority of their data, they can be converted to Metro Mirror. This is a common technique when establishing Metro Mirror pairs to minimize the host impact when synchronizing or resynchronizing volumes where there is a large amount of data to copy. By keeping the pairs in a Global

Copy asynchronous relationship, there is little or no host response time impact while the initial pass of data copy is being performed.

- Periodic copies

  Various techniques can be used to make data consistent at the secondary, such as converting the relationship to Metro Mirror and then suspending the pair. The resultant copy will be consistent until the time of the suspension and can be used as a practice copy, for data mining, as a backup, or for other purposes.

- Migration

  Global Copy can be used to asynchronously copy data over a long distance without affecting the host response time. After the completion of the initial data copy, various techniques can be used to make the data consistent, such as converting the relationship to Metro Mirror or using the Multiple Target PPRC Incremental Resynchronization feature to establish a relationship between the two remote sites. The Incremental Resynchronization feature is described in section 3.2, “Multiple Target Incremental Resynchronization” on page 33.

As with a two Metro Mirror topology, the Incremental Resynchronization capability allows for quickly resuming Metro Mirror from H2 to H3 in the event of an outage at site H1.

Figure 2-4 shows Metro Mirror H1:H2 and Global Copy H1:H3.

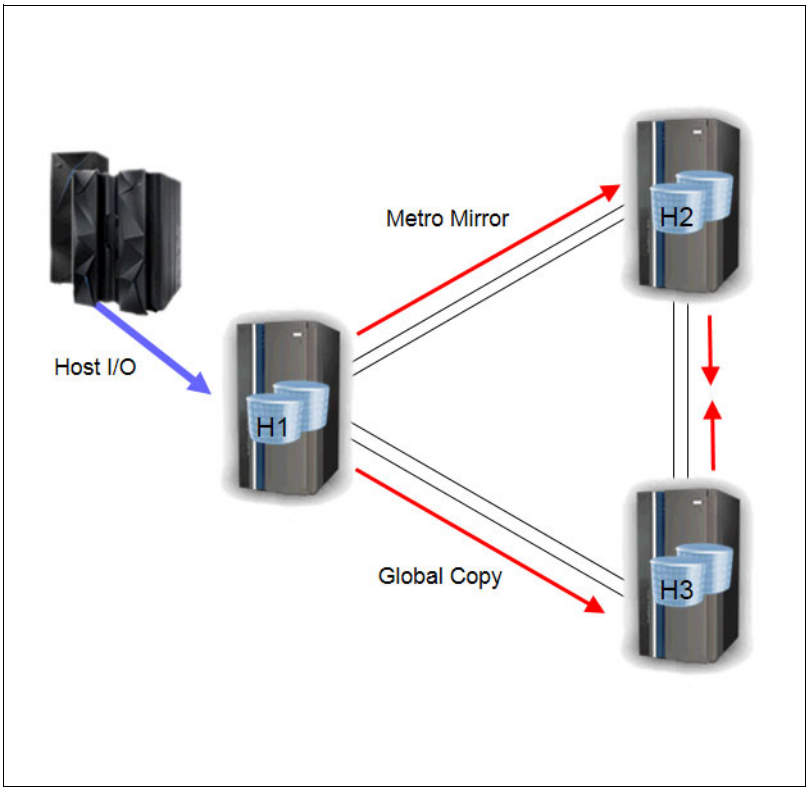

Figure 2-4 Metro Mirror and Global Copy

2.5 Global Copy plus Global Mirror

Multiple Target PPRC supports a topology with two asynchronous copies. Although Global Mirror does not allow for a primary volume to be in two different Global Mirror sessions at the same time, Multiple Target PPRC does allow for a Global Copy and a Global Mirror relationship on the same volume.

With this topology, there is mirroring among three sites in different regions at a long distance from each other. In the event of a failure at any one of the three sites, Global Mirror can be established between the two remaining sites by using the Incremental Resynchronization function described in section 3.2, “Multiple Target Incremental Resynchronization” on page 33.

In Figure 2-5, there are Global Copy relationships, H1:H2, and Global Mirror relationships, H1:H3, with the Global Mirror journal volumes at J3. Global Mirror journal volumes J1 and J2 are also shown at the H1 and H2 sites. This is so they are available for situations where the Global Mirror sessions are moved between different sites. If a site will never be used as a Global Mirror secondary site, these journal volumes are not required.

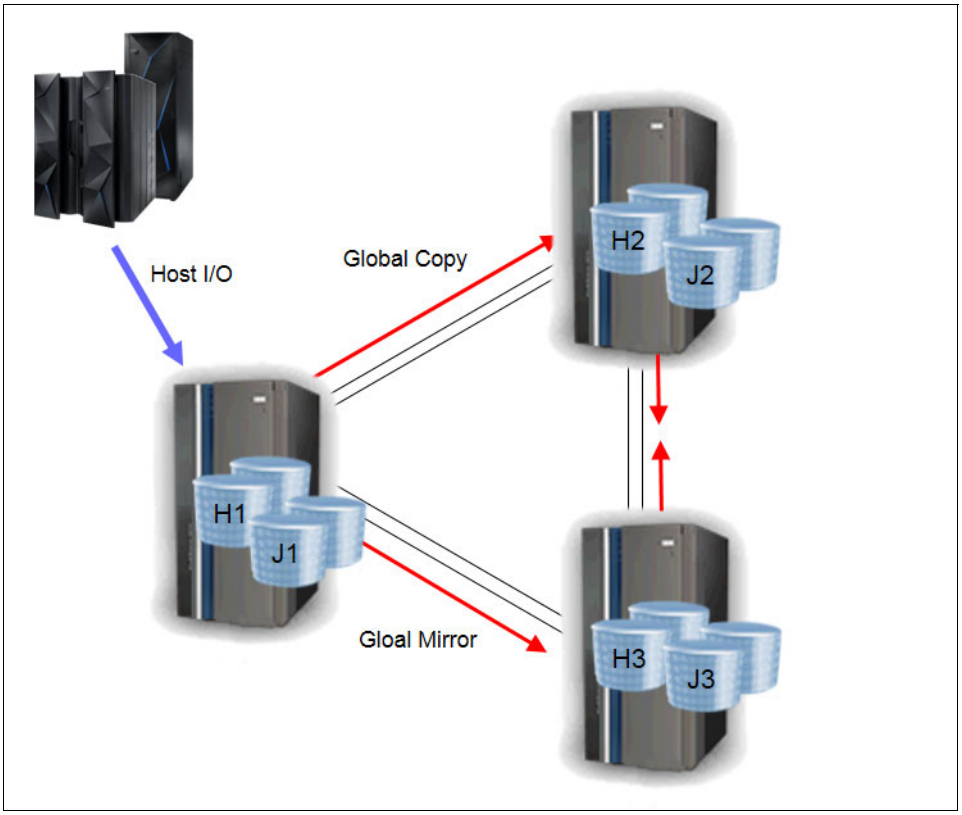

Figure 2-5 Global Copy plus Global Mirror

The specifics of using this topology are described in detail in Chapter 6, “Implementing a Global Copy plus Global Mirror topology” on page 77.

2.6 Metro Global Mirror plus Metro Mirror

A cascaded Metro Global Mirror (MGM) configuration consists of Global Mirror relationship cascaded from a Metro Mirror. Multiple Target PPRC allows for an additional Metro Mirror pair to be added, as shown in Figure 2-6. The Metro Mirror H1:H2 and Global Mirror H2:H4 create an MGM configuration. The additional Metro Mirror pair H1:H3 can be added to create an MGM plus Metro Mirror configuration.

In the event of a failure at site H1, the production applications can be moved to run at H3, and the Incremental Resynchronization capabilities of Multiple Target PPRC can be used to establish an active Metro Mirror relationship H3:H2. This situation results in an MGM configuration where there is Metro Mirror H3:H2 and Global Mirror H2:H4. Even after a failure of the primary production site, there is still the full protection of an MGM environment where Metro Mirror H3:H2 provides a high availability capability and the Global Mirror H2:H4 provides for long-distance disaster recovery.

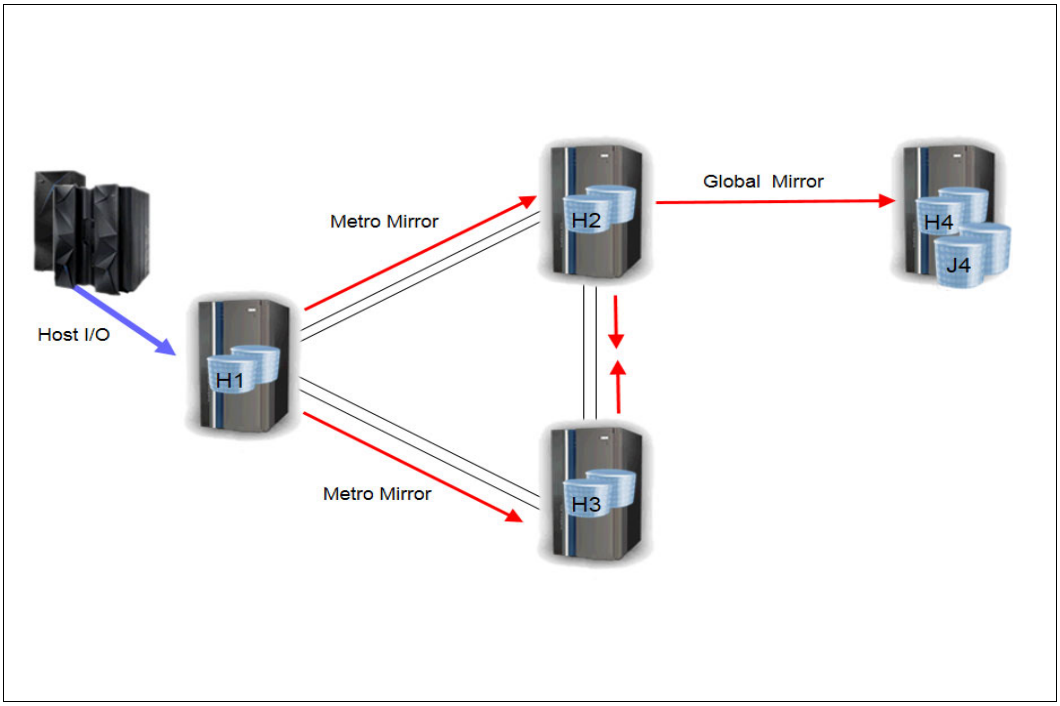

Figure 2-6 Metro Global Mirror plus Metro Mirror

2.7 z/OS Global Mirror and two Metro Mirror pairs

A topology with z/OS Global Mirror and one Metro Mirror pair is referred to as Mz/GM. With Multiple Target PPRC, it is possible to add a second Metro Mirror pair to an Mz/GM configuration. Figure 2-7 on page 30 shows such a configuration. There is Metro Mirror H1:H2 and z/OS Global Mirror H1:H3. These H1, H2, and H3 sites make up an Mz/GM configuration. With Multiple Target PPRC, a second Metro Mirror pair H1:H4 can be added.

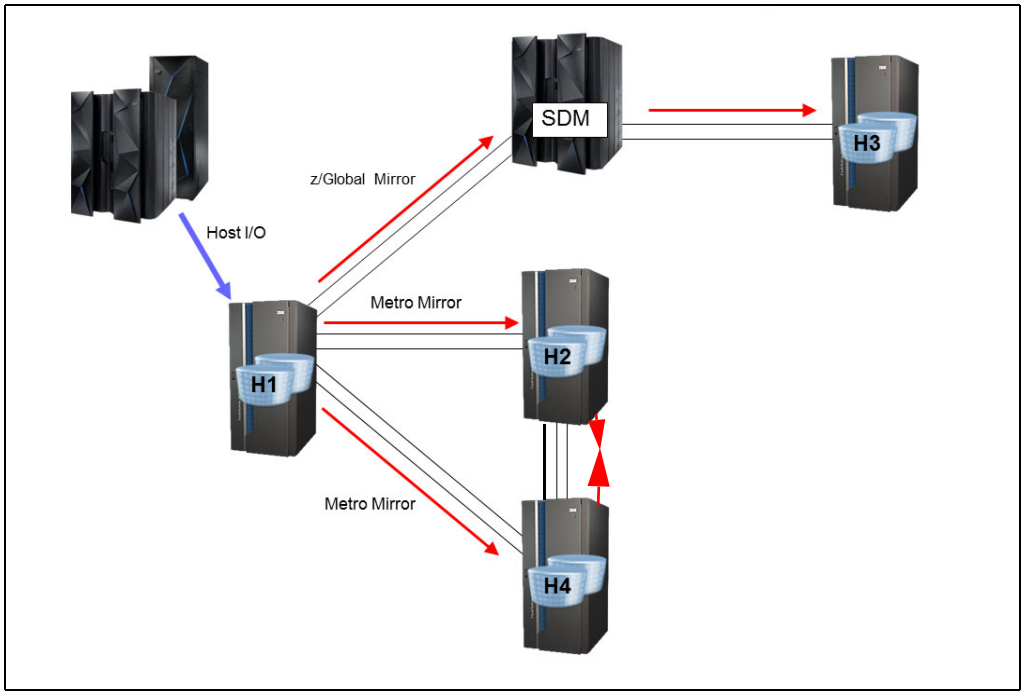

Figure 2-7 z/OS Global Mirror and two Metro Mirror pairs

3

Architecture and design

This chapter provides information about the architecture and design of Multiple Target PPRC. Knowing the design details helps the users understand the reasons for the various steps and procedures that are detailed in other sections of this paper. With this background, it is possible to modify the scenarios and sequences to fit the users' specific needs.

The following topics are covered in this chapter:

- Multiple Target PPRC pairs
- Multiple Target Incremental Resynchronization
- PPRC path states
- Remote Pair FlashCopy

3.1 Multiple Target PPRC pairs

In the Multiple Target PPRC design, each PPRC pair operates independently. As an example, in Figure 3-1, suspending, resuming, or terminating the H1:H3 pair does not change the state of or otherwise affect the H1:H2 pair. Similarly, if there is an existing H1:H2 pair, establishing an H1:H3 pair does not affect the previously existing H1:H2 pair.

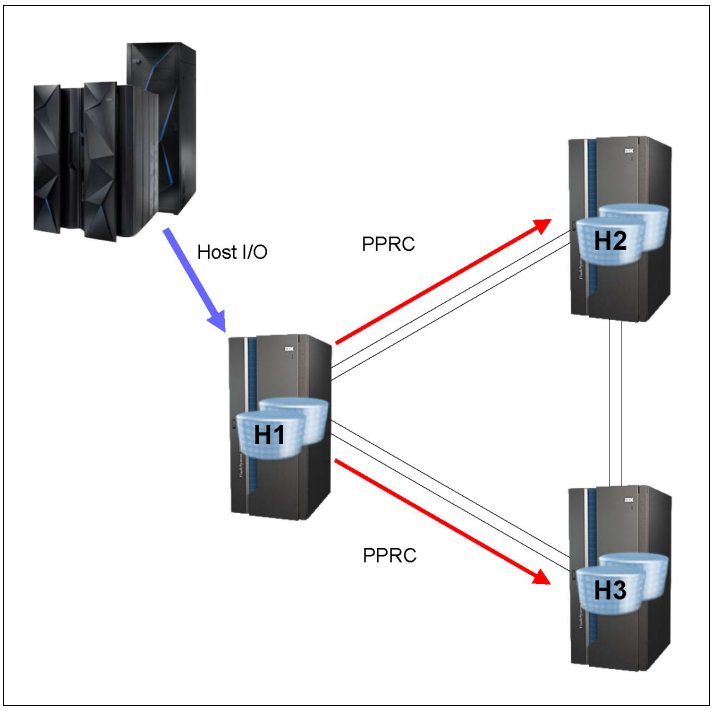

Figure 3-1 Example of a Multiple Target PPRC configuration

Note: Attempting to establish a third pair on H1 is not allowed and will result in an error.

Example 3-1 shows the message displayed by the DS CLI when attempting to create a third PPRC relationship on a volume.

Example 3-1 Attempt to establish three pairs on a volume

```
dscli> mkpprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYM31 -type mmir 5000:6000
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship
IBM.2107-75CZM21/5000:IBM.2107-75CYM31/6000 successfully created.

dscli> mkpprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYK71 -type mmir 5000:7000
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship
IBM.2107-75CZM21/5000:IBM.2107-75CYK71/7000 successfully created.

dscli> mkpprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYM31 -type mmir 5000:6100
CMUN80520E mkpprc: IBM.2107-75CZM21/5000:IBM.2107-75CYM31/6100: The Peer-to-Peer Remote Copy (PPRC)
multitarget relationship could not be established because more than two secondary relationships are
already active for the primary.
```

3.2 Multiple Target Incremental Resynchronization

An essential feature of Multiple Target PPRC is the ability to establish active PPRC relationships between the two secondary sites, H2 and H3, without requiring a full copy of the data. This capability is called *Multiple Target Incremental Resynchronization*.

3.2.1 Multiple Target Incremental Resynchronization pairs

This Incremental Resynchronization capability is available for all types of PPRC replication. The relationships H1:H2 and H1:H3 can be Metro Mirror, Global Copy or Global Mirror or a combination of Metro Mirror, Global Copy, Global Mirror and an Incremental Resynchronization is possible.

The ability to establish active PPRC pairs between the two secondary sites is facilitated by the creation of *Multiple Target Incremental Resynchronization* (MTIR) pairs. These MTIR pairs are automatically created by the DS8870 whenever they are required in a Multiple Target PPRC configuration.

When two PPRC relationships are created on primary volume, the DS8870 detects that it is now part of a Multiple Target PPRC configuration and sends commands to each of the secondary volumes, instructing them to create an MTIR pair that points to the other secondary volume. These pairs are created in the suspended state.

As an example, in Figure 3-2, Metro Mirror pairs have been established for H1:H2 and then H1:H3. When the H1:H3 pair is established, the DS8870 detects that H1 is now a Multiple Target PPRC primary volume. H1 sends a command to H2 instructing it to create an MTIR pointing to H3. Similarly, H1 also sends a command to H3 instructing it to create an MTIR pair pointing to H2. Each of these pairs is created in a suspended state with a specific suspension reason identifying it as an MTIR pair.

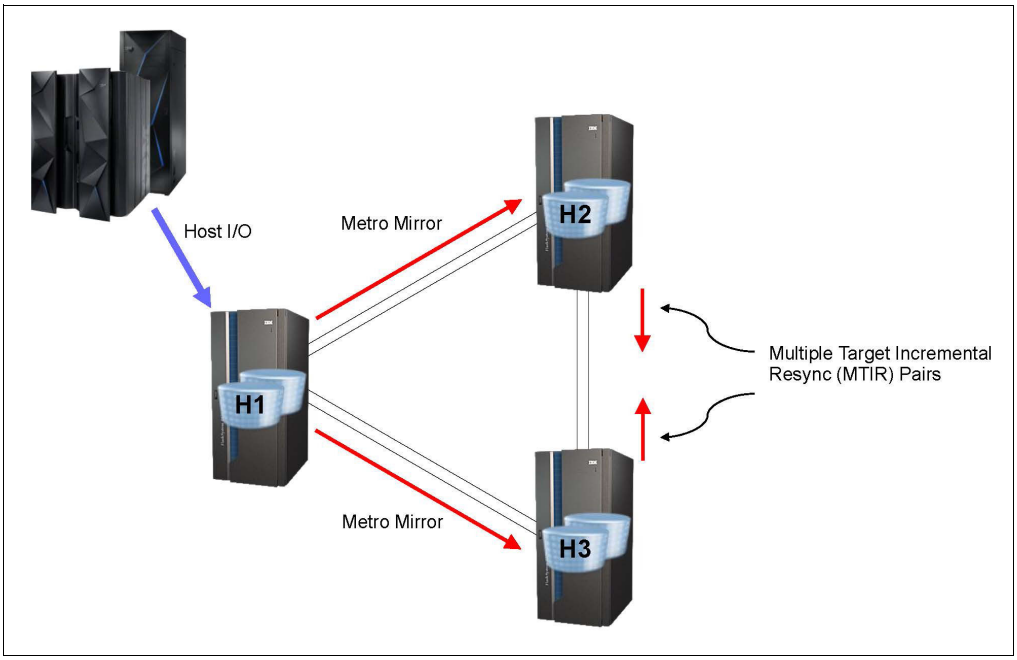

Figure 3-2 Creation of Multiple Target Incremental Resync pairs

Establishing an MTIR pair does not require connectivity to the other secondary volume. For example, referring to Figure 3-2 on page 33 again, it is not required for PPRC paths to be established as H2:H3 for the H2:H3 MTIR pair to be created. However, to later convert the MTIR pairs to active relationships does require having PPRC paths established. Establishing all of the required PPRC paths before establishing the Multiple Target PPRC pairs helps to ensure that the paths are available when they are needed.

Note: When setting up a Multiple Target PPRC configuration, it is advisable to establish the PPRC paths H2:H3 and H3:H2 before establishing the PPRC pairs H1:H2 and H1:H3.

The MTIR pairs serve two main purposes:

- They enable an active relationship to be quickly established between the two secondary volumes by converting the existing pair rather than establishing a new pair.
- They provide a change recording mechanism to track which data is potentially different between the two secondary volumes. This is what allows for the resynchronization to be an incremental copy rather than a full copy.

3.2.2 Delayed creation of MTIR pairs

If there is no connectivity between H1 and H2 at the time that H1:H3 is created, the MTIR pairs H2:H3 will not be created. These MTIR pairs will be created later after the connectivity has been restored.

As an example, consider the following sequence:

1. Establish pairs H1:H2.
2. Connectivity between H1 and H2 is lost.
   a. When data is attempted to be transferred from H1 to H2, it will not succeed and the H1:H2 pairs will suspend.
3. Establish pairs H1:H3.
4. H1 sends a command to H3 instructing it to create the H3:H2 MTIR pair.
5. H3 creates the H3:H2 MTIR pair.
6. H1 sends a command to H2 instructing it to create the H2:H3 MTIR pair. Because the connectivity to H2 is lost, the command cannot be sent and the H2:H3 MTIR pair is not created.

   At this point, there are MTIR pairs H3:H2 but no MTIR pairs H2:H3.
7. Connectivity for H1:H2 is restored and the H1:H2 pairs are resumed.
8. The MTIR pairs H2:H3 are created as part of the PPRC resume processing.

3.2.3 MTIR change recording

Each of the MTIR pairs contains structures that enable a change recording mechanism to keep track of data that is potentially out of synchronization between the two secondary volumes. This change recording mechanism is implemented by bitmaps in which each bit represents a track.

Figure 3-3 shows an MTIR pair H2:H3, including the change recording bitmap structure.

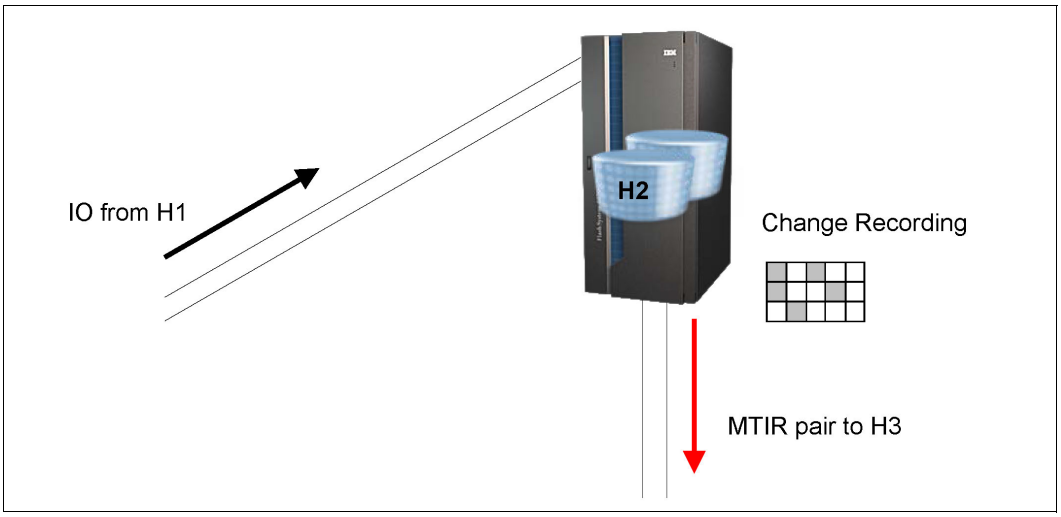

Figure 3-3 MTIR change recording

As writes are received at H2 from the primary H1 volume, the modified tracks are recorded in the change recording bitmap file to indicate that the data has been updated on H2 and is, therefore, potentially out of synchronization with the other secondary volume, H3.

The DS8870 firmware uses an internal algorithm to determine when a track has been successfully written to both H2 and H3 secondary volumes and resets the track in the bitmap. The resetting of bits in the bitmap is not done continuously. Instead, groups of tracks are collected and reset at one time. This is done to reduce the frequency of commands sent between the storage controllers. If a track is subsequently written again after a previous write had been successfully written to both secondaries, it will not be reset until the new write has also been written to both secondaries.

The same tracking is also performed on the H3 volume. Each of the secondary volumes thus contains a record of data that it has received and for which it has not yet received a notification that the data has also been written and hardened on the other secondary volume.

Notes: This description explains the effect of the change recording algorithm, although not necessarily its actual internal implementation.

The Incremental Resynchronization capability does not depend upon the type of mirror or copy. Regardless of the pair types, Metro Mirror, Global Copy, or Global Mirror, an Incremental Resynchronization between the secondary volumes is possible.

When an MTIR pair is created, the change recording structure indicates that the entire volume is out of synchronization with the other secondary volume. After the initial copy of the volume's data has been completed, the change recording mechanism starts the process of recording new updates and resetting the indications for the previous updates. It requires a few minutes before the out-of-synchronization indications are reset. Until this time, the change recording structures continue to show that the entire volume is out of synchronization.

Each MTIR pair can be queried to see how many tracks are out of synchronization between the two secondary volumes.

Example 3-2 shows an example query about the out of synchronization count for an MTIR pair, using the DS CLI. The first **lspprc** command shows that the out-of-sync track count is 81920 tracks for all volumes, which is the size of these particular volumes. A later query shows that the counts have dropped to a value lower than the full volume. As long as I/O is running to a volume, the counts will remain non-zero. However, if the host I/O is quiesced, the counts will eventually drop to zero.

Example 3-2 Query MTIR out of sync tracks

```
dscli> lspprc -l -dev IBM.2107-75CYM31 -remotedev IBM.2107-75CYK71 -multtgt 6000-600f:7000-700f
ID                                        State     Reason               Type        Out Of Sync Tracks
=============================================================================================================
IBM.2107-75CYM31/6000:IBM.2107-75CYK71/7000 Suspended Multi-target Internal Global Copy 81920               ...
IBM.2107-75CYM31/6001:IBM.2107-75CYK71/7001 Suspended Multi-target Internal Global Copy 81920               ...
IBM.2107-75CYM31/6002:IBM.2107-75CYK71/7002 Suspended Multi-target Internal Global Copy 81920               ...

dscli> lspprc -l -dev IBM.2107-75CYM31 -remotedev IBM.2107-75CYK71 -multtgt 6000-600f:7000-700f
ID                                        State     Reason               Type        Out Of Sync Tracks
=============================================================================================================
IBM.2107-75CYM31/6000:IBM.2107-75CYK71/7000 Suspended Multi-target Internal Global Copy 259                 ...
IBM.2107-75CYM31/6001:IBM.2107-75CYK71/7001 Suspended Multi-target Internal Global Copy 136                 ...
IBM.2107-75CYM31/6002:IBM.2107-75CYK71/7002 Suspended Multi-target Internal Global Copy 536                 ...
```

Note: By default, the DS CLI does not display the internal MTIR pairs in the **lspprc** output. The **-multtgt** option must be used for **lspprc** to show the MTIR pairs.

3.2.4 Synchronizing MTIR pairs

Multiple Target Incremental Resynchronization pairs are synchronized by using a PPRC failback command. This is the same failback command that is used to reverse a PPRC relationship after a PPRC failover. It is possible to use the standard failback command because Multiple Target PPRC creates the MTIR pairs in such a manner that they are in the correct state for this operation.

In Figure 3-4, the PPRC failback command for H2:H3 converts the MTIR pairs to an active pair.

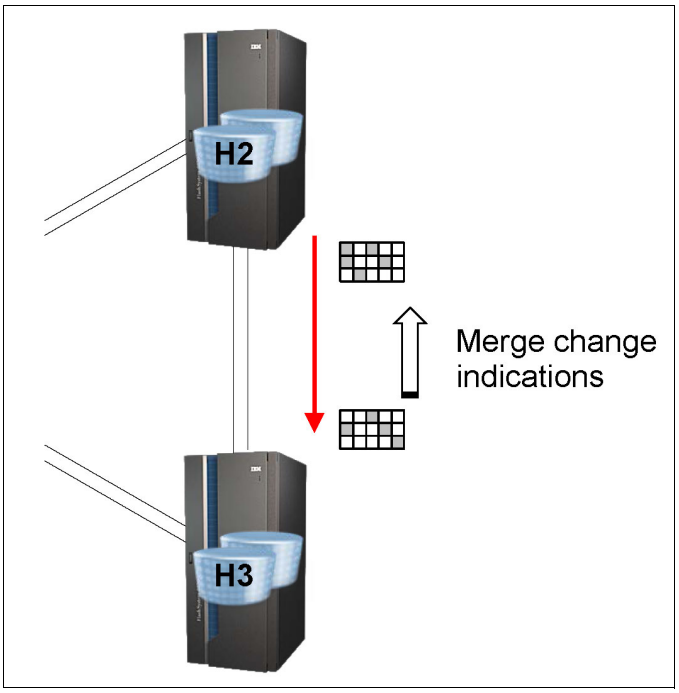

Figure 3-4 Synchronize MTIR pairs

The PPRC failback command for H2:H3 performs the following steps:

1. H2 sends a command to H3 to convert it to a secondary volume of H2 and place it in the duplex pending state.
2. H2 reads the change recording information from H3.
3. H2 merges the change recording information from H3 into its own change recording information.

   H2 now has a complete list of all changes that are potentially out of synchronization with H3.
4. H2 transfers the data indicated in the change recording structure to synchronize the pair.

   If Metro Mirror is specified on the failback pair, the pair transitions to the full-duplex state. If Global Copy is specified, the pair completes the first pass of the initial copy.

If a synchronization of the MTIR pairs is performed before the initial copy of the volume's data has been completed, the change recording structures still indicate that all tracks on the volume are out of synchronization and a full copy of data is performed.

As shown previously in Example 3-2 on page 36, the MTIR pairs can be queried to see how many tracks are recorded as being out of synchronization, which indicates the amount of data that will be copied if a synchronization is performed.

3.3 PPRC path states

Important: If a synchronization is performed before the change recording structures have reset their change indications, a full copy of the data occurs.

To keep the diagrams in this paper clear, the majority of them do not show the state and direction of the PPRC paths. For Multiple Target PPRC it is advisable to establish active paths in both directions among the H1, H2, and H3 sites. However, in some migration cases, not all of these sets of paths are required because the intermediate configurations are temporary and there is no requirement to establish active mirroring in all possible directions.

