Redbooks

ibm.com/redbooks

DS8870 Data Migration Techniques

Alexander Warmuth

Bertrand Dufrasne

Joachim Appel

Werner Bauer

Susan Douglass

Peter Klee

Miroslaw Pura

Mark Wells

Bjoern Wesselbaum

Storage

IBM

Redbooks

International Technical Support Organization

DS8870 Data Migration Techniques

April 2015

SG24-8257-00

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

First Edition (April 2015)

This edition applies to IBM System Storage® DS8000 series with DS8000 LMC 7.7.40.xx.xx (bundle version 87.40.xxx.xx), with z/OS version 2.1 and Softek TDMF for z/OS Version 5.5.0.

© Copyright International Business Machines Corporation 2015. All rights reserved.
Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices . . . vii
Trademarks . . . viii

Preface . . . ix
Authors . . . ix
Now you can become a published author, too! . . . xi
Comments welcome . . . xi
Stay connected to IBM Redbooks . . . xi

Chapter 1. Introducing disk data migration . . . 1
1.1 Overview . . . 2
1.2 Data migration and its business reasons . . . 2
1.3 Aspects of data migration . . . 3
1.4 Definition of migration processes . . . 4
1.5 Planning phase . . . 4
1.5.1 Analyzing the current situation . . . 5
1.5.2 Designing and developing migration scenarios . . . 7
1.5.3 Testing . . . 11
1.5.4 Supporting tools and products . . . 12
1.5.5 Migration plan . . . 12
1.6 Migration phase . . . 14
1.7 Post migration phase . . . 15
1.7.1 Validation . . . 15
1.7.2 Production take over . . . 16
1.7.3 Cleaning up . . . 16
1.7.4 Reports and documentation . . . 16

Part 1. Copy Services based Migration Scenarios . . . 17

Chapter 2. Copy Services based data migrations . . . 19
2.1 Terms . . . 20
2.1.1 Volume set . . . 20
2.1.2 Data replication and cut over . . . 20
2.1.3 Site . . . 20
2.1.4 Out-of-sync tracks . . . 20
2.1.5 Primary and secondary volumes . . . 21
2.1.6 Conventions in diagrams . . . 21
2.2 Requirements . . . 21
2.3 Data replication . . . 22
2.3.1 Set up Global Copy . . . 22
2.3.2 Monitor Global Copy progress . . . 23
2.3.3 Prepare for cut over . . . 24
2.4 Data validation testing . . . 24
2.4.1 Host considerations . . . 24
2.4.2 Prepare the new volume set for testing . . . 25
2.4.3 Data validation testing . . . 26
2.4.4 Data validation testing complete . . . 26
2.5 Tivoli Storage Productivity Center for Replication considerations . . . 27
2.5.1 Can Tivoli Storage Productivity Center for Replication help setting up the required

© Copyright IBM Corp. 2015. All rights reserved.

replication? 28
2.5.2 Bring new volume set into the existing Tivoli Storage Productivity Center for Replication configuration 28
2.6 HyperSwap considerations 29
2.7 Multiple Target PPRC considerations 30

Chapter 3. Data Migration with no pre-existing copy services replication 31
3.1 Initial configuration 32
3.2 Migration process 32
3.2.1 General considerations 32
3.2.2 Migrating source volume set 34
3.2.3 Using Tivoli Storage Productivity Center for Replication. 36

Chapter 4. Data Migration with 2-site Metro Mirror replication 39
4.1 Initial configuration 40
4.2 Migration process 40
4.2.1 General considerations 41
4.2.2 Replacing secondary 41
4.2.3 Replacing primary 50
4.2.4 Replacing primary and secondary DS8000. 53
4.3 Geographically Dispersed Parallel Sysplex (GDPS) considerations. 56

Chapter 5. Data Migration with 3-site Metro Mirror replication 71
5.1 Initial configuration 72
5.2 Migration process 72
5.2.1 Replacing site 2 or site 3 volume set 72
5.2.2 Replacing the site 1 volume set 76
5.2.3 Example for cut-over sequence 78
5.2.4 HyperSwap considerations 82

Chapter 6. Data Migration with 2-site Global Mirror replication 85
6.1 Initial configuration 86
6.2 Migration process 86
6.2.1 General considerations 86
6.2.2 Replacing auxiliary storage. 86
6.2.3 Replacing primary storage 91
6.2.4 Replacing primary and auxiliary storage. 94
6.3 GDPS considerations 97

Chapter 7. Data Migration with Metro/Global Mirror replication 105
7.1 Initial configuration 106
7.2 Migration process 106
7.2.1 General considerations 107
7.3 Migration Scenarios 108
7.3.1 Replacing the C volume set without automation. 108
7.3.2 Replacing the B volume set without automation 114
7.3.3 Replacing the A volume set without automation 121
7.3.4 Replacing the A and B volume set without automation. 126
7.3.5 Replace all volume sets without automation. 132
7.4 Geographically Dispersed Parallel Sysplex (GDPS) considerations. 137

Chapter 8. Data Migration with 2-site z/OS Global Mirror replication 155
8.1 Initial configuration 156
8.2 Migration process 157

8.2.1 General considerations . . . 157
8.2.2 zGM running in migration mode . . . 157
8.2.3 Replacing primary storage system in a zGM DR configuration. . . . 159
8.2.4 Replacing auxiliary storage system within a zGM DR configuration . . . 169

Chapter 9. Data Migration with GDPS/Metro and z/Global Mirror replication . . . 185
9.1 GDPS Terminology . . . 186
9.1.1 Site-1 and Site-2 . . . 186
9.1.2 PPRC Controlling System . . . 186
9.1.3 Application system . . . 186
9.1.4 PPRC Master System . . . 186
9.1.5 PPRC Freeze . . . 186
9.1.6 HyperSwap . . . 187
9.2 Initial configuration . . . 188
9.3 Migration process . . . 188
9.3.1 General considerations . . . 189
9.4 Migration Scenarios . . . 190
9.4.1 Replacing the zGM secondary . . . 190
9.4.2 Replacing the current Metro Mirror secondary . . . 193
9.4.3 Replacing the current Metro Mirror primary . . . 208

Part 2. z/OS based Migration Scenarios . . . 209

Chapter 10. Using TDMF z/OS . . . 211
10.1 General information about TDMF z/OS. . . . 212
10.1.1 Provided Functions . . . 212
10.1.2 Terminology . . . 212
10.2 Using TDMF z/OS in a local shared storage environment . . . 213
10.2.1 Overview of TDMF z/OS . . . 213
10.2.2 Example migration scenario in a shared storage environment . . . 223
10.2.3 Tips and preferred practices . . . 234
10.3 Using TDMF z/OS in a remote environment . . . 246
10.3.1 Overview of Remote Replication with TDMF z/OS . . . 246
10.3.2 General guidelines for using TDMF in a remote environment. . . . 250
10.3.3 Example Remote Replication scenario with TDMF z/OS . . . 256

Chapter 11. z/OS Dataset Mobility Facility . . . 271
11.1 z/OS Dataset Mobility Facility overview . . . 272
11.2 Terminology . . . 273
11.3 zDMF migration process . . . 273
11.3.1 zDMF migration steps . . . 273
11.3.2 Planning the migration . . . 274
11.3.3 Pre-migration tasks . . . 275
11.3.4 Tools for migrating data sets. . . . 275
11.3.5 zDMF migration planning . . . 277
11.3.6 Typical migration scenario . . . 279
11.3.7 Migration performance and scalability. . . . 282
11.3.8 Post-migration . . . 285
11.4 Important zDMF-related topics . . . 286
11.4.1 Understanding the catalog structure . . . 286
11.4.2 Catalog diagnostic recommendations . . . 286
11.4.3 Global resource serialization (GRS) . . . 288
11.5 Migration in an experimental shared DB2 test environment . . . 289
11.5.1 Overview of experimental environment. . . . 290

11.5.2 Preparing a list of data sets to be migrated. 293
11.5.3 Creating JCL for the zDMF batch utility 294
11.5.4 Starting and Monitoring the migration 297
11.5.5 Diverting the Group 302
11.5.6 Performing post migration cleanup 307

Related publications 311
IBM Redbooks 311
Other publications 311
Online resources 311
Help from IBM 311

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

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS 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.

© Copyright IBM Corp. 2015. All rights reserved.

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:

AIX®
CICS®
DB2®
DS8000®
FICON®
FlashCopy®
GDPS®
Geographically Dispersed Parallel Sysplex™
HyperSwap®
IBM®
MVS™
NetView®
OS/390®
Parallel Sysplex®
Power Systems™
RACF®
Redbooks®
Redbooks (logo) ®
RMF™
S/390®
System Storage®
System z®
TDMF®
Tivoli®
z/OS®
z/VM®
z/VSE®

The following terms are trademarks of other companies:

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

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

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

Preface

This IBM® Redbooks® publication describes data migrations between IBM DS8000® storage systems, where in most cases one or more older DS8000 models are being replaced by the newer DS8870 model. Most of the migration methods are based on the DS8000 Copy Services.

The book includes considerations for solutions such as IBM Tivoli® Productivity Center for Replication and the IBM Geographically Dispersed Parallel Sysplex™ (GDPS®) used in IBM z/OS® environments. Both offerings are primarily designed to enable a disaster recovery using DS8000 Copy Services. In most data migration cases, Tivoli Productivity Center for Replication or GDPS will not directly provide functions for the data migration itself. However, this book explains how to bring the new migrated environment back into the control of GDPS or Tivoli Productivity Center for Replication.

In addition to the Copy Services based migrations, the book also covers host-based mirroring techniques, using IBM Transparent Data Migration Facility (TDMF®) for z/OS and the z/OS Dataset Mobility Facility (zDMF).

Authors

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

Alexander Warmuth is a Senior IT Specialist for IBM at the European Storage Competence Center. Working in technical sales support, he designs and promotes new and complex storage solutions, drives the introduction of new products, and provides advice to clients, business partners, and sales. His main areas of expertise are high-end storage solutions and business resiliency. He joined IBM in 1993, and has been working in technical sales support since 2001. Alexander holds a diploma in Electrical Engineering from the University of Erlangen, Germany.

Bertrand Dufrasne is an IBM Certified IT Specialist and Project Leader for IBM System Storage® disk products at the ITSO, San Jose Center. He has worked at IBM in various IT areas. He has written many IBM publications, and has developed and taught technical workshops. Before joining the ITSO, he worked for IBM Global Services as an Application Architect. He holds a Master's degree in Electrical Engineering.

Joachim Appel is a z/OS Senior Software Engineer with the IBM Systems and Technology Group, for the TDMF/zDMF development team. He started his career in the mainframe systems programming department of a large German bank in 1983. Later he was a software engineer in the areas of tape robotic systems and mainframe storage management before he joined IBM in 2008.

Werner Bauer is a Certified Consulting IT Specialist in Germany. He has more than 30 years of experience in storage software and hardware, and with IBM S/390® and IBM z/OS. His areas of expertise include DR solutions based on IBM enterprise disk storage systems. Werner is a frequent speaker at storage conferences and GUIDE SHARE Europe (GSE) meetings. He has also written extensively in various IBM publications about the DS8000. He holds a degree in Economics from the University of Heidelberg, and in Mechanical

© Copyright IBM Corp. 2015. All rights reserved.

Engineering from FH Heilbronn. He currently works with System Vertrieb Alexander (SVA), an IBM Premier Business Partner.

Susan Douglass is a Senior IT Specialist for IBM based in the United States within the IBM Resiliency Services Group. She has implemented GDPS in numerous global organizations. Her main area of expertise is in supporting complex resiliency and high availability solutions. Before rejoining IBM in 2009, she had over 20 years of software and hardware experience in both the financial services and global distribution industries.

Peter Klee is an IBM Professional Certified IT specialist in IBM Germany. He has 20 years of experience in Open Systems platforms, SAN networks, and high end storage systems in huge data centers. He formerly worked in a large bank in Germany where he was responsible for the architecture and the implementation of the disk storage environment. This environment included installation from various storage vendors, including EMC, HDS, and IBM. He joined IBM in 2003, where he worked for Strategic Outsourcing. Since July 2004, he has worked for ATS System Storage Europe in Mainz. His main focuses are data migration, disaster recovery, and storage architecture.

Miroslaw Pura is an IBM Certified Senior IT Specialist working in IBM Poland since 2003. He provides technical presales support to IBM clients in complex enterprise storage environments. Before joining IBM, Miroslaw worked for over 15 years as a z/OS mainframe Systems Programmer. He holds an Engineer and Master of Science degree in Fundamental Technological Research at the University of Science and Technology in Cracov, Poland.

Mark Wells is an Executive IT Specialist working in the United States. As a Client Technical Specialist since 2005, he provides storage solution design, technical consulting, and implementation support to his customers. He has extensive experience on the DS8000 platform in large-scale replication environments. Mark has worked for IBM since 1997 in various roles and has worked in the IBM System z® and Enterprise environment for over 25 years.

Bjoern Wesselbaum has been an accredited Senior Product Services Specialist for IBM Germany since 1998. He has worked as a Product Field Engineer for IBM Storage Systems within the EMEA Advanced Technical Support and Product Field Engineering support team at the Storage Competence Center in Mainz. Since 2007, Bjoern has also taught IBM DS8000 Implementation and DS8000 Copy Services classes for IBM Learning Services. His areas of expertise include IBM AIX® and other UNIX-like operating systems, storage attachment, Copy Services Solutions for open systems, and DS8000 management and operations, including Remote Support Connectivity Solutions for the DS8000.

Thanks to the following people for their contributions to this project:

Charlie Burger
Nick Clayton
Ingo Dimmer
Paul Hallam
David Petersen
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 at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

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

  ibm.com/redbooks

- Send your comments in an email to:

  redbooks@us.ibm.com

- Mail your comments to:

  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 disk data migration

This chapter defines disk data migration and examines the business reasons for migrating data. It contains the following sections:

- Overview
- Data migration and its business reasons
- Aspects of data migration
- Definition of migration processes
- Planning phase
- Migration phase
- Post migration phase

© Copyright IBM Corp. 2015. All rights reserved.

1.1 Overview

Data migration is usually not as common a task as those that occur daily in a data center. Data migration typically becomes necessary when physical changes in the data center affect the storage infrastructure. Here are some situations that necessitate a migration:

- One or more storage systems must be replaced. The replacement might be indicated because the current storage system runs out of capacity or performance is degraded, or even the financial amortization for one or more storage systems is about to be completed.
- The physical footprint of the storage environment must be reduced. Newer storage system models can help to consolidate many older storage systems into one system that offers higher capacity, smaller physical dimensions, and less power consumption.
- A physical move to a new location requires the migration of data to new storage systems. Physical moves can be indicated when a new data center building is deployed or when parts of a data center facility are closed or maintained for a considerable period.

These examples also show that the data migration itself is typically only one part of a much larger ongoing project. However, when a data migration is planned, it ends up becoming a substantial project on its own and for which substantial extra skills and resources can be required.

This book provides advice based on practical experience gained with many customer engagements. It describes methods and tools that can be used for a successful data migration. It also provides guidance on how to structure a data migration project.

This book focuses mainly on data migrations between DS8000 storage systems. In most of these migrations, one or more older DS8000 models are replaced by newer DS8000 models. Most of the migration methods in this book are based on the DS8000 Copy Services. With DS8870 release 7.4, the Copy Services capabilities are expanded with the Multiple Target Peer-to-Peer Remote Copy (MT-PPRC) function. MT-PPRC allows for a number of powerful and more flexible data migration scenarios.

The book also describes software tools that are helpful in using Copy Services to perform data migrations. The most relevant software tools are Tivoli Storage Productivity Center for Replication and the IBM Geographically Dispersed Parallel Sysplex (GDPS) which is used in z/OS environments. Both software tools are designed to help validate disaster recovery scenarios using DS8000 Copy Services. Thus, in most data migration cases Tivoli Storage Productivity Center for Replication and GDPS will not directly provide functions for the data migration itself. However, this book shows for almost every example, the process for bringing the newly migrated environment back into the control of Tivoli Storage Productivity Center for Replication and GDPS.

Data migrations can also be conducted using host-based mirroring techniques. This book describes in Chapter 10, “Using TDMF z/OS” on page 211 and Chapter 11, “z/OS Dataset Mobility Facility” on page 271 two methods of host-based mirroring techniques for z/OS based environments.

1.2 Data migration and its business reasons

Data migration is the transferring of data between storage system types, formats, or computer systems. These systems can be single or multiple, and similar or dissimilar. Migration is usually performed under program control to try to make movement of data as complete and as automated as is possible.

Data migration is needed when you change computer systems or upgrade to new products. As with any data center change, you want to avoid disrupting or disabling active applications.

You might migrate data for the following reasons:

- Your organization wants to use new storage devices to increase the size of your online storage or for more flexibility.
- You want to resolve performance issues.
- You need to physically reduce the footprint of your storage subsystem within the data center to release space. Reducing the footprint reduces the costs of power consumption and air conditioning.
- You want to employ data migration technologies to implement a disaster recovery solution.
- You need the new functions and facilities offered by evolving technology to stay competitive.
- You need to relocate your data center.

1.3 Aspects of data migration

Migrating data has the following general concerns, which require that you take a holistic view of your data center:

Configuration New hardware might need to be configured into the systems being migrated before you can start.

Operations A time needs to be found when the migration can be performed. Scheduling is especially important if you use offline data migration techniques that require the applications to be stopped. After the migration, your data center operating procedures will need to be reviewed and updated.

Data Not all data can be migrated in the same way. For example, local paging data sets on IBM OS/390® cannot be moved directly. The restriction is because manual handling of local paging data sets is required to keep OS/390 systems up and running.

Infrastructure Changes to infrastructure might be required to accommodate new storage subsystem technologies. For example, new technology can require new attachment methods or enhanced data rate capabilities.

Applications You might need to modify or add installation parameters in your applications to accommodate new storage subsystem devices or infrastructures.

Performance Migration of data is an opportunity to maximize performance by carefully relocating high use data. However, if you are not careful about where you migrate data, you can create performance issues.

1.4 Definition of migration processes

Experience in migration projects with various customers has defined a common structure for the data migration process. The intention is to simplify the setup of a data migration project. Figure 1-1 shows the three different phases that are common to every data migration process.

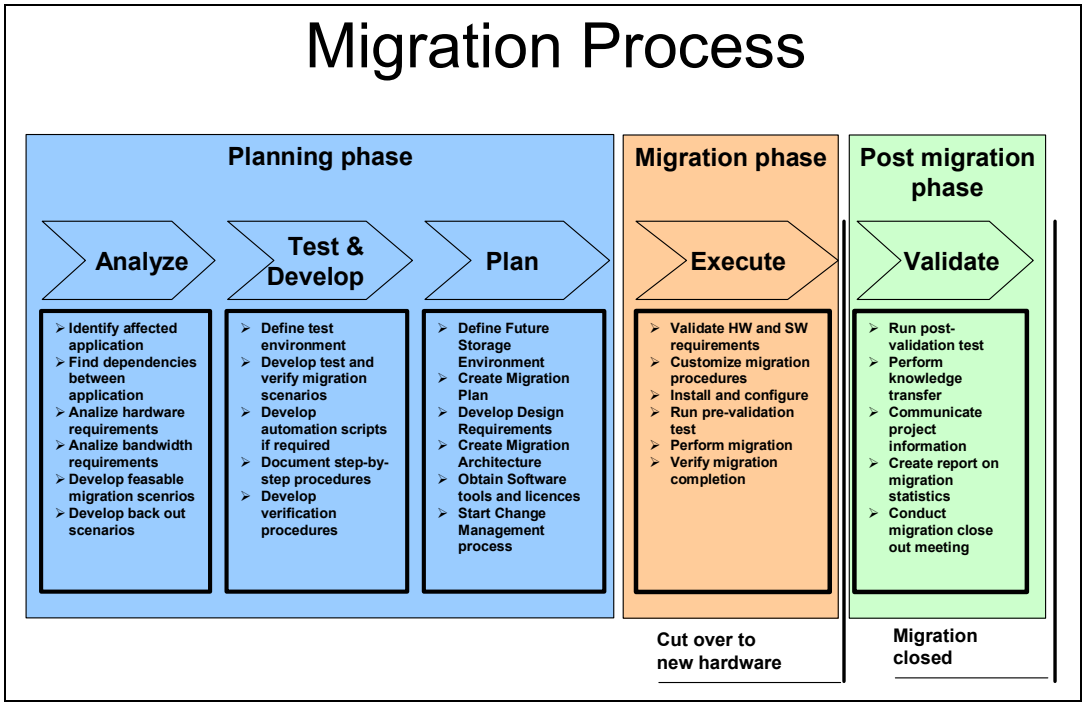

Figure 1-1 Suggested three-phase migration process

The planing phase starts with the analysis of the current situation before the migration. This phase includes determining which applications are affected and what the dependencies, requirements, and constraints are. In the next step of the planning phase, the migration scenarios are developed and tested. Finally, gather all results and findings in a project plan.

In the migration phase, the data migration is run using the tools and running scripts that were developed in the previous phase. The result of the migration phase is the cut over of the production to the new environment.

In the validation phase, the data migration is checked for completeness and to make sure that all applications are working correctly in the new environment. If successful, the production is taken over by the new environment. Otherwise, the migration is backed out to the original configuration and the root cause analysis must clarify what went wrong. When the findings indicate what must be corrected, those issues can be corrected and a new attempt can be started with the migration phase.

1.5 Planning phase

A successful data migration always requires substantial time spent on evaluation and planning. Adequate planning is the critical success factor in any migration project.

The higher the complexity of an environment and the more critical the data, the more important migration planning becomes. Careful migration planning identifies potential issues, allowing you to avoid or mitigate them. Migration planning also identifies which data to migrate first, which applications must be taken offline, and the internal and external colleagues you need to communicate with.

The planning process must cover the following considerations:

- Determine the requirements for the migration
- Assess the current storage environment
- Identify requirements for the future state of the environment
- Create a migration plan
- Document architectural decisions for specific hardware and software components
- Develop migration procedures
- Develop test plans

1.5.1 Analyzing the current situation

The first action must be a careful analysis of the current situation. In this phase, the details of the data migration subject must be worked out. To design the migration scenarios, as detailed in 1.5.2, “Designing and developing migration scenarios” on page 7, the technical environment and the infrastructure must be analyzed. Finally, the expected target solution must also be clearly defined.

Defining the migration subject

Data migration can at first seem to be an obvious task, where data must simply be transferred to a new storage hardware. However, when the subject of the data migration and its implications are examined in detail, many additional aspects and special considerations will arise, such as these examples:

- One reason for data migration can be that the current storage system must be returned because the leasing period is about to end. The return date for the storage device might be fixed. This example has a direct impact on the migration schedule because the storage system must be freed up completely by the fixed date. Other tasks can also include a secure data override of the data in the current system.
- Data migration that includes a consolidation of multiple storage systems can become quite complex because the resulting environment can be completely different. The new environment can include new host and server connectivity, which implies that the storage network must be changed as well. Migration projects that include consolidation can affect many departments and functions.
- There are in principle two approaches to conducting a migration. The fastest approach is to switch over to the new storage environment in one go, where all applications are taken over at the same time. This approach is common in mainframe environments. However, in Open System environments where many different applications are distributed across many different servers, the migration process is divided into a series of partial migrations. The staged approach implies a longer run time for the whole project.

Assessment of the environment

Migrating data to new hardware always has an impact on one or more applications because data usually are related to certain applications. Furthermore, many applications have dependencies on other applications, some of which might not be obvious. These dependencies can be at any layer of the application stack. For example, a storage device that requires a physical move to another location affects application A. Because of the resulting bandwidth restriction, the application server must also be moved to the new location.

However, Application B must be on the same IP subnet as application A, so it must be migrated as well.

To uncover and understand the implications of such dependencies, carefully assess the complete environment using the following steps:

1. Identify all applications that are directly affected by the planned migration.
2. Find application dependencies throughout the entire application stack, all the way down to the hardware layer.
3. If any additional cross-application dependencies are identified, expand the migration scope to include the newly identified application.
4. Repeat the previous steps until no additional cross-application dependencies are uncovered.
5. Identify and document the requirements for a downtime of all the applications in scope.
6. Evaluate the requirements against any restrictions for application downtime.

More constraints might be identified during the assessment process, including the need for extra capacity, connectivity, network bandwidth, and so on. Discovering these additional requirements early is vital because it gives the project team more time to address them or develop alternatives.

Successful migration planning involves more than just the IT staff. The business owners of the applications and the data being migrated must also be included. The owners are the best resource for determining how important an application or set of data is to the business. Coordinate the migration with the application owners and define the counterparts for the communication during the migration. Uncoordinated tasks can cause problems during the migration.

In this phase, it is easy to forget other implications of the migration because many questions or problems only appear later in the project. Based on IBM practical experience, we have created this checklist that you can use at the beginning of a migration project:

- Are you aware of the potential impact of Global Copy to the current replication environment?
- What are the user and application tolerance for impact and interruption?
- Is there a wanted order of replacement for the volume sets? If so, what is it and why?
- Will the existing infrastructure (network, directors, ports, and so on) support the migration?
- What replication management tools are you currently using?
- Will any training be required to develop the necessary skills to successfully complete the migration?
- What is the timeline for having the migration completed? Is it reasonable?
- What is the capacity and performance of the new system versus the old system?
- Has a performance analysis been completed to validate that the new system will handle the workload as required?
- Is this a migration of a current volume set to a new volume set with equal capacity?
- Do you plan to consolidate storage systems?
- Is this a migration of multiple current volume sets to a single new volume set with greater capacity?
- Do you understand the importance of the LSS to LSS mapping in the current configuration?

- Do you understand the scope of the various DSCLI commands and how that affects the final command set used to run the migration? Some commands are limited to the LSSs, whereas some are limited to the volumes.
- Will there be a documented process to validate the migration steps as you go?
- Do you understand the concept of Copy Once volumes?
- Are you aware that there will be Go/No Go or Recover decision points and the implications of this critical stage in the migration process?

Defining the target solution

The target solution is the final configuration and the mode of operation for the new storage environment. Ideally, for a single storage system, the new storage system will simply replace the old one. However, during the data migration both storage systems will exist in parallel. This means that a physical location and the associated infrastructure demands, such as electrical power, Fibre Channel ports, and so on, must be temporarily provided for the new system.

Sometimes the customer might combine extra requirements with the data migration. A data migration can be a good opportunity to introduce additional changes in the data center environment.

A data migration can also be necessary because of legal requirements or other regulations, not necessarily mandated by the customer. Another example might be that an existing secondary data center site that acts as a disaster recovery site is found to be too close to the primary production site. In this latter case, the data migration becomes part of a larger project where a new data center site must also be set up.

1.5.2 Designing and developing migration scenarios

In this phase, the possible ways to transfer data to the new storage system must be worked out. A set of general methods is described in Part 1, “Copy Services based Migration Scenarios” on page 17 and Part 2, “z/OS based Migration Scenarios” on page 209.

It might also be helpful to develop and provide extra tools such as scripts or programs to help the operators who perform the migration tasks. Such scripts, for example, can help to ensure that all required volumes are specified with each migration step and enforce the correct sequence of steps when performing the migration.

Designing the migration path

The migration path describes how data will be transmitted to the new storage system. The design of the migration path depends on the current production environment and the target setup.

At the beginning of this phase, several options for the tools and techniques that can be used for the migration are usually considered. Figure 1-2 shows an example. You need to evaluate the available tools based on the type of migration and the outcome of the assessment process. To do this evaluation, build a test environment that duplicates as nearly as possible the actual migration environment.

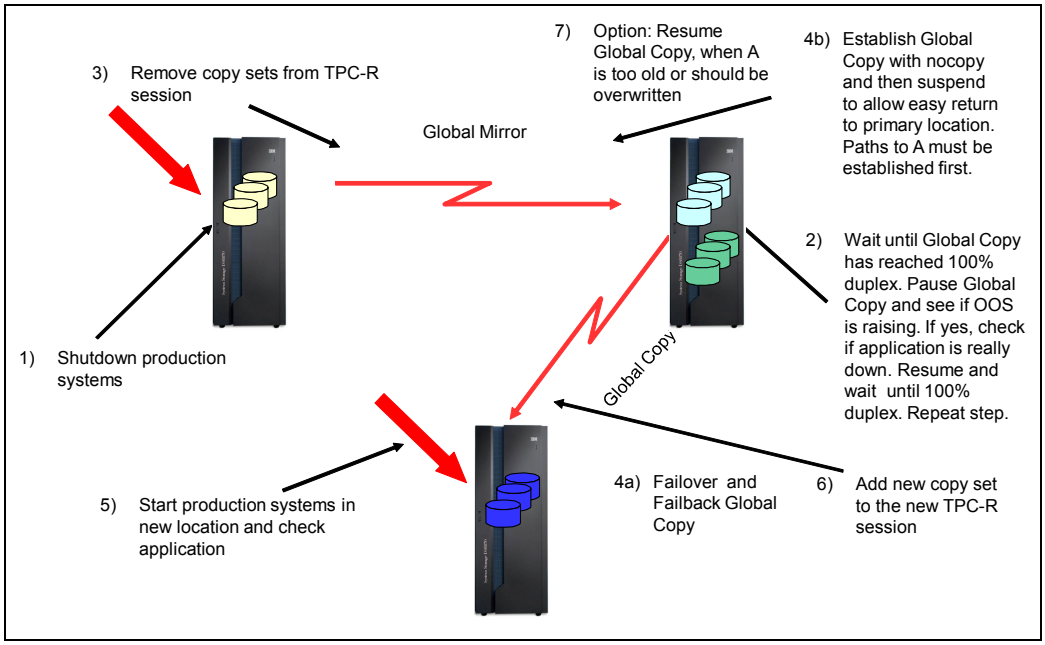

Figure 1-2 Example of a migration path to a new primary storage

If you cannot easily or economically construct such an environment, IBM has a service offering that allows you to use IBM lab equipment.

Developing Tools

Usually, more than one set of tools and techniques can be used in a data migration. You need to evaluate the available tools based on your type of migration and the outcome of the assessment process. To test these *migration scenarios*, build a test environment that duplicates as near as possible the environment on which the migration will take place. If you cannot construct such an environment, IBM has a service offering that allows you to use IBM lab equipment.

When the migration scenarios are developed in the test environment, it is helpful to plan a back-out capability into the migration path. This capability allows you to effectively return to the original production environment when errors or other problems occur during the migration. These back-out scenarios might only apply in situations where a rollback of the migration is possible. Possible exit points for a back-out must also be designed in this phase.

You might also want to develop automation scripts to be used during the migration. The following examples demonstrate situations where automation scripts are useful:

- When the time frame of the migration is weeks or months, changes to the production environment, like storage expansion, might take place during the time of the migration. These additional volumes or storage systems must be taken into account.
- The migration scenarios can get complex, so the operators will need guidance throughout the migration steps. This guidance can be provided by writing scripts to perform a set of operations automatically. When manual interactions are required, instructions can be given to the operator, which then must be acknowledged.

- When reliable status verifications are difficult to obtain, especially in situations where many volumes must be migrated. A script automating the status check can reduce errors and simplify the process, particularly when the volume addresses are not in a contiguous order.

Example 1-1 illustrates how a migration script can assist the operator with checkpoints during the migration process. For more information, see “Defining checkpoints in a migration path”. The script **MigrationCutover.pl** covers steps 4 - 6 shown in Figure 1-2 on page 8.

Example 1-1 Execution of a migration script

```
x345-lab-32b:/home/DSCLInator/DSCLInator # MigationCutover.pl -n APP_1

TPC-R username: tpcadmin
password:
Authenticating to TPC-R server @ 9.155.59.13...
   ...authentication was successful

Checking initial copy state of Global Copy between B and C...
   ...Initial copy phase has passed.

Now checking copy status of B to C relation...
   ...Copy state of B to C relation is ok!
Now checking C to D relation...
   ...Copy state of C to D relation is ok!

Are you ready to pause Global Copy B to C [Y to continue; N to quit][Y]: y
   ...Relationship successfully paused.

Now waiting for 120s before checking the OOS...
... now checking...
Total number of OOS: 0
OS: 0

Are you ready to turn around to the Global Copy to the new prodution
direciton?[N]: y
... Validating paths of direction for new production...

Now removing Global Copy B to C ...

Now turning around Global Copy C to D ...

Adding the new copy sets to the GM session and removing the current from the GM
session will now be started...

...Are you ready to add new copy sets to the GM session HRG?[Y]: y

Now adding the following copy sets:
DS8000:2107.TV181:VOL:5400,DS8000:2107.ABTV1:VOL:4200,DS8000:2107.ABTV1:VOL:4400,D
S8000:2107.ABTV1:VOL:4000
DS8000:2107.TV181:VOL:5401,DS8000:2107.ABTV1:VOL:4201,DS8000:2107.ABTV1:VOL:4401,D
S8000:2107.ABTV1:VOL:4001
DS8000:2107.TV181:VOL:5402,DS8000:2107.ABTV1:VOL:4202,DS8000:2107.ABTV1:VOL:4402,D
S8000:2107.ABTV1:VOL:4002
DS8000:2107.TV181:VOL:5403,DS8000:2107.ABTV1:VOL:4203,DS8000:2107.ABTV1:VOL:4403,D
S8000:2107.ABTV1:VOL:4003
```

```
DS8000:2107.TV181:VOL:5404,DS8000:2107.ABTV1:VOL:4204,DS8000:2107.ABTV1:VOL:4404,D
S8000:2107.ABTV1:VOL:4004
DS8000:2107.TV181:VOL:5500,DS8000:2107.ABTV1:VOL:4300,DS8000:2107.ABTV1:VOL:4500,D
S8000:2107.ABTV1:VOL:4100
DS8000:2107.TV181:VOL:5501,DS8000:2107.ABTV1:VOL:4301,DS8000:2107.ABTV1:VOL:4501,D
S8000:2107.ABTV1:VOL:4101
DS8000:2107.TV181:VOL:5502,DS8000:2107.ABTV1:VOL:4302,DS8000:2107.ABTV1:VOL:4502,D
S8000:2107.ABTV1:VOL:4102
DS8000:2107.TV181:VOL:5503,DS8000:2107.ABTV1:VOL:4303,DS8000:2107.ABTV1:VOL:4503,D
S8000:2107.ABTV1:VOL:4103
DS8000:2107.TV181:VOL:5504,DS8000:2107.ABTV1:VOL:4304,DS8000:2107.ABTV1:VOL:4504,D
S8000:2107.ABTV1:VOL:4104
Are the above CPSETs correct? Enter Y to continue[N]: y
 IWNR1000I [Apr 19, 2013 5:44:29 AM] Copy sets were created for the session named
HRG.
```

Defining checkpoints in a migration path

The whole migration scenario consists of a series of migration steps. These steps can be grouped with checkpoints after each group of steps. The status at a specific checkpoint will decide whether to go forward with the next step or a back-out.

Figure 1-3 on page 11 shows a 2-site production with Global Mirror or Metro Mirror where the primary storage system must be replaced.

The migration steps might be as follows:

1. Create a cascaded Global Copy relation from the auxiliary storage system to the new primary storage system.
2. Wait for the initial copy phase of the cascaded Global Copy to be complete.
3. Quiesce the application at the primary site.
4. Fail over the cascaded Global Copy to the new storage system.
5. Start the application at the new storage and check the application for consistency and completeness.
6. Remove the copy relation between the old primary and auxiliary storage system.
7. Fail back the Global Copy from the new primary to the auxiliary storage system.

In this migration scenario, steps 1 to 2 will result in a checkpoint in which the data has been copied as a whole to the new storage system, after the shutdown of the application, in step 5. The application is restarted with the new storage. At this point, another checkpoint has been reached where you can decide to continue at the new storage if no errors are found. Otherwise, the application must be backed out to the old storage.

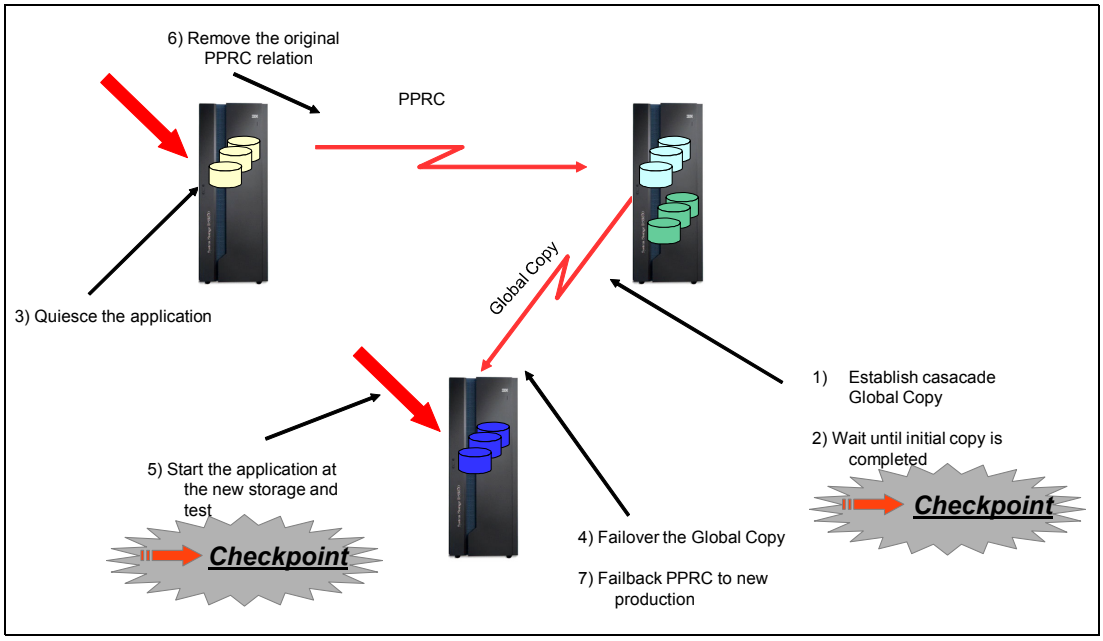

Figure 1-3 Defining Checkpoints in a migration scenario

1.5.3 Testing

Testing the migration scenarios provides confidence that the migration steps will work as expected. For this reason, a dedicated test environment is helpful. If it is not possible to provide a dedicated environment for testing at the customer site, use a subset of an advisable extent to create one, or engage IBM to assist you. The testing environment should include, at the minimum, the necessary resources to validate the migration and back-out processes to ensure that you are able to recover the environment regardless of the success or failure of the migration process.

Testing the migration scenarios

After the appropriate migration scenario has been developed, you must test and document the process extensively until you are confident that the scenario will work during the migration of live data.

As mentioned earlier, testing is a critical step in the migration process. Budget sufficient time and attention for this important phase of the project. The more completely you test and document the scenario, and provide automation where it is required, the smoother the migration will go.

Testing the migration scenarios is crucial for essentially two major reasons. At first, with testing you ensure that the scenarios are valid and are working as expected. The tests show whether all the migrations steps work as expected or reveal missing steps that were not considered in the design phase.

However, the operator who will be responsible for performing the migration steps must be familiar with every step of the overall migration scenario. With a working test environment, the operator has an opportunity to gain enough practice with the migration tools and developed scripts before performing the actual migration.

Important: Thorough development and testing of the migration process can reduce the potential impact on production and increase the success rate of the migration project.

In summary, to maximize the potential for success in a migration project, perform these steps:

- Use a dedicated test environment to develop the migration process.
- Thoroughly test and validate the chosen migration process.
- Thoroughly document the chosen migration process.
- Develop effective automation where appropriate to simplify tasks and minimize human error.
- Ensure that there is a valid, tested, and documented process to back out and reset the environment to a pre-migration state.

1.5.4 Supporting tools and products

Several tools and products are available that can be helpful during data migrations.

In Part 1, “Copy Services based Migration Scenarios” on page 17, Tivoli Storage Productivity Center for Replication and the IBM Geographically Dispersed Parallel Sysplex (GDPS) are covered. Typically, Tivoli Storage Productivity Center for Replication is used in Open System environments, and is also capable to handle the count key data (CKD) volumes that are used in z/OS environments. GDPS is a z/OS based product and service offering for z/OS customers.

Part 2, “z/OS based Migration Scenarios” on page 209 describes how to use host-based migrations on the z/OS platform. It shows how data migrations can be done with TDFM and zDMF.

GDPS and Tivoli Storage Productivity Center for Replication are management tools that offer functions for data replication with the major purpose for disaster prevention. In data migration using copy services, functions and operations are required that are beyond the set used for disaster recovery scenarios. The missing functions for a data migration must be provided using the DSCLI or TSO commands. In this situation, it is important that the new copy relations be in a state after the migration so that Tivoli Storage Productivity Center for Replication or GDPS can bring them back into their control.

When using DSCLI or TSO commands, it might be helpful to develop scripts using shell script language, Perl, Python, or REXX. The main purpose of scripting is to provide automation and consistent sequencing of migration steps. For DSCLI command scripting, a framework called DSCLIbroker is available that offers assistance with all copy services functions of the DS8000 storage system. For more information about the DSCLIbroker, see the *IBM DSCLIbroker Guide*, REDP-5155.

1.5.5 Migration plan

As part of the planning and preparation, create a high-level migration plan (see Figure 1-4 on page 14) and communicate that plan to all stakeholders.

Employ checkpoints to ensure that the migration is proceeding correctly. It can be overwhelming, for example, to consider migrating several servers in one large move. Instead, break down the migration into smaller activities to make it more manageable. Among other benefits, checkpoints allow for potential back-outs.

The plan also allows you to track schedule commitments while completing the migration. Always allocate extra time for tasks, generally 15 - 20% more time than would be required for the best case migration scenario. Doing so gives you time to resolve unexpected issues. The better you plan, the fewer unexpected or unforeseen issues will be encountered with both resource (team member) commitments and technology.

Consider the following topics, among others, to ensure the accuracy of the project plan:

- Change management

  A common institution in data centers is Change Management, which ensures that all involved departments and responsibilities are at least informed or have the opportunity to agree to or reject the changes. Many of the migration steps are certainly subject to change management because they can affect the production systems.

- Infrastructure readiness

  All components of the infrastructure must be ready to take over the production after the migration.

  - The storage connectivity (Fibre Channel/IBM FICON®) must be established before the migration scenario can begin. This includes new host connections and PPRC connections for the replication itself. When the migration is finished, the project plan should also include an action point to clean up unused connections.
  - IP network connections to all new components and devices are required.
  - Infrastructure components like air conditioning and electrical power must be validated for the new equipment
  - The logistics efforts for hardware ordering, delivery, and installation must have a dedicated timeline in the project plan.
  - You might need to reassign host and server resources like CPU, memory, and adapters, especially with logical partitioning (LPAR) and other virtualization configurations.
  - Additional temporary capacity might be needed. For configurations with disaster protection, additional capacities might be required, for example for the Disaster Recovery (DR) practice environment.
  - Licenses for storage and other software components must be reviewed and updated if required.
  - More considerations might apply.

- Job controls, scheduling

  Larger data centers typically use a job controller to automate regular jobs like batch runs, data backups or so on. Make sure that the planned migration scenarios will not collide with jobs requiring certain resources that might be compromised by the migration tasks.

- Back-out paths

  The purpose of the back-out path is to return immediately to the original situation should anything go wrong during the migration. To identify that the back-out is indicated, a verification at the application level is required.

- Post migration task

  When the migration has succeeded, further tasks are required that essentially include cleaning up path and PPRC definitions, zone definitions, and physical connections that are no longer required. When the purpose of the migration was the removal of an older storage system, the disposal of that storage system can also be part of the plan.

Figure 1-4 shows an example migration project plan.

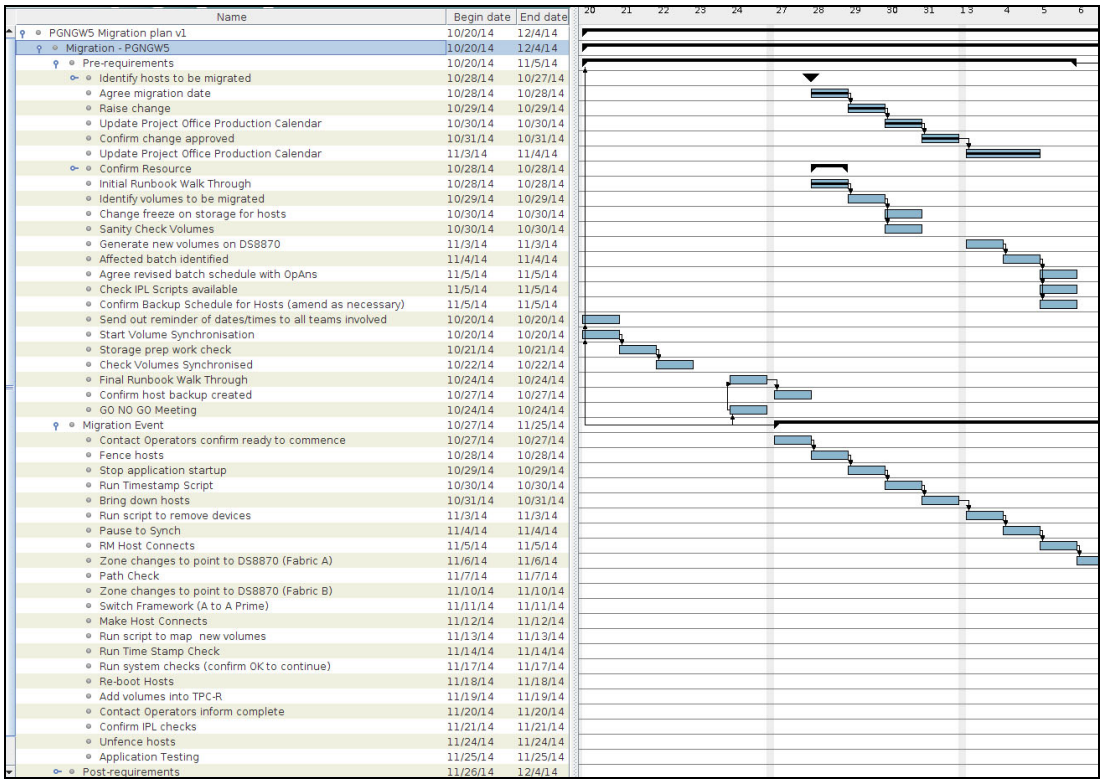

Name	Begin date	End date
PGNGW5 Migration plan v1	10/20/14	12/4/14
Migration - PGNGW5	10/20/14	12/4/14
Pre-requirements	10/20/14	11/5/14
Identify hosts to be migrated	10/28/14	10/27/14
Agree migration date	10/28/14	10/28/14
Raise change	10/29/14	10/29/14
Update Project Office Production Calendar	10/30/14	10/30/14
Confirm change approved	10/31/14	10/31/14
Update Project Office Production Calendar	11/3/14	11/4/14
Confirm Resource	10/28/14	10/28/14
Initial Runbook Walk Through	10/28/14	10/28/14
Identify volumes to be migrated	10/29/14	10/29/14
Change freeze on storage for hosts	10/30/14	10/30/14
Sanity Check Volumes	10/30/14	10/30/14
Generate new volumes on DS8870	11/3/14	11/3/14
Affected batch identified	11/4/14	11/4/14
Agree revised batch schedule with OpAns	11/5/14	11/5/14
Check IPL Scripts available	11/5/14	11/5/14
Confirm Backup Schedule for Hosts (amend as necessary)	11/5/14	11/5/14
Send out reminder of dates/times to all teams involved	10/20/14	10/20/14
Start Volume Synchronisation	10/20/14	10/20/14
Storage prep work check	10/21/14	10/21/14
Check Volumes Synchronised	10/22/14	10/22/14
Final Runbook Walk Through	10/24/14	10/24/14
Confirm host backup created	10/27/14	10/27/14
GO NO GO Meeting	10/24/14	10/24/14
Migration Event	10/27/14	11/25/14
Contact Operators confirm ready to commence	10/27/14	10/27/14
Fence hosts	10/28/14	10/28/14
Stop application startup	10/29/14	10/29/14
Run Timestamp Script	10/30/14	10/30/14
Bring down hosts	10/31/14	10/31/14
Run script to remove devices	11/3/14	11/3/14
Pause to Synch	11/4/14	11/4/14
RM Host Connects	11/5/14	11/5/14
Zone changes to point to DS8870 (Fabric A)	11/6/14	11/6/14
Path Check	11/7/14	11/7/14
Zone changes to point to DS8870 (Fabric B)	11/10/14	11/10/14
Switch Framework (A to A Prime)	11/11/14	11/11/14
Make Host Connects	11/12/14	11/12/14
Run script to map new volumes	11/13/14	11/13/14
Run Time Stamp Check	11/14/14	11/14/14
Run system checks (confirm OK to continue)	11/17/14	11/17/14
Re-boot Hosts	11/18/14	11/18/14
Add volumes into TPC-R	11/19/14	11/19/14
Contact Operators inform complete	11/20/14	11/20/14
Confirm IPL checks	11/21/14	11/21/14
Unfence hosts	11/24/14	11/24/14
Application Testing	11/25/14	11/25/14
Post-requirements	11/26/14	12/4/14

Figure 1-4 Example of a migration project plan

As shown in Figure 1-4, a well developed project plan can consist of many tasks. It also indicates how many resources and participants can be required for the migration project.

1.6 Migration phase

The migration itself starts when the data transmission to the new storage is started. The starting point is a clear checkpoint in the project plan. The whole migration phase is a two stage process. The first stage is the *data transmission stage*. Depending on the available bandwidth, this stage is typically the most time consuming task because you must wait until the new storage is synchronized. Synchronized means that when using host-based migration, both mirrors are in sync, and when using Copy Services that the initial copy of the copy relations has completed. For more information, see 2.3.2, “Monitor Global Copy progress” on page 23 for details.

When the data is synchronized, the new storage is ready for the second stage of the migration phase, which is the *take over stage*. At this stage, you must provide data consistency at the new storage system because typically the production continues to run during the transmission stage. There are several ways to provide consistency:

- When using host-based mirroring (for example when using a logical volume manager in AIX, Linux, or other UNIX-based operating system), the transmission is done by synchronizing the new physical volumes to the existing physical volumes. When synchronized, the logical volume manager (LVM) provides functions to remove the mirror with the old physical volumes and the migration is considered finished.

- WWhen using DS8000 Copy Services , the new storage volumes are migrated using a new copy relation. The new copy relation can be direct when no other copy replication is in place, or be a cascaded relation from an existing copy relation. To make the data consistent to the new volumes, the application must be shut down. In Open System environments, you must then unmount the volumes so that all host cached data is written to the volumes as well. When all copy relations show no more out-of-sync tracks, the data is consistent.
- When using DS8000 Copy Services with IBM HyperSwap®, the consistency is provided by the Metro Mirror. In migrations, HyperSwap can be used to provide a nearly seamless take over to a new storage system. This is only possible when HyperSwap is supported by the operating system of the attached host application server (System z with HyperSwap, or IBM Power Systems™ with Power HyperSwap)
- With TDMF or zDMF, the data consistency is directly provided by TDMF or zFMF respectively. For more information, see Chapter 10, “Using TDMF z/OS” on page 211 and Chapter 11, “z/OS Dataset Mobility Facility” on page 271

Before the production can be started with the new storage system, perform a post migration phase follow up.

1.7 Post migration phase

After migrating data to the new storage system, ensure that the target system contains a complete and accurate copy of the source. You also need to make sure that the migrated data is still usable by applications. It is important to validate that you have the same data and functions of the application after the migration. Make sure that the application runs with the new volumes, that performance is still adequate, and that operations and scripts work with the new system.

1.7.1 Validation

The validation after the data migration is the final criteria for the decision to run production on the new storage. Validation tasks can also be assigned to multiple responsibilities and departments. For example, after a successful migration using Copy Services, done by the storage management department, the following validations can take place:

- The server management department must check that all volumes are accessible. If yes, all volume groups and file systems must go online and check that all applications are starting correctly.
- A middleware department can be responsible to check that all databases instances can be started and no corruption occurs. They should also perform some database functions or run some reports as additional checks.
- Another department responsible for business to business transactions can also be involved in the validation phase.

It is important that during the validation phase the back-out scenarios can still be applied. If there are problems with the data or the infrastructure after the migration, either fix the problem or revert to the original state.

A good practice is to keep the data at the original site available and usable as long as possible. Maintaining this data allows you to discard the migrated data and restart production at the original site if something goes wrong with the migration.

When a secondary or a tertiary storage system is replaced, the host connection and volume mappings must also be verified. The easiest way to verify this is to perform a practice failover using Tivoli Storage Productivity Center for Replication or GDPS. When no practice volumes are defined, a full failover to the disaster recovery site and start of the applications should be done instead.

1.7.2 Production take over

The production take over is essentially the positive decision to stay at the new storage after validation and bringing the applications online. A positive decision can only be made when all other instances and departments have also concurred that the migration was successful.

1.7.3 Cleaning up

Obviously the storage system that has been replaced must be prepared for the disposal. Beside the dismounting of all cabling, the logical clean-up tasks must apply before the storage system can be powered off. For example, remaining copy and path relations from or to the old storage system should be removed.

You should also perform a secure data overwrite of whole storage system, which ensures that all data are deleted securely and all logical volume definitions are deleted before the storage system leaves the customer site.

1.7.4 Reports and documentation

After the migration is completed, compile migration statistics and prepare a report to highlight what worked and what did not work, and what potential risks were uncovered during the migration. Also, document the lessons learned from the problems encountered.

These reports will help you build a repeatable and consistent process by building on what worked, and fixing or changing what did not. Also, documenting the migration process can help you train your staff, and simplify or streamline your next migration, reducing both expense and risk.

Part 1

Copy Services based Migration Scenarios

This part examines typical data migration scenarios based on DS8000 Copy Services. It starts from the most common initial configurations and describe procedures to replace single or multiple instances of the data within them. The following initial configurations are covered:

- No initial copy services in place
- 2-site Metro Mirror
- 3-site Metro Mirror (MT-PPRC)
- 3-site Metro Global Mirror
- 2-site Global Mirror
- 2-site z/OS Global Mirror (zGM)

The first chapter in this section describes some general considerations that apply to all or most of the procedures. The subsequent chapters describe the procedures in detail, giving examples for the most significant and migration relevant steps. Each example uses the management interface that is most appropriate for that configuration.

A detailed description of the Copy Services features used in the scenarios is beyond the scope of this book. If you need more detailed information, see the following resources:

- *IBM DS8870 Copy Services for Open Systems*, SG24-6788
- *IBM DS8870 Copy Services for IBM System z*, SG24-6787

© Copyright IBM Corp. 2015. All rights reserved.

2

Copy Services based data migrations

This chapter describes some general considerations for Copy Services based data migration scenarios. It provides the basis for the scenarios in the following chapters, and defines common terms, conventions, and procedures.

This chapter includes the following sections:

- Terms
- Requirements
- Data replication
- Data validation testing
- Tivoli Storage Productivity Center for Replication considerations
- HyperSwap considerations
- Multiple Target PPRC considerations

© Copyright IBM Corp. 2015. All rights reserved.

2.1 Terms

This section defines some of the terms used in describing the scenarios.

2.1.1 Volume set

The term *volume set* is used to designate the data that is to be migrated. A volume set can be any of these things:

- A number of volumes within a single storage system
- A single complete storage system
- A number of volumes spread across several storage systems
- Several complete storage systems

2.1.2 Data replication and cut over

The DS8000 Copy Services based data migration scenarios generally consist of two phases:

- Data replication

  Data on the original volume set is replicated to the new volume set, with as little impact to normal production as possible.

- Cut over

  Switching the production from the original volume set to the new volume set after all data has been replicated. As prerequisites for cut over you must ensure that these goals are met:

  – All data is replicated to the new volume set.
  – The replicated data is consistent and valid.

Important: It is common practice to run a data validation test before the final cut over of the production systems, when the data migration is complete. With such a test you verify that these goals are met:

- All necessary data is included in the migrated volume set.
- The method to create consistency on the new data set is working.
- The cut over procedure is working.

2.1.3 Site

The term *site* is used to designate the physical location of a volume set. A site typically is a computer room with one or more storage systems that contain all the volumes of a data set. If a volume set is distributed over several rooms, this still counts as one site. There might be configurations where two volume sets (such as a Metro Mirror primary and secondary volume sets) are in the same computer room. In these terms, this counts as two sites.

2.1.4 Out-of-sync tracks

The amount of data that the data replication solution (Global Copy in most cases) has not yet replicated from a primary to a secondary volume is measured and displayed in *out-of-sync* (OOS) tracks. When the term out-of-sync tracks is used for a volume set, it means the sum of OOS tracks of all volumes in a volume set.

2.1.5 Primary and secondary volumes

The terms primary or secondary volumes, or *primaries* and *secondaries*, designate the volumes sets that participate in a specific Copy Services relation.

2.1.6 Conventions in diagrams

This part of the book uses a consistent notation in diagrams illustrating the procedures. As shown in Figure 2-1 the following standard symbols are used:

- Straight thick blue lines for host access
- Straight thin red lines for synchronous data replication
- “Lightning” shaped thin red lines for asynchronous data replication
- Dashed lines for connections that are configured, but not actively used

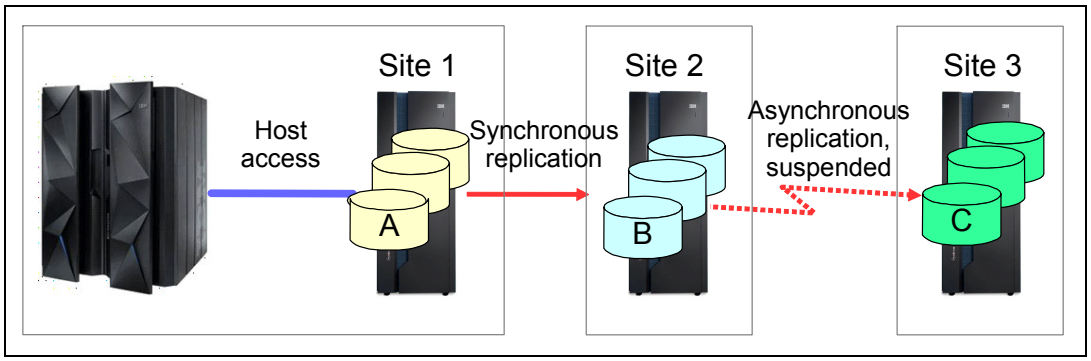

Figure 2-1 Examples for notations in diagrams

Volume sets are represented by a picture of a storage system and a group of disk symbols, specified with a volume set identifier (A, B, C) if required. The data migration target volume set (the new volume set) is marked with a prime (A’, B’, C’). A *site* is designated by a thin box and a site specifier.

2.2 Requirements

The following requirements must be in place to run a Copy Services based data migration:

- Almost all of the procedures in subsequent chapters use IBM DS8000 specific features to migrate data. Therefore, they are only applicable if both the existing and the new volume set are on a member of the DS8000 storage system family. An exception is the scenario that replaces the secondary volume set of a System z Global Mirror (zGM) configuration. Other vendors' storage systems are available that are compatible with zGM and therefore can be used to migrate from or to with this method.
- The necessary DS8000 Copy Services licenses must be available on all affected storage systems.
- Connectivity
  - To migrate data with Global Copy or Metro Mirror, you need Fibre Channel connection from the primary to the secondary devices.
  - To migrate data with System z Global Mirror, you need FICON connectivity from the primary devices to the System Data Mover and from there to the secondary devices.

2.3 Data replication

With one exception (in the zGM section), *DS8000 Global Copy* is used as the means to replicate the data to the new volume set. Global Copy is a non-synchronous replication method that can replicate data without impact to the primary storage system over unlimited distances effectively.

In many cases, a Copy Services replication configuration is already in place. To replace any of the existing volume sets, use *cascaded Global Copy* and replicate the data to the new volumes from the last instance of data in the existing configuration.

Using cascaded Global Copy has several advantages for when moving data to the new copy set:

- Unlimited distance possible
- No application downtime
- Existing disaster recovery configuration can remain active
- Replication can be stopped and restarted if necessary without full copy
- Data validation tests are possible without full copy

2.3.1 Set up Global Copy

Figure 2-2 shows a typical example where you want to replace the copy set A in site 1 of a Metro Global Mirror configuration. To replicate the data, configure Global Copy C to A' as a cascade to the existing Global Copy relation (B - C). The new volume set A' set is contained in the Global Copy secondary volumes. The Global Copy primary volume set C is not the one you want to replace in this data migration.

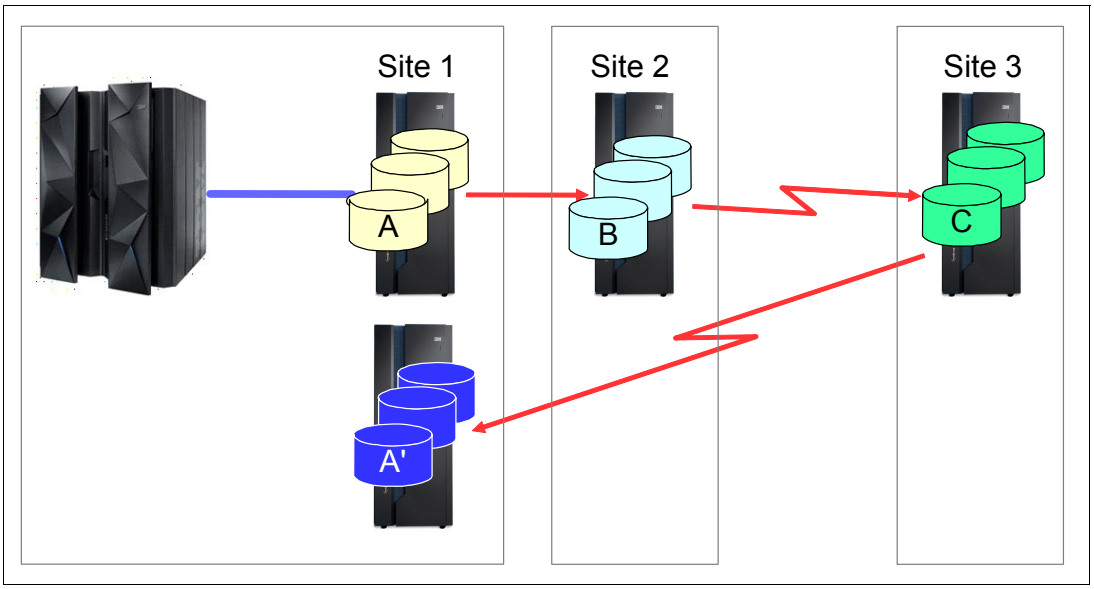

Figure 2-2 Example of a data migration setup using Global Copy

Note: This example does not consider the Global Mirror journal volumes. They are not relevant for the data replication.

Global Copy does not enforce data consistency by itself. Before the cut over, you must make sure that the migrated data in the Global Copy secondary volume set (A' in this example) is

consistent. This is specific for each data migration scenario. Therefore, refer to the description of the individual scenarios in the following chapters.

Example 2-1 shows the DSCLI command that you can use to start the Global Copy that replicates the data to the new volume set. You run the command (**mkpprc**) on the storage system that contains the Global Copy primary volume set.

Example 2-1 DSCLI command to start cascaded Global Copy for data migration

```
dscli> mkpprc -type gcp -cascade -remotedev IBM.2107-75CYK71 8000-8003:8800-8803
CMUC00153I mkpprc: Remote Mirror and Copy volume pair ... successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair ... successfully created.
...
```

The following are the important options in this command:

- **-type gcp**: This option specifies that the replication type is Global Copy. It is required because the DS8000 Copy Services do not allow a synchronous replication as the second part of a cascaded configuration.
- **-cascade**: This option allows the Global Copy primaries to be the secondaries of another remote mirror relationship, and therefore enables the cascaded copy.

Note: If the storage system that contains the Global Copy primaries is a DS8870 running micro code version 7.4 or later, the **-cascade** option is not needed.

Note: In case of very high application write workload caused by write collisions (updating the same data again and again), there can be an effect due to the attached Global Copy cascade. If you observe or suspect such an effect, consider a staged start of the Global Copy relations. Start it for a group of volumes in the volume set, wait until the initial copy is done, and then start the next group.

2.3.2 Monitor Global Copy progress

After the Global Copy is running, you monitor the progress by observing the *First Pass Status* and the amount of Out-of-Sync tracks of the Global Copy relations.

Example 2-2 shows the DSCLI command that you can use to monitor both parameters. After all relations have the `First Pass Status` status `True`, the Initial Copy is complete and your data migration is merely transferring updates made to the original volume sets. The `Out of Sync Tracks` count indicates how much data still must be replicated to the new volumes set. While production I/O is running, its total number will not reach zero permanently.

Example 2-2 DSCLI command to check the initial copy status

```
dscli> lspprc -l 8000-8003
ID        State              Reason Type         Out Of Sync Tracks ... First Pass Status
==================================================================...==================
8000:8800 Copy Pending       -      Global Copy  0                  ... True
8001:8801 Copy Pending       -      Global Copy  0                  ... True
8002:8802 Copy Pending       -      Global Copy  412                ... False
8003:8803 Copy Pending       -      Global Copy  0                  ... True
```

It is a characteristic feature of Global Copy that the relations will always stay in `Copy Pending` state. By definition, a Global Copy is never finished. After the total amount of OOS tracks is

close to zero and doesn't change much over time, you can consider the data replication nearly complete and start the cut over phase.

2.3.3 Prepare for cut over

To prepare for cut over, make sure that these statements are true:

- All data is replicated to the new volume set.
- The replicated data is consistent and valid.

Global Copy does not maintain data consistency by itself. To make the data on the Global Copy secondaries complete and consistent, you must make sure that the primaries and the secondaries are identical, and that the data on the primaries is consistent.

There are different ways to achieve consistency of the new data set, depending on the initial configuration and the volume set you are replacing. See the individual scenarios in the following chapters for the appropriate methods.

In this example, consistency is achieved by stopping all production I/O and waiting until the total OOS tracks of the B to C and C to A' relations are zero.

You can now start the replacement of the original volume set with the new one. This process is also unique to each scenario.

2.4 Data validation testing

It is common practice to run a data validation test before the final cut over of the production systems, when the data migration is complete. With such a test you verify that these statements are true:

- All necessary data is included in the migrated volume set.
- Your method to create consistency on the new data set is working.
- The cut over procedure is working.

This section describes a procedure that is applicable to any migration scenario that is using a cascaded Global Copy to perform the data migration. The example in Figure 2-2 on page 22 is used as the starting scenario.

You can perform one or several rounds of testing before you decide that it is time for the cut over. With Global Copy for data transport, you can make use of change recording. After a test, you only have to transmit changes and not run the full replication again.

2.4.1 Host considerations

The data validation testing requires a host system to be available and connected to the new volume set. The setup and preparation of this host is different depending on its normal role:

- If it is normally running active production workload, you must halt operations on that system during the tests
- If it is dedicated for the testing process, you must stop production I/O for only a short time, to create a consistent copy of the replicated data. Then, you can continue normal operations while running the tests.

After the testing has been completed, all data that are written to the new volume set during the test are backed out.

2.4.2 Prepare the new volume set for testing

To prepare the new volume set for testing, complete the following steps:

1. Make the data contained in the new volume set (A' in this example) valid and consistent. To achieve system and application level consistency, you must briefly stop all production I/O. Wait until *all copies* in the Copy Services configuration, including the replication to the new volume set, are identical. Confirm this by checking that the OOS tracks of all relations are zero.

   Note: The exact method to stop production I/O and make your data consistent depends on your application. In some cases, it might be enough to switch applications to their backup mode, whereas in other cases you might have to shut down the application systems completely.

   If a crash consistent copy of the data is good enough, or if you cannot stop production I/O to prepare for testing, you can also perform a freeze and suspend for the Metro Mirror A to B. Then, wait until the OOS tracks of all cascaded Global Copy relations are zero. Suspend the Global Copy to your new volume set and resume Metro Mirror A - B. Testing with such a copy is like running a disaster recovery test, and might require operating system and application recovery.

2. Pause the Global Copy relationships: After the data in the new volume set is valid, pause the Global Copy relations that are used to replicate the data. This action enables change recording on the Global Copy primaries. You can now continue with your normal operations on the old volume sets.
3. Perform a failover operation on the Global Copy Secondaries to prepare the new volume set for host access. It also enables change recording on the Global Copy secondaries. Thus, all updates to secondaries during this test will be reversed after it is finished and you restart the Global Copy relations.

Example 2-3 shows a DSCLI command that you can use to pause the Global Copy relationship to the new volume set. Issue the command on the storage systems that contain the Global Copy primaries (volume set C in this example).

Example 2-3 DSCLI command to pause Global Copy relations

```
pausepprc -remotedev IBM.2107-75ZYK71 8000-8003:8800-8803
CMUC00157I pausepprc: Remote Mirror and Copy volume pair ... successfully paused.
CMUC00157I pausepprc: Remote Mirror and Copy volume pair ... successfully paused.
...
```

Example 2-4 shows the DSCLI command that you use to perform the failover operation. Issue this command on the storage systems that contain the new volume set (in this example A').

Example 2-4 DSCLI command to fail over Global Copy relations

```
failoverpprc -type gcp -remotedev ibm.2107-75ZA571 8800-8803:8000-8003
CMUC00196I failoverpprc: Remote Mirror and Copy pair 8800:8000 successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 8801:8001 successfully reversed.
...
```

After the pause and failover operations are complete, validate the status of the volumes on the Global Copy primary and secondary data sets, using the **lspprc-l** command. Each of the volumes must be in a *Suspended Host Source* state. This is a smoke test to ensure that the volumes are in the correct state before proceeding.

2.4.3 Data validation testing

Figure 2-3 illustrates a test configuration example. The dashed line between the Global Copy primary and secondary volume sets C and A' indicates that this connection is physically and logically available. However, during the test cycle, no data is transmitted on it. This example uses a dedicated host system for testing. Therefore, you can continue normal operations, as soon as you are done preparing the new volume set for testing.