Note: A knowledge of the following PPRC design points is required to understand the state of the PPRC paths after different types of operations.

3.3.1 Effects of freeze on path state

A freeze command changes the state of a set of PPRC paths to Failed but does not totally remove their existence.

For example, Example 3-3 uses a DS CLI command to show how the state of a set of paths is changed by a freeze command. Before the freeze, there are two logical paths established from IBM.2107-75CZM21/50 to IBM.2107-75CYM31/60. The freeze command removes the logical path connections. However, the relationship between the two storage systems remains, as shown in the second **lspprcpath** query in the example.

Example 3-3 Path state before and after a freeze

```
dscli> lspprcpath -dev IBM.2107-75CZM21 50
Src                 Tgt                 State   SS   Port                  Attached Port          Tgt WWNN
=============================================================================================================
IBM.2107-75CZM21/50 IBM.2107-75CYM31/60 Success FF60 IBM.2107-75CZM21/I0201 IBM.2107-75CYM31/I0232 5005076305FFD71E
IBM.2107-75CZM21/50 IBM.2107-75CYM31/60 Success FF60 IBM.2107-75CZM21/I0234 IBM.2107-75CYM31/I0332 5005076305FFD71E

dscli> freezepprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYM31 50:60
CMUC00161I freezepprc: Remote Mirror and Copy consistency group IBM.2107-75CZM21/50:IBM.2107-75CYM31/60 successfully
created.

dscli> lspprcpath -dev IBM.2107-75CZM21 50
Src                 Tgt                 State   SS   Port                  Attached Port          Tgt WWNN
=============================================================================================================
IBM.2107-75CZM21/50 IBM.2107-75CYM31/60 Failed  FF60 -                     -                      5005076305FFD71E
```

A command to remove the paths, such as the DS CLI **rmpprcpath** command, will completely delete the relationship, as demonstrated in Example 3-4.

Note: All PPRC pairs using the paths must be removed before removing the paths.

Example 3-4 Remote PPRC paths

```
dscli> rmpprcpath -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYM31 -remotewwnn 5005076305FFD71E -quiet 50:60
CMUC00150I rmpprcpath: Remote Mirror and Copy path IBM.2107-75CZM21/50:IBM.2107-75CYM31/60 successfully removed.

dscli> lspprcpath 50
CMUC00234I lspprcpath: No Remote Mirror and Copy Path found.
```

3.3.2 Effects of failover on path state

PPRC paths are not required to be established before performing a failover operation, although, in most cases, it is best that the paths be established beforehand. A failover when there are no paths established causes the storage controller to internally create a path structure that can later be replaced with the actual paths that are to be used.

As an example, consider a pair H1:H2, where paths are established in the H1:H2 direction but not in the opposite H2:H1 direction. The example uses the DS8000 identifications shown in Table 3-1.

Table 3-1 Identifications used in DS CLI examples

Site	Role	Dev	WWNN	Volume range
H1	Metro Mirror primary	IBM.2107-75CZM21	5005076305FFD75A	5000 - 500F
H2	Metro Mirror secondary	IBM.2107-75CYM31	5005076305FFD71E	6000 - 600F

Even though there are no paths established for H2:H1, a failover of H2:H1 is still allowed. The DS8870 storage system will create a placeholder for the paths and mark the paths as failed, as shown in Example 3-5.

Example 3-5 Failover creates failed paths is required

```
dscli> lspprcpath 60
CMUC00234I lspprcpath: No Remote Mirror and Copy Path found.

dscli> failoverpprc -dev IBM.2107-75CYM31 -remotedev IBM.2107-75CZM21 -type mmir 6000-600f:5000-500f
CMUC00196I failoverpprc: Remote Mirror and Copy pair IBM.2107-75CYM31/6000:IBM.2107-75CZM21/5000
successfully reversed.
...
CMUC00196I failoverpprc: Remote Mirror and Copy pair IBM.2107-75CYM31/600F:IBM.2107-75CZM21/500F
successfully reversed.

dscli> lspprcpath 60
Src                 Tgt                 State  SS   Port Attached Port Tgt WWNN
=============================================================================
IBM.2107-75CYM31/60 IBM.2107-75CZM21/50 Failed FFFF -    -             5005076305FFD75A
```

Before performing a PPRC failback command to create active pairs, the PPRC paths must be successfully established. If no failback is planned, it is not required to establish the paths.

This creation of failed PPRC paths is noted here because there are some migration scenarios where these paths are created. Having an awareness of these paths is required to understand why they appear and when they can be removed. These paths can be removed in the same manner other established paths.

3.4 Remote Pair FlashCopy

FlashCopy can be used to create a nearly instantaneous copy of a volume or an extent range (for z/OS) from a source location to a target location within the same DS8000 storage system.

When both the source and target volumes are in a Metro Mirror relationship that meets certain configuration and state requirements, *Remote Pair FlashCopy* can be used to mirror the FlashCopy command from the Metro Mirror primary volumes to the secondary volumes.

Note: The terms *Remote Pair FlashCopy* and *Preserve Mirror* are used interchangeably. Preserve Mirror is the IBM z/OS software function that uses IBM Remote Pair FlashCopy.

Figure 3-5 shows a configuration where there is a Metro Mirror pair S1:S2 and also T1:T2. Assuming that configuration and states permit the use of Remote Pair FlashCopy, a FlashCopy establish from S1 to T1 will result in a FlashCopy command being transferred over the PPRC paths from site H1 to site H2 and performing the same FlashCopy command for S2 to T2 at the H2 site.

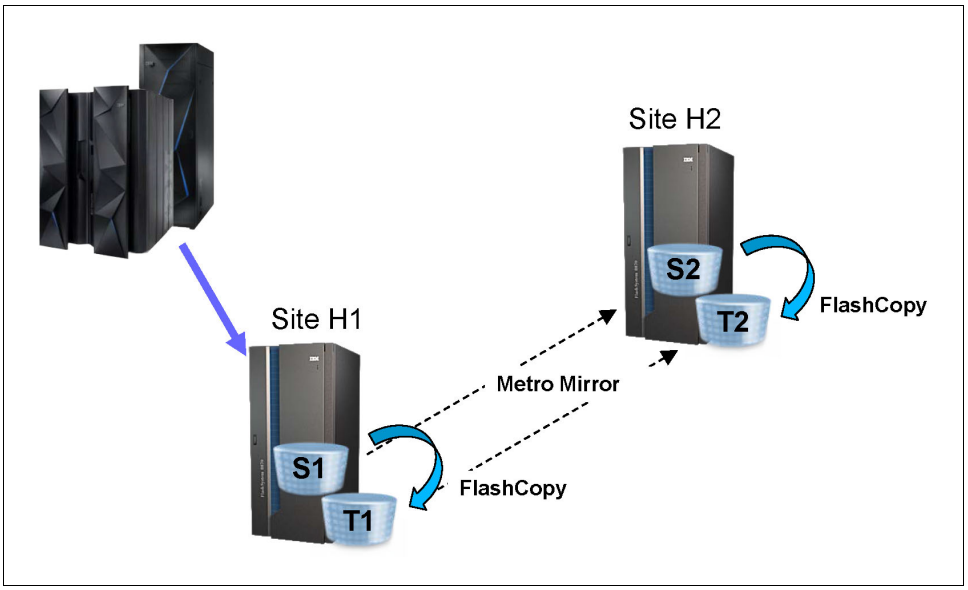

Figure 3-5 Remote Pair FlashCopy

In some instances where the configuration or state requirements do not allow a Remote Pair FlashCopy operation to be performed, the FlashCopy command from S1 to T1 will instead change the pair state for the target Metro Mirror pair T1:T2 to duplex pending and mark the tracks for the FlashCopy operation as being out of synchronization. Metro Mirror will then copy the changed data from T1 to T2 and, when the copy is complete, the T1:T2 pair will again reach the full duplex state.

Note: *FlashCopy* is a logical write of data from a source volume to a target volume. Only the target volume's data is modified, so it is only the target volume that will transition to duplex that is pending on a FlashCopy operation.

For more information about the use of Remote Pair FlashCopy, see the IBM Redpaper publication titled *IBM System Storage DS8000: Remote Pair FlashCopy (Preserve Mirror)*, REDP-4504.

With Multiple Target PPRC, the S1 and T1 volumes can now have two Metro Mirror relationships on them, as shown in Figure 3-6.

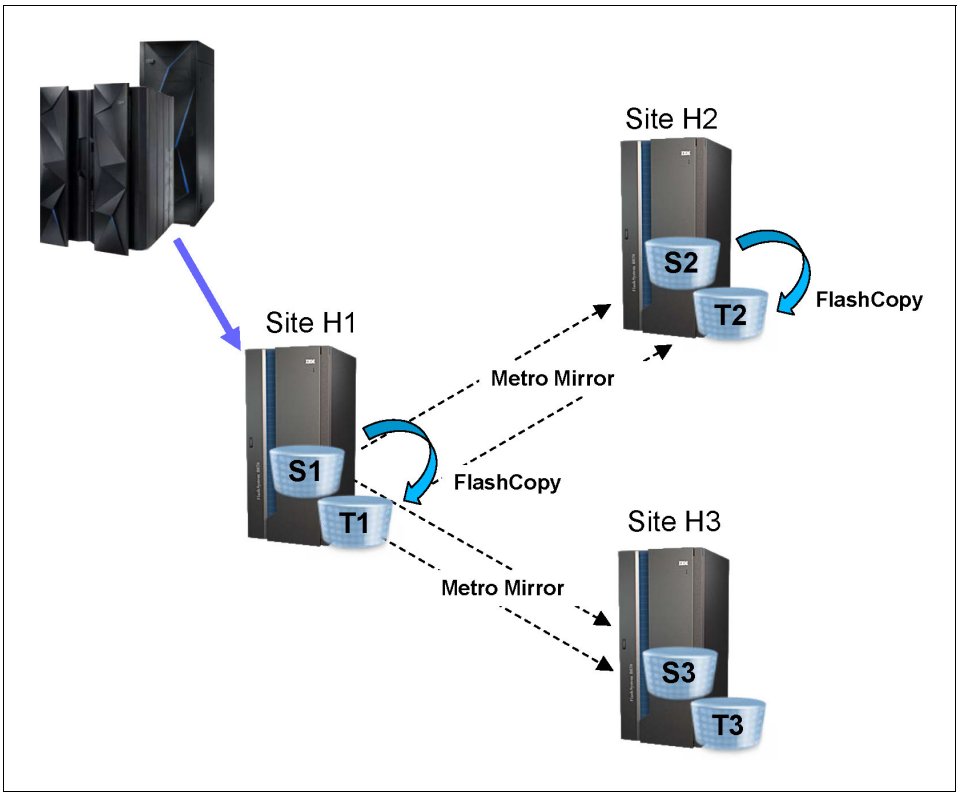

Figure 3-6 Remote Pair FlashCopy with Multiple Target PPRC

A Remote Pair FlashCopy command can be performed for only one of the Metro Mirror pairs. The other pair will transition back to duplex pending and mirror the copied data to the remote volume.

In a configuration with two Metro Mirror pairs, the user must identify which of the pairs to use for the Remote Pair FlashCopy operation. The Set PPRC Characteristics command is used to set or reset *Use Remote Pair FlashCopy* for a Metro Mirror pair to indicate whether Remote Pair FlashCopy is to be used for that pair. DS CLI, TSO, or ICKDSF can be used to set or reset this indication for a specific pair. In the case of a failure or suspension of the pair that is enabled for Remote Pair FlashCopy, the Use Remote Pair FlashCopy command can be used to enable the function on the surviving Metro Mirror pairs.

Example 3-6 on page 42 shows a DS CLI command to set Use Remote Pair FlashCopy for a pair. The terms Remote Pair FlashCopy and *Preserve Mirror* are used interchangeably. The DS CLI command uses the Preserve Mirror term, so the option is **`pmir`**. With DS CLI, the option to set Use Remote Pair FlashCopy is **`-action enable -ctrl pmir`**.

In this example, there are two Metro Mirror pairs on the primary volume, and the first **`lspprc`** query shows that `Pmir=Disabled` for both pairs. The **`chpprc`** command to enable Pmir for one relationship is demonstrated, followed by another **`lspprc`** command to show that Pmir has been changed to `Enabled` for the specified pair.

The output data from the **`lspprc`** command is lengthy, so parts have been edited out in this example to better show the fields of interest.

Example 3-6 Setting Use Remote Pair FlashCopy (Preserve Mirror)

```
dscli> lspprc -l -dev IBM.2107-75CZM21 5000
ID                                        State       Reason Type          ... Pmir     DisableAutoResync
=========================================================================== ... ========================
IBM.2107-75CZM21/5000:IBM.2107-75CYK71/7000 Full Duplex -      Metro Mirror  ... Disabled -
IBM.2107-75CZM21/5000:IBM.2107-75CYM31/6000 Full Duplex -      Metro Mirror  ... Disabled -

dscli> chpprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYK71 -action enable -ctrl pmir 5000:7000
CMUC00156I chpprc: Remote Mirror and Copy volume pair IBM.2107-75CZM21/5000:IBM.2107-75CYK71/7000
relationship successfully modified.

dscli> lspprc -l -dev IBM.2107-75CZM21 5000
ID                                        State       Reason Type          ... Pmir     DisableAutoResync
=========================================================================== ... ========================
IBM.2107-75CZM21/5000:IBM.2107-75CYK71/7000 Full Duplex -      Metro Mirror  ... Enabled  -
IBM.2107-75CZM21/5000:IBM.2107-75CYM31/6000 Full Duplex -      Metro Mirror  ... Disabled -
```

When a Remote Pair FlashCopy command is issued to a Multiple Target PPRC primary volume and the conditions allow such an operation, the Metro Mirror pair that is enabled for Remote Pair FlashCopy will mirror the FlashCopy command, and the other pair will transition back to duplex pending and copy the data to the remote site using Metro Mirror.

In Example 3-7, there are full-duplex Metro Mirror pairs in the correct state for a Remote Pair FlashCopy operation. As can be seen from the edited **lspprc** output, for each volume, one pair has `Pmir=Enabled` and the other pair has `Pmir=Disabled`. The Remote Pair FlashCopy command is successful. The Metro Mirror pair for the FlashCopy target volume with `Pmir=Enabled` stays full duplex, and the pair for the FlashCopy target volume with `Pmir=Disabled` changes to duplex pending.

Example 3-7 Remote Pair FlashCopy to Multiple Target PPRC primary

```
dscli> lspprc -l 5000-5001
ID                                        State       Reason Type         Out Of Sync Tracks        Pmir
=============================================================================================      =========
IBM.2107-75CZM21/5000:IBM.2107-75CYK71/7000 Full Duplex -      Metro Mirror 0                   ... Enabled
IBM.2107-75CZM21/5000:IBM.2107-75CYM31/6000 Full Duplex -      Metro Mirror 0                   ... Disabled
IBM.2107-75CZM21/5001:IBM.2107-75CYK71/7001 Full Duplex -      Metro Mirror 0                   ... Enabled
IBM.2107-75CZM21/5001:IBM.2107-75CYM31/6001 Full Duplex -      Metro Mirror 0                   ... Disabled

dscli> mkflash -dev IBM.2107-75CZM21 -pmir required 5000:5001
CMUC00137I mkflash: FlashCopy pair IBM.2107-75CZM21/5000:IBM.2107-75CZM21/5001 successfully created.

dscli> lspprc -l 5000-5001
ID                                        State        Reason Type         Out Of Sync Tracks   ... Pmir
============================================================================================== ... =========
IBM.2107-75CZM21/5000:IBM.2107-75CYK71/7000 Full Duplex  -      Metro Mirror 0                 ... Enabled
IBM.2107-75CZM21/5000:IBM.2107-75CYM31/6000 Full Duplex  -      Metro Mirror 0                 ... Disabled
IBM.2107-75CZM21/5001:IBM.2107-75CYK71/7001 Full Duplex  -      Metro Mirror 0                 ... Enabled
IBM.2107-75CZM21/5001:IBM.2107-75CYM31/6001 Copy Pending -      Metro Mirror 44067             ... Disabled
```

If neither pair or if both pairs are enabled for Remote Pair FlashCopy and a Remote Pair FlashCopy command is received, the command will be rejected because the storage system cannot determine which of the Metro Mirror pairs to use for the command.

In Example 3-8 on page 43, both Metro Mirror pairs are enabled for Remote Pair FlashCopy, which causes the Remote Pair FlashCopy command to fail. The result is similar if neither pair is enabled for Remote Pair FlashCopy.

Example 3-8 Remote Pair FlashCopy when Pmir=Enabled for both pairs

```
dscli> chpprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYM31 -action enable -ctrl pmir
5000-5001:6000-6001
CMUC00156I chpprc: Remote Mirror and Copy volume pair IBM.2107-75CZM21/5000:IBM.2107-75CYM31/6000
relationship successfully modified.
CMUC00156I chpprc: Remote Mirror and Copy volume pair IBM.2107-75CZM21/5001:IBM.2107-75CYM31/6001
relationship successfully modified.

dscli> lspprc -l 5000-5001
ID                                        State       Reason Type         Out Of Sync Tracks  ... Pmir
================================================================================================ ... =======
IBM.2107-75CZM21/5000:IBM.2107-75CYK71/7000 Full Duplex -      Metro Mirror 0                  ... Enabled
IBM.2107-75CZM21/5000:IBM.2107-75CYM31/6000 Full Duplex -      Metro Mirror 0                  ... Enabled
IBM.2107-75CZM21/5001:IBM.2107-75CYK71/7001 Full Duplex -      Metro Mirror 0                  ... Enabled
IBM.2107-75CZM21/5001:IBM.2107-75CYM31/6001 Full Duplex -      Metro Mirror 0                  ... Enabled

dscli> mkflash -dev IBM.2107-75CZM21 -pmir required 5000:5001
CMUN80530E mkflash: IBM.2107-75CZM21/5000:IBM.2107-75CZM21/5001: The FlashCopy relationship cannot be
established because it conflicts with other relationships of the target device.
```

4

Implementing a two Metro Mirror topology

This chapter describes the creation of a two Metro Mirror configuration and explains recovery scenarios for failures at the remote or local site.

The following topics are covered:

- Overview of a two Metro Mirror topology
- Creating a two Metro Mirror topology
- Outage at H2 or H3
- Outage at H1 site
- Return production to H1
- Tivoli Storage Productivity Center for Replication

4.1 Overview of a two Metro Mirror topology

In a two Metro Mirror configuration, there are two Metro Mirror pairs with the same source volume. They mirror data synchronously to both secondary sites.

These are basic steps of creating a Metro Mirror topology of H1:H2 and H1:H3:

1. Establish PPRC paths between all sites.
2. Establish the H1:H2 pairs.
3. Establish the H1:H3 pairs.

The terms H1, H2, and H3 refer to the different sites. Each site can contain many volumes and many volume pairs. Each site can also contain more than one storage system. There is no required order for establishing the pairs. “H1:H2” and “H1:H3” refer to the relationships between the sites and are independent of the order in which they are established.

To simplify the descriptions, these sites are often referred to as though they were a single volume, but it must be understood that the terms refer to the entire set of volumes being mirrored by PPRC. For example, “Failover H2:H1” means to failover all of the H2 volumes to their corresponding H1 volumes.

Note: The scenarios in this chapter refer to an outage at one of the sites. This can be either a failure condition or a planned event for testing or maintenance.

4.2 Creating a two Metro Mirror topology

This section describes the detailed steps to create a two Metro Mirror topology.

4.2.1 Terms used in examples

For the examples in this section, the following terms are used:

- **H1** is the current primary site, where the production applications are running.
- **H2** is the first Metro Mirror secondary site, to which H1 is mirroring data.
- **H3** is the second Metro Mirror secondary site, to which H1 is mirroring data.

Table 4-1 identifies the DS8870 storage controllers used in the examples in this chapter.

Note: The volume range on each DS8870 is different in the following examples only to help clarify the different sites used. It is not a requirement that they be different.

Table 4-1 Identifications used in DS CLI examples

Site	Role	Dev	WWNN	Volume range
H1	Current primary	IBM.2107-75CZM21	5005076305FFD75A	5000 - 500F
H2	Metro Mirror secondary	IBM.2107-75CYM31	5005076305FFD71E	6000 - 600F
H3	Metro Mirror secondary	IBM.2107-75CYK71	5005076305FFD71A	7000 - 700F

4.2.2 Establish PPRC paths

PPRC paths are required before the PPRC pairs can be established. It is best to establish all paths that might be needed so that they are available when they are required.

For example, as shown in Figure 4-1, a two Metro Mirror topology has Metro Mirror relationships H1:H2 and H1:H3. There are also Multiple Target Incremental Resynchronization (MTIR) relationships H2:H3 and H3:H2. In most cases, after a move or HyperSwap to H2, a failback H2:H1 is done after the H1 site is recovered, so paths for H2:H1 are needed. Similarly, in the event of a move or HyperSwap to H3, a failback of H3:H1 requires paths for H3:H1.

Therefore, it is advisable to establish the PPRC paths in all possible directions among the H1, H2, and H3 sites. For simplicity, the direction of the PPRC paths is not indicated in these diagrams.

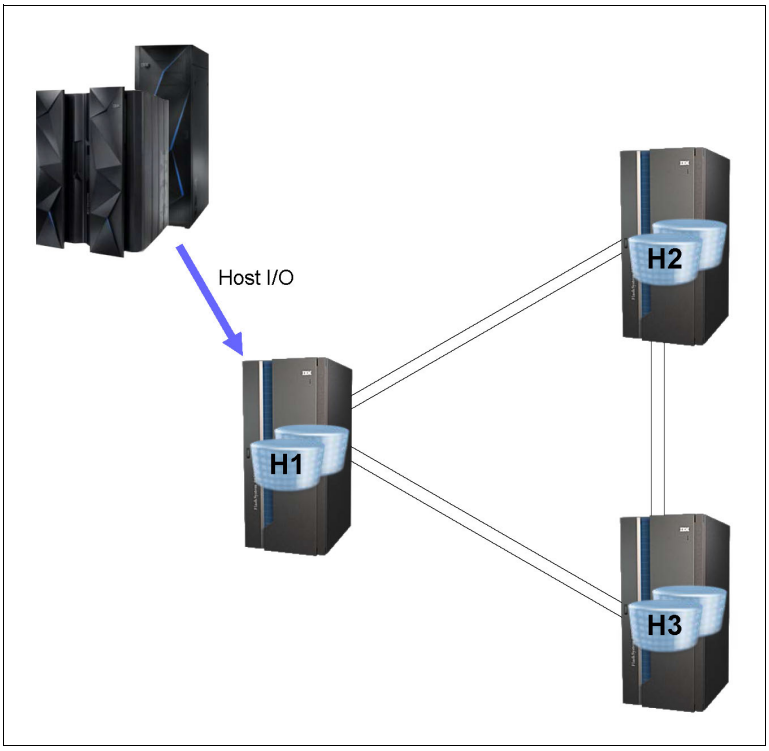

Figure 4-1 Establish PPRC paths

Create the following PPRC paths:

- Establish PPRC Paths H1:H2
- Establish PPRC Paths H1:H3

- Establish PPRC Paths H2:H1
- Establish PPRC Paths H2:H3

- Establish PPRC Paths H3:H1
- Establish PPRC Paths H3:H2

Example 4-1 shows an example of a DS CLI command to create PPRC paths between H1 and H2 and then H1:H3. Similar commands are used to create the remaining PPRC paths.

Example 4-1 Establish PPRC paths H1:H2 and H1:H3

```
dscli> mkpprcpath -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYM31 -remotewwnn 5005076305FFD71E -srclss 50
-tgtlss 60
CMUC00149I mkpprcpath: Remote Mirror and Copy path 50:60 successfully established.

dscli> mkpprcpath -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYK71 -remotewwnn 5005076305FFD71A -srclss 50
-tgtlss 70
CMUC00149I mkpprcpath: Remote Mirror and Copy path 50:70 successfully established.
```

4.2.3 Create H1:H2 Metro Mirror pairs

After establishing all of the required PPRC paths, create the Metro Mirror volume pairs H1:H2. These pairs can be initially created as asynchronous Global Copy pairs and later converted to Metro Mirror, or they can be created directly as Metro Mirror pairs.

In some large configurations, it is preferable to start the replication as Global Copy because it is less likely to impact the production systems during the initial copy phase. After all pairs have copied nearly 100% of the data, the pairs can be converted to synchronous Metro Mirror.

Example 4-2 shows the use of the DS CLI **mkpprc** command to create Metro Mirror pairs H1:H2.

Example 4-2 Establish PPRC pairs H1:H2

```
dscli> mkpprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYM31 -type mmir 5000-500f:6000-600f
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 5000:6000 successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 5001:6001 successfully created.
. . .
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 500E:600E successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 500F:600F successfully created.
```

At this point, Metro Mirror is running H1:H2.

4.2.4 Establish H1:H3 pairs

After establishing the H1:H2 pairs, establish the H1:H3 pairs, using the same procedure used for establishing H1:H2. There is no requirement to wait for the H1:H2 pairs to reach the full duplex state before creating the H1:H3 pairs. These pairs can also be established as Global Copy and later converted to Metro Mirror, or they can be established directly as Metro Mirror pairs.

After the creation of the H1:H2 and H1:H3 pairs, a two Metro Mirror topology exists.

As noted earlier, the order of establishing the pairs does not matter, and H1:H3 can be established before H1:H2.

Creation of MTIR pairs

As described in 3.2, “Multiple Target Incremental Resynchronization” on page 33, when the DS8870 detects that a Multiple Target PPRC configuration exists, it creates MTIR pairs between H2 and H3.

The creation of these MTIR pairs is initiated by the DS8870 at the primary H1 site by sending commands to the secondary H2 and H3 storage systems.

After establishing the PPRC paths and creating the pairs, the two Metro Mirror topology appears, as shown in Figure 4-2.

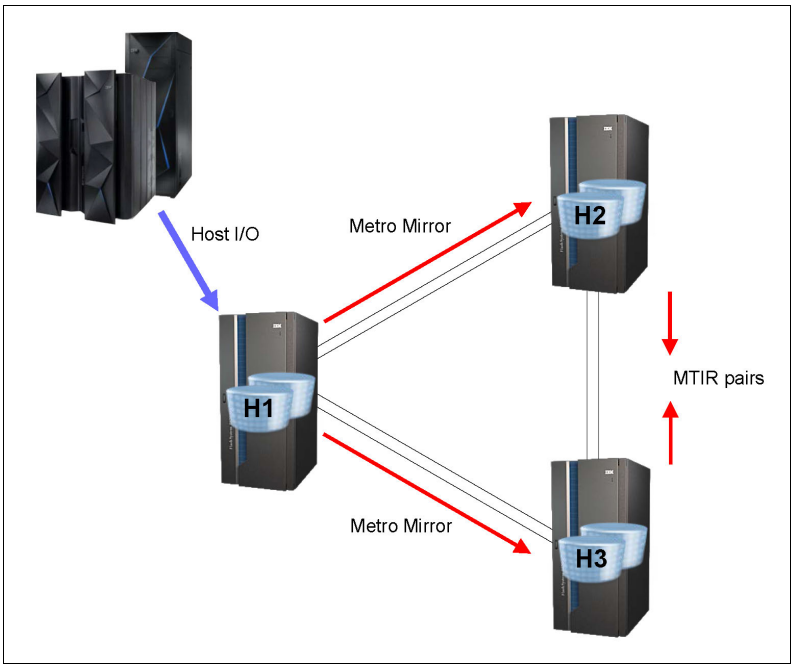

Figure 4-2 Two Metro Mirror pairs

Example 4-3 is an **lspprc** query to a volume, with the primary H1 showing that it has two Metro Mirror relationships, one to H2 and the other to H3. (The output has been truncated on the right in the example output for better readability.)

Example 4-3 lspprc for H1 showing two Metro Mirror pairs on each primary volume

```
dscli> lspprc -dev IBM.2107-75CZM21 5000-500f
ID                                        State       Reason Type         SourceLSS  ...
==========================================================================================================
IBM.2107-75CZM21/5000:IBM.2107-75CYK71/7000 Full Duplex -      Metro Mirror IBM.2107-75CZM21/50 ...
IBM.2107-75CZM21/5000:IBM.2107-75CYM31/6000 Full Duplex -      Metro Mirror IBM.2107-75CZM21/50 ...
...
IBM.2107-75CZM21/500F:IBM.2107-75CYK71/700F Full Duplex -      Metro Mirror IBM.2107-75CZM21/50 ...
IBM.2107-75CZM21/500F:IBM.2107-75CYM31/600F Full Duplex -      Metro Mirror IBM.2107-75CZM21/50 ...
```

Example 4-4 shows an **lspprc** query to the secondary H2, first without the **-multtgt** option and again with the **-multtgt** option to list the MTIR pairs. Notice that the MTIR pairs are shown as Global Copy suspended pairs in addition to being marked as *Multi-target Internal*.

Example 4-4 lspprc query to the secondary H2

```
dscli> lspprc -dev IBM.2107-75CYM31 6000-600f
ID                                        State              Reason Type         SourceLSS
==========================================================================================================
IBM.2107-75CZM21/5000:IBM.2107-75CYM31/6000 Target Full Duplex -      Metro Mirror IBM.2107-75CZM21/50...
...
IBM.2107-75CZM21/500F:IBM.2107-75CYM31/600F Target Full Duplex -      Metro Mirror IBM.2107-75CZM21/50...

dscli> lspprc -dev IBM.2107-75CYM31 -multtgt 6000-600f
ID                                        State              Reason                Type         SourceLSS
Timeout (secs) Critical Mode First Pass Status
==========================================================================================================================
IBM.2107-75CYM31/6000:IBM.2107-75CYK71/7000 Suspended          Multi-target Internal Global Copy  IBM.2107-75CYM31/60...
IBM.2107-75CYM31/600F:IBM.2107-75CYK71/700F Suspended          Multi-target Internal Global Copy  IBM.2107-75CYM31/60...
IBM.2107-75CZM21/5000:IBM.2107-75CYM31/6000 Target Full Duplex -                     Metro Mirror IBM.2107-75CZM21/50...
IBM.2107-75CZM21/500F:IBM.2107-75CYM31/600F Target Full Duplex -                     Metro Mirror IBM.2107-75CZM21/50...
```

If there is no connectivity between H1 and H2 at the time that H1:H3 is created, the creation of the MTIR pairs from H2:H3 will be delayed, as explained in 3.2.2, “Delayed creation of MTIR pairs” on page 34.

4.3 Outage at H2 or H3

An outage at one of the secondary sites, either H2 or H3, can cause the PPRC pairs to that site to suspend. Because of the design, a suspension of one relationship on a volume does not affect other relationships on the volume, so the second Metro Mirror pairs will remain active.

For example, Figure 4-3 shows a two Metro Mirror configuration where there has been an error at the H3 site, causing the H1:H3 pairs to suspend.

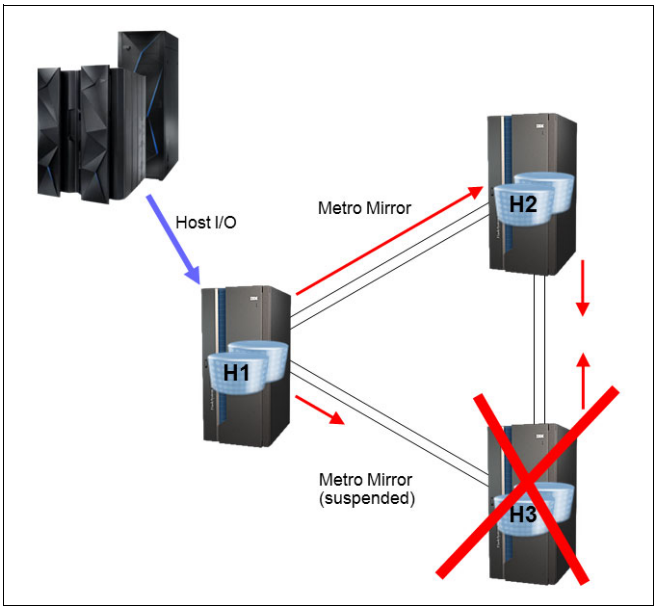

Figure 4-3 Failure at H3, causing a suspension of H1:H3

A query of the H1 volumes shows that the pairs to H2 are still active and the pairs to H3 are suspended, as shown in Example 4-5. The H1:H2 pairs remain full-duplex.

Example 4-5 lspprc showing pairs H1:H2 and H3

```
dscli> lspprc -dev IBM.2107-75CZM21 5000-500f
ID                                        State       Reason      Type         SourceLSS
==========================================================================================================
IBM.2107-75CZM21/5000:IBM.2107-75CYK71/7000 Suspended   Host Source Metro Mirror IBM.2107-75CZM21/50...
IBM.2107-75CZM21/5000:IBM.2107-75CYM31/6000 Full Duplex -           Metro Mirror IBM.2107-75CZM21/50...
...
IBM.2107-75CZM21/500F:IBM.2107-75CYK71/700F Suspended   Host Source Metro Mirror IBM.2107-75CZM21/50...
IBM.2107-75CZM21/500F:IBM.2107-75CYM31/600F Full Duplex -           Metro Mirror IBM.2107-75CZM21/50...
```

When the remote site is recovered, the suspended pairs can be resumed, restoring the Multiple Target PPRC configuration. The MTIR pairs are not used for this case.

Example 4-6 shows an example of a DS CLI **resumepprc** command.

Example 4-6 Resume suspended pairs H1:H3

```
dscli> resumepprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYK71 -type mmir 5000-500f:7000-700f
CMUC00158I resumepprc: Remote Mirror and Copy volume pair IBM.2107-75CZM21/5000:IBM.2107-75CYK71/7000
relationship successfully resumed. This message is being returned before the copy completes.
...
CMUC00158I resumepprc: Remote Mirror and Copy volume pair IBM.2107-75CZM21/500F:IBM.2107-75CYK71/700F
relationship successfully resumed. This message is being returned before the copy completes.
```

4.4 Outage at H1 site

This section describes the steps that are taken in the case of a failure or planned outage at the local production H1 site.

In a two Metro Mirror configuration, a swap can be performed to either H2 or H3. It is the user's choice to decide which of the sites is the preferred site for the swap. In this example, the swap is being performed to H2.

4.4.1 Recover at H2

The first step in the process is to move production to H2 so that the host application I/O can continue. Either HyperSwap or a traditional restart can be used.

The steps involved in this process follow.

Freeze H1:H2 and H1:H3

Freeze commands for H1:H2 and H1:H3 remove the PPRC paths and suspend all PPRC pairs for H1:H2 and H1:H3. A separate freeze command is required for each LSS to LSS PPRC relationship. With the use of consistency groups, the freeze command creates consistent data at the remote H2 and H3 sites by using extended long busy or queue full to temporarily queue dependent writes.

Depending on the type of failure at H1, the freeze commands might not complete successfully, and some or all H1 pairs can remain full duplex. If the H1 site has completely failed, these commands might not have even been issued or executed.

Example 4-7 Freeze PPRC Group H1:H2 and H1:H3

```
dscli> freezepprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYM31 50:60
CMUC00161I freezepprc: Remote Mirror and Copy consistency group IBM.2107-75CZM21/50:IBM.2107-75CYM31/60
successfully created.

dscli> freezepprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYK71 50:70
CMUC00161I freezepprc: Remote Mirror and Copy consistency group IBM.2107-75CZM21/50:IBM.2107-75CYK71/70
successfully created.
```

Unfreeze and run H1:H2 and H1:H3.

When all freeze commands have completed, the secondary H2 and H3 sites contain a consistent copy of data. The unfreeze or run command (also known as a consistency group created command) removes the extended long busy or queue full condition at the local H1 site. An unfreeze command is required for each LSS to LSS PPRC relationship.

As with the freeze command, depending on the type of failure at H1, the unfreeze commands might not complete successfully.

Example 4-8 shows an example of a DS CLI unfreeze command.

Example 4-8 Consistency Group Created, H1:H2 and H1:H3

```
dscli> unfreezepprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYM31 50:60
CMUC00198I unfreezepprc: Remote Mirror and Copy pair IBM.2107-75CZM21/50:IBM.2107-75CYM31/60 successfully
thawed.
dscli> unfreezepprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYK71 50:70
CMUC00198I unfreezepprc: Remote Mirror and Copy pair IBM.2107-75CZM21/50:IBM.2107-75CYK71/70 successfully
thawed.
```

Failover H2:H1 and H3:H1

The failover command for H2:H1 coverts the H2 volumes to suspended Metro Mirror primary volumes whose secondary volumes are H1. Similarly, the failover H3:H1 converts the H3 volumes to suspended primary volumes where secondary volumes are H1.

Example 4-9 and Example 4-10 show DS CLI examples of the failover commands for H2:H1 and H3:H1, respectively.

Example 4-9 Failover H2:H1

```
dscli> failoverpprc -dev IBM.2107-75CYM31 -remotedev IBM.2107-75CZM21 -type mmir 6000-600f:5000-500f
CMUC00196I failoverpprc: Remote Mirror and Copy pair IBM.2107-75CYM31/6000:IBM.2107-75CZM21/5000
successfully reversed.
...
CMUC00196I failoverpprc: Remote Mirror and Copy pair IBM.2107-75CYM31/600F:IBM.2107-75CZM21/500F
successfully reversed.
```

Example 4-10 Failover H3:H1

```
dscli> failoverpprc -dev IBM.2107-75CYK71 -remotedev IBM.2107-75CZM21 -type mmir 7000-700f:5000-500f
CMUC00196I failoverpprc: Remote Mirror and Copy pair IBM.2107-75CYK71/7000:IBM.2107-75CZM21/5000
successfully reversed.
...
CMUC00196I failoverpprc: Remote Mirror and Copy pair IBM.2107-75CYK71/700F:IBM.2107-75CZM21/500F
successfully reversed.
```

At this point, each volume on H1, H2, and H3 is a suspended primary volume pointing to the other two sites, as shown in Figure 4-4 on page 54. In other words:

- H1 is a suspended primary to H2
- H1 is a suspended primary to H3

- H2 is a suspended primary to H1
- H2 is a suspended primary to H3 (internally created MTIR pair)

- H3 is a suspended primary to H1
- H3 is a suspended primary to H2 (internally created MTIR pair)

Example 4-11 shows that the H1:H2 and H1:H3 pairs are suspended due to the freeze commands. If H1 is inaccessible due to the failure, the freeze commands might not be possible and a query might not be possible..

Example 4-11 lspprc for H1 pairs

```
dscli> lspprc -dev IBM.2107-75CZM21 5000-500f
ID                                        State     Reason Type         SourceLSS
==========================================================================================
IBM.2107-75CZM21/5000:IBM.2107-75CYK71/7000 Suspended Freeze Metro Mirror IBM.2107-75CZM21/50 ...
IBM.2107-75CZM21/5000:IBM.2107-75CYM31/6000 Suspended Freeze Metro Mirror IBM.2107-75CZM21/50 ...
...
IBM.2107-75CZM21/500F:IBM.2107-75CYK71/700F Suspended Freeze Metro Mirror IBM.2107-75CZM21/50 ...
IBM.2107-75CZM21/500F:IBM.2107-75CYM31/600F Suspended Freeze Metro Mirror IBM.2107-75CZM21/50 ...
```

Example 4-12 shows that the H2:H1 pairs are suspended from the failover H2:H1 command. The **-multtgt** option is used to also include the H2:H3 MTIR pairs in the output. As previously noted, the example shows that the suspended states H2:H3 and H3:H2 are created internally, not as the result of a user command.

Example 4-12 lspprc for H2 pairs

```
dscli> lspprc -dev IBM.2107-75CYM31 -multtgt 6000-600f
ID                                        State     Reason                Type         SourceLSS
==============================================================================================================
IBM.2107-75CYM31/6000:IBM.2107-75CYK71/7000 Suspended Multi-target Internal Global Copy  IBM.2107-75CYM31/60 ...
IBM.2107-75CYM31/6000:IBM.2107-75CZM21/5000 Suspended Host Source           Metro Mirror IBM.2107-75CYM31/60 ...
...
IBM.2107-75CYM31/600F:IBM.2107-75CYK71/700F Suspended Multi-target Internal Global Copy  IBM.2107-75CYM31/60 ...
IBM.2107-75CYM31/600F:IBM.2107-75CZM21/500F Suspended Host Source           Metro Mirror IBM.2107-75CYM31/60 ...
```

Example 4-13 shows that the H3:H1 pairs are suspended from the failover H3:H1 command. The **-multtgt** option is used to include the H3:H2 MTIR pairs in the output.

Example 4-13 lspprc for H3 pairs

```
dscli> lspprc -dev IBM.2107-75CYK71 -multtgt 7000-700f
ID                                        State     Reason                Type         SourceLSS
==============================================================================================================
IBM.2107-75CYK71/7000:IBM.2107-75CYM31/6000 Suspended Multi-target Internal Global Copy  IBM.2107-75CYK71/70 ...
IBM.2107-75CYK71/7000:IBM.2107-75CZM21/5000 Suspended Host Source           Metro Mirror IBM.2107-75CYK71/70 ...
...
IBM.2107-75CYK71/700F:IBM.2107-75CYM31/600F Suspended Multi-target Internal Global Copy  IBM.2107-75CYK71/70 ...
IBM.2107-75CYK71/700F:IBM.2107-75CZM21/500F Suspended Host Source           Metro Mirror IBM.2107-75CYK71/70 ...
```

Resume host systems at H2

Host systems can be restarted at H2. If HyperSwap is used, the host I/O is directed to H2.

At this point, the host systems are running at H2, and the configuration appears as in Figure 4-4.

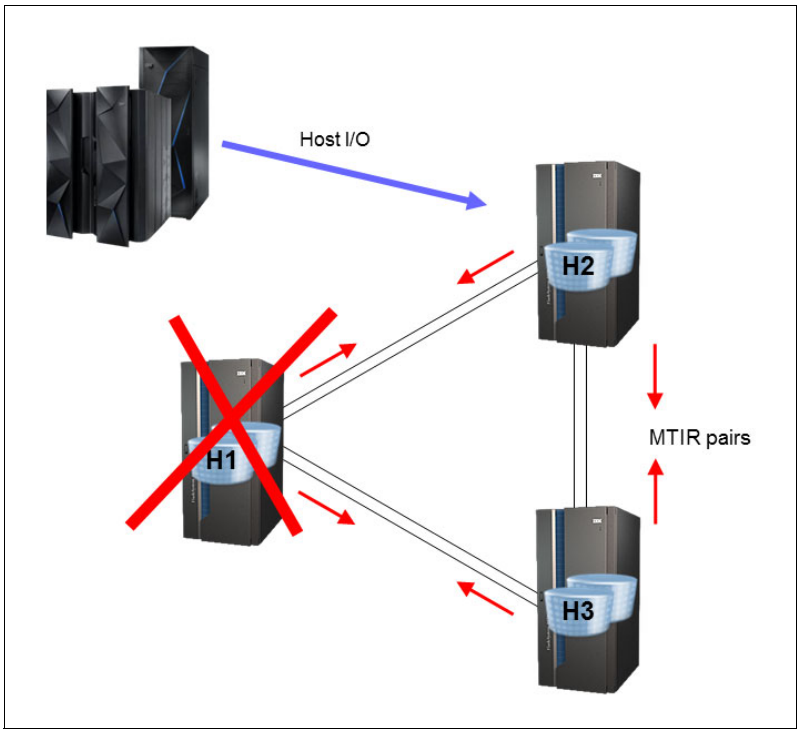

Figure 4-4 After recovery at H2

4.4.2 Start replication H2:H3

The Incremental Resynchronization capability is used to start replication H2:H3. With the change recording function of Multiple Target PPRC, only the tracks that are potentially out of synchronization between H2 and H3 are transferred and a full copy of data is not required.

Failback H2:H3

A failback command for H2:H3 merges the change recording bitmap files for H2 and H3 and initiates the transfer of the out-of-synchronization data.

Example 4-14 shows a sample DS CLI for the failback command.

Example 4-14 Failback H2:H3

```
dscli> failbackpprc -dev IBM.2107-75CYM31 -remotedev IBM.2107-75CYK71 -type mmir 6000-600f:7000-700f
CMUC00197I failbackpprc: Remote Mirror and Copy pair IBM.2107-75CYM31/6000:IBM.2107-75CYK71/7000
successfully failed back.
...
CMUC00197I failbackpprc: Remote Mirror and Copy pair IBM.2107-75CYM31/600F:IBM.2107-75CYK71/700F
successfully failed back.
```

By transferring only the out-of-synchronization data, the pairs reach the full duplex state, as shown in Example 4-15.

Example 4-15 lspprc showing full duplex after incremental resync

```
dscli> lspprc -dev IBM.2107-75CYM31 -remotedev IBM.2107-75CYK71 6000-600f:7000-700f
ID                                        State       Reason Type         SourceLSS           Timeout
(secs) Critical Mode First Pass Status
=======================================================================================================
IBM.2107-75CYM31/6000:IBM.2107-75CYK71/7000 Full Duplex -      Metro Mirror IBM.2107-75CYM31/60 ...
...
IBM.2107-75CYM31/600F:IBM.2107-75CYK71/700F Full Duplex -      Metro Mirror IBM.2107-75CYM31/60 ...
```

At this point in the process, the H2:H3 pairs have been resumed. The state of the configuration is that shown in Figure 4-5.

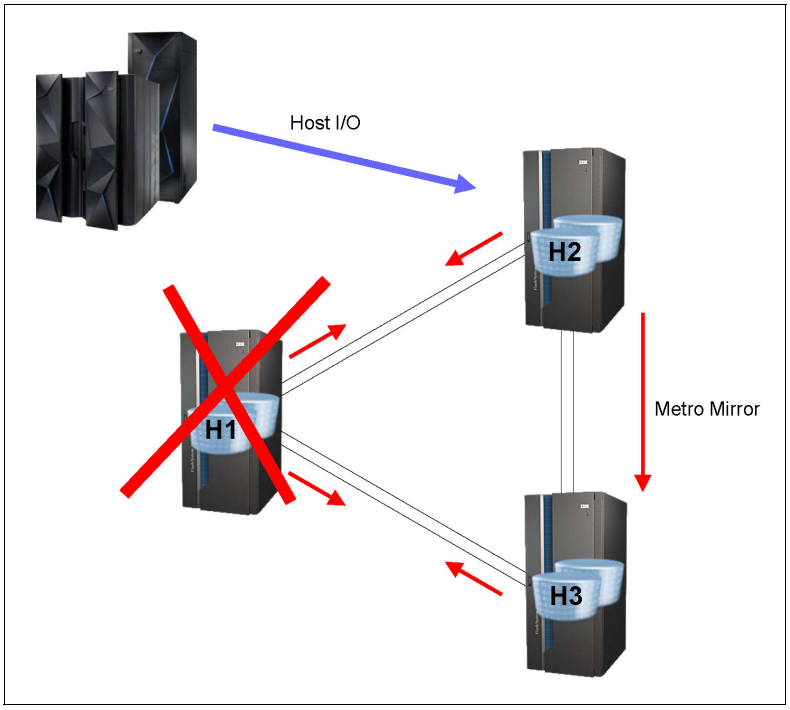

Figure 4-5 After failback H2:H3 with Incremental Resynchronization

The data protection of Metro Mirror has now been restored after a failure of the primary site. Without Multiple Target PPRC, host access could have been restored at site H2, but there would have been no mirroring protection until the H1 site was recovered and the H2:H1 pairs resumed.

4.4.3 H1 recovered

When site H1 is recovered, mirroring can be resumed to H2 to restore the Multiple Target PPRC configuration.

Establish paths H1:H2 and H1:H3

The paths for H1:H2 and H1:H3 were removed by the freeze commands. After H1 is recovered, these paths need to be reestablished. The paths are not required until the time when mirroring is later resumed for H1:H2 and H1:H3, but establishing them as soon as H1 is recovered helps to ensure that they are available when needed.

The paths are reestablished by using the same command used to originally create the paths, as shown again in Example 4-16.

Example 4-16 Reestablish paths H1:H2 and H1:H3

```
dscli> mkpprcpath -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYM31 -remotewwnn 5005076305FFD71E -srclss 50
-tgtlss 60 -consistgrp I0201:I0232 I0234:I0332
CMUC00149I mkpprcpath: Remote Mirror and Copy path 50:60 successfully established.

dscli> mkpprcpath -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYK71 -remotewwnn 5005076305FFD71A -srclss 50
-tgtlss 70 -consistgrp I0201:I0032 I0234:I0306
CMUC00149I mkpprcpath: Remote Mirror and Copy path 50:70 successfully established.
```

Failback H2:H1

A failback command for H2:H1 resumes the H2:H1 Metro Mirror pairs, as shown in Example 4-17.

Example 4-17 Failback H2:H1

```
dscli> failbackpprc -dev IBM.2107-75CYM31 -remotedev IBM.2107-75CZM21 -type mmir 6000-600f:5000-500f
CMUC00197I failbackpprc: Remote Mirror and Copy pair IBM.2107-75CYM31/6000:IBM.2107-75CZM21/5000
successfully failed back.
...
CMUC00197I failbackpprc: Remote Mirror and Copy pair IBM.2107-75CYM31/600F:IBM.2107-75CZM21/500F
successfully failed back.
```

When these pairs reach full duplex, the Multiple Target PPRC configuration has been restored, as shown in Figure 4-6.

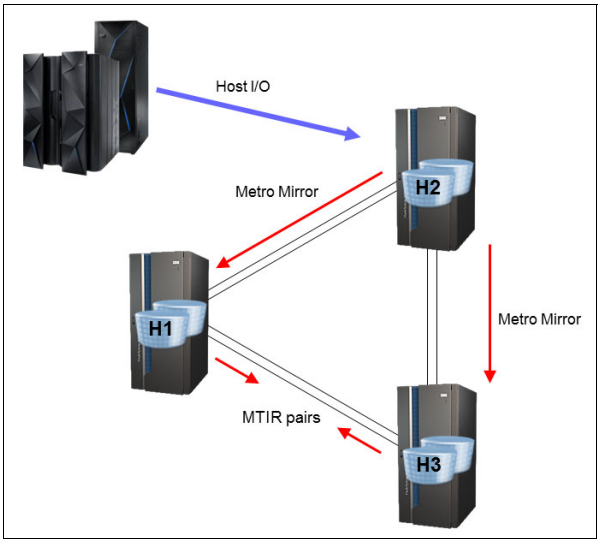

Figure 4-6 Multiple Target PPRC topology restored

The suspended pairs between H1 and H3 are automatically converted by the DS8870 storage system to become MTIR pairs. The change recording mechanism is started to track the out-of-synchronization data between H1 and H3. This is now the same topology as the initial state, with two Metro Mirror pairs. However, the host systems are now running at H2 rather than H1.

Depending upon the specific circumstances, for example if H1 and H2 are local to each other, there might be no preference for running at one site versus the other. In other situations, there might be reasons to prefer that the host systems run at site H1. Section 4.5, “Return production to H1” on page 57 describes a method to return to the original H1 site.

4.5 Return production to H1

Returning to site H1 is similar to a move to H2 previously described in 4.4, “Outage at H1 site” on page 51. The main difference is that this is a planned move and the original move to H2 might have been the result of an error event at H1.

4.5.1 Move host systems to H1

The first step in the process is to move production back to H1. As with the move to H2, either HyperSwap or a traditional restart can be used. The steps involved in this process follow.

Query out of-synchronization tracks

Query the out-of-synchronization tracks from H1:H3 and H3:H1. As described in 3.2.3, “MTIR change recording” on page 34, after the H2:H1 pair reaches the full-duplex state, there can be a period of time before the out-of-sync tracks counts is reduced from the full volume amount. It is important to verify that these counts have dropped from the full volume amount to ensure that failback H1:H3 for the MTIR pairs does not perform a full copy of data.

Freeze H2:H1 and H2:H3

Freeze commands for H2:H1 and H2:H3 remove the PPRC paths and suspend the H1:H2 and H1:H3 PPRC pairs. A separate freeze command is required for each LSS to LSS PPRC relationship. The use of consistency groups ensures consistent data at H1 and H2.

Unfreeze and run H2:H1 and H2:H3

When all freeze commands have completed, the secondary H1 and H3 sites contain a consistent copy of data. The unfreeze or run command (also known as a consistency group created command) removes the extended long busy or queue-full condition at the primary H2 site. An unfreeze command is required for each LSS to LSS PPRC relationship.

Example 4-18 shows an example of the DS CLI commands.

Example 4-18 Freeze and unfreeze H2:H1 and H2:H3

```
dscli> freezepprc -dev IBM.2107-75CYM31 -remotedev IBM.2107-75CZM21 60:50
CMUC00161I freezepprc: Remote Mirror and Copy consistency group IBM.2107-75CYM31/60:IBM.2107-75CZM21/50
successfully created.

dscli> freezepprc -dev IBM.2107-75CYM31 -remotedev IBM.2107-75CYK71 60:70
CMUC00161I freezepprc: Remote Mirror and Copy consistency group IBM.2107-75CYM31/60:IBM.2107-75CYK71/70
successfully created.

dscli> unfreezepprc -dev IBM.2107-75CYM31 -remotedev IBM.2107-75CZM21 60:50
CMUC00198I unfreezepprc: Remote Mirror and Copy pair IBM.2107-75CYM31/60:IBM.2107-75CZM21/50 successfully
thawed.

dscli> unfreezepprc -dev IBM.2107-75CYM31 -remotedev IBM.2107-75CYK71 60:70
CMUC00198I unfreezepprc: Remote Mirror and Copy pair IBM.2107-75CYM31/60:IBM.2107-75CYK71/70 successfully
thawed.
```

Failover H1:H2 and H3:H2

The failover command for H1:H2 coverts the H1 volumes to suspended Metro Mirror primary volumes whose secondary volumes are H2. Similarly, the failover H3:H2 converts the H3 volumes to suspended primary volumes whose secondary volumes are H2.

Example 4-19 shows a DS CLI example of these failover commands.

Example 4-19 Failover H1:H2 and H1:H3

```
dscli> failoverpprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYM31 -type mmir 5000-500f:6000-600f
CMUC00196I failoverpprc: Remote Mirror and Copy pair IBM.2107-75CZM21/5000:IBM.2107-75CYM31/6000
successfully reversed.
...
CMUC00196I failoverpprc: Remote Mirror and Copy pair IBM.2107-75CZM21/500F:IBM.2107-75CYM31/600F
successfully reversed.

dscli> failoverpprc -dev IBM.2107-75CYK71 -remotedev IBM.2107-75CYM31 -type mmir 7000-700f:6000-600f
CMUC00196I failoverpprc: Remote Mirror and Copy pair IBM.2107-75CYK71/7000:IBM.2107-75CYM31/6000
successfully reversed.
...
CMUC00196I failoverpprc: Remote Mirror and Copy pair IBM.2107-75CYK71/700F:IBM.2107-75CYM31/600F
successfully reversed.
```

Resume host systems at H1

The host systems can now be restarted at H1. If HyperSwap is used, the host systems are switched to H1.

At this point, the host systems are running at H1. The configuration appears as in Figure 4-7.

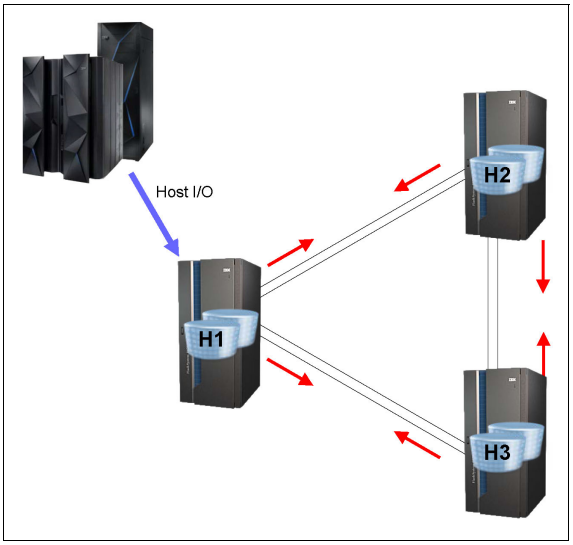

Figure 4-7 Move production back to H1

Each volume on H1, H2, and H3 is a suspended primary volume pointing to the other two sites.

4.5.2 Start H1:H2 and H1:H3 replication

Because this is a planned swap, there is no need to wait for the secondary system to be recovered. The replication H1:H2 and H1:H3 can be immediately resumed.

Establish H2:H1 and H2:H3 paths

The paths from H2:H1 and H2:H3 were removed by the previous freeze commands. reestablishing these paths as soon as possible helps to ensure that they are available when needed.

Example 4-20 shows sample DS CLI commands to establish the paths for H2:H1 and H2:H3.

Example 4-20 Reestablish paths for H2:H1 and H2:H3

```
dscli> mkpprcpath -dev IBM.2107-75CYM31 -remotedev IBM.2107-75CZM21 -remotewwnn 5005076305FFD75A -srclss 60
-tgtlss 50 -consistgrp I0232:I0201 I0332:I0234
CMUC00149I mkpprcpath: Remote Mirror and Copy path 60:50 successfully established.

dscli> mkpprcpath -dev IBM.2107-75CYM31 -remotedev IBM.2107-75CYK71 -remotewwnn 5005076305FFD71A -srclss 60
-tgtlss 70 -consistgrp I0300:I0300 I0330:I0330
CMUC00149I mkpprcpath: Remote Mirror and Copy path 60:70 successfully established.
```

Failback H1:H2 and H1:H3

A failback command for H1:H2 and H1:H3 starts replication for all pairs. Example 4-21 shows sample commands for the failback.

Example 4-21 Failback H1:H2 and H1:H3

```
dscli> failbackpprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYM31 -type mmir 5000-500f:6000-600f
CMUC00197I failbackpprc: Remote Mirror and Copy pair IBM.2107-75CZM21/5000:IBM.2107-75CYM31/6000
successfully failed back.
...
CMUC00197I failbackpprc: Remote Mirror and Copy pair IBM.2107-75CZM21/500F:IBM.2107-75CYM31/600F
successfully failed back.

dscli> failbackpprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYK71 -type mmir 5000-500f:7000-700f
CMUC00197I failbackpprc: Remote Mirror and Copy pair IBM.2107-75CZM21/5000:IBM.2107-75CYK71/7000
successfully failed back.
...
CMUC00197I failbackpprc: Remote Mirror and Copy pair IBM.2107-75CZM21/500F:IBM.2107-75CYK71/700F
successfully failed back.
```

The suspended pairs between H2 and H3 are changed to MTIR pairs.

The original configuration is now restored, as show in Figure 4-8 on page 60.

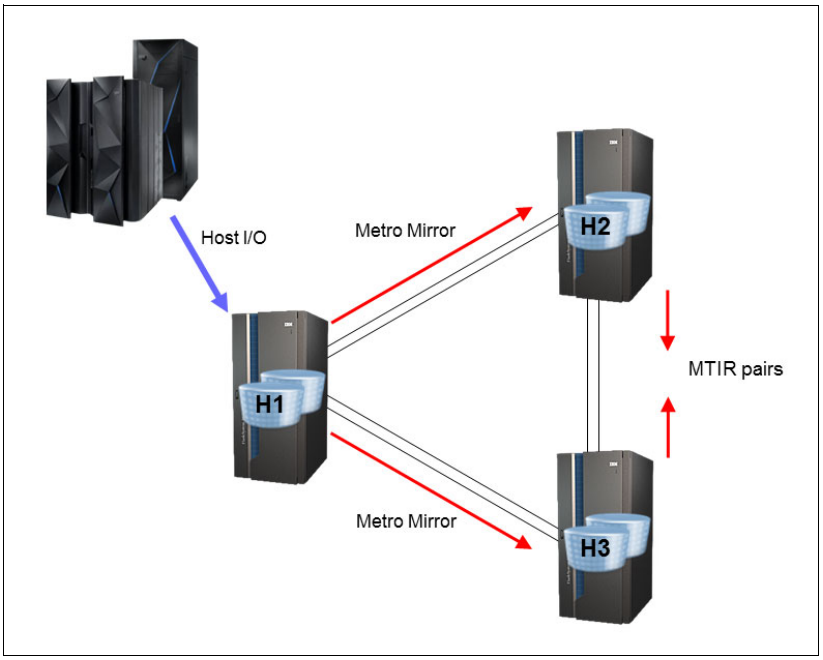

Figure 4-8 Return to original configuration

A query of the H1 volumes shows that all pairs have returned to the full-duplex state, as shown in Example 4-22.

Example 4-22 lspprc showing H1:H2 and H1:H3 in full duplex state

```
dscli> lspprc -dev IBM.2107-75CZM21 5000-500f
ID                                        State       Reason Type         SourceLSS
==========================================================================================
IBM.2107-75CZM21/5000:IBM.2107-75CYK71/7000 Full Duplex -      Metro Mirror IBM.2107-75CZM21/50 ...
IBM.2107-75CZM21/5000:IBM.2107-75CYM31/6000 Full Duplex -      Metro Mirror IBM.2107-75CZM21/50 ...
...
IBM.2107-75CZM21/500F:IBM.2107-75CYK71/700F Full Duplex -      Metro Mirror IBM.2107-75CZM21/50 ...
IBM.2107-75CZM21/500F:IBM.2107-75CYM31/600F Full Duplex -      Metro Mirror IBM.2107-75CZM21/50 ...
```

4.6 Tivoli Storage Productivity Center for Replication

The IBM Tivoli Storage Productivity Center for Replication provides an easy and reliable user interface to manage large copy service configurations, including a Multiple Target PPRC configuration with two Metro Mirror pairs. It simplifies and automates the tasks for recovery after both planned and unplanned outages.

This paper does not attempt to describe the operation or use of Tivoli Storage Productivity Center for Replication but gives a few examples to show the power and simplicity of this product in a Multiple Target PPRC configuration.