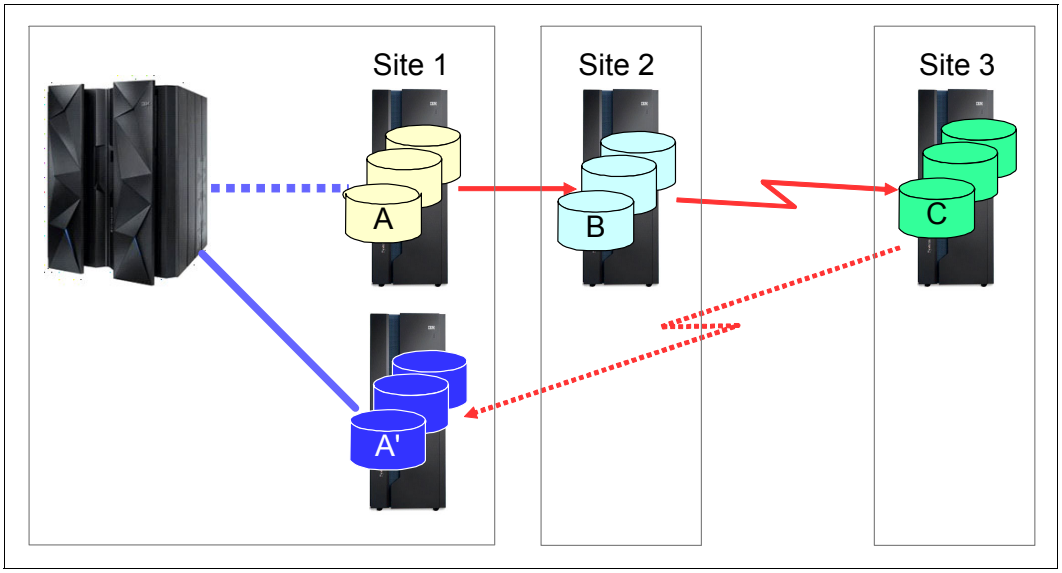

Figure 2-3 Configuration for validation testing before completing the migration

After the suspend and failover operations have completed, you can begin the data validation testing on the new volume set. This process is not described in detail here because every client configuration is unique and the process must be tailored to your specific environment. Whatever process you use, it is imperative that it be well tested and documented as discussed in 1.5.3, "Testing" on page 11.

Important: Do not allow any real production updates to the new volumes set (A') during the validation tests. All updates made to this volume set in the test will be overwritten when Global Copy is restarted.

2.4.4 Data validation testing complete

When the validation tests are complete, return to the normal data migration configuration, as it was before the tests:

1. Verify that no volumes of the new volume set are online to a host or have a reserve set on them. If any volumes are in this state, you must resolve the situation before moving forward. The **`failbackpprc`** command provides the **`-tgtonline`** and **`-resetreserve`** options to reset or override these conditions. Use them only if you fully understand the reason for the condition to make sure that the fail back operation does not overwrite valid data.
2. Perform a fail back operation at the Global Copy primary volume set. This action restarts the cascaded Global Copy relationships and override all changes made to the new volume set during the data tests. Updates made by active production hosts during the validation testing are accumulated by the change recording and are also propagated to the new volume set.

Example 2-5 shows the DSCLI command that you can use to fail back the Global Copy relationship from the Global Copy primary to the secondary volume set. You issue the command on the storage systems that hold the Global Copy primaries. It restarts the Global Copy in the original direction.

Example 2-5 DSCLI Commands to fail back the data migration Global Copy relationships

```
failbackpprc -type gcp -remotedev IBM.2107-75ZYK71 8000-8003:8800-8803
CMUC00197I failbackpprc: Remote Mirror and Copy pair 8000:8800 successfully failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 8001:8801 successfully failed back.
...
```

After the cascaded Global Copy relationships have been restarted, monitor the Out-of-Sync tracks as described in 2.3, “Data replication” on page 22. You are ready for the next test or the final data migration steps as soon as they are close to zero. For more information about these final steps, see the descriptions of the individual scenarios in the next chapters.

Note: Verify the Out-of-Sync track status for all replication relationships in the migration environment before continuing. Doing so helps ensure that no unexpected activity has occurred in the environment during the testing phase.

2.5 Tivoli Storage Productivity Center for Replication considerations

Because disaster recovery procedures are critical and complex, many clients have management and automation solutions, such as *Tivoli Storage Productivity Center for Replication*, installed to manage their existing Copy Services configuration. These solutions are designed to manage Copy Services during normal operations, test, and during failure or even a disaster. They are usually not designed to help with data migrations. However, you must consider an existing DR automation solution when you plan a data migration.

This section explains these considerations using the example of Tivoli Storage Productivity Center for Replication:

- Can Tivoli Storage Productivity Center for Replication help set up the required replication?
- How do I replace the old with the new volume set in the Tivoli Storage Productivity Center for Replication configuration after data migration is complete?

For more information about how to perform tasks described in this chapter, see the Tivoli Storage Productivity Center for Replication documentation:

- IBM Knowledge Center: Tivoli Storage Productivity Center for Replication for System z

  http://www.ibm.com/support/knowledgecenter/SSU9FR/welcome
- IBM Knowledge Center: Tivoli Storage Productivity Center in the section *Managing replication*

  http://www.ibm.com/support/knowledgecenter/SSNE44/welcome
- These IBM Redbooks publications:
  - *IBM Tivoli Storage Productivity Center for Replication for System z*, SG24-7563
  - *Tivoli Storage Productivity Center for Replication for Open Systems*, SG24-8149

2.5.1 Can Tivoli Storage Productivity Center for Replication help setting up the required replication?

In some cases, you can use Tivoli Storage Productivity Center for Replication to set up the data replication that is needed for a migration scenario. However, Tivoli Storage Productivity Center for Replication does not support all the steps that are required for data validation, cut over, and back out processing. There always is some low-level intervention that is required, where you must use other interfaces, like the DSCLI, or z/OS TSO or ICKDSF commands.

Regardless of whether you have Tivoli Storage Productivity Center for Replication installed or not, you need a sound knowledge of DS8000 Copy Services and the interfaces you plan to use for the critical steps.

Leave the Tivoli Storage Productivity Center for Replication configuration that supports the management of the existing environment intact and unchanged until cut over. Make sure that the existing documented DR procedures remain valid until this point in the case you must back out.

Note: You can use the existing Tivoli Storage Productivity Center for Replication configuration (for example an exported list of Copy Sets) to specify and check the volume sets for your data migration.

Note: Usually, the copy sets in a single Tivoli Storage Productivity Center for Replication session depend on each other. Therefore, migrate them together as one volume set. If you want to migrate only a subset of an existing session, pull this subset into a separate session and thoroughly test whether it is really independent from the rest before starting the migration.

2.5.2 Bring new volume set into the existing Tivoli Storage Productivity Center for Replication configuration

After successful data migration, you must replace the old volume set in your Tivoli Storage Productivity Center for Replication configuration with the one you migrated to. You can do this by completing these steps:

1. Export the Copy Sets of your running Tivoli Storage Productivity Center for Replication session to a *comma-separated values* (CSV) file. Load this file into an editor or a spreadsheet and replace the references to the old volume set with the definitions of the new volume set. Save the new Copy Set definition to a new file.
2. Immediately before cut over, log on to Tivoli Storage Productivity Center for Replication and remove all copy sets from the session that contain the volume set that is migrated. Select the option to leave the Copy Services relations intact. When all copy sets are removed, Tivoli Storage Productivity Center for Replication automatically ends the session. You can no longer manage your DR solution with Tivoli Storage Productivity Center for Replication.
3. Perform the cut over as described in the individual procedures in the following chapters.
4. Import the new Copy Set definition file into your Tivoli Storage Productivity Center for Replication session.
5. Start the session. Tivoli Storage Productivity Center for Replication queries the defined Copy Services relations and after it detects that they are already running, it immediately

sets the session to *Prepared* state. From now on, you can manage your DR solution with Tivoli Storage Productivity Center for Replication again.

6. Review and test all your DR procedures with the new volume set.

2.6 HyperSwap considerations

HyperSwap, or an equivalent solution that allows you to move production I/O operations from one volume set to another transparently, enables you to avoid the downtime for the cut over.

Figure 2-4 illustrates the setup that has been used throughout this chapter. In the initial setup, production I/O operations run against volume set A, with Metro Mirror A to B and Global Mirror B to C. The setup is HyperSwap enabled and can transparently swap production I/O from the A to the B volume set. If you want to replace volume set A, you now have two choices:

- Use one of the “classical” procedures to replace the Metro Mirror primary volume set A as described in the next chapters. They generally require a downtime of the production I/O for the cut over.
- HyperSwap production I/O to volume set B and use the *Incremental Resync* capabilities of the DS8000 storage systems to switch the replication from an A - B - C to a B - A - C configuration. Now you can replace A using the procedure to replace the Metro Mirror secondary volume set without stopping production I/O.

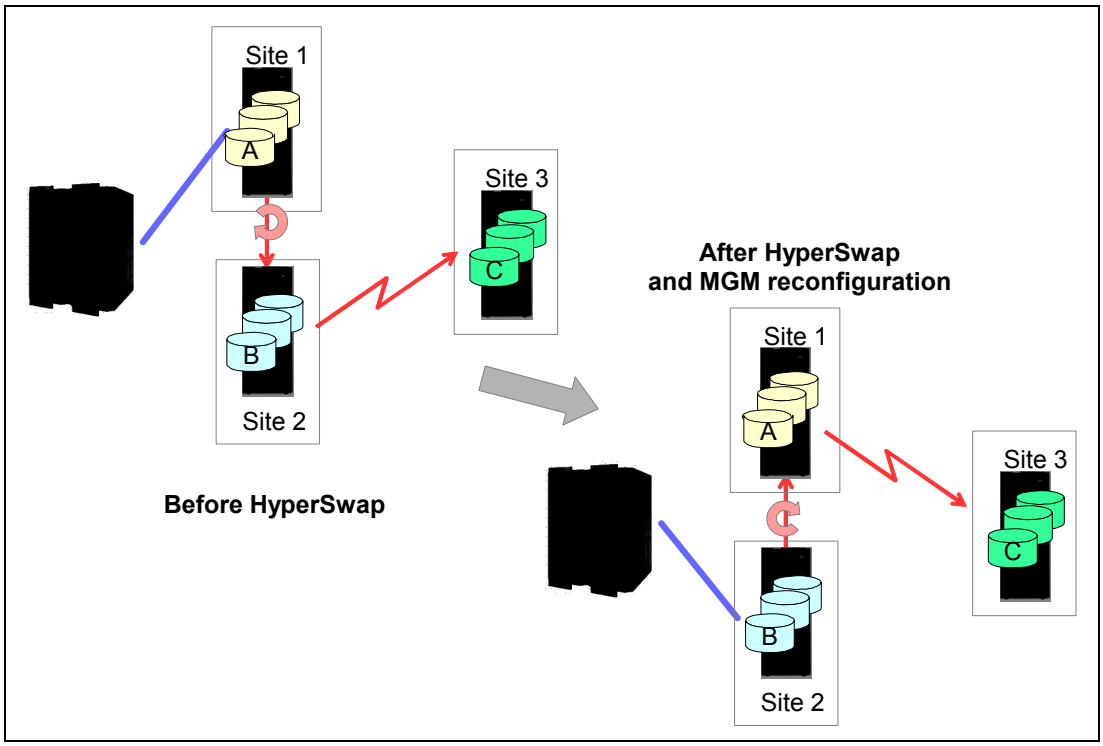

Figure 2-4 Use HyperSwap to avoid migrating active production data

2.7 Multiple Target PPRC considerations

In an initial configuration where all storage systems are DS8870s or later, running microcode release 7.4 or later, you can also use the *Multiple Target Peer-to-Peer Remote Copy* (MT-PPRC) feature for data migration. Instead of using the usual cascaded Global Copy, you can replicate the data to the new volume with Metro Mirror or Global Copy in a second relation from the same primary volumes.

Figure 2-5 illustrates this situation for the example used in the previous sections. Instead of using a cascaded Global Copy from the C to the A' volume set, set up a second Metro Mirror relation from A to A'. After this relation is synchronized, cut over from A to A' with the same method as in the "traditional" scenarios.

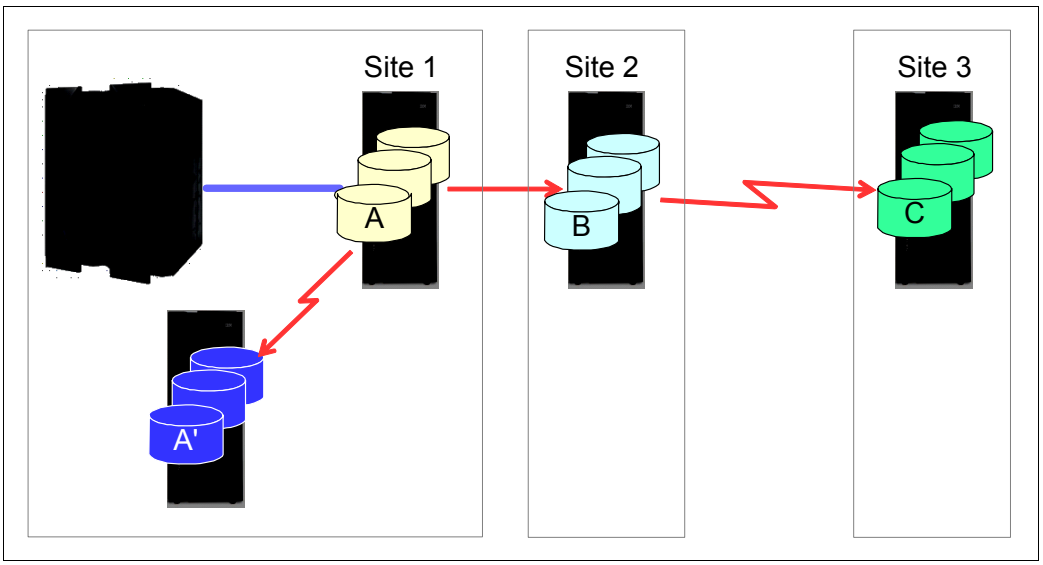

Figure 2-5 Example of a data migration using Multiple Target PPRC

Using MT PPRC has some significant advantages:

- The data replication is local. You can use existing data center infrastructure instead of long-distance connections:
  - Higher bandwidth results in shorter replication time
  - Avoid cost for additional long-distance connection
- With Metro Mirror, you have a synchronous copy of your original volume set and do not have to care about OOS tracks.

Note: Currently, MT PPRC limits the number of secondaries for each primary to two. It is not possible to add a third replication to an existing MT PPRC configuration.

3

Data Migration with no pre-existing copy services replication

The migration process described in this chapter can be used when the following two conditions are met in your environment:

- The storage unit that is the source of the migration (old storage unit) is any model of the IBM DS8000 series.
- The storage unit that is the source of the migration is not in any Metro Mirror, Global Copy, or Global Mirror relationship.

This chapter describes the following data migration phases:

- Creating target volumes on a new DS8870 (the migration target)
- Defining Global Copy / Metro Mirror relations between old and new storage units
- Establish Global Copy or Metro Mirror sessions
- Wait until all the data is synchronized and validate results
- When validation is completed, cut-over to new DS8870

This chapter includes the following sections:

- Initial configuration
- Migration process

© Copyright IBM Corp. 2015. All rights reserved.

3.1 Initial configuration

It is assumed that the DS8000 systems initially have no copy-services-based replication in place (see Figure 3-1). The new DS8870 is connected to the existing DS8000 over a storage area network (SAN) and physically installed in the same computer room or in a remote location.

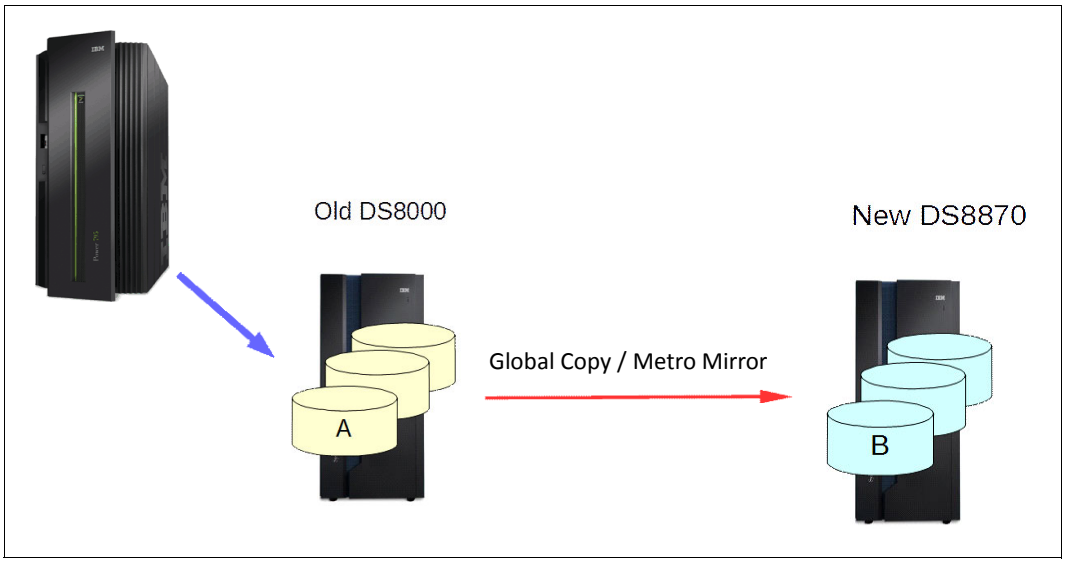

Figure 3-1 Initial configuration for 1-site environment and new DS8870 installed and connected to existing DS8000

Creating paths for Metro Mirror or Global Copy for data migration requires physical Fibre Channel (FC) connectivity between the existing storage and the new DS8870. These FC paths can be through a switch (including hardware for extended distances), or through direct physical connections between the storage units.

3.2 Migration process

Appropriate license codes to use the Remote Mirror and Copy functions must be purchased and applied to each of the systems in the migration scenario. Old and new storage systems are connected through a redundant SAN.

3.2.1 General considerations

The Copy Services functions of the DS8000 are a set of flexible data mirroring solutions that allow replication between volumes of disk storage systems. These functions are used to implement remote data backup and disaster recovery solutions.

For more information about these topics, see the following books:

- *IBM System Storage DS8000: Copy Services in Open Environments*, SG24-6788
- *DS8000 Copy Services for IBM System z*, SG24-6787

This chapter specifically deals with Metro Mirror and Global Copy.

Exception: The Global Mirror function is not addressed because using IBM Global Mirror is not a practical approach for data migration. Global Mirror is intended to be a long-distance business continuity solution. The cost and setup complexity of IBM Global Mirror are not typically justified for a one-time data migration project.

Remote Mirror and Copy migration methods replicate data between IBM storage systems, and are sometimes called *hardware-based replication*. Remote Mirror and Copy migration offers the following advantages:

- High performance
- Operating system-independence
- Does not consume application host resources

For migration, use Global Copy, which copies data non-synchronously for both open and z/OS systems. Global Copy is a great tool for data migrations, especially for longer distances than are supported by Metro Mirror. With Global Copy, write operations complete on the source disk system before they are received by the target disk system. This capability is designed to prevent the local performance from being affected by wait time of writes on the remote system. Therefore, the source and target copies can be separated by any distance. The Metro Mirror migration method can be used if the distance between the two storage system is less than 300 km and if the synchronous replication does not affect existing production. Global Copy migration method requires you to stop host IO after the first pass while the remaining data is moved.

Both Global Copy and Metro Mirror require the application host operating system to acquire the new storage system volumes. The following factors must be accounted for when planning the data migration:

- Acquire new migration volumes and make them known to the operating system.
- Modify host and application configuration files.
- Perform data integrity checks.

When operating in Global Copy mode, the source volume sends a periodic, incremental copy of updated tracks to the target volume instead constant updates. Incremental updates cause less impact to application writes for source volumes and less demand for bandwidth resources, allowing a more flexible use of the available bandwidth.

Data Migration using Global Copy requires a consistent copy of the data to move applications and servers to the new system. By design, the data on the remote system is a *fuzzy copy*, with the volume pairs in a copy pending state. To create a consistent copy of the migrated data, the application must be quiesced and the volume pairs suspended. If the pairs are terminated before quiescing I/O and suspending the pairs, you might lose transactions at the remote site.

There are two ways to ensure data consistency during a migration using Global Copy:

- Shut down all applications at the primary site and allow the out-of-sync sectors to drain completely.
- Change the Global Copy relationships to synchronous mode when the out-of-sync sectors are approaching or at zero. When the out-of-sync sectors are fully drained, the pairs are in Full Duplex mode and there is a consistent relationship. Using **`go to sync`** is the more reliable method.

3.2.2 Migrating source volume set

The DS8000 command-line interface (DS8000 CLI) provides a full set of commands to perform all copy services functions. The DS8000 CLI code is installed on one of the supported open systems platforms (AIX, Windows, Solaris, Linux, or HP-UX). Although DS8000 CLI is not installed on a System z server, it can create and manage System z volumes in an IBM FlashCopy®, and Remote Mirror and Copy environment. The DS8000 CLI can also create volumes for all supported open systems servers. Use the latest version of DS8000 CLI for data migration because it is compatible with earlier versions of the DS8000.

For the migration process, define a remote mirror relation between existing storage system and new IBM DS8870. Before defining the Global Copy or Metro Mirror session, be sure to properly define the connectivity, which involves managing the logical paths. The logical paths define the relationship between a source logical subsystem (LSS) and a target LSS that is created over a physical path (IO port). In the following examples, creating and monitoring logical path are described as first four steps.

To create Remote Mirror and Copy paths, complete the following steps:

1. Before the paths can be created, you need to determine the remote WWNN and the available paths. Use the **lssi** command to determine the remote WWNN as shown in Example 3-1.

   Example 3-1 Using the lssi command to obtain the WWNN of the remote system

   ```
   dscli> lssi
   Name  ID               Storage Unit     Model WWNN             State  ESSNet
   ==========================================================================
   TUC02 IBM.2107-75CYM31 IBM.2107-75CYM30 961   5005076305FFD71E Online Enabled
   ```

2. Display the available PPRC ports between the two storage systems using the **lsavailpprcport** command as seen in Example 3-2. The **-remotewwnn** parameter is the WWNN determined in the previous step. Use the existing DS8000 (IBM.2107-75ZA571 in this example) as the **-dev** parameter, and the new DS8870 (IBM.2107-75CYM31 in this example) as the **-remotedev** parameter.

   Example 3-2 Query available PPRC ports

   ```
   dscli> lsavailpprcport -dev IBM.2107-75ZA571 -remotedev IBM.2107-75CYM31
   -remotewwnn 5005076305FFD71E -fullid 3B:82
   Local Port             Attached Port          Type
   ================================================
   IBM.2107-75ZA571/I0036 IBM.2107-75CYM31/I0030 FCP
   IBM.2107-75ZA571/I0336 IBM.2107-75CYM31/I0330 FCP
   ```

3. Use the WWNN and a port pair to create a path using the **mkpprcpath** command from the existing storage system (Example 3-3). In this example, two paths are created between LSS 3B on the existing DS8000 (IBM.2107-75ZA571) and LSS 82 on the new DS8870 (IBM.2107-75CYM31).

Example 3-3 Create path

```
dscli> mkpprcpath -dev IBM.2107-75ZA571 -remotedev IBM.2107-75CYM31 -remotewwnn
5005076305FFD71E -srclss 3B -tgtlss 82 I0036:I0030 I0336:I0330
CMUC00149I mkpprcpath: Remote Mirror and Copy path 3B:82 successfully established.
```

4. Query and verify the paths using the **lspprcpath** command as shown in Example 3-4. If you need to make changes, remove the path and re-create or modify it.

Example 3-4 Query PPRC paths between existing and new DS8000

```
dscli> lspprcpath -dev IBM.2107-75ZA571 3B
Src Tgt State   SS   Port  Attached Port Tgt WWNN
=======================================================
3B  82  Success 8201 I0036 I0030         5005076305FFD71E
3B  82  Success 8201 I0336 I0330         5005076305FFD71E
```

Creating the Remote Mirror and Copy pairs can be tedious task if many volumes are involved in the migration. An easy way to create many pairs is to use the DS8000 CLI in a scripting mode or use Tivoli Storage Productivity Center for Replication as described in 3.2.3, “Using Tivoli Storage Productivity Center for Replication” on page 36.

Migration with Global Copy mode is used when asynchronous replication is required. If you do not need use asynchronous data replication to migrate your volume set you can omit steps 5 and 6 and use Metro Mirror synchronous replication as shown in Example 3-7.

5. To start a Global Copy asynchronous replication from source volume 3B00 to target volume 8200 use the **mkpprc** command with **-type gcp** parameter as shown in Example 3-5.

Example 3-5 sing the mkpprc command to start asynchronous replication

```
dscli> mkpprc -dev IBM.2107-75ZA571 -remotedev IBM.2107-75CYM31 -type gcp -mode
full 3B00:8200
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 3B00:8200
successfully created.
```

6. After the Remote Mirror and Copy pairs are established, the data will start copying (Copy Pending state). Monitoring asynchronous replication on existing DS8000 is shown in Example 3-6.

Example 3-6 Using the lspprc command to monitor asynchronous replication

```
dscli> lspprc -dev IBM.2107-75ZA571 -remotedev IBM.2107-75CYM31 3B00:8200
ID        State        Reason Type        SourceLSS Timeout (secs) Critical Mode First Pass Stat
===============================================================================================
3B00:8200 Copy Pending -      Global Copy 3B        120            Disabled      True
```

7. For Metro Mirror, wait for the pairs to enter a Full Duplex state so that the data migration can be completed. For Global Copy, there is an intermediate state that is required to get to the Full Duplex state. You need to convert Global Copy into Metro Mirror mode by running the **mkpprc** command with **-type mmir** parameter as shown in Example 3-7. This operation is known as **go-to-sync**.

Example 3-7 Using the mkpprc command to convert Global Copy to Metro Mirror

```
dscli> mkpprc -dev IBM.2107-75ZA571 -remotedev IBM.2107-75CYM31 -type mmir 3B00:8200
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 3B00:8200 successfully created.
dscli> lspprc -dev IBM.2107-75ZA571 -remotedev IBM.2107-75CYM31 3B00:8200
ID        State       Reason Type         SourceLSS Timeout (secs) Critical Mode First Pass Stat
===============================================================================================
3B00:8200 Full Duplex -      Metro Mirror 3B        120            Disabled      Invalid
```

When all volumes relationships are in Full Duplex state, you can stop host IO, suspend Metro Mirror pairs, and perform data validation from the server and application. When the target volume set is validated, you can resume Metro Mirror relationships and start copying the data

to new IBM DS8870. Only tracks modified during validation process are sent to secondary volumes because DS8000 keeps a record of all changed tracks.

Finally, when validation is successful and cut-over decision is made, you can switch your application from existing storage to the new IBM DS8870 System Storage.

3.2.3 Using Tivoli Storage Productivity Center for Replication

Tivoli Storage Productivity Center for Replication provides management for the DS8000 copy services in various combinations. Tivoli Storage Productivity Center for Replication version 5.2 can also manage z/OS HyperSwap enabled sessions through an IP connection to a z/OS server. The following section provides an overview of the Tivoli Storage Productivity Center for Replication capabilities for DS8000 Global Copy and Metro Mirror relations that are involved in data migration scenarios.

For the migration process, define a Metro Mirror relation between existing storage system and new IBM DS8870. Before you start to operate Tivoli Storage Productivity Center for Replication with a DS8000 session, you must properly define the replication connectivity schema to implement. This process involves managing the logical paths. The logical paths define the relationship between a source LSS and a target LSS that is created over a physical path (IO port).

Tivoli Storage Productivity Center for Replication includes the Path Manager feature to provide control of logical paths when relationships between source and target storage systems are established. Path Manager helps you control the port pairing that Tivoli Storage Productivity Center for Replication uses when the logical paths are established and ensure redundant port combinations. It also keeps that information persistent for use when the path is terminated because of a suspended operation. Tivoli Storage Productivity Center for Replication provides the following options to create the logical paths and specify port pairing:

- Adding logical paths automatically: Tivoli Storage Productivity Center for Replication automatically picks the paths or uses paths that were established.
- Adding logical paths and creating port pairing by using a CSV file.
- Adding logical paths by using Tivoli Storage Productivity Center for Replication GUI.

Metro Mirror technology requires at least one logical path per LSS pair to establish a mirroring relationship. Tivoli Storage Productivity Center for Replication does not check the path availability before starting a mirroring relationship. If no paths are available between two DS8000s, this leads to an error at the first attempt to establish a pair relationship. Although the physical Fibre Channel link is bidirectional, the logical path definition is not. Therefore, to use the same port pair in both directions (which is not needed for migration process, but is required for disaster recovery) define a logical path in both storage systems (for example, one path from storage system A to B and another from storage system B to A).

With the Metro Mirror Single Direction session type, Metro Mirror replication is available only from the primary site and does not allow any action that inverts the replication direction. With this type of session, Tivoli Storage Productivity Center for Replication allows the user to perform the following tasks:

- Start the Metro Mirror.
- Pause and resume the Metro Mirror.

- Recover the Metro Mirror secondary site volumes, which makes them available for the use at the remote site.
- Restart the Metro Mirror following a recovery. This is accomplished by performing an incremental copy.

Figure 3-2 shows a defined Metro Mirror Single Direction session that is ready to start.

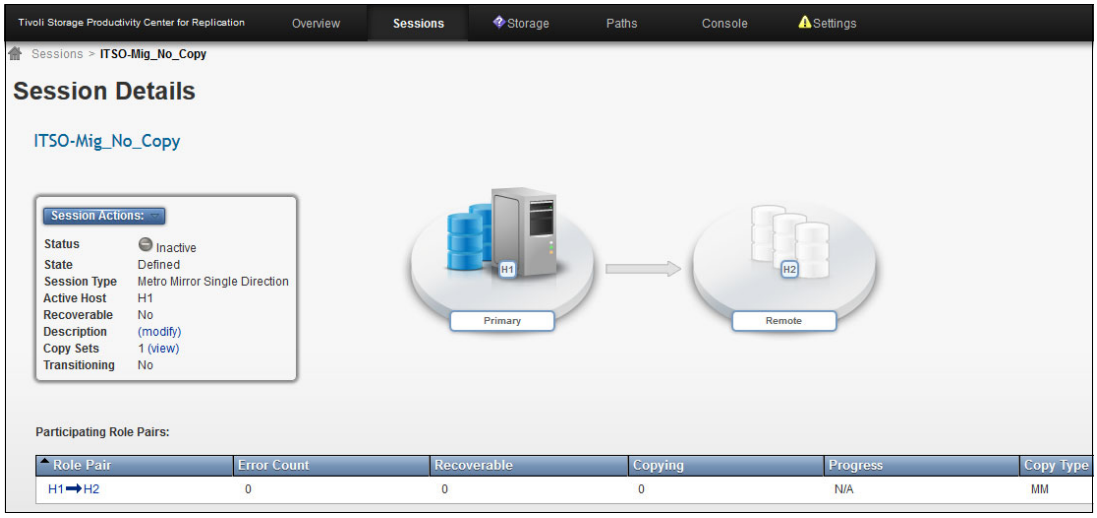

Figure 3-2 TPC-R Metro Mirror session defined and prepared to start.

When all pairs of volumes are defined in the copy set, the session named ITSO-Mig_No_Copy is started as shown in Figure 3-3.

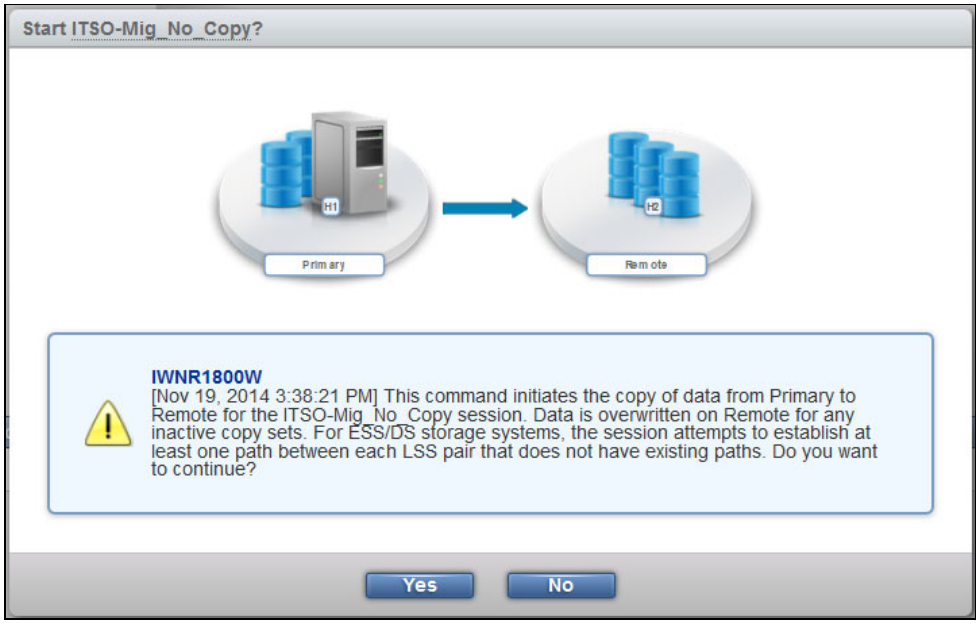

Figure 3-3 Metro Mirror session is starting

After the session starts, the source and target volume pairs are established and the data starts copying. For Metro Mirror, wait for the pairs to enter a Full Duplex state, at which time the data migration is complete as shown in Figure 3-4.

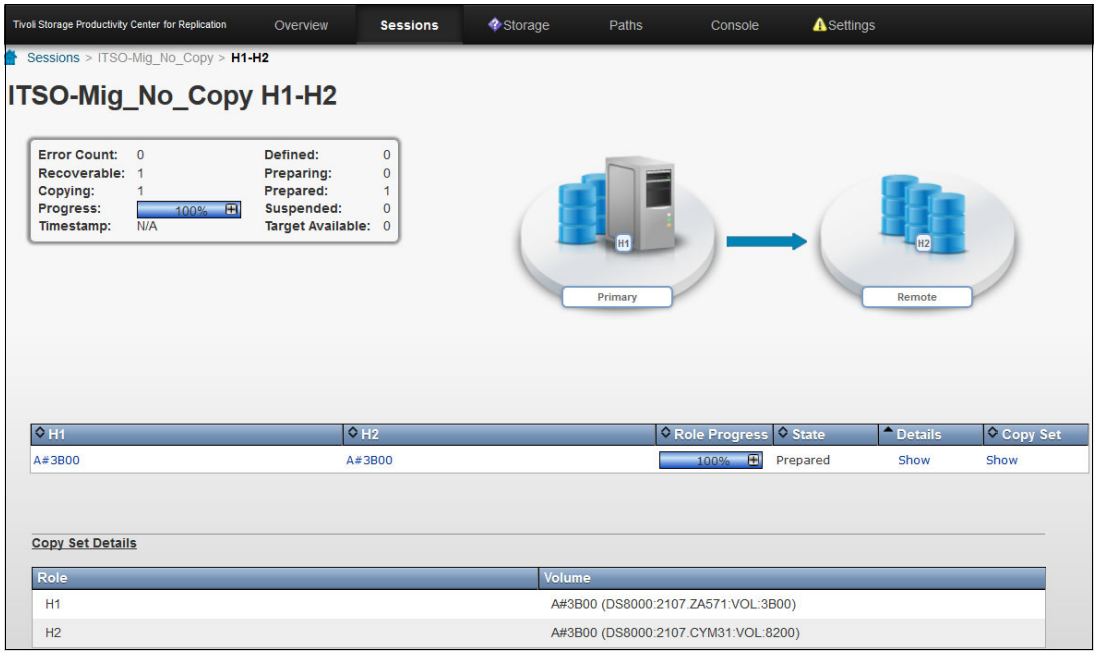

Figure 3-4 TPC-R session in full duplex state

4

Data Migration with 2-site Metro Mirror replication

This chapter discusses how to migrate data between DS8000 systems using the DS8000 series Metro Mirror function.

The following scenarios and considerations are presented:

- Replacing secondary
- Replacing primary
- Replacing primary and secondary DS8000
- Geographically Dispersed Parallel Sysplex (GDPS) considerations

This chapter includes the following sections:

- Initial configuration
- Migration process
- Geographically Dispersed Parallel Sysplex (GDPS) considerations

© Copyright IBM Corp. 2015. All rights reserved.

4.1 Initial configuration

The initial configuration consists of two DS8000 systems with a Metro Mirror relationship established as shown in Figure 4-1. During normal operation, the host accesses the primary volume set (A) in Site 1. The target volume set (B) is in Site 2. The Copy relation between volume set A and volume set B is a synchronous Metro Mirror relation. Volume set B is configured for the Standby or Disaster Recovery (DR) System to be accessed after an operation switch to site 2.

A volume set can consist a subset or all of the volumes in a single storage system, or even be spread across multiple systems. The granularity for considering the data migration should be set at the application level. Doing so ensures that you do not need to migrate all applications at the same time, and that you can schedule the migration at a time that suits the application availability requirements.

For simplification, these scenarios and illustrations represent a unidirectional Metro Mirror environment. If the site 1 and site 2 storage systems are serving primary volume sets to different set of applications or hosts, the migration case might be different for each volume set.

Access to the data is represented with a solid line, whereas the dashed line represents a configured but currently inactive connection.

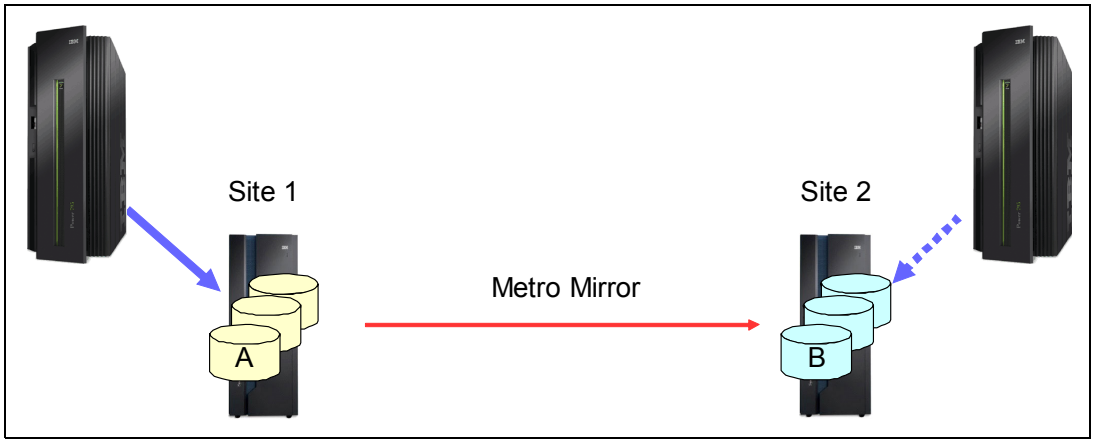

Figure 4-1 Initial 2 Site synchronous Mirror configuration

4.2 Migration process

The migration process uses the DS8000 Metro Mirror and Global Copy functions. The following migration scenarios are valid for all models and types of the IBM DS8000 series.

The starting point for the migration is always the setup as described in 4.1, “Initial configuration” on page 40.

The migration is always set up as a cascaded and asynchronous Global Copy relationship.

4.2.1 General considerations

The scenarios presented assume that the production is running on the host at site 1 and if changes are required to the host at site 1, those are performed during a scheduled change window.

The target audience for these cases is clients that do not provide the same processing power on the DR site (Site 2) than on Production (Site 1).

If moving the production between the sites is acceptable or standard business practice, the scenario described in 4.2.2, “Replacing secondary” on page 41 can be followed to replace the storage system in site 1 or site 2.

4.2.2 Replacing secondary

This section describes how data is migrated from one DS8000 to another where the volume set B is in the volume state “target”.

To complete these tasks, the copy methods Global Copy and Metro Mirror with incremental resync are used. By using the incremental resync feature, the cut-over from site 1 to site 2 can be performed without interrupting the I/O on the production system.

To replace the Secondary, the following steps are required:

1. Initial copy
2. Validate data
3. Cut-over

Initial copy phase

To initially copy the data to B’, an asynchronous copy, Global Copy, is established between B and B’. Figure 4-2 illustrates the volume set configuration for the migration process.

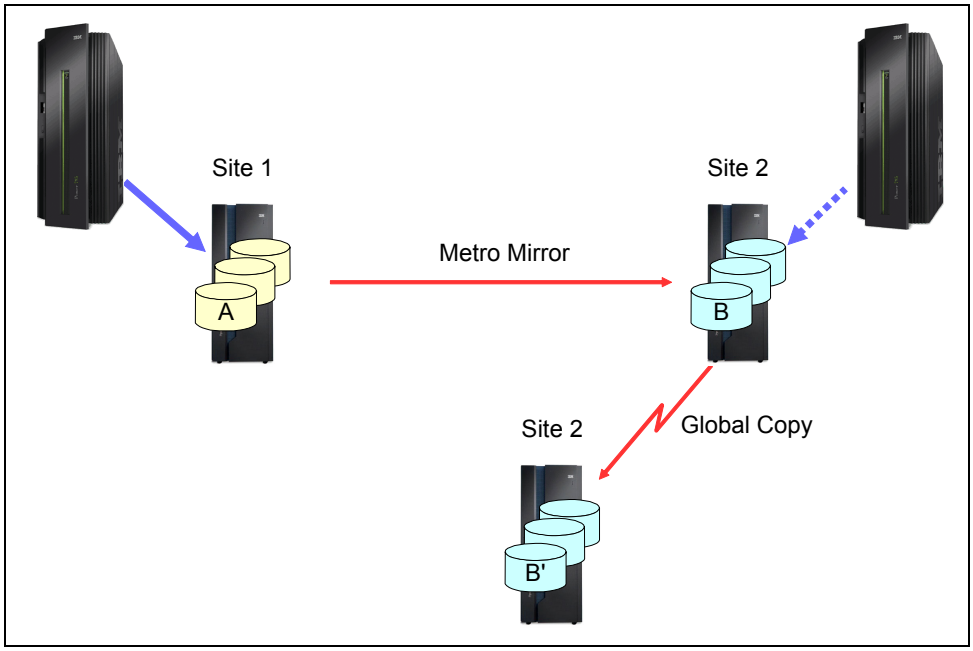

Figure 4-2 Migration target DS8000 B’ was added to the environment on Site 2

To attach the new storage system that hosts volume set B', the copy relation needs to be established with the following parameters:

- **-cascade** to allow the primary volume to be the secondary of another copy relationship.
- **-type gcp** to establish the relation as an asynchronous relation. A cascading synchronous relation is not allowed.

> **Note:** For the examples in this chapter, the following convention applies:
>
> Volume set A is hosted on IBM.2107-75ZA571. This storage system is retained.
> Volume set B is hosted on IBM.2107-75CYK71. This storage system is removed.
> Volume set B' is hosted on IBM.2107-75CZM21. This storage system is added.

Example 4-1 shows the **mkpprc** command to establish the Global Copy relationship and to specify the **-cascade** option.

Example 4-1 Creating a Global Copy relationship

```
mkpprc -dev IBM.2107-75CYK71 -remotedev IBM.2107-75CZM21 -cascade -type gcp 7900-7903:7900-7903
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 7900:7900 successfully ...
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 7901:7901 successfully ...
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 7902:7902 successfully ...
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 7903:7903 successfully ...
```

After creating the Global Copy relationship, wait until the value in the Column `First Pass Status` becomes `True`. The status can be checked using the dscli command **lspprc** as shown in Example 4-2. Depending on the amount of data and the bandwidth of the physical links, the time needed to get to the `True` status can vary. The `Out Of Sync Tracks` column indicates in this example also that all data have been copied between volume set B and B' and that they are in sync at that point.

Example 4-2 Status of the Global Copy relation with First Pass Status True

```
dscli> lspprc -l 7900-7903:7900-7903
ID        State        Reason Type        Out Of Sync Tracks Tgt Read Src Cascade ...... First Pass Status
================================================================================...... ==================
7900:7900 Copy Pending -      Global Copy 0                  Disabled Enabled     ...... True
7901:7901 Copy Pending -      Global Copy 0                  Disabled Enabled     ...... True
7902:7902 Copy Pending -      Global Copy 0                  Disabled Enabled     ...... True
7903:7903 Copy Pending -      Global Copy 0                  Disabled Enabled     ...... True
```

Validation phase

After the initial copy completed and the *First Pass Status* is *True* on volume set B and B', only actual data changes are transferred.

At this stage the validation of the migrated data can be performed. The data copied during this migration can be validated.

The validation phase is the break-out point to either continue with the migration or to return to the original environment.

For validation, it is required to get a consistent data state on the volume set B'. To that end, the application running on volume set A needs to create a consistency point. This can be achieved by either switching the application in backup mode, stopping the application, or shutting down the host.

After data consistency on volume set A can be ensured, pause the Metro Mirror connection. The application can then be resumed or restarted.

While the Metro Mirror connection is still in a paused state, wait until all data is copied from volume set B to volume set B'. After all data is copied and the `Out of Sync` counter is zero, issue a **`failoverpprc`** to allow host access to volume set B'.

After the failover (B' to B) is complete, the Metro Mirror (A to B) can be resumed to ensure DR readiness.

When using the DR server at Site 2 as a validation system, prepare the SAN zoning for accessing volume set B and volume set B'.

When starting the actual validation testing, enable the new SAN zoning to connect the server to volume set B'. After the testing is finished, shut down the server and change the SAN zoning back to the original volume set B until the final cut-over.

Figure 4-3 illustrates the final configuration for the validation phase.

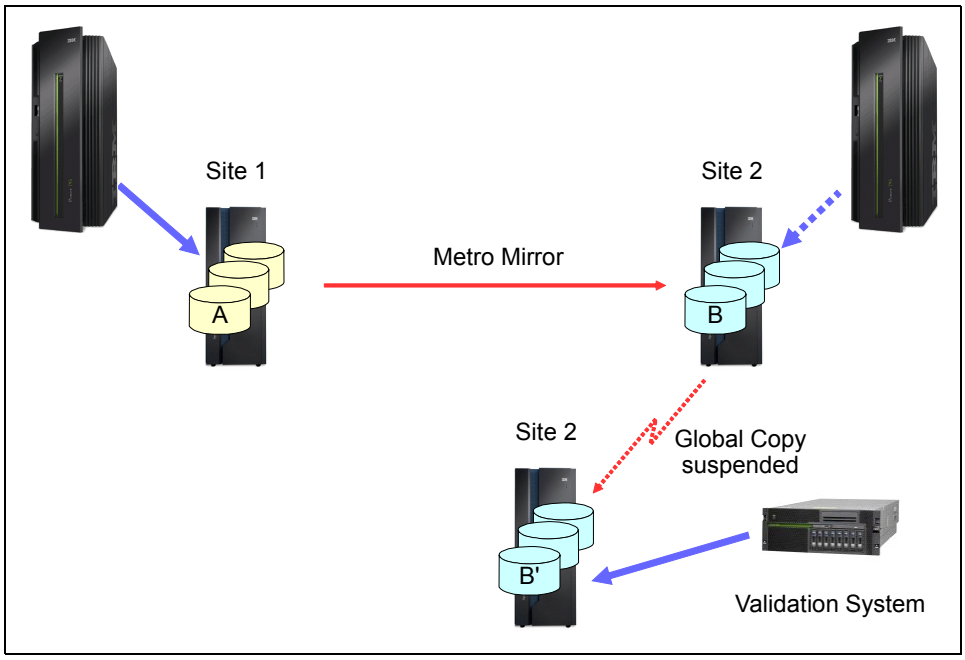

Figure 4-3 Validating the migration process and data

The following steps are performed during in the validation phase on the DS8000:

1. Suspend the Metro Mirror between volume set A and B
2. Suspend the Global Copy between volume set B and B'
3. Resume the Metro Mirror between volume set A and B
4. Failover to volume set B'
5. Run the validation tests on volume set B'
6. Failback volume set B to B' to update and invalidate changes written to volume set B'.

These are those general steps in detail:

1. Wait for the application or host accessing volume set A to write all data from the file system cache to the DS8000 to create a consistent data state. For more information, see 2.4, "Data validation testing" on page 24.

2. Pause the Metro Mirror relation by issuing a **pausepprc** command to pause the synchronous copy process between volume set A and volume set B, as shown in Example 4-3.

Example 4-3 Pausing the Metro Mirror relation between volume set A and volume set B

```
pausepprc -dev IBM.2107-75ZA571 -remotedev IBM.2107-75CYK71 7900-7903:7900-7903
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 7900:7900 relationship successfully paused.
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 7901:7901 relationship successfully paused.
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 7902:7902 relationship successfully paused.
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 7903:7903 relationship successfully paused.
```

3. I/O to volume set A can be started again. The Metro Mirror will not be resumed at this time. The resume is performed at step 6.
4. Verify that the `Out Of Sync-Tracks` counter for the Global Copy between volume sets B and B' are zero before continuing as shown in Example 4-4.

Example 4-4 lspprc output showing the Out Of Sync Tracks are zero

```
dscli> lspprc -l -dev IBM.2107-75CYK71 -remotedev IBM.2107-75CZM21 7900-7903:7900-7903
ID        State        Reason Type        Out Of Sync Tracks Tgt Read Src Cascade ......
=========================================================================================.....
7900:7900 Copy Pending -      Global Copy 0                  Disabled Enabled     ......
7901:7901 Copy Pending -      Global Copy 0                  Disabled Enabled     ......
7902:7902 Copy Pending -      Global Copy 0                  Disabled Enabled     ......
7903:7903 Copy Pending -      Global Copy 0                  Disabled Enabled     ......
```

5. Pause the Global Copy between volume set B and B' to allow a graceful failover. When pausing the Global Copy, the volume state of all volumes in volume set B should be in the `Suspended` state. Example 4-5 shows the **pausepprc** of the Global Copy relation and state of the volumes in the volume set B. The pause is issued against the source volume set B.

Example 4-5 Pausing the Global copy between B and B' and verifying the Suspend state for the volume sets

```
pausepprc -dev IBM.2107-75CYK71 -remotedev IBM.2107-75CZM21 7900-7903:7900-7903
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 7900:7900 relationship successfully paused.
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 7901:7901 relationship successfully paused.
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 7902:7902 relationship successfully paused.
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 7903:7903 relationship successfully paused.

lspprc -dev IBM.2107-75CYK71 -remotedev IBM.2107-75CZM21 7900-7903:7900-7903
ID        State     Reason      Type        SourceLSS Timeout (secs) Critical Mode First Pass Status
==========================================================================================================
7900:7900 Suspended Host Source Global Copy 79        60             Disabled      True
7901:7901 Suspended Host Source Global Copy 79        60             Disabled      True
7902:7902 Suspended Host Source Global Copy 79        60             Disabled      True
7903:7903 Suspended Host Source Global Copy 79        60             Disabled      True
```

6. Resuming the Metro Mirror connection between volume set A to volume set B reestablishes the DR readiness. The environment is fully DR capable after all changes that occurred on volume set A during the pause have been copied to B. How long it takes depends on the amount of Write I/O and time between the **pausepprc** and **resumepprc** operation shown in Example 4-6. The **resumepprc** is issued on volume set A.

Example 4-6 Resuming the Metro Mirror between volume set A and volume set B

```
resumepprc -type mmir -dev IBM.2107-75ZA571 -remotedev IBM.2107-75CYK71 7900-7903:7900-7903
CMUC00158I resumepprc: Remote Mirror and Copy volume pair 7900:7900 relationship successfully resumed. T...
CMUC00158I resumepprc: Remote Mirror and Copy volume pair 7901:7901 relationship successfully resumed. T...
```

```
CMUC00158I resumepprc: Remote Mirror and Copy volume pair 7902:7902 relationship successfully resumed. T...
CMUC00158I resumepprc: Remote Mirror and Copy volume pair 7903:7903 relationship successfully resumed. T...
```

7. To allow the validation system to access the volumes in volume set B' as shown in Figure 4-3 on page 43, the volumes must be in the state `Suspended` with the reason `Host Source`. To change the volume state and reason for volume set B', a failover operation on the storage system that hosts B' needs to be performed. Example 4-7 shows the volume state and reason before and after the failover and the dscli **`failoverpprc`** command. All of these commands need to be issued on the storage system that hosts B'.

Example 4-7 Failover of Global Copy relation between volume set B and B' and the different states of the volume set B.

```
lspprc  -dev IBM.2107-75CZM21  7900-7903
ID        State            Reason        Type        SourceLSS Timeout (secs) Critical Mode First Pass ...
==========================================================================================================
7900:7900 Target Suspended Update Target Global Copy 79        unknown        Disabled      Invalid
7901:7901 Target Suspended Update Target Global Copy 79        unknown        Disabled      Invalid
7902:7902 Target Suspended Update Target Global Copy 79        unknown        Disabled      Invalid
7903:7903 Target Suspended Update Target Global Copy 79        unknown        Disabled      Invalid

failoverpprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYK71  -type gcp 7900-7903:7900-7903
CMUC00196I failoverpprc: Remote Mirror and Copy pair 7900:7900 successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 7901:7901 successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 7902:7902 successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 7903:7903 successfully reversed.

lspprc -remotedev IBM.2107-75CYK71 -dev IBM.2107-75CZM21 7900-7903:7900-7903
ID        State     Reason      Type        SourceLSS Timeout (secs) Critical Mode First Pass Status
=================================================================================================
7900:7900 Suspended Host Source Global Copy 79        60             Disabled      True
7901:7901 Suspended Host Source Global Copy 79        60             Disabled      True
7902:7902 Suspended Host Source Global Copy 79        60             Disabled      True
7903:7903 Suspended Host Source Global Copy 79        60             Disabled      True
```

8. After the volumes in volume set B' are in *state* `Suspended` and *reason* `Host Source`, the validation system can access the data and the environment as illustrated in Figure 4-3 on page 43.

After performing these steps, the data validation and testing can be performed.

The outcome of the validation and testing is the decision point to proceed the migration or to stop the efforts and start over again.

When the validation and testing procedures have completed successfully, continue with the "Cut-over phase".

The backout process to return to the initial configuration is as described in 4.1, "Initial configuration" on page 40:

1. Remove the Global Copy relationship between volume set B and B'
2. Remote the associated copy paths for the volume sets.

No additional cleanup is required.

After a successful validation, the validation system is disconnected. Volume sets B and B' need to be synchronized again. The synchronization can be performed by issuing **`failbackpprc`** from volume set B to volume set B' as shown in Example 4-8. The **`failbackpprc`** command is issued on volume set B.

> **Caution:** Issuing the fail-back command to the incorrect system causes the wrong data to be kept and replicated to the peer system.

Example 4-8 Successful failback of the Global Copy from volume set B to volume set B' for synchronization'

```
failbackpprc -dev IBM.2107-75CYK71 -remotedev IBM.2107-75CZM21 -type gcp 7900-7903:7900-7903
CMUC00197I failbackpprc: Remote Mirror and Copy pair 7900:7900 successfully failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 7901:7901 successfully failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 7902:7902 successfully failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 7903:7903 successfully failed back.

lspprc -dev IBM.2107-75CYK71 -remotedev IBM.2107-75CZM21 7900-7903:7900-7903
ID        State        Reason Type        SourceLSS Timeout (secs) Critical Mode First Pass Status
==============================================================================h=====================
7900:7900 Copy Pending -      Global Copy 79        60             Disabled      True
7901:7901 Copy Pending -      Global Copy 79        60             Disabled      True
7902:7902 Copy Pending -      Global Copy 79        60             Disabled      True
7903:7903 Copy Pending -      Global Copy 79        60             Disabled      True
```

Cut-over phase

As soon as all changes are replicated from volume set B to B', the environment is ready for the final cut-over. Figure 4-4 on page 47 illustrates the environment before the migration.

Ensure that the Out-of-Sync-Tracks number for the relation B to B' is zero or near zero before starting the next step.

During the cut-over phase, the following steps are performed:

1. Enable Incremental Resync for the Metro Mirror on volume set A
2. Pausing the Metro Mirror between A and B
3. Pausing the Global Copy relation B and B'
4. Making volume set B' accessible for the host by performing a failover between B' and B
5. Creating a new Metro Mirror relation from A to B.

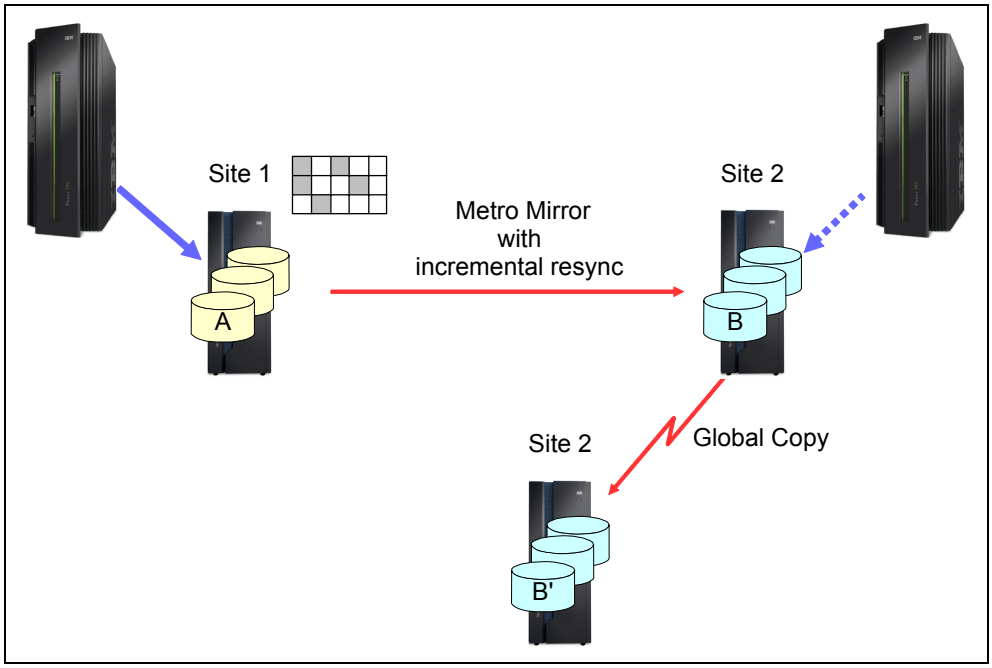

Figure 4-4 Using Metro Mirror with incremental resync between volume set A and volume set B

Just before the cut-over is performed, the Metro Mirror connection is configured for Incremental Resync. The incremental resync allows you to track data that is modified on A and needs to be copied to B' after the cut-over is complete. The DS8000 creates an internal bitmap to mark all the modified tracks that need to be copied.

> **Note:** The Incremental Resync option is usually used in Metro / Global Mirror (MGM) configurations. In the MGM configuration, the DS8000 internal microcode records in a bitmap, the track changes based on the Global Mirror Consistency Groups.
>
> Because Global Mirror is not used in this migration process, the bitmap does not get reset. Enabling the incremental resync option too far in advance might result in a lot of data being written twice to the System B', unnecessarily affecting the DR readiness.
>
> For more information about the Incremental Resync, see *IBM System Storage DS8000 Copy Services for Open Systems*, SG24-6788 or *IBM System Storage DS8000 Copy Services for IBM System z*, SG24-6787.

Incremental Resync option is enabled for this migration scenario with the **`-incrementalresync enablenoinit`**. Because there is no option for enabling the Incremental resync with **`resumepprc`**, **`mkpprc`** with the option **`-mode nocp`** is used as shown in Example 4-9. With **`lspprc -l`**, the **Incremental Resync** flag can be verified.

Example 4-9 Enable and verify Incremental Resync on the Metro Mirror relation between volume sets A and B

```
mkpprc -dev IBM.2107-75ZA571 -remotedev IBM.2107-75CYK71 -type mmir -incrementalresync enablenoinit -mode
nocp 7900-7903:7900-7903
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 7900:7900 successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 7901:7901 successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 7902:7902 successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 7903:7903 successfully created.

lspprc -l -dev IBM.2107-75ZA571 -remotedev IBM.2107-75CYK71 7900-7903:7900-7903
ID        State       Reason Type         Out Of Sync Tracks ... Status Incremental RESYNC...
```

```
==============================================================...====================...
7900:7900 Full Duplex -      Metro Mirror 0                  ...        Enabled     ...
7901:7901 Full Duplex -      Metro Mirror 0                  ...        Enabled     ...
7902:7902 Full Duplex -      Metro Mirror 0                  ...        Enabled     ...
7903:7903 Full Duplex -      Metro Mirror 0                  ...        Enabled     ...
```

After Incremental Resync has been enabled on the Metro Mirror copy relation, complete the following remaining steps:

1. Pause the Metro Mirror between A and B as shown in Example 4-10

Example 4-10 Pausing the Metro Mirror between volume set A and volume set B

```
pausepprc -dev IBM.2107-75ZA571 -remotedev IBM.2107-75CYK71 7900-7903:7900-7903
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 7900:7900 relationship successfully paused.
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 7901:7901 relationship successfully paused.
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 7902:7902 relationship successfully paused.
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 7903:7903 relationship successfully paused.
```

2. Wait for the out of sync tracks from B to B’ to be zero, as shown in Example 4-11.

Example 4-11 Verifying that Out-Of-Sync-Tracks between volume set B and B’ are zero

```
lspprc -l -dev IBM.2107-75CYK71 -remotedev IBM.2107-75CZM21 7900-7903:7900-7903
ID        State        Reason Type        Out Of Sync Tracks Tgt Read Src Cascade ...
==============================================================================...
7900:7900 Copy Pending -      Global Copy 0                  Disabled Enabled     ...
7901:7901 Copy Pending -      Global Copy 0                  Disabled Enabled     ...
7902:7902 Copy Pending -      Global Copy 0                  Disabled Enabled     ...
7903:7903 Copy Pending -      Global Copy 0                  Disabled Enabled     ...
```

3. Issue a failover for the Global Copy relationship B’ to B (Example 4-12). When you issue the failover, volume set B’ becomes accessible, so it can be configured as secondary volume set for the Metro Mirror of the new relationship A to B’.

Example 4-12 Performing the failover on volume sets B’ and B

```
failoverpprc -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYK71  -type gcp 7900-7903:7900-7903
CMUC00196I failoverpprc: Remote Mirror and Copy pair 7900:7900 successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 7901:7901 successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 7902:7902 successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 7903:7903 successfully reversed.
```

Attention: After performing step 4, you can no longer back out.

4. Create a new Metro Mirror session between the volume set A and B’ with the **`-incrementalresync override`** option. The **`lspprc`** option **`-fullid`** provides the individual device and remote device ID in the output. Verify that all volumes in the volume set A now have a Metro Mirror relationship to the storage system that hosts volume set B’. Refer to Example 4-13.

Note: For an IBM i PowerHA HyperSwap setup between A and B, check for any possibly existing "Multi-target internal" incremental resync pairs between A and B' using the **`lspprc`** command with the **`-fullid -multtgt`** parameters.

Remove these internal pairs only from A and B' using the **`rmpprc`** command prior to performing the incremental override from A to B', which would otherwise fail.

Example 4-13 Creating the new Metro Mirror relationship between volume set A and volume set B'

```
mkpprc -type mmir -remotedev IBM.2107-75CZM21 -incrementalresync override -cascade 7900-7903:7900-7903
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 7900:7900 successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 7901:7901 successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 7902:7902 successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 7903:7903 successfully created.

lspprc -dev IBM.2107-75ZA571 -fullid 7900-7903
ID                                          State       Reason Type         SourceLSS          ...
==============================================================================================...
IBM.2107-75ZA571/7900:IBM.2107-75CZM21/7900 Full Duplex -      Metro Mirror IBM.2107-75ZA571/79...
IBM.2107-75ZA571/7901:IBM.2107-75CZM21/7901 Full Duplex -      Metro Mirror IBM.2107-75ZA571/79...
IBM.2107-75ZA571/7902:IBM.2107-75CZM21/7902 Full Duplex -      Metro Mirror IBM.2107-75ZA571/79...
IBM.2107-75ZA571/7903:IBM.2107-75CZM21/7903 Full Duplex -      Metro Mirror IBM.2107-75ZA571/79...
```

5. Reconfigure the host access to the new storage system now hosting volume set B'.

Figure 4-5 illustrates the final configuration. When creating the Metro Mirror relationship between A and B', the old relation between A and B is withdrawn by the DS8000 microcode, and the migration is completed.

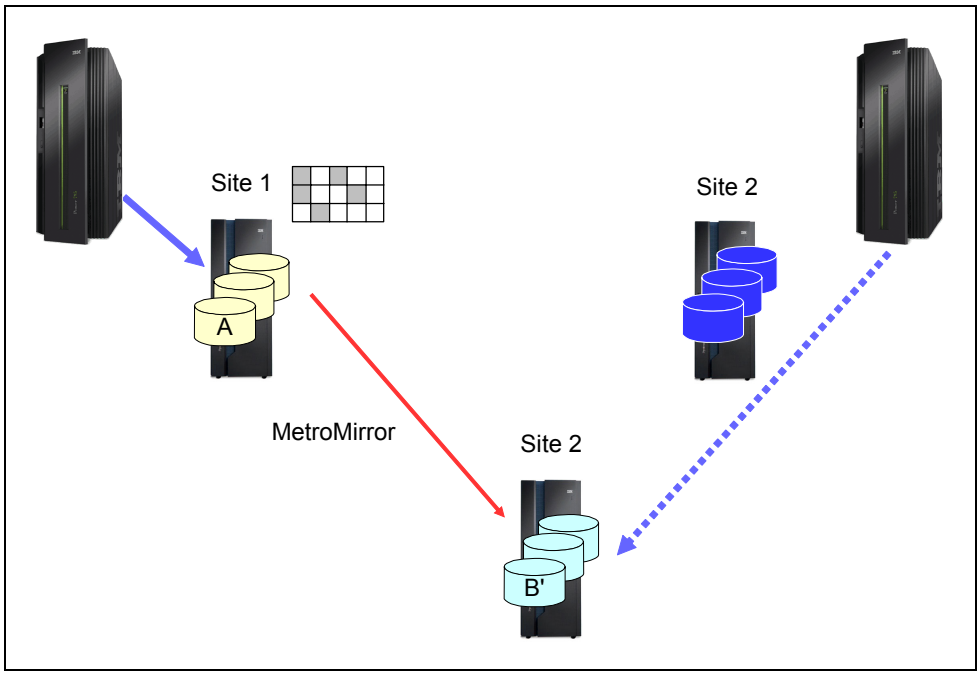

Figure 4-5 Final configuration with the new Storage System B

As post migration step, a cleanup on B' is needed.

Volume set B still has the configuration for the Global Copy relation with volume set B'. This relation needs to be removed, as it is not needed any more. To remove it, issue **`rmpprc -at src`** option. To identify the Global Copy relation that needs to be removed, the **`lspprc`** options **`-type gcp`** and **`-fullid`** can be used as shown in Example 4-14.

Example 4-14 removing the Global Copy relation between volume set B' and volume set B.

```
lspprc -fullid -type gcp 7900-7903
ID                                          State     Reason      Type        SourceLSS           ...
==============================================================================================...
IBM.2107-75CZM21/7900:IBM.2107-75CYK71/7900 Suspended Host Source Global Copy IBM.2107-75CZM21/79 ...
IBM.2107-75CZM21/7901:IBM.2107-75CYK71/7901 Suspended Host Source Global Copy IBM.2107-75CZM21/79 ...
IBM.2107-75CZM21/7902:IBM.2107-75CYK71/7902 Suspended Host Source Global Copy IBM.2107-75CZM21/79 ...
IBM.2107-75CZM21/7903:IBM.2107-75CYK71/7903 Suspended Host Source Global Copy IBM.2107-75CZM21/79 ...

rmpprc -at src -remotedev IBM.2107-75CYK71 7900-7903:7900-7903
CMUC00160W rmpprc: Are you sure you want to delete the Remote Mirror and Copy volume pair relationship
7900-7903:7900-7903:? [y/n]:y
CMUC00155I rmpprc: Remote Mirror and Copy volume pair 7900:7900 relationship successfully withdrawn.
CMUC00155I rmpprc: Remote Mirror and Copy volume pair 7901:7901 relationship successfully withdrawn.
CMUC00155I rmpprc: Remote Mirror and Copy volume pair 7902:7902 relationship successfully withdrawn.
CMUC00155I rmpprc: Remote Mirror and Copy volume pair 7903:7903 relationship successfully withdrawn.
```

The Metro Mirror and Global Copy relation on volume set B storage system can also be removed.

After all Volumes have been migrated, and no host is accessing the system any more, the logical configuration can be removed on the volume set B storage system and the system can be discontinued or reconfigured for a different use.

4.2.3 Replacing primary

This section describes a method for replacing the primary system that hosts volume set A in Site 1. Because Site 1 is serving the volumes for the active host, a downtime is required for reconfiguring the host connection.

If the application can easily be swapped to run from Site 2 with volume set B, use the migration scenario described in 4.2.2, “Replacing secondary” on page 41.

If you prefer to stay with the production in Site 1 and volume set A, the following scenario applies.

Initial Copy

The setup for the initial copy, shown in Figure 4-6 is identical to the initial copy setup used in 4.2.2, “Replacing secondary” on page 41, except that the new storage system is at Site 1 instead of Site 2.