Figure 4-9 on page 61 is a sample of the GUI interface showing Metro Mirror pairs H1:H2 and H1:H3, including icons indicating that both relationships are HyperSwap-capable.

The role pairs displayed at the bottom of the window show the internal MTIR pairs that are created between the H2 and H3 secondary volumes.

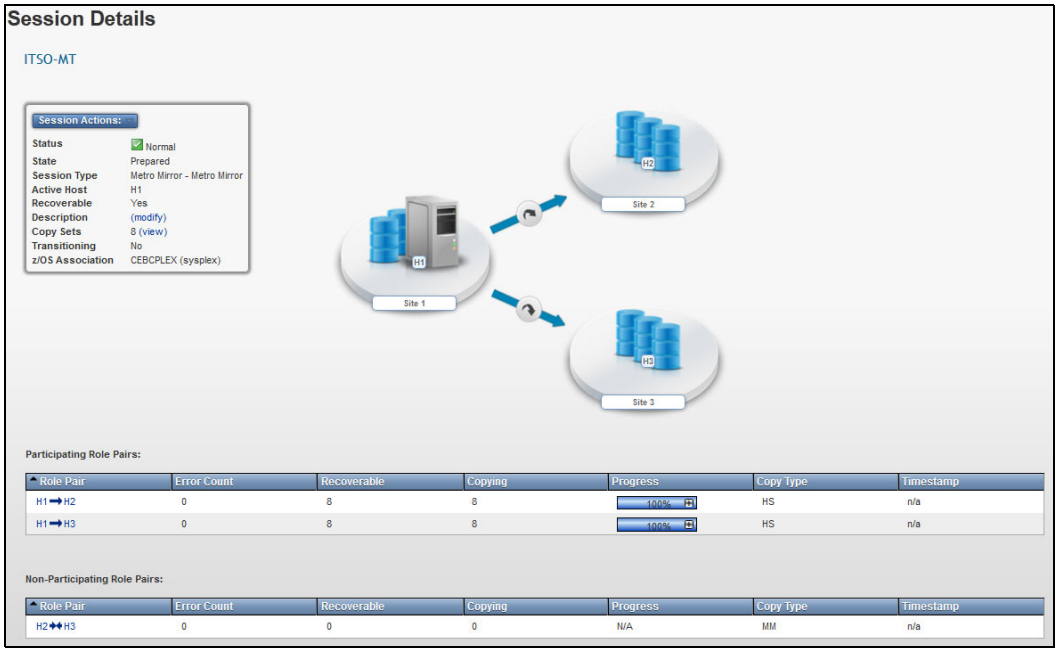

Figure 4-9 Tivoli Storage Productivity Center for Replication

Complex actions can be performed in a simple and intuitive manner. For example, a planned HyperSwap can be initiated by selecting **Commands** and then **HyperSwap**, as shown in Figure 4-10.

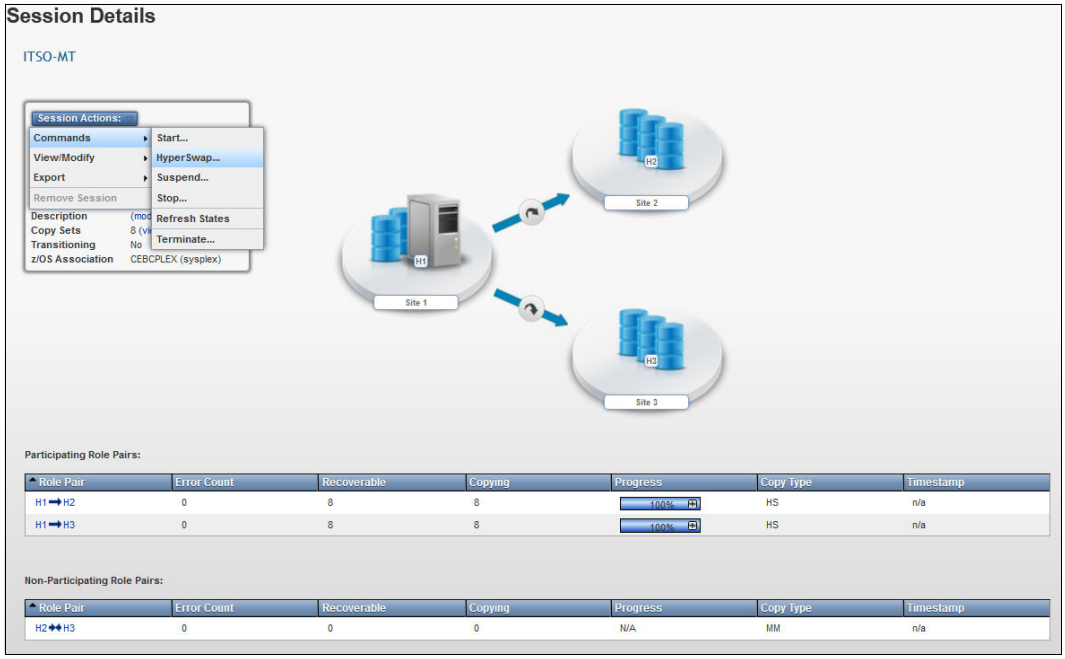

Figure 4-10 Selecting HyperSwap

With just a few simple and intuitive actions, Tivoli Storage Productivity Center for Replication can HyperSwap the activity either of the remote sites. After moving production to one of the secondary sites, replication can be started to either or both of the other sites.

As an example, Figure 4-11 shows the results after a HyperSwap to H2 and restarting replication for both H2:H1 and H2:H3.

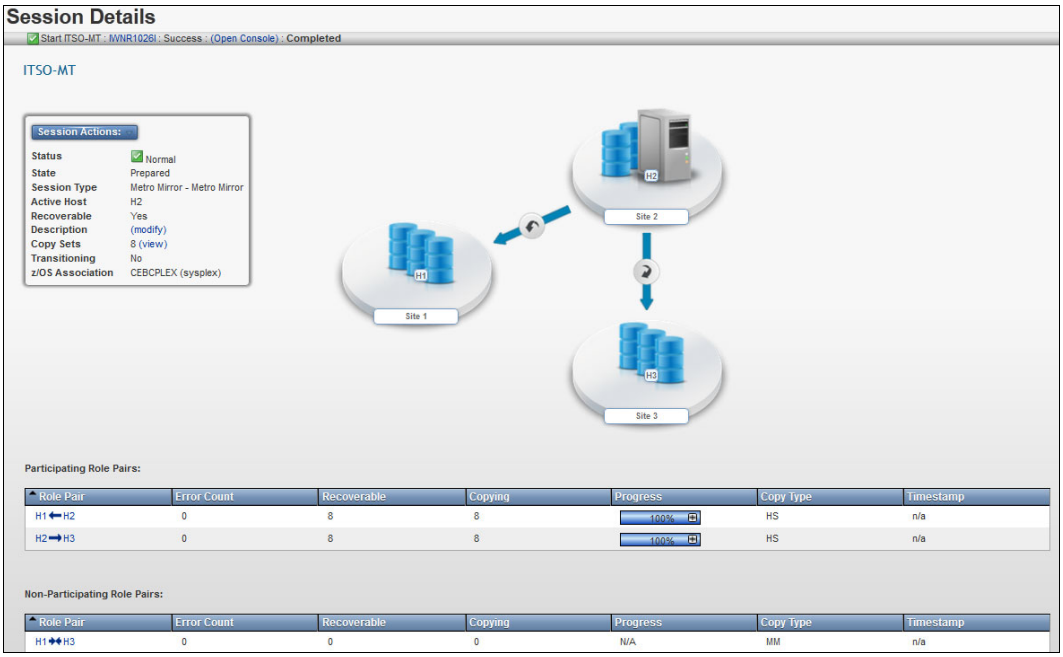

Figure 4-11 After HyperSwap to H2 and restarting H2:H1, H2:H3

5

Implementing a Metro Mirror and Global Mirror topology

This chapter describes the creation of a Multiple Target PPRC topology with Metro Mirror and Global Mirror. It covers the following topics:

- Overview of a Metro Mirror and Global Mirror topology
- Creating a Metro Mirror and Global Mirror topology
- Outage at H3
- Outage at H2
- Outage at H1

5.1 Overview of a Metro Mirror and Global Mirror topology

A topology with Metro Mirror and Global Mirror provides both a local high availability capability and, at the same time, a long-distance disaster recovery capability. In such a configuration, data is synchronously mirrored to one secondary site and is asynchronously mirrored to a separate disaster recovery site.

Figure 5-1 shows Metro Mirror H1:H2 and Global Mirror H1:H3.

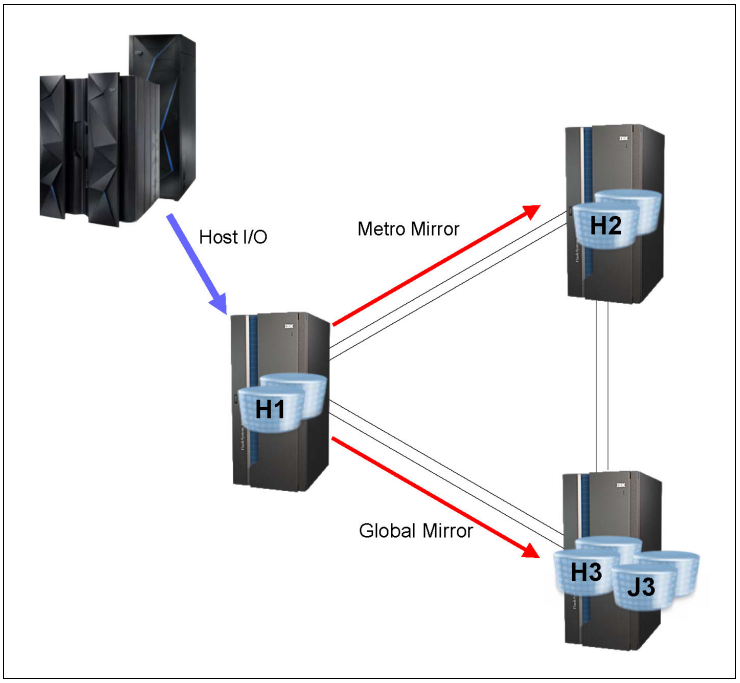

Figure 5-1 Metro Mirror and Global Mirror

These are the basic steps of creating a Metro Mirror H1:H2 and Global Mirror H1:H3 topology:

1. Establish all required paths.
2. Establish the H1:H2 Metro Mirror pairs.
3. Establish Global Mirror H1:H3.
   a. Establish H1:H3 Global Copy pairs.
   b. Create FlashCopy H3:J3.
   c. Create and start the Global Mirror session for H1:H3.

The scenarios in this chapter describe recovery actions after an outage at one of the sites. This can be either a failure condition or a planned event for testing or maintenance.

To simplify the descriptions, these sites are often referred to as though they were a single volume, but it must be understood that the terms refer to the entire set of volumes being mirrored by PPRC. For example, “Failover H2:H1” means to failover all of the H2 volumes to their corresponding H1 volumes.

5.2 Creating a Metro Mirror and Global Mirror topology

This section describes the steps in creating a Multiple Target PPRC configuration with one Metro Mirror and one Global Mirror relationship on each primary volume.

5.2.1 Terms used in examples

For the examples in this section, the following terms are used:

- **H1** is the current primary site where the production applications are running.
- **H2** is the Metro Mirror secondary site to which H1 is mirroring data.
- **H3** is the Global Mirror secondary site to which H1 is mirroring data asynchronously.

Table 5-1 identifies the DS8870 storage controllers used in the examples in this chapter.

Note: The volume range on each DS8870 is different in these examples only to help clarify the different sites used. It is not a requirement that they be different.

Table 5-1 Identifications used in DS CLI examples

Site	Role	Dev	WWNN	Volume range
H1	Current primary	IBM.2107-75CZM21	5005076305FFD75A	5000 - 500F
H2	Metro Mirror secondary	IBM.2107-75CYM31	5005076305FFD71E	6000 - 600F
H3	Global Mirror secondary	IBM.2107-75CYK71	5005076305FFD71A	7000 - 700F
J3	Global Mirror FlashCopy targets	IBM.2107-75CYK71	5005076305FFD71A	7200 - 720F

5.2.2 Establish PPRC Paths

PPRC paths are required before the PPRC pairs can be established. It is best to establish all paths that might be needed so that they are available at the time when they are required.

The procedure for establishing these paths is similar to that for the two Metro Mirror topology that is described in 4.2.2, “Establish PPRC paths” on page 47, so it is not repeated here.

5.2.3 Create H1:H2 Metro Mirror pairs

After establishing all of the required PPRC paths, create the Metro Mirror volume pairs H1:H2. These pairs can be initially created as asynchronous Global Copy pairs and later converted to Metro Mirror, or they can be created directly as Metro Mirror pairs. The procedure for establishing the Metro Mirror pairs is the same as that described in 4.2.3, “Create H1:H2 Metro Mirror pairs” on page 48, for the two Metro Mirror topology.

5.2.4 Create H1:H3 Global Copy pairs

Example 5-1 shows a sample DS CLI command to establish the Global Copy pairs.

Example 5-1 Create H1:H3 Global Copy pairs

```
dscli> mkpprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYK71 -type gcp 5000-500f:7000-700f
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship
IBM.2107-75CZM21/5000:IBM.2107-75CYK71/7000 successfully created.
...
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship
IBM.2107-75CZM21/500F:IBM.2107-75CYK71/700F successfully created.
```

Creation of MTIR pairs

As described in 3.2, “Multiple Target Incremental Resynchronization” on page 33, when the primary DS8870 detects that a Multiple Target PPRC configuration exists, it creates MTIR pairs between H2 and H3. These pairs are created regardless of the type of PPRC relationships for H1:H2 and H1:H3.

After establishing the PPRC paths and creating the Metro Mirror and Global Copy pairs, the topology appears as shown in Figure 5-2.

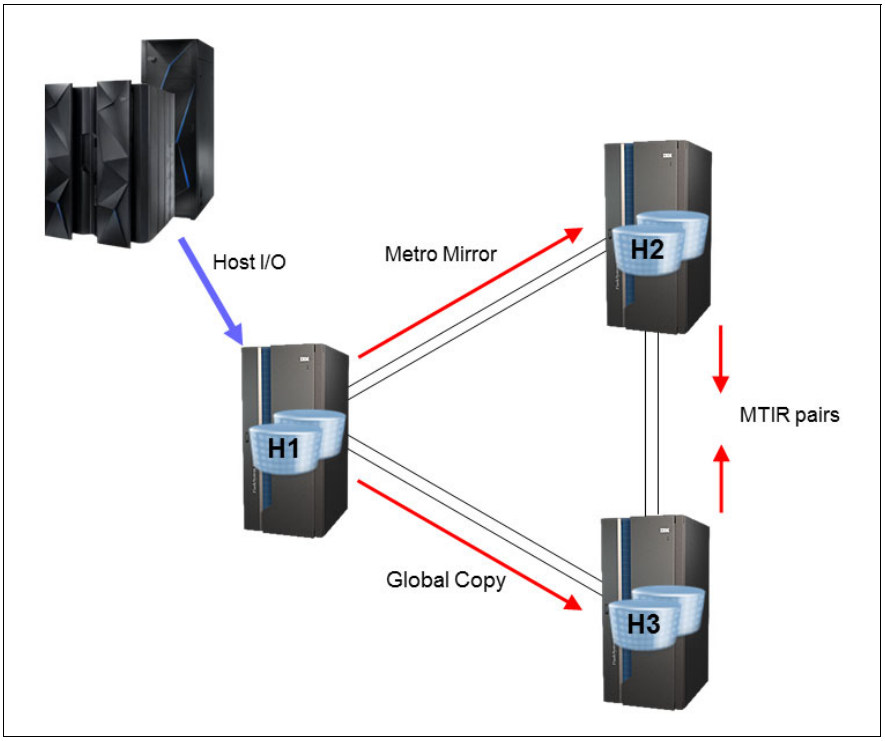

Figure 5-2 Metro Mirror and Global Copy

5.2.5 Create FlashCopy H3:J3

The next step is to establish the FlashCopy relationships from the Global Copy secondary volumes to create the journal volumes used by Global Mirror. These FlashCopy relationships are created with the following attributes:

Inhibit target writes Protects the FlashCopy target volume from being modified by host writes. Only the Global Copy primary can write to the target volume.

Start change recording Starts recording which tracks have changed on either volume in the FlashCopy pair.

Persistent Keeps the FlashCopy relationship until it is explicitly or implicitly terminated.

No copy Does not perform a background copy of the data from the FlashCopy source volume to the target.

Example 5-2 shows a sample DS CLI command to create FlashCopy relationships for Global Mirror to use.

Example 5-2 Create FlashCopy H3:J3

```
dscli> mkflash -dev IBM.2107-75CYK71 -tgtinhibit -record -persist -nocp 7000-700f:7200-720f
CMUC00137I mkflash: FlashCopy pair IBM.2107-75CYK71/7000:IBM.2107-75CYK71/7200 successfully created.
...
CMUC00137I mkflash: FlashCopy pair IBM.2107-75CYK71/700F:IBM.2107-75CYK71/720F successfully created.
```

The creation of the FlashCopy relationships can be verified as shown in Example 5-3.

Example 5-3 List FlashCopy relationships

```
dscli> lsflash -dev IBM.2107-75CYK71 7000-700f
ID                                        SrcLSS             SequenceNum Timeout ActiveCopy Recording
Persistent Revertible SourceWriteEnabled TargetWriteEnabled BackgroundCopy
==========================================================================================================
IBM.2107-75CYK71/7000:IBM.2107-75CYK71/7200 IBM.2107-75CYK71/70 0           60      Disabled   Enabled
Enabled    Disabled   Enabled            Disabled           Disabled
...
IBM.2107-75CYK71/700F:IBM.2107-75CYK71/720F IBM.2107-75CYK71/70 0           60      Disabled   Enabled
Enabled    Disabled   Enabled            Disabled           Disabled
```

5.2.6 Create and start Global Mirror session

The Global Mirror session is created. Example 5-4 shows sample DS CLI commands to create a session, start the session, and then query to show the session status. In this simple example, there are no external subordinates.

Example 5-4 Create Global Mirror session

```
dscli> mksession -volpair 5000-500f:7000-700f -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYK71
-lss 50 01
CMUC00145I mksession: Session 01 opened successfully.

dscli> mkgmir -dev IBM.2107-75CZM21 -lss 50 -session 01
CMUC00162I mkgmir: Global Mirror for session 01 successfully started.

dscli> lsgmir -dev IBM.2107-75CZM21 -session 01 50
SessionID MasterID         ID State   %Success CGtime

0x01      IBM.2107-75CZM21 50 Running 100      11/17/2014 23:24:42 CETv
```

At this point, there is a Metro Mirror and Global Mirror Multiple Target PPRC configuration, which is illustrated in Figure 5-3.

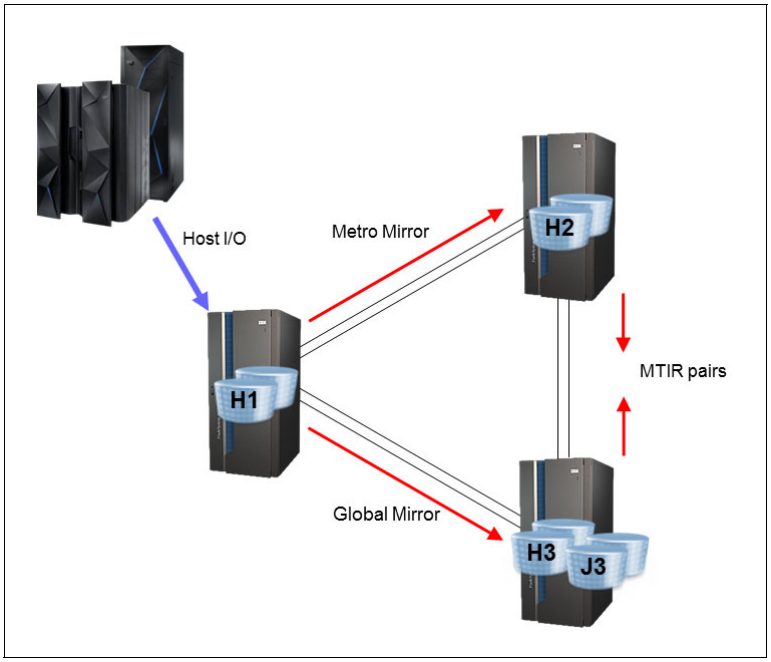

Figure 5-3 Metro Mirror and Global Mirror

5.3 Outage at H3

A failure or planned outage at the Global Mirror remote site H3 can cause the Global Copy pairs to suspend and the Global Mirror session to stop forming consistency groups. The Metro Mirror pairs to H2 continue to run and provide protection in case of a outage at H1. When site H3 is recovered, Global Mirror H1:H3 can be resumed to restore the Global Mirror protection.

5.4 Outage at H2

A failure or planned outage at the secondary site H2 can cause the Metro Mirror pairs H1:H2 to suspend. Global Mirror to H3 continues to provide disaster recovery protection in case of a failure at H1. When site H2 is recovered, Metro Mirror H1:H2 can be resumed to restore the Metro Mirror protection.

5.5 Outage at H1

The sections that follow describe the steps that are taken in the case of a failure or planned outage at the local production H1 site.

In a Metro Mirror and Global Mirror configuration a swap to the Metro Mirror secondary site H2 can be performed.

5.5.1 Recover at H2

The first step in the process is to move production to H2 so that the host application I/O can continue.

Either HyperSwap or a traditional restart can be used. The steps involved in this process are `freeze` H1:H2 and H1:H3.

Freeze commands for H1:H2 and H1:H3 remove the PPRC paths and suspend all PPRC pairs for H1:H2 and H1:H3. A separate `freeze` command is required for each LSS-to-LSS PPRC relationship.

With the use of consistency groups, the `freeze` command creates consistent data at the remote H2 site by using extended long *busy* or *queue full* to temporarily queue dependent writes.

Depending on the type of failure at H1, the `freeze` commands might not complete successfully and some or all H1 pairs can remain full-duplex.

Example 5-5 Freeze PPRC Group H1:H2 and H1:H3

```
dscli> freezepprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYM31 50:60
CMUC00161I freezepprc: Remote Mirror and Copy consistency group IBM.2107-75CZM21/50:IBM.2107-75CYM31/60
successfully created.

dscli> freezepprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYK71 50:70
CMUC00161I freezepprc: Remote Mirror and Copy consistency group IBM.2107-75CZM21/50:IBM.2107-75CYK71/70
successfully created.
```

Unfreeze and run H1:H2 and H1:H3

When all freeze commands have completed, the secondary H2 and H3 sites contain a consistent copy of data. The unfreeze or run command (also known as a *consistency group created* command) removes the extended long busy or queue-full condition at the local H1 site. An unfreeze command is required for each LSS-to-LSS PPRC relationship.

As with the freeze command, depending on the type of failure at H1, the unfreeze commands might not complete successfully.

Example 5-6 shows an example of a DS CLI unfreeze command.

Example 5-6 Consistency Group Created, H1:H2 and H1:H3

```
dscli> unfreezepprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYM31 50:60
CMUC00198I unfreezepprc: Remote Mirror and Copy pair IBM.2107-75CZM21/50:IBM.2107-75CYM31/60 successfully
thawed.

dscli> unfreezepprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYK71 50:70
CMUC00198I unfreezepprc: Remote Mirror and Copy pair IBM.2107-75CZM21/50:IBM.2107-75CYK71/70 successfully
thawed.
```

Failover H2:H1 and H3:H1

The failover command for H2:H1 coverts the H2 volumes to suspended Metro Mirror primary volumes whose secondary volumes are H1. Similarly, the failover H3:H1 converts the H3 volumes to suspended primary Global Copy volumes whose secondary volumes are H1.

Example 5-7 on page 70 and Example 5-8 on page 70 show DS CLI examples of the failover commands for H2:H1 and H3:H1, respectively.

Example 5-7 Failover H2:H1

```
dscli> failoverpprc -dev IBM.2107-75CYM31 -remotedev IBM.2107-75CZM21 -type mmir 6000-600f:5000-500f
CMUC00196I failoverpprc: Remote Mirror and Copy pair IBM.2107-75CYM31/6000:IBM.2107-75CZM21/5000
successfully reversed.
...
CMUC00196I failoverpprc: Remote Mirror and Copy pair IBM.2107-75CYM31/600F:IBM.2107-75CZM21/500F
successfully reversed.
```

Example 5-8 Failover H3:H1

```
dscli> failoverpprc -dev IBM.2107-75CYK71 -remotedev IBM.2107-75CZM21 -type gcp 7000-700f:5000-500f
CMUC00196I failoverpprc: Remote Mirror and Copy pair IBM.2107-75CYK71/7000:IBM.2107-75CZM21/5000
successfully reversed.
...
CMUC00196I failoverpprc: Remote Mirror and Copy pair IBM.2107-75CYK71/700F:IBM.2107-75CZM21/500F
successfully reversed.
```

At this point, each volume on H1, H2, and H3 is a suspended primary volume pointing to the other two sites as shown in Figure 5-4 on page 71. Specifically, the pair states are:

- H1 suspended primary to H2
- H1 suspended primary to H3

- H2 suspended primary to H1
- H2 suspended primary to H3

- H3 suspended primary to H1
- H3 suspended primary to H2

Example 5-9 shows that the H1:H2 and H1:H3 pairs are suspended due to the freeze commands. If H1 is inaccessible due to a failure, the freeze commands might not be possible and a query might not be possible, either.

Example 5-9 lspprc for H1 pairs

```
dscli> lspprc -dev IBM.2107-75CZM21 5000-500f
ID                                        State     Reason Type         SourceLSS
==========================================================================================
IBM.2107-75CZM21/5000:IBM.2107-75CYK71/7000 Suspended Freeze Global Copy  IBM.2107-75CZM21/50 ...
IBM.2107-75CZM21/5000:IBM.2107-75CYM31/6000 Suspended Freeze Metro Mirror IBM.2107-75CZM21/50 ...
...
IBM.2107-75CZM21/500F:IBM.2107-75CYK71/700F Suspended Freeze Global Copy  IBM.2107-75CZM21/50 ...
IBM.2107-75CZM21/500F:IBM.2107-75CYM31/600F Suspended Freeze Metro Mirror IBM.2107-75CZM21/50 ...
```

Example 5-10 shows that the H2:H1 pairs are suspended from the **failover** H2:H1 command. The **-multtgt** option is used to include the H2:H3 MTIR pairs in the output.

Example 5-10 lspprc for H2 pairs

```
dscli> lspprc -dev IBM.2107-75CYM31 -multtgt 6000-600f
ID                                        State     Reason               Type         SourceLSS
==============================================================================================
IBM.2107-75CYM31/6000:IBM.2107-75CYK71/7000 Suspended Multi-target Internal Global Copy  IBM.2107-75CYM31/60 ...
IBM.2107-75CYM31/6000:IBM.2107-75CZM21/5000 Suspended Host Source          Metro Mirror IBM.2107-75CYM31/60 ...
...
IBM.2107-75CYM31/600F:IBM.2107-75CYK71/700F Suspended Multi-target Internal Global Copy  IBM.2107-75CYM31/60 ...
IBM.2107-75CYM31/600F:IBM.2107-75CZM21/500F Suspended Host Source          Metro Mirror IBM.2107-75CYM31/60 ...
```

Example 5-11 shows that the H3:H1 pairs are suspended from the **failover** H3:H1 command. The **-multtgt** option is used to also include the H3:H2 MTIR pairs in the output.

Example 5-11 lspprc for H3 pairs

```
dscli> lspprc -dev IBM.2107-75CYK71 -multtgt 7000-700f
ID                                        State     Reason               Type        SourceLSS
==========================================================================================================
IBM.2107-75CYK71/7000:IBM.2107-75CYM31/6000 Suspended Multi-target Internal Global Copy IBM.2107-75CYK71/70 ...
IBM.2107-75CYK71/7000:IBM.2107-75CZM21/5000 Suspended Host Source          Global Copy IBM.2107-75CYK71/70 ...
...
IBM.2107-75CYK71/700F:IBM.2107-75CYM31/600F Suspended Multi-target Internal Global Copy IBM.2107-75CYK71/70 ...
IBM.2107-75CYK71/700F:IBM.2107-75CZM21/500F Suspended Host Source          Global Copy IBM.2107-75CYK71/70 ...
```

Resume host systems at H2

Host systems can be restarted at H2. If HyperSwap is used, the host I/O is switched to H2.

At this point, the host systems are running at H2, and the configuration appears as in Figure 5-4.

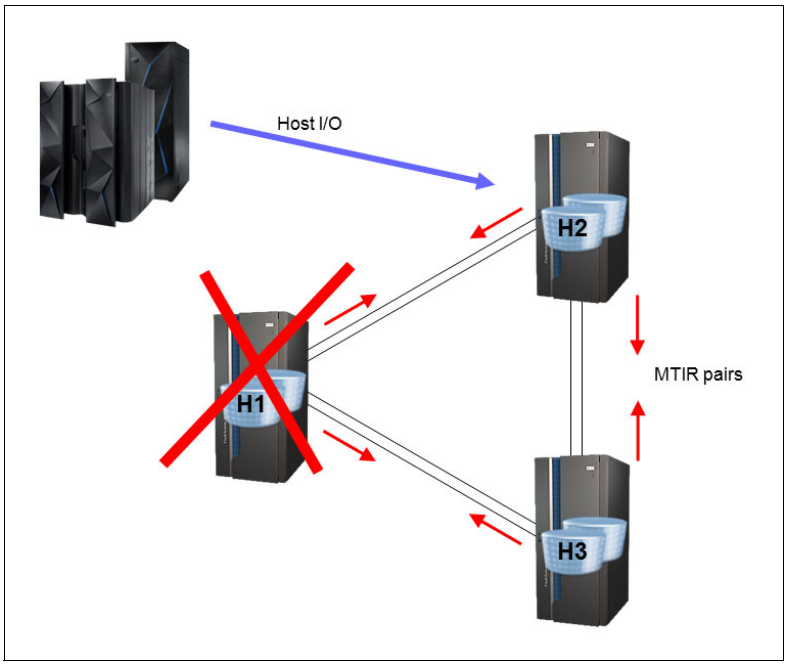

Figure 5-4 After recovery at H2

Start replication H2:H3

The Incremental Resynchronization capability is used to start replication H2:H3. With the change recording function of Multiple Target PPRC, only the tracks that are potentially out of synchronization between H2 and H3 are transferred, and a full copy of data is not required.

This Incremental Resynchronization capability operates the same for Global Copy and Global Mirror pairs as it does for Metro Mirror. Even though H2 was a Metro Mirror secondary and H3 was a Global Mirror scenario, the Incremental Resynchronization is still possible..

Failback H2:H3

A failback command for H2:H3 merges the change recording bitmap files from H2 and H3 and initiates the transfer of the out-of-synchronization data.