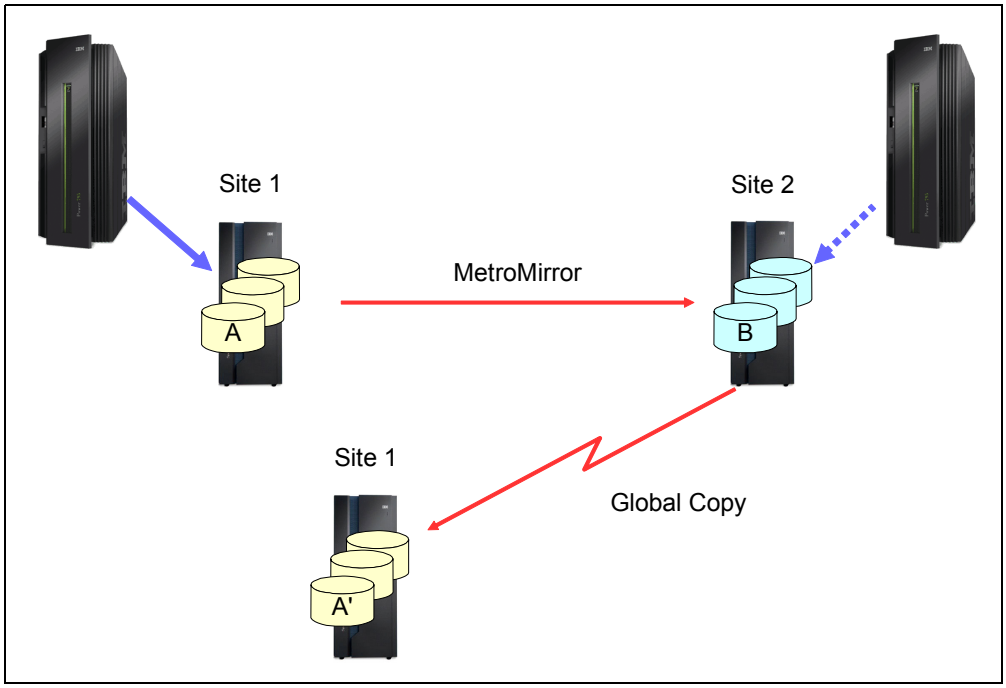

Figure 4-6 Setting up a new DS8000 in Site 1 to replace the DS8000 hosting volume set A.

Validation Phase

The validation is equal to the validation phase for replacing the secondary as detailed in "Validation phase" on page 42. The major difference is that the production system needs to be used as validation system. In this case, a downtime or site switch is required for the time of the validation test.

The validation test is performed in the following sequence:

1. Stop the production on the site 1 server.
2. Wait for all out of sync tracks between A to B and B to A' are zero.
3. Perform a failover from A' to B to make the volumes host accessible.
4. Verify that the out of sync tracks the relation B to A' do not change and remain zero.
5. Change the SAN Zoning that the server can access A' instead of A.
6. Start the server with the volume set A' and run the validation test.

After performing the validation testing, restore the original SAN configuration and reestablish the Global copy connection between B and A'.

Cut-over phase

After the validation has been performed on the volume set A', the Global Copy between volume set B and A' is resumed using the failback command as seen in Example 4-8 on page 46.

> **Important:** During the subsequent migration steps, no write I/O is allowed to volume set A. Otherwise, the data integrity between old volume set A and the new volume set B' is not guaranteed.

After the Global Copy relation between B and A' shows zero Out-of-Sync tracks:

1. Stop I/O to volume set A.

2. Verify that all Out-of-Sync-Tracks are zero for the relation between volume set A to volume set B.
3. Remove the Metro Mirror connection from A to B as shown in Example 4-15.

> **Note:** For the examples in this section, the following convention applies:
>
> Volume set A is hosted on IBM.2107-75ZA571. This storage system is removed.
> Volume set A' is hosted on IBM.2107-75CZM21. This storage system is added.
> Volume set B is hosted on IBM.2107-75CYK71. This storage system is retained.

Example 4-15 Removing the Metro Mirror relation between A and B

```
rmpprc -dev IBM.2107-75ZA571 -remotedev IBM.2107-75CYK71 7900-7903:7900-7903
CMUC00160W rmpprc: Are you sure you want to delete the Remote Mirror and Copy volume pair relationship
7900-7903:7900-7903:? [y/n]:y
CMUC00155I rmpprc: Remote Mirror and Copy volume pair 7900:7900 relationship successfully withdrawn.
CMUC00155I rmpprc: Remote Mirror and Copy volume pair 7901:7901 relationship successfully withdrawn.
CMUC00155I rmpprc: Remote Mirror and Copy volume pair 7902:7902 relationship successfully withdrawn.
CMUC00155I rmpprc: Remote Mirror and Copy volume pair 7903:7903 relationship successfully withdrawn.
```

4. Verify that all Out-of-Sync-Tracks for Global Copy relation between the volume set B and volume set A' are zero. Refer to Example 4-16.

Example 4-16 Verify zero Out-of-Sync-Tracks between volume set B and volume set A'

```
lspprc -l -dev IBM.2107-75CYK71 -remotedev IBM.2107-75CZM21 7900-7903:7900-7903
ID        State        Reason Type        Out Of Sync Tracks Tgt Read Src Cascade...
==========================================================================...
7900:7900 Copy Pending -      Global Copy 0                  Disabled Enabled    ...
7901:7901 Copy Pending -      Global Copy 0                  Disabled Enabled    ...
7902:7902 Copy Pending -      Global Copy 0                  Disabled Enabled    ...
7903:7903 Copy Pending -      Global Copy 0                  Disabled Enabled    ...
```

5. Fail over volume set A' to volume set B to allow the host to access the volumes in volume set A', as shown in Example 4-17.

Example 4-17 Making volume set A host accessible

```
dscli> failoverpprc -type mmir -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYK71 7900-7903:7900-7903
CMUC00196I failoverpprc: Remote Mirror and Copy pair 7900:7900 successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 7901:7901 successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 7902:7902 successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 7903:7903 successfully reversed.
```

6. Fail back volume set A' to volume set B with **-type mmir** as shown in Example 4-18. This action reverses the copy direction of the previous Global Copy and also converts the Global Copy into a synchronous Metro Mirror connection.

Example 4-18 Establishing the Metro Mirror between volume Set A' and volume Set B again.

```
failbackpprc -type mmir -dev IBM.2107-75CZM21 -remotedev IBM.2107-75CYK71 7900-7903:7900-7903
CMUC00197I failbackpprc: Remote Mirror and Copy pair 7900:7900 successfully failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 7901:7901 successfully failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 7902:7902 successfully failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 7903:7903 successfully failed back.
```

7. Configure the host to run I/O on volume set A' instead of A

8. Restart I/O on volume set A' as new primary as illustrated in Figure 4-7.

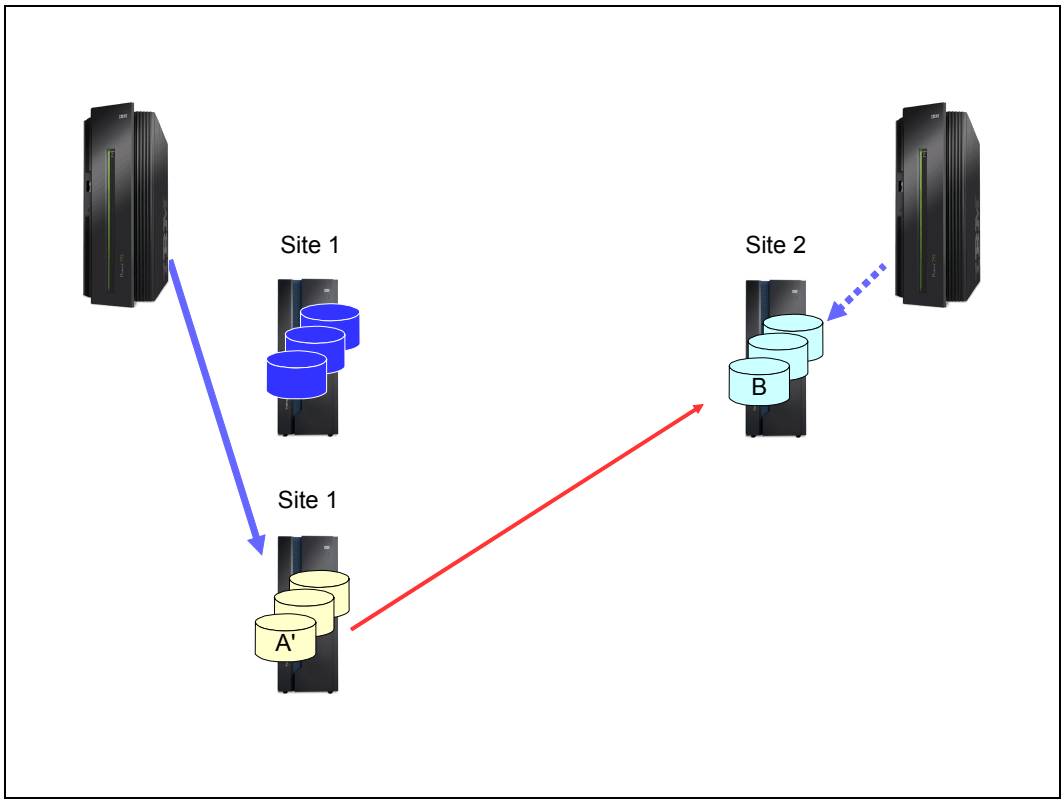

Figure 4-7 Final configuration with replacing the DS8000 hosting the primary volume set A with A'

4.2.4 Replacing primary and secondary DS8000

This section describes how to replace the primary and auxiliary storage systems hosting volume set A and volume set B.

Initial copy

The initial data copy uses a cascading configuration, with a second cascading Global Copy that is used to copy the data from volume set B' to volume set A'. Figure 4-8 provides an overview of the initial copy environment with two cascading Global Copy relations between volume set B to B' and then to A'.

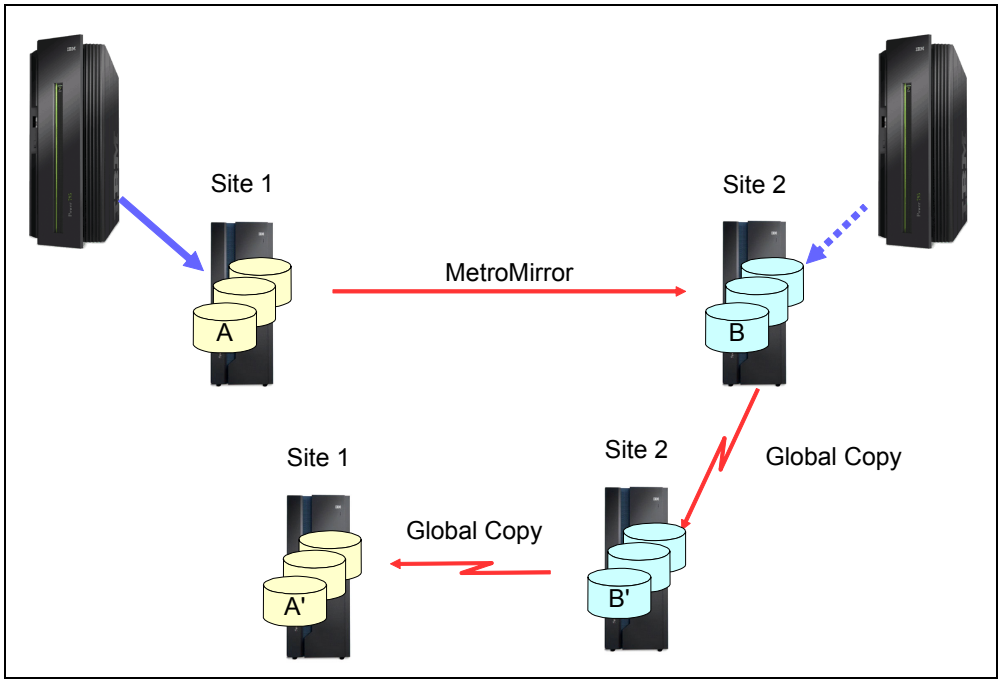

Figure 4-8 Initial copy setup for replacing the primary and secondary DS8000.

Validation phase

As soon as the initial copy has been completed, the verification and testing phase can be started. The validation and testing can be performed on volume set A' at the end of the cascaded copy operations, as depicted in Figure 4-9.

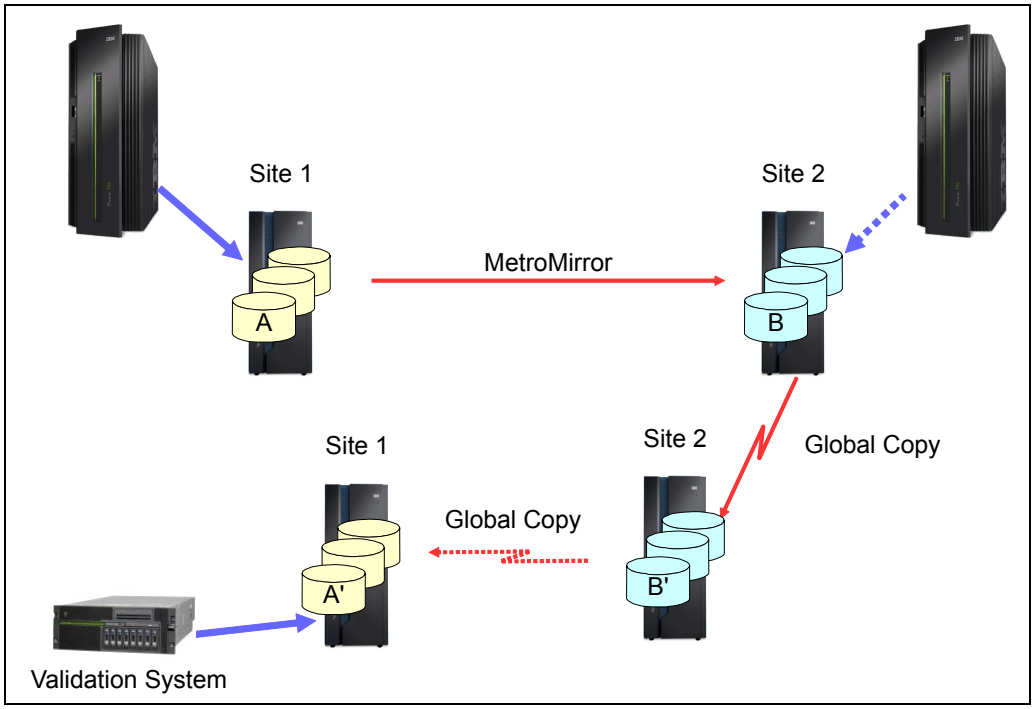

Figure 4-9 Validation with volume set A'

If you want to perform a validation that includes the volume set A' and B', then the Global Copy needs to be stopped, which requires an extra initial copy performed after the validation and testing is completed.

Cut-over phase

Migrating to the final configuration as shown in Figure 4-10 requires the same steps as described in “Cut-over phase” on page 51:

1. Stop I/O to volume set A.
2. Remove the Global Copy relation between volume set B and volume set B’.
3. Perform the failover from volume set A’ to B’.
4. Perform a failback from volume set A’ to B’.

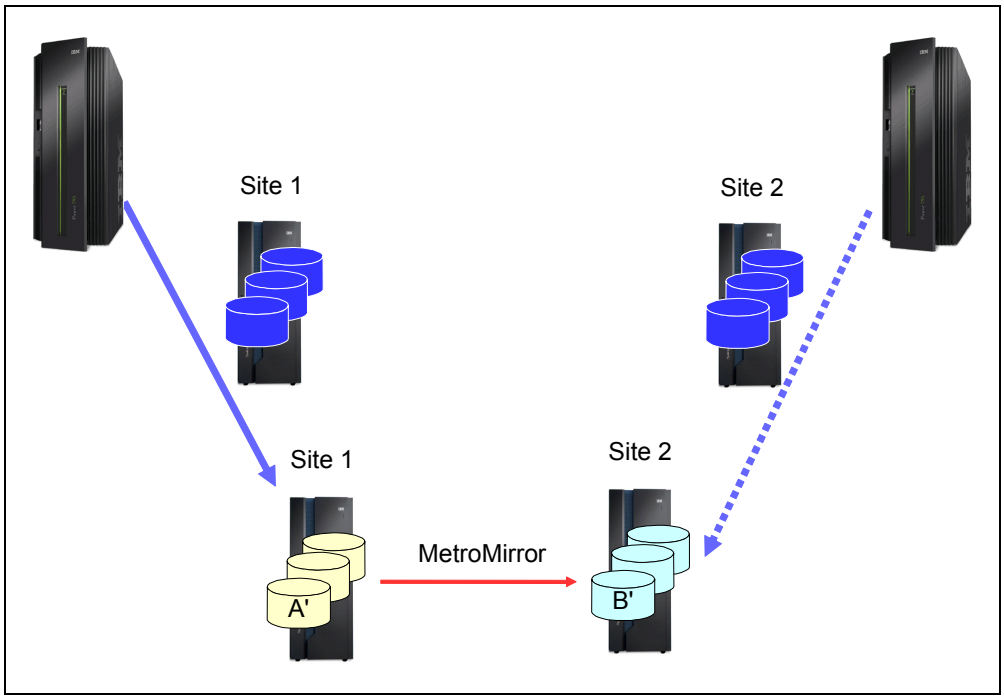

Figure 4-10 Cut-over phase

4.3 Geographically Dispersed Parallel Sysplex (GDPS) considerations

In addition to the general considerations discussed in Chapter 1, “Introducing disk data migration” on page 1, some additional considerations need to be taken into account when replacing storage systems in a Metro Mirror environment with GDPS managing the environment. Copy Services Incremental Resynchronization (IR) is required to replace the A and B volume sets in the following examples. Using GDPS and IR allows the applications systems to remain active and HyperSwap is only disabled for a short time. For more information about GDPS terminology, see 9.1, “GDPS Terminology” on page 186.

To minimize the number of configuration changes required, configure the new volume sets as similar to the one being removed. This means, if at all possible, use the same LSS range, SSID range, UCB range, and PPRC links for the new volume set that were in use on the old volume set.

For all migrations, this example makes use of the cascade option for Global Copy to complete the initial data migration to the new volume set.

If Remote Pair FlashCopy (RPFC) is in use, the FlashCopy target volumes cannot be in a cascaded global copy relationship. The FlashCopy must be withdrawn before establishing the cascaded relationship.

In a GDPS/Metro Mirror environment, there are at least one and in most cases two GDPS controlling systems. This book refers to these controlling systems as follows:

- Site 1 Kp-sys, which runs on non-PPRCed disk on a Site 1 server on Site 1 disk (optional, but highly recommended).
- Site2 Kp-sys in Site2 running on a Site2 server from local disk in Site2.

Important: Special consideration needs to be taken into account when migrating a volume set, which includes the GDPS Kx-systems. For instance, if the Site 1 Kp-sys system volumes are on volume set A that you are migrating, a Global Copy relationship needs to be created from A to A' for these devices. The system will need to be shut down and the relationships deleted. After the copy completes, IPL this Kp-sys on the new A' volume set after it is made available. The same is true for the Site 2 Kp-sys.

Initial configuration

Figure 4-1 on page 40 illustrates the basic layout of a 2-site Metro Mirror configuration. The host system accesses the primary volume set in site 1. There is a synchronous Metro Mirror relationship to the site 2 volume set.

Note: It is advised that you consult with your IBM representative during the planning process of the migration project to validate your replication network requirements. The correct network infrastructure helps you to avoid inadvertently impacting the production and replication environments during the data migration.

Replacing the B volume set

The approach described here to replace the B volume set uses both GDPS scripts and the Copy Services IR function to complete the migration. A cascaded relationship is created from the B volume set to the new B' volume set. It is assumed that the B to B' cascaded relationship has completed First Pass (Initial Copy) before the actual migration, so the time required to obtain 100% copy status is minimal.

Additional Fibre Channel (FC) connectivity is required between B to B' and A to B'.

B' initial copy phase

Establish paths and a cascaded Global Copy (PPRC/XD) relationship from the B devices to the B' devices. These relationships should be established well in advance of your actual migration cut-over. The copy time to complete first pass depends on your network bandwidth and the configuration size of the volume set you are migrating from/to. Figure 4-11 illustrates the Metro Mirror configuration with the cascaded copy.

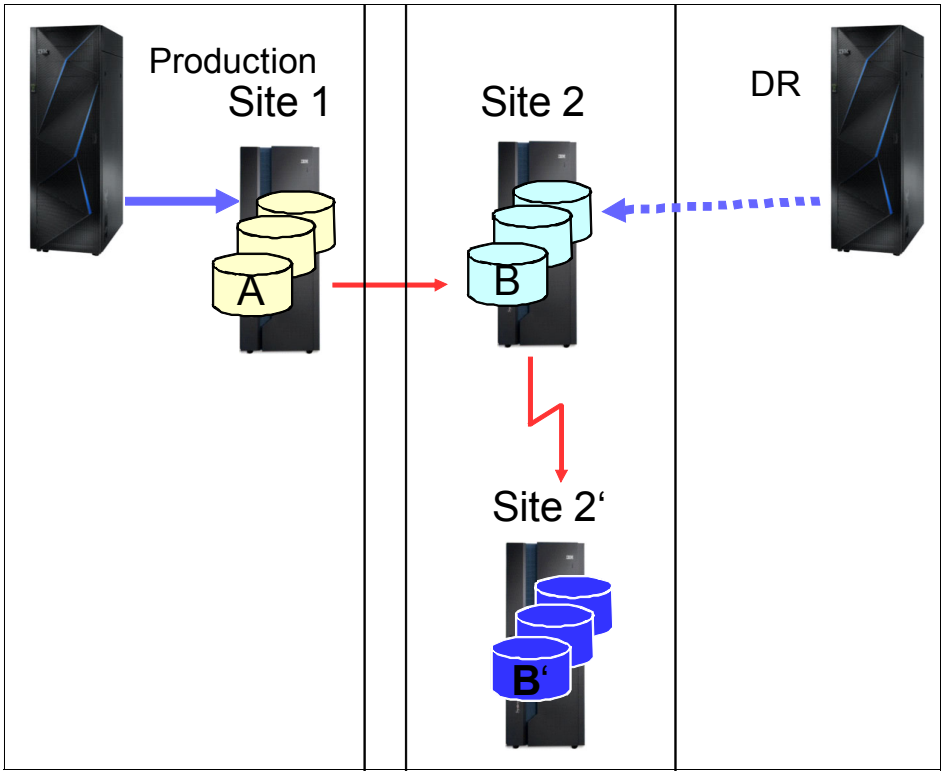

Figure 4-11 MM with Cascaded Global Copy to B'

Start IR for the A to B relationship with the NOCOPY option. Example 4-19 illustrates the TSO Copy Services command to start IR.

Example 4-19 Start IR A to B with NOCOPY

```
CESTPAIR DEVN(X'Adevn') -
          PRIM(X'Assid' Aserial X'Acca' X'Alss') -
          SEC(X'Bssid' Bserial X'Bcca' X'Bclss')  -
          MODE(NOCOPY) OPTION(SYNC) INCRESYNC(NOCOPY)
```

B' Storage Migration Phase

Establish the paths from A to B' and move all couple data sets (CDSs) to Site 1 by running the GDPS script statement **CDS SITE1**. After the CDSs have been moved, vary all of the B CDS devices offline on all systems.

Freeze the A to B relationship by running GDPS/PPRC script statement **STOP SECONDARY** to stop all I/O to the B volume set. At this point, the A devices are being updated and changes are being recorded in the Change Record (CR) bitmap. The B devices are not being updated because Metro Mirror (A to B) is frozen.

Using standard operating procedures, shut down the Site 2 Kp-sys and any other systems that are on the B volume set that is being replaced.

Wait for the copy to complete B to B' validating so that the out-of-sync tracks=0 before moving forward. After the copy is 100% complete, issue a recover to the B devices for the A to B relationship.

Reverse the relationship between B and B'. First, establish the paths from B' to B. Issue a failover to the B' devices, B' to B followed by a failback B' to B.

Remove the B volume set and add the B' volume set into the configuration. Use standard operating procedures that you would normally perform to swap out hardware such as swapping cables or activating an input/output definition file (IODF) to remove the old volume set and add the new volume set.

Establish the A to B' relationship using IR. Example 4-20 illustrates the TSO command to accomplish the resynchronization from A to B' without a full copy.

Example 4-20 Resync A to B' with IR

```
CESTPAIR DEVN(X'Adevn') -
          PRIM(X'Assid' Aserial X'Acca' X'Alss') -
          SEC(X'B'ssid' B'serial X'B'cca' X'B'clss')  -
          MODE(INCRES) OPTION(SYNC) FORCE(YES)
```

Monitor the A to B' copy progress. After the pairs are in DUPLEX state, start IR from A to B' with NOCOPY. Example 4-21 illustrates the TSO command to start the IR recording.

Example 4-21 Start IR A to B' with NOCOPY

```
CESTPAIR DEVN(X'Adevn') -
          PRIM(X'Assid' Aserial X'Acca' X'Alss') -
          SEC(X'B'ssid' B'serial X'B'cca' X'B'clss')  -
          MODE(NOCOPY) OPTION(SYNC) INCRESYNC(NOCOPY)
```

The configuration is now running A to B' to B with IR recording turned on A to B'.

Perform data validation on the new B' volume set. If the Site 2 Kp-sys is on the B volume set being replaced, IPL it on B' and validate the configuration. The Site 2 Kp-sys requires full access to all A and B volume sets. Therefore, whether the Site 2 Kp-sys is on the B volume set being replaced or not, it can be used to validate the new volume set being replaced. This is your break-out (GO/NOGO) decision point. If you determine there are errors during the validation process and you need to back out, see "Back-out scenario for B'" on page 60.

Cut-over to the new volume set

Load the GEOPARM configuration by using GDPS/PPRC Configuration Management panel (option C) to find the new B' volume set. If the configuration between the B and B' volume sets changed, the GEOPARM needs to be updated with the changes before it is loaded.

Vary the B' CDS volumes online to all systems and move the CDSs to the Normal configuration with the primaries on A and alternates on B'. Moving the CDSs can be done with GDPS script statement **`CDS NORMAL`**.

Delete the B’ to B relationships and paths. The configuration is now A to B’ as shown in Figure 4-12.

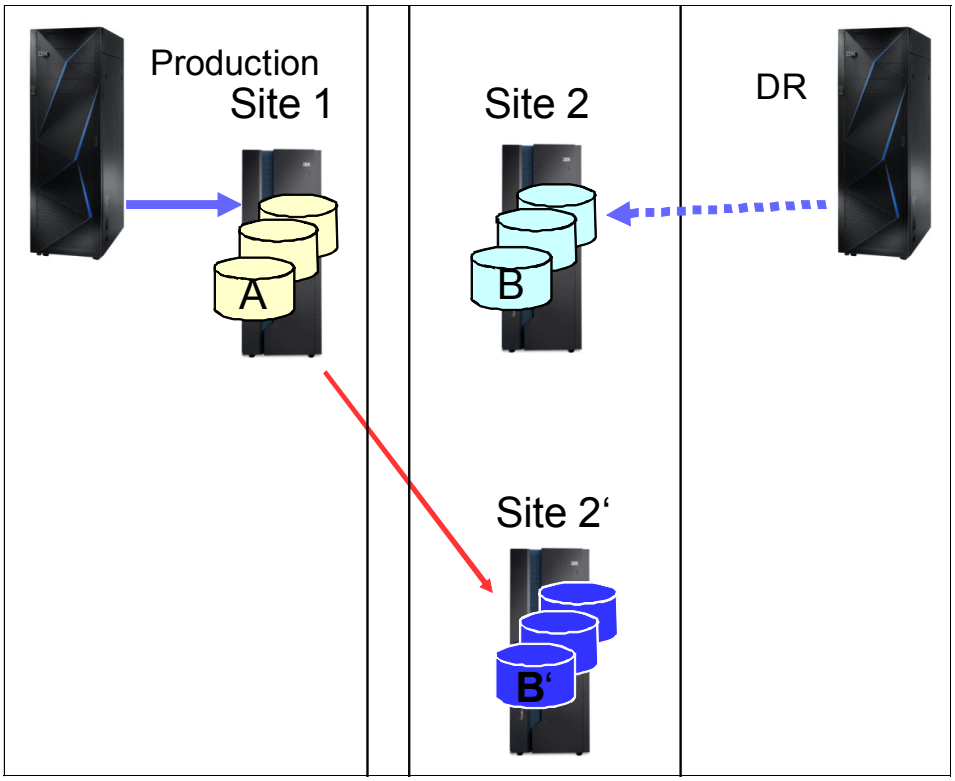

Figure 4-12 Final Metro Mirror Configuration with B’

Back-out scenario for B’

Complete the following steps to back out of B’:

1. Freeze the A to B’ relationship by running GDPS/PPRC script statement **STOP SECONDARY** to stop all I/O to the B’ volume set. At this point, the A devices are being updated and changes are being recorded in the CR bitmap. The B’ devices are not being updated because Metro Mirror (A to B’) is frozen.
2. Wait for the copy to complete B’ to B so that the out-of-sync tracks=0 before moving forward. After the copy is 100% complete, issue a recover to the B’ devices for the A to B’ relationship.
3. Reverse the relationship between B and B’. First, establish the paths from B’ to B. Issue a failover to the B devices, B to B’ followed by a failback B to B’.
4. Remove the B’ volume set and add the B volume set into the configuration. Use standard operating procedures that you would normally perform to swap out hardware such as swapping cables or activating an IODF to revert to your original B volume set.
5. Reestablish the paths A to B and then establish the A to B relationship using IR. Example 4-22 illustrates the TSO command to accomplish the resynchronization from A to B without a full copy.

Example 4-22 Resync A to B with IR

```
CESTPAIR DEVN(X'Adevn') -
         PRIM(X'Assid' Aserial X'Acca' X'Alss') -
```

```
          SEC(X'Bssid' Bserial X'Bcca' X'Bclss')  -
          MODE(INCRES) OPTION(SYNC) FORCE(YES)
```

6. Monitor the A to B copy progress to ensure that it reaches full DUPLEX. Start IR from A to B. Example 4-23 illustrates the TSO command to start the IR recording.

Example 4-23 Start IR A to B

```
CESTPAIR DEVN(X'Adevn') -
          PRIM(X'Assid' Aserial X'Acca' X'Alss') -
          SEC(X'Bssid' Bserial X'Bcca' X'Bclss')  -
          MODE(NOCOPY) OPTION(SYNC) INCRESYNC(NOCOPY)
```

The configuration is now running A to B to B' with IR recording turned on A to B.

7. IPL the Site 2 Kp-sys on the B. Delete the B to B' relationships and paths.

Replacing the A volume set

The approach described here to replace the A volume set uses GDPS scripts, HyperSwap, and the Copy Services IR function to complete the migration. A cascaded relationship is created from the B volume set to the new A' volume set. It is assumed that the B to A' cascaded relationship has completed First Pass before the actual migration, so the time required to obtain 100% copy status is minimal.

Additional FC connectivity is required between A to A' and B to A'.

A' Initial copy phase

If you are replacing the volume set where your Site 1 Kp-sys volumes are, establish paths and then a Global Copy (PPRC/XD) relationship from A to A' for these volumes. These relationships should be established well in advance of your actual migration cut-over. The copy time to complete first pass depends on your network bandwidth and the configuration size of the volume set you are migrating from/to.

Figure 4-13 illustrates the MM configuration with the cascaded copy.

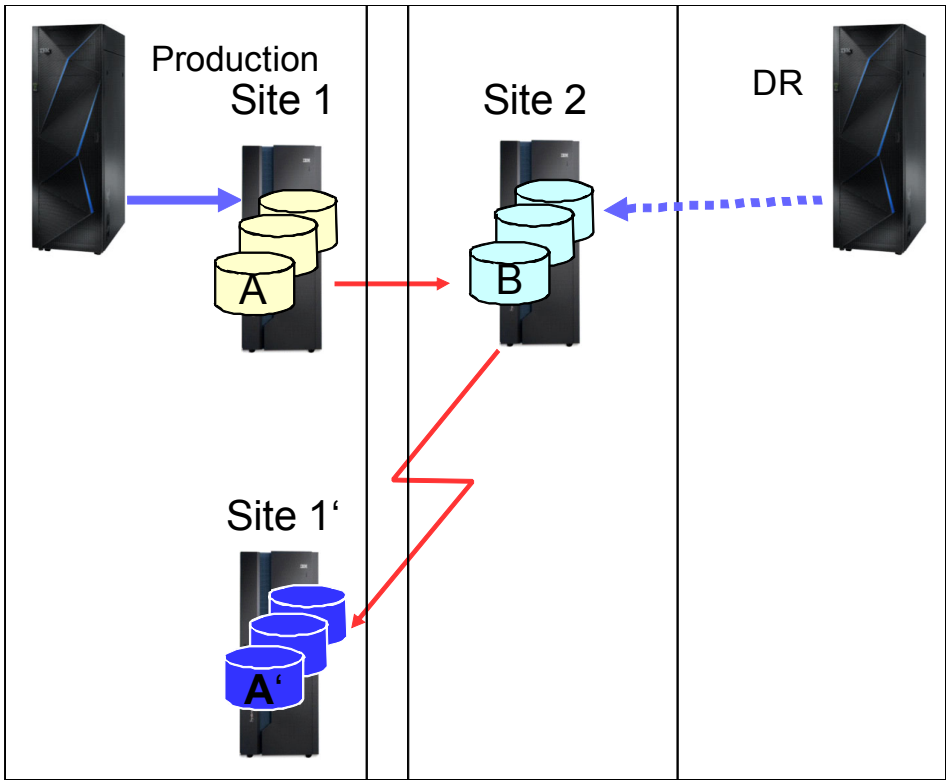

Figure 4-13 MM with Cascaded Global Copy to B’

A’ migration phase

HyperSwap from A to B by running the GDPS script statement **`SWTICH HYPERSWAP SUSPEND`** from the Master Kp-sys.

Using standard operating procedures, shut down the GDPS Site 1 Kp-sys system. Wait for the copy A to A’ to complete (out-of-sync tracks=0). Delete the A to A’ relationship.

You are now ready to remove the A volume set and add the A’ volume set into the configuration. Use standard operating procedures to complete this task, which you would normally perform to swap out hardware such as swapping cables or activating an IODF to remove the A system and add the A’ system.

IPL the GDPS Site 1 Kp-sys on A’. This is the GO/NOGO decision point. Perform data validation on the new A’ volume set. If you determine there are errors during the validation process and you need to back out, see “Back-out scenario for A’” on page 63.

Cut over to the new volume set

Recover the B device for the old A to B relationship. Load a GEOPARM with the new A’ volume set. Move all CDSs to the B devices by using GDPS script statement **`CDS SITE2`**.

Convert the Global Copy (XD) copy B to A’ relationships to synchronous mode. Run the GDPS script statement **`START SECONDARY`**. Run GDPS monitors and validate the environment is HyperSwap ready and green OK.

HyperSwap from B to A’ by running the GDPS script statement **`SWITCH HYPERSWAP SUSPEND`** from the Master Kp-sys.

Resync A' to B by running the GDPS script statement **START SECONDARY** from the Master Kp-sys.

Move the CDS primary and secondaries back to their normal configuration with primaries on Site 1 and alternates on Site 2. To do so, run GDPS script statement **CDS NORMAL**. The configuration is now A' to B as shown in Figure 4-14.

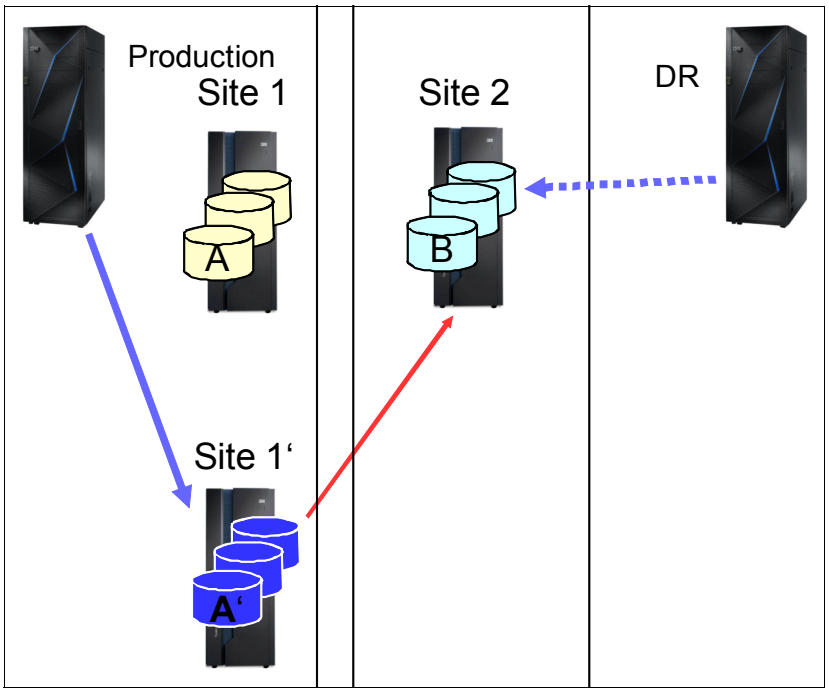

Figure 4-14 Final MM Configuration with A'

Back-out scenario for A'

Suspend the B to A' relationships. Resync the B to A relationships by running GDPS script statement **START SECONDARY**. Run GDPS monitors and validate the environment is HyperSwap ready and green OK.

HyperSwap from B to A by running GDPS script statement **SWITCH HYPERSWAP SUSPEND** from the Master Kp-sys.

Resync A to B by running GDPS script statement **START SECONDARY** from the Master Kp-sys.

Move the CDS primary and secondaries back to their normal configuration with primaries on Site 1 and alternates on Site 2. To accomplish this, run the GDPS script statement **CDS NORMAL**. The configuration is back to your original A to B configuration. At this point, you can either delete or resync the B to A' relationship depending on the reason the back out was required.

Replacing both the A and B Volume Sets

This section describes the procedure to replace both the primary and secondary volume set hosting volume set A' and volume set B'.

The initial data copy is using the same cascading method as described in the previous chapters, with one exception. A second cascading Global Copy is used to copy the data from volume set B' to volume set A'.

The approach described here to replace both the A and B volume sets uses GDPS scripts, HyperSwap, and the Copy Services IR function to complete the migration. A cascaded relationship is created from the B to B' to A' volume sets. It is assumed that the B to B' and B' to A' cascaded relationship has completed First Pass (initial copy) before the actual migration, so the time required to obtain 100% copy status is minimal.

Additional FC connectivity is required between A to A', B to B', and B' to A'.

A' and B' initial copy phase

If you are replacing the volume set where your Site 1 Kp-sys volumes are, establish paths and then a Global Copy (PPRC/XD) relationship from A to A' for these volumes. Establish the paths and PPRC/XD relationships from B to B' and B' to A'. These relationships should be established well in advance of your actual migration cut-over. The copy time to complete first pass depends on your network bandwidth and the configuration size of the volume set you are migrating from/to. Figure 4-15 illustrates the cascaded copy relationships from A to B to B' to A'.

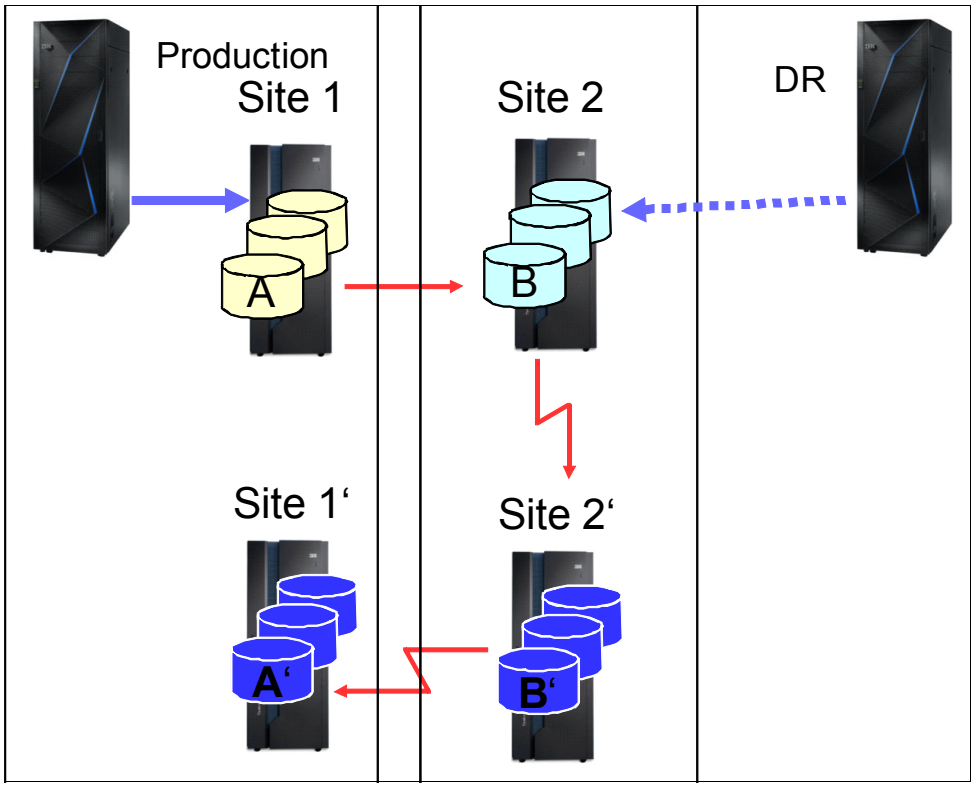

Figure 4-15 MM with Cascaded copy to B' and A'

A’ and B’ migration phase

Enable IR on the A to B relationships by running GDPS script statement **INCRESYN SETGM**. This script statement is only allowed to run when the GDPS environment is green OK and will briefly set the mirroring status to NOK and perform a suspend/resume of the PPRC pairs. In a HyperSwap environment, a Monitor2 or a **START SECONDARY** must be used to re-enable HyperSwap.

Freeze the A to B relationship by running GDPS/PPRC script statement **STOP SECONDARY** to stop all I/O to the B volume set. At this point, the A devices are being updated and changes are being recorded in the CR bitmap. The B devices are not being updated because Metro Mirror (A to B) is frozen.

Using standard operating procedures, shut down the Site 2 Kp-sys and any other systems that might be on the volume set being replaced.

Wait for the copy to complete B to B’ validating that the out-of-sync tracks are at 0 before moving forward. After the copy is 100% complete, issue a **delpair** to the B devices for the B to B’ relationship.

Remove the B volume set and add the B’ volume set into the configuration. Use standard operating procedures that you would normally perform to swap out hardware such as swapping cables or activating an IODF to remove the old volume set and add the new volume set.

Establish the A to B’ relationship using IR. Example 4-24 illustrates the TSO command to accomplish the resynchronization from A to B’ without a full copy.

Example 4-24 Resync A to B’ with IR

```
CESTPAIR DEVN(X'Adevn') -
          PRIM(X'Assid' Aserial X'Acca' X'Alss') -
          SEC(X'B’ssid' B’serial X'B’cca' X'B’clss')  -
          MODE(INCRES) OPTION(SYNC) FORCE(YES)
```

Monitor the A to B’ copy progress to ensure that it reaches full DUPLEX state before moving forward. After the pairs are in DUPLEX state, start IR from A to B’ with NOCOPY. Example 4-25 illustrates the TSO command to start the IR recording.

Example 4-25 Start IR A to B’ with NOCOPY

```
CESTPAIR DEVN(X'Adevn') -
          PRIM(X'Assid' Aserial X'Acca' X'Alss') -
          SEC(X'B’ssid' B’serial X'B’cca' X'B’clss')  -
          MODE(NOCOPY) OPTION(SYNC) INCRESYNC(NOCOPY)
```

The configuration is now running A to B' to A' with IR recording turned on A to B' as illustrated in Figure 4-16.

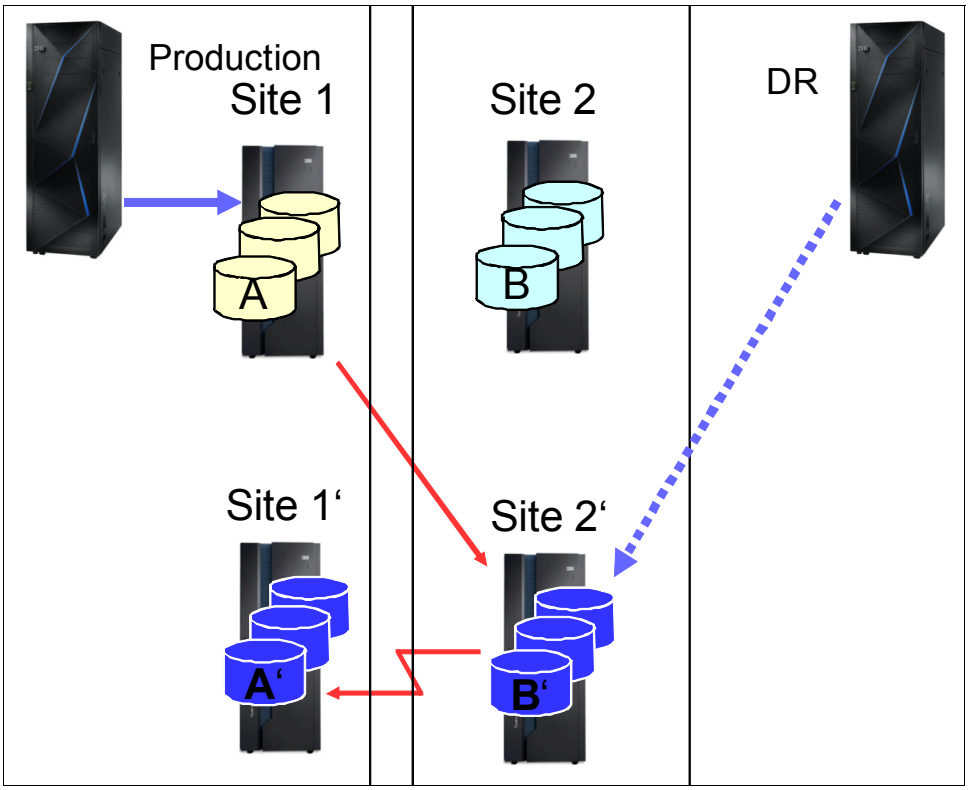

Figure 4-16 MM A to B' Sync with Cascaded A'

Perform data validation on the new B' volume set. If the Site 2 Kp-sys is on the B' volume set, IPL it and validate the configuration. This is your GO/NOGO decision point. If you determine there are errors during the validation process and you need to back out, see "Back-out scenario for B'" on page 67.

Cut over to the new B' volume set

Load the GEOPARM configuration by using GDPS/PPRC 'Configuration Management' panel (option C) to find the new B' volume set. If the configuration between the B and B' volume sets changed, the GEOPARM needs to be updated with the changes before it is loaded.

HyperSwap from A to B' by running the GDPS script statement **`SWITCH HYPERSWAP SUSPEND`** from the Master Kp-sys.

Move the CDSs to SITE 2 by running the GDPS script statement **`CDS SITE2`**.

Cleanup is required for on the B' devices for the A to B' relationships to allow the B' to A' to be converted from ASYNC to SYNC. To accomplish this, issue the Copy Services command recover to B' for the A to B' relationships.

If the Site 1 Kp-sys is on the A volume set being replaced, shut down this system. Wait for the copy to complete A to A', at which point the out-of-sync tracks=0, before moving forward. After the copy is complete, issue a **`delpair`** to the A devices for the A to A' relationship.

Convert the B' to A' ASYNC relationship to a SYNC relationship with Copy Services establish pair command with OPTION(SYNC) MODE(RESYNC).

Remove the A volume set and add the A' volume set into the configuration. Use standard operating procedures that you would normally perform to swap out hardware such as swapping cables or activating an IODF to remove the old volume set and add the new volume set.

Load the GEOPARM configuration by using GDPS/PPRC 'Configuration Management' panel (option C) to find the new A' volume set. If the configuration between the A and A' volume sets changed, the GEOPARM needs to be updated with the changes before it is loaded.

Perform data validation on the new A' volume set. IPL the Site 1 Kp-sys on the A' and validate the configuration. This is your GO/NOGO decision point. If you determine there are errors during the validation process and you need to back out, see "Back-out scenario for A'" on page 68.

Cut over to the new A' volume set

Attention: If a backout is needed required at this stage, a full copy is required from B' to A. This is required because the A to B' relationship had to be removed to establish the B' to A' synchronous relationship.

Ensure that GDPS is HyperSwap enabled and green OK. Run the GDPS monitors. HyperSwap from B' to A' by running the GDPS script statement **`SWITCH HYPERSWAP SUSPEND`** from the Master Kp-sys.

Resync A' to B' by running the GDPS script statement **`START SECONDARY`** from the Master Kp-sys.

Move the CDS primary and secondaries back to their normal configuration with primaries on Site 1 and alternates on Site 2. To accomplish this, run the GDPS script statement **`CDS NORMAL`**. The configuration is now A' to B' as shown in Figure 4-10 on page 56.

Back-out scenario for B'

Load the GEOPARM configuration by using the GDPS/PPRC 'Configuration Management' panel (option C) with the B' volume set. If the configuration between the B and B' volume sets changed, the GEOPARM needs to be updated with the changes before it is loaded.

Start IR from A to B' with NOCOPY. Example 4-25 on page 65 illustrates the TSO command to start the IR recording.

Freeze the A to B' relationship by running GDPS/PPRC script statement **`STOP SECONDARY`** to stop all I/O to the B' volume set. At this point, the A devices are being updated and changes are being recorded in the CR bitmap. The B' devices are not being updated because Metro Mirror (A to B') is frozen.

Establish A to B with IR. Load the original GEOPARM configuration by using the GDPS/PPRC 'Configuration Management' panel (option C) with the B volume set. Run the GDPS monitors.

Move the CDS primary and secondaries back to their normal configuration with primaries on Site 1 and alternates on Site 2. To accomplish this, run the GDPS script statement **CDS NORMAL**. The configuration is back to the original as shown in Figure 4-17.

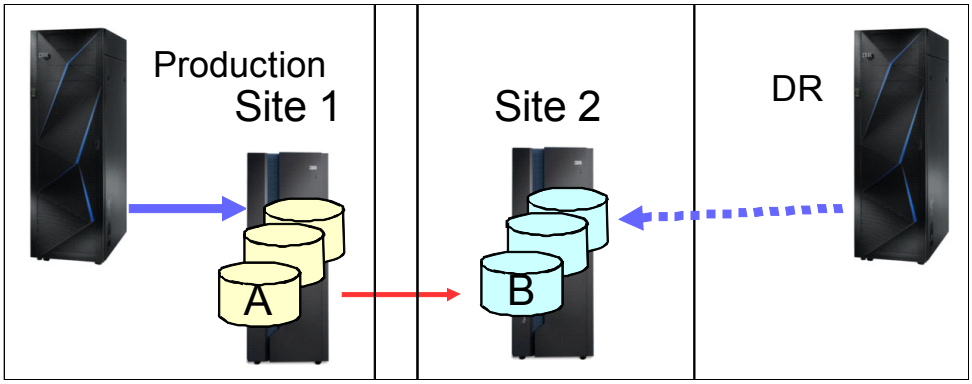

Figure 4-17 Metro Mirror original configuration after back out of B'

Back-out scenario for A'

Freeze the B' to A' relationship by running GDPS/PPRC script statement **STOP SECONDARY** to stop all I/O to the A' volume set.

> **Attention:** A full copy is required from B' to A. This is required because the A to B' relationship had to be removed to establish the B' to A' synchronous relationship.

Establish the paths and the pairs for the B' to A relationships. Careful planning is required to ensure that ample time is allotted for this copy to complete.

After the copy from B' to A is complete, load the GEOPARM configuration by using the GDPS/PPRC 'Configuration Management' panel (option C) with the A and B' volume sets. Run the GDPS monitors.

Ensure that GDPS is HyperSwap enabled and green OK. Run the GDPS monitors. HyperSwap from B' to A by running GDPS script statement **SWITCH HYPERSWAP SUSPEND** from the Master Kp-sys.

Resync A to B' by running GDPS script statement **START SECONDARY** from the Master Kp-sys.

Move the CDS primary and secondaries back to their normal configuration with primaries on Site 1 and alternates on Site 2. To accomplish this, run the DPS script statement **CDS NORMAL**. The configuration is now A to B' as illustrated in Figure 4-18.

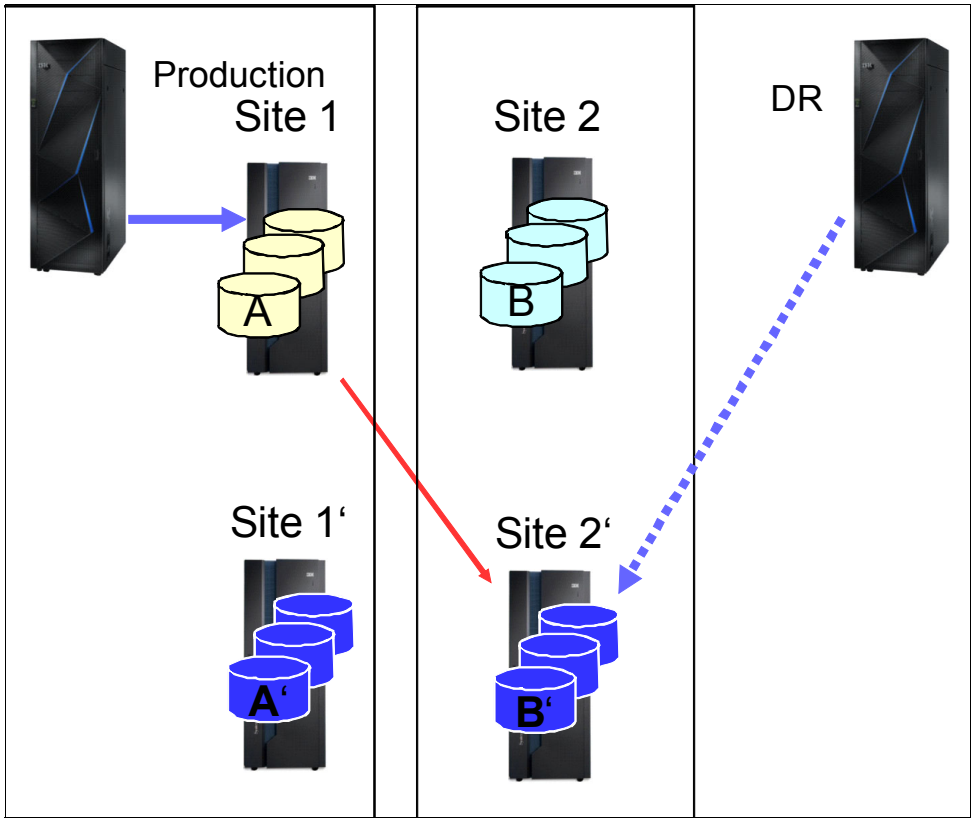

Figure 4-18 Metro Mirror configuration with A' backed out

5

Data Migration with 3-site Metro Mirror replication

This chapter describes various data migration scenarios for installations where an existing Multiple Target Peer-to-Peer Remote Copy (MT-PPRC) configuration is in place. It explains the sequence of tasks and provides examples on how to replace either the secondary or the primary volume set. The chapter also includes special considerations when the HyperSwap capability is installed.

MT-PPRC is a new feature of the IBM DS8870 storage system and some of the terms and procedures that are used in this chapter might sound unfamiliar. To learn more about MT-PPRC, see *DS8870 Multiple Target PPRC*, REDP5151.

This chapter includes the following sections:

- Initial configuration
- Migration process

© Copyright IBM Corp. 2015. All rights reserved.

5.1 Initial configuration

Figure 5-1 illustrates the basic layout of a 3-site Metro Mirror configuration. The host system accesses the primary volume set A in the site 1 storage systems. There is a synchronous replication relation to each of the two secondary volume sets, B in site 2 and C in site 3.

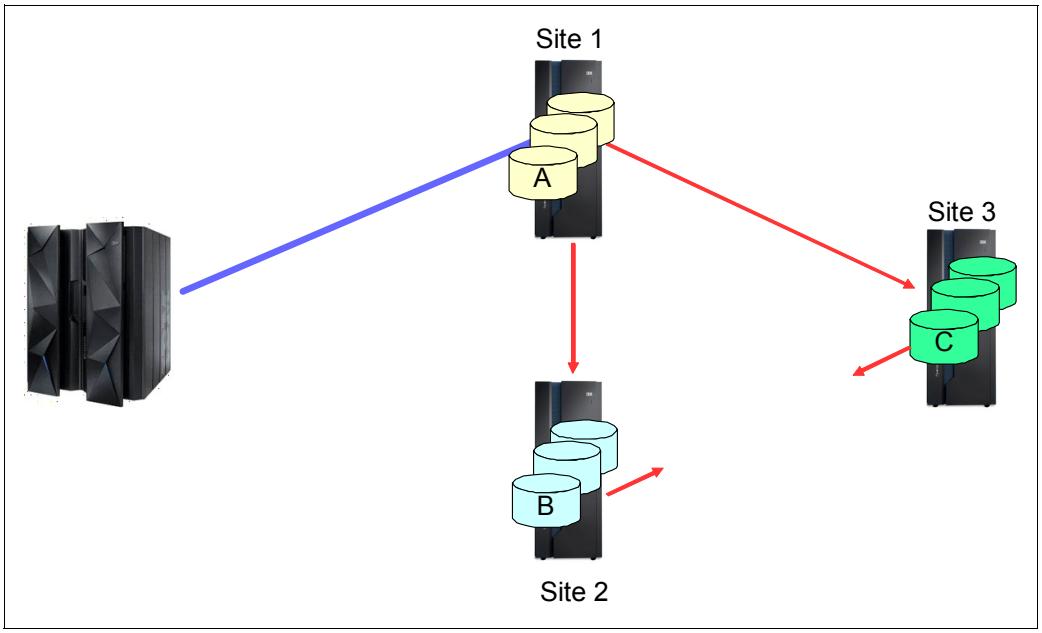

Figure 5-1 Basic 3-site Metro Mirror configuration

The copy services relation between the site-2 and the site-3 volume sets that is indicated by the two facing arrows does not exist during normal 3-site operations. However, it can be established quickly without the need for a full copy if one or both of the regular relations are suspended.

To keep the description of the migration procedures simple, assume that all host I/O operations happen on the primary storage systems, and all auxiliary storage systems IOs result from a replication. There are additional considerations in a configuration where the auxiliary storage systems are used as primaries for other host systems or applications.

5.2 Migration process

This section discusses various scenarios for the migration process.

5.2.1 Replacing site 2 or site 3 volume set

In a three-site Metro Mirror configuration, the two secondary volume sets can be replaced without impact to host I/O operations. The configuration is symmetrical from a replication point of view. Therefore, this section only describes the replacement of the site 2 volume set B. The replacement of the volume set C in site 3 follows the same process.

There are two general methods to replace one of the secondary volume sets in this configuration:

1. If it is sufficient to have only one consistent and synchronous replication during the migration phase, you can simply stop the replication to the volume set that you want to replace, remove it, provide the new volumes, and start the replication to the new volumes.
2. If it is required to retain the two synchronous replications for as long as possible, you must create a fourth copy of the data using the cascading feature of Metro Mirror. After this process is complete, replace the old volume set with the new copy.

Both procedures are described in more detail in the following sections.

One replication relation is sufficient during migration

If you can do without the third copy of your data during the migration, you can follow this sequence to replace the site 2 volume set:

1. Remove the Metro Mirror relationships between the site 1 and the site 2 volumes.
2. Remove the original site 2 volume set.
3. Provide the new site 2 volume set.
4. Establish the Metro Mirror relationships between the site 1 and the site 2 volumes.

After the new Metro Mirror relations are synchronized, you are back in a full three-site configuration. The operations mentioned are all straight forward Copy Services operations with no migration-specific considerations.

Both replication relations are required during migration

If both replications must remain available and synchronized while the migration takes place, you must take a different approach. See Figure 5-1 on page 72 to visualize the initial configuration.

Set up replication

To prepare the new volume set for site 2, set up a cascaded Global Copy relation that replicates all the data from the site 2 volume set B to a new one called B'. This is illustrated in Figure 5-2.

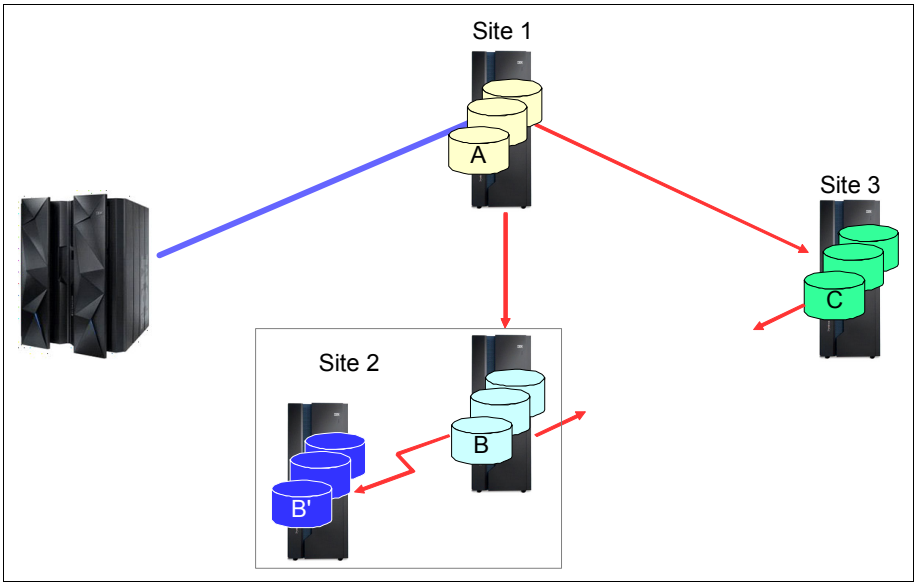

Figure 5-2 Replicate data using Global Copy cascade.

Global Copy replicates the data in a non-synchronous fashion, normally without affecting the I/O operations to its primaries. Therefore, normal operations in site 1, site 2, and site 3 continue without impact. Global Copy will never reach a synchronized state. After an initial copy, almost all of the data is replicated, but there will always be some that is out of date while changes occur on site 2. For more information about cascading Global Copy and example commands to set up and monitor the replication, see 2.3, “Data replication” on page 22.

Note: A major advantage to a “traditional” (cascaded) 3-site configuration is that you do not have to set up the Global Copy replication to fill the new volume set from the “end” of the existing configuration. Therefore, the replication happens locally, which is more cost effective (no additional network needed) and faster (local Fibre Channel connection).

Validation testing

After the replication to the new volume set B’ is complete, you have the opportunity to run a data validation test, as described in 2.4, “Data validation testing” on page 24. Make the data in the volume set B’ consistent using one of these methods:

- Stop host I/O operations to the A volume set and wait until the Out-Of-Sync (OOS) tracks between B to B’ are zero. Pause Global Copy B to B’ and resume I/O operations on A. This method creates a system or application consistent copy of the data on B’.
- Perform a freeze and suspend operation for the Metro Mirror relation between the A and B volume sets. This action creates a *crash consistent* copy of the data on B. Wait until the OOS tracks between B and B’ are zero. Then, suspend the Global Copy B to B’ and resume Metro Mirror A to B. Testing with such a copy is like running a Disaster Recovery test, and might require operating system and application recovery.

Cut over

When your tests have completed with the wanted results, you can perform the cut-over sequence. It consists of the following steps:

1. Prepare for *incremental resync* (IR) on the primary volume set A.
2. Suspend the Metro Mirror replication from A to B.
3. Wait until the OOS tracks for Global Copy B to B’ are zero.
4. Suspend Global Copy B to B’.
5. Fail over Global Copy.
6. Start mirroring A to B’ with IR.
7. Monitor the OOS tracks until A to B’ is synchronized.

There are no specific steps required to ensure consistency because there is a second synchronous copy available throughout the whole process.

This section describes the steps in more detail and with examples. To start the incremental resync feature for the primary volume set (A), run a **mkpprc** command with the **-incrementalresync enablenoinit** and **-nocopy** options on the storage systems that contain this volume set, as shown in Example 5-1. Specify the storage systems that contain the volume set B in the **-remotedev** parameter.

Example 5-1 DSCLI command to start IR on volume set A

```
dscli> mkpprc -type mmir -remotedev IBM.2107-75CYM31 -incrementalresync enablenoinit -mode
nocp 3800-3803:8000-8003
CMUC00153I mkpprc: Remote Mirror and Copy ... relationship 3800:8000 successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy ... relationship 3801:8001 successfully created.
...
```

This **mkpprc** command does not initiate a replication. It turns on change recording for the existing relations A to B. From now on, all updates that are made to volume set A are recorded in the IR bitmaps. Perform this step as close to cut over time as possible. The longer you wait for the final cut over, the more data must be copied to A' later. You can repeat the command if there is an unexpected delay.

When final cut over time arrives, suspend the Metro Mirror relations A to B, as shown in Example 5-2. Again, run this command on the storage systems that contain the volume set A.

Example 5-2 DSCLI command to suspend Metro Mirror A to B

```
dscli> pausepprc -remotedev IBM.2107-75CYM31 3800-3803:8000-8003
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 3800:8000 ... successfully paused.
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 3801:8001 ... successfully paused.
...
```

If you now check the Metro Mirror relations on the primary side with the **lspprc -l** command, you see that OOS tracks for the A to B relations are increasing, whereas A to C remains in sync.

Monitor the Global Copy that replicates data from B to B' until the OOS tracks for all relations are zero. Use the **lspprc -l** command, issued on the storage systems that hold volume set B. As a result, B and B' are identical. Now, suspend the Global Copy B to B' by running the **pausepprc** command on the storage systems containing volume set B, as shown in Example 5-3.

Example 5-3 DSCLI command to suspend Global Copy B to B'

```
dscli> pausepprc -remotedev IBM.2107-75CYK71 8000-8003:8800-8803
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 8000:8800 ... successfully paused.
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 8001:8801 ... successfully paused.
...
```

The failover operation on the B' volumes is required to allow them to enter a new relation. Issue the **pprcfailover** command on the storage systems that hold the B' volumes, as shown in Example 5-4. The volume pairs must be entered in the opposite direction (B':B) and the storage systems holding the B volumes are specified in the **-remotedev** parameter.

Example 5-4 DSCLI command to fail over Global Copy B' to B

```
dscli> failoverpprc -type gcp -remotedev IBM.2107-75CYM31 8800-8803:8000-8003
CMUC00196I failoverpprc: Remote Mirror and Copy pair 8800:8000 successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 8801:8001 successfully reversed.
...
```

The final step in the cut over sequence is to restart the mirroring from A, this time to B', using the IR change recording bitmaps that were created in the beginning. Use the **mkpprc** command with the `-incrementalresync override` option on the storage systems that contain volume set A as illustrated in Example 5-5. The **-remotedev** parameter specifies the B' storage systems.

Example 5-5 DSCLI command to start Metro Mirror A to B' with IR

```
dscli> mkpprc -type mmir -remotedev IBM.2107-75CYK71 -incrementalresync override
3800-3803:8800-8803
CMUC00153I mkpprc: Remote Mirror and Copy volume pair ... 3800:8800 successfully created.
```

```
CMUC00153I mkpprc: Remote Mirror and Copy volume pair ... 3801:8801 successfully created.
...
```

Now you only have to monitor the OOS tracks and status of the A to B' relations. As soon as all relations are in the Full Duplex state, the cut over is complete.

5.2.2 Replacing the site 1 volume set

To replace the primary volume set A in a 3-site Metro Mirror configuration, again create a cascaded Global Copy to replicate the data from one of the secondary volume sets B or C to the new volume set A'. After the replication is complete, you perform the cut over from A to A'.

> **Note:** If you can move your host I/O away from volume set A to one of the secondary sets B or C, using HyperSwap or a similar capability, the procedure to replace A is the same as replacing B or C. For more information, see 5.2.4, "HyperSwap considerations" on page 82.

Which of the two Metro Mirror secondary volume sets B or C you use as origin of the replication to A' does not matter from a functionality point of view. Select the one that makes the replication easier, faster, or more cost effective in your particular situation. The example uses volume set C.

Set up replication

To prepare the new volume set for site 1, set up a cascaded Global Copy relation that replicates all the data from the site 2 volume set C to a new one called A', as shown in Figure 5-3. For more information about cascading Global Copy and example commands to set up and monitor the replication, see 2.3, "Data replication" on page 22.

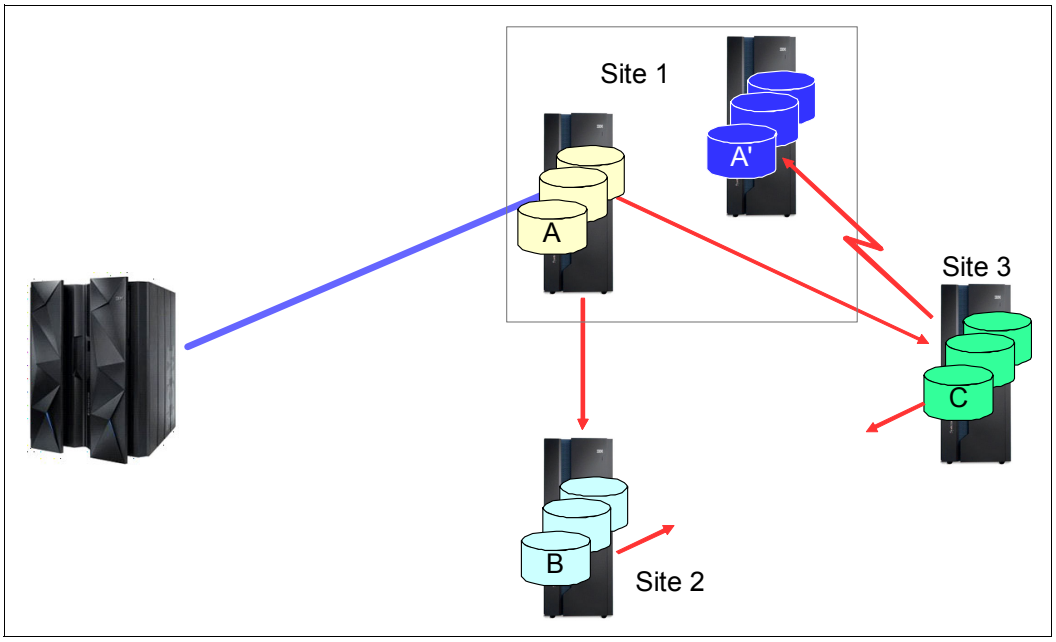

Figure 5-3 Replicate data C to A' using Global Copy cascade.

Validation testing

This scenario replaces the volume set that is used for production I/O operations. A good practice is to run data validation testing as described in 2.4, “Data validation testing” on page 24 before the cut over to the new volume set.

To make the data in the volume set A’ valid and consistent, stop host I/O operations to the A volume set and wait until the OOS tracks between C and A’ are zero. Pause Global Copy C to A’ and resume I/O operations on A. You now have a system or application consistent copy of the data on A’.

> **Note:** If a crash consistent copy of the data is good enough or if you cannot stop production I/O to prepare the data for testing, you can also perform a freeze and suspend for the Metro Mirror A to C. Then, wait until the OOS tracks between C and A’ are zero. Suspend the Global Copy C to A’ and resume Metro Mirror A to C. Testing with such a copy is like performing a Disaster Recovery test, and might require operating system and application recovery.

Preparations for cut-over

The cut-over process consists of two major blocks of actions. First, change the configuration to move the three remaining volume sets (A’, B, C) into a Multi-Target configuration, as shown in Figure 5-4. There is still a Metro Mirror relation from A to B, and two new Global Copy relations from B to C and from B to A’. This allows you later to use Multi-Target PPRC Incremental Resync to move to the final configuration. Production I/O can continue to run during this phase.

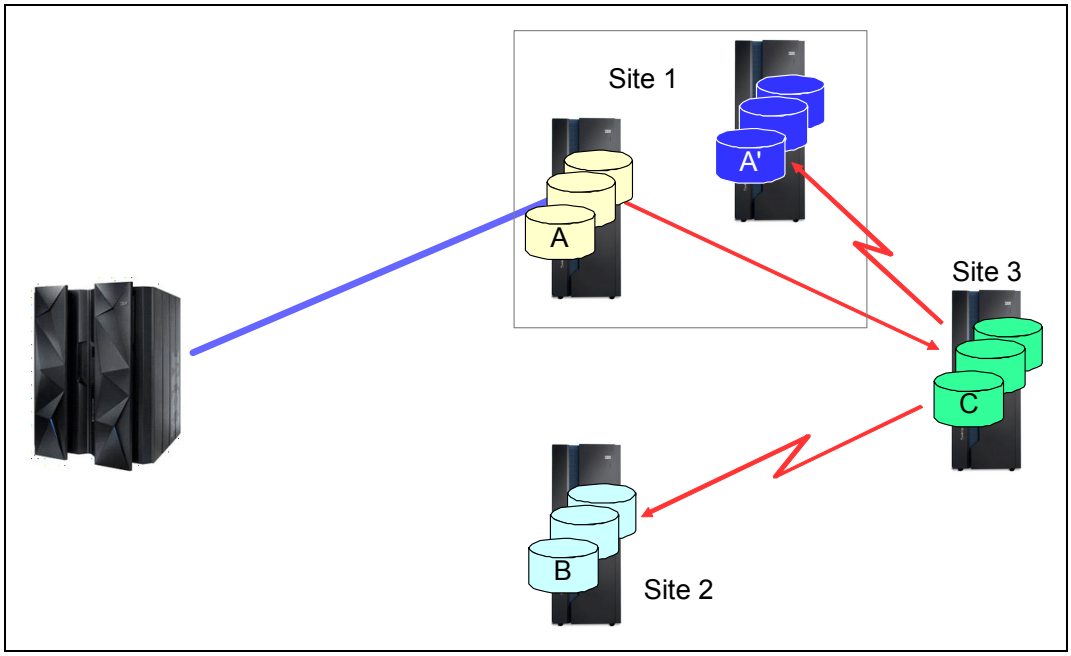

Figure 5-4 Temporary configuration to enable the volume sets A’, B and C for Multi-Target IR.

This first block consists of the following steps:

1. Check that Multiple Target Incremental Resync (MTIR) bitmap toggle runs on B and C.
2. Fail over Metro Mirror B to A.
3. Fail back C to B as Global Copy.
4. Remove Metro Mirror A to B relations.
5. Check that MTIR bitmap toggle runs on A’ and B.

In this temporary configuration, you have only one synchronous copy of your active data. Proceed to the next steps as soon as possible. These steps are described in more detail in 5.2.3, “Example for cut-over sequence” on page 78.

Cut over

With the second block of steps, you perform the final cut-over of the production I/O operations from the A to the A’ volume set. Replacing the volume set that is used for production I/O requires a short outage to the host processing for the cut-over process. Perform all Copy Services reconfiguration using DS8000 incremental resync capabilities. Doing so, the outage is only as long as it takes to do the reconfiguration.

The following steps are required for the final cut-over, based on the configuration that is shown in Figure 5-4 on page 77:

1. Quiesce I/O.
2. Remove the Metro Mirror A to C relations.
3. Wait until OOS tracks or Global Copy C to A' are zero.
4. Fail over A' to C.
5. Fail over B to C.
6. Fail back Metro Mirror A’ to B.
7. Fail back Metro Mirror A’ to C.
8. Resume I/O.

These steps are described in more detail in 5.2.3, “Example for cut-over sequence” on page 78.

5.2.3 Example for cut-over sequence

This section describes the whole cut-over sequence in more detail and provide examples for the DSCLI commands.

Example for cut-over preparation

Starting with the configuration as shown in Figure 5-3 on page 76, transform the setup to the temporary one shown in Figure 5-4 on page 77.

Step 1: Check MTIR bitmap toggle

To make sure that the subsequent steps can all be performed with incremental resync, check whether the bitmap toggle process of the MTIR pairs for the MT relations A to B and A to C are taking place. As shown in Example 5-6, use the **`lspprc`** command with the parameter **`-multtgt`** to display the MTIR relations and their OOS tracks. MTIR bitmap toggle is working if the OOS tracks of the MTIR pairs increase for periods of one minute and then drop to about 50%, start increasing again, and so on. In the **`lspprc`** output, you recognize the MTIR relations with their status `Suspended` with reason `Multi-target Internal`.

Example 5-6 DSCLI command to check MTIR pairs.

```
dscli> lspprc -l -multtgt 8000-8003
ID        State              Reason                Type         Out Of Sync Tracks
===============================================================================
3800:8000 Target Full Duplex -                     Metro Mirror 0
3801:8001 Target Full Duplex -                     Metro Mirror 0
...
8000:A000 Suspended          Multi-target Internal Global Copy  3860
8001:A001 Suspended          Multi-target Internal Global Copy  3859
...
```

Perform this check for both MT PPRC secondary volume sets B and C. If the bitmap toggle is not running, you can force its restart by suspending and restarting the affected Metro Mirror relations.

Step 2: Fail over Metro Mirror B to A

In this step, run a **failoverpprc** command on the storage systems that contain volume set B, against the Metro Mirror relation A to B, as shown in Example 5-7. This command stops replication from A to B and places the volumes in volume set B in a Primary state, and allows you to start mirroring from C to B in the next step.

Example 5-7 DSCLI command to fail over the A to B relations to B

```
dscli> failoverpprc -type mmir -multtgt -remotedev IBM.2107-75ZA571 8000-8003:3800-3803
CMUC00196I failoverpprc: Remote Mirror and Copy pair 8000:3800 successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 8001:3801 successfully reversed.
...
```

You must specify the volume set A storage systems in the **-remotedev** parameter. The volume IDs appear in the order as though the relations are seen from B to A.

Step 3: Fail back C to B as Global Copy

You can now start the replication from C to B, using the **failbackpprc** command on the volume set C storage systems as in Example 5-8.

Example 5-8 DSCLI command to start Global Copy C to B with IR

```
dscli> failbackpprc -type gcp -remotedev IBM.2107-75CYM31 A000-A003:8000-8003
MUC00197I failbackpprc: Remote Mirror and Copy pair A000:8000 successfully failed back.
MUC00197I failbackpprc: Remote Mirror and Copy pair A001:8001 successfully failed back.
...
```

You must start these relations in Global Copy mode (**-type gcp**) because they are cascaded on the A to C relations.

Step 4: Remove Metro Mirror A to B relations

The last configuration step is to remove the suspended Metro Mirror relations between volume sets A and C. You must do this for both volume sets separately, because you did a failover operation for these relations earlier and both volume sets are in the Primary state. Example 5-9 shows the command to issue on the volume set A storage systems.

> **Note:** The commands in the following examples both show error messages. They are related to the fact that the Metro Mirror relations were failed over earlier, and can be ignored. Use the **lspprc** command on both sides (A and B storage systems) to verify that all Metro Mirror relations are successfully removed.

Example 5-9 DSCLI command to remove A to B relations on A

```
dscli> rmpprc -quiet -remotedev IBM.2107-75CYM31 3800-3803:8000-8003
CMUN03012E rmpprc: 3800:8000: An error occurred for a Remote Mirror and ...
CMUN03012E rmpprc: 3801:8001: An error occurred for a Remote Mirror and ...
...
```

Example 5-10 shows the one for the B storage systems. You must specify the A storage systems in the **-remotedev** parameter. The volume IDs appear in the order as though the relations are seen from C to A.

Example 5-10 DSCLI command to remove A to B relations on B

```
dscli> rmpprc -quiet  -remotedev IBM.2107-75ZA571 8000-8003:3800-3803
CMUN03012E rmpprc: 8000:3800: An error occurred for a Remote Mirror and ...
CMUN03012E rmpprc: 8001:3801: An error occurred for a Remote Mirror and ...
...
```

Step 5: Check MT PPRC internal pairs bitmap toggle

Now you are running the temporary configuration as seen in Figure 5-4 on page 77. To be on the safe side, check the MTIR pairs again, as in step 1, but this time on the A' and B storage systems.

Note: Now you have the opportunity to run another data validation test before continuing with the final cut-over steps.

Example for cut over

In the following example steps, you perform the final cut-over to the new site 1 volume set A'. This phase requires an application downtime and reconfiguration of the host access from volume set A to A'.

Step 1: Quiesce I/O

You must take any required measures to stop all host I/O operations to the old volume set A. Do not configure the host to access A' volumes yet. Ensure that there will be no further I/O attempts to both the old (A) and new (A') volume set.

Note: If something goes wrong during the steps that follow, you can always return to volume set A for production I/O operations, up until step 8 (Resume I/O). If you do not start production I/O on the new volume set A', the old volume set A contains the most recent production data.

Note: The safest way to avoid unwanted I/O activity is to shut down all hosts that access volume set A.

Step 2: Remove Metro Mirror A to C relations

Remove volume set A from the configuration by removing the Metro Mirror relations from A to C. Currently, the relations are still intact and it is sufficient to run the **rmpprc** command on the volume set A storage systems, as shown in Example 5-11.

Example 5-11 DSCLI command to remove the Metro Mirror relations A to C

```
dscli> rmpprc -quiet -remotedev IBM.2107-75CZM21 3800-3803:A000-A003
CMUC00155I rmpprc: Remote Mirror and ... 3800:A000 relationship successfully withdrawn.
CMUC00155I rmpprc: Remote Mirror and ... 3801:A001 relationship successfully withdrawn.
...
```

Step 3: Wait until OOS tracks or Global Copy C to A' are zero

Because you stopped host I/Os, the DS8000 systems in the configuration are draining OOS tracks of the running Global Copy relations (C to A' and C to B). Check the progress using the **`lspprc -l`** command on the volume set C storage systems. After the OOS track count is zero for all relations, all three volume sets (A', B, and C) are identical.

In the next step, reverse the mirroring direction of the C to A' relations and switch the C to B relations to A' to B to achieve the final configuration.

Step 4: Fail over A' to C

To reverse the mirroring direction of the C to A' Global Copy relations, perform a **`failoverpprc`** operation against the Global Copy secondaries. Run this command, as shown in Example 5-12, on the storage systems that hold volume set A'.

Example 5-12 DSCLI command to fail over the C to A' Global Copy relations.

```
dscli> failoverpprc -type gcp -remotedev IBM.2107-75CZM21 8800-8803:A000-A003
CMUC00196I failoverpprc: Remote Mirror and Copy pair 8800:A000 successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 8801:A001 successfully reversed.
...
```

Note: After this operation, the A' volumes are primaries and host accessible. You can resume your host I/O at this time (to A'), if necessary. However, there is no synchronous and consistent replication yet.

Step 5: Fail over B to C

To prepare the C to B Global Copy relations for the switch, perform a **`failoverpprc`** operation against the secondaries. Run this command, as shown in Example 5-13, on the storage systems that hold volume set B.

Example 5-13 DSCLI command to fail over the C to B Global Copy relations.

```
dscli> failoverpprc -type mmir -multtgt -remotedev IBM.2107-75CZM21  8000-8003:A000-A003
CMUC00196I failoverpprc: Remote Mirror and Copy pair 8000:A000 successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 8001:A001 successfully reversed.
...
```

Now you are ready to restart mirroring such that A' becomes the primary, and B and C the two secondary volume sets.

Step 6: Fail back Metro Mirror A' to B

Start mirroring for the A' to B relations by issuing a **`failbackpprc`** command on the storage systems that hold volume set A', as shown in Example 5-14. Use the **`-type mmir`** option for synchronous mirroring (Metro Mirror).

Example 5-14 DSCLI command to fail back Metro Mirror A' to B

```
dscli> failbackpprc -type mmir -remotedev IBM.2107-75CZM21 8800-8803:A000-A003
CMUC00197I failbackpprc: Remote Mirror and Copy pair 8800:A000 successfully failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 8801:A001 successfully failed back.
...
```

Step 7: Fail back Metro Mirror A' to C

Start the second set of Metro Mirror relations (A' to C) again with a **`failbackpprc`** command, also on the storage systems that hold volume set A', but with C storage systems as the remote devices. This is shown in Example 5-15.

Example 5-15 DSCLI command to fail back Metro Mirror A' to C

```
dscli> failbackpprc -type mmir -remotedev IBM.2107-75CYM31 8800-8803:8000-8003
CMUC00197I failbackpprc: Remote Mirror and Copy pair 8800:8000 successfully failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 8801:8001 successfully failed back.
...
```

Step 8: Resume I/O

The reconfiguration of your Multiple Target PPRC setup is no complete. You have the new volume set A' as MT PPRC primary, mirroring to the two secondary volume sets B and C. You can now configure your host systems to access volume set A' and resume your production workload.

5.2.4 HyperSwap considerations

In a 3-site Metro Mirror configuration with HyperSwap, you have two choices for the HyperSwap. You can either swap to the site 2 or the site 3 volume set. You can use this capability to replace the primary volume set without impact to host I/O operations while still maintaining a synchronous replication with HyperSwap capability almost all the time.

Figure 5-5 shows the initial situation of a 3-site Metro Mirror configuration with HyperSwap. The red circles with arrowheads depict the HyperSwap capable relations. In this configuration, you can perform a HyperSwap from site 1 to site 2 or from site 1 to site 3.

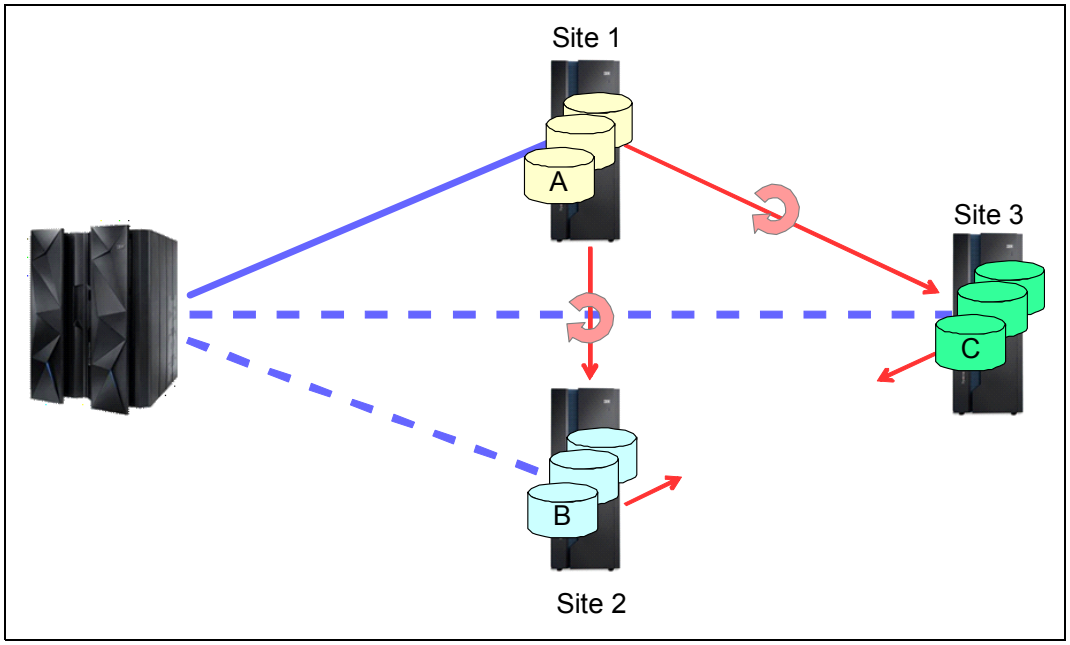

Figure 5-5 3-site Metro Mirror configuration with HyperSwap

Using HyperSwap, you never have to perform data migration for a production I/O volume set. If you must replace volume set A in this example, simply swap your workload to one of the secondary volume sets (B or C), and restart the Multiple Target relations from there, as shown in Figure 5-6. Thus A becomes a Metro Mirror secondary and can be migrated using one of the methods described in 5.2.1, “Replacing site 2 or site 3 volume set” on page 72.

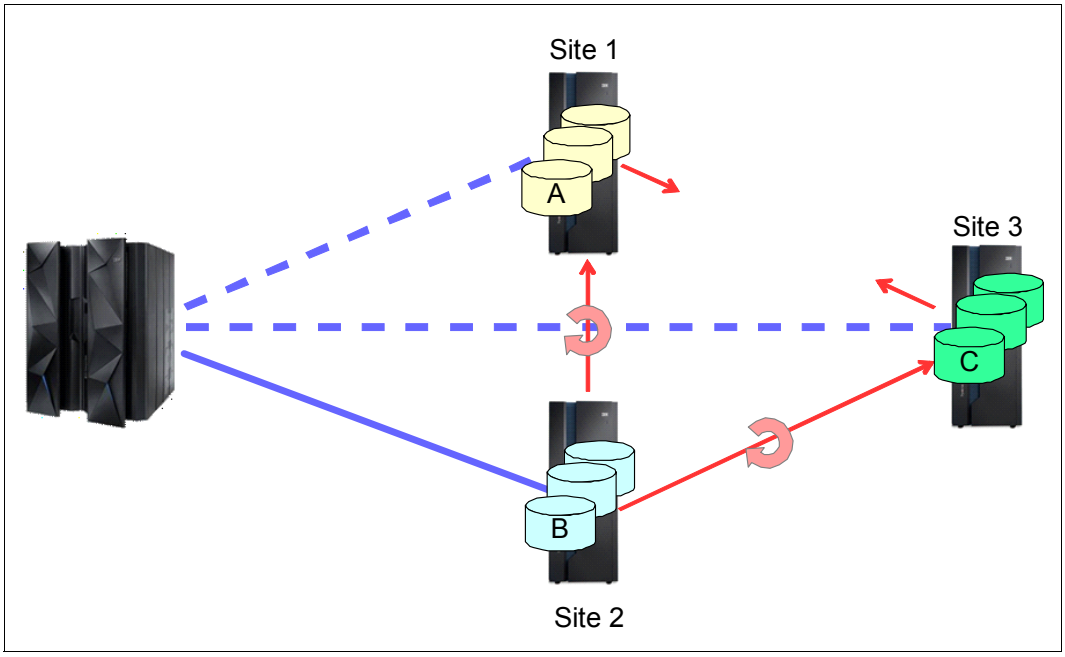

Figure 5-6 Migrate A as secondary volume set after HyperSwap

6

Data Migration with 2-site Global Mirror replication

This chapter provides a way how to migrate data between DS8000 systems using the DS8000 Series Global Mirror function.

The following scenarios are discussed:

- Replacing auxiliary storage
- Replacing primary storage
- Replacing primary and auxiliary storage
- GDPS considerations

This chapter includes the following sections:

- Initial configuration
- Migration process
- GDPS considerations

© Copyright IBM Corp. 2015. All rights reserved.

6.1 Initial configuration

The initial configuration of the 2-site Global Mirror, as shown in Figure 6-1, consists like a 2-site Metro Mirror of two DS8000 storage systems. The Global Mirror is composed by a Global Copy relationship between the primary and the auxiliary storage system. In addition, a FlashCopy relation from the Global Copy target volumes to FlashCopy target volumes in the secondary box. The FlashCopy relation holds the most recent consistency group that was formed by the Global Mirror process at the primary side.

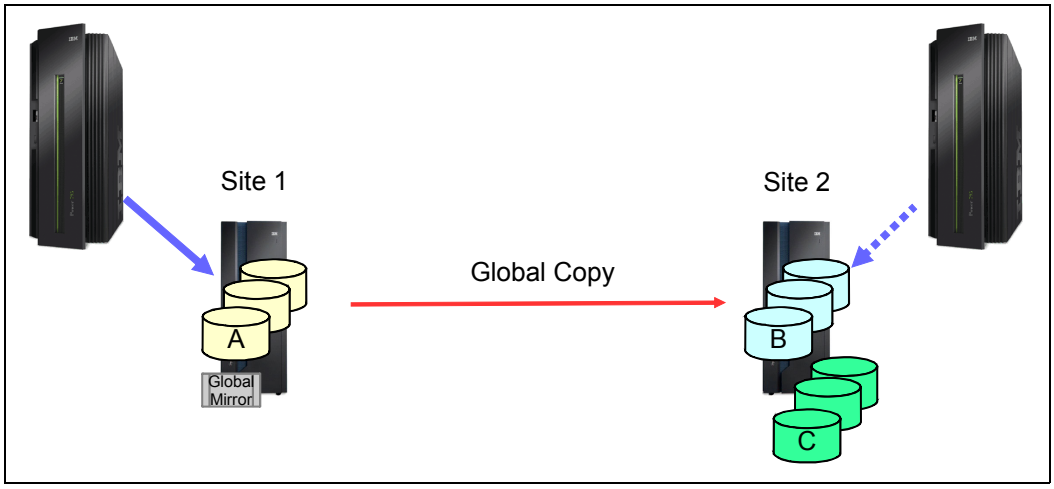

Figure 6-1 Initial 2-site Global Mirror configuration

6.2 Migration process

The migration process in a 2-site Global Mirror is essentially the same as in a 2-site Metro Mirror configuration. The following description refers to the migration scenarios described in Chapter 4, “Data Migration with 2-site Metro Mirror replication” on page 39.

6.2.1 General considerations

The transmission of the data to the new storage in a 2-site Global Mirror is orchestrated with an addition cascade Global Copy relation from the secondary site. However, in a Global Mirror the data is cascaded from the B-Volumes, which are the target volumes of the productive Global Copy relation. The C-Volumes cannot be used because they are exclusively under the control of the Global Mirror.

6.2.2 Replacing auxiliary storage

This section discusses how to replace the auxiliary storage system in a Global Mirror configuration

Initial copy phase

The initial copy phase for Global Mirror is the same as in a Metro Mirror. A cascaded Global Copy relation is established from the B volumes to the new storage system B’ volumes. For more information about the DSCLI commands that apply, see “Initial copy phase” on page 41.

Remember that the cascaded Global Copy relation is established to the B volumes and not the C volumes (Figure 6-2).

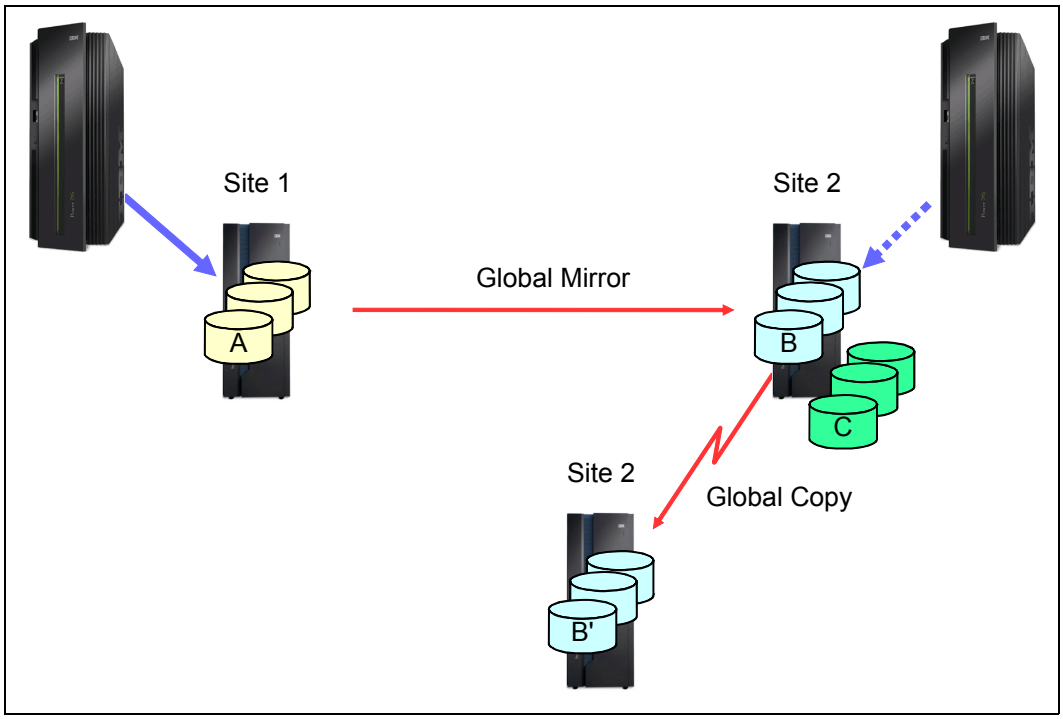

Figure 6-2 Establish a cascaded Global Copy from B to B'

The following are the steps for setting up the initial copy:

1. Establish the paths from the auxiliary storage being replaced to the new one.
2. Establish a cascaded Global Copy relation between the B and B' volumes.

The relevant DSCLI commands are shown in Example 4-1 on page 42 to Example 4-2 on page 42 and following. During this phase, the Global Mirror continues with normal operation. The initial copy phase is completed when all out-of-sync tracks have been copied to the B' volumes. You can monitor this process using the command **`lspprc -l`**, as shown in Example 4-2 on page 42.

Validation phase

In the validation phase, the data is checked using a validation host system. The purpose of this task is to ensure that all data has been copied completely. When replacing the auxiliary storage, the validation system is typically represented by the host systems at the disaster recovery site.

It is essential for the validation that the data on B' volumes be in a consistent state. It is not necessarily required for the most recent data to be available. The advantage with this approach is that the production at the primary site does not have to be stopped during validation.

Figure 6-3 shows the validation configuration.

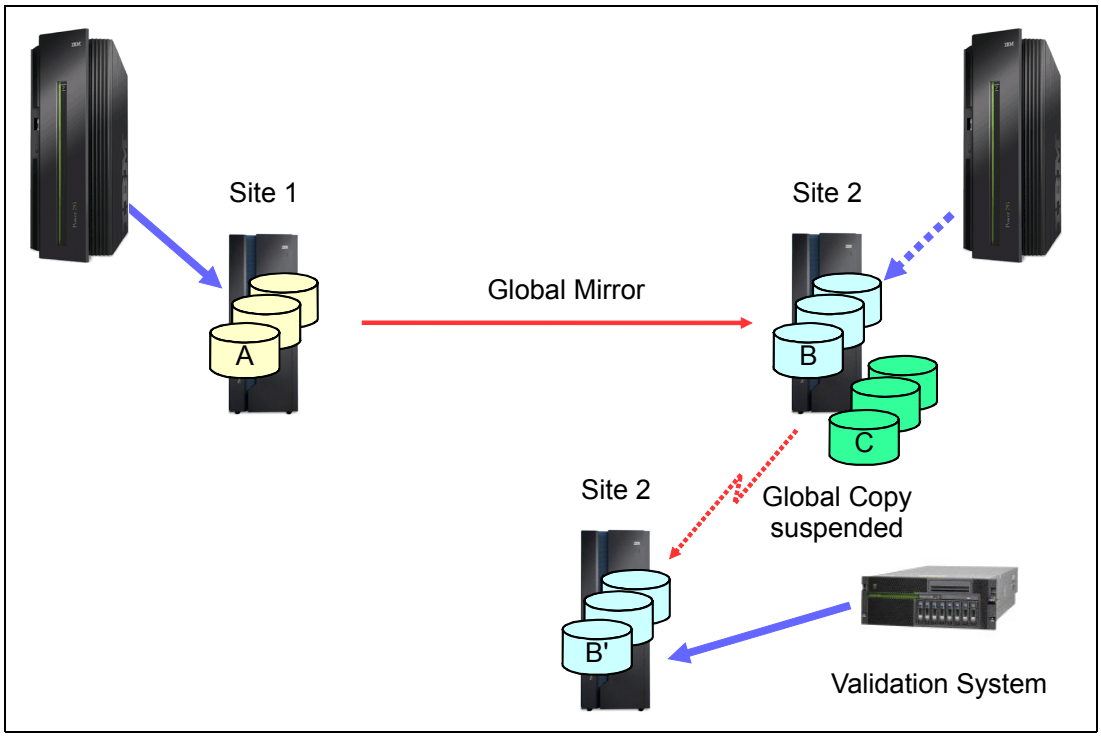

Figure 6-3 Validation of the Global Mirror migration

Data consistency can be provided by using the Global Mirror itself. A Global Mirror failover procedure provides consistency to the B volumes. In this state, the Global Copy relation from site 1 to site 2 is in the failed over state, and there will be no writes against the B volumes. All remaining tracks will drain to the B' volumes. Assuming that the Global Mirror failover is performed by using Tivoli Storage Productivity Center for Replication, the remaining steps for the validation are described from step 4 on page 44 through step 8 on page 45.

In summary, the validation phase comprises the following steps:

1. Failing over Global Mirror to H2 using Tivoli Storage Productivity Center for Replication
2. Wait until all tracks are copied from B to B'
3. Pause the Global Copy between volume set B and B' to allow a graceful failover
4. Restart Global Mirror H1 to H2 to continue normal production operation
5. Fail over cascaded Global Copy to B' volumes
6. Start validation test at B'
7. When test is done, stop validation systems and release the B' volumes
8. Fail back cascaded Global Copy B to B' to prepare the migration phase

When no Tivoli Storage Productivity Center for Replication is available, stop the primary production at step 1 to get consistent data. In this case, the whole scenario as described in "Validation phase" on page 42 must be run.

Note: In this scenario, the consistency was provided by the Global Mirror, which means that the data might not have most recent state.

If the validation test fails, the errors must be corrected. No special back-out scenario is required.

Possible errors are as follows:

- The failover of the Global Mirror in step 1 on page 88 did not apply the correct volume set. The result is inconsistent data at B' and the validation system not able to start.
- The failover to the cascaded Global Copy to B' did not apply to the correct volumes set, with the same effect as above.
- Not all tracks have been copied to the B' before the failover in step 5 on page 88 took place. Make sure, using the `lspprc -l` command, that the no data are out-of-sync. Repeat the command several times as needed.

When all errors have been corrected, repeat the validation phase. Only when the validation scenario has passed completely without any errors, are you prepared for the migration itself.

Migration phase

When the validation phase has passed successfully, migration is ready for the cut over to the new storage system.

The main target in this phase is to switch over the replication from the primary site to the new storage system. This is achieved by using the incremental resynchronization function as shown in Figure 6-4. When incremental resynchronization is enabled at the primary storage system, all new writes to A volumes are tracked in a change recording bitmap.

Note: Make sure that the C' volumes and eventually the D' volumes (which are the Global Mirror practice volumes) are created.

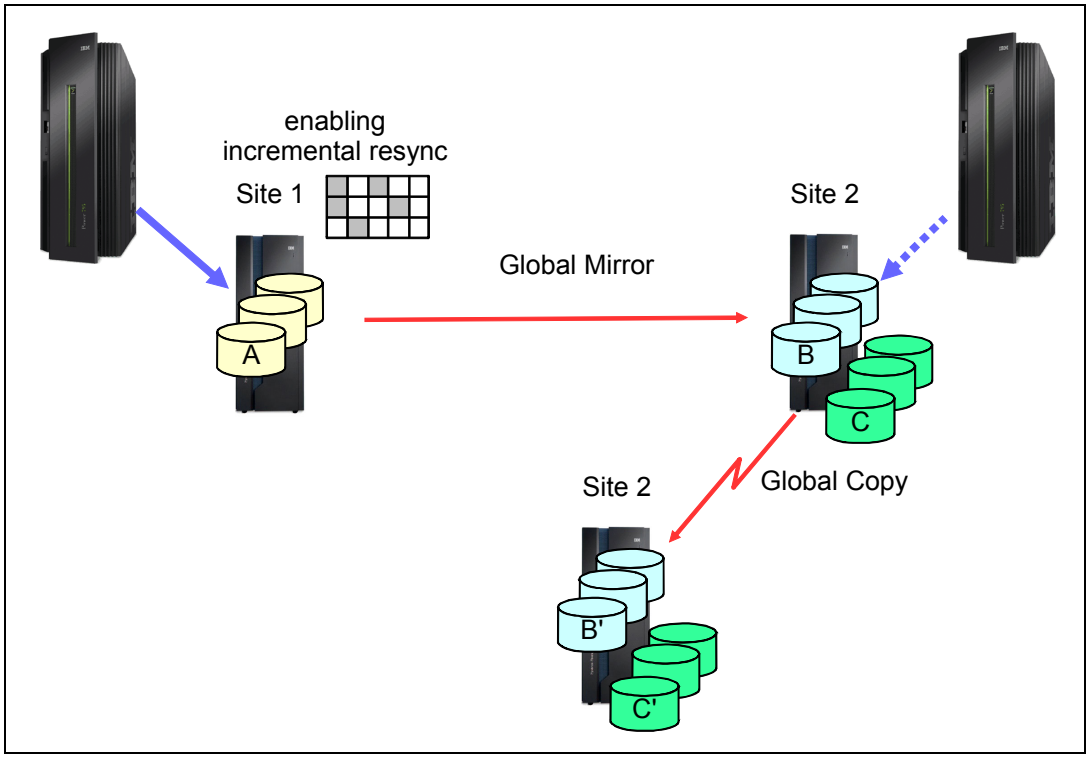

Figure 6-4 Using incremental resynchronization between Site 1 and Site 2

The cut-over scenario to the new storage system consists of the following steps:

1. Enable incremental resynchronization at A volumes.
2. Suspend the Global Mirror between site 1 and site 2.
3. Wait all track been copied to the B' volumes.
4. Fail over cascaded Global Copy to B' volumes.
5. Create a new Global Copy relation between A and B' using incremental resynchronization.
6. Remove the original Global Mirror copy set between A and B in Tivoli Storage Productivity Center for Replication.
7. Starting the new Global Mirror session between A and B'.

After the incremental resynchronization has been enabled, all writes are tracked in the change recording bitmap. To keep the number of tracks that are copied incrementally in step 5 as low as possible, run this scenario within a reasonable time frame. For more information about how to enable the incremental resynchronization at the primary site, see Example 4-9 on page 47.

When Tivoli Storage Productivity Center for Replication is managing the current Global Mirror, steps 2, 6, and 7 can be done using it. In step 2, the suspend of the Global Mirror session is presented in the Tivoli Storage Productivity Center for Replication GUI.

When the new Global Copy relation has been established, the old Global Mirror copy set must be removed from the Tivoli Storage Productivity Center for Replication session. This must be done with the option to keep the base hardware relationships on the storage system because you have already changed the copy direction from the A volumes to the new B' volumes.

Finally, when the old Global Mirror copy set has been removed from the old session the new copy set A to B' to C' must be added to the new Global Mirror session from A to B'. When this process is complete, the new Global Mirror session can be started from H1 to H2.

The DSCLI commands for the steps 3 to 5 are as shown in the Metro Mirror section in Example 4-11 on page 48 to Example 4-13 on page 49 and following.

Post migration phase

When replacing the auxiliary storage system, the post migration consists of cleaning up the remaining parts in the old DS8000 storage system. Complete the following steps:

1. Remove the FlashCopy relations between B and C.
2. Remove the Global Copy target relation of the original Global Copy between A and B at B.
3. Remove the cascaded Global Copy source relation between B and B' at B.
4. Remove all path definitions.
5. Delete the B volumes and the C volumes.
6. Remove obsolete zoning definitions in the storage area network.

6.2.3 Replacing primary storage

This section describes the replacement of the Global Mirror primary storage system.

Initial copy phase

The initial copy phase for a replacement of the primary storage system is following the same steps as in 6.2.2, “Replacing auxiliary storage” on page 86. The only difference here is that the new storage is at the primary site as shown in Figure 6-5.

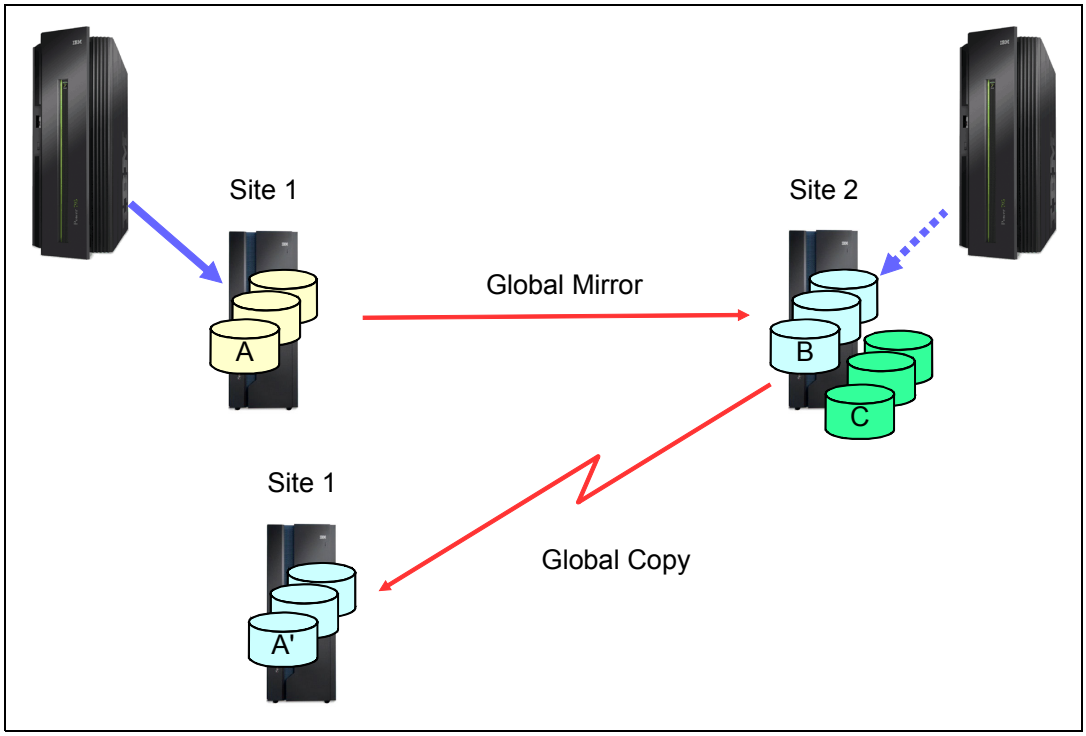

Figure 6-5 Initial copy setup for migrating to new primary storage

To set up the initial copy, complete these steps:

1. Establish the paths from the auxiliary storage being replaced to the new one.
2. Establish a cascaded Global Copy relation between the B and A’ volumes.

The corresponding DSCLI commands are shown in Example 4-1 on page 42 to Example 4-2 on page 42 and following.

The question in this configuration is whether sufficient bandwidth can be provided for the migration. For most data replication environments, the bandwidth is provided in full duplex. This configuration means that when the replication from the primary site to the secondary site is using a certain bandwidth, the same bandwidth is also available in the opposite direction. However, verify whether the link from the secondary site to the primary site is used by other resources, and if yes, that it can (temporarily) be made available for the migration exclusively.

Validation phase

During the validation phase, the data is tested on the new storage system at the primary site. Usually there is no additional server hardware available, and the validation must be conducted with the available primary servers. This is a planned outage of the production during which the server must be attached to the new storage.

For this reason, the cascade Global Copy must be failed over to the A' volumes, so that these are now accessible for the servers. The servers are then stopped and access to the new storage is provided by changing the zoning of the storage area network and the appropriate host connection at the new DS8870 storage system. Finally, the servers are started and the tests against the new storage can begin. See Figure 6-6.

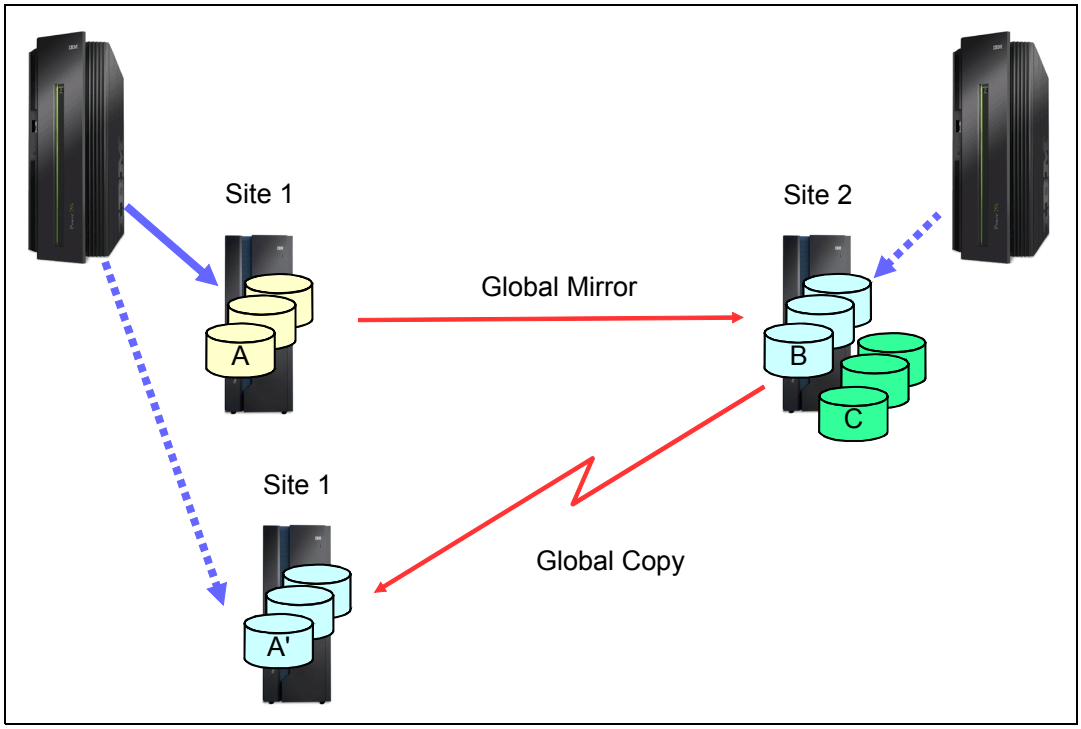

Figure 6-6 Validating the data migration at the primary site

At this point, there are two possible approaches on how to continue with the migration after the validation is successful:

- Continue with production on the new storage

  In this approach, the replication must be turned around and the new Global Mirror session started. The advantage of this approach is that there is no further disruption of the production operations. However, if errors occur during the tests, the outage of the production lasts until all problems are solved.

- Moving back to the original configuration

  When going back to the original configuration, you can plan plenty of time to examine the test results and fix potential problems uncovered during the test. This approach gives more security in planning the whole migration. However, it also requires a second outage of the production during the cut-over phase.

The following example assume that the production will go back to the original configuration when the validation tests have been completed.

The validation at the primary site is done by completing the following tasks:

1. Shut down the production.
2. Wait until all tracks are copied from B to A'.
3. Pause the Global Copy between volume set B and A' to allow a graceful failover.
4. Fail over cascaded Global Copy to A' volumes.
5. Change the zoning to allow host access to A'.

6. Start validation test at A'.
7. When the tests are complete, stop the validation systems and release the A' volumes.
8. Change zoning back to original configuration and start production at A.
9. Fail back cascaded Global Copy B to A' to prepare the migration phase.

Corresponding DSCLI commands are show in Example 4-1 on page 42 to Example 4-8 on page 46 and following.

If the validation fails, the errors must be corrected. There are no special back-out scenarios required. Possible errors can include:

- The zoning to the new storage system is not correct.
- The failover to the cascaded Global Copy to A' did not apply to the correct volumes set.
- Not all tracks have been copied to the A' before the failover in step 5 on page 92 is applied. Make sure with multiple **`lsprc -l`** queries that there are no more out-of-sync tracks.

To repeat the validation phase after all errors have been corrected might be a problem because it involves an additional production outage. However, if the criticality of the migration justifies another downtime, running another validation is by far the safest option.

Migration phase

In the migration phase, the production server is attached to the new primary storage systems (refer to Figure 6-7). The Global Copy replication direction must be turned around from A' to B, and this new volume set must be added to the new Global Mirror session in Tivoli Storage Productivity Center for Replication.

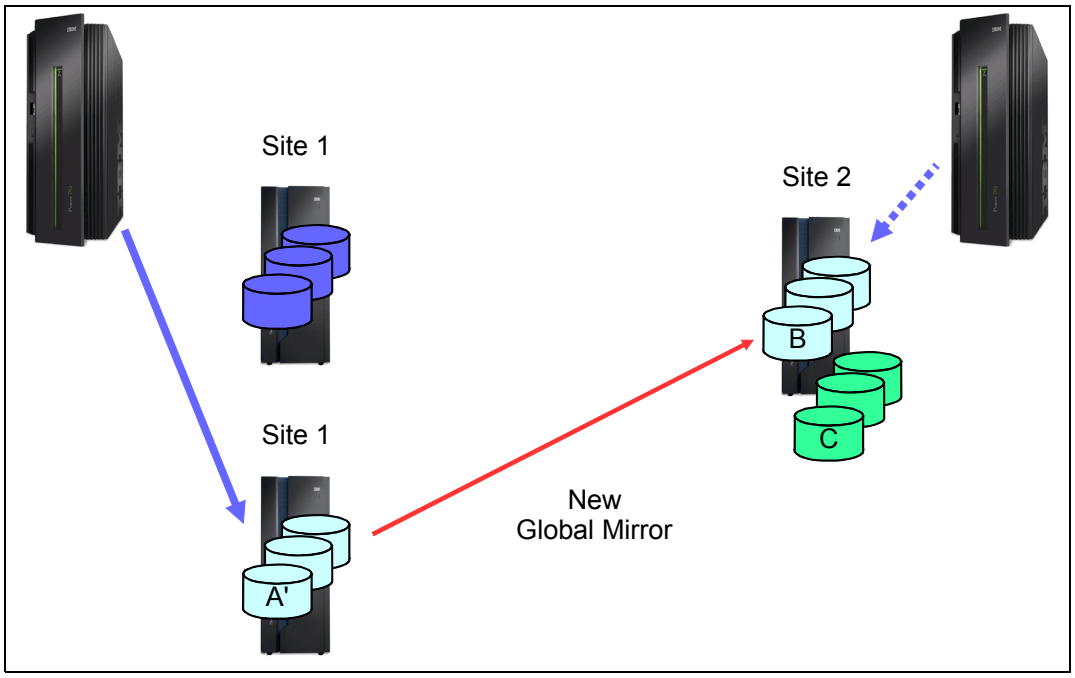

Figure 6-7 Cut-over to the new primary storage system

Complete these additional tasks to make A’ the production primary storage:

1. Remove the Global Mirror copy set of the original replication A to B from the old Tivoli Storage Productivity Center for Replication Global Mirror session.
2. Fail back the cascade Global Copy from A’ to B.
3. Add the new copy set A’, B and C to the new Global Mirror session in Tivoli Storage Productivity Center for Replication.

Post migration phase

In the post migration phase, the replaced storage system A must be prepared for disposal.

Do not forget to clean up obsolete zoning definitions in the storage area network.

6.2.4 Replacing primary and auxiliary storage

When both storage systems (primary and auxiliary) are to be replaced at the same time, the data is copied in a doubled cascaded Global Copy. Apart from a few extra tasks to handle the two cascaded Global Copy relations, the cut-over mechanism is comparable with that of the primary storage replacement.

Initial copy phase

To copy the data of the current Global Mirror environment to two new storage systems at each site, two cascaded Global Copy relations are required that are attached to the B volumes at the secondary site. However, to avoid two subsequent initial copy phases, the second Global Copy from B’ to A’ is established first as a no-copy relation. This means that only the copy relation is defined, and no data are copied at this moment.

The first cascaded Global Copy from B to B’ is established as a full copy, and all data in this initial copy phase are drained to the new primary storage A’. See Figure 6-8.

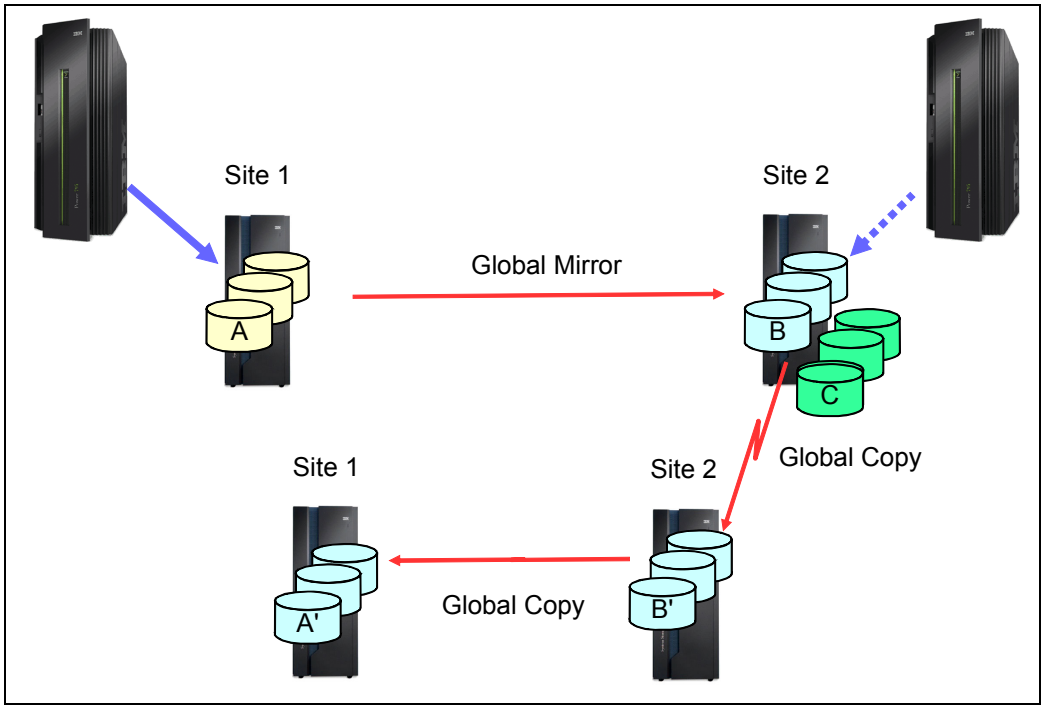

Figure 6-8 Initial setup with a doubled cascaded Global Copy

Complete these steps to set up the initial copy phase:

1. Establish the Global Copy with cascaded allowed and no-copy between B' and A'.
2. Establish the cascaded Global Copy between B and B' as a full copy relation.
3. Monitor the copy progress at the second cascaded Global Copy at B'.

During the initial copy phase of the Global Copy relation between B and B', all tracks received at B' are forwarded through the Global Copy relation between B' and A' to the primary site. When all tracks have been copied once from B to B', only a few tracks might be left to copy to A' because the link speed between B' and A' is presumably lower than the local speed at the secondary. To check whether the initial copy phase has successfully passed, issue the **`lspprc -l`** command to see that the relation between B' and A' indicates `First Pass Status` as shown in Example 4-2 on page 42. The migration is ready for the next phase only when the first pass status is `True` for all pairs.

Validation phase

The validation phase is done in the same way as in 6.2.3, "Replacing primary storage" on page 91. The data migration can only be validated with the primary servers. This means that the current production operation must be interrupted and the servers must access the new storage. Refer to Figure 6-9.

To provide host access to the A' volumes, the second Global Copy between B' and A' must be paused and failed over after the primary servers are stopped. Make sure that all tracks have been copied all the way to the A' volumes. To check, use the **`lspprc -l`** command for the Global Copy relation between B' and A'.

As assumed in the validation phase in 6.2.3, "Replacing primary storage" on page 91, the production servers are attached back to the original storage when the validation is finished.

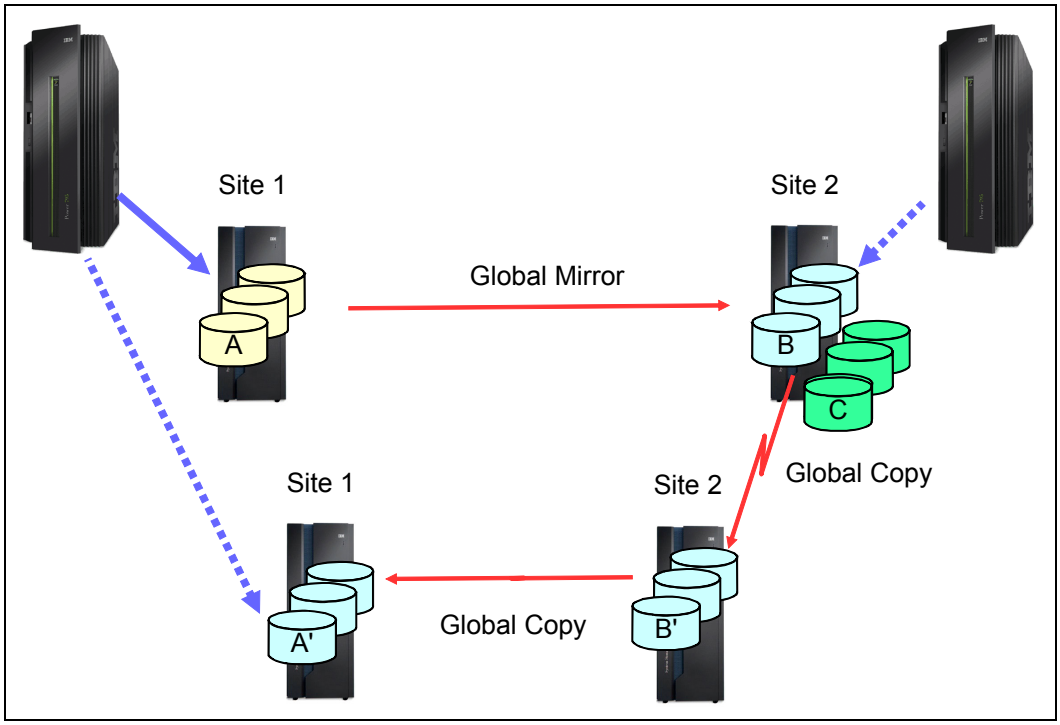

Figure 6-9 Validation of the migration when both storage systems are replaced

The validation at the primary site is done by completing the following sequence of steps:

1. Shut down the production.
2. Wait until all tracks are copied from B' to A'.
3. Pause the Global Copy between volume set B' and A' to allow a graceful failover.
4. Fail over cascaded Global Copy to the A' volumes.
5. Change the zoning to allow host access to A'.
6. Start validation test at A'.
7. When tests are done, stop validation systems and release the A' volumes.
8. Change zoning back to original configuration and start production at A.
9. Fail back cascaded Global Copy B to A' to prepare the migration phase.

Examples of corresponding DSCLI commands are shown in Example 4-1 on page 42 to Example 4-8 on page 46 and following.

Tip: When the migration is operated application-wise, only a subset of the whole application volume set is migrated in one go. To make sure that the correct volume sets in on scope, the cascaded Global Copy between B and B' can be paused for a while to see whether the number of out-of-sync tracks will increase. If so, the volume set does not meet the scope of the current migration.

Migration phase

In the migration phase, moving production to the new storage initially requires the same steps as performed in the validation phase. Furthermore, the first Global Copy must be removed and the cascaded Global Copy must be turned around, so that the production data are replicated from the new primary storage to the new auxiliary storage. Refer to Figure 6-10. Finally, the new Global Mirror session in Tivoli Storage Productivity Center for Replication must be populated with the new copy sets and the Global Mirror session must be started, if this has not happened yet.

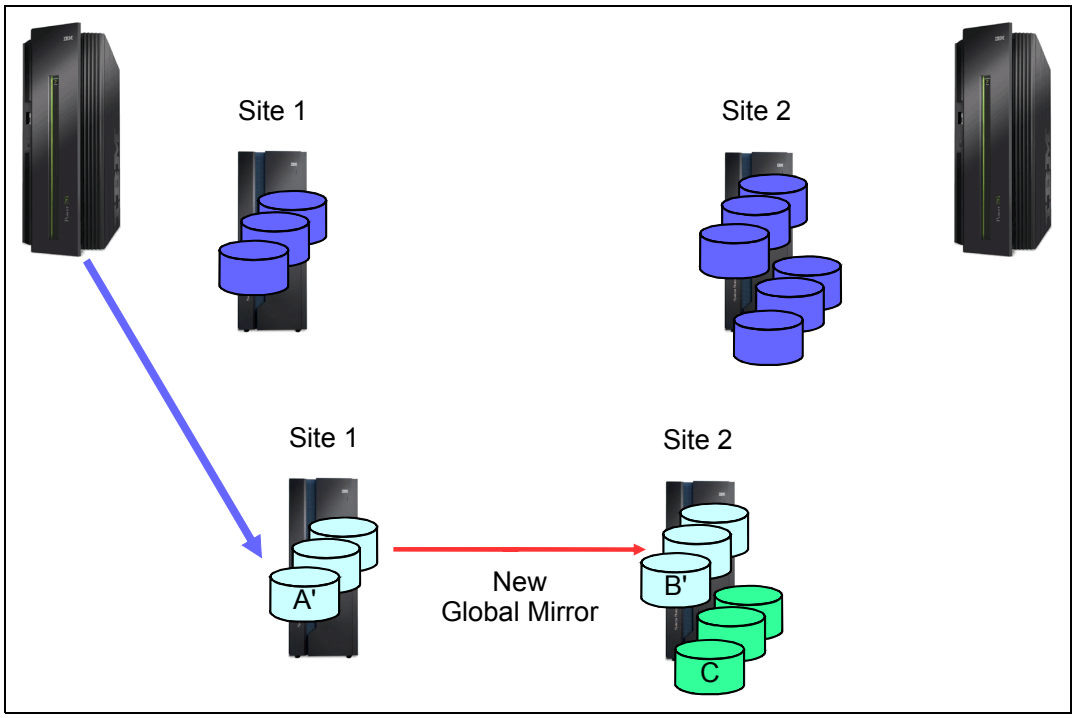

Figure 6-10 Cut-over to the new storage systems

In addition to the step shown previously in the validation phase, complete these tasks to cut over the production to the new storage systems:

1. Remove the cascaded Global Copy from B to B'.
2. Fail back the cascaded Global Copy from A' to B.
3. Add the new copy set A', B', and C' to the new Global Mirror session in Tivoli Storage Productivity Center for Replication.

Post migration phase

The post migration phase consists of cleaning up the configuration of the old DS8000 storage system:

1. Remove the FlashCopy relations between B and C.
2. Remove the Global Copy target relation of the original Global Copy between A and B at both sites. The status of the Global Copy is most likely failed over to B. Therefore, the easiest way to remove the Global Copy completely is to fail it back to A to B and then remove the Global Copy relationship.
3. Remove the cascaded Global Copy source relation between B and B' at B.
4. Remove all obsolete path definitions.
5. Delete the A, B, and C volumes.
6. Remove obsolete zoning definitions in the storage area network.

6.3 GDPS considerations

In addition to the general considerations discussed in Chapter 1, "Introducing disk data migration" on page 1, some considerations should to be taken into account when replacing storage systems in a GM environment with GDPS managing the environment.

The appropriate license codes must be purchased and applied to each of the systems in the migration scenario to use the Remote Mirror and Copy functions. For information about licensing requirements, speak to your IBM storage marketing representative.

To minimize the number of configuration changes required, configure the new storage systems as symmetrical as possible with the one being removed. This means, if at all possible, use the same LSS range, SSID range, UCB range, and PPRC links for the new storage system that were used on the old storage system.

For all migrations, use the cascade option for Global Copy to complete the initial data migration to the new storage system.

In a GDPS/GM environment, there are two GDPS controlling systems. This document refer to these controlling systems as follows:

- Kg-sys, which is the GDPS/GM local controlling system. This Kg-sys has a few different options as to its placement that are described in the *GDPS/GM 3.11 Planning and Installation Guide*, ZG24-6739-17.
- Kr-sys is the GDPS/GM remote controlling system. The Kr-sys is in the R-site (remote) on a stand-alone system.

Important: Special consideration needs to be taken into account when migrating a volume set that includes the GDPS Kx systems. For instance, if the Kg-sys system volumes are on a B volume set you are migrating, a Global Copy relationship needs to be created from C to B' for these devices. The system will need to be shut down, the relationships deleted, wait for the copy to complete and then IPL this Kg-sys on the new B' volume set after it is made available. The same is true for the Kr-sys.

Initial configuration

Figure 6-11 illustrates the basic layout of a Global Mirror environment. The host system accesses the primary volume set. There is an asynchronous Global Mirror (GM) relationship to the remote volume set.

Note: Consult with your IBM representative during the planning process of the migration project to validate your replication network requirements. This helps you to avoid inadvertently impacting the production and replication environments during the data migration.

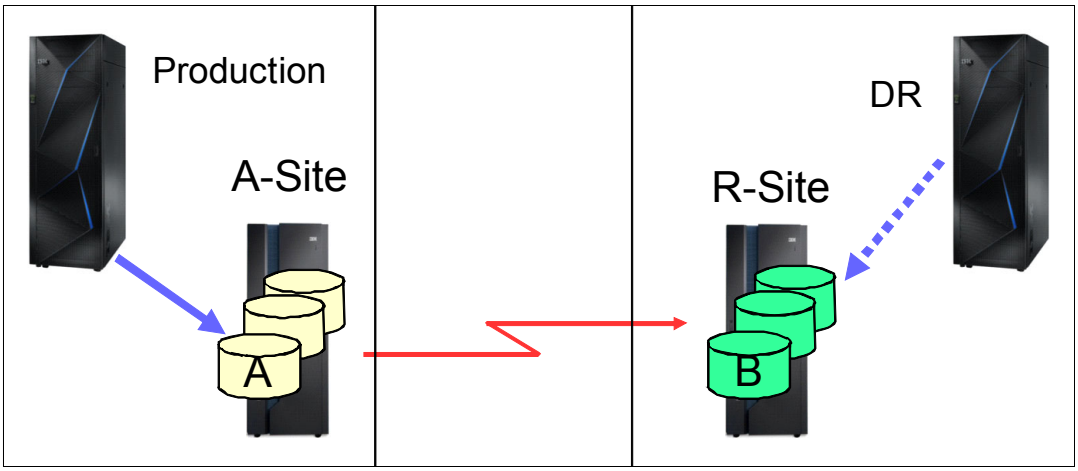

Figure 6-11 Initial GM Configuration

Replacing the B Volume Set with Automation

The approach described here to replace the B volumes set uses GDPS scripts and Copy Services to complete the migration. This process requires an interruption to host processing. If you cannot take a host system outage to replace the B volume set, a full copy needs to be completed from A to B' followed by rebuilding the GDPS/GM environment.

By using the cascaded Global Copy to migrate the data before the transition, you minimize the time that is required to fully recover the mirrored environment.

Additional Fibre Channel (FC) connectivity is required between B to B' and A to B'.

It is assumed that the B to B' cascaded relationship has completed First Pass before the actual migration, so the time required to obtain 100% copy status is minimal.

B' Initial Copy Phase

The initial setup of the cascaded replication environment is described in detail in 2.3, "Data replication" on page 22. See this section to complete this process and then return here after that is accomplished.

All write activity to the A devices must be stopped to begin the moving of the Global Mirror targets to the new populated B' volume set. The fundamental concept behind the migration plan is if there are no write updates to B and all out of sync tracks are copied to B'.

Create the cascade copy by establishing paths from B to B' and then establish an XD relationship from the B to the B' devices. These relationships should be established well in advance of your actual migration cut over. The copy time to complete the first pass depends on your network bandwidth and the configuration size of the volume set being migrated from/to. Figure 6-12 illustrates the GM configuration with the cascaded copy.

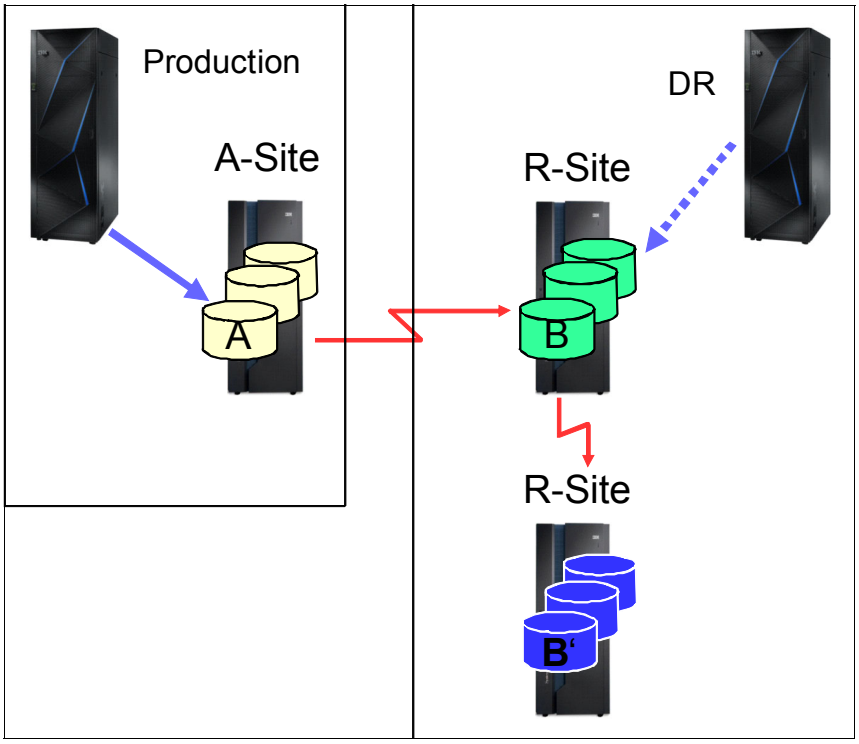

Figure 6-12 GM Configuration with Cascaded Copy to B'

B' Storage Migration Phase

Shut down the application systems standard operating procedures.

Stop the Global Mirror session and suspend the A to B relationship by running GDPS script statement '*session* STOP SECONDARY'.

If your GDPS Kg-sys system volumes have an XD relationship to the B volume set being replaced, shut it down using standard operating procedures.

Because there are no updates to the A devices at this point, validate that all changed tracks on the A devices have been copied to the B devices (out-of-sync tracks = 0). Failure to perform this action would corrupt the target volumes even though there would be no indication of this problem.

Wait for the copy to complete from A to B and B to B'.

Delete the Global Copy relationship from A to B and remove the paths from A to B.

Validate that all updates to the B devices have been copied to the B' devices (out-of-sync tracks = 0). As with the A to B validation, failure to ensure this relationship will corrupt the Global Mirror session at the end of this migration.

Delete the Global Copy relationships and paths from B to B’.

Remove the Global Mirror FlashCopy from the B devices by running GDPS script statement ‘*session* CLEANUPR GMFC’.

To complete the integration of the B’ volume set into the environment, install and validate all physical cabling required, including the replication links to the A volume sets. Use standard operating procedures to complete this task that you would normally perform to swap out hardware such as swapping cables or activating an IODF to remove the old box and add the new box.

Validate that the Kr-sys now references the B' devices instead of the old B devices. Issue additional **CQUERY** and other IBM MVS™ commands to validate as required.

Cut-over to the new volume set

This is the GO/NOGO decision point. Perform data validation on the new B’ volume set. If you determine there are errors during the validation process and you need to back out, refer to the “Back Out Scenario for B’” on page 101.

Define the PPRC paths and establish the Global Copy relationships between the A and B’ devices. Because the A and B' devices are identical, the NOCOPY option is used to prevent a full copy to the B' devices.

If your Kg-sys was shut down during the migration phase, IPL the Kg-sys using standard procedures.

Establish the Global Mirror B’ to C’ ASYNC FlashCopy relationships.

Create a GEOMPARM with the B’ volume set and load this GEOMPARM configuration by using the GDPS/GM “Configuration Management” panel (option C).

Run all GDPS/Global Mirror monitors and verify the accuracy of the recently loaded configuration.

Using the GDPS/Global Mirror DASD Management panels (option 1), start the Global Mirror session.

Validate that Global Mirror is forming consistency groups by using the GDPS DASD panels and the Global Mirror Monitor is started.

Perform necessary DR testing to certify that the application Disaster Recovery images are now accessing the new B' volume set.

Figure 6-13 shows the final GM configuration with the new B’ volume set.