Example 5-12 shows a sample DS CLI for the failback command

Example 5-12 Failback H2:H3

```
dscli> failbackpprc -dev IBM.2107-75CYM31 -remotedev IBM.2107-75CYK71 -type gcp 6000-600f:7000-700f
CMUC00197I failbackpprc: Remote Mirror and Copy pair IBM.2107-75CYM31/6000:IBM.2107-75CYK71/7000
successfully failed back.
...
CMUC00197I failbackpprc: Remote Mirror and Copy pair IBM.2107-75CYM31/600F:IBM.2107-75CYK71/700F
successfully failed back.
```

At this point in the process, the H2:H3 pairs have been resumed. The state of the configuration is as shown in Figure 5-5, with Global Copy now active from H2 to H3.

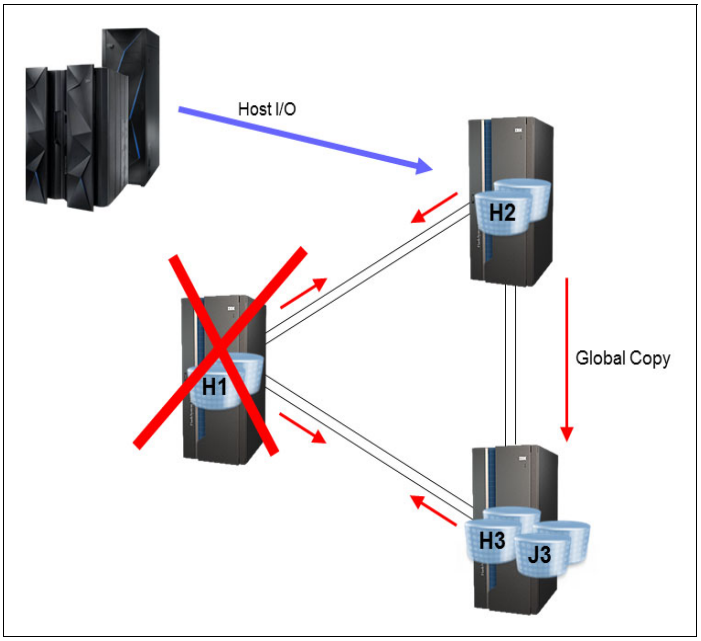

Figure 5-5 After failback H2:H3 with Incremental Resynchronization

Start Global Mirror H2:H3

Global Copy is running for H2:H3 and transferring data. The FlashCopy journal volumes at J3 contain the last consistency group that was formed before the failure at H1.

Example 5-13 shows the next step, which is to start Global Mirror from H2:H3 and resume the formation of Global Mirror consistency groups.

Example 5-13 Start Global Mirror H2:H3

```
dscli> mksession -volpair 6000-600f:7000-700f -dev IBM.2107-75CYM31 -remotedev IBM.2107-75CYK71 -lss 60 01
CMUC00145I mksession: Session 01 opened successfully.

dscli> mkgmir -dev IBM.2107-75CYM31 -lss 60 -session 01
CMUC00162I mkgmir: Global Mirror for session 01 successfully started.
```

Global Mirror is started for H2:H3 by using the same FlashCopy volumes on H3, and consistency group formation is resumed, as shown in Example 5-14.

Example 5-14 Query showing consistency group formation

```
dscli> lsgmir -dev IBM.2107-75CYM31 -session 01 60
SessionID MasterID         ID State   %Success CGtime
=====================================================================
0x01      IBM.2107-75CYM31 60 Running 100      mm/dd/yyyy hh:mm:ss TZ
```

5.5.2 H1 recovered

When site H1 is recovered, mirroring can be resumed to H2 to restore the Multiple Target PPRC configuration.

Pause or remove Global Mirror H1:H3

H1 has Global Mirror running for H1:H3. Because the Global Copy pairs H1:H3 are suspended, Global Mirror is unable to form consistency groups. The Global Mirror session can be either paused or removed.

Example 5-15 shows a DS CLI command to pause the Global Mirror session, followed by a query showing that the state is now paused.

Example 5-15 Pause Global Mirror H1:H3

```
dscli> pausegmir -dev IBM.2107-75CZM21 -lss 50 -session 01
CMUC00163I pausegmir: Global Mirror for session 01 successfully paused.
...
dscli> lsgmir -dev IBM.2107-75CZM21 -session 01 50
SessionID MasterID         ID State  %Success CGtime
=====================================================================
0x01      IBM.2107-75CZM21 50 Paused 95       11/18/2014 10:11:31 CET
```

Alternatively, the Global Mirror H1:H3 session can be ended by removing the Global Copy pairs from the session and then removing the session, as shown in Example 5-16.

Example 5-16 Stop Global Mirror H1:H3

```
dscli> chsession -volpair 5000-500f:7000-700f -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYK71 -lss 50
-action remove 01
CMUC00147I chsession: Session 01 successfully modified.

dscli> rmsession -dev IBM.2107-75CZM21 -lss 50 -quiet 01
CMUC00146I rmsession: Session 01 closed successfully.
```

Establish paths H1:H2 and H1:H3

The paths H1:H2 and H1:H3 were removed by the freeze commands issued in 5.5.1, “Recover at H2”. After H1 is recovered, these paths need to be reestablished. The paths are not required until the time when mirroring is later resumed from H1:H2 and H1:H3, but establishing them as soon as H1 is recovered helps to ensure that they are available when needed.

The paths are reestablished by using the same commands that were used to create the paths.

Failback H2:H1

A failback command for H2:H1 resumes the H2:H1 Metro Mirror pairs, as shown in Example 5-17.

Example 5-17 Failback H2:H1

```
dscli> failbackpprc -dev IBM.2107-75CYM31 -remotedev IBM.2107-75CZM21 -type mmir 6000-600f:5000-500f
CMUC00197I failbackpprc: Remote Mirror and Copy pair IBM.2107-75CYM31/6000:IBM.2107-75CZM21/5000
successfully failed back.
...
CMUC00197I failbackpprc: Remote Mirror and Copy pair IBM.2107-75CYM31/600F:IBM.2107-75CZM21/500F
successfully failed back.
```

Site H2 is now a Multiple Target PPRC primary with Metro Mirror H2:H1 and Global Mirror H2:H3. The suspended pairs that existed between the two secondary sites H1 and H3 are automatically converted by the DS8870 to MTIR pairs. Example 5-17 shows an **lspprc** command issued at H1, showing that it is now a Metro Mirror target of H2 and also has MTIR pairs to H3. Compare this to the query of H1 shown in Example 5-9 on page 70, which was issued before the failback H2:H1, to see how the DS8870 has created the MTIR pair relationships.

Example 5-18 lspprc at H1 showing H1:H3 MTIR pairs

```
dscli> lspprc -dev IBM.2107-75CZM21 -multtgt 5000-500f
ID                                        State              Reason                Type         SourceLSS
==========================================================================================================================
IBM.2107-75CYM31/6000:IBM.2107-75CZM21/5000 Target Full Duplex -                    Metro Mirror IBM.2107-75CYM31/60...
...
IBM.2107-75CYM31/600F:IBM.2107-75CZM21/500F Target Full Duplex -                    Metro Mirror IBM.2107-75CYM31/60...
IBM.2107-75CZM21/5000:IBM.2107-75CYK71/7000 Suspended          Multi-target Internal Global Copy  IBM.2107-75CZM21/50...
...
IBM.2107-75CZM21/500F:IBM.2107-75CYK71/700F Suspended          Multi-target Internal Global Copy  IBM.2107-75CZM21/50...
```

After the H2:H3 pairs reach full-duplex, the Multiple Target PPRC configuration has been restored, as shown in Figure 5-6.

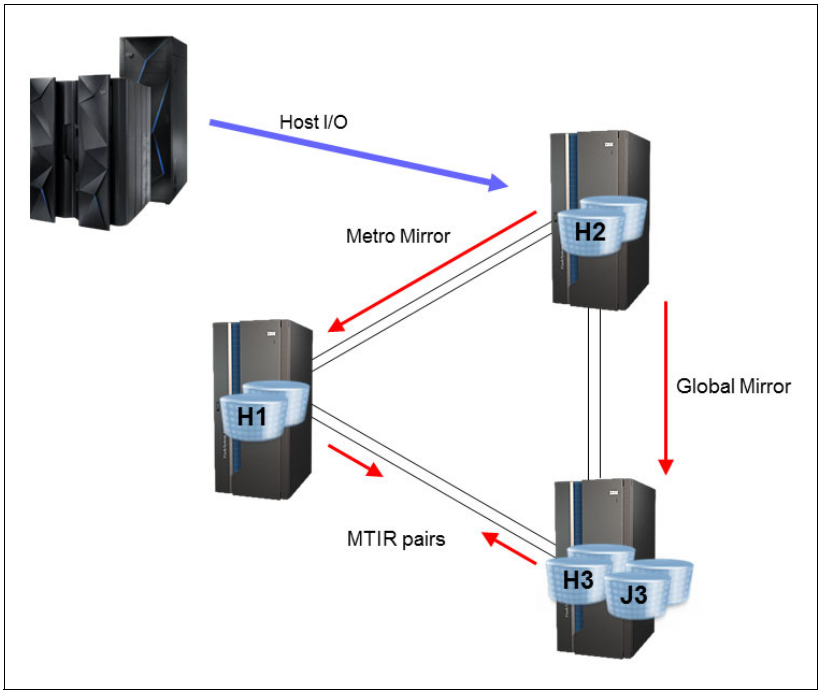

Figure 5-6 Metro Mirror and Global Mirror topology restored

This is now the same topology as the initial state with Metro Mirror and Global Mirror.

You can compare Figure 5-3 on page 68 with Figure 5-6 to see that it is the same general topology, but now the host systems are now running at H2 rather than H1. Depending upon the specific circumstances, for example if H1 and H2 are local to each other, there might be no preference for running at one site versus the other.

In other situations, there might be reasons to prefer that the host systems run at site H1. Returning to H1 involves the same sequence of steps as the move from H1 to H2. This move back will be a planned move, although the original move from H1 to H2 might have been performed due to a failure event at H1.

6

Implementing a Global Copy plus Global Mirror topology

This chapter describes the Global Copy plus Global Mirror topology and its use in providing a disaster recovery solution with three sites in three different regions separated by long distances. It covers these topics:

- Overview of Global Copy plus Global Mirror topology
- Creating a Global Copy plus Global Mirror topology
- Outage at Global Copy secondary H2
- Outage at Global Mirror H3
- Outage at Global Mirror H3

6.1 Overview of Global Copy plus Global Mirror topology

A topology consisting of a Global Copy and a Global Mirror relationship provides a disaster recovery solution with three sites in three different regions that are separated by long distances. Figure 6-1 shows an example of this topology, where H1:H2 is Global Copy and H1:H3 is Global Mirror. The J3 volumes are the Global Mirror FlashCopy journal volumes.

In the case of a failure at H3, the Global Copy H1:H2 pairs can be put into a Global Mirror session for H1:H2. Therefore, FlashCopy journal volumes at J2 are also required.

As with all types of Multiple Target PPRC topologies, there are Multiple Target Incremental Resync (MTIR) pairs between the two secondary sites, H2 and H3, as shown in Figure 6-1.

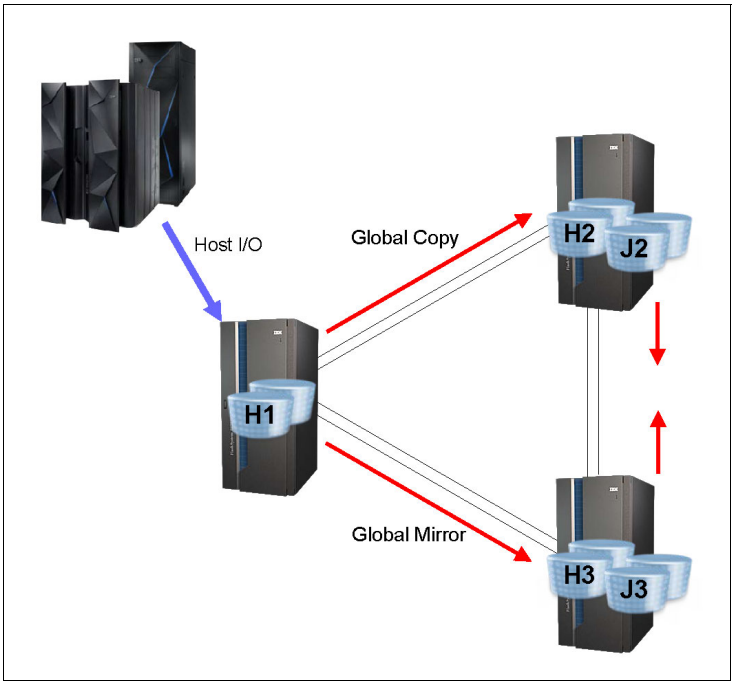

Figure 6-1 Global Copy plus Global Mirror topology

With this topology, a long-distance disaster recovery configuration remains available after an outage at any one of the three sites.

6.2 Creating a Global Copy plus Global Mirror topology

The subsections that follow explain these steps for creating a Global Copy H1:H2 and a Global Mirror H1:H3 topology:

1. Establish PPRC paths between all sites.
2. Establish the H1:H2 Global Copy pairs.
   a. Create FlashCopy H2:J2. (These are required to run Global Mirror H1:H2.)
3. Establish Global Mirror H1:H3.
   a. Establish H1:H3 Global Copy pairs.
   b. Create FlashCopy H3:J3.
   c. Create and start the Global Mirror session for H1:H3.

The details of creating these relationships are similar to creating the Metro Mirror and Global Mirror topology described in 5.2, “Creating a Metro Mirror and Global Mirror topology” on page 65. The one difference is that the H1:H2 relationships are Global Copy rather than Metro Mirror.

The scenarios in this section describe an outage at each of the three sites. Each of these could be either a failure condition or a planned event for testing or maintenance.

To simplify the descriptions, these sites are often referred to as though they were a single volume, but it must be understood that the terms refer to the entire set of volumes being mirrored by PPRC. For example, “Failover H2:H1” means to failover all of the H2 volumes to their corresponding H1 volumes.”

6.2.1 Establish PPRC paths

PPRC paths are required before the PPRC pairs can be established. It is best that paths among all three sites be established so that they are available when they are required.

The procedure for establishing these paths is similar to that for the two Metro Mirror topology that is described in 4.2.2, “Establish PPRC paths” on page 47.

6.2.2 Create H1:H2 Global Copy pairs

After establishing all of the required PPRC paths, create the Global Copy volume pairs H1:H2.

To provide for the possibility of starting Global Mirror H1:H2 in the event of an outage at the Global Mirror remote H3 site, Global Mirror FlashCopy journal volumes J2 are required. These FlashCopy relationships from H2:J2 are created with “inhibit target writes”, “start change recording”, “persistent”, and “no copy” options.

6.2.3 Create H1:H3 Global Mirror

Global Mirror H1:H3 is established in the same manner as described in 5.2, “Creating a Metro Mirror and Global Mirror topology” on page 65.

After both the H1:H2 and H1:H3 relationships have been created, the DS8870 storage system creates the MTIR pairs between H2 and H3, as described in 3.2, “Multiple Target Incremental Resynchronization” on page 33. At this point, there is Global Copy running from H1 to H2, and Global Mirror is running from H1 to H3.

The topology is now as shown in Figure 6-2 on page 80.

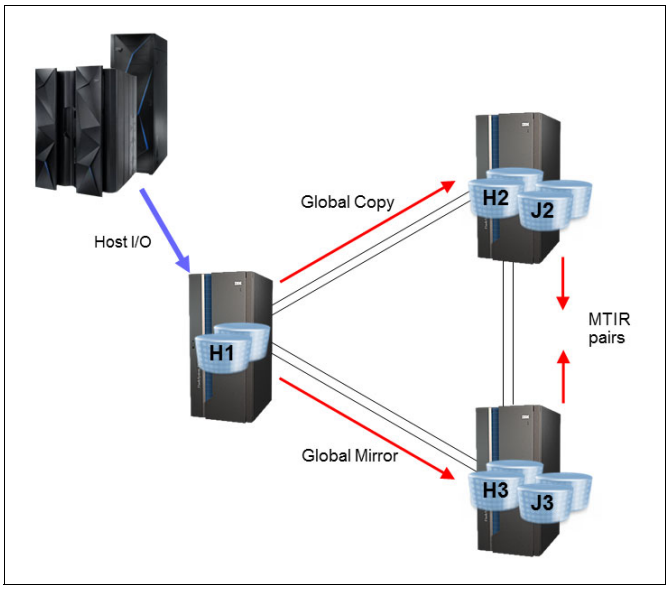

Figure 6-2 Global Copy and Global Mirror, including MTIR pairs

6.3 Outage at Global Copy secondary H2

An outage at the Global Copy secondary H2 site can cause the Global Copy H1:H2 pairs to suspend. Because of the design, a suspension of one relationship on a volume does not affect other relationships on the volume, so Global Mirror H1:H3 continues to run and provide a disaster recovery capability, as shown in Figure 6-3.

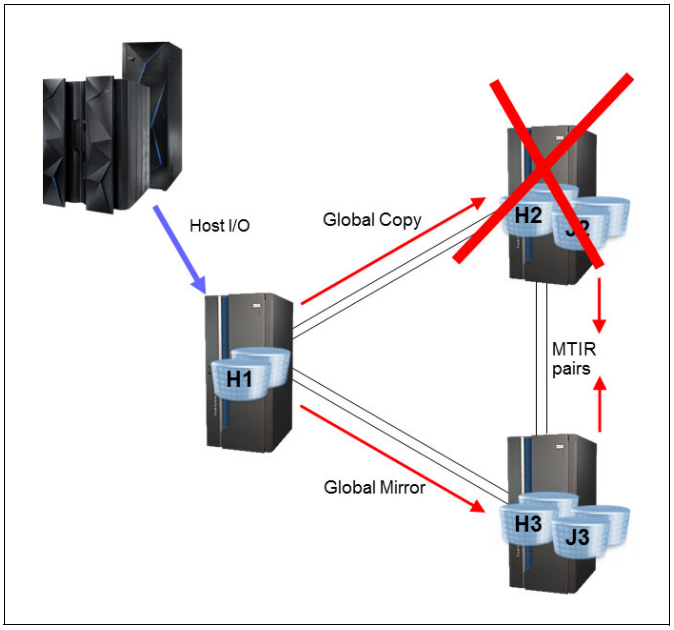

Figure 6-3 Outage at Global Copy secondary H2

When site H2 is recovered, the Global Copy H1:H2 pairs are resumed and data is copied from H1 to H2 again.

6.4 Outage at Global Mirror H3

An outage at the Global Mirror remote site H3 causes the loss of the Global Mirror disaster recovery capability.

The Global Copy H1:H2 pairs can be placed into a Global Mirror session to restore the disaster recovery capability.

6.4.1 Start Global Mirror H1:H2

Global Copy for H1:H2 remains active. The disaster recovery capability can be restored by converting the Global Copy H1:H2 to Global Mirror. This conversion requires removing the H1:H3 Global Copy pairs from the H1:H3 Global Mirror session, removing the GM H1:H3 session, and adding the Global Copy pairs into a new H1:H2 Global Mirror session.

Because the H1:H2 Global Copy pairs have already been running and transferring data from H1 to H2, there is no additional data transfer required. The H1:H2 pairs are already past the first round of copy, so they will be immediately joined to the session, and consistency groups will begin to be formed at the H2 site.

The resulting configuration is shown in Figure 6-4.

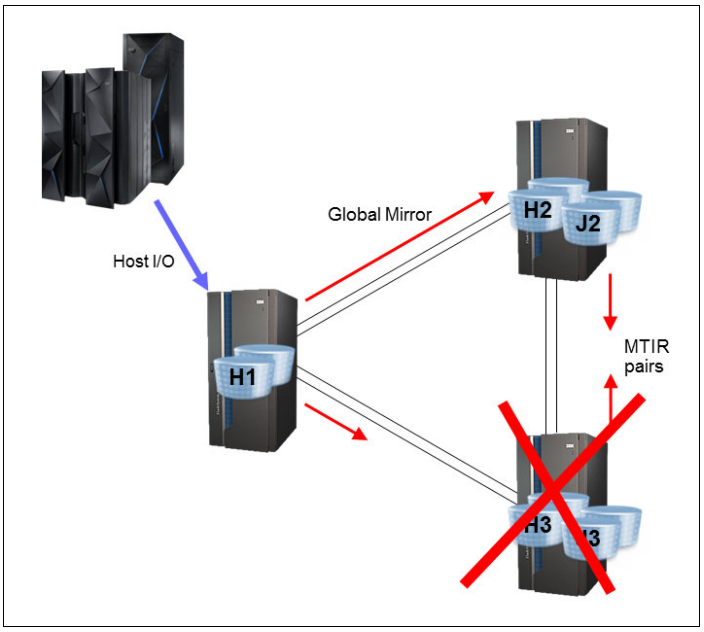

Figure 6-4 Global Mirror H1:H2 after outage at H3

6.4.2 H3 recovered

When the H3 site is recovered, the Global Copy H1:H3 pairs are resumed, resulting in the configuration shown in Figure 6-5 on page 82.

This is the same topology as in the initial starting case with Global Copy plus Global Mirror. However, Global Mirror is now running from H1 to H2 rather than from H1 to H3. If desired, the Global Mirror can be moved back to H1:H3 with a steps similar to those described previously.

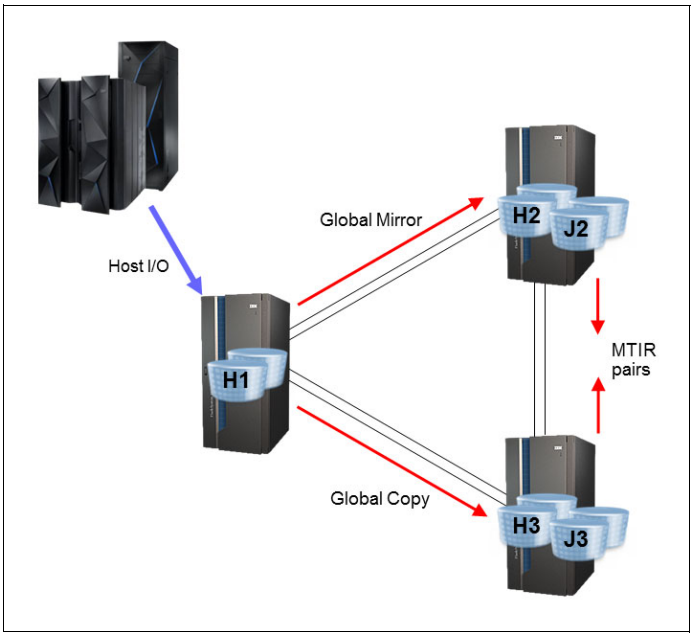

Figure 6-5 Resume Global Copy H1:H3

6.5 Outage at H1

In the case of an outage at the H1 site, a recovery at the remote Global Mirror H3 site is required. This requires a Global Mirror recovery and restart of production at the disaster recovery site, H3.

With Multiple Target PPRC, the MTIR pairs between H2 and H3 are used to perform an Incremental Resynchronization between the two sites.

The initial configuration is shown in Figure 6-6, where there has been a failure at the H1 site.

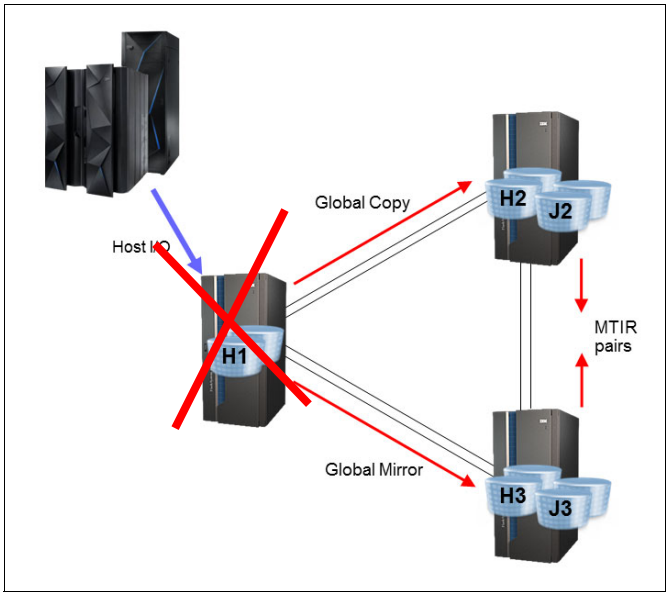

Figure 6-6 Global Copy plus Global Mirror topology

6.5.1 Recover at Global Mirror remote H3 site

The Global Mirror recovery at the H3 site is the same as in a non-Multiple Target PPRC environment. Further information about the process for this is described in the following related IBM Redbooks publications:

- *IBM System Storage DS8000 Copy Services for IBM System z*, SG24-6787
- *IBM System Storage DS8000 Copy Services for Open Systems*, SG24-6788

6.5.2 Restore Global Mirror disaster recovery capability

After recovering at the remote Global Mirror location, the disaster recovery protection of Global Mirror can be restored.

Failover H2:H1

A PPRC failover H2:H1 causes H2 to be converted from a secondary of H1 to become a suspended primary to H1. This is required so that H2 can become a secondary of H3.

Incremental Resynchronization for H3:H2

The Multiple Target PPRC configuration created MTIR pairs between H2 and H3. As with the previously described Multiple Target PPRC configurations, the MTIR pairs can be changed to active pairs with a PPRC failback command. In this case, the recovery is performed at the H3 site, and Global Copy is restored with a failback H3:H2 command that specifies Global Copy.

6.5.3 Start Global Mirror H3:H2

A Global Mirror session is started for H3:H2, and then the H3:H2 Global Copy pairs are added to this Global Mirror session. This results in the configuration shown in Figure 6-7.

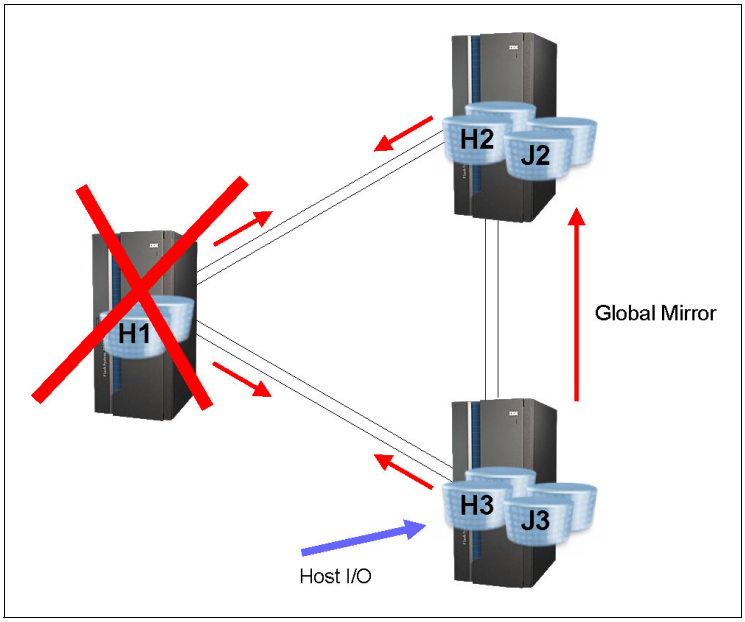

Figure 6-7 Restore Global Mirror capability

At this point, the disaster recovery protection of Global Mirror as been restored, with H2 acting as the remote recovery site.

6.5.4 H1 recovered

When the H1 site has been recovered, Global Copy H3:H1 can be resumed with a PPRC failback H3:H1. This results in the configuration shown in Figure 6-8.

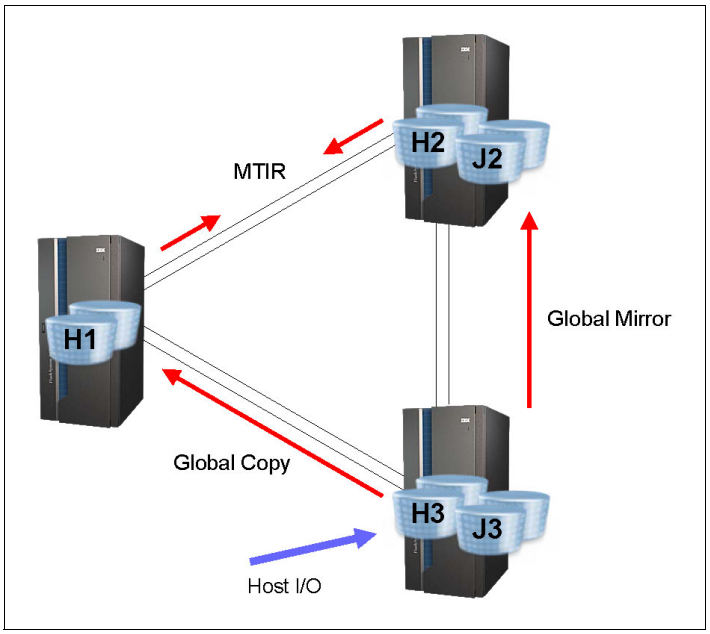

Figure 6-8 Failback H3:H1, Global Copy

The original topology of Global Copy plus Global Mirror has now been restored. The host I/O is running at the H3 site rather than the original H1 site. If you want to return to the H1 site, you can follow a similar sequence of steps to do a planned move back to the H1 site location.

7

Implementing a Metro Global Mirror and Metro Mirror topology

This chapter describes a Metro Global Mirror (MGM) topology with an additional Metro Mirror relationship at the primary site. It includes these topics:

- Overview of MGM and Metro Mirror topology
- Scenario: Outage at primary site H1

7.1 Overview of MGM and Metro Mirror topology

A topology of Metro Global Mirror (MGM) and Metro Mirror adds an additional synchronous Metro Mirror secondary site to an MGM topology, as shown in Figure 7-1.

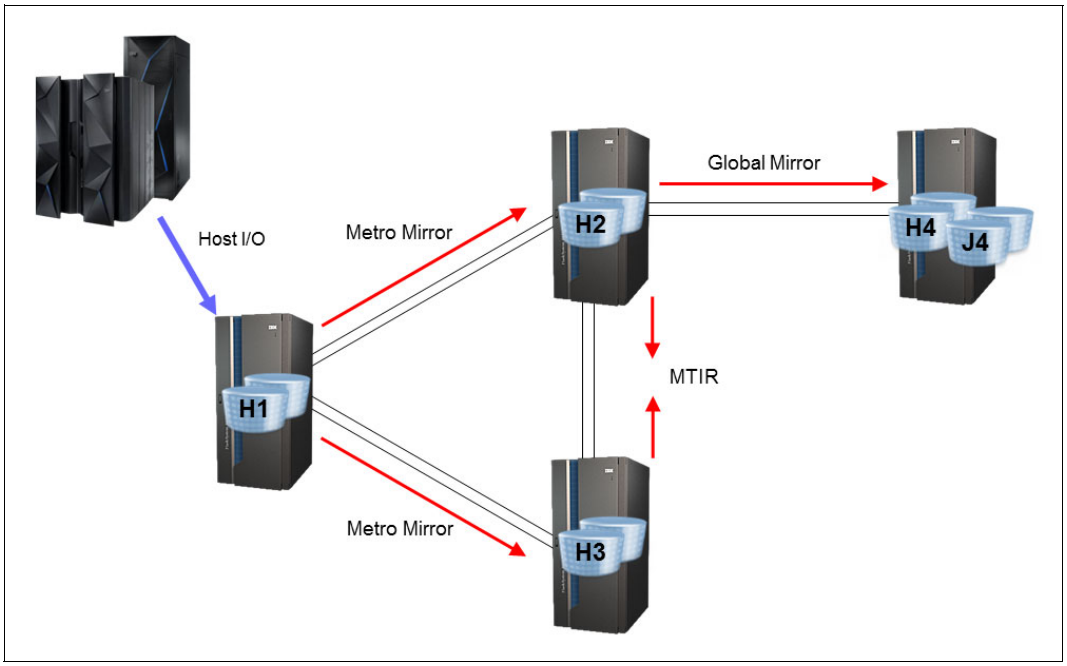

Figure 7-1 MGM and Metro Mirror topology

The relationships across the top of the illustration show an MGM configuration of Metro Mirror H1:H2 and Global Mirror H2:H4. Metro Mirror H1:H3 has been added to this configuration.