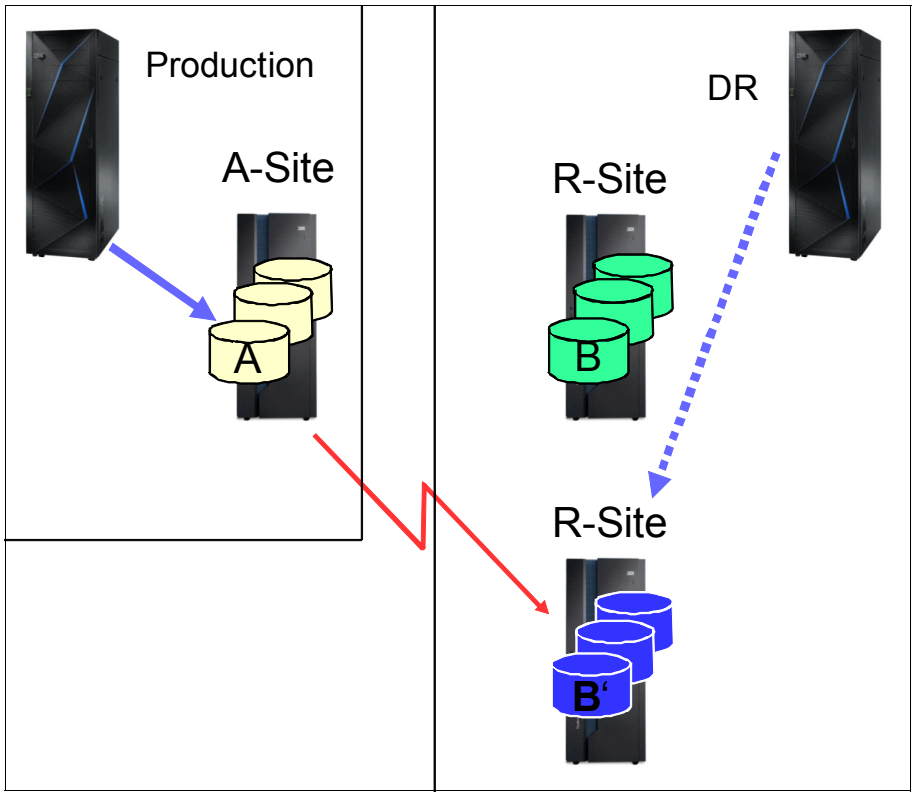

Figure 6-13 GM final configuration with B’

Back Out Scenario for B’

If during the validation testing you encounter a situation where you need to back out, follow these procedures to insert the original B volume set into the configuration:

1. At the point of GO/NOGO during validation testing, your configuration A to B does not exist, nor does B to B’.
2. Remove the B’ volume set and add the B volume set back into your configuration by using the same technique such as swapping cables or reverting to the original IODF. Verify that your Kr-sys can access the original B devices.
3. If your Kg-sys was shut down during the migration phase, IPL the Kg-sys using standard procedures.
4. Blackout the GEOMPARM that references the new B’ volume set by loading the old GEOMPARM with the B volume set.
5. Run all GDPS/GM monitors to validate the configuration.
6. Establish the Global Copy paths from A to B.
7. Establish the Global Copy pairs from A to B.
8. Create the Global Mirror B to C ASYNC FlashCopy relationships.
9. GDPS/Global Mirror should be green OK and forming Consistency Groups.

   The environment has now been restored to the state before the B' migration. Global Mirror is running A to B.

Replacing the A Volume Set with Automation

The approach described here to replace the A volumes set uses GDPS scripts and Copy Services to complete the migration. This process requires an interruption to host processing.

By using the cascaded Global Copy to migrate the data before the transition, you minimize the time that is required to fully recover the mirrored environment.

Additional FC connectivity is required between B to A' and A' to B.

It is assumed that the B to A' cascaded relationship has completed First Pass before the actual migration, so the time required to obtain 100% copy status is minimal.

A' Initial Copy Phase

The initial set up of the cascaded replication environment is described in detail in 2.3, "Data replication" on page 22. Complete the process in that section and then return here after that is accomplished.

All write activity to the A devices must be stopped to begin the moving of the Global Mirror source volumes to the new populated A' volume set. The fundamental concept behind the migration plan is if there are no write updates to A and all out of sync tracks are copied to A'.

Create the cascade copy by establishing paths from B to A' and then establish an XD relationship from the B to the A' devices. These relationships should be established well in advance of your actual migration cut over. The copy time to complete first pass depends on your network bandwidth and the configuration size of the volume set being migrated from/to. Figure 6-14 illustrates the GM configuration with the cascaded copy.

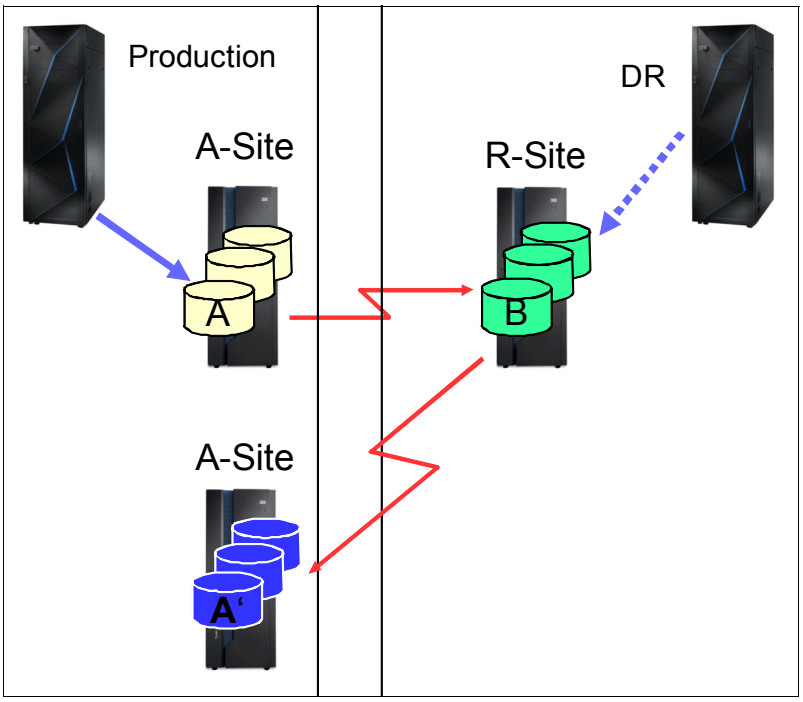

Figure 6-14 GM Configuration with Cascaded Copy to A'

A’ Storage Migration Phase

Complete these steps to migrate the A’ storage:

1. Shut down the application systems using standard operating procedures.
2. If your GDPS Kg-sys system volumes are on the A volume set being replaced, shut it down using standard operating procedures.
3. Wait for the copy to complete from A to B.
4. Stop the Global Mirror session and suspend the A to B relationship. If your GDPS Kg-sys did not need to be shut down in the previous step, run the GDPS script statement ‘*session* STOP SECONDARY’. If it was shut down, use Copy Services commands to suspend the A to B relationships and stop the Global Mirror session.
5. Delete the Global Copy relationship from A to B and remove the paths from A to B.
6. Validate that all updates to the B devices have been copied to the A’ devices (out-of-sync tracks = 0). Failure to ensure this relationship will corrupt the Global Mirror session at the end of this migration.
7. To make the A’ devices available for validation, define the PPRC paths from A’ to B and then issue a failover of the Global Copy relationships from A’ to B.

   Attention: Data validation should be accomplished with a stand-alone non-application system to prevent any production data from being written. Any updates to the new A’ volume set should be throw away in the event that a blackout is required.

8. To complete the integration of the A’ volume set into the environment, install and validate all physical cabling required, including the replication links to the B volume sets. Use standard operating procedures to complete this task, which you would normally perform to swap out hardware such as swapping cables or activating an IODF to remove the old box and add the new box.

Cut over to the new volume set

This is the GO/NOGO decision point. Perform data validation on the new A’ volume set. If you determine there are errors during the validation process and you need to back out, see the “Back Out Scenario for A’” on page 104. Complete these steps:

1. Reverse the direction of the copy from B to A’ to A’ to B. Issue a fail back A’ to B using the OPTION(SYNC) and FORCE(YES) options.
2. If your Kg-sys was shut down during the migration phase, IPL the Kg-sys using standard procedures.
3. Create a GEOMPARM with the A’ volume set and load this GEOMPARM configuration by using the GDPS/GM “Configuration Management” panel (option C). Define the new A’ LSS’ and join the A’ volumes to the Global Mirror session.
4. Run all GDPS/Global Mirror monitors and verify the accuracy of the recently loaded configuration.
5. Using the GDPS/Global Mirror DASD Management panel (option 1), start the Global Mirror session.
6. Validate that Global Mirror is forming consistency groups by using the GDPS DASD panels and that the Global Mirror Monitor is started.
7. Perform necessary DR testing to certify that the application Disaster Recovery images are now accessing the new B' volume set.

Figure 6-15 shows the final GM configuration with the new A’ volume set.

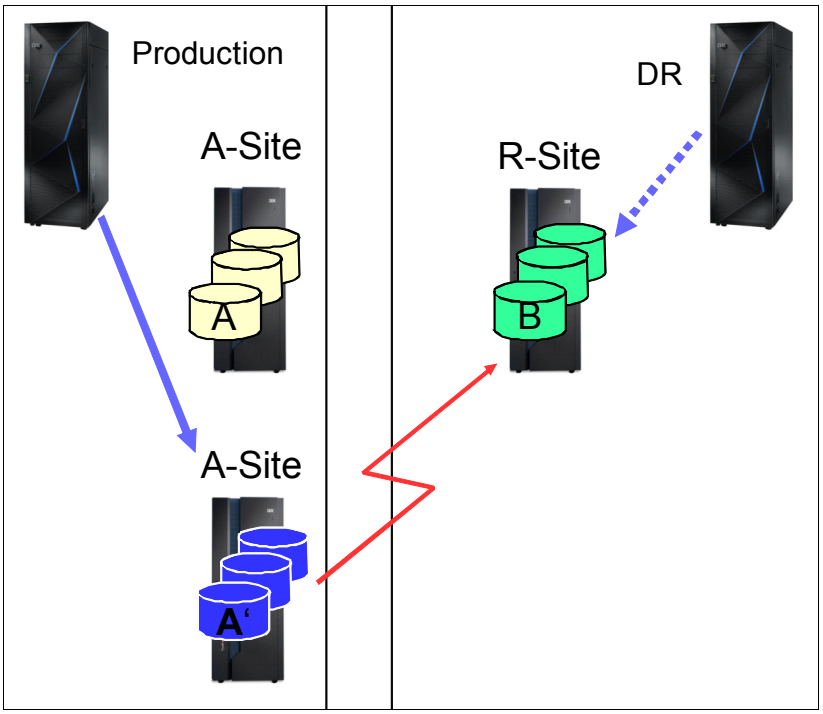

Figure 6-15 Final GM Configuration with A’

Back Out Scenario for A’

If during the validation testing you encounter a situation where you need to back out, follow these procedures to insert the original B volume set into the configuration:

1. At the point of GO/NOGO during validation testing, your configuration A to B does not exist and A’ is failed over to B. There should have been no application I/O to the environment because the GO/NOGO was before starting any application systems.
2. Remove the A’ volume set and add the A’ volume set back into your configuration using the same technique such as swapping cables or reverting to the original IODF.
3. Establish the Global Copy paths from A to B.
4. Establish the Global Copy pairs from A to B with the NOCOPY option.
5. If your Kg-sys was shut down during the migration phase, IPL the Kg-sys using standard procedures.
6. Blackout the GEOMPARM that references the new A’ volume set by loading the old GEOMPARM with the A volume set.
7. Run all GDPS/GM monitors to validate the configuration.
8. GDPS/Global Mirror should be green OK and forming Consistency Groups.
9. The environment has now been restored to the state before the A' migration. Global Mirror is running A to B.

7

Data Migration with Metro/Global Mirror replication

This chapter discusses the data migration considerations for a three site Metro/Global Mirror (MGM) environment using the DS8000 Global Copy feature. The existing environment can be managed by Geographically Dispersed Parallel Sysplex (GDPS) or Tivoli Storage Productivity Center for Replication, but does not have to be. The physical location of the participating disk systems does not affect the processes discussed below, but might have an impact on the infrastructure requirements.

This chapter includes the following sections:

- Initial configuration
- Migration process
- Migration Scenarios
- Geographically Dispersed Parallel Sysplex (GDPS) considerations

© Copyright IBM Corp. 2015. All rights reserved.

7.1 Initial configuration

Figure 7-1 illustrates the basic layout of a 3-site Metro Global Mirror (GM) configuration. The host system accesses the primary volume set A. There is a synchronous Metro Mirror replication relationship to the B volume sets and a cascaded asynchronous with consistency Global Mirror replication relationship to the C volume sets.

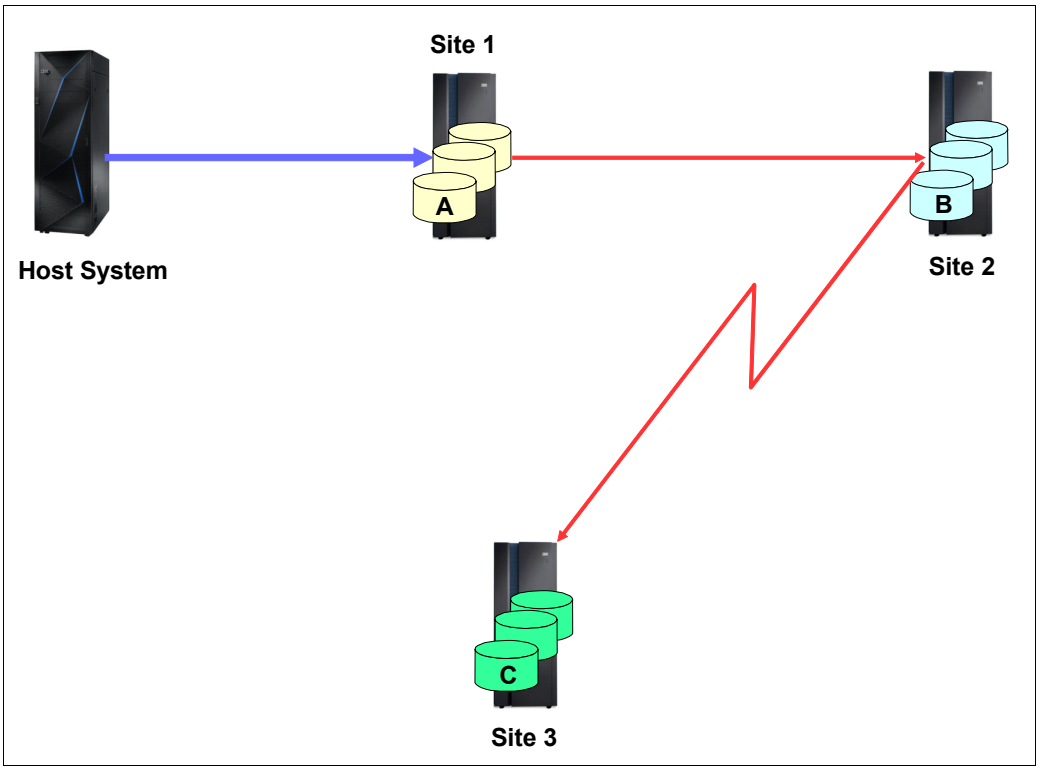

Figure 7-1 Basic 3-Site Metro-Global Mirror Environment

7.2 Migration process

The following sections describe the processes, assumptions, and general considerations for migrating data in an existing MGM environment.

It is a preferred practice to replace volume sets in an MGM environment as follows:

- Starting with the C volume set first. Doing so minimizes the potential for performance issues that might result by installing a storage system with higher throughput capacity for the A or B volume sets that might overwhelm the storage system containing the C volume set.
- Replace all three systems simultaneously, eliminating the possibility of the performance issues mentioned above.

The scenarios included in this chapter follow that practice by starting the migration with the C volume set, then moving to the B and A volume sets. There is also a scenario provided to replace all three volume sets simultaneously.

The following assumptions are made with regard to the migration scenarios in this chapter:

- The storage systems being migrated from do not support Multiple-Target PPRC
- Cascaded Global Copy will be used to complete the initial migration of the data
- Normal host processing operations will continue whenever possible
- The availability of the MGM environment will be maintained whenever possible
- The migrated data will remain in the MGM environment
- The volume sets and scope of the migration project are clearly defined and documented
- The applications affected by the migration project are clearly defined and documented
- The post-migration environment is clearly defined and documented

The examples in this chapter reference the migration of an entire storage system, but the process is valid for multiple storage systems or even a subset of a single storage system. If the plan is to migrate a subset of the data on one or multiple storage systems, the master migration plan must clearly define the applications and volume sets that will be affected.

To minimize the impact of the migration process, complete these steps:

- Use Global Copy to create an initial copy of data for each volume set in the migration environment without interruption to the host processing.
- Validate that there are no Out-of-Sync tracks before moving forward with the migration. This process ensures that each volume set in the environment is logically identical before it is safe to proceed.
- Use the **`mkpprc`** command with the **`-type nocp`** option when required to maintain the logically identical state. This option creates the Global Copy or Metro Mirror relationship without transferring any data from the primary volume set to the secondary volume set when the pairs are created.

Note: The examples included in this section are intended to be used as the template for creating a migration process. They are not to be used ‘as is’.

Each storage system referenced in the DSCLI examples use the generic storage image ID of 75xxxxx to refer to the -dev device and 75yyyyy to refer to the -remotedev device.

7.2.1 General considerations

This topic is covered in detail in Chapter 1, “Introducing disk data migration” on page 1 and Chapter 2, “Copy Services based data migrations” on page 19.

Important: It is essential to the success of any migration project that the information discussed in the two chapters referenced above is thoroughly reviewed before proceeding.

7.3 Migration Scenarios

The various migration scenarios described are for a three site MGM environment that does not use automation to manage the replication. All migration activities are performed using DS8000 DSCLI commands.

7.3.1 Replacing the C volume set without automation

The process to replace the C volume set in an existing MGM environment can be completed without interruption to the host processing at the primary site. The MGM environment remains intact during the initial copy process. Your specific migration might require additional steps, depending on the environment.

The migration process is broken down into five major steps:

1. Initial setup of the migration environment
2. Data validation testing
3. Preparing to migrate
4. Cut over to the new volume set
5. Recovering the pre-migration environment

The steps assume that the new volume set is installed and configured as required. The cascade option for Global Copy is used to complete the initial data migration to the new volume set. DS8000 DSCLI commands are used to build the migration environment and cut-over to the new volume set.

Important: The potential for human error is significantly increased if there is no automation software to help the migration process. It is strongly suggested that the DSCLI command set to be used in the migration process be reviewed and validated by the local IBM representative for correct syntax and sequence before setting up the migration environment.

If this is a GDPS managed environment, see 7.4, “Geographically Dispersed Parallel Sysplex (GDPS) considerations” on page 137.

Initial setup of the migration environment

The initial setup of the cascaded replication environment is described in 2.3, “Data replication” on page 22. Refer to that section to complete this process and then return here after that is accomplished. For this migration scenario, the cascaded Global Copy is extended to an additional volume set. The migration path is from A → B → C → C’

Figure 7-2 on page 109 illustrates the basic layout of the migration configuration. In addition to the existing MGM relationship between A, B, and C, an asynchronous Global Copy leg from the current C volume set to the new C’ volume set has been added. This relationship will allow the data from the C volume set to be copied to the C’ volume set without impact to the production applications or the existing MGM environment.

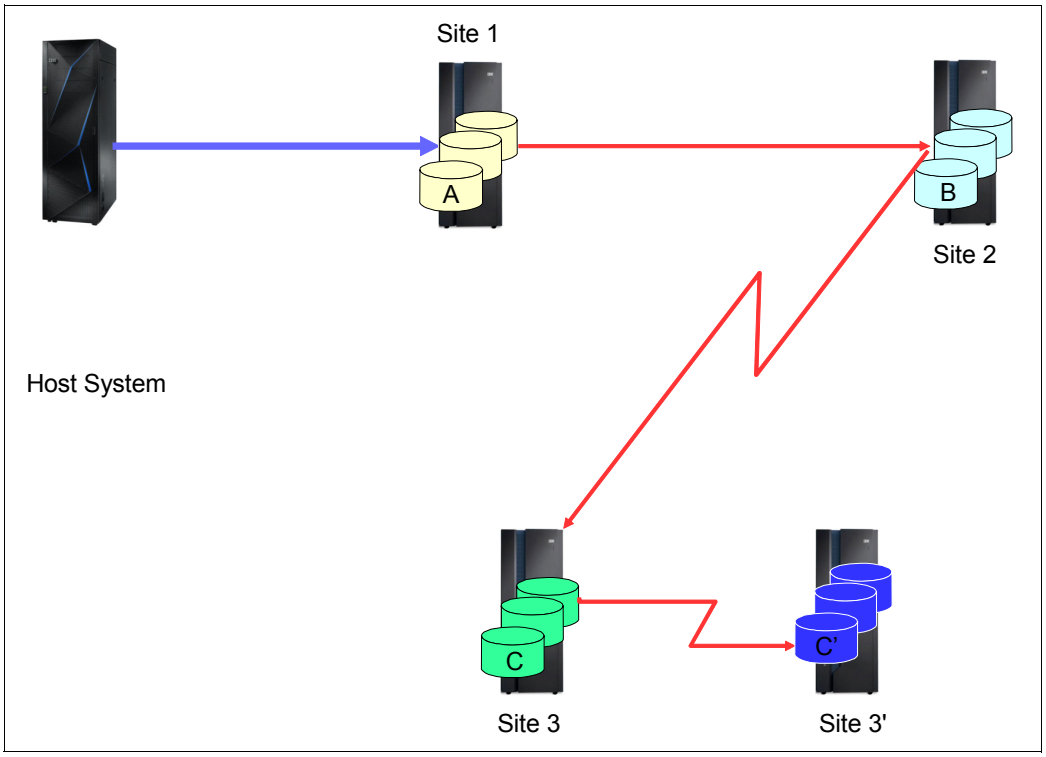

Figure 7-2 Cascaded Global Copy with existing 3-Site Metro-Global Mirror Environment

Figure 7-3 illustrates the post-migration environment with C' as the new secondary device in the MGM environment.

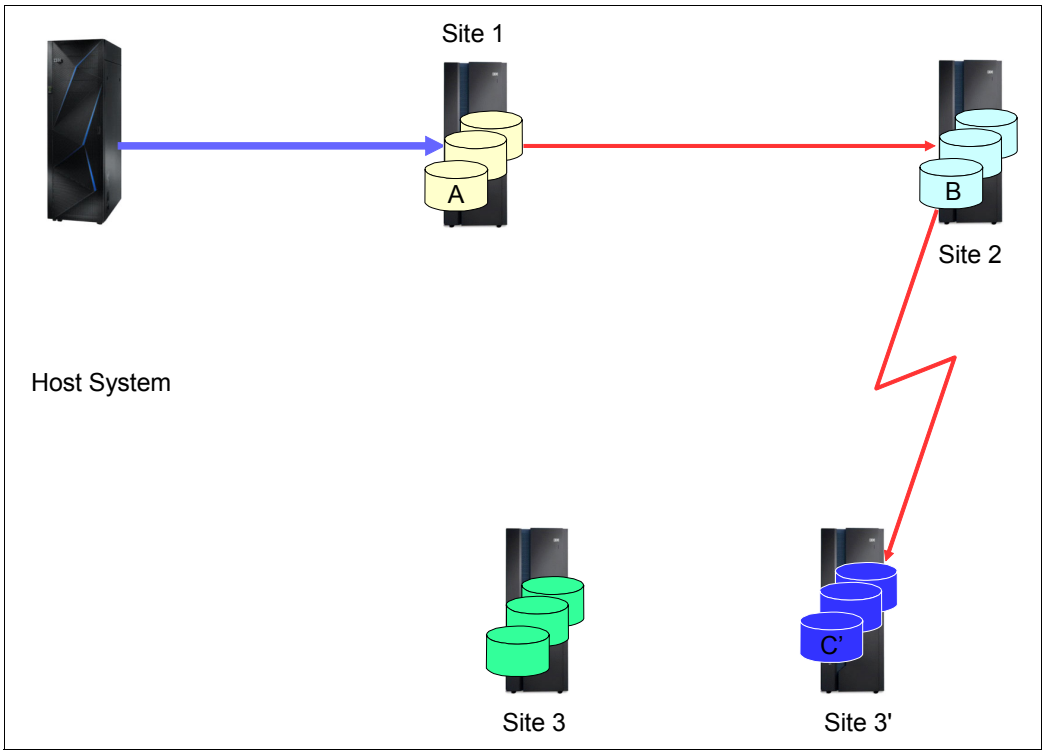

Figure 7-3 Post migration 3-Site Metro-Global Mirror Environment

Data validation testing

An initial data validation test can be performed on the C' volume set before proceeding with the final cut-over. An available test environment with access to the C' volume set is required.

This process can be accomplished with minimum impact to the current environment, allowing a quick fall back to the pre-migration configuration should it be necessary. For more information, see 2.4, "Data validation testing" on page 24.

Prepare to migrate to the new volume set

This process can be completed without interruption to host processing. The MGM environment will be temporarily affected. By using the cascaded Global Copy to migrate the data before the transition, the time required to fully recover the mirrored environment is minimized.

To complete the integration of the C' volume set into the environment, install and validate all physical cabling required, including the replication links to the B volume sets. The C volume set should also be removed from the environment now to ensure that there is no possibility of data being written to this volume set.

After the physical connectivity has been completed and validated, the MGM environment can be rebuilt with the C' volume set as the new secondary volume set using the following process.

The first step in this process is to temporarily suspend the Metro Mirror relationship between the A volume set and the B volume set. Doing so ensures that no updates are made to the C volume set during the migration process.

This process accomplishes several important objectives:

- Place the A volume in a state where no host I/O can be directed to them, either Extended Long Busy or Queue Full
- Remove the remote mirror and copy paths between the A and B volume sets.
- Create a remote mirror and copy consistency group
- Ensure all updates that are made to the A volume set are tracked. These updates will be written to the B and C' volume sets after the C' volume set has been added to the replication environment.
- Reset the condition for the A volume set to allow the host I/O activity to resume.

This task is accomplished using the DSCLI **`freeze`** and **`unfreeze`** commands at the A volume set.

> **Important:** The **`freezepprc`** command causes a temporary disruption to host I/O processing. Ensure the potential impact to the environment, and how to use the **`unfreezepprc`** command, is understood before proceeding.

Example 7-1 shows the DSCLI commands (and resulting output) used to freeze and then unfreeze the volume sets. In this example, the command is being issued at the A volume set.

Example 7-1 DSCLI commands to freeze and unfreeze the A to B Metro Mirror relationship

```
freezepprc -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy 76:76
CMUC00161I freezepprc: Remote Mirror and Copy consistency group 76:76 successfully created.
```

```
unfreezepprc -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy 76:76
CMUC00198I unfreezepprc: Remote Mirror and Copy pair 76:76 successfully thawed.
```

After the Metro Mirror relationship between the A and B volume sets has been suspended, the Out-of-Sync tracks status on the B and C volume sets must be verified.

Example 7-2 shows the DSCLI commands (and resulting output) used to check the Out-of-Sync tracks status for the copy relationships. Each query must show zero Out-of-Sync tracks before proceeding with the next step.

Example 7-2 DSCLI Commands to verify Out-of Sync track status for Global Copy, Metro Mirror, and Global Mirror

```
lspprc -l -dev ibm.2107-75xxxxx 7600-7601 Verify zero Out-of-Sync tracks at B
lspprc -l -dev ibm.2107-75xxxxx 7A00-7A01 Verify zero Out-of-Sync tracks at C

Verify zero Out-of-Sync tracks for GM session at B
dscli> showgmiroos -dev ibm.2107-75xxxxx -lss 76 -scope si 2
Scope           IBM.2107-75xxxxx
Session         02
OutOfSyncTracks 0
```

Important: If any of the Out-of-Sync track counts do not decrement to zero, data must still be written to the B' volume set. This situation must be resolved before proceeding.

Now remove the Global Mirror relationship between the B and C volume sets. After this process is completed, completely remove the C volume set from the MGM environment. This process requires six steps:

1. Remove the volumes from the Global Mirror session.
2. Close the GM session.
3. End the GM session.
4. Remove the Global Copy pairs between the B volume set and the C volume set.
5. Remove the replication paths between the B volume set and the C volume set.
6. Remove the GM FlashCopy pairs on the C volume set. This step is not essential because the machine will be taken out of the environment, but it ensures that the FlashCopy relationships will not interfere with any plans for this volume set.

Example 7-3 shows the DSCLI commands (and resulting output) used to complete this step. Some of these commands must be validated.

Example 7-3 DSCLI Commands to remove the Global Mirror relationship from B to C

```
chsession -dev ibm.2107-75xxxxx -lss 76 -action remove -volume 7600-7601 2
CMUC00147I chsession: Session 2 successfully modified.

rmsession -dev ibm.2107-75xxxxx -lss 76 2
CMUC00148W rmsession: Are you sure you want to close session 2? [y/n]:y
CMUC00146I rmsession: Session 2 closed successfully.

rmgmir -dev ibm.2107-75xxxxx -lss 76 -session 2
CMUC00166W rmgmir: Are you sure you want to stop the Global Mirror session 2:? [y/n]:y
CMUC00165I rmgmir: Global Mirror for session 2 successfully stopped.

rmpprc -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy 7600-7601:7A00-7A01
```

```
CMUC00160W rmpprc: Are you sure you want to delete the Remote Mirror and Copy volume pair relationship
7600-7604:7A00-7A04:? [y/n]:y
CMUC00155I rmpprc: Remote Mirror and Copy volume pair 7600:7A00 relationship successfully withdrawn.
CMUC00155I rmpprc: Remote Mirror and Copy volume pair 7601:7A01 relationship successfully withdrawn.

rmpprcpath -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy 76:7A
CMUC00152W rmpprcpath: Are you sure you want to remove the Remote Mirror and Copy path 76:7A:? [y/n]:y
CMUC00150I rmpprcpath: Remote Mirror and Copy path 76:7A successfully removed.


rmflash -dev ibm.2107-75xxxxx 7A00-7A01:7A10-7A11
CMUC00144W rmflash: Are you sure you want to remove the FlashCopy pair 7A00-7A01:7A10-7A11:? [y/n]:y
CMUC00140I rmflash: FlashCopy pair 7800:7802 successfully removed.
CMUC00140I rmflash: FlashCopy pair 7801:7803 successfully removed.
```

Cut-over to the new volume set

For this step, establish the replication paths and set up the Global Mirror pairs from the B volume set to the C' volume set. This process requires six steps:

> **Note:** Only step 1 is performed at the C' volume set. All other commands are performed at the B volume set.

1. Create the Global Mirror FlashCopy relationships for the C' volume set.
2. Establish the replication paths between the B volume set and the C' volume set.
3. Create the Global Copy pairs between the B volume set and the C' volume set.
4. Open the Global Mirror session.
5. Add the volumes to the Global Mirror session.
6. Start the GM session

Example 7-4 shows the DSCLI commands (and resulting output) used to complete this step.

Example 7-4 DSCLI Commands to create the Global Mirror relationship from B to C'

```
mkflash -dev IBM.2107-75xxxxx -tgtinhibit -record -persist -nocp 7800-7801:C800-C801
CMUC00137I mkflash: FlashCopy pair 7800:C800 successfully created.
CMUC00137I mkflash: FlashCopy pair 7801:C801 successfully created.


mkpprcpath -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy -remotewwnn 5005076305FFD71A
-srclss 78 -tgtlss 7A I0201:I0032 I0234:I0306
CMUC00149I mkpprcpath: Remote Mirror and Copy path 78:7A successfully established.

mkpprc -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy -type gcp -mode nocp -cascade
7800-7801:7a00-7a01
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 7600:7A00 successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 7601:7A01 successfully created.

mksession -dev ibm.2107-75xxxxx -lss 78 2
CMUC00145I mksession: Session 2 opened successfully.

chsession -dev ibm.2107-75xxxxx -lss 78 -action add -volume 7800-7801 2
CMUC00147I chsession: Session 2 successfully modified.
```

```
mkgmir -dev ibm.2107-75xxxxx -lss 78 -session 2
CMUC00162I mkgmir: Global Mirror for session 2 successfully started.
```

> **Important:** This is the Go/No Go decision point for this migration scenario. Proceeding with the following step prohibits a return to the initial pre-migration state.
>
> To halt the migration process and return to the original configuration, see “Return to the pre-migration configuration” on page 113.

Finally, resume the replication process between A and B. This allows all of the updates to drain and complete the rebuilding of the MGM environment with the C’ volume set. The replication paths between A and B must be re-established before resuming the Metro Mirror.

Example 7-5 shows the DSCLI commands (and resulting output) used to complete this step. Note the inclusion of the **-consistgrp** option with the **mkpprcpath** command. This option creates a consistency group for the remote mirror and copy volume pairs being defined to this Metro Mirror session.

Example 7-5 DSCLI Commands to resume the Metro Mirror relationship from A to B

```
mkpprcpath -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy -remotewwnn 5005076305FFD75A
-srclss 78 -tgtlss 76 -consistgrp I0003:I0201 I0133:I0234
CMUC00149I mkpprcpath: Remote Mirror and Copy path 78:76 successfully established.

resumepprc -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy -type mmir 7600-7600:7600-7601
CMUC00158I resumepprc: Remote Mirror and Copy volume pair 7600:7600 relationship successfully
resumed. This message is being returned before the copy completes.
CMUC00158I resumepprc: Remote Mirror and Copy volume pair 7601:7601 relationship successfully
resumed. This message is being returned before the copy completes.
```

The MGM environment is now rebuilt with the new C’ volume set. The storage migration process for this site is now completed.

Return to the pre-migration configuration

The process below will return the environment to the pre-migration configuration for this scenario. The specific process should be documented, tested, and included in the master migration plan.

> **Attention:** The **-dev** and **-remotedev** statements in the DSCLI commands need to be changed to refer to the appropriate volume set.

Remove the Global Mirror relationship between the B and C' volume sets as described in Example 7-3 on page 111.

Rebuild the Global Mirror relationship between the B and C volume sets as described in Example 7-4 on page 112

Resume the Metro Mirror relationship between the A and B volume sets as described in Example 7-5

The pre-migration MGM environment has now been restored.

7.3.2 Replacing the B volume set without automation

This section describes the process to replace the B volume set in an existing MGM environment. The process can be completed without interruption to the host processing at the primary site. The MGM environment remains intact during the initial copy process. Your specific migration might require more steps, depending on the environment.

The migration process is broken down into five major steps:

1. Initial setup of the migration environment
2. Data validation testing
3. Preparing to migrate
4. Cut over to the new volume set
5. Recovering the pre-migration environment

The steps assume that the new volume set is installed and configured as required. The cascade option for Global Copy is used to complete the initial data migration to the new volume set. DS8000 DSCLI commands are used to build the migration environment and cut over to the new volume set.

Important: The potential for human error is significantly increased if you do not use automation software to help during the migration process. Therefore, the DSCLI command set to be used in the migration process should be reviewed and validated by the local IBM representative for correct syntax and sequence before setting up the migration environment.

If this is a GDPS managed environment, see 7.4, “Geographically Dispersed Parallel Sysplex (GDPS) considerations” on page 137.

Initial setup of the migration environment

The initial setup of the cascaded replication environment is described in detail in 2.3, “Data replication” on page 22. Refer to that section to complete this process and then return here after that is accomplished. For this migration scenario, the cascaded Global Copy is extended to an additional volume set. The migration path is from A → B → C → B’.

Note: Consult an IBM representative during the planning process of the migration project to validate the replication network requirements. This helps avoid inadvertently impacting the production and replication environments during the data migration.

Figure 7-4 on page 115 illustrates the basic layout of the migration configuration. In addition to the existing MGM relationship between A, B, and C, an asynchronous Global Copy leg from the current C volume set to the new B’ volume set has been added. This relationship allows the data from the C volume set to be copied to the B’ volume set without impact to the production applications or the existing MGM environment.

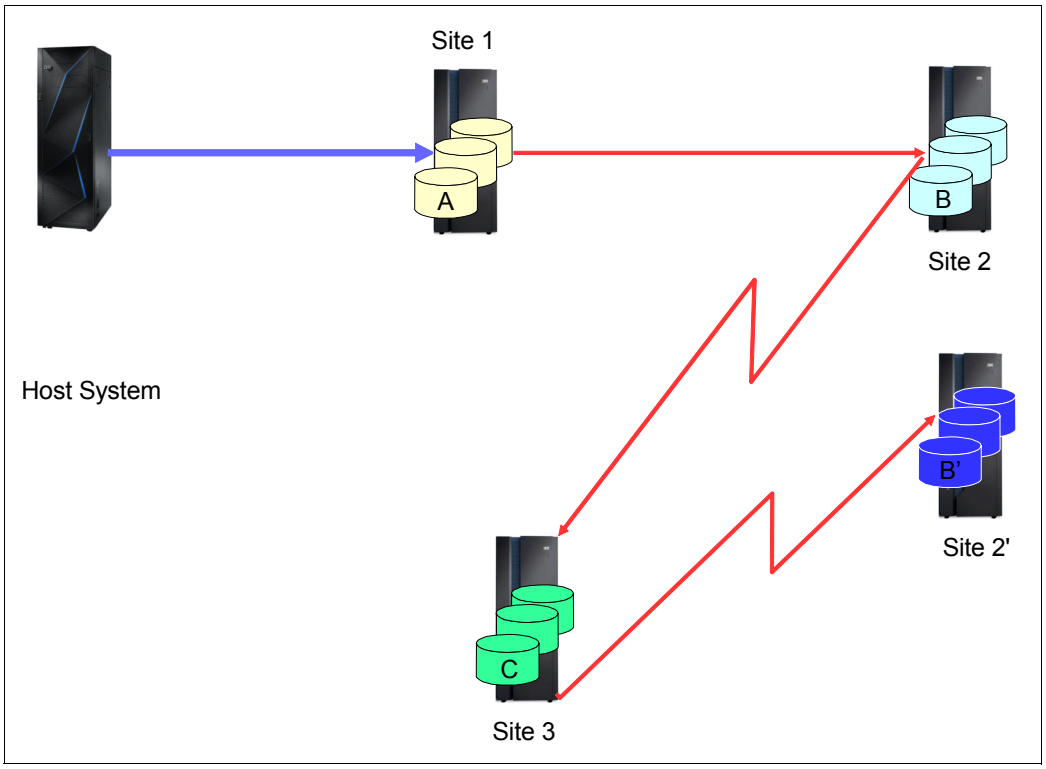

Figure 7-4 Cascaded Global Copy with existing 3-Site Metro-Global Mirror Environment

Figure 7-5 illustrates the post-migration environment with B' as the new secondary device in the MGM environment.

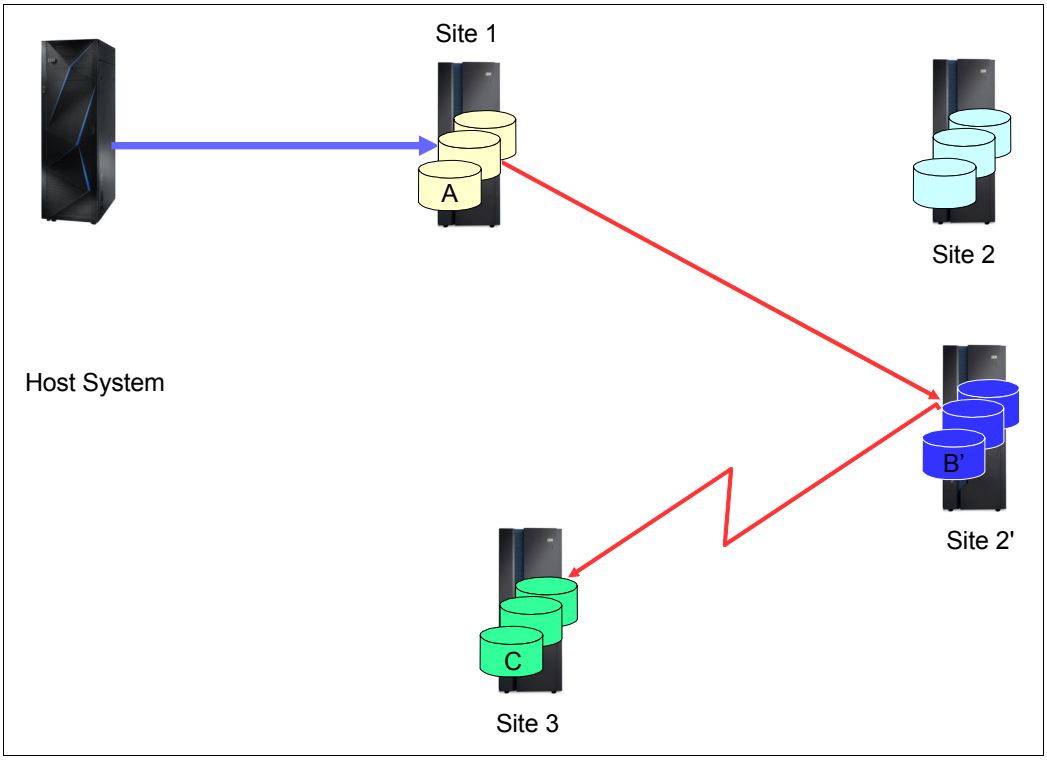

Figure 7-5 Post migration 3-Site Metro-Global Mirror Environment

Data validation testing

Perform an initial data validation test on the B' volume set before proceeding with the final cut over. An available test environment with access to the B' volume set is required.

This process can be accomplished with minimum impact to the current environment, allowing a quick fall back to the pre-migration configuration if necessary. For more information, see 2.4, "Data validation testing" on page 24.

Prepare to migrate to the new volume set

This process can be completed without interruption to host processing. The MGM environment will be temporarily affected. By using the cascaded Global Copy to migrate the data before the transition, the time required to fully recover the mirrored environment is minimized.

To complete the integration of the B' volume set into the environment, install and validate all physical cabling required, including the replication links to the A and C volume sets. The B volume set should also be removed from the environment now to ensure that no data is written to this volume set.

After the physical connectivity has been completed and validated, the MGM environment can be rebuilt with the B' volume set as the new secondary volume set using the following process.

The first step in this process is to validate that the Metro Mirror relationship between the A and B volume sets has the incremental resync option enabled. This ensures that any updates made to the A volume set can be tracked while the migration is taking place.

This can be accomplished by using the **`lspprc -l`** command to verify the status of the Incremental Resync option. The option should be enabled if the A to B Metro Mirror pairs were created using the correct options for the MGM environment, but this cannot be assumed.

> **Important:** The status of Incremental Resync on the A volume set must be validated and enabled if necessary before proceeding.

Example 7-6 shows the DSCLI commands (and resulting output) used to verify the status of incremental resync at the A volume set. In this example, the option is enabled.

Example 7-6 DSCLI command to verify the incremental resync status

```
lspprc -l -dev ibm.2107-75xxxxx 7600-7601
ID        State       Reason Type         Out Of Sync Tracks Tgt Read Src Cascade Tgt Cascade Date
Suspended SourceLSS Timeout (secs) Critical Mode Fi
rst Pass Status Incremental Resync Tgt Write GMIR CG PPRC CG isTgtSE Pmir     DisableAutoResync
==========================================================================================7600:7800
Full Duplex -      Metro Mirror 0                  Disabled Disabled    Invalid     -             76
60             Disabled      In
valid           Enabled            Disabled  N/A     Enabled Unknown Disabled -
7601:7801 Full Duplex -      Metro Mirror 0                  Disabled Disabled    Invalid     -
76        60             Disabled      In
valid           Enabled            Disabled  N/A     Enabled Unknown Disabled -
```

If the status of the Incremental Resync option shows enabled, continue to the next step. If the status of the Incremental Resync option shows disabled, it must be enabled now before proceeding.

Example 7-7 shows the DSCLI commands (and resulting output) used to verify and enable incremental resync at the A volume set when it is required. After this is completed, continue to the next step.

Example 7-7 DSCLI command to verify and enable incremental resync for the A volume set when required

```
lspprc -l -dev ibm.2107-75xxxxx 7600-7601
ID        State       Reason Type         Out Of Sync Tracks Tgt Read Src Cascade Tgt Cascade Date
Suspended SourceLSS Timeout (secs) Critical Mode Fi
rst Pass Status Incremental Resync Tgt Write GMIR CG PPRC CG isTgtSE Pmir     DisableAutoResync
==========================================================================================7600:7800
Full Duplex -      Metro Mirror 0                 Disabled Disabled    Invalid     -              76
60             Disabled      In
valid           Disabled Disabled  N/A     Enabled Unknown Disabled -
7601:7801 Full Duplex -      Metro Mirror 0                 Disabled Disabled    Invalid     -
76        60             Disabled      In
valid           Disabled Disabled  N/A     Enabled Unknown Disabled -

mkpprc -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy -type mmir -mode nocp
-incrementalresync enablenoinit 7600-7601:7800-7801
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 7600:7800 successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 7601:7801 successfully created.
```

Next, the Metro Mirror relationship between the A volume set and the B volume set is temporarily halted. This ensures that no updates are made to the B or C volume sets during the migration process.

This process accomplishes several important objectives:

- Place the A volume set in a state where no host I/O can be directed to them, which can be Extended Long Busy or Queue Full.
- Remove the remote mirror and copy paths between the A and B volume sets.
- Create a remote mirror and copy consistency group.
- Ensure all updates made to the A volume set are tracked. These updates will be written to the B and C' volume sets after the C' volume set has been added to the replication environment.
- Reset the condition for the A volume set to allow the host I/O activity to resume.

This task is accomplished using the DSCLI **freeze** and **unfreeze** commands at the A volume set.

> **Important:** The **freezepprc** command causes a temporary disruption to host I/O processing. Ensure the potential impact to the environment, and how to use the **unfreezepprc** command, is understood before proceeding.

Example 7-8 shows the DSCLI commands (and resulting output) used to freeze and then unfreeze the volume sets. In this example, the command is being issued at the A volume set.

Example 7-8 DSCLI commands to freeze and unfreeze the A to B Metro Mirror relationship

```
freezepprc -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy 76:76
CMUC00161I freezepprc: Remote Mirror and Copy consistency group 76:76 successfully created.
```

```
unfreezepprc -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy 76:76
CMUC00198I unfreezepprc: Remote Mirror and Copy pair 76:76 successfully thawed.
```

After the Metro Mirror relationship between the A and B volume sets is suspended, the Out-of-Synch tracks status on the B, C, volume sets must be verified.

Example 7-9 shows the DSCLI commands (and resulting output) used to check the Out-of-Sync track status for the copy relationships. Each query must show zero Out-of-Sync tracks before proceeding with the next step.

Example 7-9 DSCLI Commands to verify Out-of -Sync track status for Global Copy, Metro Mirror, and Global Mirror

```
lspprc -l -dev ibm.2107-75xxxxx 7600-7601 Verify zero Out-of-Sync tracks at B
lspprc -l -dev ibm.2107-75xxxxx 7A00-7A01 Verify zero Out-of-Sync tracks at C

Verify zero Out-of-Sync tracks for GM session at B
dscli> showgmiroos -dev ibm.2107-75xxxxx -lss 76 -scope si 2
Scope           IBM.2107-75xxxxx
Session         02
OutOfSyncTracks 0
```

Important: If any of the Out-of-Sync track counts do not decrement to zero, data must still be written to the B' volume set. This must be resolved before proceeding.

Cut-over to the new volume set

To begin, remove the Global Mirror relationship between the B volume set and the C volume set. The process requires five steps:

1. Remove the B volume set from the Global Mirror session.
2. Close the GM session.
3. End the GM session.
4. Remove the Global Copy pairs between the B volume set and the C volume set.
5. Remove the replication paths between the B volume set and the C volume set.

Example 7-10 shows the DSCLI commands (and resulting output) used to complete this step. Some of these commands must be validated.

Example 7-10 DSCLI Commands to remove the Global Mirror relationship from B to C

```
chsession -dev ibm.2107-75xxxxx -lss 76 -action remove -volume 7600-7601 2
CMUC00147I chsession: Session 2 successfully modified.

rmsession -dev ibm.2107-75xxxxx -lss 76 2
CMUC00148W rmsession: Are you sure you want to close session 2? [y/n]:y
CMUC00146I rmsession: Session 2 closed successfully.

rmgmir -dev ibm.2107-75xxxxx -lss 76 -session 2
CMUC00166W rmgmir: Are you sure you want to stop the Global Mirror session 2:? [y/n]:y
CMUC00165I rmgmir: Global Mirror for session 2 successfully stopped.

rmpprc -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy 7600-7601:7A00-7A01
CMUC00160W rmpprc: Are you sure you want to delete the Remote Mirror and Copy volume pair relationship
7600-7604:7A00-7A04:? [y/n]:y
CMUC00155I rmpprc: Remote Mirror and Copy volume pair 7600:7A00 relationship successfully withdrawn.
CMUC00155I rmpprc: Remote Mirror and Copy volume pair 7601:7A01 relationship successfully withdrawn.
```

```
rmpprcpath -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy 76:7A
CMUC00152W rmpprcpath: Are you sure you want to remove the Remote Mirror and Copy path 76:7A:? [y/n]:y
CMUC00150I rmpprcpath: Remote Mirror and Copy path 76:7A successfully removed.
```

To complete the removal of the B volume set, the A to B Metro Mirror relationships must be removed. Do this at the B volume set to preserve the A Metro Mirror relationships that are used to track any changes made to that system during the migration.

> **Important:** This is a critical step. Do not remove the relationships at the A volume set. This must be preserved to track the updates using the incremental resync bitmap.

Example 7-11 shows the DSCLI commands (and resulting output) used to complete this step. This command requires the action to be validated.

Example 7-11 DSCLI Command to remove the Metro Mirror relationship at B only

```
rmpprc -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy -at tgt -unconditional 7600-7604
CMUC00160W rmpprc: Are you sure you want to delete the Remote Mirror and Copy volume pair
relationship 7600-7604:? [y/n]:y
CMUC00155I rmpprc: Remote Mirror and Copy volume pair :7600 relationship successfully withdrawn.
CMUC00155I rmpprc: Remote Mirror and Copy volume pair :7601 relationship successfully withdrawn.
```

After this process is completed, the B volume set will be removed from the MGM environment.

The Global Copy relationship between the C volume set and the B’ volume sets are now reversed. This is the first step in building the new Global Mirror relationship between B’ and C. This process requires two steps:

1. Fail over the Global Copy session between C and B' at the B’ volume set.
2. Fail back the Global Copy session between C and B' at the B’ volume set.

Example 7-12 shows the DSCLI commands (and resulting output) used to complete this step.

Example 7-12 DSCLI Commands to reverse the Global Copy relationship between C and B’

```
failoverpprc -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy -type gcp 7800-7801:7a00-7a01
CMUC00196I failoverpprc: Remote Mirror and Copy pair 7800:7A00 successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 7801:7A01 successfully reversed.

failbackpprc -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy -type gcp 7800-7801:7a00-7a01
CMUC00197I failbackpprc: Remote Mirror and Copy pair 7800:7A00 successfully failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 7801:7A01 successfully failed back.
```

> **Important:** This is the Go/No Go decision point for this migration scenario. Proceeding with the following step prohibits a return to the initial pre-migration state.
>
> For information about halting the migration process and returning to the original configuration, see “Return to the pre-migration configuration” on page 121

The Metro Mirror relationship between the A and B' volume sets is now created. This step ensures any updates that are made to the A volume set during the migration will be replicated to the B' and C volume sets. This process requires three steps:

1. Establish the replication paths between the A volume set and the B' volume set.
2. Create the Metro Mirror pairs between the A volume set and the B' volume set.
3. Reenable Incremental Resync on the A volume set.

Example 7-13 shows the DSCLI commands (and resulting output) used to complete this step.

> **Note:** The use of the **`-incrementalresync override`** option ensures the updates made to the A volume set during the migration are replicated to the B' and C. It also resets the `incrementalresync` flag on the A volume set, requiring it to be reenabled it as the last step in this process.

Example 7-13 DSCLI Commands to create the Metro Mirror relationship from A to B'

```
mkpprcpath -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy -remotewwnn 5005076305FFD75A
-srclss 76 -tgtlss 76 -consistgrp I0003:I0201 I0133:I0234
CMUC00149I mkpprcpath: Remote Mirror and Copy path 76:76 successfully established.

mkpprc -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy -type mmir -incrementalresync override
7600-7604:7800-7804
Date/Time: November 10, 2014 4:16:01 AM PST IBM DSCLI Version: 7.7.40.335 DS: IBM.2107-75xxxxx
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 7600:7800 successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 7601:7801 successfully created.

mkpprc -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy -type mmir -mode nocp
-incrementalresync enable 7600-7604:7800-7804
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 7600:7800 successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 7601:7801 successfully created.
```

Finally, the Global Mirror from the B' volume set to the C volume set is started. You do not need to establish the FlashCopy pairs required for the Global Mirror because they were not removed. This process requires three steps:

1. Open the GM session.
2. Add the volumes to the Global Mirror session.
3. Start the GM session.

Example 7-14 shows the DSCLI commands (and resulting output) used to complete this step.

Example 7-14 DSCLI Commands to create the Global Mirror relationship from B' to C

```
mksession -dev ibm.2107-75xxxxx -lss 76 2
CMUC00145I mksession: Session 2 opened successfully.

chsession -dev ibm.2107-75xxxxx -lss 76 -action add -volume 7600-7604 2
CMUC00147I chsession: Session 2 successfully modified.

mkgmir -dev ibm.2107-75xxxxx -lss 76 -session 2
CMUC00162I mkgmir: Global Mirror for session 2 successfully started.
```

The MGM environment is now rebuilt with the new B' volume set. The storage migration process for this site is now completed.

Return to the pre-migration configuration

The process below returns the environment to the pre-migration configuration for this scenario. The specific process should be documented, tested, and included in the master migration plan.

> **Attention:** The `-dev` and `-remotedev` statements in the DSCLI commands must be changed to refer to the appropriate volume set.

Remove the Global Mirror relationship between the B' and C volume sets as described in Example 7-10 on page 118.

Reverse the Global Copy relationship between the B' and C volume sets as described in Example 7-12 on page 119.

Rebuild the Global Mirror relationship between the B and C volume sets as described in Example 7-4 on page 112.

Rebuild the Metro Mirror relationship between the A and B volume sets as described in Example 7-13 on page 120.

The pre-migration MGM environment has now been restored.

7.3.3 Replacing the A volume set without automation

This section describes the process to replace the A volume set in an existing MGM environment. The process requires an interruption to the host processing at the primary site. The MGM environment remains intact during the initial copy process. Your specific migration might require more steps, depending on the environment.

The migration process is broken down into five major steps:

1. Initial setup of the migration environment
2. Data validation testing
3. Preparing to migrate
4. Cut over to the new volume set
5. Recovering the pre-migration environment

The following steps assume that the new volume set is installed and configured as required. The cascade option for Global Copy is used to complete the initial data migration to the new volume set. DS8000 DSCLI commands are used to build the migration environment and cut over to the new volume set.

> **Important:** The potential for human error is significantly increased if you do not use automation software to help the migration process. Therefore, have the DSCLI command set to be used in the migration process reviewed and validated by the local IBM representative for correct syntax and sequence before setting up the migration environment.
>
> If this is a GDPS managed environment, see 7.4, "Geographically Dispersed Parallel Sysplex (GDPS) considerations" on page 137.

Initial setup of the migration environment

The initial setup of the cascaded replication environment is described in detail in 2.3, “Data replication” on page 22. See that section to complete this process and then return here after that is accomplished. For this migration scenario, the cascaded Global Copy is extended to an additional volume set. The migration path will be from A → B → C → A’.

> **Note:** Consult an IBM representative during the planning process of the migration project to validate the replication network requirements. This helps avoid inadvertently impacting the production and replication environments during the data migration.

Figure 7-6 illustrates the basic layout of the migration configuration. In addition to the existing MGM relationship between A, B, and C, an asynchronous Global Copy leg is added from the current C volume set to the new A’ volume set. This relationship allows the data from the C volume set to be copied to the A’ volume set without affecting the production applications or the existing MGM environment.

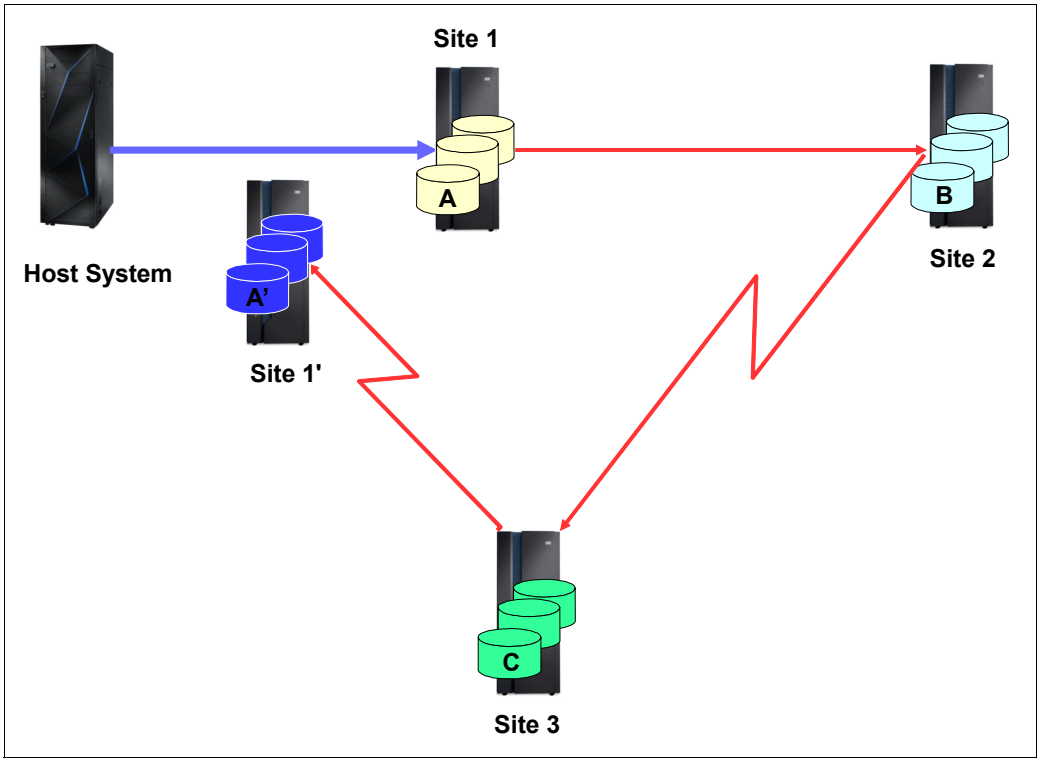

Figure 7-6 Cascaded Global Copy with existing 3-Site Metro-Global Mirror Environment

Figure 7-7 illustrates the post-migration environment with A' as the new primary device in the MGM environment.

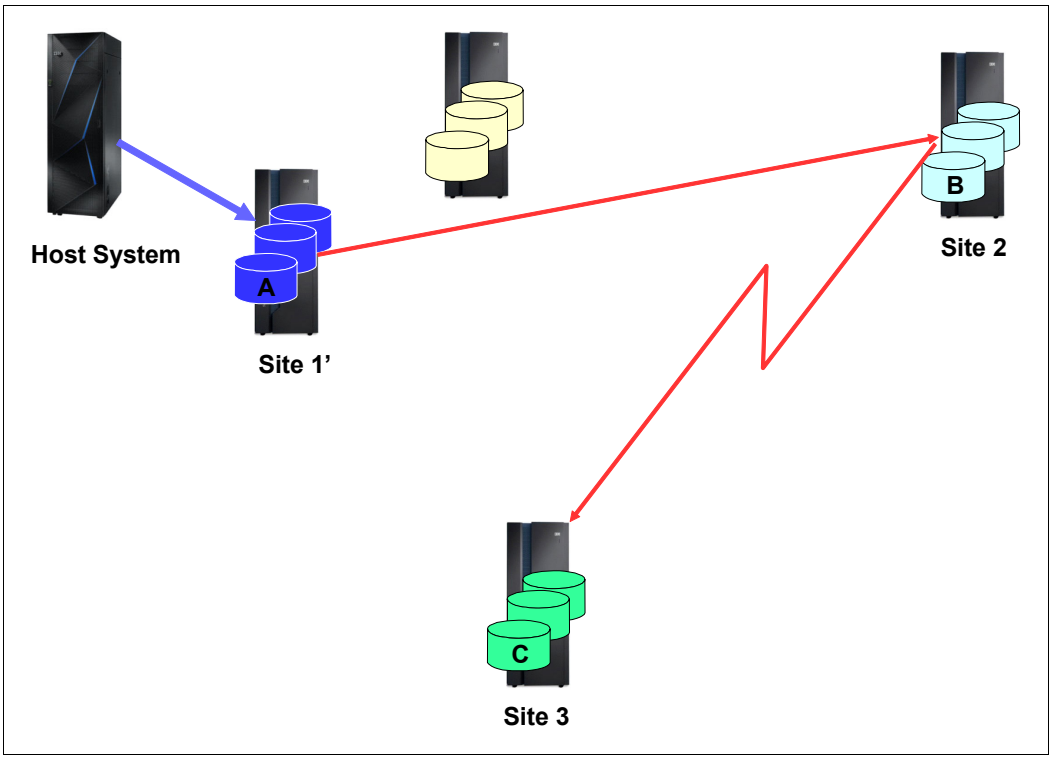

Figure 7-7 Post migration 3-Site Metro-Global Mirror Environment

Data validation testing

An initial data validation test should be performed on the A' volume set before proceeding with the final cut over.

This process can be accomplished with minimum impact to the current environment, allowing a quick fall back to the pre-migration configuration if necessary. For more information, see 2.4, "Data validation testing" on page 24.

If the validation testing is successful, proceed with the migration process. If the data validation process is not successful, determine why the validation testing failed before proceeding.

Prepare to migrate to the new volume set

This process requires an outage to the host processing for the volume sets affected by the migration. By using the cascaded Global Copy to migrate the data before the transition, the time required to fully recover the environment is minimized.

The first step in this process is to stop all I/O to the current primary volume set. There are two standard options to accomplish this:

- Shut down all host systems that are attached to the A volume set. This is the best choice to ensure a clean transition as there are no residual host relationships with the volumes on that system.
- Quiesce all host systems that are attached to the A volume set. This is a workable alternative, but additional steps might be required to integrate the Site 1' volume set in place of the A volume set. These steps can include, but are not limited to, importing the

volumes defined on the A' volume set or removing the A volume set and its volumes from the host configuration.

After all I/O is halted, it is critical to validate that all updated data has been written successfully to the new A' volume set. This can be accomplished by verifying the Out-of-Synch track status for all of the existing copy relationships at A, B, and C. There must be zero Out-of-Sync tracks for each session before it is safe to proceed to the next step.

Example 7-15 shows the DSCLI commands (and resulting output) used to check the Out-of-Sync track status for all of the copy relationships in the cascaded MGM environment. This includes each Metro Mirror, Global Copy, and Global Mirror relationship. Each query must show zero Out-of-Sync tracks before proceeding with the next step.

Example 7-15 DSCLI Commands to verify Out-of Sync track status for Global Copy, Metro Mirror, and Global Mirror

```
lspprc -l -dev ibm.2107-75xxxxx 7600-7601 Verify zero Out-of-Sync tracks at A
lspprc -l -dev ibm.2107-75xxxxx 7600-7601 Verify zero Out-of-Sync tracks at B
lspprc -l -dev ibm.2107-75xxxxx 7A00-7A01 Verify zero Out-of-Sync tracks at C

Verify zero Out-of-Sync tracks for GM session at B
dscli> showgmiroos -dev ibm.2107-75xxxxx -lss 76 -scope si 2
Scope           IBM.2107-75xxxxx
Session         02
OutOfSyncTracks 0
```

Important: If any of the Out-of-Sync track counts do not decrement to zero, data must still be written to the A' volume set. This situation must be resolved before proceeding.

To complete the integration of the A' volume set into the environment, install and validate all physical cabling required, including the host links from the production systems, and the replication links to the B volume set. The A volume set should also be removed from the environment now to ensure that no data are written to this volume set. After the physical connectivity is completed and validated, the MGM environment can be rebuilt with the A' volume set as the new primary volume set using the following process.

Cut-over to the new volume set

Begin by removing the Global Copy relationship between the C volume set and the A' volume sets. This relationship is no longer required to complete the migration. This process requires two steps:

1. Remove the Global Copy pairs between the C volume set and the A' volume set.
2. Remove the replication paths between the C volume set and the A' volume set.

Example 7-16 shows the DSCLI commands (and resulting output) used to complete this step. Note that each command requires the action to be validated.

Example 7-16 DSCLI Commands to remove the Global Copy relationship from C to A'

```
rmpprc -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy 7A00-7A01:7800-7801
CMUC00160W rmpprc: Are you sure you want to delete the Remote Mirror and Copy volume pair relationship
7a00-7a04:7800-78
04:? [y/n]:y
CMUC00155I rmpprc: Remote Mirror and Copy volume pair 7A00:7800 relationship successfully withdrawn.
CMUC00155I rmpprc: Remote Mirror and Copy volume pair 7A01:7801 relationship successfully withdrawn.

rmpprcpath -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy 7A:78
```

```
CMUC00152W rmpprcpath: Are you sure you want to remove the Remote Mirror and Copy path 7A:78:? [y/n]:y
CMUC00150I rmpprcpath: Remote Mirror and Copy path 7A:78 successfully removed.
```

Next, remove the Metro Mirror relationship between the A volume set and the B volume set. After the completion of this stage, the A volume set will be removed from the MGM environment. This process requires two steps:

1. Remove the Metro Mirror pairs between the A volume set and the B volume set.
2. Remove the replication paths between the A volume set and the B volume set.

Example 7-17 shows the DSCLI commands (and resulting output) used to complete this step. Note that each command requires the action to be validated.

Example 7-17 DSCLI Commands to remove the Metro Mirror relationship from A to B

```
rmpprc -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy 7600-7604:7600-7604
CMUC00160W rmpprc: Are you sure you want to delete the Remote Mirror and Copy volume pair relationship
7600-7604:7600-76
04:? [y/n]:y
CMUC00155I rmpprc: Remote Mirror and Copy volume pair 7600:7600 relationship successfully withdrawn.
CMUC00155I rmpprc: Remote Mirror and Copy volume pair 7601:7601 relationship successfully withdrawn.

rmpprcpath -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy 76:76
CMUC00152W rmpprcpath: Are you sure you want to remove the Remote Mirror and Copy path 76:76:? [y/n]:y
CMUC00150I rmpprcpath: Remote Mirror and Copy path 76:76 successfully removed.
```

Finally, establish the replication paths and create the Metro Mirror pairs from the A' volume set to the B volume set. The A' volume set is the primary in the MGM environment so the required options with the **mkpprc** command for that specific role are used. Those options are **-type mmir** and **-incrementalresync enable**

Example 7-18 shows the DSCLI commands (and resulting output) used to complete this step. Note the inclusion of the **-consistgrp** option with the **mkpprcpath** command. This option creates a consistency group for the remote mirror and copy volume pairs being defined to this Metro Mirror session.

Example 7-18 DSCLI Commands to create the Metro Mirror relationship from A' to B

```
mkpprcpath -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy -remotewwnn 5005076305FFD75A
-srclss 78 -tgtlss 76 -consistgrp I0003:I0201 I0133:I0234
CMUC00149I mkpprcpath: Remote Mirror and Copy path 78:76 successfully established.

mkpprc -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy -type mmir -mode nocp
-incrementalresync enable 7800-7801:7600-7601
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 7800:7600 successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 7801:7601 successfully created.
```

Important: This is the Go/No Go decision point for this migration scenario. Proceeding with the following step prohibits a return to the initial pre-migration state.

To halt the migration process and return to the original configuration, see "Return to the pre-migration configuration" on page 126.

The MGM environment has now been rebuilt with the new A' volume set. The storage migration process for this site is now completed. The remaining steps in the process are related to IPLing or resuming the host systems to resume normal operations. Refer to your migration plan to continue with these steps.

Return to the pre-migration configuration

The process below returns the environment to the pre-migration configuration for this scenario. The specific process should be documented, tested, and included in the master migration plan.

> **Attention:** The `-dev` and `-remotedev` statements in the DSCLI commands need to be changed to refer to the appropriate volume set.

Remove the Metro Mirror relationship between the A and B' volume sets as described in Example 7-17 on page 125.

Rebuild the Metro Mirror relationship between the A and B volume sets as described in Example 7-18 on page 125.

Return all host cabling to the A volume set as required.

The pre-migration MGM environment has now been restored. You can resume normal operations.

7.3.4 Replacing the A and B volume set without automation

This section describes the process to simultaneously replace the A and B volume sets in an existing MGM environment. The process requires an interruption to the host processing at the primary site. The MGM environment remains intact during the initial copy process. Your specific migration might require more steps, depending on the environment.

The migration process is broken down into five major steps:

1. Initial setup of the migration environment
2. Data validation testing
3. Preparing to migrate
4. Cut-over to the new volume set
5. Recovering the pre-migration environment

The steps assume that the new volume sets are installed and configured as required. The cascade option for Global Copy is used to complete the initial data migration to the new volume set. DS8000 DSCLI commands are used to build the migration environment and cut over to the new volume set.

> **Important:** The potential for human error is significantly increased if there is no automation software to help the migration process. The DSCLI command set to be used in the migration process should be reviewed and validated by the local IBM representative for correct syntax and sequence before setting up the migration environment.
>
> If this is a GDPS managed environment, see 7.4, "Geographically Dispersed Parallel Sysplex (GDPS) considerations" on page 137.

Initial setup of the migration environment

The initial setup of the cascaded replication environment is described in 2.3, "Data replication" on page 22. Refer to this section to complete this process and then return here after that is accomplished. For this migration scenario, the cascaded Global Copy is extended to two additional volume sets. The migration path is from A → B → C → A' → B'.

Note: Consult an IBM representative during the planning process of the migration project to validate the replication network requirements. This helps avoid inadvertently impacting the production and replication environments during the data migration.

Figure 7-8 illustrates the basic layout of the migration configuration. In addition to the existing MGM relationship between A, B, and C, two asynchronous Global Copy legs from the current C volume set to the new A' and B' volume sets are added. This relationship allows the data from the C volume set to be copied to the A' and B' volume sets without impact to the production applications or the existing MGM environment.

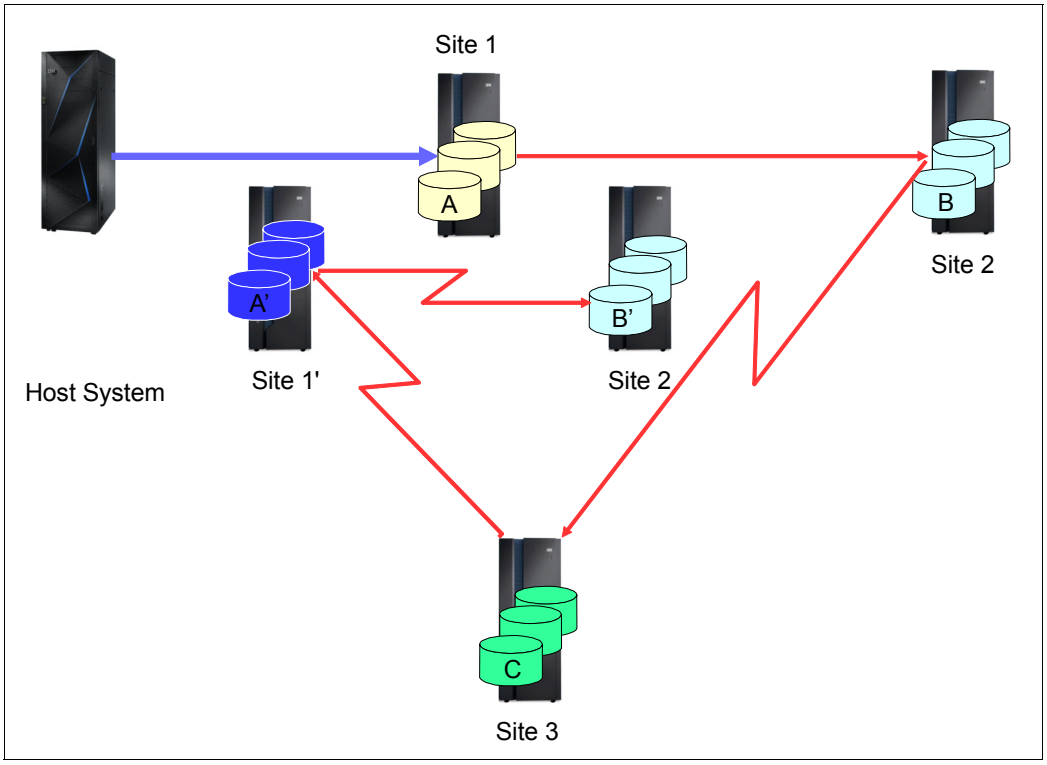

Figure 7-8 Multiple cascaded Global Copy with existing 3-Site Metro-Global Mirror Environment

Figure 7-9 illustrates the post-migration environment with A’ and B’ as the new primary and secondary in the MGM environment.

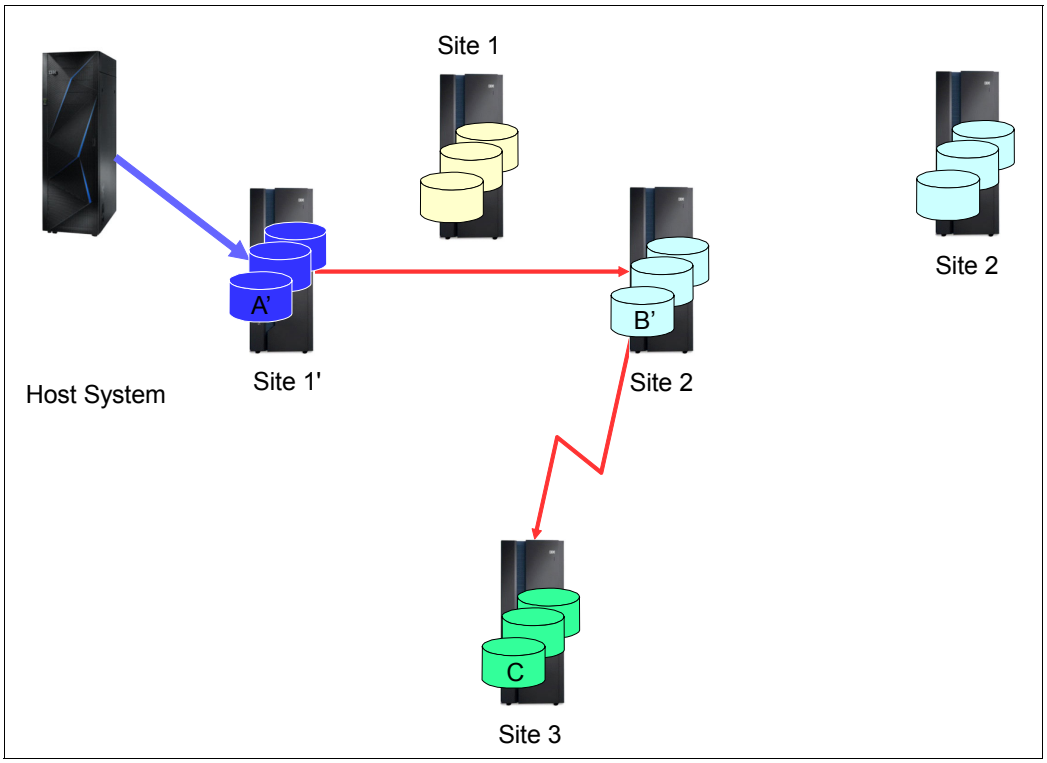

Figure 7-9 Post migration 3-Site Metro-Global Mirror Environment

Data validation testing

It is good practice to run an initial data validation test on the A’ volume set before proceeding with the final cut over.

This process can be accomplished with minimum impact to the current environment, allowing a quick fall back to the pre-migration configuration if necessary. For more information, see 2.4, “Data validation testing” on page 24.

If the validation testing is successful, proceed with the migration process. If the data validation process is not successful, determine why the validation testing has failed before proceeding.

Prepare to migrate to the new volume set

This process requires an outage to the host processing for the volume sets affected by the migration. By using the cascaded Global Copy to migrate the data before the transition, the time required to fully recover the environment is minimized.

The first step in this process is to stop all I/O to the current primary volume set. There are two standard options to accomplish this:

- Shut down all host systems that are attached to the A volume set. This is the best alternative to ensure a clean transition as there are no residual host relationships with the volumes on that system.
- Quiesce all host systems that are attached to the A volume set. This is a workable alternative but extra steps might be required to integrate the Site 1’ volume set in place of the A volume set. These steps might include, but are not limited to, importing the volumes

defined on the A' volume set or removing the A volume set and its volumes from the host configuration.

After all I/O is halted, validate that all updated data has been written successfully to the new A' and B' volume sets. This can be accomplished by verifying the Out-of-Sync track status for all of the existing copy relationships at A, B, C, and A'. There must be no Out-of-Sync tracks for each session before it is safe to proceed to the next step.

Example 7-19 shows the DSCLI commands (and resulting output) used to check the Out-of-Sync track status for all of the copy relationships in the cascaded MGM environment. This includes each Metro Mirror, Global Copy, and Global Mirror relationship. Each query must show zero Out-of-Sync tracks before proceeding with the next step.

Example 7-19 DSCLI Commands to verify Out-of-Sync track status for Global Copy, Metro Mirror, and Global Mirror

```
lspprc -l -dev ibm.2107-75xxxxx 7600-7601 Verify zero Out-of-Sync tracks at A
lspprc -l -dev ibm.2107-75xxxxx 7600-7601 Verify zero Out-of-Sync tracks at B
lspprc -l -dev ibm.2107-75xxxxx 7A00-7A01 Verify zero Out-of-Sync tracks at C

Verify zero Out-of-Sync tracks for GM session at B
dscli> showgmiroos -dev ibm.2107-75xxxxx -lss 76 -scope si 2
Scope           IBM.2107-75xxxxx
Session         02
OutOfSyncTracks 0
```

Important: If any of the Out-of-Sync track counts do not decrement to zero, data must still be written to the A' volume set. This situation must be resolved before proceeding.

Cut-over to the new volume set

To complete the integration of the A' and B' volume sets into the environment, install and validate all physical cabling required, including the host links from the production systems to the A' volume set and the replication links from the B' volume set to the C volume set. The A volume set should also be removed from the environment now to ensure that no data is written to this volume set.

Before rebuilding the MGM environment with the new A' and B' volume sets, remove the copy relationships to C that are no longer required. This process requires seven steps as shown in the following two examples:

1. Remove the B volume set from the Global Mirror session.
2. Close the GM session.
3. End the GM session.
4. Remove the Global Copy pairs between the B volume set and the C volume set.
5. Remove the replication paths between the B volume set and the C volume set.

   Example 7-20 shows the DSCLI commands (and resulting output) used to complete the first five steps. Note that some commands require the action to be validated.

Example 7-20 DSCLI Commands to remove the unnecessary copy relationship from B to C

```
chsession -dev ibm.2107-75xxxxx -lss 76 -action remove -volume 7600-7601 2
CMUC00147I chsession: Session 2 successfully modified.

rmsession -dev ibm.2107-75xxxxx -lss 76 2
CMUC00148W rmsession: Are you sure you want to close session 2? [y/n]:y
CMUC00146I rmsession: Session 2 closed successfully.
```

```
rmgmir -dev ibm.2107-75xxxxx -lss 76 -session 2
CMUC00166W rmgmir: Are you sure you want to stop the Global Mirror session 2:? [y/n]:y
CMUC00165I rmgmir: Global Mirror for session 2 successfully stopped.

rmpprc -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy 7600-7601:7A00-7A01
CMUC00160W rmpprc: Are you sure you want to delete the Remote Mirror and Copy volume pair relationship
7600-7604:7A00-7A04:? [y/n]:y
CMUC00155I rmpprc: Remote Mirror and Copy volume pair 7600:7A00 relationship successfully withdrawn.
CMUC00155I rmpprc: Remote Mirror and Copy volume pair 7601:7A01 relationship successfully withdrawn.

rmpprcpath -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy 76:7A
CMUC00152W rmpprcpath: Are you sure you want to remove the Remote Mirror and Copy path 76:7A:? [y/n]:y
CMUC00150I rmpprcpath: Remote Mirror and Copy path 76:7A successfully removed.
```

6. Remove the Global Copy pairs between the C volume set and the A’ volume set.
7. Remove the replication paths between the C volume set and the A’ volume set.

   Example 7-21 shows the DSCLI commands (and resulting output) used to complete the last two steps. Note that some commands require the action to be validated.

Example 7-21 DSCLI Commands to remove the unnecessary copy relationship from A to C

```
rmpprc -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy 7A00-7A01:7800-7801
CMUC00160W rmpprc: Are you sure you want to delete the Remote Mirror and Copy volume pair relationship
7a00-7a04:7800-78
04:? [y/n]:y
CMUC00155I rmpprc: Remote Mirror and Copy volume pair 7A00:7800 relationship successfully withdrawn.
CMUC00155I rmpprc: Remote Mirror and Copy volume pair 7A01:7801 relationship successfully withdrawn.

rmpprcpath -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy 7A:78
CMUC00152W rmpprcpath: Are you sure you want to remove the Remote Mirror and Copy path 7A:78:? [y/n]:y
CMUC00150I rmpprcpath: Remote Mirror and Copy path 7A:78 successfully removed.
```

The MGM environment can now be rebuilt using the new A’ and B’ volume sets.

Begin by creating the B’ to C Global Copy relationship. Use the **-cascade** option with the **mkpprc** command as required for the B’ volume set in an MGM environment. This step assumes that the required physical cabling has already been installed and tested.

Example 7-22 shows the DSCLI commands (and resulting output) used to complete this step.

Example 7-22 DSCLI Commands to create the Global Copy relationship from B’ to C

```
mkpprcpath -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy -remotewwnn 5005076305FFD71A
-srclss 78 -tgtlss 7A I0201:I0032 I0234:I0306
CMUC00149I mkpprcpath: Remote Mirror and Copy path 78:7A successfully established.

mkpprc -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy -type gcp -mode nocp -cascade
7800-7801:7a00-7a01
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 7600:7A00 successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 7601:7A01 successfully created.
```

Next, convert the A’ to B’ Global Copy relationship to a synchronous Metro Mirror relationship. The A’ volume set will be the primary in the MGM environment, so the required options with the **mkpprc** command for that specific role are used. Those options are **-type mmir** and **-incrementalresync enable.**

Example 7-23 shows the DSCLI commands (and resulting output) used to complete this step.

Example 7-23 DSCLI Commands to convert the Global Copy relationship to a Metro Mirror relationship

```
mkpprc -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy -type mmir -mode nocp
-incrementalresync enable 7800-7801:7800-7801
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 7800:7800 successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 7801:7801 successfully created.
```

Finally, the Global Mirror from the B' volume set to the C volume set is created. There is no need to establish the FlashCopy pairs that are required for the Global Mirror because they were not removed. This process requires three steps:

1. Open the GM session.
2. Add the volumes to the Global Mirror session.
3. Start the GM session.

Example 7-24 shows the DSCLI commands (and resulting output) used to complete this step.

Example 7-24 DSCLI Commands to create the Global Mirror relationship from B' to C

```
mksession -dev ibm.2107-75xxxxx -lss 78 2
CMUC00145I mksession: Session 2 opened successfully.

chsession -dev ibm.2107-75xxxxx -lss 78 -action add -volume 7800-7801 2
CMUC00147I chsession: Session 2 successfully modified.

mkgmir -dev ibm.2107-75xxxxx -lss 78 -session 2
CMUC00162I mkgmir: Global Mirror for session 2 successfully started.
```

Important: This is the Go/No Go decision point for this migration scenario. Proceeding with the following step prohibits a return to the initial pre-migration state.

To halt the migration process and return to the original configuration, see "Return to the pre-migration configuration" on page 131.

The MGM environment has now been rebuilt with the new A' and B' volume sets. The storage migration process for this site is now completed. The remaining steps in the process are related to IPLing or resuming the host systems to resume normal operations. Refer to your migration plan to continue with these steps.

Return to the pre-migration configuration

The process below returns the environment to the pre-migration configuration for this scenario. The specific process should be documented, tested, and included in the master migration plan.

Attention: The `-dev` and `-remotedev` statements in the DSCLI commands must be changed to refer to the appropriate volume set.

Remove the Global Mirror relationship between the B' and C volume sets as described in Example 7-20 on page 129.

Rebuild the Global Copy relationship between the B and C volume sets as described in Example 7-22 on page 130.

Rebuild the Metro Mirror relationship between the A and B volume sets as described in Example 7-23 on page 131.

Rebuild the Global Mirror relationship between the B and C volume sets as described in Example 7-24 on page 131.

Return all host cabling to the A volume set as required.

The pre-migration MGM environment has now been restored so you can resume normal operations.

7.3.5 Replace all volume sets without automation

This section describes the process to simultaneously replace all volume sets in an existing MGM environment. The process requires an interruption to the host processing at the primary site. The MGM environment remains intact during the initial copy process. Your specific migration might require more steps, depending on the environment.

The migration process is broken down into five major steps:

1. Initial setup of the migration environment
2. Data validation testing
3. Preparing to migrate
4. Cut over to the new volume set
5. Recovering the pre-migration environment

The following steps assume that the new volume sets are installed and configured as required. The cascade option for Global Copy is used to complete the initial data migration to the new volume set. DS8000 DSCLI commands are used to build the migration environment and cut over to the new volume set.

Important: The potential for human error is significantly increased if you do not use automation software to help during the migration process. Therefore, the DSCLI command set to be used in the migration process should be reviewed and validated by the local IBM representative for correct syntax and sequence before setting up the migration environment.

If this is a GDPS managed environment, see 7.4, “Geographically Dispersed Parallel Sysplex (GDPS) considerations” on page 137.

Initial setup of the migration environment

The initial setup of the cascaded replication environment is described in 2.3, “Data replication” on page 22. See this section to complete this process and then return here after that is accomplished. For this migration scenario, the cascaded Global Copy is extended to three additional volume sets. The migration path will be from A → B → C → A’ → B → C’.

Note: Consult with an IBM representative during the planning process of the migration project to validate the replication network requirements. This helps avoid inadvertently affecting the production and replication environments during the data migration.

Figure 7-10 illustrates the basic layout of the migration configuration. In addition to the existing MGM relationship between A, B, and C, three asynchronous Global Copy legs from the current C volume set to the new A’, B’, and C’ volume sets have been added. This relationship allows the data from the C volume set to be copied to the A’, B’, and C’ volume sets without affecting the production applications or the existing MGM environment.

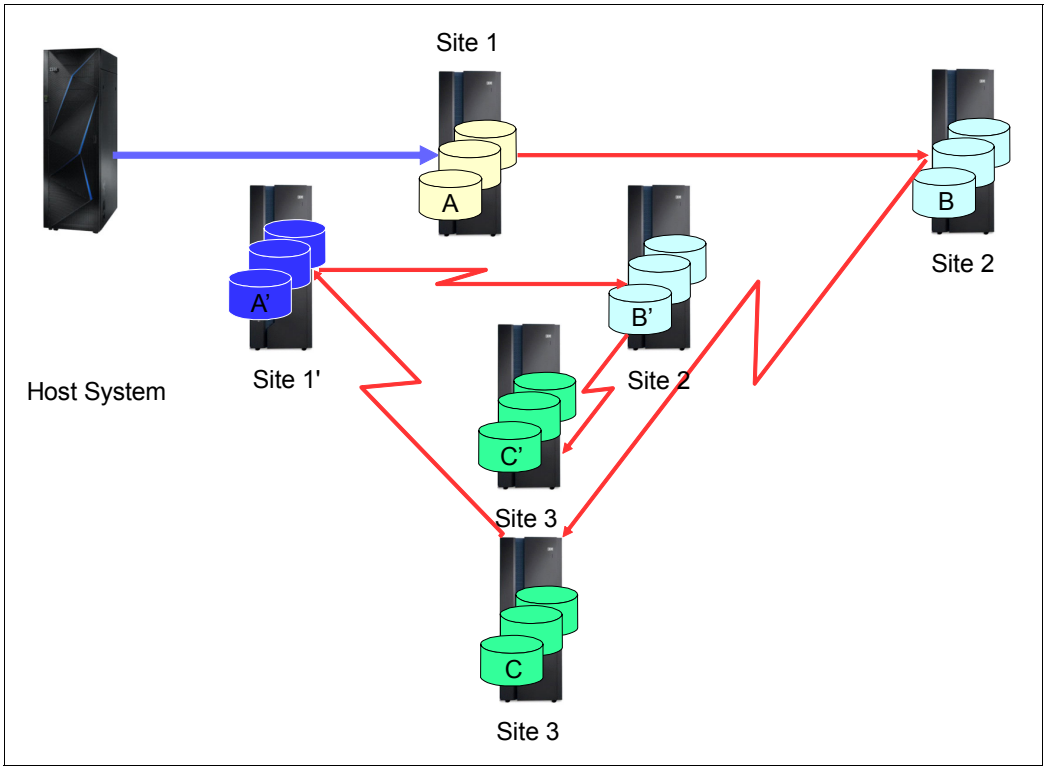

Figure 7-10 Multiple cascaded Global Copy with existing 3-Site Metro-Global Mirror Environment

Figure 7-11 illustrates the post-migration environment with A’, B’, and C’ as the new members of the MGM environment.

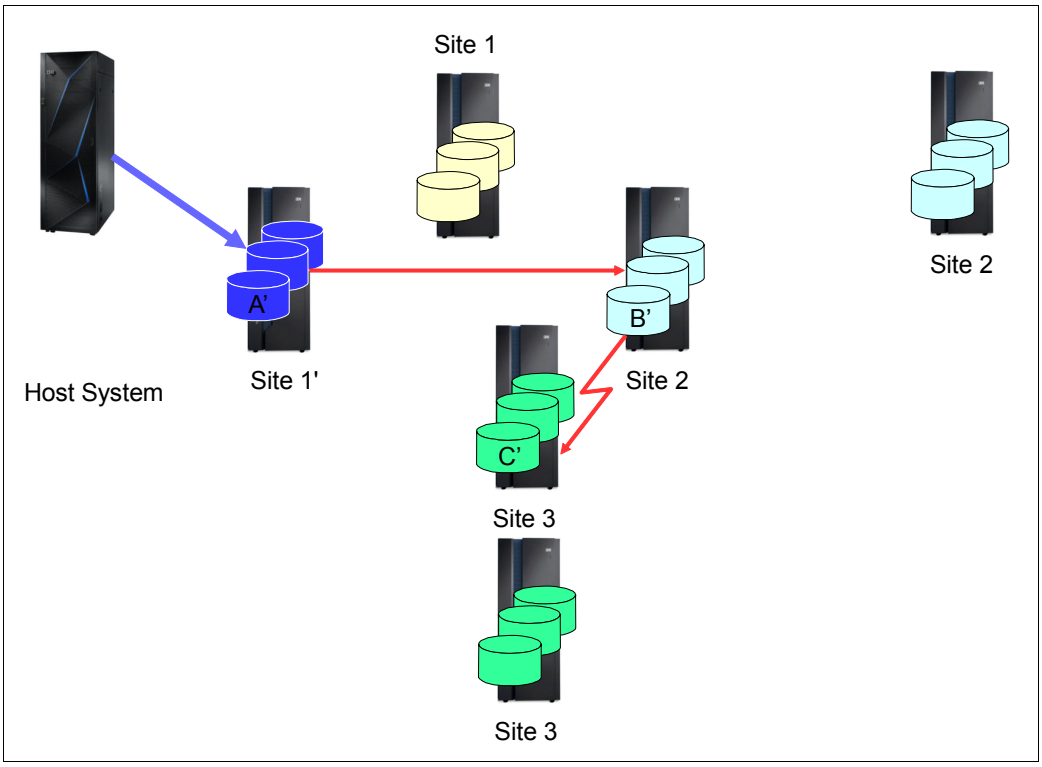

Figure 7-11 Post migration 3-Site Metro-Global Mirror Environment

Data validation testing

An initial data validation test should be run on the A’ volume set before proceeding with the final cut over.

This process can be accomplished with minimum impact to the current environment, allowing a quick fall back to the pre-migration configuration if necessary. For more information, see 2.4, “Data validation testing” on page 24.

If the validation testing is successful, proceed with the migration process. If the data validation process is not successful, determine why the validation testing failed before proceeding.

Prepare to migrate to the new volume sets

This process requires an outage to the host processing for the volume sets affected by the migration. By using the cascaded Global Copy to migrate the data before the transition, the time that is required to fully recover the environment is minimized

The first step in this process is to stop all I/O to the current primary volume set. There are two standard options to accomplish this:

- Shut down all host systems that are attached to the A volume set. This is the best alternative to ensure a clean transition because there are no residual host relationships with the volumes on that system.
- Quiesce all host systems that are attached to the A volume set. This is a workable alternative, but extra steps might be required to integrate the Site 1’ volume set in place of the A volume set. These steps can include, but are not limited to, importing the volumes

defined on the A' volume set or removing the A volume set and its volumes from the host configuration.

After all I/O is halted, validate that all updated data has been written successfully to the new A' and B' volume sets. This can be accomplished by verifying the Out-of-Sync track status for all of the existing copy relationships at A, B, C, and A'. There must be zero Out-of-Sync tracks for each session before it is safe to proceed to the next step.

Example 7-25 shows the DSCLI commands (and resulting output) used to check the Out-of-Sync track status for all of the copy relationships in the cascaded MGM environment. This includes each Metro Mirror, Global Copy, and Global Mirror relationship. Each query must show zero Out-of-Sync tracks before proceeding with the next step.

Example 7-25 DSCLI Commands to verify Out-of-Sync track status for Global Copy, Metro Mirror, and Global Mirror

```
lspprc -l -dev ibm.2107-75xxxxx 7600-7601 Verify zero Out-of-Sync tracks at A
lspprc -l -dev ibm.2107-75xxxxx 7600-7601 Verify zero Out-of-Sync tracks at B
lspprc -l -dev ibm.2107-75xxxxx 7A00-7A01 Verify zero Out-of-Sync tracks at C

Verify zero Out-of-Sync tracks for GM session at B
dscli> showgmiroos -dev ibm.2107-75xxxxx -lss 76 -scope si 2
Scope           IBM.2107-75xxxxx
Session         02
OutOfSyncTracks 0
```

Important: If any of the Out-of-Sync track counts do not decrement to zero, data must still be written to the A' volume set. This situation must be resolved before proceeding.

Cut-over to the new volume set

To complete the integration of the A', B', and C' volume sets into the environment, install and validate all physical cabling that is required, including the host links from the production systems to the new A' volume set. The A volume set should also be removed from the environment now to ensure no data is being written to this volume set.

Next, create the MGM environment on the new A', B', and C' volume sets. The benefit of creating the cascaded environment in the earlier steps is now revealed. The replication paths are already established in the correct direction for each of the three volume sets.

Start by converting the A' to B' Global Copy relationship to a synchronous Metro Mirror relationship. The A' volume set will be the primary in the MGM environment, so the required options with the **mkpprc** command for that specific role are used. Those options are **-type mmir** and **-incrementalresync enable.**

Example 7-26 shows the DSCLI commands (and resulting output) used to complete this step.

Example 7-26 DSCLI Commands to convert the Global Copy relationship to a Metro Mirror relationship

```
mkpprc -dev ibm.2107-75xxxxx -remotedev ibm.2107-75yyyyy -type mmir -mode nocp
-incrementalresync enable 7800-7801:7800-7801
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 7800:7800 successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 7801:7801 successfully created.
```

Next, create the Global Mirror from the B' volume set to the C' volume set. This process requires four steps:

Note: Only step 1 is performed at the C' volume set. All other commands are performed at the B volume set.

1. Establish the Global Mirror FlashCopy relationships for the C' volume set.
2. Open the GM session.
3. Add the volumes to the Global Mirror session.
4. Start the GM session.

Example 7-27 shows the DSCLI commands (and resulting output) used to complete this step.

Example 7-27 DSCLI Commands to create the Global Mirror relationship from B' to C'

```
mkflash -dev IBM.2107-75xxxxx -tgtinhibit -record -persist -nocp 7800-7801:C800-C801
CMUC00137I mkflash: FlashCopy pair 7800:C800 successfully created.
CMUC00137I mkflash: FlashCopy pair 7801:C801 successfully created.


mksession -dev ibm.2107-75xxxxx -lss 78 2
CMUC00145I mksession: Session 2 opened successfully.

chsession -dev ibm.2107-75xxxxx -lss 78 -action add -volume 7800-7801 2
CMUC00147I chsession: Session 2 successfully modified.

mkgmir -dev ibm.2107-75xxxxx -lss 78 -session 2
CMUC00162I mkgmir: Global Mirror for session 2 successfully started.
```

Important: This is the Go/No Go decision point for this migration scenario. Proceeding with the following step prohibits a return to the initial pre-migration state.

To halt the migration process and return to the original configuration, see "Return to the pre-migration configuration" on page 136.

The MGM environment has now been rebuilt with the new A', B', and C' volume sets. The storage migration process for this site is now completed. The remaining steps in the process are related to IPLing or resuming the host systems to resume normal operations. Refer to the master migration plan to continue with these steps.

Return to the pre-migration configuration

The process below returns the environment to the pre-migration configuration for this scenario. The specific process should be documented, tested, and included in the master migration plan.

Return all host cabling to the A volume set as required.

The pre-migration MGM environment is still intact, so you do not need to rebuild any of the relationships to recover in this scenario.

Resume normal operations.

7.4 Geographically Dispersed Parallel Sysplex (GDPS) considerations

In addition to the general considerations discussed in Chapter 1, “Introducing disk data migration” on page 1, some considerations should to be taken into account when replacing volume sets in an MGM environment with GDPS managing the environment. Copy Services Incremental Resynchronization (IR) is required to replace the A and B volume sets in the following examples. GDPS can be used for some steps in the migration process. Using GDPS and IR allow the applications systems to remain active, impact to the current recovery point objective (RPO) minimized and HyperSwap is only disabled for a short period.

To minimize the number of configuration changes required, the recommended approach is to configure the new volume sets as symmetrical as the one being removed. This means, if at all possible, use the same LSS range, SSID range, UCB range, and PPRC links for the new volume set that was in use on the old volume set.

For all migrations, use the cascade option for Global Copy to complete the initial data migration to the new volume set.

If Remote Pair FlashCopy (RPFC) is in use, the FlashCopy Target volumes cannot be in a cascaded global copy relationship. The FlashCopy must be withdrawn before establishing the cascaded relationship.

In a GDPS/MGM environment, there are at least 3 and in most cases 4 GDPS controlling systems. This document refers to these controlling systems:

- Site 1 Kp-sys, which runs on non-PPRCed disk on a Site 1 server (optional, but highly recommended).
- Site2 Kp-sys in Site2 running on a Site2 server from local disk in Site2.
- Kg-sys, which is the GDPS/GM local controlling system. This Kg-sys has a few different options as to its placement that are described in the *GDPS/MGM 3.11 Planning and Installation Guide,* ZG24-6726-13. In this guide, the Kg-sys is running on non-PPRCed disk embedded in one of the Kp-systems (under a second IBM NetView® instance, separate from the Tivoli NetView for z/OS running the GDPS/PPRC code) as per the MGM IR configurations.
- Kr-sys is the GDPS/GM remote controlling system. The Kr-sys is in the R-site (remote) on a stand-alone system.

Important: Special consideration needs to be taken into account when migrating a volume set that includes the GDPS Kx-systems. For instance, if the Site 1 Kp-sys system volumes are on a A volume set you are migrating, a Global Copy relationship needs to be created from A to A’ for these devices. The system will need to be shut down, the relationships deleted, wait for the copy to complete, and then IPL this Kp-sys on the new A’ volume set after it is made available. The same is true for the Site 2 Kp-sys, Kg-sys, and Kr-sys.

Initial configuration

Figure 7-1 on page 106 illustrates the basic layout of a 3-site Metro Global Mirror. The following migrations describe an environment with Incremental Resynchronization (MGM IR) configuration. The host system accesses the primary volume set in site 1. There is a synchronous Metro Mirror (MM) relationship to the site 2 volume set, and an asynchronous Global Mirror relationship to the site 3 volume sets.

Note: Consult with your IBM representative during the planning process of the migration project to validate your replication network requirements. The correct network infrastructure helps you to avoid inadvertently affecting the production and replication environments during the data migration.

In a GDPS MGM IR configuration, there are asymmetric devices. All couple data sets (CDS) except for the Logger data set are asymmetric (ASYM). This means the primary CDS has a Site 1 to Site 3 asynchronous relationship with a Site 2 empty place holder device. The alternate CDS has a Site 2 to Site 3 asynchronous relationship with a Site 1 empty place holder device. This type of configuration is required in an MGM IR environment and also required to complete volume set migrations.

Figure 7-12 is an example of how the CDS devices must be configured for these migration procedures to work.

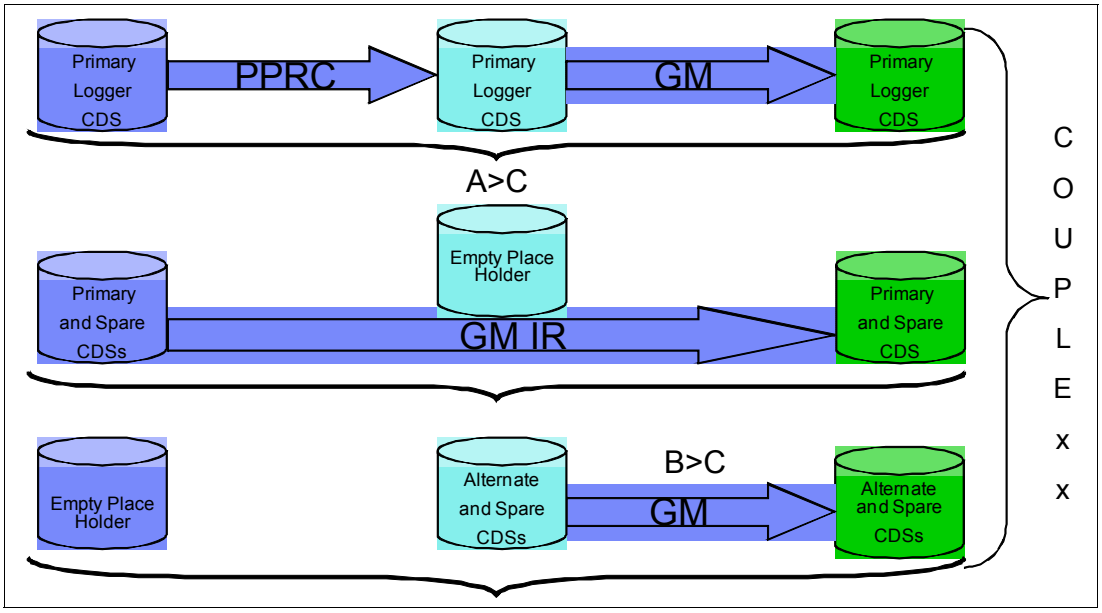

Figure 7-12 Couple data set configuration with MGM IR

Replacing the C Volume Set with Automation

The approach described here to replace the C volumes set uses GDPS scripts and Copy Services to complete the migration. This process can be completed without interruption to host processing. By using the cascaded Global Copy to migrate the data before the transition, you minimize the time that is required to fully recover the mirrored environment.

Additional Fibre Channel (FC) connectivity is required between C to C', A to C', and B to C'.

If the GDPS Site 2 Kp-sys mirrors data to the C volume set you are migrating or in the GM session, it must be shut down at the start of the migration and brought back up at the end of the migration.

It is assumed that the C to C' cascaded relationship has completed First Pass before the actual migration, so the time required to obtain 100% copy status is minimal.

C' Initial copy phase

The initial setup of the cascaded replication environment is described in 2.3, "Data replication" on page 22. Refer to that section to complete this process and then return here after that is accomplished.

All write activity to the B devices must be stopped to begin the moving of the Global Mirror targets to the new populated C' volume set. The fundamental concept behind the migration plan is if there are no write updates to B and all out of sync tracks are copied from B to C to C', then all three of devices are equal. This configuration allows the relationship from B to C to be moved to B to C' without performing a full copy.

Create the cascade copy by establishing paths from C to C', and then establish an XD relationship from the C to the B' devices. These relationships should be established well in advance of your actual migration cut over. The copy time required to complete the first pass depends on your network bandwidth and the configuration size of the volume set being migrated from/to.

Figure 7-13 illustrates the MGM IR configuration with the cascaded copy.

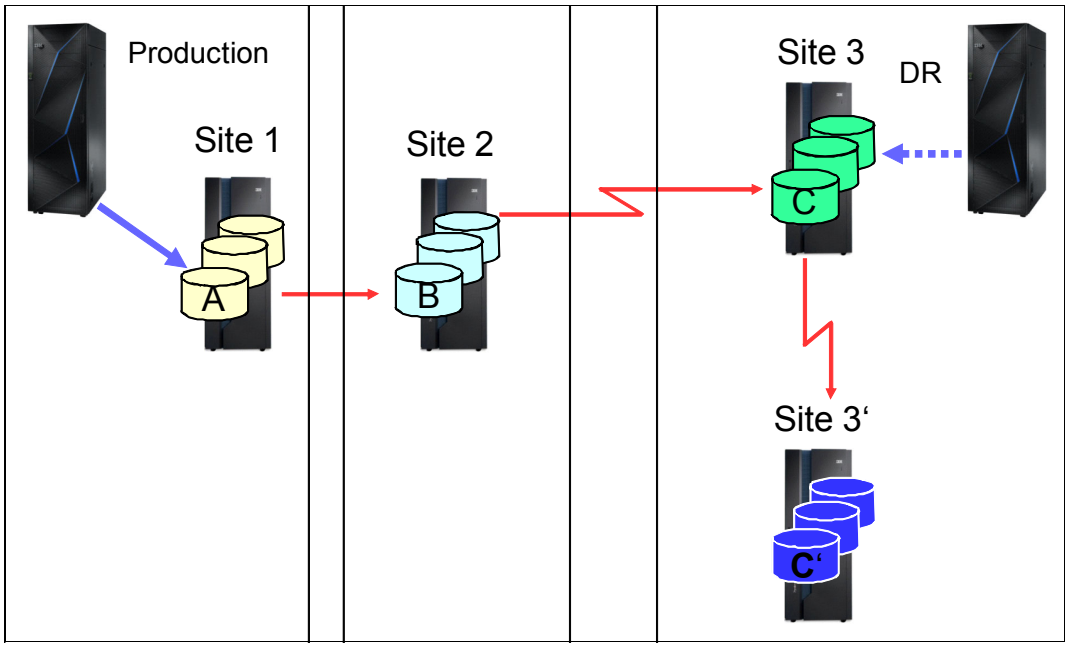

Figure 7-13 MGM IR with Cascaded XD Copy to C'

C' storage migration phase

Freeze the Metro Mirror A to B relationship by using the standard GDPS/PPRC script statement **`STOP SECONDARY`**.

Move all CDS to Site 1 by running GDPS script statement **`CDS SITE1`**.

Shut down the GDPS Kp-sys running on Site 2 devices using standard operating procedures.

Because there are no updates to the B devices currently, validate that all changed tracks on the B devices have been copied to the C devices (out-of-sync tracks = 0). Failure to do this could corrupt the target volumes without any indication of the problem.

Wait for the copy to complete from B to C and C to C'.

The B and C devices are identical. Stop the Global Mirror session and suspend the B to C relationship by running GDPS script statement *session* **STOP SECONDARY**.

Delete the Global Copy relationship from B to C, and remove the paths from B to C.

Validate that all updates to the C devices have been copied to the C’ devices (out-of-sync tracks = 0). As with the B to C validation, failure to ensure this relationship corrupts the Global Mirror session at the end of this migration.

At this point, the A devices are being updated and changes are being recorded in the Change Record (CR) bitmap. The B devices are not being updated because Metro Mirror (A to B) is frozen. Because Global Mirror is stopped and the Global Copy relationship between B and C has been deleted, there are also no updates to the C devices. Verify that all changed tracks have been copied to C'. If so, B is equal to C, which is equal to C'.

Delete the Global Copy relationships and paths from C to C’.

Remove the Global Mirror FlashCopy from the C devices by running GDPS script statement *session* **CLEANUPR GMFC**.

To complete the integration of the C’ volume set into the environment, install and validate all physical cabling required, including the replication links to the B volume sets. Use standard operating procedures to complete this task, which you would normally perform to swap out hardware such as swapping cables or activating an input/output definition file (IODF) to remove the old box and add the new box.

Validate that the Kr-sys now references the C' devices instead of the old C devices. Issuing more CQUERYs and other MVS commands to validate is advised. You should run an initial data validation test on the C’ volume set before proceeding with the final cut over.

Cut over to the new volume set

Establish the Global Mirror C’ to D’ ASYNC FlashCopy relationships.

Create a GEOMPARM with the C’ volume set and load this GEOMPARM configuration by using the GDPS/GM “Configuration Management” panel (option C).

Run all GDPS/Global Mirror monitors and verify the accuracy of the configuration.

This is the GO/NOGO decision point. Perform data validation on the new C’ volume set. If you determine there are errors during the validation process and you need to back out, see “Back Out Scenario for C’” on page 141.

Define the PPRC paths and establish the Global Copy relationships between the B and C’ devices. Because the B and C' devices are identical, use the NOCOPY option to prevent a full copy to the C' devices.

Because the A to B Metro Mirror is still frozen, and no writes are going to the B devices, you can create a data consistent copy of the environment on the E' devices (FC1). This Point in Time (PiT) copy can be used during Disaster Recovery (DR) certification.

The following steps return the environment to “Steady State” where HyperSwap is enabled and Global Mirror is running from to C'.

Using standard operational procedures, IPL K2-Sys on the B devices.

Using the GDPS/Global Mirror DASD Management panels (option 1), start the Global Mirror session.

Validate that Global Mirror is forming consistency groups by using the GDPS DASD panels and the Global Mirror Monitor is started. Only updates from the GDPS/GM Local controlling system Site 2 devices are copied so the RPO and Out-of-Sync (OOS) tracks should be very low.

Resync Metro Mirror A to B using the GDPS/PPRC script statement **START SECONDARY**.

The duration of this resynchronization depends on the amount of time Metro Mirror has been frozen, and the number of changes that must be copied from the A devices.

Move your CDS back to a normal configuration with the primary's running on Site 1 and alternates are on Site 2. The movement of the CDS can be accomplished by using the GDPS script statement **CDS NORMAL**.

Validate Global Mirror is forming consistency groups by using the GDPS DASD panels and the Global Mirror Monitor is started.

Run all GDPS monitors to validate the environment. GDPS/PPRC should be green OK. HyperSwap should be enabled, and the CDS configuration should have Primaries at Site 1, and Alternates at Site 2 (Normal Config). GDPS/Global Mirror should also be green OK and forming Consistency Groups.

Perform necessary DR testing to certify that the application Disaster Recovery images are now accessing the new C' volume set.

See Figure 7-14 for the final MGM IR configuration with the new C' volume set.

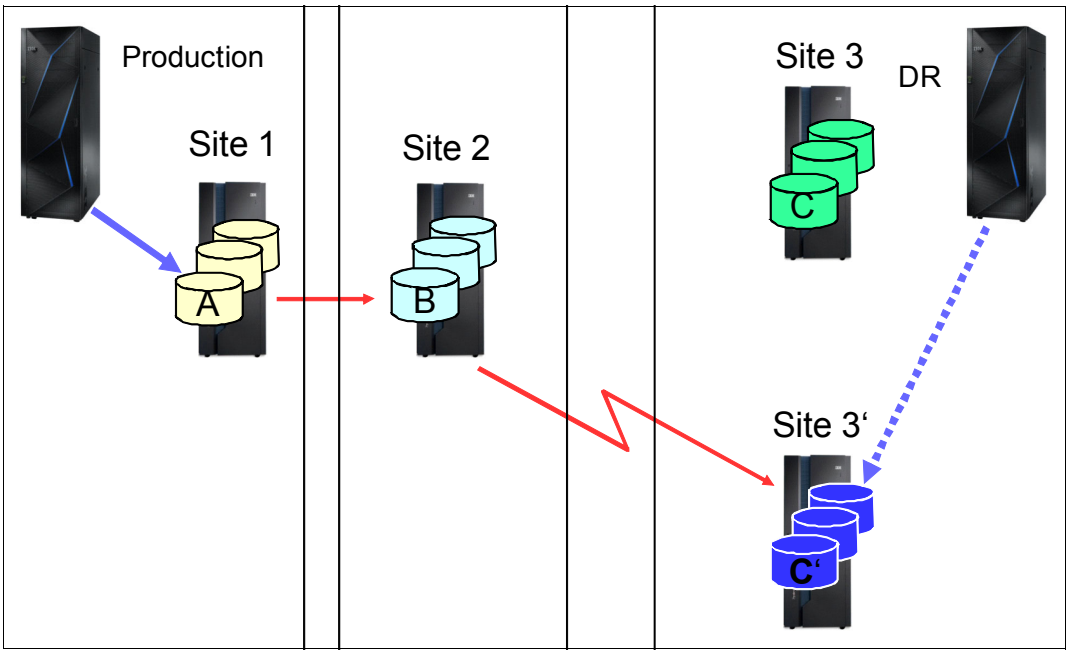

Figure 7-14 Final MGM IR Configuration with C'

Back Out Scenario for C'

If during the validation testing you encounter a situation where you need to back out, follow these procedures to insert the original B volume set into the configuration.

At the point of GO/NOGO during validation testing, your configuration A to B is suspended, and neither B to C nor B to C' exist. The Global Mirror Flashcopies C' to D' were created.

Withdraw the Global Mirror C' to D' FlashCopy.

Withdraw the C' to E' (FC1) test FlashCopy.

Remove the C’ volume set and add the C volume set back into your configuration using the same technic such as swapping cables or reverting to the original IODF. Verify that your Kr-sys can access the original C devices.

Back out the GEOMPARM that references the new C’ volume set by loading the old GEOMPARM with the C volume set.

Run all GDPS/GM monitors to validate the configuration.

Establish the Global Copy paths from B to C.

Establish the Global Copy pairs from B to C.

Create the Global Mirror C to D FlashCopy relationships.

Using standard operational procedures, IPL the Site 2 Kp-sys on the B devices.

Resync Metro Mirror A to B using the GDPS/PPRC script statement ‘START SECONDARY’.

Move your CDS back to a normal configuration with the primaries running on Site 1’ and alternates on Site 2. The movement of the CDS can be accomplished by using a GDPS script statement **CDS NORMAL**.

Run all GDPS monitors to validate the environment. GDPS/PPRC should be green OK. HyperSwap should be enabled, and the CDS configuration should have Primaries at Site 1, and Alternates at Site 2 (Normal Config). GDPS/Global Mirror should also be green OK and forming Consistency Groups.

The environment has now been restored to the state before the C' migration. Metro Mirror A to B is active and Global Mirror B to C running.

Replacing the B Volume Set with Automation

The approach described here to replace the B volume set uses both GDPS scripts and the Copy Services IR function to complete the migration. A cascaded relationship is created from the C volume set to the new B’ volume set.

Additional FC connectivity is required between C to B’, B to B’, and A to B’.

If you have Copy Once volumes, add them to the Global Mirror session after the first step in the migration process and then remove them after the migration.

It is assumed that the C to B’ cascaded relationship has completed First Pass before the actual migration, so the time required to obtain 100% copy status is minimal.

B’ Initial Copy Phase

The initial setup of the cascaded replication environment is described in 2.3, “Data replication” on page 22. Refer to that section to complete this process and then return here after that is accomplished.

If you have added Copy Once volumes back into the GM session, IR needs to be enabled on these volumes to complete the migration. Run GDPS/PPRC script statement ‘INCRESYN SETGM’. This script statement is only allowed to run when the GDPS/PPRC configuration is green OK and briefly sets the mirroring status to NOK and runs a suspend/resume of the

PPRC pairs. In a HyperSwap environment, a Monitor2 or a START SECONDARY must be used to re-enable HyperSwap

Establish paths from C to B'.

Establish a cascaded Global Copy (PPRC/XD) relationship from the C devices to the B' devices. If you are replacing the volume set where your Site 2 Kp-sys volumes are and it is not part of the GM or does not have an XD relationship B to C, establish paths and then a Global Copy relationship from the B to B' for these volumes. These relationships should be established well in advance of your actual migration cut over. The copy time required to complete the first pass depends on your network bandwidth and the configuration size of the volume set you are migrating from/to.

Figure 7-15 illustrates the MGM IR configuration with the cascaded copy.

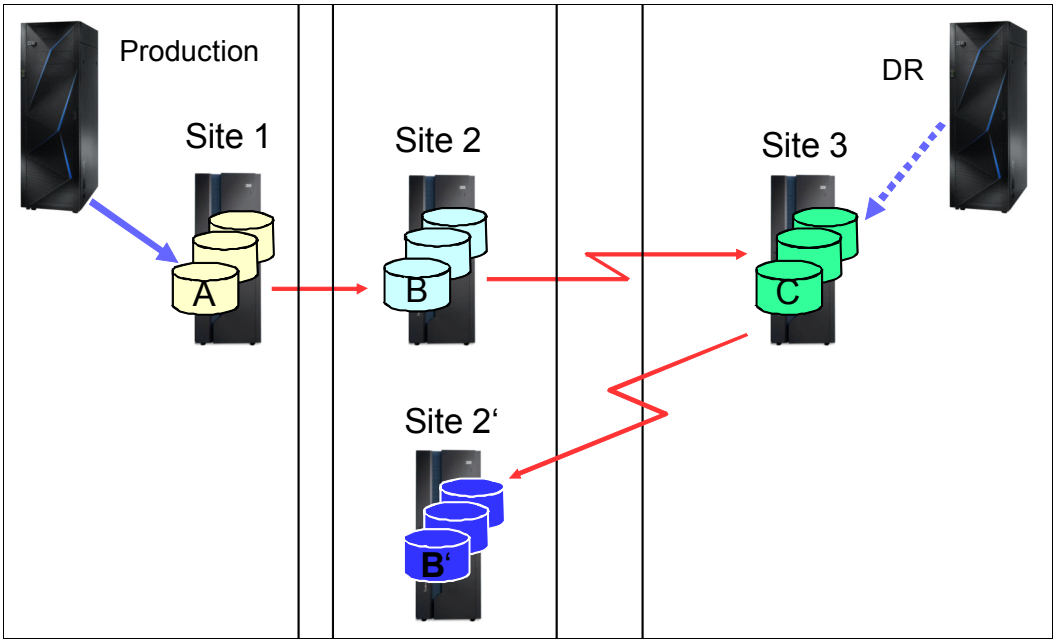

Figure 7-15 MGM IR with Cascaded XD Copy to B'

Important: Do not start the next step until at least six good Global Mirror consistency groups have been formed after adding the Copy Once volumes into the GM session and enabling IR. Failure to wait will cause a full copy.

B' Storage migration phase

Using standard operating procedures, shut down the Site 2 Kp-sys.

Freeze the A to B relationship by running GDPS/PPRC script statement **`STOP SECONDARY`** to stop all I/O to the B volume set.

Move all CDS to Site 1 by running the GDPS script statement **`CDS SITE1`**.

Vary the B CDS devices offline on all systems.

The A devices are now being updated and changes are being recorded in the CR bitmap. The B devices are not being updated because Metro Mirror (A to B) is frozen.

Stop the Global Mirror session by using the GDPS DASD Remote Copy Panel, *not* a GDPS script.

Prepare to move the GM session from B to A. Run the GDPS script statement *session* **PREPMOVE1**. This results in the relationship from A to B devices being suspended, if not already suspended, and starting the PPRC paths from A to C.

Clean up the C devices for the B to C relationship by issuing the Copy Services recover command.

Change the GM site indicator from Intermediate (B to C) to Local (A to C) by using the GDPS script statement *session* **CHGGMSITE**.

Establish the A to C relationship with MODE(INCRES) causing the resynchronization to occur using the CR bitmap limiting the tracks that must be copied. See Example 7-28 for the command to establish this relationship using IR.

Example 7-28 Establish A to C with IR

```
CESTPAIR DEVN(X'Adevn')                                          -
         PRIM(X'Assid' Aserial X'Acca' X'Alss')      -
         SEC(X'Cssid' Cserial X'Ccca' X'Cclss')  -
         MODE(INCRES) OPTION(XD) FORCE(YES)
```

Start Global Mirror by using the GDPS DASD Remote Copy panel for the Local (A) session.

Run the GDPS/GM monitors and validate good consistency groups are being formed.

Clean up the B devices for the A to B relationship by issuing the **recover** command.

Clean up the B devices for the B to C relationship by issuing the **delpair** command.

You are now ready to remove the B volume set and add the B' volume set into the configuration. Use standard operating procedures to complete this task, which you would normally perform to swap out hardware such as swapping cables or activating an IODF to remove the old box and add the new box. Perform data validation on the new B' volume set. This is your GO/NOGO decision point. If you determine there are errors during the validation process and you need to back out, see "Back out scenario for B'" on page 146.

Cut-over to the new volume set

Load the current GEOPARM configuration by using the GDPS/PPRC 'Configuration Management' panel (option C) to find the new B' volume set. If the PPRC links changed, the GEOPARM must be updated with the changes before it is loaded.

Establish the A to C relationship with INCRESYNC(NOCOPY) and MODE(NOCOPY) to start IR from the A to the C devices. Example 7-29 shows the command used to establish this relationship.

Example 7-29 Start IR from A to C

```
CESTPAIR DEVN(X'Adevn')                                          -
         PRIM(X'Assid' Aserial X'Acca' X'Alss')      -
         SEC(X'Cssid' Cserial X'Ccca' X'Cclss')  -
         MODE(NOCOPY) OPTION(XD) INCRESYNC(NOCOPY)
```

Stop the GM session running A to C with the GDPS script statement *session* **STOP SECONDARY**.

Issue the recovery command to clean up the non-ASYM C devices for the A to C relationship.

Define the new B’ LSS’ and join the B’ volumes to the Global Mirror session.

Establish the Global Copy paths from B’ to C and the PPRC paths from A to B’.

Fail over and then fail back the B’, B’ to C devices including the ASYM devices to reverse the direction of the C to B’ relationship to a B’ to C relationship.

Monitor the B’ to C copy to ensure that it completes before moving on to the next step.

Resync the A to B’ relationship with using IR with MODE(INCRES) and FORCE(YES). Example 7-30 illustrates the IR command to perform this resync.

Example 7-30 Resync A to B’ with IR

```
CESTPAIR DEVN(X'Adevn')                                          -
         PRIM(X'Assid' Aserial X'Acca' X'Alss')      -
         SEC(X'B’ssid' B’serial X'B’cca' X'B’clss')  -
         MODE(INCRES) OPTION(SYNC) FORCE(YES)
```

Monitor the copy process A to B’ to ensure that it goes full duplex.

Start IR A to B’ with the NOCOPY option. Example 7-31 illustrates the IR command to perform the start IR.

Example 7-31 Start IR A to B’

```
CESTPAIR DEVN(X'Adevn')                                          -
         PRIM(X'Assid' Aserial X'Acca' X'Alss')      -
         SEC(X'B’ssid' B’serial X'B’cca' X'B’clss')  -
         MODE(NOCOPY) OPTION(SYNC) INCRESYNC(NOCOPY)
```

Change the GDPS/GM session site indicator to point to B’ from A. Issue the GDPS script statement **session CHGGMSITE**.

Create a GEOMPARM with the new B’ volume set. Load this new GEOMPARM configuration by using the GDPS/GM ‘Configuration Management’ panel (option C) to find the new B’ volume set.

Resync the A to C ASYM devices. Issue the GDPS script statement *session* **START SECONDARY RESYNC ASYM**.

Start the GM session by using the GDPS ‘DASD Remote Copy’ panel.

Using standard operating procedures, IPL the Site 2 Kp-sys on the B’ volume set.

You are now running A to B' to C with the new volume set installed. See Figure 7-16 for the final MGM IR configuration.

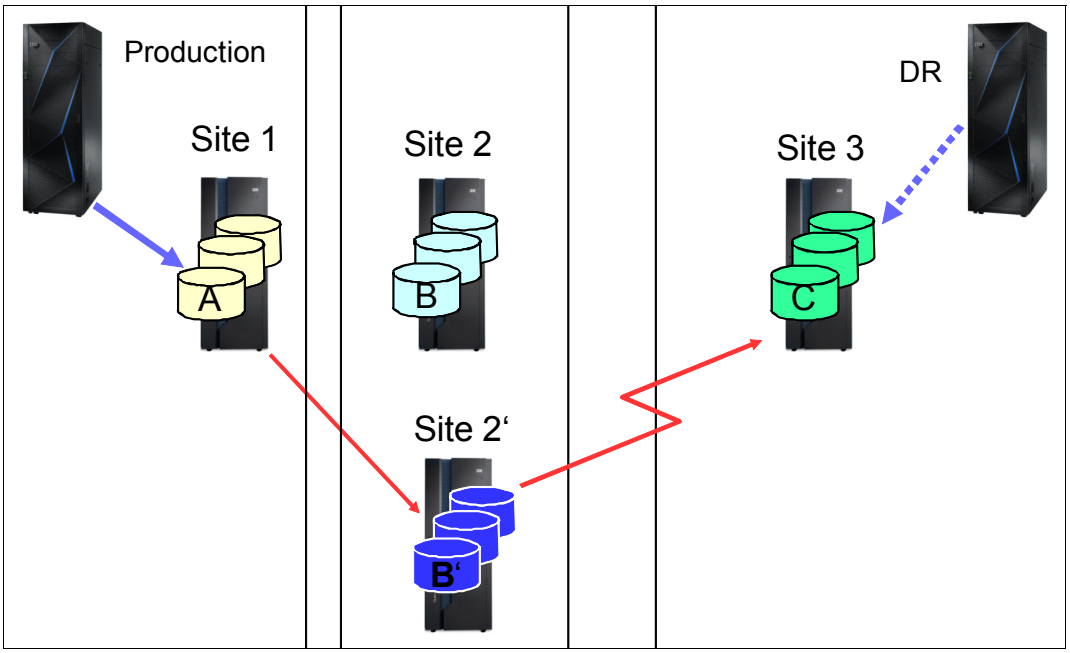

Figure 7-16 Final MGM IR Configuration with B'

Back out scenario for B'

If during the validation testing you encounter a situation where you need to back out, the same procedures should be used to insert the original B volume set into the configuration.

Remove the B' volume set and add the B volume set back into your configuration using the same techniques such as swapping cables or reverting back to the original IODF. At the point of GO/NOGO during validation testing, your configuration is Global Mirror session running A to C. After the B volume set is again accessible by the Kg-sys, continue with the next step after the GO/NOGO step and complete all steps through the remainder of the procedures with your original B volume set.

Replacing the A volume set with automation

This approach to replace the A primary (A) volume set uses GDPS scripts, GDPS HyperSwap, and the Copy Services IR function to complete the migration.

Additional FC connectivity is required between C to A', A to A', and A' to B.

If your Site 2 Kp-sys or Kg-sys are included in the GM session and are on a triplet where the A devices are empty and are used as a place holder, no further action is required. If it is not, remove it from the Global Mirror session before the migration and add it back in after the final step.

If you have Copy Once volumes, add them to the GM session before the migration and remove them after the final step to ensure that they get copied to your A' volume set.

It is assumed that the C to A' cascaded relationship has completed First Pass before the actual migration, so the time required to obtain 100% copy status is minimal.

Note: You must wait for at least six good Global Mirror consistency groups to be created after enabling the IR bitmap before you can use the bitmap during the replacement of the A' volume set. Failure to wait the appropriate amount of time causes a full copy to occur.

A' Initial Copy Phase

The initial setup of the cascaded replication environment is described in 2.3, "Data replication" on page 22. See this section to complete this process and then return here after that is accomplished.

If you have Copy Once volumes, add them into the B to C and A to C GM sessions. After you add Copy Once volumes back into the GM session, IR needs to be enabled on them for this migration. Run GDPS/PPRC script statement **`INCRESYN SETGM`**. This script statement is only allowed to run when the GDPS/PPRC configuration is green OK and briefly sets the mirroring status to NOK and runs a suspend/resume of the PPRC pairs. In a HyperSwap environment, a Monitor2 or a **`START SECONDARY`** must be used to re-enable HyperSwap.

Establish the paths and then the cascaded PPRC/XD relationship from the C devices to the A' Devices. If you are replacing the volume set where your Site 1 Kp-sys volumes are, establish a PPRC/XD relationship from the A to A' for these volumes. These relationships should be established well in advance of your actual migration cut-over. The copy time required to complete the first pass depends on your network bandwidth and the configuration size of the volume set you are migrating from/to. Figure 7-17 illustrates the configuration with the cascaded copy to A'.

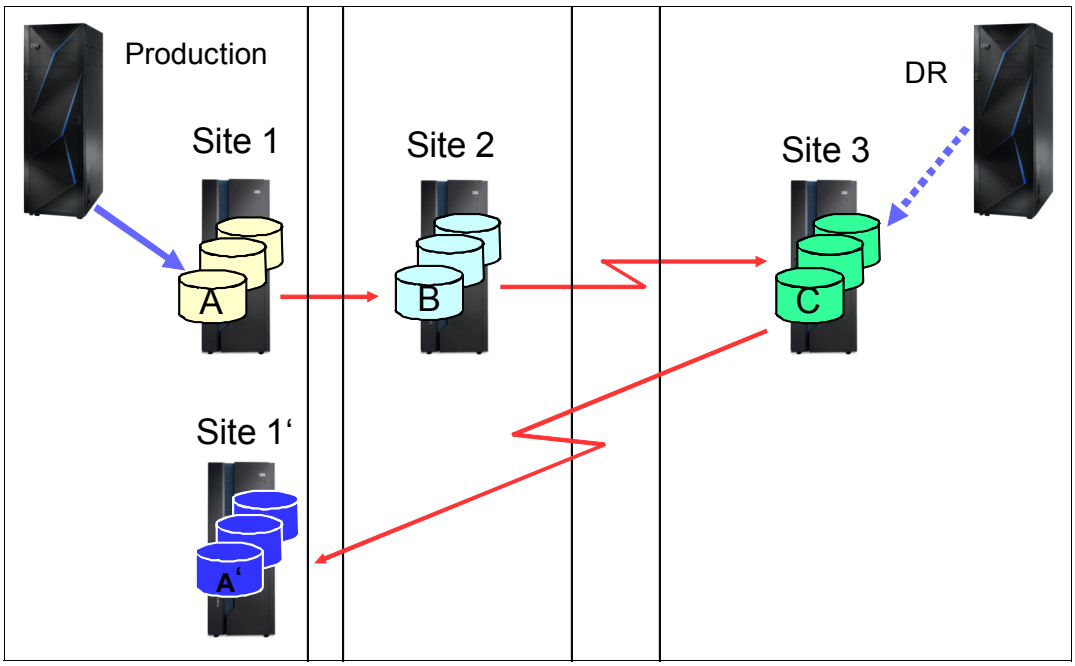

Figure 7-17 MGM IR Configuration with Cascaded Copy to A'

A' Storage Migration Phase

Ensure that your GDPS/PPRC environment is HyperSwap enabled and green OK. Run GDPS script statement **`SWITCH HyperSwap SUSPEND`** to the B devices to free up the A volume set you are replacing.

Move all of your CDS to Site 2 and vary all of your A CDS devices offline. The movement of the CDSs can be accomplished by issuing the GDPS script **CDS SITE2**. If your Site 1 Kp-sys volumes are on the volume set that you are migrating, it must be shut down.

After your Site 1 Kp-sys is down and the out-of-sync tracks=0 from A to A', delete the A to A' relationships and paths.

Recover the B devices, A to B.

Enable IR on B to C to allow a resynchronization from B to A' without a full copy.
Example 7-32 is the command to start IR with the NOCOPY option that indicates the change recording bitmaps do not need to be initialized with all bits on.

Example 7-32 Start IR with NOCOPY

```
CESTPAIR DEVN(X'Bdevn')                                          -
         PRIM(X'Bssid' Bserial X'Bcca' X'Blss')  -
         SEC(X'Cssid' Cserial X'Ccca' X'Cclss')  -
         MODE(NOCOPY) OPTION(XD) INCRESYNC(NOCOPY)
```

Stop the intermediate GM session (B to C) by running GDPS script statement **session STOP SECONDARY**.

Cleanup should be completed on the A volume set you are removing. Delete the relationships from A to B and A to C. You might receive an RC=8 from the delpair A to B because there is no path between the two devices, but it does clean up the relationship. Delete the paths between A to B and A to C.

You should now be ready to remove the A volume set and add the A' volume set into the configuration. This must include any tasks that you typically perform to swap out hardware such as swapping cables or activating an IODF to remove the old volume set and add the new volume set.

Perform a configuration and data validation test on the A' volume set before the final cut-over. The Site 1 Kp-sys can be IPLed on A' and used for data and configuration validation. This process is not described in detail here because every customer is unique and the process must be tailored to your specific environment. This is a GO/NOGO decision point. If there are errors during the validation process and you need to back out, see "Back Out scenario for A'" on page 151.

Cut over to the new volume set

To introduce the A' volume set to GDPS, load the GDPS/PPRC GEOPARM. Updates might be required for this GEOPARM if there are differences between the configuration of the A and A' for the LSS range, SSID range, UCB range, or PPRC links.

The next steps will convert the configuration to a B to A' to C MGM IR environment:

1. Validate that the copy from C to A' has out-of-sync tracks=0. If the out-of-sync tracks are zero, suspend the C to A' relationships including the ASYM devices.
2. Recover the C ASYM devices, A to C.
3. Define all LSS' and join the volumes to the A' Global Mirror session.
4. Establish paths from A' to C and B to A'.
5. Issue a failover for the volume pairs from A' to C including the ASYM devices.
6. Issue a failback for the volume pairs, A' to C including the ASYM devices.

7. Monitor the copy from A’ to C to ensure that first pass has completed.
8. After first pass has completed, resynchronize the B to A’ relationships using IR. Example 7-33 shows the command to complete the resynchronization using IR.

Example 7-33 Resync with IR

```
CESTPAIR DEVN(X'Bdevn')                                          -
         PRIM(X'Bssid' Bserial X'Bcca' X'Blss')      -
         SEC(X'A’ssid' A’serial X'A’cca' X'A’clss')  -
         MODE(INCRES) OPTION(SYNC) CASCADE(YES)                    -
         FORCE(YES)
```

9. Start IR on B, B to A’. The command to start IR can be found in Example 7-34.

Example 7-34 Start IR

```
CESTPAIR DEVN(X'Bdevn')                                          -
          PRIM(X'Bssid' Bserial X'Bcca' X'Blss')       -
           SEC(X'A’ssid' A’serial X'A’cca' X'A’clss')  -
          MODE(NOCOPY) CASCADE(YES) OPTION(SYNC) INCRESYNC(START)
```

10. Monitor the copy process until B to A’ goes full duplex and A’ to C completes first pass before moving on to the next step.
11. Resync the ASYM devices with the CASCADE option, B to C.
12. Change the Global Mirror SITE indicator from running B to C to A’ to C by running GDPS/GM script statement *session* **CHGGMSITE**.
13. Load a new GDPS/Global Mirror GEOMPARM with the updated A’ storage subsystem and start the Global Mirror session on LOCAL (A’). Your applications are now running on your B volume set and the MGM configuration should be B to A’ to C as displayed in Figure 7-18.

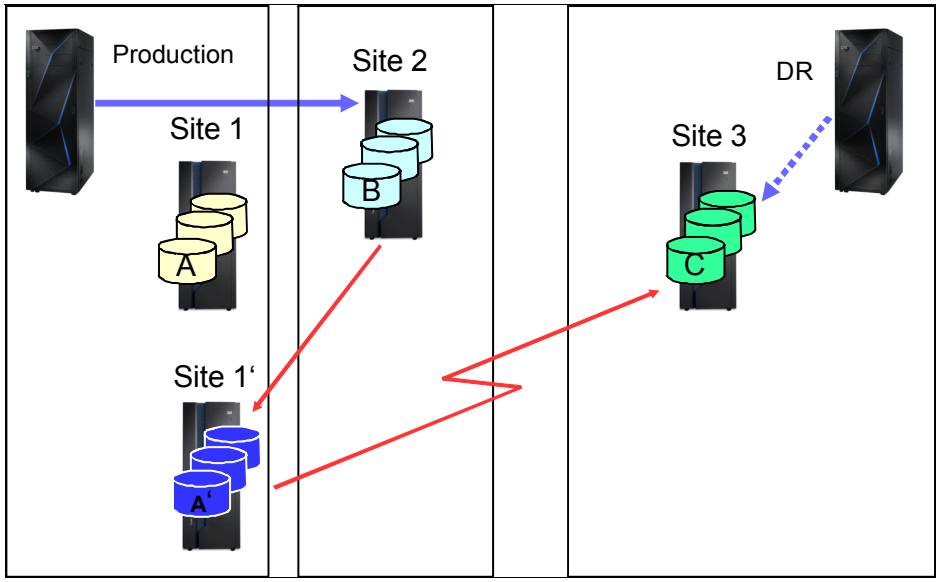

Figure 7-18 MGM IR Configuration B to A’ to C

The next set of steps will convert from a B → A’ → C configuration to an A’ → B → C configuration.

> **Important:** Do not start the conversion process of moving to an A’ to B to C configuration until at least six good Global Mirror consistency groups have been formed since starting the Global Mirror session running A’ to C. Failure to wait causes a full copy.

Establish the C to B paths.

Fail over the C to B relationships including the A-SYM devices using the FORCE(YES) option. This failover must be completed before you HyperSwap to A’ to prevent a full copy from C to B.

Vary the A’ CDS UCB’s online and add your spare CDSs by using the GDPS panels.

Check GDPS/PPRC to ensure that your environment is HyperSwap enabled and green OK. Issue the GDPS script command **`SWITCH HyperSwap SUSPEND`** to HyperSwap to the A’ devices.

Move all of the CDS to the A’ devices and vary the B CDS devices offline. The movement of the CDSs can be accomplished by issuing a GDPS script using the statement **`CDS SITE1`**. All application I/O is now running on your new A’ volume set.

Fail back the C to B relationships, including the ASYM devices using the FORCE(YES) option. Wait for the first pass to complete, C to B.

Recover the A’ devices, B to A’.

Start IR on the A’ devices, A’ to C with the NOCOPY option. See Example 7-35 for the start IR command.

Example 7-35 Start IR on A’

```
CESTPAIR DEVN(X'Adevn’')                                          -
           PRIM(X'A’ssid' A’serial X'A’cca' X'A’lss')  -
            SEC(X'Cssid' Cserial X'Ccca' X'Cclss')     -
           MODE(NOCOPY) OPTION(XD) INCRESYNC(NOCOPY)
```

Issue the GDPS script command ***`session`*** **`STOP SECONDARY`** to stop the Global Mirror session running from A’ to C.

Suspend the C to B relationships after ensuring first pass has completed, including the ASYM devices.

Issue a recover of the C devices, B to C.

Fail over all of the B to C relationships including the ASYM devices with the CASCADE option.

Fail back all of the B to C relationships including the ASYM devices with the CASCADE option.

Monitor the copy process from B to C and wait for the first pass to complete before moving on to the next step.