An alternative but equivalent way to view this topology is to consider it as a Multiple Target PPRC topology with Metro Mirror H1:H2 and H1:H3 to which a Global Mirror H2:H4 has been added.

An outage at one of the secondary sites is handled in a similar manner as in a non-Multiple Target PPRC configuration.

7.2 Scenario: Outage at primary site H1

In the case of an outage at the primary site H1, a decision can be made to move host production I/O to site H3. Either IBM HyperSwap or a traditional restart can be used. The steps for this are described in the subsections that follow.

7.2.1 Freeze H1:H2 and H1:H3

A freeze is issued for both the H1:H2 and H1:H3 relationships. A separate **`freeze`** command is required for each relationship. The production systems are stopped or quiesced. Depending upon the type of failure at H1, the **`freeze`** commands might not be able to be issued or they might fail. If the **`freeze`** commands are not successful, the pairs may remain duplex.

7.2.2 Failover H3:H1

A PPRC failover for H3:H1 converts H3 to become a suspended primary to H1. The production systems can then be restarted on H3. All writes that are received on H3 are tracked for later resynchronization. At this point, the configuration is as shown in Figure 7-2.

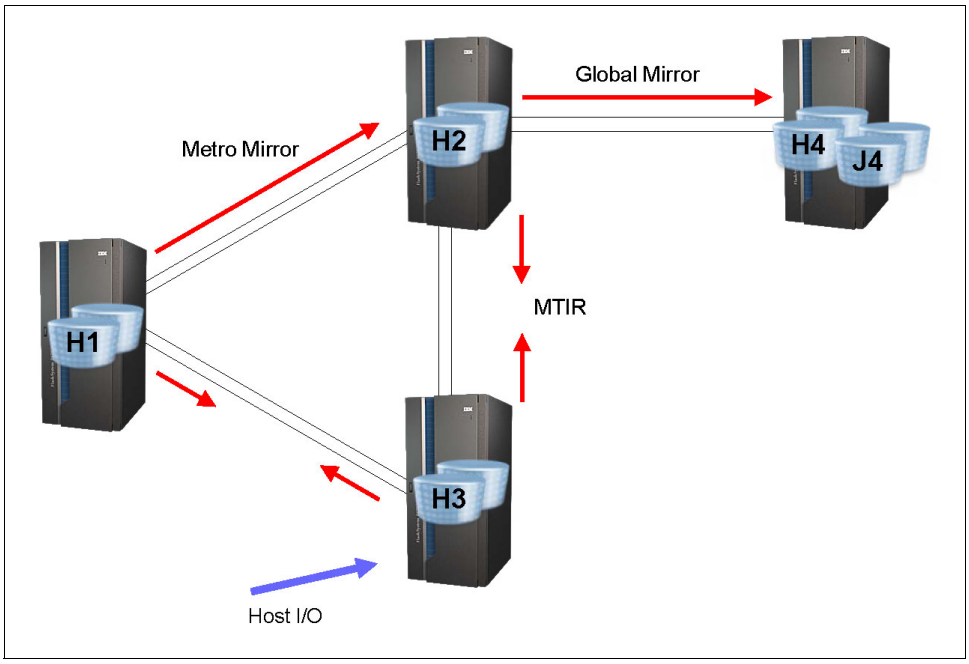

Figure 7-2 Failover H3:H1 and move host I/O

7.2.3 Start Metro Mirror H3:H2

The Multiple Target Incremental Resync (MTIR) pairs between H2 and H3 are used to restore the Metro Mirror replication. The steps involved follow.

Failover H2:H1

A PPRC failover H2:H1 causes H2 to be converted from a secondary of H1 to become a suspended primary to H1. This is required so that it is available to become a secondary of H3.

After this step H1, H2, and H3 are all suspended primary volumes to each of the other sites.

Failback H3:H2

A failback H3:H2 uses the Multiple Target PPRC Incremental Resynchronization capability to resume the H3:H2 pairs, restoring Metro Mirror replication and creating the configuration shown in Figure 7-3 on page 88.

At this point, there is the full protection of a cascaded MGM configuration. The Metro Mirror from H3:H2 provides local synchronous data protection and the Global Mirror H2:H4 provides long-distance, asynchronous disaster recovery protection.

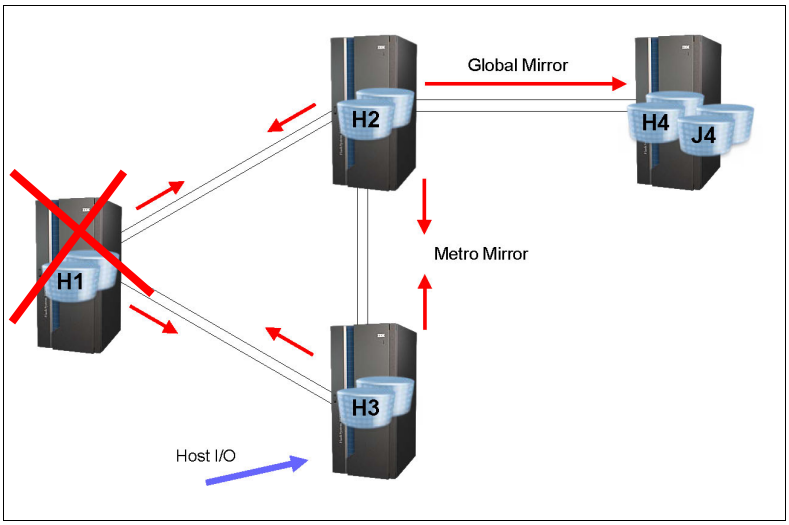

Figure 7-3 Failback H2:H3 to restore MGM configuration

7.2.4 H1 recovered

When site H1 is recovered, mirroring can be resumed to H2 to restore the full MGM plus Metro Mirror topology.

Establish paths H1:H2 and H1:H3

The paths H1:H2 and H1:H3 were removed by the **`freeze`** commands issued in 7.2.1, “Freeze H1:H2 and H1:H3” on page 86. After H1 is recovered, these paths can be reestablished so that they are available when needed. The paths are not required until the time when mirroring is later resumed from H1:H2 and H1:h3, but establishing them as soon as H1 is recovered helps to ensure that they are available when needed.

Failback H3:H1

A failback command for H3:H1 resumes the H3:H1 Metro Mirror pairs, as shown in Figure 7-4.

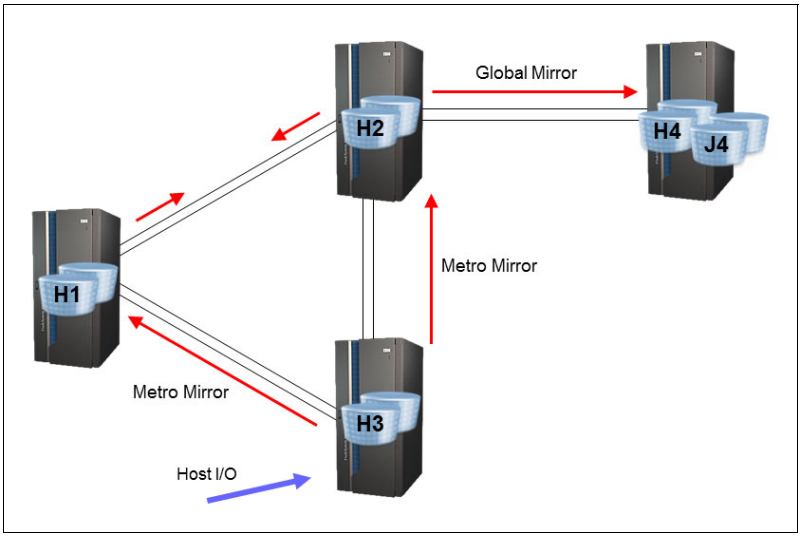

Figure 7-4 Failback H3:H1, restoring MGM + MM

This is now the same topology as the initial starting state. There is Metro Mirror H3:H1 and Metro Global Mirror H3:H2:H4. The host I/O is now running to H3 rather than H1. Depending upon the specific circumstances, for example if H1 and H3 are local to each other, there might be no advantage to running at one site versus the other. In other situations, there might be reasons to prefer that the host systems run at site H1. If wanted, production I/O can be returned to site H1 using the same sequence of steps as used for the move to H3.

8

Cascaded Metro Global Mirror improvement

This chapter describes how Multiple Target PPRC simplifies the use of a cascaded Metro Global Mirror (MGM) environment and provides added flexibility. It covers the following topics:

- Cascaded Metro Global Mirror topology
- Outage at H3
- Outage at H2
- Outage at H1
- Cascaded MGM and Multiple Target PPRC MGM

8.1 Cascaded Metro Global Mirror topology

Multiple Target PPRC provides several benefits and simplifications in a cascaded Metro Global Mirror (MGM) topology. Figure 8-1 shows a cascaded MGM topology where there is Metro Mirror H1:H2 and Global Mirror H2:H3.

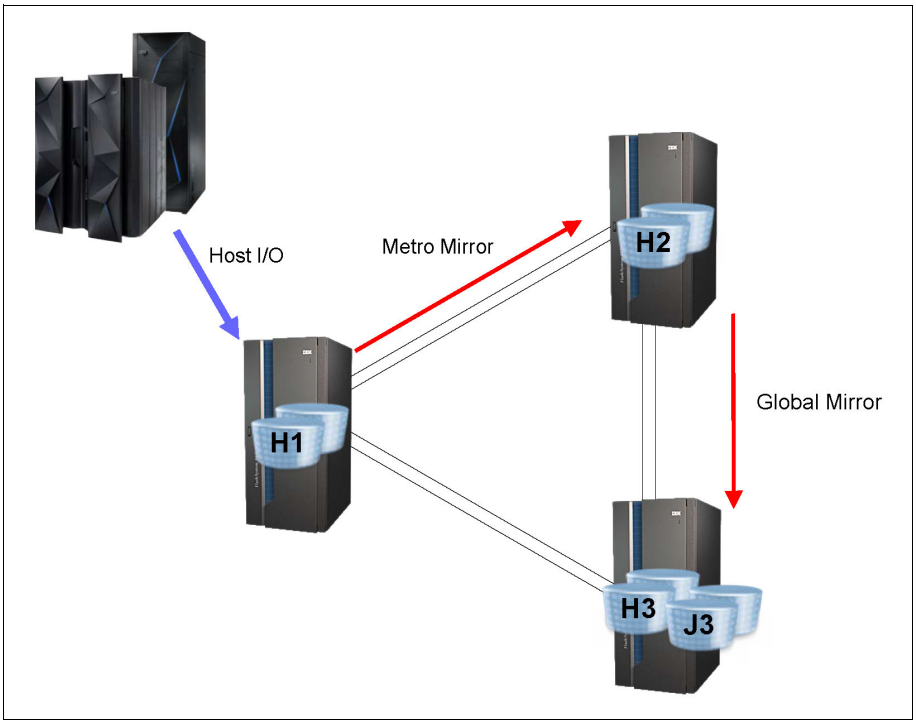

Figure 8-1 Cascaded MGM configuration

Like Multiple Target PPRC with Metro Mirror and Global Mirror, this topology provides both a local high availability capability and a long-distance disaster recovery capability at the same time. Data is synchronously mirrored from H1 to H2 and asynchronously mirrored from H2:H3. Because H2:H3 is an asynchronous copy, H3 can be located in a different region at a long distance from the local primary site. In the event of a widespread disaster that affects both H1 and H2, production can be restarted at the H3 remote location.

The scenarios in this section describe outages at one of the sites. These can be either a failure condition or a planned event for testing or maintenance purposes.

To simplify the descriptions, these sites are often referred to as though they were a single volume, but it must be understood that the terms refer to the entire set of volumes being mirrored by PPRC. For example, “Failover H2:H1” means to failover all of the H2 volumes to their corresponding H1 volumes.”

8.2 Outage at H3

A failure or outage at the Global Mirror remote site H3 can cause the Global Copy pairs to suspend and the Global Mirror session to stop forming consistency groups. The Metro Mirror pairs H1 continue to run and provide protection in case of a failure at H1. When site H3 is recovered, Global Mirror H2:H3 can be resumed to restore the Global Mirror protection.

8.3 Outage at H2

A failure or outage at the secondary site H2 can cause the Metro Mirror pairs H1:H2 to suspend.

Global Mirror with Incremental Resynchronization from H1 to H3 can be used to start Global Mirror H1:H3 and continue to provide disaster recovery protection. See one of the following IBM Redbooks publications for details about the use of MGM Incremental Resynchronization:

- *IBM System Storage DS8000 Copy Services for IBM System z*, SG24-6787
- *IBM System Storage DS8000 Copy Services for Open Systems*, SG24-6788

8.4 Outage at H1

Multiple Target PPRC can be used to simplify the recovery procedures in the case of a failure or outage at site H1 in a cascaded MGM configuration.

After a failure at the local site H1, production can be moved to the H2 Metro Mirror secondary site. The following sections describe the steps that are taken in the case of a failure or planned outage at the local production H1 site, including the use of the Multiple Target PPRC capabilities.

8.4.1 Terms used in this example

For the examples in this section, the following terms are used:

- **H1** is the current primary site where the production applications are running.
- **H2** is the intermediate site which is a Metro Mirror secondary site from H1 and a Global Mirror primary to H3.
- **H3** is the remote Global Mirror secondary site.

Table 8-1 identifies the DS8870 storage controllers used in the examples.

Note: The volume range on each DS8870 is different in these examples only to help clarify the different sites used. It is not a requirement that they be different.

Table 8-1 Identifications used in DS CLI examples

Site	Role	Dev	WWNN	Volume range
H1	Local primary	IBM.2107-75CZM21	5005076305FFD75A	5000 - 500F
H2	Intermediated cascaded site	IBM.2107-75CYM31	5005076305FFD71E	6000-600F
H3	Remote Global Mirror site	IBM.2107-75CYK71	5005076305FFD71A	7000-700F

8.4.2 Recover at H2

The first step in the process is to move production to H2 so that the host application I/O can continue. Either HyperSwap or a traditional restart can be used. After performing the freeze and unfreeze operations for H1:H2, a failover H2:H1 is performed to prepare the H2 site for receiving application host I/O.

Freeze H1:H2

Freeze commands for H1:H2 remove the PPRC paths and suspend all PPRC pairs for H1:H2. A separate freeze command is required for each LSS-to-LSS PPRC relationship. With the use of consistency groups, the freeze command creates consistent data at the H2 site by using extended long busy or queue full to temporarily queue dependent writes.

Depending on the type of failure at H1, the freeze commands might not complete successfully and some or all H1 pairs can remain full duplex. If the H1 site has completely failed, these commands might not have even been issued or executed.

Failover H2:H1 without Multiple Target PPRC

Without Multiple Target PPRC in a cascaded configuration, a failover H2:H1 cannot cause H2 to become a suspended primary to H1 because it is already a primary to H3, and a PPRC primary volume can be a source to only one secondary. In this case, H2 detects that it is a cascaded volume, and the processing is different than for a conventional failover in a noncascaded configuration. In a cascaded failover, the H2 volumes remain as secondaries of H1 and their secondary pair state is changed to target suspended. The H2 volumes remain as active primary volumes to H3. The states of the H2 volumes are changed to allow host I/O to be received.

Example 8-1 shows a failover H2:H1 without the **-multtgt** option. This prevents the failover command from creating a Multiple Target PPRC configuration. As this example shows, H2 is left as a target of H1.

Example 8-1 Failover H2:H2 without -multtgt

```
dscli> failoverpprc -dev IBM.2107-75CYM31 -remotedev IBM.2107-75CZM21 -type mmir 6000-600f:5000-500f
CMUC00196I failoverpprc: Remote Mirror and Copy pair IBM.2107-75CYM31/6000:IBM.2107-75CZM21/5000 successfully reversed.
...

dscli> lspprc -dev IBM.2107-75CYM31 6000-600f
ID                                        State            Reason      Type         SourceLSS
=============================================================================================================
IBM.2107-75CYM31/6000:IBM.2107-75CYK71/7000 Copy Pending     -           Global Copy  IBM.2107-75CYM31/60 ...
...
IBM.2107-75CYM31/600F:IBM.2107-75CYK71/700F Copy Pending     -           Global Copy  IBM.2107-75CYM31/60 ...
IBM.2107-75CZM21/5000:IBM.2107-75CYM31/6000 Target Suspended Host Target Metro Mirror IBM.2107-75CZM21/50 ...
...
IBM.2107-75CZM21/500F:IBM.2107-75CYM31/600F Target Suspended Host Target Metro Mirror IBM.2107-75CZM21/50 ...
```

Figure 8-2 shows the configuration state after a failover H2:H1.

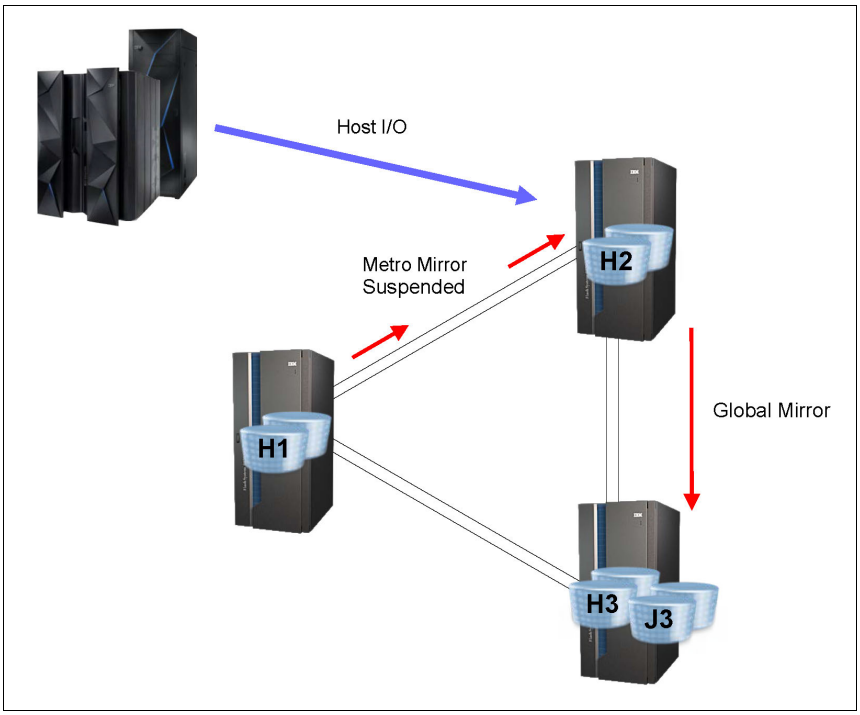

Figure 8-2 Failover H2:H1 without Multiple Target PPRC

It is not possible to failback H2:H1 without first suspending the H2:H3 pairs.

Failover H2:H1 using Multiple Target PPRC

Multiple Target PPRC target allows for a simpler handing of a PPRC failover in a cascaded configuration. With Multiple Target PPRC, a PPRC failover H2:H1 behaves in a similar manner to a failover in a noncascaded configuration and causes H2 to become a suspended primary to the H1 volumes.

In some circumstances, it is necessary for a failover in a cascaded configuration to behave in the same manner as without Multiple Target PPRC support. These are among the reasons for requiring this previous behavior:

- Not all of the DS8000 storage systems have Multiple Target PPRC support yet, so a Multiple Target PPRC configuration is not allowed.
- Not all host systems have been updated to have Multiple Target PPRC support.
- Customer processes and procedures have not yet been updated and still expect the previous behavior.

A keyword on the PPRC failover command instructs the DS8870 storage system that creating a Multiple Target PPRC configuration is allowed. If this keyword is not specified, the previous cascaded failover behavior is used. If the keyword is included on the PPRC failover command, a Multiple Target PPRC configuration is created.

Example 8-2 on page 96 shows a sample DS CLI command to failover to the intermediate H2 site by using the **`-multtgt`** option to allow the creation of a Multiple Target PPRC configuration.

Example 8-2 Failover H2:H1 with -multtgt

```
dscli> failoverpprc -dev IBM.2107-75CYM31 -remotedev IBM.2107-75CZM21 -type mmir -multtgt
6000-600f:5000-500f
CMUC00196I failoverpprc: Remote Mirror and Copy pair IBM.2107-75CYM31/6000:IBM.2107-75CZM21/5000
successfully reversed.
...
CMUC00196I failoverpprc: Remote Mirror and Copy pair IBM.2107-75CYM31/600F:IBM.2107-75CZM21/500F
successfully reversed.
```

As Example 8-3 shows, the failover H2:H1 caused H2 to become a Multiple Target PPRC primary to both H1 and H3. Compare this to the results shown in Example 8-1 on page 94 when the **-multtgt** option was not used.

Example 8-3 Query after failover H2:H1 with -multtgt

```
dscli> lspprc -dev IBM.2107-75CYM31 6000-600f
ID                                        State        Reason      Type         SourceLSS
==========================================================================================
IBM.2107-75CYM31/6000:IBM.2107-75CYK71/7000 Copy Pending -           Global Copy  IBM.2107-75CYM31/60 ...
IBM.2107-75CYM31/6000:IBM.2107-75CZM21/5000 Suspended    Host Source Metro Mirror IBM.2107-75CYM31/60 ...
...
IBM.2107-75CYM31/600F:IBM.2107-75CYK71/700F Copy Pending -           Global Copy  IBM.2107-75CYM31/60 ...
IBM.2107-75CYM31/600F:IBM.2107-75CZM21/500F Suspended    Host Source Metro Mirror IBM.2107-75CYM31/60 ...
```

The host applications can then run at the H2 site by using either HyperSwap or a traditional restart method.

The configuration at this point in the process is shown in Figure 8-3.

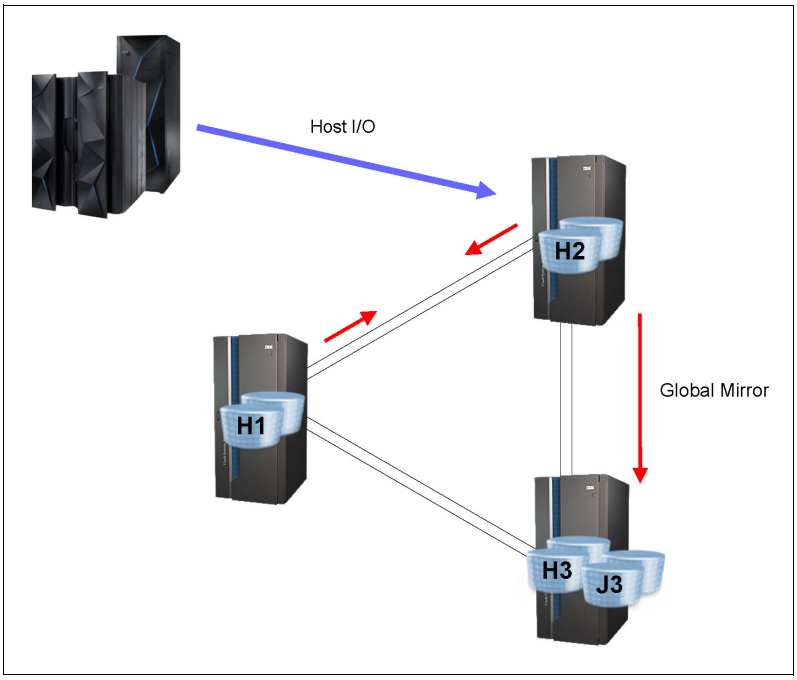

Figure 8-3 Multiple Target PPRC failover H2:H1

8.4.3 H1 recovered

This scenario is for the case where Multiple Target PPRC was used for the failover H2:H1 after the outage at H1.

When site H1 is recovered, mirroring can be resumed to H2 to create a Multiple Target PPRC MGM configuration.

Establish paths H1:H2

The paths for H1:H2 were removed by the freeze commands. After H1 is recovered, these paths can be reestablished so that they are available when needed. The paths are not required until the time when mirroring is later resumed for H1:H2, but establishing them as soon as H1 is recovered helps to ensure that they are available when needed. The paths are reestablished using the same command that was used to originally create the paths.

Failback H2:H1

When the local H1 site has been recovered, a failback H2:H1 resumes Metro Mirror from H2:H1 and copies all data that was updated since the time of the failover to H2. Example 8-4 shows an example DS CLI command to failback H2:H1.

Example 8-4 Failback H2:H1

```
dscli> failbackpprc -dev IBM.2107-75CYM31 -remotedev IBM.2107-75CZM21 -type mmir 6000-600f:5000-500f
CMUC00197I failbackpprc: Remote Mirror and Copy pair IBM.2107-75CYM31/6000:IBM.2107-75CZM21/5000
successfully failed back.
...
CMUC00197I failbackpprc: Remote Mirror and Copy pair IBM.2107-75CYM31/600F:IBM.2107-75CZM21/500F
successfully failed back.
```

The cascaded MGM configuration is now a Multiple Target PPRC MGM configuration with Metro Mirror H2:H1 and Global Mirror H2:H3. Because H2 is now a Multiple Target PPRC primary, MTIR pairs are created between the secondary H1 and H3 sites. The resulting configuration is shown in Figure 8-4.

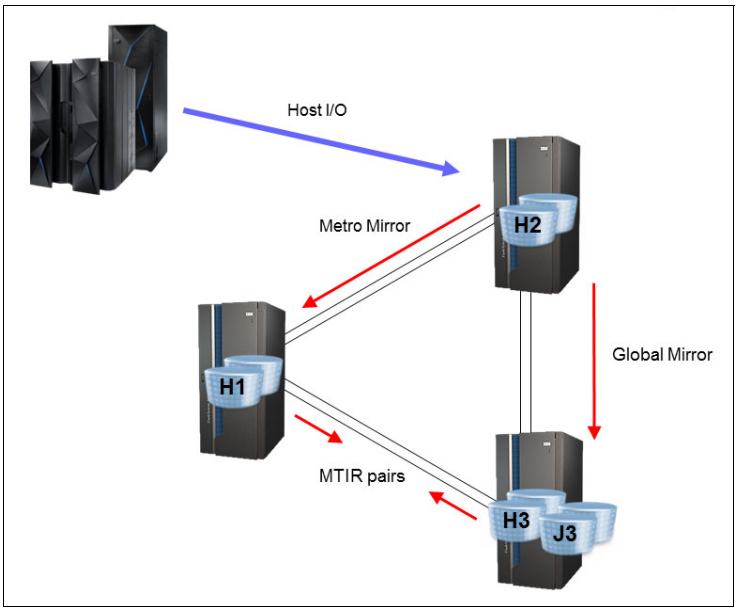

Figure 8-4 Failback H2:H1 creating Multiple Target PPRC MGM configuration

8.4.4 Return production to H1

After the H1 site has been recovered, the original cascaded configuration of Metro Mirror H1:H2 and Global Mirror H2:H3 can be restored easily. The first step in the process is to move production back to H1. The process and commands are similar to those described for the two Metro Mirror case in 4.5, “Return production to H1” on page 57, which contains examples of the DS CLI commands.

As with the move to H2, either HyperSwap or a traditional restart can be used. The steps involved in this process follow.

Freeze H2:H1

A freeze command for H2:H1 removes the PPRC paths for H2:H1 and suspends the PPRC H2:H1 pairs. A separate freeze command is required for each LSS-to-LSS PPRC relationship. With the use of consistency groups, this creates consistent data at the H1 site.

Unfreeze and run H2:H1

When all freeze commands have completed, the secondary H1 site contains a consistent copy of data. The unfreeze or run command (also known as a *consistency group created* command) removes the extended long busy or queue full condition at the primary H2 site. A separate unfreeze command is required for each LSS-to-LSS PPRC relationship.

Failover H1:H2

The failover command for H1:H2 coverts the H1 volumes to suspended Metro Mirror primary volumes whose secondary volumes are H2.

Resume host systems at H1

The host systems can now be restarted at H1. If HyperSwap is used, the host systems are switched to H1.

8.4.5 Start replication H1:H2

Because this is a planned swap, there is no need to wait for a secondary system to be recovered. The replication of H1:H2 can be immediately resumed.

Failback H1:H2

A failback command for H1:H2 resumes replication for the H1:H2 pairs.

Establish paths H2:H1

The paths from H2:H1 were removed by the freeze commands that were performed before the failover H2:H1. Reestablishing these paths as soon as possible helps to ensure that they are available when needed.

The configuration has now been restored to original configuration as shown in Figure 8-1 on page 92.

8.5 Cascaded MGM and Multiple Target PPRC MGM

This scenario described in 8.1, “Cascaded Metro Global Mirror topology” on page 92 started with a cascaded MGM configuration, converted it to a Multiple Target PPRC MGM configuration and then went back to the original cascaded configuration again. This demonstrates how cascaded and Multiple Target PPRC MGM configurations complement each other.

In the course of operations, a configuration can change from a cascaded configuration to Multiple Target PPRC and back again. The ability to move between cascaded and Multiple Target MGM configurations as needed gives additional flexibility and options to manage different requirements and circumstances.

9

Using Multiple Target PPRC for migration

This chapter demonstrates how Multiple Target PPRC can be used for data migration between DS8870 storage systems within an existing PPRC environment.

Scenarios are presented to give a detailed description for replacing a Metro Mirror secondary or a Metro Mirror primary volume without the loss of replication protection.

The following topics are covered:

- General considerations
- Replacement of Metro Mirror secondary
- Replacement of Metro Mirror primary
- Replacement of Metro Mirror primary, alternative method

9.1 General considerations

Multiple Target PPRC can be used to migrate data from a primary or secondary DS8870 storage system in PPRC configuration. The use of Multiple Target PPRC allows for migration procedures with either few or no periods of time where the system is not protected by mirroring.

In a Multiple Target PPRC configuration, all DS8870 storage systems must have Multiple Target PPRC support and the correct license features.

Note: All DS8870 storage systems must support Multiple Target PPRC to use these techniques. This means that these methods cannot be used to migrate from previous code levels. But after all systems have Multiple Target PPRC support, the technique can be used for future migrations.

9.2 Replacement of Metro Mirror secondary

The general method of this migration is to use the Multiple Target PPRC capabilities to start Metro Mirror from the existing H1 primary site to the new H2' secondary site. After all of the volume pairs H1:H2' have reached full duplex, the original H2 site can be removed.

The new H2' storage system can be at the same location as the existing H2 system, as for a hardware replacement. Alternatively, the new H2' system can be at a different location, as for a data center move.

Metro Mirror H1:H2 remains active until Metro Mirror H1:H2' is fully operational, so there is no loss of recovery capability during this process.

9.2.1 Requirements

The Multiple Target PPRC license is required for all DS8870 storage systems at the H1 primary site, the current H2 secondary site, and the new H2' secondary site.

Fibre Channel connectivity is required from the H1 site to the new H2 site.

9.2.2 Terms used in examples

For the examples in this section, the following terms are used:

- **H1** is the current primary site where the production applications are running.
- **H2** is the current Metro Mirror secondary site to which H1 is mirroring data.
- **H2'** is the new Metro Mirror secondary site which is replacing the current H2.

Table 9-1 on page 103 identifies the DS8870 storage controllers used in the examples in this chapter.

Note: The volume range on each DS8870 is different in these examples only to help clarify the different sites used. It is not a requirement that they be different.

Table 9-1 Identifications used in DS CLI examples

Site	Role	Dev	WWNN	Volume range
H1	Current primary	IBM.2107-75CZM21	5005076305FFD75A	5000 - 500F
H2	Current secondary	IBM.2107-75CYM31	5005076305FFD71E	6000 - 600F
H2’	New secondary (to replace H2)	IBM.2107-75CYK71	5005076305FFD71A	7000 - 700F

9.2.3 Initial configuration

The examples use DS CLI commands. Equivalent TSO and ICKDSF commands are listed in section 1.8, “Commands” on page 21.

The initial configuration consists of Metro Mirror pairs H1:H2, as shown in Figure 9-1.

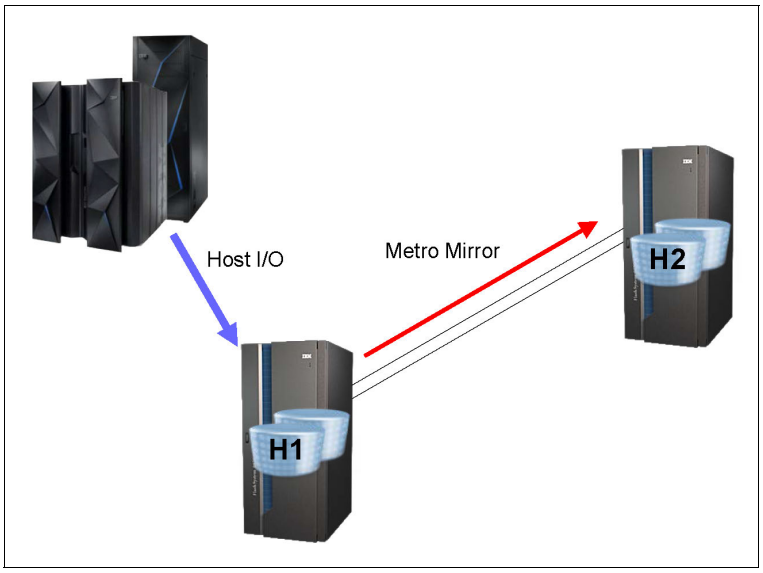

Figure 9-1 Initial configuration for migration of Metro Mirror secondary

Site H1 contains the Metro Mirror primary volumes, and site H2 contains the Metro Mirror secondary volumes that are to be migrated to a different DS8870 system. H2 may be only some of the volumes on a DS8870, all volumes on a DS8870, or volumes across multiple DS8870 systems.

The H1:H2 pairs are in the full-duplex state at the beginning of this scenario.

9.2.4 Installation of new DS8870 at H2’

Install, configure, and physically connect the new H2’ storage system, including any host system requirements and the Fibre Channel connections to H1 to be used for Metro Mirror.

9.2.5 PPRC paths H1:H2’

PPRC paths from H1 to the new H2’ are required before the PPRC pairs can be established. Example 9-1 on page 104 shows a DS CLI command used to create PPRC paths from H1 to

H2'. Separate **mkpprcpath** commands are required for each primary-to-secondary LSS relationship.

Example 9-1 Establish paths H1:H2'

```
dscli> mkpprcpath -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYK71 -remotewwnn 5005076305FFD71A -srclss 50
-tgtlss 70 -consistgrp I0201:I0032 I0234:I0306
CMUC00149I mkpprcpath: Remote Mirror and Copy path 50:70 successfully established.
```

9.2.6 Create H1:H2' Metro Mirror pairs

After establishing all of the required PPRC paths, create the H1:H2' Metro Mirror volume pairs. These pairs can be initially created as asynchronous Global Copy pairs and later converted to Metro Mirror, or they can be created directly as Metro Mirror pairs.

In some large configurations, it is preferable to start the replication as Global Copy because it is less likely to affect the production systems during the initial copy phase. After all pairs have copied nearly 100% of the data, the pairs can be converted to synchronous Metro Mirror.

Example 9-2 shows a sample **mkpprc** command to create the H1:H2' Metro Mirror pairs.

Example 9-2 Create H1:H2' Metro Mirror pairs

```
dscli> mkpprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYK71 -type mmir 5000-500f:7000-700f
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship
IBM.2107-75CZM21/5000:IBM.2107-75CYK71/7000 successfully created.
...
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship
IBM.2107-75CZM21/500F:IBM.2107-75CYK71/700F successfully created.
```

Figure 9-2 shows the configuration after establishing Metro Mirror H1:H2'. Notice that MTIR pairs are established between H2 and H2' even though they are not required for this migration scenario. It is not required to establish PPRC paths between H2 and H2' for these MTIR pairs because there is no intent to create active pairs between the two sites.

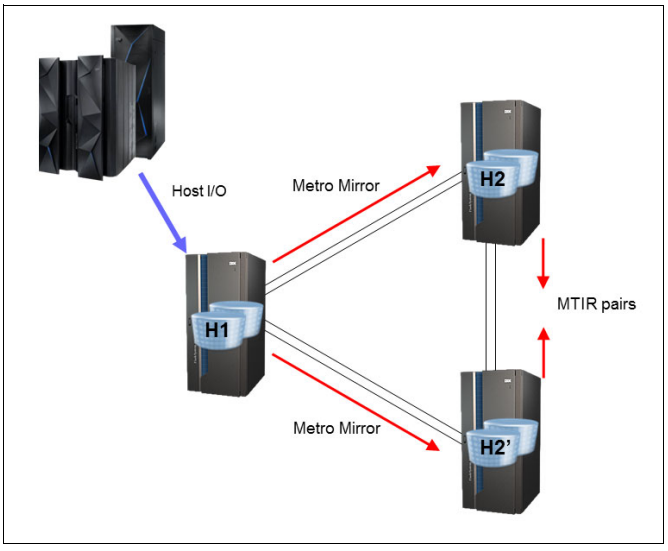

Figure 9-2 Metro Mirror established to H2'

After the Metro Mirror pairs H1:H2' have been established, the data at site H1 will be replicated to the new H2' site. During this time, the original H1:H2 pairs remain active and H2 is still available as a recovery site in case of an error at H1.

9.2.7 Monitor for H1:H2' full duplex

The H1:H2' pairs copy all data from H1 to the new H2' site. During this time, the progress of the copy can be monitored. Example 9-3 is a sample DS CLI command that shows the pair state and the number of out-of-sync tracks left to copy for each pair. When the number of out-of-sync tracks for a pair drops to zero, the pair state changes from Copy Pending to Full Duplex. The output of the **lspprc** command can be long, especially when the **-l** (long) option is specified, so the output in this example is truncated on the right side to better show the fields of interest.

Example 9-3 Monitoring for full duplex

```
dscli> lspprc -l -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYK71 5000-500f:7000-700f
ID                                        State        Reason Type         Out Of Sync Tracks Tgt Read
=============================================================================================================
IBM.2107-75CZM21/5000:IBM.2107-75CYK71/7000 Full Duplex  -      Metro Mirror 0                 Disabled...
IBM.2107-75CZM21/5001:IBM.2107-75CYK71/7001 Copy Pending -      Metro Mirror 38750             Disabled...
IBM.2107-75CZM21/5002:IBM.2107-75CYK71/7002 Copy Pending -      Metro Mirror 81920             Disabled...
...
```

After all H1:H2' pairs reach full duplex, the migration procedure can proceed.

9.2.8 Remove H1:H2

After all of the H1:H2' pairs have reached full duplex, the H1:H2 pairs can be removed, because the H1:H2' pairs are now available to provide the Metro Mirror protection.

Example 9-4 shows a DS CLI command to remove the H1:H2 pairs.

Example 9-4 Remote H1:H2 Metro Mirror pairs

```
dscli> rmpprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYM31 -quiet 5000-500f:6000-600f
CMUC00155I rmpprc: Remote Mirror and Copy volume pair IBM.2107-75CZM21/5000:IBM.2107-75CYM31/6000
relationship successfully withdrawn.
...
CMUC00155I rmpprc: Remote Mirror and Copy volume pair IBM.2107-75CZM21/500F:IBM.2107-75CYM31/600F
relationship successfully withdrawn.
```

After removing the H1:H2 pairs, there is no longer a Multiple Target PPRC configuration, and the DS8870 storage systems automatically remove the MTIR pairs that were created between H2 and H2'. The PPRC paths between H1 and the original H2 are no longer needed and should be removed. In most instances, the PPRC paths will be established in both directions between H1 and H2.

Example 9-5 shows a sample DS CLI command for removing these paths.

Example 9-5 Remove PPRC paths between H1 and H2

```
dscli> rmpprcpath -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYM31 -remotewwnn 5005076305FFD71E -quiet
50:60
CMUC00150I rmpprcpath: Remote Mirror and Copy path IBM.2107-75CZM21/50:IBM.2107-75CYM31/60 successfully
removed.
```

```
dscli> rmpprcpath -dev IBM.2107-75CYM31 -remotedev IBM.2107-75CZM21 -remotewwnn 5005076305FFD75A -quiet
60:50
CMUC00150I rmpprcpath: Remote Mirror and Copy path IBM.2107-75CYM31/60:IBM.2107-75CZM21/50 successfully
removed.
```

The resulting configuration is shown in Figure 9-3.

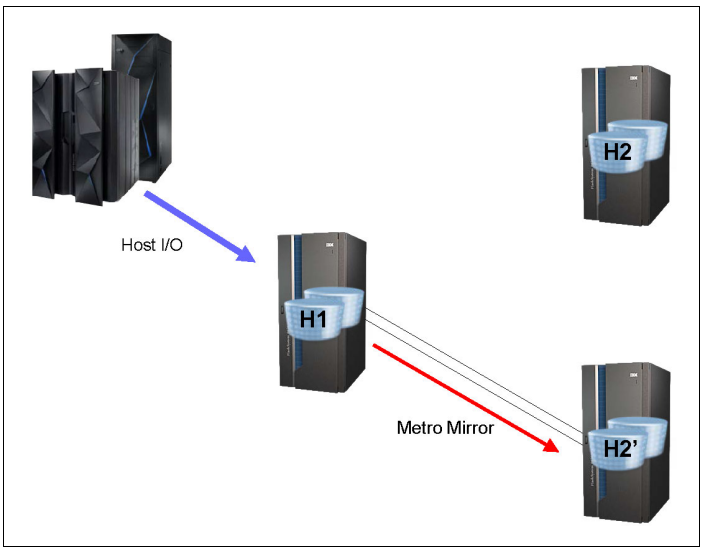

Figure 9-3 Metro Mirror H1:H2' after migration

The migration is now complete.

9.3 Replacement of Metro Mirror primary

The general method of this migration is to use the Multiple Target PPRC capabilities to start Metro Mirror from the existing H1 primary site to the new H1' primary site. After all of the H1:H1' volume pairs have reached full duplex, the host systems can be cut over to run at the new H1' site. Multiple Target PPRC Incremental Resync is then used to start Metro Mirror H1' to H2 and restore the Metro Mirror configuration.

9.3.1 Requirements

The Multiple Target PPRC license is required for all DS8870 storage systems at the H1 primary site, the current H2 secondary site, and the new H1' primary site.

Fibre Channel connectivity is required from the new H1' site to H2.

9.3.2 Terms used in examples

For the examples in this section, the following terms are used:

- **H1** is the current primary site where the production applications are running.
- **H2** is the current Metro Mirror secondary site to which H1 is mirroring data.
- **H1'** is the new Metro Mirror secondary site which is replacing H1.

Table 9-2 identifies the DS8870 storage controllers used in the examples in this section.

Note: The volume range on each DS8870 is different only to help clarify the different sites in the examples. It is not a requirement that they be different.

Table 9-2 Identifications used in DS CLI examples

Site	Role	Dev	WWNN	Volume range
H1	Current primary	IBM.2107-75CZM21	5005076305FFD75A	5000 - 500F
H2	Current secondary	IBM.2107-75CYM31	5005076305FFD71E	6000 - 600F
H1'	New primary (to replace H1)	IBM.2107-75CYK71	5005076305FFD71A	7000 - 700F

9.3.3 Initial configuration

The initial configuration consists of a Metro Mirror pair H1:H2, as shown in Figure 9-4.

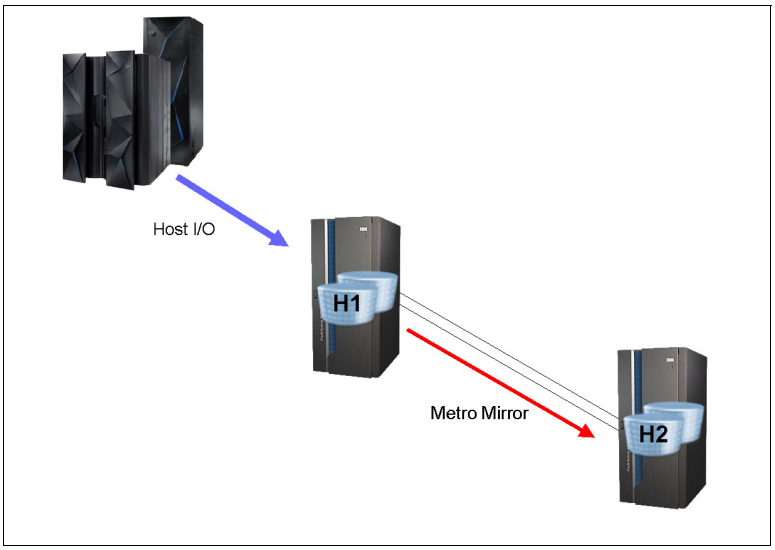

Figure 9-4 Initial configuration for migration of Metro Mirror primary

Site H1 contains the Metro Mirror primary volumes that are to be migrated to a different DS8870 storage system, and site H2 contains the Metro Mirror secondary volumes. H1 may be only some of the volumes on a DS8870 storage system or it may be volumes across multiple DS8870 storage systems.

The H1:H2 pairs are in the full-duplex state at the beginning of this scenario.

9.3.4 Installation of new DS8870 at H1'

Install, configure, and physically connect the new H1' storage system, including any host system requirements and the Fibre Channel connections to H2 to be used for the Metro Mirror PPRC paths.

9.3.5 Start Metro Mirror H1:H1'

PPRC H1:H1' paths are required for the replication of data from the original H1 to the new H1' system. In addition, the new Metro Mirror configuration will be H1' to H2, so PPRC paths are

required for H1' to H2. Finally, in the event of a swap to H2, paths H2:H1' are required to be able to mirror back to H1'.

The following commands are required:

- Establish PPRC paths H1:H1'
- Establish PPRC paths H1':H2
- Establish PPRC paths H2:H1'

Example 9-6 shows a sample DS CLI command to establish paths from H1 to the new H1'.

Example 9-6 Establish PPRC paths H1:H1'

```
dscli> mkpprcpath -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYK71 -remotewwnn 5005076305FFD71A -srclss 50
-tgtlss 70 -consistgrp IO201:IO032 IO234:IO306
CMUC00149I mkpprcpath: Remote Mirror and Copy path 50:70 successfully established.
```

Example 9-7 shows sample DS CLI commands to establish PPRC paths from H1' to H2 and then in the opposite direction, from H2 to H1'.

Example 9-7 Establish PPRC paths H1':H2 and H2:H1'

```
dscli> mkpprcpath -dev IBM.2107-75CYK71 -remotedev IBM.2107-75CYM31 -remotewwnn 5005076305FFD71E -srclss 70
-tgtlss 60 -consistgrp IO300:IO300 IO330:IO330
CMUC00149I mkpprcpath: Remote Mirror and Copy path 70:60 successfully established.

dscli> mkpprcpath -dev IBM.2107-75CYM31 -remotedev IBM.2107-75CYK71 -remotewwnn 5005076305FFD71A -srclss 60
-tgtlss 70 -consistgrp IO300:IO300 IO330:IO330
CMUC00149I mkpprcpath: Remote Mirror and Copy path 60:70 successfully established.
```

Create Metro Mirror H1:H1'

After required PPRC paths have been established, the H1:H1' pairs can be created. Example 9-8 shows a sample DS CLI command to create the H1:H1' Metro Mirror pairs.

Example 9-8 Create Metro Mirror H1:H1'

```
dscli> mkpprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYK71 -type mmir 5000-500f:7000-700f
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship
IBM.2107-75CZM21/5000:IBM.2107-75CYK71/7000 successfully created.
...
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship
IBM.2107-75CZM21/500F:IBM.2107-75CYK71/700F successfully created.
```

This is now a Multiple Target PPRC configuration, so MTIR pairs are automatically created by the DS8870 storage system for H1':H2 and H2:H1'.

This is also now a Multiple Target PPRC configuration, so MTIR pairs will be automatically created for H1':H2 and H2:H1'.

Figure 9-5 on page 109 shows the configuration after creating the Metro Mirror pairs H1:H1'.

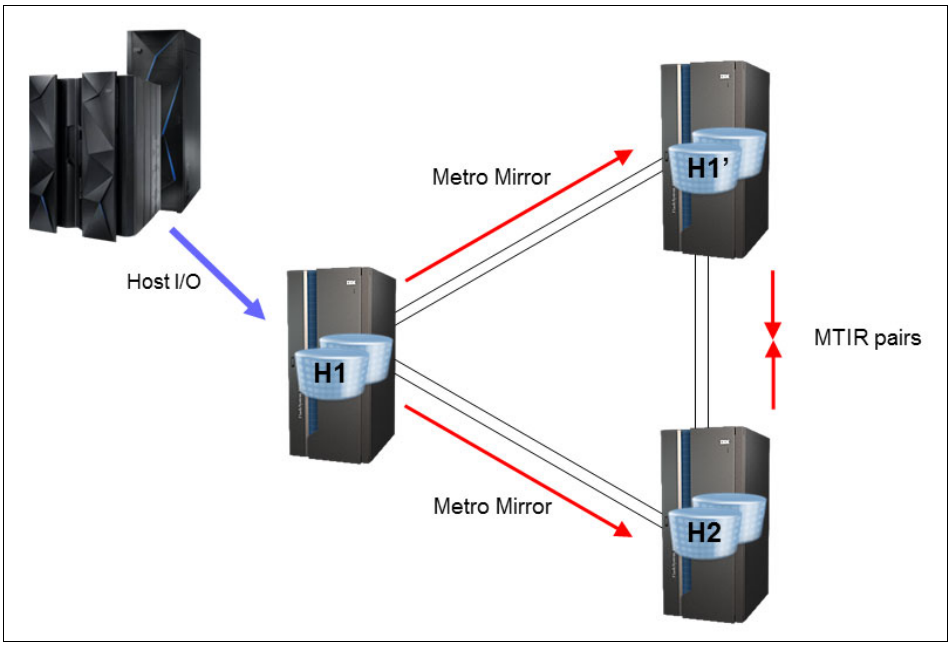

Figure 9-5 Metro Mirror established to H1'

Notice that MTIR pairs are established between H1' and H2. These MTIR pairs are used for the Incremental Resync H1':H2, which allows for mirroring to be resumed after moving the host systems to run on the new H1' without requiring a full copy of data.

The correct configuration can be verified by querying the volumes within the Multiple Target PPRC configuration. Example 9-9 shows a query to volumes at H1 to verify that they are Multiple Target Metro Mirror volumes that are being mirrored to both H2 and H1'. While the initial copy of data is being performed, some of the H1:H1' pairs will be Copy Pending, while others are Full Duplex.

Example 9-9 Query to primary H1 volumes

```
dscli> lspprc -dev IBM.2107-75CZM21 5000-500f
ID                                        State        Reason Type         SourceLSS
==========================================================================================
IBM.2107-75CZM21/5000:IBM.2107-75CYK71/7000 Full Duplex  -      Metro Mirror IBM.2107-75CZM21/50 ...
IBM.2107-75CZM21/5000:IBM.2107-75CYM31/6000 Full Duplex  -      Metro Mirror IBM.2107-75CZM21/50 ...
IBM.2107-75CZM21/5001:IBM.2107-75CYK71/7001 Copy Pending -      Metro Mirror IBM.2107-75CZM21/50 ...
IBM.2107-75CZM21/5001:IBM.2107-75CYM31/6001 Full Duplex  -      Metro Mirror IBM.2107-75CZM21/50 ...
...
IBM.2107-75CZM21/500F:IBM.2107-75CYK71/700F Copy Pending -      Metro Mirror IBM.2107-75CZM21/50 ...
IBM.2107-75CZM21/500F:IBM.2107-75CYM31/600F Full Duplex  -      Metro Mirror IBM.2107-75CZM21/50 ...
```

A query of the secondary volumes can be performed to verify that they are in the expected state.

A sample DS CLI query of the H1' volumes is shown in Example 9-10. This query shows that there are MTIR pairs H1':H2 and also that H1' is the target of the original H1.

Example 9-10 Query of new H1' volumes

```
dscli> lspprc -dev IBM.2107-75CYK71 -multtgt 7000-700f
ID                                        State              Reason               Type         SourceLSS
==========================================================================================================
IBM.2107-75CYK71/7000:IBM.2107-75CYM31/6000 Suspended        Multi-target Internal Global Copy  IBM.2107-75CYK71/70...
...
IBM.2107-75CYK71/700F:IBM.2107-75CYM31/600F Suspended        Multi-target Internal Global Copy  IBM.2107-75CYK71/70...
IBM.2107-75CZM21/5000:IBM.2107-75CYK71/7000 Target Full Duplex -                    Metro Mirror IBM.2107-75CZM21/50...
...
IBM.2107-75CZM21/500F:IBM.2107-75CYK71/700F Target Full Duplex -                    Metro Mirror IBM.2107-75CZM21/50...
```

During this time, the original H1:H2 pairs remain active and H2 remains available as a recovery site in case of an outage at H1.

9.3.6 Monitor for H1:H1' full duplex

The H1:H1' pairs will copy all data from H1 to the new H1' box. During this time, the progress of the copy can be monitored. Example 9-11 is a sample DS CLI command that shows the pair state and the number of out-of-sync tracks left to copy for each pair. When the number of out-of-sync tracks for a pair goes to zero, the pair state changes from Copy Pending to Full Duplex. The output of the **lspprc** command can be long, especially when the **-l** (long) option is specified, so the output has been truncated on the right side to better show the fields of interest.

Example 9-11 Monitoring for full duplex

```
dscli> lspprc -l -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYK71 5000-500f:7000-700f
ID                                        State        Reason Type         Out Of Sync Tracks Tgt Read
==========================================================================================================
IBM.2107-75CZM21/5000:IBM.2107-75CYK71/7000 Full Duplex  -      Metro Mirror 0                 Disabled...
IBM.2107-75CZM21/5001:IBM.2107-75CYK71/7001 Copy Pending -      Metro Mirror 41662             Disabled...
IBM.2107-75CZM21/5002:IBM.2107-75CYK71/7002 Copy Pending -      Metro Mirror 486               Disabled...
...
```

9.3.7 Monitor out of sync tracks for H1':H2 and H2:H1'

As each H1:H1' pair goes full duplex, the MTIR change recording mechanism will start as described in 3.2.3, "MTIR change recording" on page 34. It typically takes several minutes after reaching full duplex for the number of out-of-sync tracks for the MTIR pairs to drop from the total number of tracks on the volume to only the most recently updated tracks.

Note: If the move to H1' is performed while the MTIR out-of-synchronization track count is still at the full volume level, then the replication for H1':H2 will require a full copy of the volumes and the benefit of the Multiple Target PPRC Incremental Resynchronization will not be realized.

Example 9-12 on page 111 shows a query of the MTIR pairs H1':H2 and H2:H1' displaying the full volume counts, followed by a later query of the same relationships after the counts have been dropped to represent only the most recently updated tracks.

Because the different pairs reach full duplex at different times, the out-of-synchronization track counts also change at different times.

Example 9-12 Query MTIR pair out-of-synchronization track counts

```
dscli> lspprc -l -dev IBM.2107-75CYM31 -remotedev IBM.2107-75CYK71 -multtgt 6000-600f:7000-700f
ID                                        State     Reason               Type        Out Of Sync Tracks
=============================================================================================================
IBM.2107-75CYM31/6000:IBM.2107-75CYK71/7000 Suspended Multi-target Internal Global Copy 96 ...
IBM.2107-75CYM31/6001:IBM.2107-75CYK71/7001 Suspended Multi-target Internal Global Copy 81920 ...
IBM.2107-75CYM31/6002:IBM.2107-75CYK71/7002 Suspended Multi-target Internal Global Copy 81920 ...
...

dscli> lspprc -l -dev IBM.2107-75CYK71 -remotedev IBM.2107-75CYM31 -multtgt 7000-700f:6000-600f
ID                                        State     Reason               Type        Out Of Sync Tracks
=============================================================================================================
IBM.2107-75CYK71/7000:IBM.2107-75CYM31/6000 Suspended Multi-target Internal Global Copy 81920 ...
IBM.2107-75CYK71/7001:IBM.2107-75CYM31/6001 Suspended Multi-target Internal Global Copy 81920 ...
IBM.2107-75CYK71/7002:IBM.2107-75CYM31/6002 Suspended Multi-target Internal Global Copy 81920 ...
...

dscli> lspprc -l -dev IBM.2107-75CYM31 -remotedev IBM.2107-75CYK71 -multtgt 6000-600f:7000-700f
ID                                        State     Reason               Type        Out Of Sync Tracks
=============================================================================================================
IBM.2107-75CYM31/6000:IBM.2107-75CYK71/7000 Suspended Multi-target Internal Global Copy 112   ...
IBM.2107-75CYM31/6001:IBM.2107-75CYK71/7001 Suspended Multi-target Internal Global Copy 111   ...
IBM.2107-75CYM31/6002:IBM.2107-75CYK71/7002 Suspended Multi-target Internal Global Copy 111   ...
...

dscli> lspprc -l -dev IBM.2107-75CYK71 -remotedev IBM.2107-75CYM31 -multtgt 7000-700f:6000-600f
ID                                        State     Reason               Type        Out Of Sync Tracks
=============================================================================================================
IBM.2107-75CYK71/7000:IBM.2107-75CYM31/6000 Suspended Multi-target Internal Global Copy 137   ...
IBM.2107-75CYK71/7001:IBM.2107-75CYM31/6001 Suspended Multi-target Internal Global Copy 136   ...
IBM.2107-75CYK71/7002:IBM.2107-75CYM31/6002 Suspended Multi-target Internal Global Copy 136   ...
...
```

9.3.8 Move production to H1’

Either HyperSwap or a traditional restart can be used. The steps in this process follow.

Freeze H1:H1’ and H1:H2

Freeze commands for H1:H1’ and H1:H2 remove the PPRC paths, and suspend all H1:H1’ and H1:H2 Metro Mirror pairs. A separate freeze command is required for each LSS-to-LSS PPRC relationship. The use of consistency groups creates consistent data at the new H1’ site by using extended long busy or queue full to temporarily queue dependent writes. Alternatively, all application I/Os to the set of volumes being migrated can be quiesced.

Example 9-13 Freeze H1:H1’ and H1:H2

```
dscli> freezepprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYM31 50:60
CMUC00161I freezepprc: Remote Mirror and Copy consistency group IBM.2107-75CZM21/50:IBM.2107-75CYM31/60
successfully created.

dscli> freezepprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYK71 50:70
CMUC00161I freezepprc: Remote Mirror and Copy consistency group IBM.2107-75CZM21/50:IBM.2107-75CYK71/70
successfully created.
```

Unfreeze and run H1:H1' and H1:H2

When all freeze commands have completed, the secondary H1' and H2 sites contain consistent copies of data. The unfreeze or run command (also known as a *consistency group created* command) removes the extended long busy or queue-full condition at the local H1 site. An unfreeze command is required for each LSS-to-LSS PPRC relationship. A sample DS CLI unfreeze command is shown in Example 9-14.

Example 9-14 Unfreeze H1:H1' and H1:H2

```
dscli> unfreezepprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYM31 50:60
CMUC00198I unfreezepprc: Remote Mirror and Copy pair IBM.2107-75CZM21/50:IBM.2107-75CYM31/60 successfully
thawed.

dscli> unfreezepprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYK71 50:70
CMUC00198I unfreezepprc: Remote Mirror and Copy pair IBM.2107-75CZM21/50:IBM.2107-75CYK71/70 successfully
thawed.
```

Failover H1':H1 and H2:H1

The failover command for H1':H1 coverts the H1' volumes to suspended Metro Mirror primary volumes whose secondary volumes are H1. Similarly, the failover H2:H1 converts the H2 volumes to suspended primary volumes whose secondary volumes are H1.

Example 9-15 Failover H1':H1 and H2:H1

```
dscli> failoverpprc -dev IBM.2107-75CYK71 -remotedev IBM.2107-75CZM21 -type mmir 7000-700f:5000-500f
CMUC00196I failoverpprc: Remote Mirror and Copy pair IBM.2107-75CYK71/7000:IBM.2107-75CZM21/5000
successfully reversed.
...
CMUC00196I failoverpprc: Remote Mirror and Copy pair IBM.2107-75CYK71/700F:IBM.2107-75CZM21/500F
successfully reversed.

dscli> failoverpprc -dev IBM.2107-75CYM31 -remotedev IBM.2107-75CZM21 -type mmir 6000-600f:5000-500f
CMUC00196I failoverpprc: Remote Mirror and Copy pair IBM.2107-75CYM31/6000:IBM.2107-75CZM21/5000
successfully reversed.
...
CMUC00196I failoverpprc: Remote Mirror and Copy pair IBM.2107-75CYM31/600F:IBM.2107-75CZM21/500F
successfully reversed.
```

Note: The H1' site becomes the new production site, and the H1:H2 pairs are no longer used. If you wonder why a failover of H2:H1 is performed rather than terminating the H1:H2 pair, the reason for performing a failover H2:H1 is that a termination of H1:H2 would result in ending the Multiple Target PPRC configuration. This, in turn, would cause the MTIR pairs between H1' and H2 to be removed. These pairs are required to perform the Incremental Resynchronization from H1':H2.