Resync the A' to B relationships using IR. See Example 7-36 for the IR command to complete the resync.

Example 7-36 Resync A' to B

```
CESTPAIR DEVN(X'Adevn'')                                          -
           PRIM(X'A'ssid' A'serial X'A'cca' X'A'lss')  -
            SEC(X'Bssid' Bserial X'Bcca' X'Bclss')     -
           MODE(INCRES) OPTION(SYNC) FORCE(YES)
```

Start IR on the A' to B relationship. See Example 7-37 for the start IR command.

Example 7-37 Start IR A' to B

```
CESTPAIR DEVN(X'Adevn'')                                          -
           PRIM(X'A'ssid' A'serial X'A'cca' X'A'lss')  -
            SEC(X'Bssid' Bserial X'Bcca' X'Bclss')     -
           MODE(NOCOPY) OPTION(SYNC) INCRESYNC(START)
```

Monitor the A' to B relationship to ensure that it goes full DUPLEX and the B to C relationship waiting for first pass to complete.

Resync the A' to C ASYM devices with the CASCADE option.

Move your CDS back to a normal configuration with the primary's running on A' and alternates on Site 2. The movement of the CDS can be accomplished by using a GDPS script statement **CDS NORMAL**.

If you added Copy Once volumes to the GM session, remove them. If you removed the Site 2 Kp-sys or Kg-sys from the GM session, add it back in.

You are now running A' to B to C with the new volume set installed. See Figure 7-19 for the final MGM IR configuration.

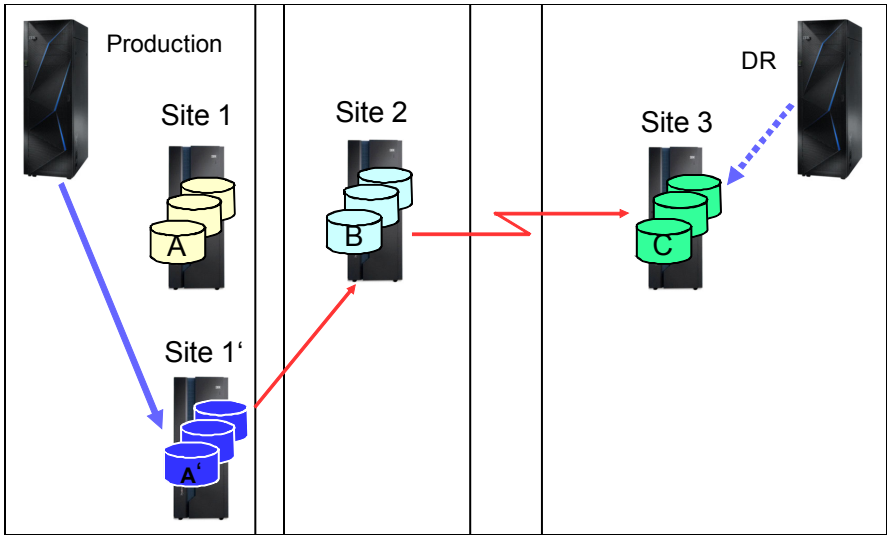

Figure 7-19 Final MGM IR Configuration

Back Out scenario for A'

If during the validation testing you encounter a situation where you need to back out, the same procedures should be used to insert the original A volume set into the configuration.

Remove the A’ volume set and add the A volume set back into your configuration using the same technique such as swapping cables or reverting to the original IODF. At the point of GO/NOGO during validation testing, your configuration is a suspended B to C relationship and Global Mirror is stopped. After the A system is accessible by the Kg-sys, start the next step after the GO/NOGO step and complete all steps through the remainder of the procedures with your original A volume set to get back to an A to B to C configuration.

Replacing the A and B volume sets

Replacing both the A and B volume sets uses the approach described in 7.3.4, “Replacing the A and B volume set without automation” on page 126 with a couple exceptions relating to GDPS managing the MGM environment.

The Site 1 Kp-sys, Site 2 Kp-sys, and Kg-sys need to be shut down at the same time as all of production systems using these volume sets.

The Global Mirror session needs to be stopped by using the GDPS Remote Copy panel at the beginning of the migration, not by using a GDPS script.

If you have Copy Once volumes, add them to the GM session before the migration and ensure that a Global Copy relationship has been established from C to A’ to B’ before the start of the migration. Remove the Copy Once volumes after the final cut-over step.

IPL the GDPS Kx-systems before establishing the Global Mirror environment, but after establishing the B' to C relationships with the NOCOPY option, in the following steps provided in “Cut-over to the new volume set” on page 112.

IPL the GDPS Kx-systems. Load the GDPS/PPRC GEOPARM and run the GDPS monitors. Ensure the A’ to B’ relationship is SYNC. Load the GDPS/GM GEOMPARM. Establish the GM FlashCopy relationships. If your configuration is MGM IR, define the local subsystems (LSSs) for both the Local (A’ to C’) and the Intermediate (B’ to C’) GM sessions. If you have multiple volume sets for A’, B’, or both, establish the Master-Subordinate links and then join the volumes to both the Local and the Intermediate GM sessions. Then start the GM session (B’ to C).

Replacing all volume sets (A, B & C)

The same approach can be used as described in 7.3.5, “Replace all volume sets without automation” on page 132 to replace all volumes sets with some exceptions with GDPS managing the MGM environment.

The GDPS Kx-systems must be shut down when all application systems are shut down. This migration requires a sysplex-wide IPL, so in a GDPS environment, the last system that was stopped must be the first one IPLed after migration. If you do not follow this recommendation, misleading Takeover prompts (GEO112E/GEO113A messages) might occur. The preference is to take down the current GDPS Master controlling Kp-sys last, which should then be the first one brought back up after the migration.

To prevent alerts and errors, new GDPS/PPRC GEOPARM and GDPS/GM GEOMPARM should be created specifying the new volume set configuration before the Site 1 Kp-sys, Site 2 Kp-sys, and Kg-sys GDPS are started.

Establish the GM FlashCopy pairs, C’ to D’. After GDPS is started, define the LSS’ for both the Local (A’ to C’) and the Intermediate (B’ to C’) GM sessions. If you have multiple volume sets for A’, B’, or both, establish the Master-Subordinate links and then join the volumes to both the Local and the Intermediate GM sessions. Start the intermediate, B’ to C’ GM session and validate good consistency groups are being created.

If you are running an MGM IR configuration, start IR A' to B' by issuing the GDPS script statement **INCRESYN SETGM**.

8

Data Migration with 2-site z/OS Global Mirror replication

This chapter describes several approaches to perform a volume-based migration to a new storage system in a z/OS Global Mirror (zGM) environment.

zGM, previously known as Extended Remote Copy (XRC), is a 2-site disaster recovery solution based on system software that has the potential to provide a self-contained solution that does not require any additional complex or costly management framework.

Because zGM is driven by the host server (System z), you cannot migrate the storage system concurrently and without any impact to applications hosted on that server.

The focus of this chapter is how to migrate to either a new primary storage system or replace the auxiliary storage system in a 2-site zGM configuration without GDPS management.

zGM provides incremental resync support with a 3-site zGM / Metro Mirror configuration (MzGM) and HyperSwap. This is discussed in Chapter 9, “Data Migration with GDPS/Metro and z/Global Mirror replication” on page 185.

This chapter includes the following sections:

- Initial configuration
- Migration process

© Copyright IBM Corp. 2015. All rights reserved.

8.1 Initial configuration

Figure 8-1shows a general zGM configuration. zGM is typically represented by the system data mover (SDM), which is a DFSMSdfp software component within z/OS.

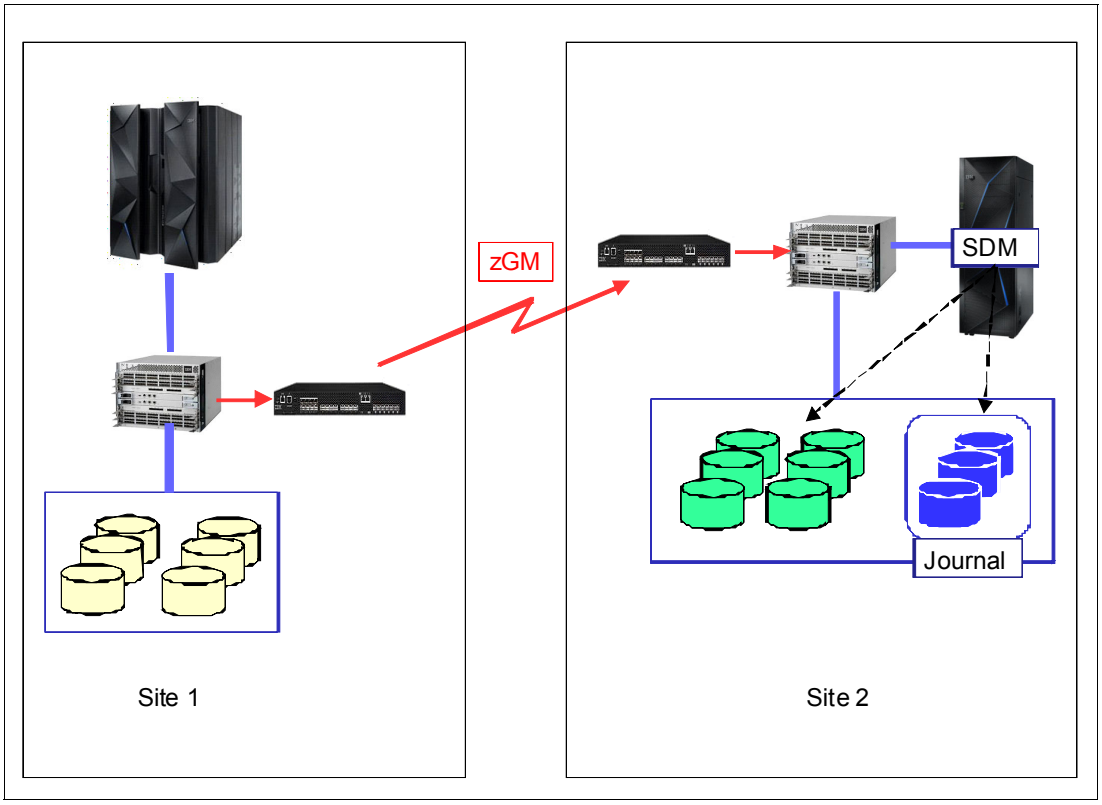

Figure 8-1 zGM D/R configuration

Usually Site 1 hosts the active application server that updates the local volumes. All local writes (that must be replicated to Site 2) are also temporarily buffered in the Site 1 storage system cache. They are then picked up by SDM and copied over to Site 2.

At Site 2, SDM pulls the data out of the storage cache at Site 1 and copies the data twice:

1. Journal writes of all replicated data for data consistency reasons and restart capabilities in case of a failure of SDM during the data replication processing
2. Update writes to all replicated application volumes

This duplicate write is the reason to place SDM at Site 2, among other reasons like providing CPU resources at Site 2 in case of an application failover to Site 2. This configuration allows you to run needed business applications out of Site 2 (on the SDM system).

The volumes labeled Journal in Figure 8-1 also include the required State Data Set and Control Data Set of a zGM session as well as the Couple Data Set and Master Data Set for coupled zGM configurations.

8.2 Migration process

In most cases, volume migration to a new storage system is required to replace an older or depreciated storage system with a new one. Other migration options might include just a set of volumes that must be moved to another disk storage system.

The migration approach depends on the number of volumes that must be moved off the existing storage system and onto a new storage system. It also depends on which tools are available for volume migration. For instance, the Transparent Data Migration Facility (TDMF) can be used, as described in Chapter 10, “Using TDMF z/OS” on page 211.

This chapter assumes an existing 2-site zGM Disaster Recovery (DR) configuration is in place and an existing storage system must be replaced by a new storage system. Figure 8-1 on page 156 generically shows a group of volumes in Site 1 and Site 2 and can represent an entire disk storage system or just a set of volumes that need to be moved to a new disk storage system. Site 1 and Site 2 are distant and are at least in different buildings. In the context of moving onto new devices, this chapter uses a group of volumes with A-, B-, and C-volumes with A- and B-volumes being the existing primary and secondary volumes. C-volumes represent volumes in the new disk storage system.

8.2.1 General considerations

Based on the configuration shown in Figure 8-1 on page 156, you could have the following volume migration scenarios:

- Replace the storage system or parts of it at Site 1
- Replace the storage system or parts of it at Site 2

The goal is to perform the migration with the least impact to the applications at Site 1, or preferably without any impact at all.

Another goal is to contain the costs as much as possible, which is possible with zGM because it does not require any additional copy services licenses and feature codes for what might just be a one time migration effort.

Usually, it is also desirable to provide a fall back possibility. When testing and verification of the data on the new storage system show inconsistencies or a fall back is wanted for any other reason, measures must be in place to provide such a fail back possibility.

8.2.2 zGM running in migration mode

zGM can be used in migration mode to concurrently replace a storage system (or a set of secondary volumes).

Therefore, this chapter covers the difference when zGM runs in migration mode compared to a zGM session with runs in DR mode triggered through the **XSTART** parameter **SESSIONTYPE(XRC)**.

zGM, when running in migration mode, can still handle application I/O and apply changes at Site 1 to remain consistent at Site 2, using consistency group formation.

zGM running in migration mode simplifies the zGM configuration and SDM does not write changed data twice at Site 2. This behavior only makes sense for migration purposes when using zGM to perform the actual data movement. Figure 8-2 depicts such a zGM-based configuration. In zGM migration mode, there is no need to additionally write to Journal volumes like in zGM through the session parameter **SESSIONTYPE(XRC)**.

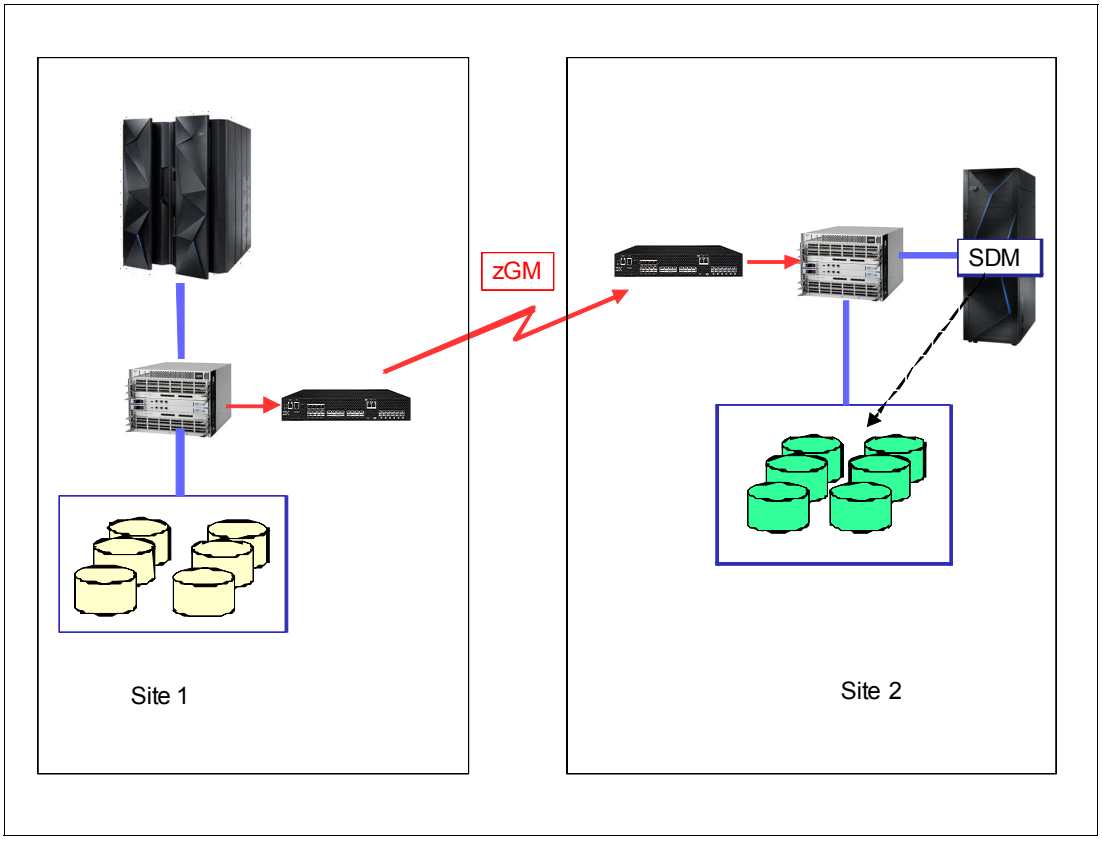

Figure 8-2 zGM configuration running in migration mode

Figure 8-2 can also be read as a migration between distant sites. The zGM-based migration can be particularly interesting when the zGM primary is from a non-IBM storage vendor that also supports zGM functionality in their storage system. In such a case proprietary functions, such as copy services, cannot be used.

zGM in migration mode is going to be the choice when replacing the auxiliary storage system in a D/R zGM 2-site configuration.

8.2.3 Replacing primary storage system in a zGM DR configuration

There are different alternatives to replace the primary storage system. zGM would be a natural choice to avoid the extra cost for additional copy services feature codes. Figure 8-3 shows a cascaded zGM configuration with two zGM sessions:

1. Session 1, SDM1, is the existing D/R configuration. SDM1 is pulling the data through its utility devices from Site 1 and applying the data to the local storage in Site 2, including the journal writes.
2. Session 2, SDM2 is a zGM session that runs in migration mode and is just active while the migration lasts. SDM2 most likely also requires utility devices in Site 2 and is pulling the data from Site 2 to Site 1. Then SDM just applies the data locally in Site 1 to the new storage system.

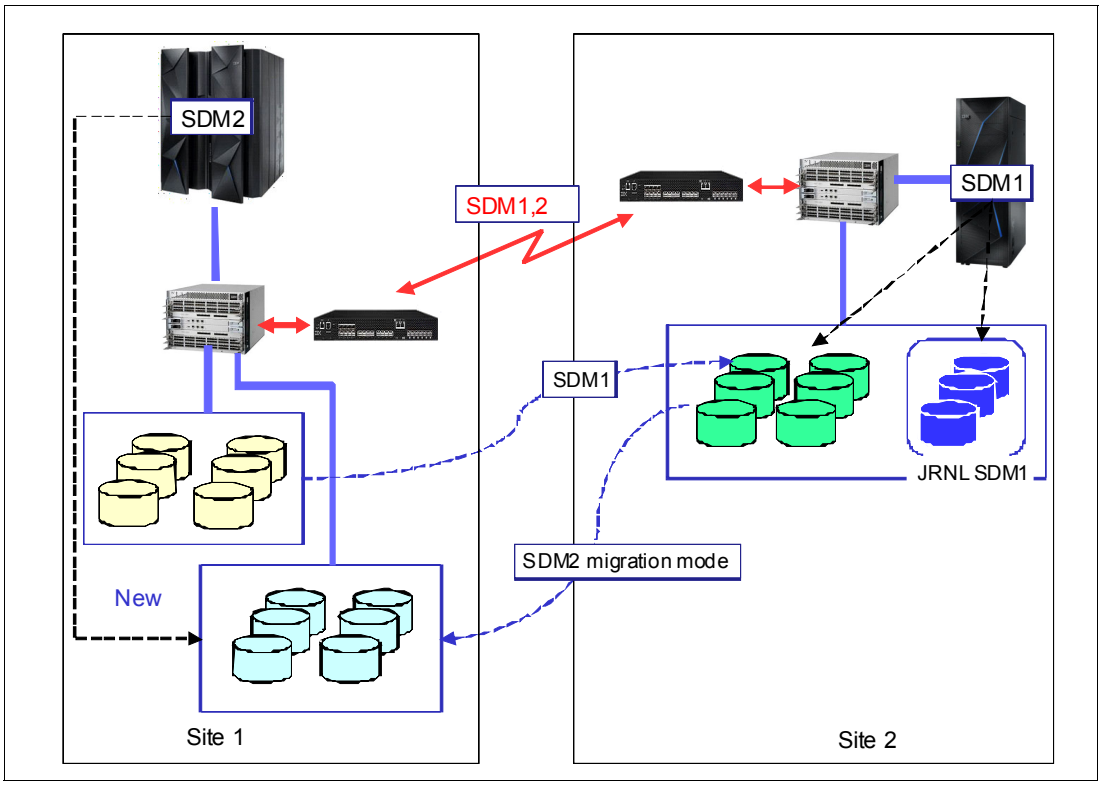

Figure 8-3 Replace zGM primary storage system through cascaded zGM sessions

However, such approach is usually cost prohibitive because you need additional bandwidth between sites (about twice as much data throughput between sites when compared to the normal replication mode). Therefore, it makes more sense to acquire a Copy Services Global Copy or Metro Mirror license and use the scheme depicted in Figure 8-4.

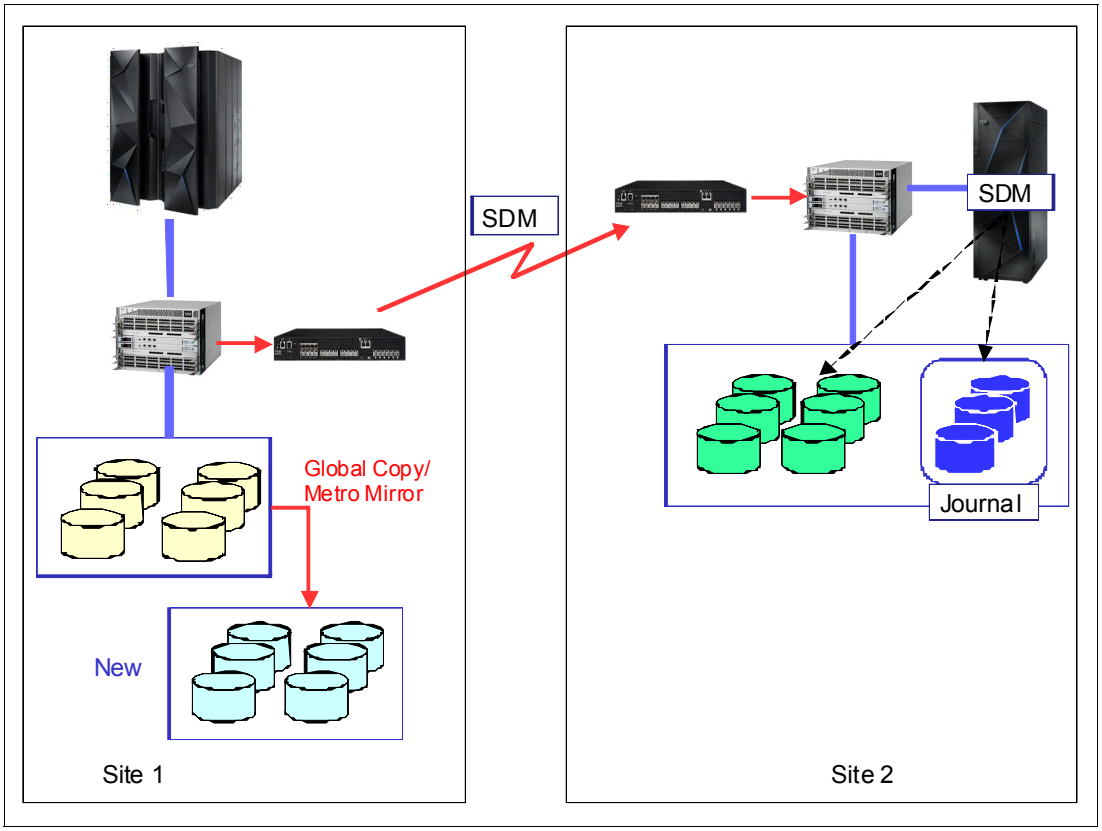

Figure 8-4 Replacing primary storage system using Copy Service GC/MM

SDM2 can also be based in Site 2 and just pushing the data over to the new storage system in Site 1. Because there are no journal writes when running zGM in migration mode, the amount of data moving through the infrastructure is the same as with SDM2 in Site 1.

When replacing the primary storage system or part of it, only application volumes are affected along with potential zGM utility devices in Site 1. Application volumes in this context include also all system volumes such as spool volumes, ML1 volumes, and so on. Generally, exclude paging volumes and move these volumes using z/OS capabilities to dynamically add paging volumes and remove paging volumes from the system.

Migration steps for zGM primary volume migration

To replace a storage system or part of it in Site 1, complete these steps:

1. Establish Global Copy between the existing storage system and the new storage system in Site 1. This assumes that the required PPRC paths have been previously defined.
2. When Global Copy finishes its first full initial copy phase, you are ready to perform the actual swap process. Select a point in time with the least application I/O activity, if possible.
3. Go to sync and change from Global Copy to Metro Mirror.
4. After all Metro Mirror volume pairs are in FULL DUPLEX quiesce or shut down the application servers that access the zGM primary volumes.

5. End the zGM session at a session level to drain the primary subsystem side files and then end the zGM session.
6. Fail over to the new storage system volumes, which are the Metro Mirror secondary volumes. This failover places the Metro Mirror secondary volumes into the state of PRIMARY SUSPENDED. Now these volumes are accessible by the host systems.
7. Because SDM1 is connected to all volumes in Site 1, prepare SDM1 to restart in its corresponding LPAR through these steps:
   a. Vary the old zGM primary volumes in SDM1 LPAR system OFFLINE.
   b. Vary ONLINE the zGM primary Volumes in the new storage system, which are also Metro Mirror primary volumes in suspended state.
8. Now the zGM primary volumes in the new storage system are available and ready to be used.
9. Start the SDM1 session again.
10. Add all zGM volume pairs with the parameter NOCOPY to the zGM session.
11. You can then restart the application systems at Site 1 with an IODF that defines the old zGM primary volumes as OFFLINE and the volumes in the new storage system as ONLINE. Note that Metro Mirror replicated the zGM primary volumes, including their VOLSER. This action brings the application server up and running, and any potential write now happens to the new storage system.

The volumes in the new storage system remain in Metro Mirror PRIMARY SUSPENDED state to allow a fallback, if necessary, to the old storage server in Site 1.

After you decide that the migration process is finished, remove the Metro Mirror relationships in the new storage system and in the old storage system.

The following sections describe each individual step and provide the relevant commands to process each step.

Establish Global Copy between storage systems in Site 1

As soon as the new storage system is up and running and ready to be used by the systems at Site 1, establish a Global Copy relationship between both storage systems as shown in Figure 8-4 on page 160.

Doing so creates a working Copy Services infrastructure between both storage systems, at least for the duration of the migration. This infrastructure requires Fibre Channel Protocol (FCP). You might not have FCP ports available in both storage systems because you have a pure FICON configuration. In such a case, there are two options:

1. When there are sufficient Host Adapter (HA) ports in the existing storage system, two FICON ports can be temporarily taken away as host connections and changed to FCP ports to provide a link connection for Global Copy. Similar considerations apply to the new storage systems.
2. When there are not enough HA FICON ports are available to be take away, plan for two FCP ports when ordering the new storage system or even install an additional HA in the existing storage system with at least two ports to be set as FCP ports.

The number of required FCP ports for the Global Copy setup also depends on the potential write workload to the existing storage system at Site 1. For example, a 730 TB 4-frame DS8800 with 192 ranks and 1,536 DDMs with 600 GB each have a write potential that might not be sufficiently covered by two FCP PPRC links even if only used for migration. In this case, you need at least four FCP links between the storage systems.

To create a Global Copy configuration in a z/OS environment, use TSO commands to manage DS8000 Copy Services operations. Also, use TSO commands to manage the zGM configuration. Note that zGM cannot be managed anyway through the DS8000 DSCLI nor through the DS8000 Storage Management GUI.

TSO commands tend to be robust and run through FICON channels straight into the DS8000, bypassing the DS8000 HMC and their software stack. Similar considerations apply to ICKDSF, although ICKDSF also does not support zGM.

Establish PPRC paths between storage systems in Site 1 and Site 2

Because you are using TSO commands, all volumes in the new storage system are defined to z/OS with valid device numbers and volume serial numbers, VOLSER.

First, establish PPRC paths. Defining PPRC paths through CESTPATH requires to know the worldwide node name (WWNN) of both storage systems. The WWNN can be obtained through a CQUERY to any device in the storage system as shown in Example 8-1.

Example 8-1 The WWNN of the concerned DS8000s

```
CQUERY   DEVN(X'3B00')
ANTP0090I CQUERY FORMATTED LVL 6
VOLUME REPORT
************** PPRC REMOTE COPY CQUERY - VOLUME ********************
*                                          (PRIMARY)   (SECONDARY) *
*                                         SSID CCA LSS SSID CCA LSS*
*DEVICE   LEVEL      STATE     PATH STATUS  SERIAL#     SERIAL#    *
*------ ---------  ----------  ----------- ---------    ---------  *
* 3B00  .........  SIMPLEX...   INACTIVE   3B01 00 3B   .......... *
* SCH(0) ...............      ........... 0000000KAB31 ............*
*       ...........           ...............                      *
* PATHS SAID DEST STATUS: DESCRIPTION                                *
* ----- --------- ------  -------------------                       *
*   0   ---- ----   00    NO PATH............                       *
*       ---- ----   00    NO PATH............                       *
*       ---- ----   00    NO PATH............                       *
*       ---- ----   00    NO PATH............                       *
*  SUBSYSTEM         WWNN                   LIC LEVEL               *
* -----------  ----------------            -----------             *
* PRIMARY....  5005076303FFD5AA            7.7.40.335              *
**********************************************************************
ANTP0001I CQUERY COMMAND COMPLETED FOR DEVICE 3B00. COMPLETION CODE: 00
```

In the test environment, the device number **3B00** is a device number in the existing DS8000 at Site 1. Determine in the same way the WWNN of the new DS8000 storage system.

Save these WWNNs in a text file so you can easily copy and paste their values in respective commands.

Example 8-1 also reveals information that is required for other PPRC TSO commands. Under the PRIMARY heading, you find the SSID and DS8000 LSS number that are required when establishing Global Copy or Metro Mirror volume pairs. In the test environment, this is SSID **3B01** and LSS number **3B** in the existing DS8000 storage system.

Define PPRC paths between the existing and the new DS8000 storage system as shown in Example 8-2. For more information about syntax and semantic meaning of the command parameters, see *DFSMS Advanced Copy Services*, SC35-0428.

Example 8-2 Establish PPRC path from existing to new DS8000.

```
//* ------------------------------------------------------------ ***
//* ESTABLISH PATHS A -> C                                       ***
//* ------------------------------------------------------------ ***
//CESTPATH  EXEC PGM=IKJEFT01
//SYSTSPRT DD SYSOUT=*
//SYSTSIN  DD *

 CESTPATH DEVN(X'3B00') PRIM(X'3B01' 5005076303FFD5AA X'3B') +
                        SEC (X'A301' 5005076305FFD75A X'A3') +
                        LINK(X'00030201' X'01330234')
/*
//* ------------------------------------------------------------ ***
```

Example 8-2 shows two link connections and, therefore, allows two PPRC paths for every LSS pair. In Figure 8-4 on page 160, this would be the PPRC connection between the existing disk storage system and the new disk storage system.

Establish Global Copy volume pairs in Site 1

Next, establish Global Copy relationships between the existing zGM primary volumes, which then become Global Copy primary volumes, and the secondary Global Copy volume in the new storage system. See Example 8-3.

Example 8-3 Establish Global Copy relationships in Site 1

```
//* ------------------------------------------------------------ ***
//* ESTABLISH PAIRS A -> C (C BEING THE NEW STORAGE SYSTEM)      ***
//* ------------------------------------------------------------ ***
//CESTPAIR  EXEC PGM=IKJEFT01
//SYSTSPRT DD SYSOUT=*
//SYSTSIN  DD *

 CESTPAIR DEVN(X'3B00') PRIM(X'3B01' ZA571 X'00'      X'3B') +
                        SEC (X'A301' CZM21 X'00'      X'A3') +
                        OPTION (XD)                         +
                        MODE   (COPY) ONLINSEC (YES)
..............................
/*
//* ------------------------------------------------------------ ***
```

Example 8-3 shows just one volume pair. Again, check the meaning of the required parameters of the **CESTPAIR** command in *DFSMS Advanced Copy Services*, SC35-0428. In the test environment, SSID **3B01** in **ZA571** represents the existing storage system. SSID **A301** in **CZM21** is the new storage system that is going to receive the zGM primary volumes in the course of the actual volume migration.

As soon as these CESTPAIR commands are submitted through a TSO batch job, the first initial full copy starts to replicate the entire primary volume to the Global Copy secondary volume.

Go to sync at Site 1

When you are about to perform the actual swap to the new storage system, alter the Global Copy state to Metro Mirror and change from asynchronous replication to synchronous replication. This can be done through the same set of **CESTPAIR** commands used before. Copy these commands to a new member and slightly modify each command as shown in Example 8-4.

Example 8-4 Go to sync in Site 1

```
//* ------------------------------------------------------------ ***
//* GOTO SYNC PAIRS A -> C (C BEING THE NEW STORAGE SYSTEM)       ***
//* ------------------------------------------------------------ ***
//CESTPAIR  EXEC PGM=IKJEFT01
//SYSTSPRT DD SYSOUT=*
//SYSTSIN  DD *

 CESTPAIR DEVN(X'3B00') PRIM(X'3B01' ZA571 X'00'      X'3B') +
                        SEC (X'A301' CZM21 X'00'      X'A3') +
                        OPTION (SYNC)                        +
                        MODE   (COPY) ONLINSEC (YES)
..............................
/*
//* ------------------------------------------------------------ ***
```

The only parameter that is changed is the OPTION parameter. This brings all addressed Global Copy volume pairs from PENDING to FULL DUPLEX state.

The applications are not affected yet and keep running.

Prepare for actual volumes swap at Site 1

When you are ready for the actual swap, stop all application servers that access the zGM primary volumes in question.

> **Important:** A HyperSwap through this approach is not possible here because the required zGM incremental resynch support is only possible when GDPS manages such a configuration. An important component in zGM is the storage control session, SCSESS construct in each primary LCU and their associated side file in storage system cache. Global Copy and Metro Mirror are not aware of these structure elements of zGM, and therefore will not migrate these elements to the new storage system. zGM incremental resynch does provide this support, but the interface to this support is not externalized and only available to GDPS managed configurations.

Not having GDPS leaves you with the bare TSO command based possibilities that require an application outage here. You still need to be able to access the volumes through TSO. This might happen from another system like a test or maintenance system, which has all volumes defined in their configuration file.

After all application systems are stopped, you need to end the zGM session and force zGM to completely drain the local side files that are cached on the existing storage system.

As shown in Example 8-5, an **XQUERY** command before ending the zGM DR session might reveal that all changed data is already drained and copied to the zGM secondary Site 2.

Example 8-5 End zGM d/r session

```
//* ------------------------------------------------------------ ***
//* end zGM D/R session                                          ***
//* ------------------------------------------------------------ ***
//XEND1     EXEC PGM=IKJEFT01
//SYSTSPRT DD SYSOUT=*
//SYSTSIN  DD *

     XQUERY SESS1   V(*) DETAIL
     XEND   SESS1   NOVERIFY     DRAIN
/*
//* ------------------------------------------------------------ ***
```

Swap to volumes in new storage system

Example 8-6 shows again a **CESTPAIR** command to fail over to the Metro Mirror secondary volumes.

Example 8-6 Failover to Metro Mirror secondary volumes in Site 1

```
//* ------------------------------------------------------------ ***
//* FAILOVER  PAIRS C -> A (C BEING THE NEW STORAGE SYSTEM)        ***
//* ------------------------------------------------------------ ***
//CESTPAIR  EXEC PGM=IKJEFT01
//SYSTSPRT DD SYSOUT=*
//SYSTSIN  DD *

 CESTPAIR DEVN(X'A300') SEC (X'3B01' ZA571 X'00'      X'3B') +
                        PRIM(X'A301' CZM21 X'00'      X'A3') +
                        OPTION (SYNC)                        +
                        ACTION (FAILOVER) ONLINSEC (YES)
.............................
/*
//* ------------------------------------------------------------ ***
```

Because a failover always addresses secondary volumes, use the same set of **CESTPAIR** commands from before, copy them to a new member, and modify the commands. Note that the **SEC** and **PRIM** parameters are swapped, **PRIM** is now the new storage system, and **SEC** points to the existing storage system.

This action changes the Metro Mirror secondary volumes to the new state of **PRIMARY SUSPENDED**, along with implications such as turning on change recording and allowing the host to now access these volumes.

Figure 8-5 shows a small test configuration with the concerned MVS device numbers and VOLSERs so you can orient yourself when following the MVS system commands in Example 8-7.

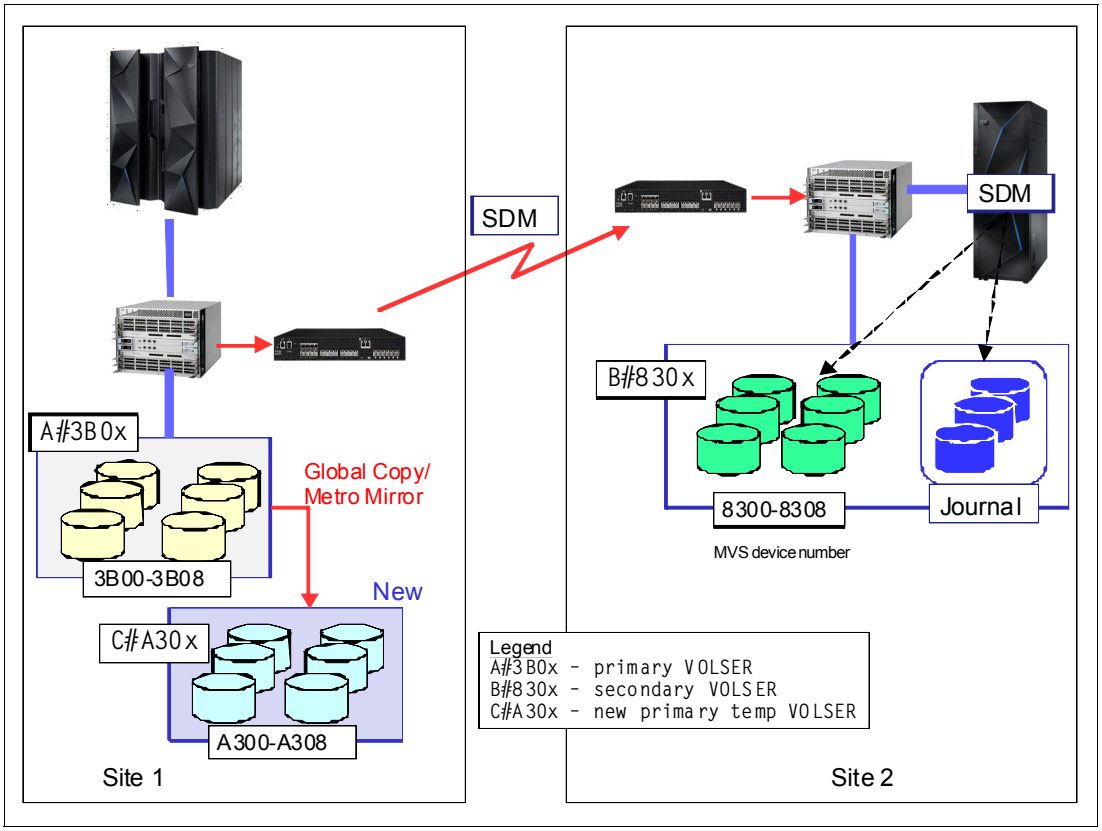

Figure 8-5 Swap to new primary volumes in Site 1 - with device numbers and VOLSERs

Example 8-7 shows a sequence of MVS system commands to finalize the swap and end the migration steps at Site 1.

Example 8-7 MVS command sequence to swap to new storage system in Site 1

```
V 3B00-3B08,OFFLINE
IEE457I 12.14.28 UNIT STATUS 234
UNIT TYPE STATUS        VOLSER     VOLSTATE
3B00 3390 F-NRD                        /RSDNT
3B01 3390 F-NRD                        /RSDNT
................
V A300-A308,ONLINE
IEE457I 12.14.39 UNIT STATUS 236
UNIT TYPE STATUS        VOLSER     VOLSTATE
A300 3390 O             A#3B00     PRIV/RSDNT
A301 3390 O             A#3B01     PRIV/RSDNT
................
DS QD,3B00,9
IEE459I 12.14.51 DEVSERV QDASD 238
 UNIT VOLSER SCUTYPE DEVTYPE       CYL  SSID SCU-SERIAL DEV-SERIAL EFC
03B00 ------ 2107961 2107900      1113  3B01 0175-ZA571 0175-ZA571 BYP
03B01 ------ 2107961 2107900      1113  3B01 0175-ZA571 0175-ZA571 BYP
................
DS QD,A300,9
```

```
IEE459I 12.15.01 DEVSERV QDASD 260
 UNIT VOLSER SCUTYPE DEVTYPE       CYL  SSID SCU-SERIAL DEV-SERIAL EFC
0A300 A#3B00 2107961 2107900      1113  A301 0175-CZM21 0175-CZM21 *OK
0A301 A#3B01 2107961 2107900      1113  A301 0175-CZM21 0175-CZM21 *OK
................
DS QD,8300,9
IEE459I 12.14.57 DEVSERV QDASD 240
 UNIT VOLSER SCUTYPE DEVTYPE       CYL  SSID SCU-SERIAL DEV-SERIAL EFC
08300 B#8300 2107961 2107900      1113  8301 0175-CYM31 0175-CYM31 *OK
08301 B#8301 2107961 2107900      1113  8301 0175-CYM31 0175-CYM31 *OK
```

The actions in Example 8-7 on page 166 are as follows:

- Vary the old zGM primary volumes offline.
- Address the volumes in the new storage system and turn them online. These are the volumes with the A3xx MVS addresses in the test example.
- Because the VOLSER contains the original MVS device numbers and an indication on where the volumes are (in this example, volume A), the subsequent DEVSERV commands prove the actual swap to the new storage system occurred. The last DEVSERV command in Example 8-7 on page 166 displays the zGM secondary volumes that did not change their state or their location.

Restart zGM session and application hosts

Now you are ready to start zGM again as well as the application hosts, if required adjustments have been made to the z/OS configuration files in the application server and in the SDM1 host. When the restart of the application servers succeeds, meaning that all data are correct and consistent, you can discard the recorded changes by deleting the suspended Metro Mirror relationship between the existing volumes and the volumes in the new storage system at Site 1.

The restart sequence of zGM is a 2-step process:

- Restart the zGM session
- Add all relevant zGM volumes pairs from before with the NOCOPY parameter to the zGM session.

Example 8-8 shows a TSO in batch example to restart the SDM session.

Example 8-8 Restart zGM session

```
//* ------------------------------------------------------------ ***
//XSTART1   EXEC PGM=IKJEFT01
//SYSTSPRT DD SYSOUT=*
//SYSTSIN  DD DDNAME=SYSIN
 XSTART SESS1 SESSIONTYPE(XRC) ERRLVL(SESSION) HLQ(TEAM31)
 XSET   SESS1 TIMEOUT(00.00.20)
 XSET   SESS1 PAGEFIX(999)
 XSET   SESS1 SYNCH(45)
 XSET   SESS1 SCSYNCH(4,4)
 XSET   SESS1 PRIORITY(LOAD)
 XSET   SESS1 UTILITY(FIX)
 XQUERY SESS1 SET
 XQUERY SESS1 ENV(PARM)
/*
```

```
//* ============================ JOB END ======================== ***
//
```

You must end the SDM session before a conventional start of SDM. Example 8-8 on page 167 has some additional settings that are useful to consider. Good practice is to query the session parameters so that you can document them.

The next step is to add all volume pairs from before going back into the zGM session. You can do this without replicating the volumes again if you can guarantee that change happened to the A-volumes in Site 1 after the application server shut down in Site 1.

Example 8-9 shows a concerned job with TSO in batch processing. Besides the actual **XADDPAIR** command with the parameter **SUSPENDED**, a few query commands will document which volume pairs are in the actual configuration to compare with the session configuration from before.

Example 8-9 Add all zGM volume pairs back to the session

```
//* ------------------------------------------------------------ ***
//XADDSUSP  EXEC PGM=IKJEFT01
//SYSTSPRT  DD SYSOUT=*
//SYSPRINT  DD SYSOUT=*
//SYSTSIN   DD *


 XADDPAIR SESS1  SUSPENDED          /* add back all suspended volumes */

 XQUERY   SESS1  CONFIG
 XQUERY   SESS1  SSID(ALL)   V(*)
 XQUERY   SESS1  STOR
 XQUERY   SESS1  V(*) DETAIL
/*
//* ============================ JOB END ======================== ***
//
```

After all volume pairs are back in session, you can restart all concerned application servers that access zGM primary volumes.

You can verify that all volume pairs are successfully contained in the zGM session when the zGM message shown in Example 8-10 appears in the SDM system log.

Example 8-10 Message that all volume pairs in zGM session are processed

```
ANTX8120I xx VOLUMES IN SESSION(SESS1) ARE NOW DUPLEX
```

This message ANTX8120I is only issued when all volume pairs are processed. As a side effect, you get the actual number **xx** of all volumes pairs in the concerned zGM session.

Now y restart all application servers again. Eventually you will clean up the configuration and dynamically adjust the configuration reflecting the new storage system.

Cleanup Metro Mirror relationships

After you consider the move as successful and finished, there is no need to keep the storage system, the old and the new storage system in Metro Mirror **PRIMARY SUSPENDED** state.

Example 8-11 shows how to end the Metro Mirror state in the old primary storage system. Here it is the first volume in LSS **3B**. You must also specify the former auxiliary storage system and its related LSS and device number.

Example 8-11 Clean up suspended Metro Mirror relationship in old primary storage system

```
//* ------------------------------------------------------------ ***
//CDELPAIR  EXEC PGM=IKJEFT01
//SYSTSPRT DD SYSOUT=*
//SYSTSIN  DD *

 CDELPAIR DEVN(X'3B00') PRIM(X'3B01' ZA571 X'00'      X'3B') +
                        SEC (X'A301' CZM21 X'00'      X'A3')
/*
//* ============================ JOB END ======================== ***
//
```

You also need to end the PRIMARY SUSPENDED state in the new zGM primary storage system. You can follow the explicit device numbers in Figure 8-5 on page 166, which match all the examples.

Note in Example 8-12. the parameters SEC and PRIM are swapped due to the failover from before to keep the editing effort as little as possible.

Example 8-12 Clean up suspended Metro Mirror relationship in new zGM primary storage system

```
//* ------------------------------------------------------------ ***
//CDELPAIR  EXEC PGM=IKJEFT01
//SYSTSPRT DD SYSOUT=*
//SYSTSIN  DD *

 CDELPAIR DEVN(X'A300') SEC (X'3B01' ZA571 X'00'      X'3B') +
                        PRIM(X'A301' CZM21 X'00'      X'A3')

/*
//* ============================ JOB END ======================== ***
//
```

This concludes the swap to the new storage system in Site 1.

8.2.4 Replacing auxiliary storage system within a zGM DR configuration

This section describes the replacement of a auxiliary storage system at Site 2.

Because zGM is in place in the DR configuration, it is natural to use zGM also for migrating from the existing auxiliary storage system to the new auxiliary storage system.

Figure 8-6 shows a cascaded zGM configuration that enables this configuration to migrate to the new storage system without impacting application availability at Site 1.

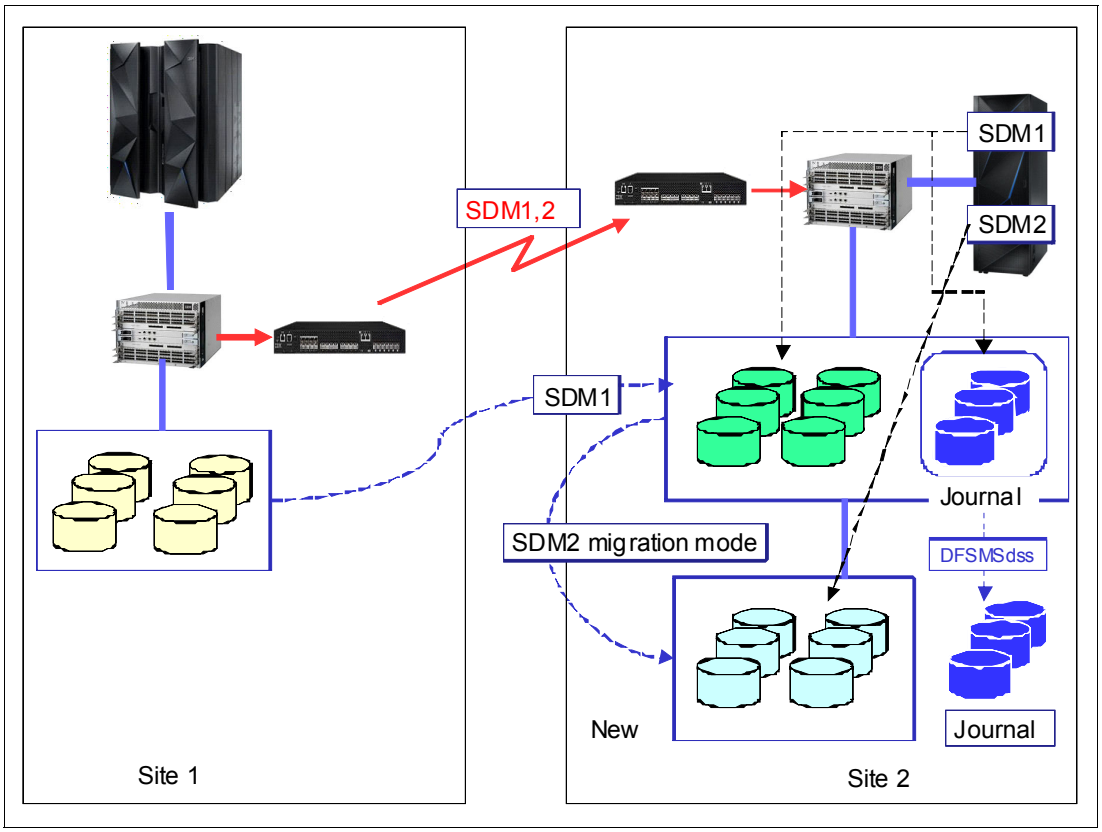

Figure 8-6 Replace storage system in site 2 using zGM

DR readiness might be affected for a brief period when finally swapping in Site 2 from the existing storage system to the new storage system. This is expressed through a session suspend of the D/R zGM session **SDM1** to free up allocations from **SDM1** on devices in the existing auxiliary storage system to allow the actual volume swap and moving the Journal, State, and Control data set to the new storage system. But the applications in Site 1 can continue to run and are not affected.

Actual migration steps to replace auxiliary storage system

After the new storage system is up and running, the z/OS I/O configuration in SDM2 LPAR at Site 2 must be updated to reflect the new storage system. This requires new LCUs, new SSIDs, and new device numbers from an z/OS viewpoint. In a switched environment with FICON switches or FICON directors, additional ports might be required to connect the new storage system to the host at Site 2. The same requirement can also apply to FICON channels from the host to the FICON switches or FICON directors.

1. All potential new zGM secondary volumes must be initialized with a valid and unique volume serial number, VOLSER in SDM2 LPAR, so that these volumes can be varied online to the hosts at Site 2, at least to the LPAR that is hosting SDM2. See Figure 8-6.
2. Define a STATE data set for the zGM migration session SDM2. A STATE data set is only required when running zGM in migration mode, **SESSIONTYPE(MIGRATE)**. Note that zGM in migration mode cannot use coupled SDMs. This constraint can require more planning when a large number of volumes in a coupled DR SDM configuration must be migrated. In this context, the hint is again to provide as much real storage as possible to the SDM2 session to allow a fast move of the data.

3. Start zGM session SDM2 in migration mode and provide sufficient real storage to SDM2, in the range of 1+GB of real storage for SDM2. A typical maximum is 2 GB of real storage to be available and fixed for the SDM2 LPAR. Although CPU resources are not as crucial as real storage. You must still allow sharing of a sufficient number of CPs without capping for the SDM2 LPAR to keep the migration time as short as possible.
4. Using **XADD** commands, add zGM volume pairs that contain all secondary volumes of SDM1 in the existing storage system as primary volumes in SDM2. The corresponding secondary devices are in the new storage system. Explicit SDM utility devices are usually used, **UTILITY(FIX)** in zGM DR configurations. Use the same device numbers within an LCU for fixed utility devices in the new storage system that are used in the existing auxiliary storage system by SDM2.
5. After all zGM SESS2 volume pairs are in FULL DUPLEX state and the message ANTX8120I appeared in MVS system log where SDM2 runs, you are ready for the swap to the new storage systems or to the volumes in the new storage system.
6. To relieve allocations from SDM1 in the current auxiliary storage system, suspend the SDM1 based session on a session level. This leaves the primary storage system with change recording on for all potential updates to the primary volumes.
7. As soon as SDM1 based session is suspended, perform the actual swap as follows:
   a. Move State, Control, and Journal data sets from SDM1 based session to the new volumes in the new storage system. This test case uses DFSMSdss (ADRDSSU) to do the move.
   b. Fail over all volumes in SDM2 based session to the new storage system. This is achieved through an XEND of the SDM2 based session followed by an XRECOVER to the zGM secondary volumes in the new storage system. When managing this fail over process out of the LPAR where SDM2 runs, you must verify that the old zGM secondary volumes are offline before the actual XRECOVER. This is because the XRECOVER will sign the volumes in the new storage system with the same VOLSERs from the SDM1 based secondary volumes. This leads to identical VOLSERs in the old storage system and in the new storage system in Site 2.
8. After all volumes are recovered through SDM2 in the new storage system in Site 2, you can restart SDM1 session.
9. Before restarting all zGM SDM1 volume pairs and resynchronizing them, adjust first the new SSID in the new storage system in Site 2. This is achieved through another XSET command with the REFRESHS parameter.
10. Finally, used XADD to add all zGM SMD1 based volumes pairs back to the DR session.

This concludes the logical steps that are required when migrating zGM secondary volumes or auxiliary storage systems to new zGM secondary volumes or to new zGM auxiliary storage systems.

The following paragraphs outline these steps in a small test configuration and show the required zGM commands and parameters.

Discrete migration steps for Site 2

The following paragraphs outline each individual migration step and show the relevant commands with their respective command parameters (Figure 8-7).

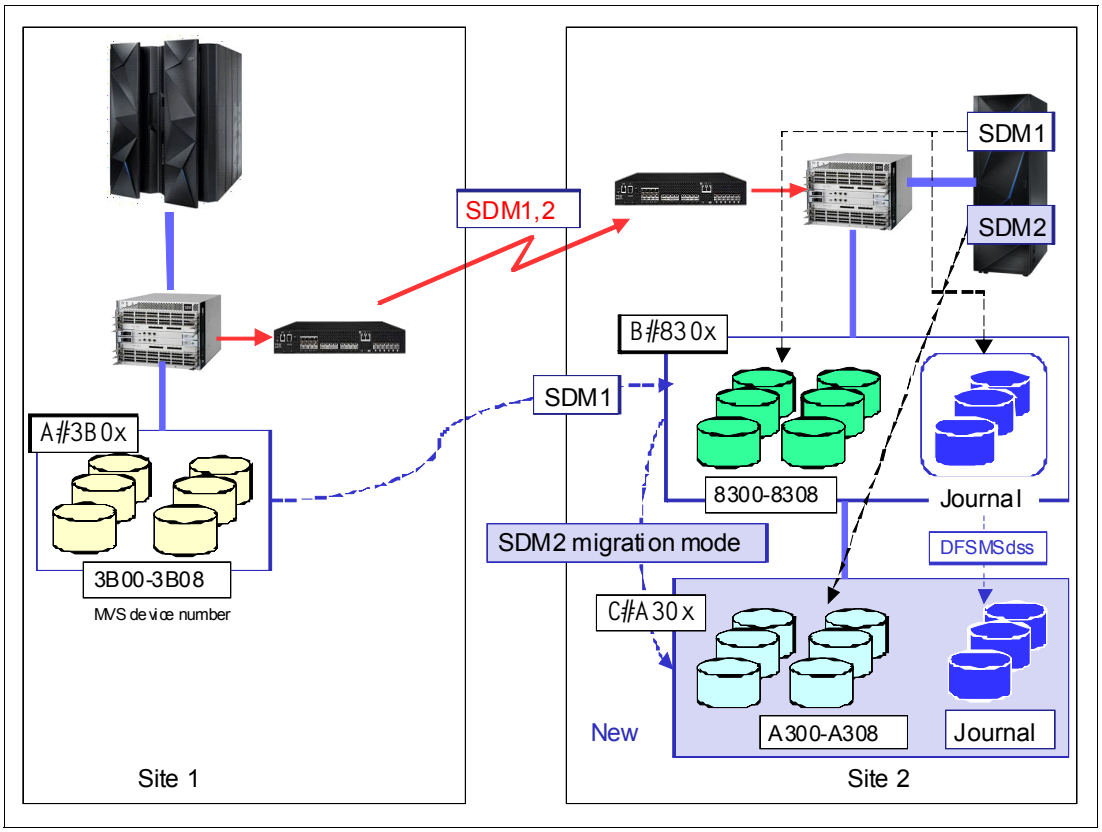

Figure 8-7 Cascaded zGM test configuration with VOLSERs and device numbers

Figure 8-7 has three storage systems:

1. The storage system in Site 1 is the zGM primary storage system. In this small test configuration, the zGM primary volumes are labeled with a leading **A** followed by the corresponding MVS device number, separated by a # between the **A** and the device number.
2. The existing storage system is at site 2 with the SDM1 based auxiliary storage system and its VOLSERs. Here a leading **B** indicates Site 2 (or Site B), followed by a separating # sign and the MVS device number. For example, **B#8300** as VOLSER and its device number **8300**.
3. New storage system in Site 2 that is supposed to replace the existing storage system at Site 2. The VOLSERs are labeled with a lead **C** indicating the new secondary volumes in the new storage system. As before, a # separator is between the **C** and the concerned MVS device number, such as **C#A300** as VOLSER and its MVS device number **A300**.

Define STATE Data Set for SESS2

All volumes in the new storage system are initialized using a leading **C** for the new zGM secondary volumes. As seen in Example 8-13, for the new State Data Set for **SESS2,** use a volume **XCA306** in the new storage system on device number **A306**. This is just the STATE data set for SESS2 and this STATE data set is not related to SESS1.

Example 8-13 Define STATE data set for SESS2

```
//* ------------------------------------------------------------ ***
 //DELETE  EXEC PGM=IDCAMS
 //SYSPRINT  DD SYSOUT=*
 //SYSIN     DD *
    DELETE         TEAM31.XCOPY.SESS2.STATE
    SET MAXCC = 0
 /*
 //* --------- SEE ALSO PARMLIB MEMBER ANTXIN00 ------------------- ***
 //INIT   EXEC PGM=IEFBR14
 //STATE    DD DSN=TEAM31.XCOPY.SESS2.STATE,DISP=(,CATLG),
 //            STORCLAS=NONSMS,
 //            UNIT=3390,VOL=SER=XCA306,SPACE=(CYL,(010,00)),
 //            RECFM=FB,LRECL=04096,BLKSIZE=04096,DSORG=PO,
 //            DSNTYPE=LIBRARY
 //* --------------------- JOB END -------------------------------- ***
 //
```

Start SESS2 in migration mode.

Example 8-14 shows the zGM command to start the migration session, named SESS2.

Example 8-14 Start migration session SESS2

```
//* ------------------------------------------------------------ ***
//XSTART2 EXEC PGM=IKJEFT01
//SYSTSPRT  DD SYSOUT=*
//SYSTSIN   DD *
 XSTART SESS2 SESSIONTYPE(MIGRATE) ERRLVL(VOLUME) HLQ(TEAM31)
 XSET   SESS2 TIMEOUT(00.00.20)
 XSET   SESS2 PAGEFIX(999)
 XSET   SESS2 SYNCH(45)
 XSET   SESS2 SCSYNCH(8,8)
 XSET   SESS2 PRIORITY(LOAD)
 XSET   SESS2 UTILITY(FIX)
 XQUERY SESS2 SET
 XQUERY SESS2 ENV(PARM)
/*
//* ------------------------  JOB END --------------------------- ***
```

Usually migration sessions are fine with the scope of **VOLUME** concerning the **ERROR LEVEL** of the session.

In addition to the session start command, the example shows typical settings. For more information about these settings, see *DFSMS Advanced Copy Services*, SC35-0428.

Add zGM volume pairs to SDM2 session SESS2

The next step is to establish all zGM volumes pairs to the migration session SESS2 with the following volumes, as shown in Example 8-15:

- SESS2 primary volumes are all SDM1 secondary volumes, which in this example are all volumes that start with B# in their VOLSERs. See Figure 8-7 on page 172.
- SESS2 secondary volumes, which are all corresponding volume that start with C# in their VOLSERs and going to be potentially the new secondary volumes for SESS1. See Figure 8-7 on page 172.

Example 8-15 Add zGM copy pairs to migration session SESS2

```
//* ------------------------------------------------------------ ***
//XADDUD    EXEC PGM=IKJEFT01
//SYSTSPRT DD SYSOUT=*
//SYSTSIN  DD *
     XADD     SESS2  V(XU8307,XRCUTL)
/*
//* ------------------------------------------------------------ ***
//XADDPAIR  EXEC PGM=IKJEFT01
//SYSTSPRT DD SYSOUT=*
//SYSTSIN  DD *
     XADDPAIR SESS2  V(B#8300,C#A300)
     XADDPAIR SESS2  V(B#8301,C#A301)
   . . . . . . .
/*
//* ================  JOB END ==================================== ***
//
```

Example 8-15 shows two types of zGM volume pairs:

1. Because fixed zGM utility volumes are used, establish those volumes in the existing storage system and in each relevant LCU. You might need to configure such a volume beforehand. It might be a small volume because all that is required is a dedicated MVS device address for SDM2 to read all data from the SESS2 primary volumes. In our test case, it is a volume with the device address **8307** and **XU8307** as VOLSER. It is a volume in the zGM primary storage system.
2. Next step contains all zGM volume pairs that must be migrated to the new storage system. In our test configuration, these are all **B**-volumes as primary volumes and all corresponding **C**-Volumes as secondary volumes. See Figure 8-7 on page 172.

The job that is shown in Example 8-15 starts to copy all SDM1 secondary Volumes to the SDM2 secondary volumes. All applications in Site 1 remain running and available.

All SESS2 volume pairs are in full duplex mode as soon as the message **ANTX8120I** appears in the SYSLOG where SDM2 runs. You are now ready to swap to the new storage system.

Suspend DR zGM session SESS1

To free up SDM1 allocations and allow it to swap its secondary devices to the new storage system, suspend the SDM1 session SESS1. See Figure 8-8 on page 175. This stops SDM1 from fetching the data from the SDM1 primary storage system in Site 1.

Example 8-16 shows the concerned TSO command to suspend the D/R session on a session level.

Example 8-16 Session suspend of D/R session SESS1

```
//* ************************************************************ ***
//* SUSPEND D/R session SESS1                                    ***
//* ************************************************************ ***
//XSUSP   EXEC PGM=IKJEFT01
//SYSTSPRT  DD SYSOUT=*
//SYSPRINT  DD SYSOUT=*
//SYSTSIN   DD DDNAME=SYSIN

     TIME
     XSUSPEND  SESS1 TIMEOUT(00.00.25)  NOVERIFY  IMMEDIATE
     TIME

/*
//* ============================ JOB END ======================== ***
//
```

Session SESS1 and data mover SDM1 end as shown in Figure 8-8. But the application at Site 1 can continue to update the SDM1 primary volumes at Site 1. The DS8000 provides adequate mechanisms to record all subsequent changed tracks for a later resync process.

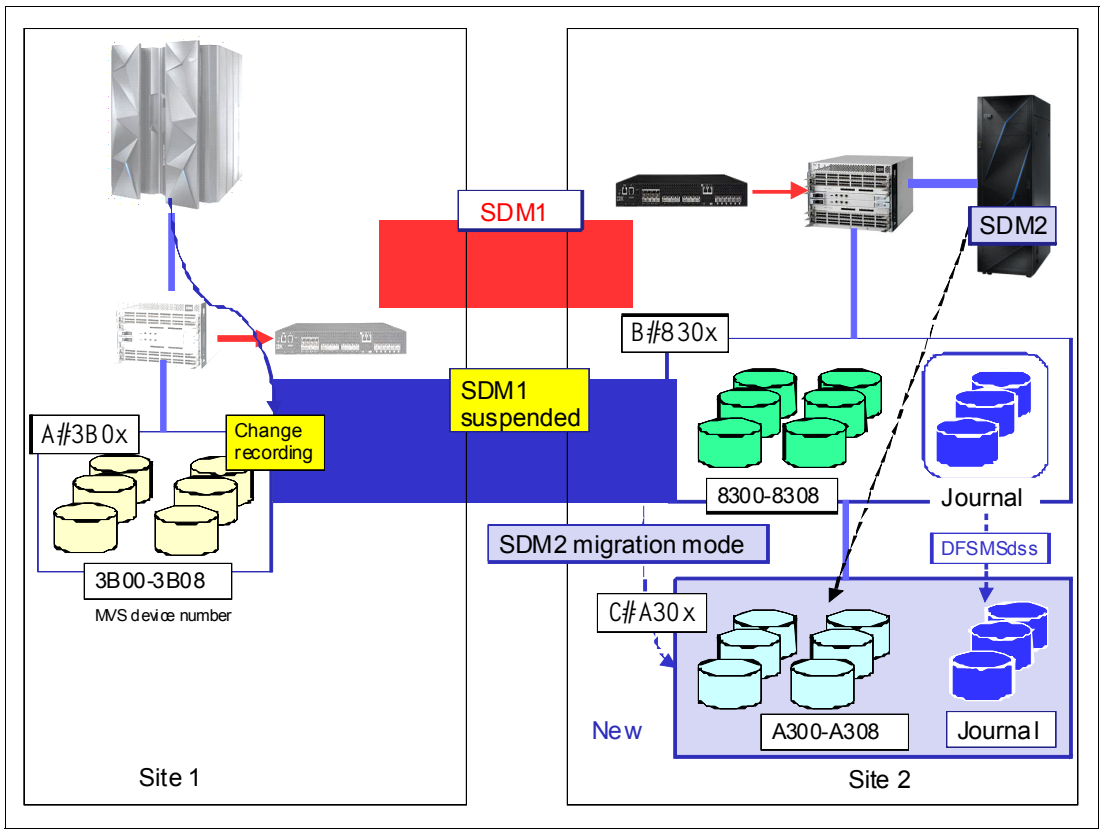

Figure 8-8 D/R session SDM1 suspended, change recording on storage system level in Site 1

The migration session SDM2 is still draining the changed data from the existing storage system, here from the B# volumes, as shown in Figure 8-8 on page 175.

Because a session suspend to SDM1 ends the SDM1 address space, all concerned allocations are also relieved and you can move the Journal environment from the existing storage system to the new storage system.

Move SDM1 Journal environment

You can use DFSMSdss to move the SDM1 control and state data set as well as entire journal volumes.

This can happen with SDM2 still being active although most likely by now all data on **B#** and **C#** volumes are identical because no new data are arriving from Site 1.

Example 8-17 shows an example on how to move all related SDM1 journal and control data to the new storage system. In this example, device numbers **830x** are in the existing storage system and device numbers **A30x** are device numbers in the new storage system:

- XJ8305: Volume in the existing storage system that contains journal data sets related to SDM1. Usually there is more than one journal volume. So add all relevant journal volumes to this job step as shown in Example 8-17.
- XC8306: Volume in the existing storage system that contains State and Control data sets related to SDM1, and eventually also potential Couple and Master data sets.
- XJA305: Volume in the new storage system to house the journal data sets related to SDM1. Again, add the same number of potentially new Journal volumes as coming from the existing storage system.
- XCA306: Volume in the new storage system that is going to house the moved state and control data sets related to SDM1, and eventually also potential couple and master data sets.

Example 8-17 Move SDM1 Journal and Control environment

```
//* ------------------------------------------------------------ ***
//COPY1   EXEC PGM=ADRDSSU
//SYSPRINT  DD SYSOUT=*
//IN1       DD UNIT=SYSDA,VOL=SER=XJ8305,DISP=SHR
//IN2       DD UNIT=SYSDA,VOL=SER=XC8306,DISP=SHR
//OUT1      DD UNIT=SYSDA,VOL=SER=XJA305,DISP=SHR
//OUT2      DD UNIT=SYSDA,VOL=SER=XCA306,DISP=SHR
//SYSIN     DD *
  COPY INDD(IN1) OUTDD(OUT1) ALLD(*) ALLX CANCELERROR COPYVOLID
  COPY INDD(IN2) OUTDD(OUT2) ALLD(*) ALLX CANCELERROR COPYVOLID
/*
//* ================  JOB END ==================================== ***
//
```

You can also adjust the placement of journal, state, control, master, and couple data sets using ADRDSSU as shown in Example 8-18.

Example 8-18 Move SDM1 Control data set

```
//* ------------------------------------------------------------ ***
 //COPY1   EXEC PGM=ADRDSSU
 //SYSPRINT  DD SYSOUT=*
 //IN2       DD UNIT=SYSDA,VOL=SER=XC8306,DISP=SHR
 //OUT2      DD UNIT=SYSDA,VOL=SER=XCA306,DISP=SHR
 //SYSIN     DD *
   COPY DS(INC(                                          -
           TEAM31.XCOPY.SESS1.CONTROL                    -
                                 ))                      -
           ALLD(*) ALLX                                  -
           LIDD (IN2)                                    -
           OUTDD(OUT2) STORCLAS(NONSMS)                  -
           CAT DEL TOL(ENQF) WAIT(0,0)
 /*
 //* ================  JOB END ==================================== ***
 //
```

Example 8-18 shows how to move the control data set related to SDM1 through its session name SESS1.