Resume host systems at H1'

Host systems can now be restarted at H1'. If HyperSwap is used, the host systems are switched to H1'.

At this point, the host systems are running at H2 and the configuration appears as in Figure 9-6 on page 113.

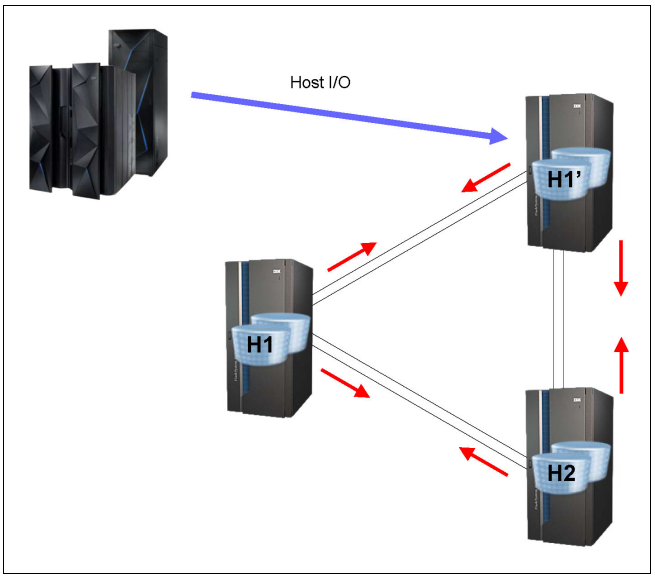

Figure 9-6 Move production to H1'

The H1, H2, and H3 sites are all suspended primary volumes to each of the other sites.

9.3.9 Start replication H1':H2

The Incremental Resynchronization capability is used to start replication for H1':H2. With the change recording function of Multiple Target PPRC, only the out-of-synchronization tracks between H1' and H2 are transferred, and a full copy of data is not required.

Failback H1':H2

The failback command for H1':H2 merges the change recording bitmap files from H1' and H2 and initiates the transfer of the out-of-sync data. A sample DS CLI command for this failback is shown in Example 9-16.

Example 9-16 Failback H1':H2

```
dscli> failbackpprc -dev IBM.2107-75CYK71 -remotedev IBM.2107-75CYM31 -type mmir 7000-700f:6000-600f
CMUC00197I failbackpprc: Remote Mirror and Copy pair IBM.2107-75CYK71/7000:IBM.2107-75CYM31/6000
successfully failed back.
...
CMUC00197I failbackpprc: Remote Mirror and Copy pair IBM.2107-75CYK71/700F:IBM.2107-75CYM31/600F
successfully failed back.
```

The pairs reach the full duplex state, as shown in the DS CLI **lspprc** command of Example 9-17.

Example 9-17 Query showing H1':H2 pairs full duplex

```
dscli> lspprc -dev IBM.2107-75CYK71 -remotedev IBM.2107-75CYM31 -multtgt 7000-700f:6000-600f
ID                                        State       Reason Type         SourceLSS
==========================================================================================================
IBM.2107-75CYK71/7000:IBM.2107-75CYM31/6000 Full Duplex -      Metro Mirror IBM.2107-75CYK71/70 ...
...
IBM.2107-75CYK71/700F:IBM.2107-75CYM31/600F Full Duplex -      Metro Mirror IBM.2107-75CYK71/70 ...
```

The state of the configuration at this point is shown in Figure 9-7.

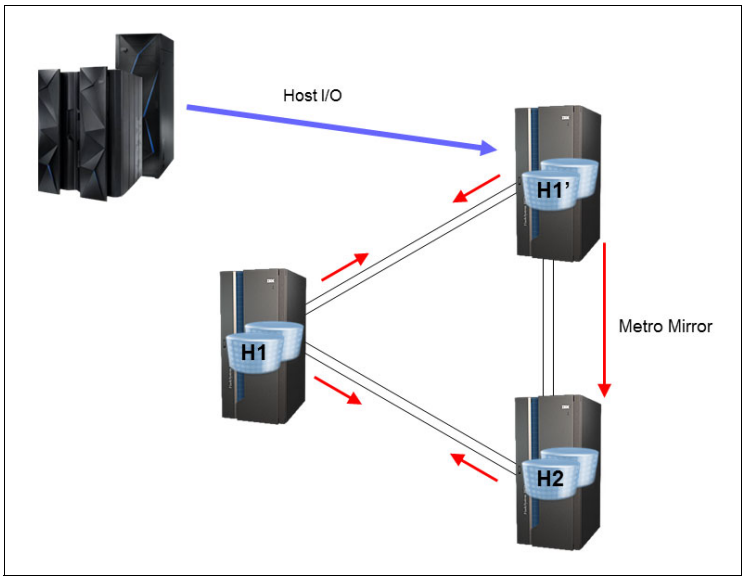

Figure 9-7 After failback H1' to H2

Metro Mirror is active from the new H1' site to the secondary H2 site.

9.3.10 Remove Metro Mirror pairs for H1

In the case of a storage system replacement, the original H1 is no longer being used and can be removed. In some instances, a migration may be performed to move only some data from the storage system, in which case other volumes at the H1site might still be in use.

For the LSSes and LCUs that are no longer used for PPRC replication, the pairs and paths should be removed.

In Figure 9-7, remove the unused pairs between H1 and H1' and also between H1 and H2.

- Remove H1:H1' Metro Mirror pairs
- Remove H1':H1 Metro Mirror pairs
- Remove H1:H2 Metro Mirror pairs
- Remove H2:H1 Metro Mirror pairs

Note: The PPRC paths do not exist between all of the storage systems. Removing PPRC when the paths do not exist results in warning messages. In this scenario, these messages are expected.

To keep the diagrams in this paper clear, the majority of diagrams do not indicate the state and direction of the PPRC paths. For most Multiple Target PPRC uses, it is best to establish active paths in both directions between the different sites involved in the replication. In some migration cases, however, not all of these paths are required, because the intermediate configurations are temporary and there is no requirement to establish active mirroring in all possible directions.

Figure 9-8 on page 115 shows the existing PPRC path relationships at this point in the migration scenario. It is important to understand which sets of paths exists so that a complete removal can be performed.

Refer to 3.3, “PPRC path states” on page 38 for a detailed description of how the PPRC path states are affected by freeze and PPRC failover commands.

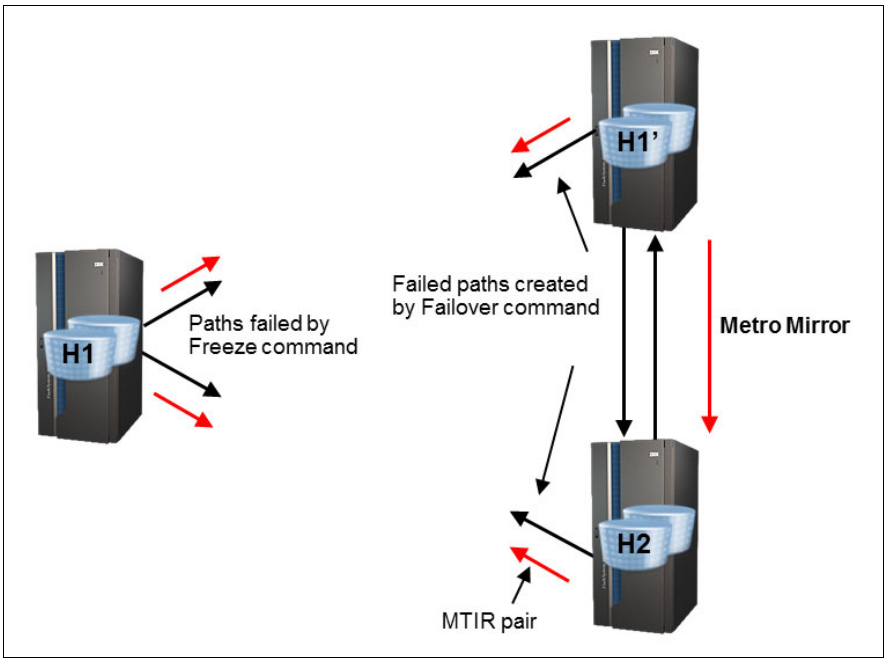

Figure 9-8 Path and pair relationships between H1, H2 and H1’

Notice that H1’ is a Multiple Target PPRC primary to H1 and H2. Because it is a Multiple Target PPRC primary, it sends commands to create the MTIR pairs H1:H2 and H2:H1. However, as explained in 3.2.2, “Delayed creation of MTIR pairs” on page 34, because there is no communication from H1’ to H1, the command cannot be sent to H1 and the suspended H1:H2 pairs remain suspended because of the freeze.

This is shown in the DS CLI **lspprc** output in Example 9-18. The H1:H2 pairs are suspended because of the freeze, and the H2:H1 pairs have been converted to MTIR pairs.

Example 9-18 Query H1:H2 and H2:H1 pairs

```
dscli> lspprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYM31 -multtgt 5000-500f:6000-600f
ID                                        State     Reason Type         SourceLSS
==========================================================================================
IBM.2107-75CZM21/5000:IBM.2107-75CYM31/6000 Suspended Freeze Metro Mirror IBM.2107-75CZM21/50  ...
...
IBM.2107-75CZM21/500F:IBM.2107-75CYM31/600F Suspended Freeze Metro Mirror IBM.2107-75CZM21/50  ...


dscli> lspprc -dev IBM.2107-75CYM31 -remotedev IBM.2107-75CZM21 -multtgt 6000-600f:5000-500f
ID                                        State     Reason                Type        SourceLSS
=============================================================================================================
IBM.2107-75CYM31/6000:IBM.2107-75CZM21/5000 Suspended Multi-target Internal Global Copy IBM.2107-75CYM31/60
...
IBM.2107-75CYM31/600F:IBM.2107-75CZM21/500F Suspended Multi-target Internal Global Copy IBM.2107-75CYM31/60
```

Remove pairs H1':H1

The H1':H1 pairs were created by the failover H1':H1 and can be removed, as shown in Example 9-19.

Example 9-19 Remove pairs H1':H1

```
dscli> rmpprc -dev IBM.2107-75CYK71 -remotedev IBM.2107-75CZM21 -quiet 7000-700f:5000-500f
CMUN03012E rmpprc: IBM.2107-75CYK71/7000:IBM.2107-75CZM21/5000: An error occurred for a Remote Mirror and
Copy (RMC) operation or its channel extender path which can result in the failure of a command
 or the incomplete removal of a relationship pair.
...
CMUN03012E rmpprc: IBM.2107-75CYK71/700F:IBM.2107-75CZM21/500F: An error occurred for a Remote Mirror and
Copy (RMC) operation or its channel extender path which can result in the failure of a command
 or the incomplete removal of a relationship pair.
```

The error messages are to inform the user that there is no communication from H1' to H1 so the pair was removed only at the local H1' site. This is expected in this case because it is only the H1':H1 relationship at H1' that is being removed.

The removal of the H1':H1 pair causes H1 to no longer be a Multiple Target PPRC primary, so it will remove the MTIR pairs between the two secondary sites, H1 and H2. There is active communication for H1':H2, and the MTIR pairs from H2:H1 are removed.

Remove H1:H1' and H1:H2 pairs

Example 9-20 shows sample DS CLI commands to remove the H1:H1' and H1:H2 pairs. As with the removal of H1':H1 pairs, because there are no active PPRC paths between the storage systems, the error messages are expected.

Example 9-20 Remove pairs H1:H1'

```
dscli> rmpprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYK71 -quiet 5000-500f:7000-700f
CMUN03012E rmpprc: IBM.2107-75CZM21/5000:IBM.2107-75CYK71/7000: An error occurred for a Remote Mirror and
Copy (RMC) operation or its channel extender path which can result in the failure of a command
 or the incomplete removal of a relationship pair.
...
CMUN03012E rmpprc: IBM.2107-75CZM21/500F:IBM.2107-75CYK71/700F: An error occurred for a Remote Mirror and
Copy (RMC) operation or its channel extender path which can result in the failure of a command
 or the incomplete removal of a relationship pair.

dscli> rmpprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYM31 -quiet 5000-500f:6000-600f
CMUN03012E rmpprc: IBM.2107-75CZM21/5000:IBM.2107-75CYM31/6000: An error occurred for a Remote Mirror and
Copy (RMC) operation or its channel extender path which can result in the failure of a command
 or the incomplete removal of a relationship pair.
 ...
CMUN03012E rmpprc: IBM.2107-75CZM21/500F:IBM.2107-75CYM31/600F: An error occurred for a Remote Mirror and
Copy (RMC) operation or its channel extender path which can result in the failure of a command
 or the incomplete removal of a relationship pair.
```

At this point, all of the pairs associated with the original H1 site have been removed.

9.3.11 Remove PPRC paths for H1

If all Metro Mirror volumes have been migrated off of H1, the PPRC paths between H1 and H1' and between H1 and H2 can be removed. If some volumes on H1 are still in Metro Mirror relationships to H2, the paths could still be required. These are the commands:

- Remove PPRC paths for H1:H2
- Remove PPRC paths for H1:H1'
- Remove PPRC paths for H2:H1
- Remove PPRC paths for H1':H1

Example 9-21 Remove unused paths

```
dscli> rmpprcpath -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYM31 -remotewwnn 5005076305FFD71E -quiet 50:60
CMUC00150I rmpprcpath: Remote Mirror and Copy path IBM.2107-75CZM21/50:IBM.2107-75CYM31/60 successfully removed.

dscli> rmpprcpath -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYK71 -remotewwnn 5005076305FFD71A -quiet 50:70
CMUC00150I rmpprcpath: Remote Mirror and Copy path IBM.2107-75CZM21/50:IBM.2107-75CYK71/70 successfully removed.

dscli> rmpprcpath -dev IBM.2107-75CYM31 -remotedev IBM.2107-75CZM21 -remotewwnn 5005076305FFD75A -quiet 60:50
CMUC00150I rmpprcpath: Remote Mirror and Copy path IBM.2107-75CYM31/60:IBM.2107-75CZM21/50 successfully removed.

dscli> rmpprcpath -dev IBM.2107-75CYK71 -remotedev IBM.2107-75CZM21 -remotewwnn 5005076305FFD75A -quiet 70:50
CMUC00150I rmpprcpath: Remote Mirror and Copy path IBM.2107-75CYK71/70:IBM.2107-75CZM21/50 successfully removed.
```

At this point, only the H1': H2 relationships exist, and H1 can be removed from the Metro Mirror configuration, as shown in Figure 9-9.

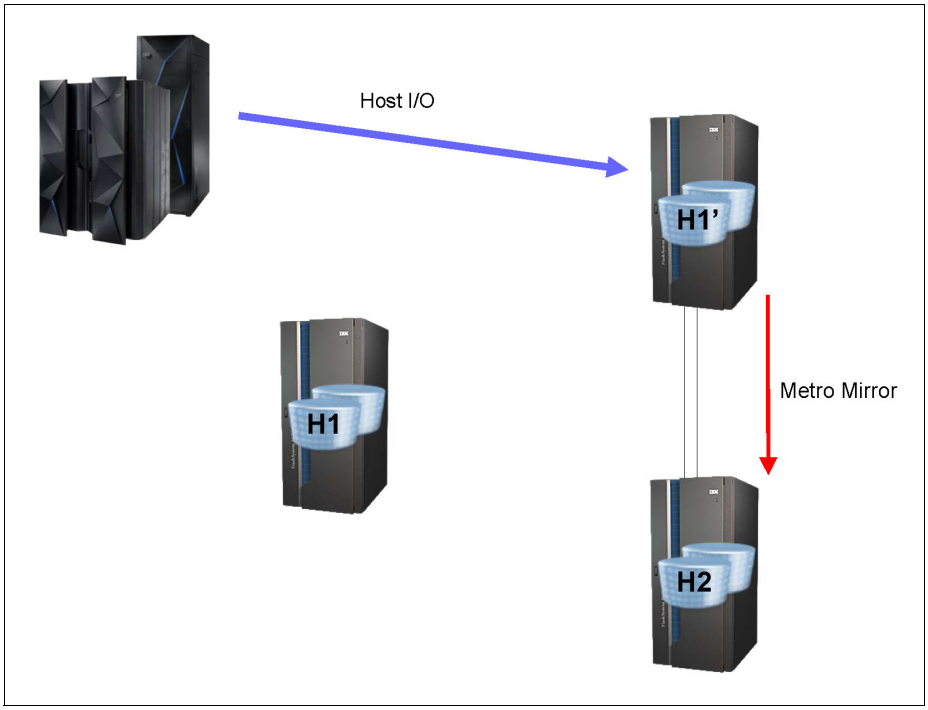

Figure 9-9 Migration from H1 to H1' is complete

9.4 Replacement of Metro Mirror primary, alternative method

An alternative method for replacing a Metro Mirror primary is to first switch the direction of the replication, and then follow the process for replacing a Metro Mirror secondary, as described in 9.2, “Replacement of Metro Mirror secondary” on page 102.

9.4.1 Initial configuration

The initial configuration consists of a Metro Mirror pair H1:H2, as shown in Figure 9-10.

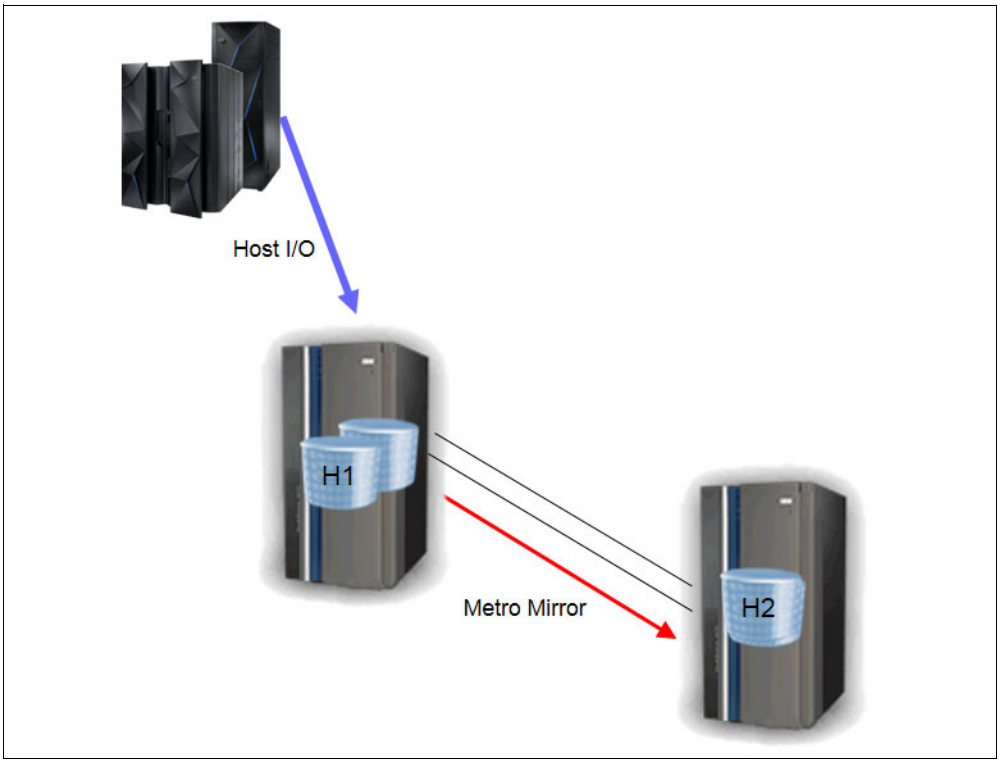

Figure 9-10 Initial configuration for migration of Metro Mirror primary

Site H1 contains the Metro Mirror primary volumes that are to be migrated to a different DS8870 storage system, and site H2 contains the Metro Mirror secondary volumes. H1 may be only some of the volumes on a DS8870 storage system or it may be volumes across multiple DS8870 storage systems.

The H1:H2 pairs are in the full-duplex state at the beginning of this scenario.

9.4.2 Move production to H2

Either HyperSwap or a traditional restart may be used to move the production workload to the H2 site. After the move to H2, a failback command H2:H1 restores active replication. Metro Mirror replication is now active in the reverse direction, as shown in Figure 9-11 on page 119.

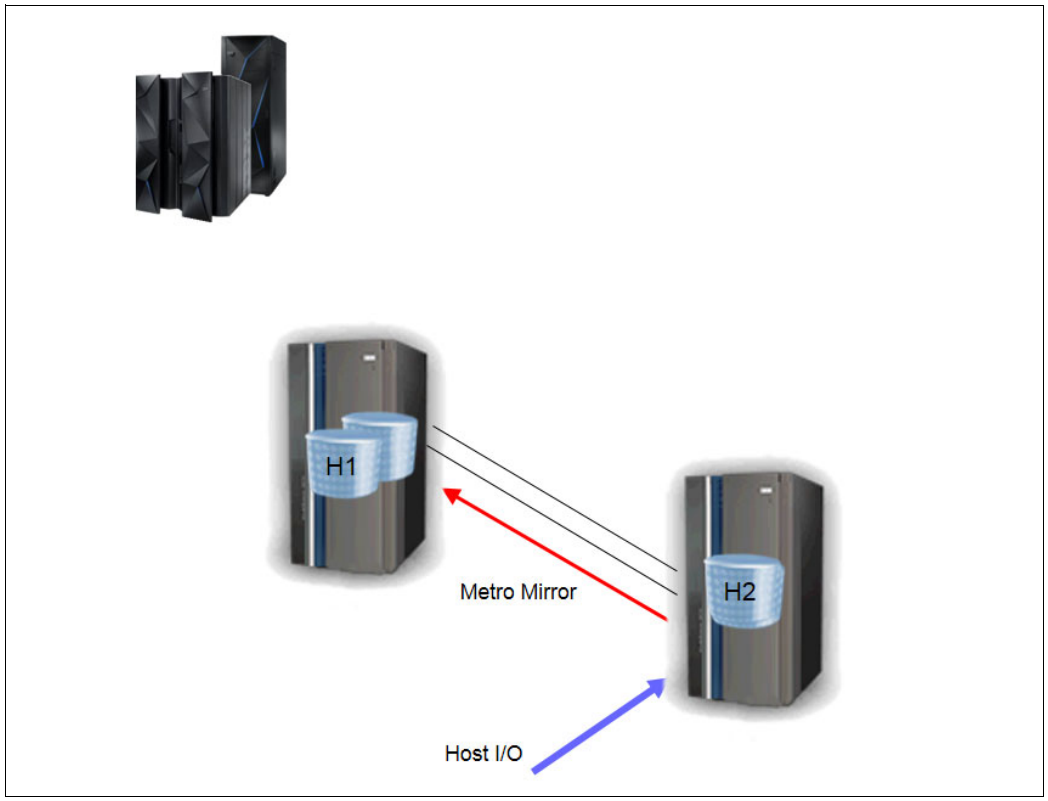

Figure 9-11 Direction of Metro Mirror is reversed

9.4.3 Installation of new DS8870 at H1'

Install, configure, and physically connect the new H1' storage system, including any host system requirements and the Fibre Channel connections to H2 to be used for the Metro Mirror paths.

9.4.4 Start Metro Mirror H2:H1'

PPRC paths are required H2:H1' for the replication of data from H2 to the new H1' system. The new Metro Mirror configuration will be H1' to H2, so PPRC paths are also required in the H1' to H2 direction.

Establish new Metro Mirror pairs H2:H1', and wait for the pairs to reach the full-duplex state. At the completion of this step, there is a Multiple Target PPRC topology of H2:H1 and H2:H1', as shown in Figure 9-12 on page 120.

Because this is now a Multiple Target PPRC topology, the storage systems automatically create MTIR pairs between the secondary volumes, H1' and H2. Because this migration scenario does not require an active relationship between H1' and H2, it is not necessary to establish PPRC paths between the two sites. This is one possible advantage over the method described in 9.3, "Replacement of Metro Mirror primary" on page 106.

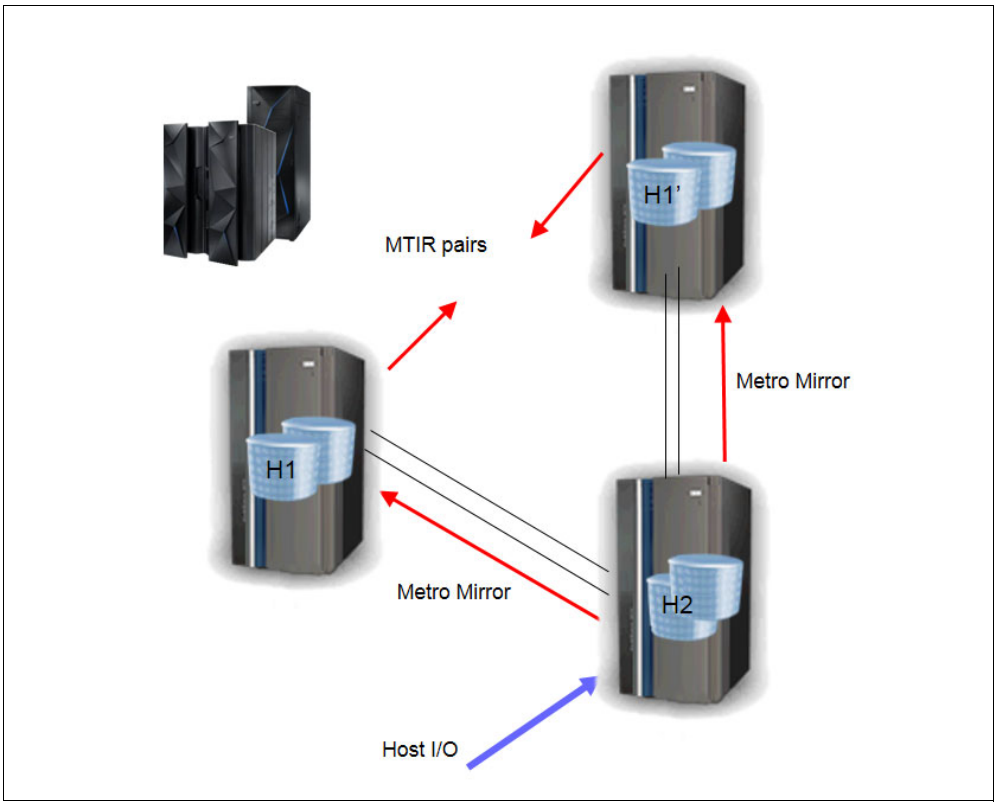

Figure 9-12 Multiple Target Metro Mirror H2:H1 and H2:H1'

9.4.5 Remove H2:H1

After all of the H2:H1' pairs have reached full duplex, the H2:H1 pairs can be removed because H2:H1' is now available to provide the Metro Mirror protection. When the H2:H1 pairs are removed, the storage system detects that there is no longer a Multiple Target PPRC topology and automatically removes the MTIR pairs between H1 and H1'. After all pairs have been removed, the PPRC paths between H1 and H2 can be removed.

9.4.6 Move production back to H1'

Either HyperSwap or a traditional restart may be used to move the production workload to the new H1' site. After the move to H1', a failback command H1':H2 restores active replication. The migration to H1' is now complete, as shown in Figure 9-13 on page 121.

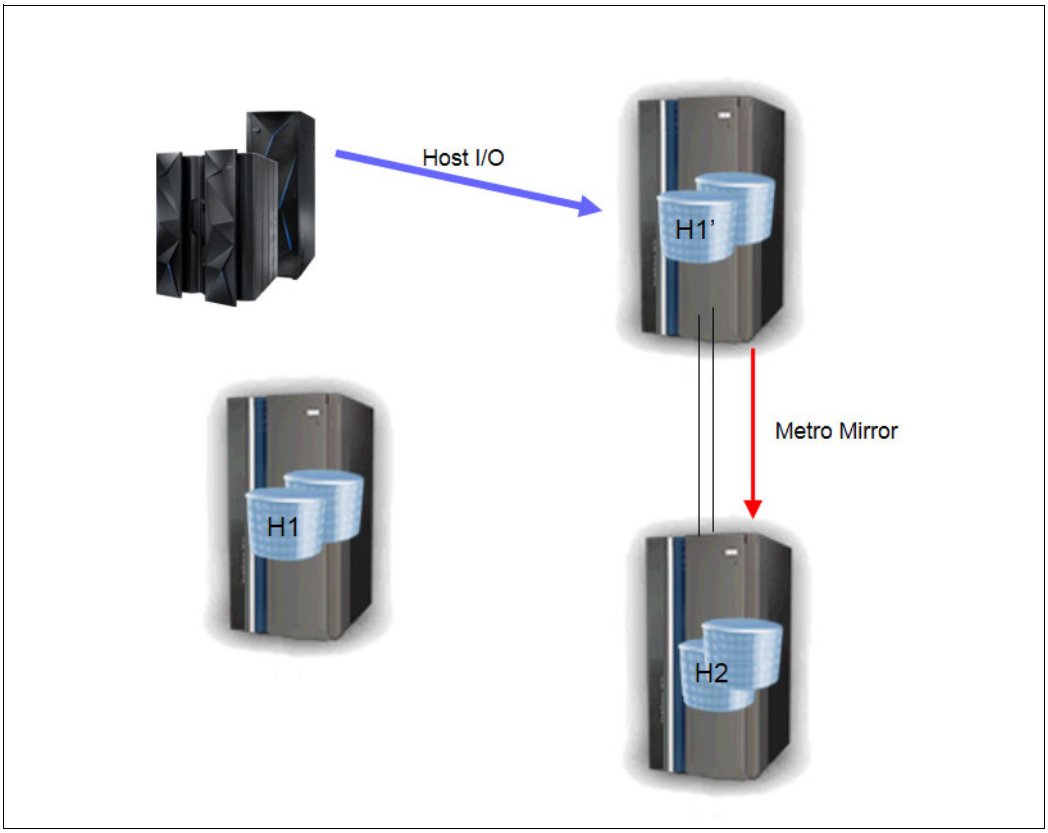

Figure 9-13 Migration to H1' is complete

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this paper.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- *IBM DS8870 Architecture and Implementation*, SG24-8085
- *IBM System Storage DS8000 Copy Services for IBM System z*, SG24-6787
- *IBM System Storage DS8000 Copy Services for Open Systems*, SG24-6788

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Other publications

These publications are also relevant as further information sources:

- *IBM System Storage DS Command-Line Interface User's Guide*, GC27-4212

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

IBM DS8870 Multiple Target Peer-to-Peer Remote Copy

Redpaper™

Enhanced resiliency functions

Architecture and design

Implementation and use

Multiple Target Peer-to-Peer Remote Copy (MT-PPRC) was introduced with IBM System Storage DS8870 Version 7 Release 4 (License Machine Code 7.7.40.xx.xx).

MT-PPRC provides the capability to have two PPRC relationships on a single primary volume. With this enhancement, there can be another target available to provide additional data protection to act as a backup if a disaster occurs.

This IBM Redpaper publication provides practical information about the characteristics and functions of the DS8000 Multiple Target Peer-to-Peer Remote Copy (MT-PPRC). The paper includes scenarios that illustrate the implementation and use of different topologies, as well as enhanced migration scenarios.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

REDP-5151-00