> **Note:** The only draw back in this test environment relates to the lack of system-managed storage volumes. In a system-managed storage environment, you might have to use a Storage Class with the **`guaranteed space`** attribute to force the placement of zGM related data sets on to certain volumes.

Similarly, you can move other SDM1 based data sets like journal, state, couple, and master data sets.

Fail over all volumes from SDM2 session SESS2

Complete these dedicated steps to perform the actual failover to the new storage system:

1. End SDM2 session SESS2
2. Recover all SESS2 secondary volumes.

These two steps are run out of the LPAR where SDM2 runs. This requires you to vary the SDM2 primary volumes OFFLINE before Step 2 is run and recovering the SESS2 secondary volumes.

End zGM migration session SESS2

Before the actual swap to the new storage system in Site 2, end the zGM migration session SESS2 as shown in Example 8-19.

Example 8-19 End migration session SESS2/SDM2

```
//* ------------------------------------------------------------ ***
//XEND2     EXEC PGM=IKJEFT01
//SYSTSPRT DD SYSOUT=*
//SYSTSIN  DD *

     XQUERY SESS2   CONFIG
     XQUERY SESS2   V(*) DETAIL
```

```
     XQUERY SESS2   STOR

     XEND   SESS2   NOVERIFY     DRAIN
/*
//
```

Good practice is to document the actual configuration through a few query commands before the XEND command, as shown in Example 8-19 on page 177.

XEND ends the zGM migration session SESS2 and you have a configuration state as Figure 8-9 shows. None of the SDMs is active any longer. This configuration is now ready to recover the zGM secondary volumes from SESS2/SDM2.

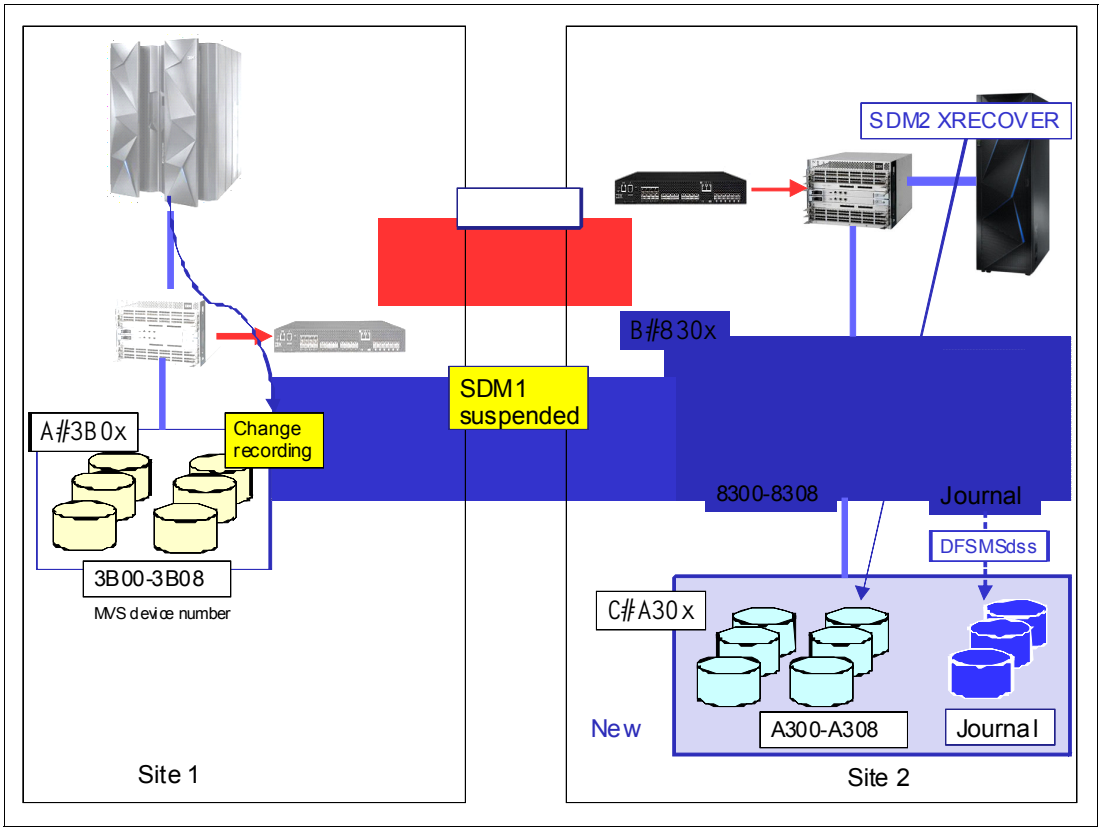

Figure 8-9 Recover zGM migration session SESS2/SDM2 to prepare volumes in new storage system

Because zGM **XRECOVER** will relabel the **C#** volumes with the label (= VOLSER) from the **B#** volumes, you must vary OFFLINE the old **B#** volumes on the MVS device addresses **830x** before the **XRECOVER**.

Recover to new storage system

To move to the new storage system, an **XRECOVER** is needed to restart SDM2 and prepare the zGM secondary volumes to become available with the VOLSER of the zGM primary volumes.

Before the actual **XRECOVER** of SESS2, set the stage in Site 2 because **XRECOVER** restarts the session SDM2/SESS2 and in the course of **XRECOVER** SDM2 relabels the SDM2 secondary volumes with the VOLSER from the SESS2 primary volumes. For that reason, you must first vary the SDM2 primary volumes OFFLINE to allow the recovered SDM2 secondary volumes to become ONLINE after the XRECOVER.

See Example 8-20.

Example 8-20 Recover zGM migration session SESS2 / SDM2

```
//  V 8300-8306,OFFLINE
//  DS P,8300,8
//  DS P,A300,8
//* ------------------------------------------------------------ ***
//* 83xx - SESS1 old secondary devices                           ***
//* A3xx - SESS1 new secondary devices                           ***
//* ------------------------------------------------------------ ***
//XRECOVER EXEC PGM=IKJEFT01
//SYSTSPRT DD SYSOUT=*
//SYSPRINT DD SYSOUT=*
//SYSTSIN  DD *

     XRECOVER  SESS2   HLQ(TEAM31)

     CALL 'TEAM31.L.LOAD(WAIT)' '/5000'
     XQUERY    SESS2
     XQUERY    SESS2 V(*) DETAIL
/*
//
```

XRECOVER restarts SDM2/SESS2 just to relabel the volumes with MVS device addresses **A30x** by the label/VOLSER of the MVS device addresses **830x**. See Figure 8-9 on page 178. All related information takes SDM2 out of its state data set.

After the recover process is done, SDM2 will end again.

Restart DR session SDM1/SESS1

Now you can restart the suspended /R session SESS1/SDM1, as shown in Example 8-21.

Example 8-21 Restart D/R session SDM1/SESS1

```
//* ------------------------------------------------------------ ***
//XSTART1 EXEC PGM=IKJEFT01
//SYSTSPRT  DD SYSOUT=*
//SYSTSIN   DD DDNAME=SYSIN
 XSTART SESS1 SESSIONTYPE(XRC) ERRLVL(SESSION) HLQ(TEAM31)
 XSET   SESS1 TIMEOUT(00.00.20)
 XSET   SESS1 PAGEFIX(999)
 XSET   SESS1 SYNCH(45)
 XSET   SESS1 SCSYNCH(4,4)
 XSET   SESS1 PRIORITY(LOAD)
 XSET   SESS1 UTILITY(FIX)
 XQUERY SESS1 SET
 XQUERY SESS1 ENV(PARM)
/*
//* ============================ JOB END ======================== ***
//
```

Now zGM manages its secondary volumes not only through their VOLSER, but also through the corresponding SSIDs of the related LCUs where the zGM secondary volumes are.

SDM1 still has the old SSID of its secondary LCUs in its state data set.

Example 8-22 shows the zGM session SDM1/SESS1 secondary volumes in the old storage system at Site 2 that shows the SSID of **8301**. You do not see the VOLSER because you varied device numbers **830x** OFFLINE.

Example 8-22 zGM SESS1/SDM1 secondary volumes

```
DS QD,8300,9                          <<<< old SESS1 secondary devices
IEE459I 16.14.05 DEVSERV QDASD 150
 UNIT VOLSER SCUTYPE DEVTYPE       CYL  SSID SCU-SERIAL DEV-SERIAL EFC
08300 ------ 2107961 2107900      1113  8301 0175-CYM31 0175-CYM31 BYP
08301 ------ 2107961 2107900      1113  8301 0175-CYM31 0175-CYM31 BYP
 . . . . .

****      x DEVICE(S) MET THE SELECTION CRITERIA
****      0 DEVICE(S) FAILED EXTENDED FUNCTION CHECKING


DS QD,A300,9                          <<<< new SESS1 secondary devices
IEE459I 16.14.12 DEVSERV QDASD 152
 UNIT VOLSER SCUTYPE DEVTYPE       CYL  SSID SCU-SERIAL DEV-SERIAL EFC
0A300 B#8300 2107961 2107900      1113  A301 0175-CZM21 0175-CZM21 *OK
0A301 B#8301 2107961 2107900      1113  A301 0175-CZM21 0175-CZM21 *OK
 . . . . .

****      x DEVICE(S) MET THE SELECTION CRITERIA
****      0 DEVICE(S) FAILED EXTENDED FUNCTION CHECKING
```

Also note in Example 8-22, under device addresses **A30x**, the VOLSER of the original zGM secondary volumes related to SDM1/SESS1 with the SSID of **A301**.

An attempt to add the suspended zGM volume pairs of zGM DR session SESS1/SDM1 would fail because the SSIDs do not match what SDM1 has in its state data set.

Before adding the suspended zGM pairs in SDM1/SESS1, alter the SSID for the SDM1 secondary volumes.

Use the XSET command as shown in Example 8-23.

Example 8-23 Alter the SSID for the zGM secondary volumes

```
ANTL8800I  XSET   SESS1 REFRESHS

ANTX8140W XSET REFRESHS IGNORED FOR SESSION(SESS1).  UPDATE ENABLEREFRES
HS IN PARMLIB.
```

Example 8-23 shows that altering zGM configuration information requires being enabled through the related zGM PARMLIB parameter.

To understand what is defined in the zGM PARMLIB, list the content through an XQUERY command as Example 8-24 shows.

Example 8-24 Verify PARMLIB content

```
ANTL8800I XQUERY SESS1 ENV(PARM)
ANTQ8200I XQUERY STARTED FOR SESSION(SESS1) ASNAME(ANTAS001) 808
ANTQ8202I XQUERY ENVIRONMENT_PARM REPORT - 001
ANTQ8251I NAME                    VALUE  NAME                    VALUE
```

```
ANTQ8203I ------------------------------------------------------------
ANTQ8253I zIIPEnable                 NO  MonitorOutput              ON
........
ANTQ8253I EnableREFRESHS             NO  ReleaseFixedPages          NO
....
ANTQ8253I MiscLow                     2
ANTQ8203I ------------------------------------------------------------
```

Per the default, altering the secondary configuration is not permitted.

You need to update the appropriate zGM PARMLIB member ANTXIN00, as illustrated in Example 8-25.

Example 8-25 Alter zGM PARMLIB member ANTXIN00

```
VOLINIT -
    HaltThreshold(5120)            -
    InitializationsPerPrimary(8)   -
    InitializationsPerSecondary(8) -
    MaxNumberInitializations(45)   -
    TracksPerRead(12)              -
    TracksPerWrite(12)             -
    EnableRefeshs(YES)
```

A subsequent zGM query confirms the update of the PARMLIB member ANTXIN00.

Note: To refuse altered zGM PARMLIB information, you can recycle the address space **ANTAS00** through an MVS **CANCEL** command. This command triggers an immediate restart of **ANTAS00**. During the restart of **ANTAS00**, the zGM PARMLIB information is also refreshed.

XEND the session SDM1/SESS1 and restart SDM1/SESS1 as shown in Example 8-21 on page 179.

Example 8-26 shows the basic results after restarting SDM1/SESS1.

Example 8-26 Restart SDM1/SESS1 again after altering zGM PARMLIB member ANTXIN00

```
ANTL8800I  XSTART SESS1 SESSIONTYPE(XRC) ERRLVL(SESSION) HLQ(TEAM31)
 ANTL8800I  XSET   SESS1 TIMEOUT(00.00.20)
 ANTL8800I  XSET   SESS1 PAGEFIX(999)
 ANTL8800I  XSET   SESS1 SYNCH(45)
 ANTL8800I  XSET   SESS1 SCSYNCH(4,4)
 ANTL8800I  XSET   SESS1 PRIORITY(LOAD)
 ANTL8800I  XSET   SESS1 UTILITY(FIX)
 ANTL8800I  XQUERY SESS1 SET
 ANTL8800I  XQUERY SESS1 ENV(PARM)

 ANTI1020I PARAMETERS FOUND IN MEMBER ANTXIN00 OF DATA SET SYS1.PARMLIB+
 HAVE BEEN APPLIED
 IKJ56228I DATA SET TEAM31.XCOPY.PARMLIB NOT IN CATALOG OR CATALOG CAN
 NOT BE ACCESSED
 IEF196I IKJ56228I DATA SET TEAM31.XCOPY.PARMLIB NOT IN CATALOG OR
 IEF196I CATALOG CAN NOT BE ACCESSED

 ANTX8120I ALL VOLUMES IN SESSION(SESS1) ARE NOW SUSPENDED
```

```
ANTS8100I XSTART COMPLETE FOR XRC SESSION(SESS1) WITH SESSIONTYPE(XRC),
ERRORLEVEL(SESSION), HLQ(TEAM31).  SESSION NOW
ANTS8100I (CONT) RESTARTED

ANTX8113I XSET(TIMEOUT(00.00.20)) ACCEPTED FOR SESSION(SESS1)
ANTX8113I XSET(PAGEFIX(999)) ACCEPTED FOR SESSION(SESS1)
ANTX8113I XSET(SYNCH(45)) ACCEPTED FOR SESSION(SESS1)
ANTX8113I XSET(SCSYNCH(4,4)) ACCEPTED FOR SESSION(SESS1)
ANTX8113I XSET(PRIORITY(LOAD)) ACCEPTED FOR SESSION(SESS1)
ANTX8113I XSET(UTILITY(FIX)) ACCEPTED FOR SESSION(SESS1)
```

Verify whether the PARMLIB change was successful through another XQUERY as illustrated in Example 8-27.

Example 8-27 Verify ANTXIIN00 PARMLIB content again

```
ANTQ8200I XQUERY STARTED FOR SESSION(SESS1) ASNAME(ANTAS001) 724
ANTQ8202I XQUERY ENVIRONMENT_PARM REPORT - 001
ANTQ8251I NAME                    VALUE  NAME                    VALUE
ANTQ8203I -------------------------------------------------------------
ANTQ8253I zIIPEnable                 NO  MonitorOutput             ON
ANTQ8253I AllowEnhancedReader       YES  MonitorWakeup          10000
 . . . . .
ANTQ8253I EnableREFRESHS            YES  ReleaseFixedPages          NO
 . . . . .
ANTQ8203I -------------------------------------------------------------
ANTQ8201I XQUERY ENVIRONMENT_PARM REPORT COMPLETE FOR SESSION(SESS1)
```

Potentially refreshing the secondary configuration is now enabled.

Example 8-28 shows another attempt to alter the SDM1/SESS1 secondary configuration.

Example 8-28 Alter SESS1/SDM secondary volumes configuration

```
ANTL8800I  XSET   SESS1 REFRESHS
ANTX8141I XSET REFRESHS UPDATED 5 VOLUMES IN SESSION(SESS1)
```

The result is an update of five volumes in SESS1. This is exactly the number of zGM volume pairs in the example configuration.

After the SDM1/SESS1 session restart, issue an XQUERY as shown in Example 8-29.

Example 8-29 Verify the altered zGM secondary configuration

```
ANTL8800I  XQUERY   SESS1  SSID(ALL)   V(*)
 ANTQ8200I XQUERY STARTED FOR SESSION(SESS1) ASNAME(ANTAS001) 891
 ANTQ8202I XQUERY CONFIGURATION REPORT - 003
 ANTQ8281I ---------PRIMARY----------  -----SECONDARY------
 ANTQ8284I      SC SC
 ANTQ8282I SSID SN ID DVCN CCA VOLSER  SSID DVCN CCA VOLSER
 ANTQ8203I -------------------------------------------------------------
 ANTQ8283I 3B01 AA 01 3B00 00  A#3B00  A301 A300 00  B#8300
 ANTQ8283I 3B01 AA 01 3B01 01  A#3B01  A301 A301 01  B#8301
 . . . . . . .
 . . . . . . .
```

Example 8-29 on page 182 shows now under the heading of SECONDARY, the new storage system through the altered SSID. But the correct zGM secondary VOLSER for session SDM1/SESS1 is still persevered.

Example 8-30 is another query to show the still suspended zGM volume pairs of session SESS1/SDM1.

Example 8-30 List suspended session volumes in SESS1/SDM1

```
ANTL8800I  XQUERY   SESS1  V(*) DETAIL
 ANTQ8200I XQUERY STARTED FOR SESSION(SESS1) ASNAME(ANTAS001) 894
 ANTQ8202I XQUERY VOLUME REPORT - 003
 ANTQ8211I PRIM   SEC     ERROR  SYNCH
 ANTQ8212I VOL    VOL     LEVEL    % STA CMD OP -------TIMESTAMP--------
 ANTQ8203I -------------------------------------------------------------
 ANTQ8213I A#3B00 B#8300 SESSION  00 SUS        2014.304 12:51:54.169881
 ANTQ8213I A#3B01 B#8301 SESSION  00 SUS        2014.304 12:51:54.169881
 . . . . . . .
 . . . . . . .
```

All SDM1/SESS1 volume pairs are still suspended and changes are recorded in the primary storage system in Site 1 because you did not stop the applications.

The final step is to add all suspended zGM volume pairs back into session and initiate the resynchronization process. Use a zGM ADD command to start the resync process for all zGM volume pairs in session SESS1. Example 8-31 is an MVS SYSLOG fragment.

Example 8-31 Start resynch process for all zGM volume pairs in SESS1

```
ANTL8800I  XADD     SESS1  SUSPENDED               /* ALL SUSPENDED */

ANTA8101I XADDPAIR COMPLETE FOR VOLUME PAIR(A#3B00,B#8300) FOR
SESSION(SESS1) WITH ERRORLEVEL(SESSION), SCSESSION(AA)
ANTA8101I XADDPAIR COMPLETE FOR VOLUME PAIR(A#3B01,B#8301) FOR
SESSION(SESS1) WITH ERRORLEVEL(SESSION), SCSESSION(AA)
ANTA8101I XADDPAIR COMPLETE FOR VOLUME PAIR(A#3B02,B#8302) FOR
SESSION(SESS1) WITH ERRORLEVEL(SESSION), SCSESSION(AA)
 . . . . . . .
ANTA8124I XADDPAIR SUSPENDED COMPLETED FOR SESSION(SESS1)

ANTI8023I RESYNC INITIALIZATION STARTED FOR XRC VOLUME PAIR(A#3B03,
B#8303)
ANTI8023I RESYNC INITIALIZATION STARTED FOR XRC VOLUME PAIR(A#3B02,
B#8302)
 . . . . . . .
```

The message in Example 8-32 indicates when all zGM volumes pairs are successfully resynchronized and back in zGM DUPLEX state:

Example 8-32 All zGM volumes pairs are resynchronized

```
ANTX8120I 5 VOLUMES IN SESSION(SESS1) ARE NOW DUPLEX
```

Figure 8-10 shows the final result after a successful migration into a new zGM auxiliary storage system at Site 2.

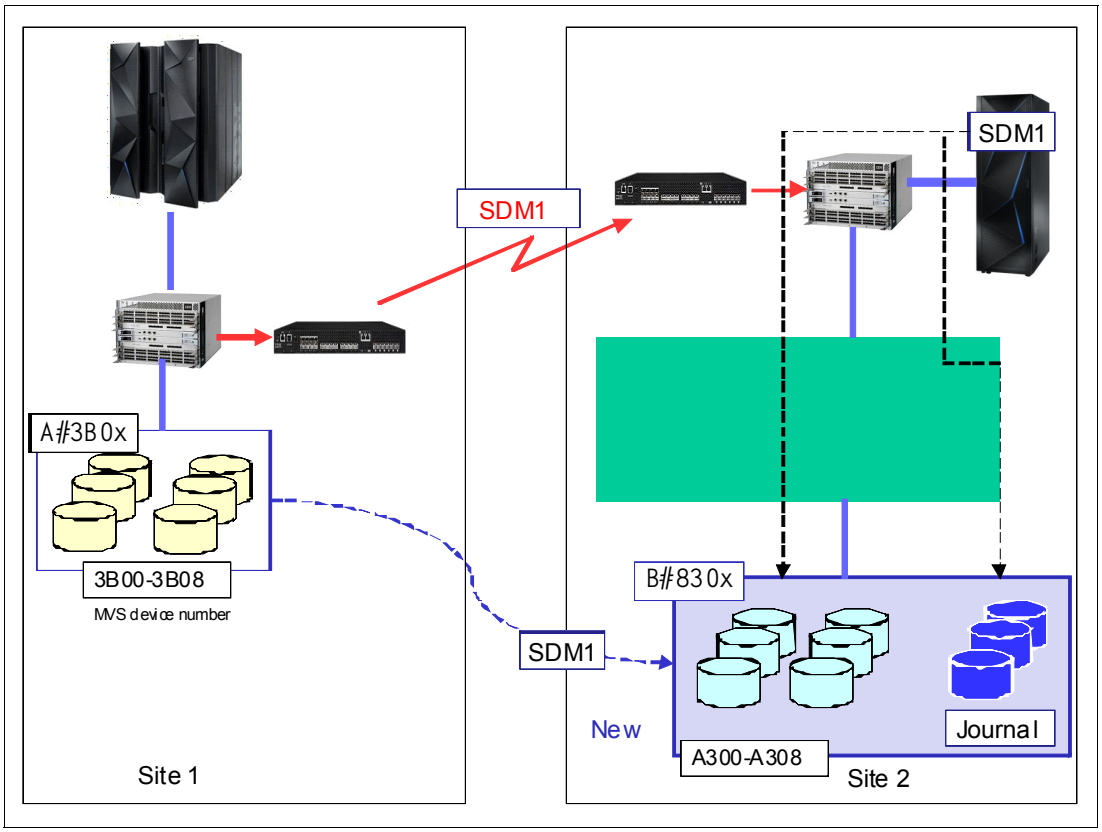

Figure 8-10 Target configuration after the migration in Site 2

In conclusion, you can see that a complete storage system migration at Site 2 is possible without affecting application availability at Site 1. Only during the actual swap in Site 1 is data replication between Site 1 and Site 2 paused, but this period is covered through the primary storage system and its capability to change record all changes during the zGM replication pause period.

Afterwards, a simple resynchronization process replicates just the incremental data that changed during the pause period.

9

Data Migration with GDPS/Metro and z/Global Mirror replication

This chapter discusses the data migration considerations for a three site GDPS/Metro with z/Global Mirror (GDPS/MzGM) environment. The physical location of the participating disk systems does not affect the processes discussed below, but can have an impact on the infrastructure requirements.

This chapter includes the following sections:

- GDPS Terminology
- Initial configuration
- Migration process
- Migration Scenarios

© Copyright IBM Corp. 2015. All rights reserved.

9.1 GDPS Terminology

Here is a brief refresher of the Geographically Dispersed Parallel Sysplex (GDPS) terminology

9.1.1 Site-1 and Site-2

GDPS/PPRC introduces the concept of Site-1 and Site-2, which are typically two physically separate sites or data centers. It can also describe more than two sites where each site is logically designated either a Site-1 or a Site-2 site within a region.

9.1.2 PPRC Controlling System

The PPRC controlling system, also known as the K-system (or Kp-system), participates in the same sysplex as the application systems for which PPRC is being managed. The primary role of the controlling system is to control and manage recovery by using the freeze and HyperSwap technologies. It is also the system through which the operator manages the IPL and shutdown activities of the application systems. Two controlling systems are recommended for availability.

9.1.3 Application system

An application system is any member of the GDPS/PPRC sysplex and is not a controlling system

9.1.4 PPRC Master System

Normally the controlling system operating in Site-2 is the master controlling system. Certain GDPS functions (for example, Configuration) can only be performed from the master system. GDPS/PPRC monitors and tracks the location of the PPRC secondary disk and attempts, wherever possible, to designate the controlling system in this site as the master controlling system.

9.1.5 PPRC Freeze

The PPRC Freeze capability is an attribute assigned to storage system SSID-pairs. All PPRC device pairs that are defined to this SSID-pair automatically inherit this attribute. A typical GDPS/PPRC configuration has many SSID-pairs that collectively form the freeze-group. GDPS/PPRC provides secondary consistency across all devices that are part of the freeze-group. Following the detection of a disk-mirroring-related error, GDPS/PPRC initiates a Freeze. As part of the Freeze process, GDPS/PPRC will temporarily cause all PPRC primary disk to enter what it is commonly known as an “Extended Long Busy” state. With primary I/O temporarily suspended, GDPS/PPRC is able to secure a consistent secondary copy. Because all I/O to the primary devices in the freeze-group is suspended, it is clear that the controlling systems that are responsible for managing the freeze process must not be dependent on any of these devices. It is for this reason that the controlling systems must be configured on disk that is not part of the freeze-group. It is also worth noting that a GDPS HyperSwap event also involves Freeze processing.

It is recommended that all disk devices in the GDPS complex are PPRC mirrored, except for the sysplex couple data sets, and that they are all defined to the freeze-group. Although all

application data should be PPRC mirrored, there are data items within the GDPS/PPRC environment that should not. Essentially data that is pivotal to the survival of the GDPS/PPRC controlling systems must be excluded from the PPRC mirror group. This means that the disk supporting the controlling system's software stacks must be excluded from the PPRC mirror group. Supplementary GDPS monitoring code will also be developed to ensure that the appropriate volumes are included or excluded from the GDPS/PPRC mirror group.

If a Freeze event occurs where no controlling system is active, one or more application systems might be locked out during the primary storage system hardware freeze timeout value (default 120 seconds). During a lockout, the freeze might not complete and therefore the secondary data might be inconsistent. To minimize the risk of a lockout during a freeze event, it is recommended that two controlling systems are configured and active.

When using synchronous disk mirroring, over any distance, two mirroring policies are available in GDPS to determine what happens should the mirroring suspend:

- Freeze and Stop
- Freeze and Go

Freeze and Stop

Freeze all disk activity and stop all LPARs in the sysplex. Does so ensures that no data are lost, but requires a restart of any database and recovery of the structures. If your recovery point objective (RPO) is zero (that is, you cannot tolerate any data loss), select the Freeze and Stop option to reset all production systems. Because all systems that can update the primary volumes are reset while disk I/O is suspended, you are sure that no updates are made to the primary volumes that are not mirrored to the secondary. However, if the event that caused GDPS to take action was a transient event rather than a real disaster, the production service will have been stopped unnecessarily.

Freeze and Go

Freeze all disk activity momentarily, but then continue processing. If the requirement (or acceptability) for RPO is higher than zero, you might decide to allow the production systems to continue operation after the secondary volumes have been protected. This is known as Freeze and Go. In this case, if the trigger was only a transient event, you avoid an unnecessary service outage. However, if the trigger was the first sign of an actual disaster, you might continue operating for some amount of time before all systems actually fail. Any updates that are made to the primary volumes during this time will not have been mirrored, and therefore are lost.

9.1.6 HyperSwap

HyperSwap extends the concept and technology of Parallel Sysplex to the level of the storage system. HyperSwap provides the capability of masking primary storage system failures by switching transparently to a secondary, synchronously mirrored, set of disk volumes. It allows you to perform disk maintenance and planned site maintenance without requiring applications or services to be stopped.

9.2 Initial configuration

Figure 9-1 illustrates the basic layout of a GDPS managed 3-site Metro + z/OS Global Mirror (GDPS/MzGM) configuration. The host system accesses the primary volume set A, which has a synchronous Metro Mirror replication relationship to the B volume set. The System Data Movers and C volume set form the asynchronous z/OS Global Mirror replication relationship with the A volume set. The A volume set is the source for both the synchronous and asynchronous replication relationships. The term volume set is used as a generalization. A volume set can consist of a subset of, or all of the volumes in, a single storage system, or even be spread across multiple storage systems.

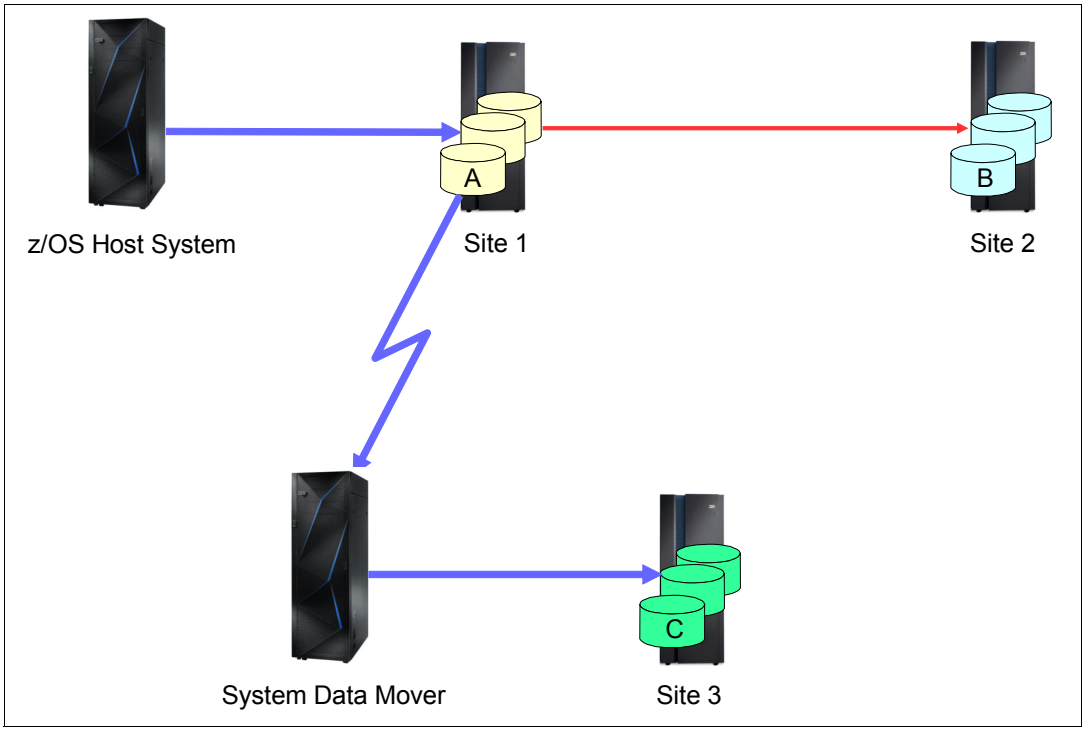

Figure 9-1 Basic 3-Site Metro + z/OS Global Mirror Environment

9.3 Migration process

The following sections discuss the typical processes, assumptions, and general considerations for migrating data in an existing GDPS/MzGM environment.

It is a preferred practice to replace volume sets in an MzGM environment as follows:

- Starting with the C volume set first. This minimizes the potential for performance issues that might result by installing a storage system with higher throughput capacity for the A or B volume sets that might overwhelm the storage system containing the C volume set.
- Replace all three systems simultaneously, eliminating the possibility of the performance issues.

The scenarios included in this chapter follow that practice by starting the migration with the C volume set, then moving to the B and A volume set. Using the HyperSwap function eliminates the disruption and risk of replacing a current MzGM primary system. This allows host

processing to continue at the primary site, and minimizes interruptions to the Disaster Recovery environment.

The following assumptions are made regarding the migration scenarios in this chapter:

- The storage systems do not support Multiple-Target PPRC
- Cascaded Global Copy is used to complete the initial migration of the data.
- HyperSwap is used to maintain host processing wherever possible.
- The availability of the MzGM environment is maintained whenever possible.
- The migrated data remains in the MzGM environment.
- The volume sets and scope of the migration project are clearly defined and documented.
- The applications affected by the migration project are clearly defined and documented.
- The post-migration environment is clearly defined and documented.

The examples in this chapter reference the migration of an entire storage system, but the process is valid for multiple storage systems or even a subset of a single storage system. If there are plans to migrate a subset of the data on one or multiple storage systems, the master migration plan must clearly define the applications and volume sets that will be affected.

To minimize the impact of the migration process, take these precautions:

- Use Global Copy to create an initial copy of data for each volume set in the migration environment without interruption to the host processing.
- Validate that there are no Out-of-Sync tracks before moving forward with the migration. This process ensures that each volume set in the environment is logically identical before you can safely proceed.
- Use the **NOCOPY** option when required to maintain the logically identical state. This option creates the Global Copy or Metro Mirror relationship without transferring any data from the primary volume set to the secondary volume set when the pairs are established.

Note: The examples included in this section are intended to be used as the template for creating a migration process. They are not to be used ‘as is’.

9.3.1 General considerations

This topic is covered in detail in Chapter 1, “Introducing disk data migration” on page 1 and Chapter 2, “Copy Services based data migrations” on page 19.

Important: It is essential to the success of any migration project that the information discussed in the two chapters referenced above is thoroughly reviewed before proceeding.

9.4 Migration Scenarios

This section describes the various migration scenarios for a three site GDPS/MzGM environment. The migration activities are performed using a combination of GDPS processes, TSO commands, and DS8000 DSCLI commands.

9.4.1 Replacing the zGM secondary

The process to replace the C volume set in an existing MzGM environment can be completed without interruption to the host processing at the primary site. The MzGM environment remains intact during the initial copy process. Your specific migration might require extra steps, depending on the environment.

The migration process is broken down into five major steps:

1. Initial setup of the migration environment
2. Data validation testing
3. Preparing to migrate
4. Cut over to the new volume set
5. Recovering the pre-migration environment

The steps assume that the new volume set is installed and configured as required. The cascade option for Global Copy is used to complete the initial data migration to the new volume set. The migration activities are performed using a combination of GDPS processes, TSO commands, and DS8000 DSCLI commands.

Initial setup of the migration environment

Note: This process is described below in its entirety. It does not refer to the generic DSCLI-based process in the 2.3, “Data replication” on page 22 section.

Note that you will not be using the cascade option for the Global Copy because the C volumes do not currently have a Metro Mirror or Global Copy relationship.

Before setting up the cascaded Global Copy environment, the GDPS and System Data Mover environments must be prepared for the migration. To do this, complete these steps:

- Build the new I/O configurations to use the R volume set instead of the B volume set as the zGM secondary
- Build the new GEOXPARM configuration to specify the R volume set instead of the B volume set as the zGM secondary
- Verify that all required software maintenance is installed

For this migration scenario, the cascaded Global Copy is extended to an additional volume set. The migration path is from C → R. This process requires two steps:

1. Establish the replication paths between the C volume set and the R volume set.
2. Create the Global Copy pairs between the C volume set and the R volume set.

Note: Consult an IBM representative during the planning process of the migration project to validate the replication network requirements. This helps avoid inadvertently affecting the production and replication environments during the data migration.

Figure 9-2 illustrates the basic layout of the migration configuration. In addition to the existing MzGM relationship between A, B, and C, an asynchronous Global Copy leg from the current C volume set to the new R volume set has been added. This relationship allows the data from the C volume set to be copied to the R volume set without impact to the production applications or the existing MzGM environment.

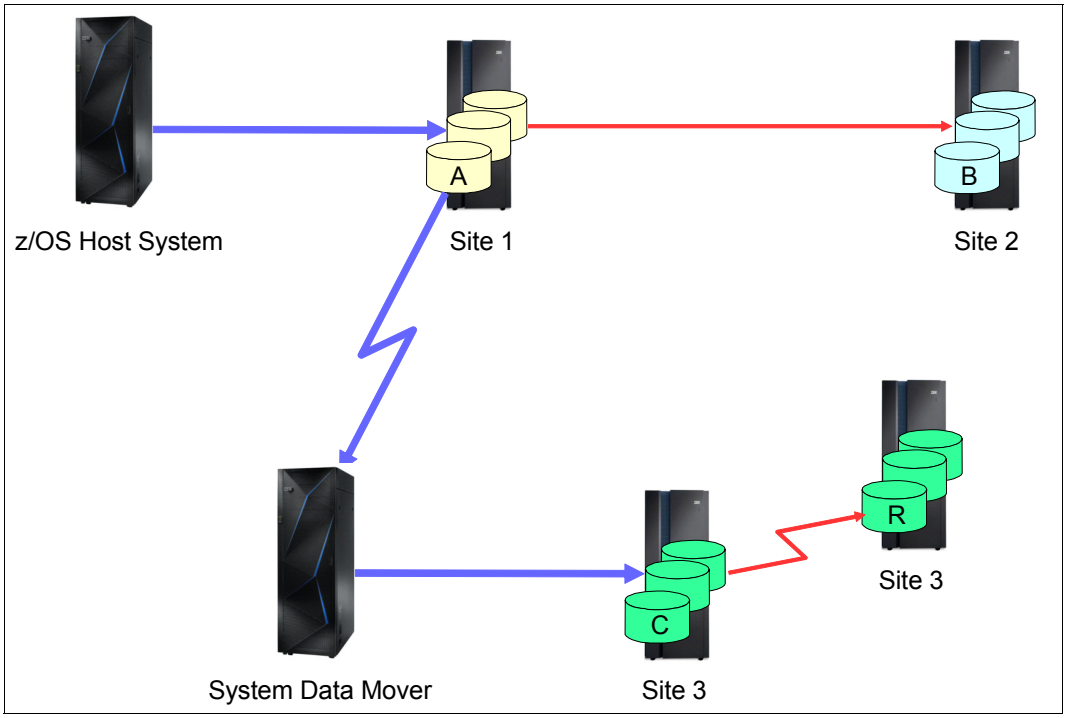

Figure 9-2 Global Copy with existing 3-Site Metro + z/Global Mirror Environment

Figure 9-3 illustrates the post-migration environment with R as the new zGM secondary device in the MzGM environment.

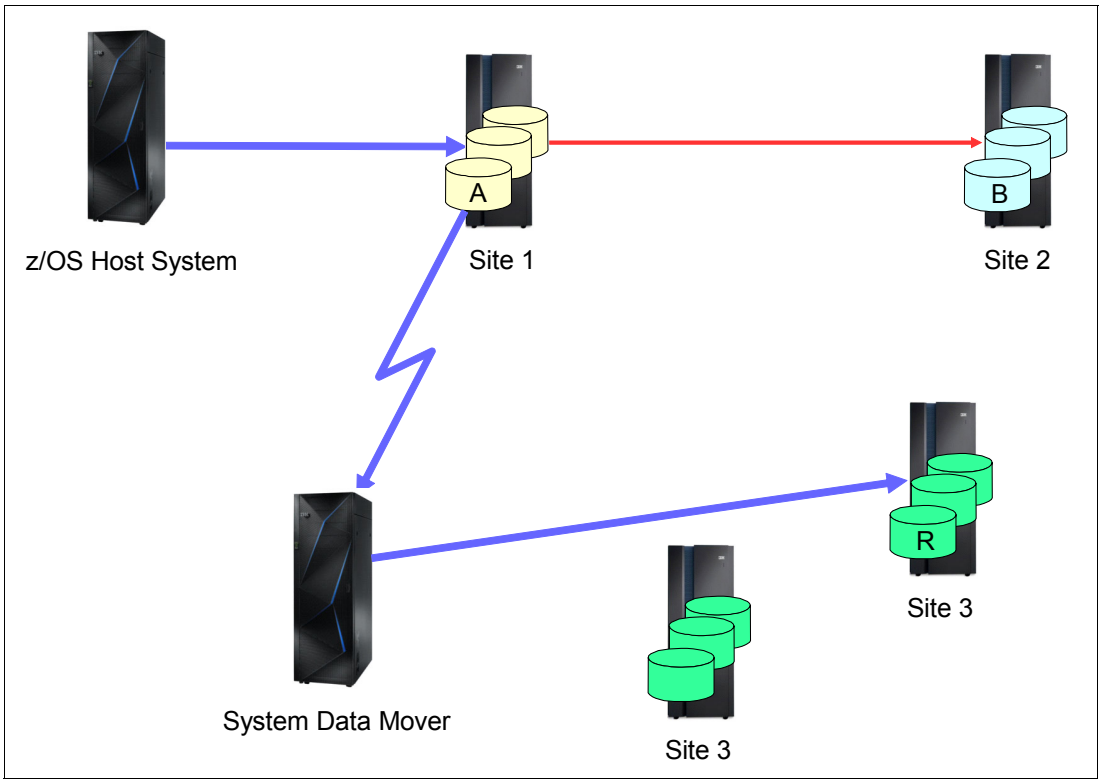

Figure 9-3 Post migration 3-Site Metro + z/Global Mirror Environment

Establish the replication paths

This step can be completed using either one of commands shown in Example 9-1 and Example 9-2.

Example 9-1 shows the TSO command used to complete this step. This command is run on the GDPS/XRC Master K-Sys TSO.

Example 9-1 TSO - Establish the replication paths between the B and L volume sets

```
CESTPATH DEVN(X'2000') +
PRIM(X'8300' CWWNN X'10') +
SEC(X'8400' RWWNN X'10') +
LINK(LINK(X'ccccrrrr') CGROUP(NO)
```

Example 9-2 shows the DSCLI command used to complete this step. This command is run on the C volume set.

Example 9-2 DSCLI - Establish the replication paths between the B and L volume sets

```
mkpprcpath -dev IBM.2107-75CCCCC -remotedev IBM.2107-75RRRRR -remotewwnn RWWNN
-srclss 10 -srcss 8300 -tgtlss 10 -tgtss 8400 Icccc:Irrrr
```

Create the Global Copy from B to L

This step can be completed using either one of commands in Example 9-3 on page 193 or Example 9-4 on page 193.

Example 9-3 shows the TSO command used to complete this step. This command is run on the GDPS/XRC Master K-Sys TSO.

Example 9-3 TSO - Create the Global Copy between the C and R volume sets

```
CESTPAIR DEVN(X'2000') +
PRIM(X'8300 CCCCC X’00' X’10') +
SEC(X'8400' RRRRR X’00' X’10') +
OPTION(XD) MODE(COPY) CASCADE(NO)
```

Example 9-4 shows the DSCLI command used to complete this step. This command is run at the C volume set.

Example 9-4 DSCLI - Create the Global Copy between the C and R volume sets

```
mkpprc -dev IBM.2107-75CCCCC -remotedev IBM.2107-75RRRRR -type gcp 1000:1000
```

Wait for initial copy to complete.

Data Validation Testing

You can now perform an initial data validation test on the R volume set before proceeding with the final cut over. This requires an available test environment with access to the R volume set.

This process is described in detail in 2.4, “Data validation testing” on page 24. It can be accomplished with minimum impact to the current environment, allowing a quick fall back to the pre-migration configuration should it be necessary.

Prepare to migrate to the new volume set

This process can be completed without interruption to host processing. The MzGM environment is temporarily affected. By using the cascaded Global Copy to migrate the data before the transition, the time required to fully recover the environment is minimized.

To complete the integration of the R volume set into the environment, complete these steps:

1. Use the XSUSPEND VOLUME(ALL) command to suspend all volume pairs from sessions that are coupled to the master session.
2. Go to sync C to R or wait for OOS to drain C to R.
3. Fail over at R to enable host access.
4. Failback at R in case you need to fall back to C.
5. Install and validate all physical cabling required, including the host links from the System Data Movers.
6. Load new I/O Gen at SDM to pick up new secondaries.
7. Load new GEOXPARM.
8. Verify or set ANTQ8253I EnableREFRESHS YES in ANTXIN00 PARMLIB.
9. Resume XRC with XSET REFRESHS.

9.4.2 Replacing the current Metro Mirror secondary

This section describes the process to replace the B volume set in an existing MzGM environment. The process can be completed without interruption to the host processing at the primary site. The MzGM environment remains intact during the initial copy process. Your specific migration might require extra steps, depending on the environment.

To minimize the risk and impact to the production environment, use the HyperSwap function to replace only Metro Mirror secondary volume sets. Do not replace an MzGM primary volume set.

The migration process is broken down into five major steps:

1. Initial setup of the migration environment
2. Data validation testing
3. Preparing to migrate
4. Cut over to the new volume set
5. Recovering the pre-migration environment

The following processes assume that the new volume set is installed and configured as required. The cascade option for Global Copy is used to complete the initial data migration to the new volume set.

The migration activities are performed using a combination of GDPS processes, TSO commands, and DS8000 DSCLI commands. Each step is described below in detail and includes the commands that are required, the status of the copy relationships after the step is completed, and any notes specific to that action.

Initial setup of the migration environment

Note: This process is described below in its entirety. This section does not refer to the generic DSCLI-based process in the 2.3, “Data replication” on page 22.

Before setting up the cascaded Global Copy environment, the GDPS and z/OS environments must be prepared for the migration. To do this, complete these steps:

- Build the new I/O configurations to use the L volume set instead of the B volume set as the Metro Mirror/HyperSwap secondary.
- Build the new GEOPARM configuration to specify the L volume set instead of the B volume set as the Metro Mirror/HyperSwap secondary.
- Verify that all required software maintenance is installed.

The following process sets up the initial cascaded replication environment.

For this migration scenario, the cascaded Global Copy is extended to an additional volume set. The migration path is from A → B → L. This process requires three steps:

1. Verify the status of the replication environment on the GDPS/PPRC panels at the GDPS/PPRC Master K-Sys NetView.
2. Establish the replication paths between the B volume set and the L volume set.
3. Create the Global Copy pairs between the B volume set and the L volume set.

Note: Consult an IBM representative during the planning process of the migration project to validate the replication network requirements. This helps to avoid inadvertently affecting the production and replication environments during the data migration.

Verify the mirror status

The GDPS/PPRC panels at the GDPS/PPRC Master K-Sys NetView should reflect the copy and volume status as shown in Example 9-5.

Example 9-5 Pre-migration mirror status

```
Copy Status
        A to B  Metro Mirror    Duplex
        B to L  n/a             n/a
DASD Status
        A       PPRC Pri        DUPLEX          Metro Mirror to B
        B       PPRC Sec        DUPLEX          Metro Mirror from A
        L       Simplex         n/a
```

Establish the replication paths

This step can be completed using either one of the following commands.

Example 9-6 shows the TSO command used to complete this step. This command is run on the GDPS/PPRC Master K-Sys TSO.

Example 9-6 TSO - Establish the replication paths between the B and L volume sets

```
CESTPATH DEVN(X'2000') +
PRIM(X'8200' BWWNN X'10') +
SEC(X'8300' LWWNN X'10') +
LINK(LINK(X'bbbbllll') CGROUP(NO)
```

Example 9-7 shows the DSCLI command used to complete this step. This command is run at the B volume set.

Example 9-7 DSCLI - Establish the replication paths between the B and L volume sets

```
mkpprcpath -dev IBM.2107-75BBBBB -remotedev IBM.2107-75LLLLL -remotewwnn LWWNN
-srclss 10 -srcss 8200 -tgtlss 10 -tgtss 8300 Ibbbb:Illll
```

After this is complete, the GDPS/PPRC Master K-Sys NetView should reflect the status that is shown in Example 9-8.

Example 9-8 Note that these activities did not affect the state of the copy relationships

```
Copy Status
        A to B  Metro Mirror    Duplex
        B to L  n/a             n/a
DASD Status
        A       PPRC Pri        DUPLEX          Metro Mirror to B
        B       PPRC Sec        DUPLEX          Metro Mirror from A
        L       Simplex         n/a
```

Note: This step is performed in advance of the cut-over. However, the physical links between B and L must be installed.

Create the Global Copy from B to L

This step can be completed using either one of commands shown in Example 9-9 on page 196 and Example 9-10 on page 196.

Example 9-9 shows the TSO command used to complete this step. This command is run on the GDPS/PPRC Master K-Sys TSO.

Example 9-9 TSO - Create the cascaded Global Copy relationship between the B and L volume sets

```
CESTPAIR DEVN(X'2000') +
PRIM(X'8200 BBBBB X’00' X’10') +
SEC(X'8300' LLLLL X’00' X’10') +
OPTION(XD) MODE(COPY) CASCADE(YES)
```

Example 9-10 shows the DSCLI command used to complete this step. This command is run at the B volume set.

Example 9-10 DSCLI - Create the cascaded Global Copy relationship between the B and L volume sets

```
mkpprc -dev IBM.2107-75BBBBB -remotedev IBM.2107-75LLLLL -type gcp -cascade
1000:1000
```

After this is completed, the GDPS/PPRC Master K-Sys NetView should reflect the status shown in Example 9-11.

Example 9-11 The copy and volume set status after the cascaded Global Copy is established to L

```
Copy Status
        A to B  Metro Mirror    Duplex
        B to L  Global Copy     Pending.XD
DASD Status
        A       PPRC Pri        DUPLEX          Metro Mirror to B
        B       PPRC Sec        DUPLEX.CAS      Metro Mirror from A
                PPRC Pri        PNDING.X.L      Global Copy to L
        L       PPRC Sec        PENDING.XD      Global Copy from B
```

Note: This step is performed in advance of the cut-over so that most of the data is already replicated to the L volume set. For a large configuration, it is MUCH quicker to use CLI commands rather than TSO commands, especially when you specify a range of source and target volumes, for example, 1000-1042:1000-1042, instead of separate commands for each mirror pair.

Figure 9-4 on page 197 illustrates the basic layout of the migration configuration. In addition to the existing MzGM relationship between A, B, and C, an asynchronous Global Copy leg from the current B volume set to the new L volume set has been added. This relationship allows the data from the B volume set to be copied to the L volume set without impact to the production applications or the existing MzGM environment. The replication links between the A and L volume sets for use during the cut over process have also been defined.

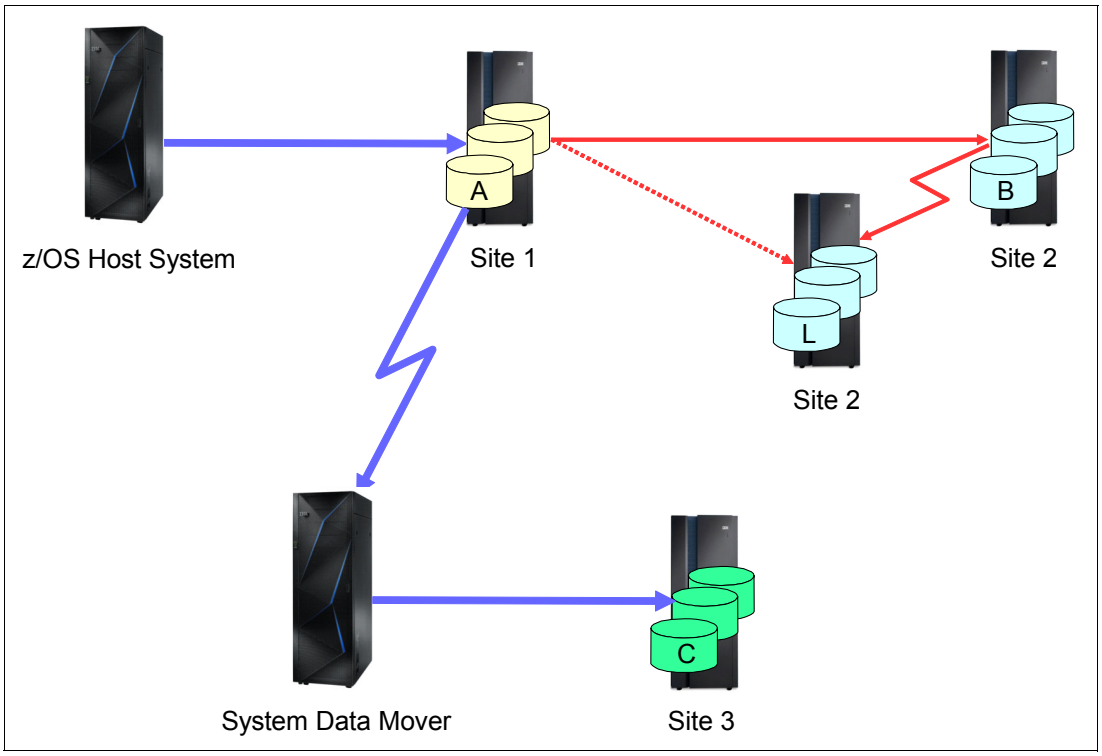

Figure 9-4 Cascaded Global Copy with existing 3-Site Metro + z/Global Mirror Environment

Figure 9-5 illustrates the post-migration environment with L as the new Metro Mirror secondary device in the MzGM environment.

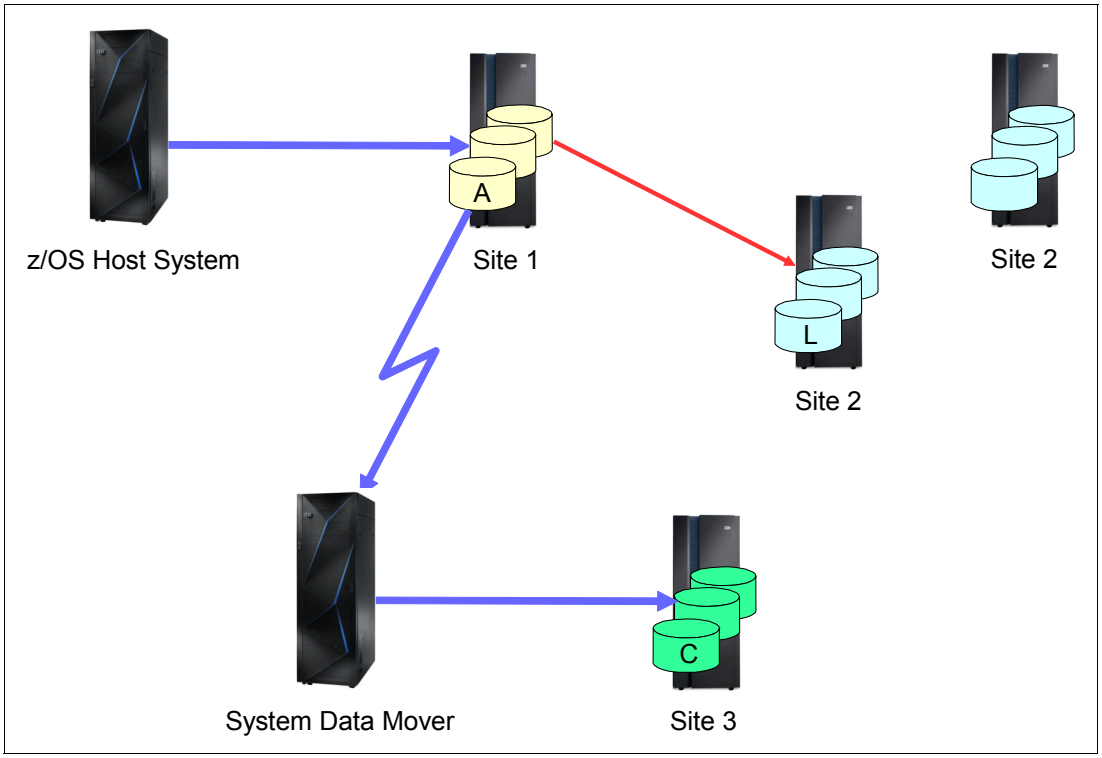

Figure 9-5 Post migration 3-Site Metro + z/Global Mirror Environment

Data Validation Testing

At this point, perform an initial data validation test on the L volume set before proceeding with the final cut over.

This process is described in 2.4, “Data validation testing” on page 24. It can be accomplished with minimum impact to the current environment, allowing a quick fall back to the pre-migration configuration should it be necessary.

If the validation testing is successful, proceed with the migration process. If the data validation process is not successful, determine why the validation testing failed before proceeding.

Prepare to migrate to the new volume set

Before moving forward with the cut over phase of the migration, the following two steps will need to be completed:

1. Install and validate any remaining physical cabling required, including the host links from the primary production and System Data Mover systems.
2. Verify the status of the Global Copy relationship between B and L using the GEOXAPMO command at the GDPS/PPRC Master K-Sys NetView.

Note: Use the Add Pair Monitor tool from the PPRC Preserve Mirror Tool to monitor the status of the Global Copy mirrors and wait until nearly all of the data have been copied.

The GEOXAPMO command at the GDPS/PPRC Master K-Sys NetView should reflect the status shown in Example 9-12. If not, any issues must be resolved before proceeding.

Example 9-12 Verify the copy and volume set status before beginning the cut over phase.

```
Copy Status
        A to B  Metro Mirror    Duplex
        B to L  Global Copy     Pending.XD
DASD Status
        A       PPRC Pri        DUPLEX          Metro Mirror to B
        B       PPRC Sec        DUPLEX.CAS      Metro Mirror from A
                PPRC Pri        PNDING.X.L      Global Copy to L
        L       PPRC Sec        PENDING.XD      Global Copy from B
```

Cut over to the new volume set

After the physical connectivity has been completed and the mirror status verified, the L volume set can be integrated into the MzGM environment as the new Metro Mirror secondary. This process can be completed without interruption to the host processing. The MzGM environment is temporarily affected. By using the cascaded Global Copy to migrate the data before the transition, the time that is required to fully recover the mirrored environment is minimized. This process requires 15 steps:

1. Remove XRC Incremental Resync.
2. Establish the Incremental Resync Bitmaps on the A volume set.
3. Freeze the Metro Mirror relationship between the A and B volume set.
4. Verify that all data has been copied to the L volume set.
5. Suspend the B to L Global Copy relationship.
6. Fail over the B to the L Metro Mirror relationship at the L volume set.

7. Establish the replication paths from the A to L volume set.
8. Create the Metro Mirror relationship from the A to L volume sets by using Incremental Resync.
9. Remove the A to B Metro Mirror relationship at the B volume set.
10. Establish the replication paths from the L to B volume set.
11. Fail back the B to L Metro Mirror relationship at the L volume set.
12. Load the new I/O Configuration.
13. Load the new GDPS/PPRC Configuration.
14. Restore XRC Incremental Resync.
15. Remove the L to B Global Copy relationship.

Remove XRC Incremental Resync

Example 9-13 shows the actions that are required to temporarily remove the XRC Incremental Resync capability. These commands are run at the GDPS/PPRC and GDPS/XRC Master K-Sys NetViews.

Example 9-13 Steps to remove XRC Incremental Resync

```
Set THREESITE='NO' in the GEOPLEX OPTIONS for GDPS/XRC
Suspend XRC
Load new SA policy for GDPS/XRC
Resync XRC
Set THREESITE='CASCADE' in the GEOPLEX OPTIONS for GDPS/PPRC
Suspend PPRC between A and B disk
Load new SA policy for GDPS/PPRC
Resync PPRC
```

After this is completed, the GDPS/PPRC Master K-Sys NetView should reflect the status shown in Example 9-14.

Example 9-14 Note that these activities did not affect the state of the copy relationships

```
Copy Status
        A to B  Metro Mirror    Duplex
        B to L  Global Copy     Pending.XD
DASD Status
        A       PPRC Pri        DUPLEX          Metro Mirror to B
        B       PPRC Sec        DUPLEX.CAS      Metro Mirror from A
                PPRC Pri        PNDING.X.L      Global Copy to L
        L       PPRC Sec        PENDING.XD      Global Copy from B
```

Note: This step is only required if the A and B disks are the source for an XRC mirror with Incremental Resynchronization enabled, that is, both the A and B disk have an XRC session.

You can use F ANTAS000,LISTSESS 2000 console command to check XRC sessions are removed on the PPRC secondary disk.

Establish the Incremental Resync Bitmaps on the A volume set

This step can be completed using either one of commands below.

Example 9-15 shows the TSO command used to complete this step. This command is run on the GDPS/PPRC Master K-Sys TSO.

Example 9-15 TSO - Establish the IR bitmaps on A

```
CESTPAIR DEVN(X'1000') +
PRIM(X'8100' AAAAA X'00' X'10') +
SEC(X'8200' BBBBB X'00' X'10') +
OPTION(SYNC) MODE(NOCOPY) CASCADE(NO) +
INCRESYNC (NOCOPY)
```

Example 9-16 shows the DSCLI command that is used to complete this step. This command is run at the A volume set.

Example 9-16 DSCLI - Establish the IR bitmaps at A

```
mkpprc -dev IBM.2107-75AAAAA -remotedev IBM.2107-75BBBBB -type mmir -mode nocp
-incrementalresync enablenoinit 1000:1000
```

After this is completed, the GDPS/PPRC Master K-Sys NetView should reflect the status shown in Example 9-17.

Example 9-17 This activity did not affect the state of the copy relationships

```
Copy Status
        A to B  Metro Mirror    Duplex
        B to B' Global Copy     Pending.XD
DASD Status
        A       PPRC Pri        DUPLEX          Metro Mirror to B
        B       PPRC Sec        DUPLEX.CAS      Metro Mirror from A
                PPRC Pri        PNDING.X.C      Global Copy to B'
        B'      PPRC Sec        PENDING.XD      Global Copy from B
```

Note: This command is issued at the start of the cut-over as all tracks changed on A from this point are resynchronized to the L volume set later in the process. It can be applied to an existing and active Metro Mirror relationship. There is no need to suspend the Metro Mirror relationship between A and B. For a large configuration, it is MUCH quicker to use CLI commands rather than TSO commands, especially when you specify a range of source and target volumes, for example, 1000-1042:1000-1042 instead of separate commands for each mirror pair. After this step, a CQUERY of disk A will show INCRES(YES).

Freeze the Metro Mirror relationship between the A and B volume sets

Example 9-18 shows the TSO command used to complete this step. This command is run on the GDPS/PPRC Master K-Sys NetView by using the script.

Example 9-18 Freeze the A to B Metro Mirror

```
DASD STOP SECONDARY
```

After this is completed, the GDPS/PPRC Master K-Sys NetView should reflect the status shown in Example 9-19.

Example 9-19 The copy and volume set status after the Freeze is issued

```
Copy Status
        A to B  Metro Mirror    Frozen
```

```
        B to L  Global Copy     Pending.XD
DASD Status
        A       PPRC Pri        SUSPEND(A)      Metro Mirror to B
        B       PPRC Sec        DUPLEX.CAS      Metro Mirror from A
                PPRC Pri        PNDING.X.L      Global Copy to L
        L       PPRC Sec        PENDING.XD      Global Copy from B
```

Note: From this point, there is no High Availability capability. You would lose systems if the A disk is lost after this suspend step until you establish Metro Mirror from A to C with incremental resync is complete.

Verify that all data has been copied to the L volume set

Example 9-20 shows the TSO commands that are used to complete this step. This command is run on the GDPS/PPRC Master K-Sys NetView.

Example 9-20 Verify that all data have been copied to the L volume set

```
GEOXAPMO
```

After this is completed, the GDPS/PPRC Master K-Sys NetView should reflect the status that is shown in Example 9-21.

Example 9-21 The copy and volume set status as any remaining updates are drained

```
Copy Status
        A to B  Metro Mirror    Frozen
        B to L  Global Copy     Pending.XD
DASD Status
        A       PPRC Pri        SUSPEND(A)      Metro Mirror to B
        B       PPRC Sec        DUPLEX.CAS      Metro Mirror from A
                PPRC Pri        PNDING.X.L      Global Copy to L
        L       PPRC Sec        PENDING.XD      Global Copy from B
```

Note: Use the Add Pair Monitor tool from the PPRC Preserve Mirror Tool to monitor the status of the PPRC XD mirrors, and wait until 100% of the data has been copied. Most of the data has already been copied, so this so should only take a few seconds.

Suspend the B to L Global Copy relationship

This step can be completed using either one of commands that are shown in Example 9-22 and Example 9-23 on page 202.

Example 9-22 shows the TSO command used to complete this step. This command is run on the GDPS/PPRC Master K-Sys TSO.

Example 9-22 TSO - Suspend the B to L Global Copy relationship

```
CGROUP DEVN(X'2000') +
PRIM(X'8200' BBBBB X'10') +
SEC(X'8300' LLLLL X'10') +
FREEZE
```

Example 9-23 shows the DSCLI command used to complete this step. This command is run at the B volume set.

Example 9-23 DSCLI - Suspend the B to L Global Copy relationship

```
freezepprc -dev IBM.2107-75BBBBB -remotedev IBM.2107-75LLLLL 10:10
```

After this is completed, the GDPS/PPRC Master K-Sys NetView should reflect the status shown in Example 9-24.

Example 9-24 The copy and volume set status after the B to L relationship is converted to Metro Mirror

```
Copy Status
        A to B  Metro Mirror    Frozen
        B to L  Global Copy     Frozen
DASD Status
        A       PPRC Pri        SUSPEND(A)      Metro Mirror to B
        B       PPRC Sec        DUPLEX.CAS      Metro Mirror from A
                PPRC Pri        SUSP(A).X.L     Global Copy to L
        L       PPRC Sec        PENDING.XD      Metro Mirror from B
```

Note: Suspend the B disks in case you need to failback later as part of a back out.

Failover the B to L relationship at the L volume set

Example 9-25 shows the command used to complete this step. This command is run on the L volume set using the DS8000 DSCLI.

Example 9-25 The DSCLI Failover command

```
failoverpprc -dev IBM.2107-75LLLLL -remotedev IBM.2107-75BBBBB -type gcp -cascade
1000:1000
```

After this is completed, the GDPS/PPRC Master K-Sys NetView should reflect the status shown in Example 9-26.

Example 9-26 The copy and volume set status after the B to L relationship is failed over

```
Copy Status
        A to B  Metro Mirror    Frozen
        B to L  Global Copy     Frozen
DASD Status
        A       PPRC Pri        SUSPEND(A)      Metro Mirror to B
        B       PPRC Sec        DUPLEX.CAS      Metro Mirror from A
                PPRC Pri        SUSP(A).X.L     Global Copy to L
        L       Simplex         SMPLEX.CAS      Inactive
                PPRC Pri        SUSP(3).CAS     Failed Over(B)
```

Note: You do this to remove the L disks secondary status, so that you can later establish Metro Mirror from A to L. At this stage, you cannot address the L disk from z/OS as they have the same addresses as the B disk until the I/O configuration change later in the process. Hence the use of the CLI and no TSO equivalent. For a large configuration, it is MUCH quicker to specify a range of source and target volumes, for example, 1000-1042:1000-1042, instead of separate commands for each mirror pair.

Establish the replication paths from the A to L volume set

This step can be completed using either one of commands below.

Example 9-27 shows the TSO command used to complete this step. This command is run on the GDPS/PPRC Master K-Sys TSO.

Example 9-27 TSO - Establish the replication paths from the A to L volume sets

```
CESTPATH DEVN(X'1000') +
PRIM(X'8100 AAAAA X'10') +
SEC(X'8300' LLLLL X'10') +
LINK(LINK(X'aaaallll') CGROUP(NO)
```

Example 9-28 shows the DSCLI command used to complete this step. This command is run on the A volume set.

Example 9-28 DSCLI - Establish the replication paths from the A to L volume sets

```
mkpprcpath -dev IBM.2107-75AAAAA -remotedev IBM.2107-75LLLLL -remotewwnn CWWNN
-srclss 10 -srcss 8100 -tgtlss 10 -tgtss 8300 Ibbbb:Illll
```

After this process is complete, the GDPS/PPRC Master K-Sys NetView should reflect the status shown in Example 9-29.

Example 9-29 This activity does not affect the state of the copy relationships

```
Copy Status
        A to B  Metro Mirror    Frozen
        B to L  Global Copy     Frozen
DASD Status
        A       PPRC Pri        SUSPEND(A)      Metro Mirror to B
        B       PPRC Sec        DUPLEX.CAS      Metro Mirror from A
                PPRC Pri        SUSP(A).X.L     Global Copy to L
        L       Simplex         SMPLEX.CAS      Inactive
                PPRC Pri        SUSP(3).CAS     Failed Over(B)
```

Note: The replication links between A and L must be installed before running this step

Create the Metro Mirror relationship from the A to L volume sets with IR

This step can be completed using either one of commands shown in Example 9-30 and Example 9-31 on page 204.

Example 9-30 shows the TSO command used to complete this step. This command is run on the GDPS/PPRC Master K-Sys TSO.

Example 9-30 TSO - Create the A to L Metro Mirror with Incremental Resync

```
CESTPAIR DEVN(X'1000') +
PRIM(X'8100' AAAAA X’00' X'10') +
SEC (X'8300' LLLLL X'00' X'10') +
OPTION(SYNC) MODE(INCRES) CASCADE(NO) +
FORCE (YES)
```

Example 9-31 shows the DSCLI command used to complete this step. This command is run at the A volume set.

Example 9-31 DSCLI - Create the A to L Metro Mirror with Incremental Resync

```
mkpprc -dev IBM.2107-75AAAAA -remotedev IBM.2107-75LLLLL -type mmir -mode nocp
-incrementalresync override 1000:1000
```

After this is completed, the GDPS/PPRC Master K-Sys NetView should reflect the status shown in Example 9-32.

Example 9-32 The copy and volume set status after the A to L mirror is created

```
Copy Status
        A to L  Metro Mirror    Duplex
        B to L  Global Copy     Frozen
DASD Status
        A       PPRC Pri        DUPLEX          Metro Mirror to L
        B       PPRC Sec        DUPLEX.CAS      Metro Mirror from A
                PPRC Pri        SUSP(A).X.L     Global Copy to L
        L       PPRC Sec        DUPLEX.CAS      Metro Mirror from A
                PPRC Pri        SUSP(3).CAS     Failed Over (B)
```

Note: This is not a full copy, but replicates only the changes made on A since “Establish the Incremental Resync Bitmaps on the A volume set” on page 199. After the mirrors come to duplex, the DR position is restored, but HyperSwap protection is still suspended. The z/OS systems do not have access to the secondary DS8870 devices on the L disk yet. Do not wait for duplex before continuing with the next step.

Remove the A to B Metro Mirror relationship at the B volume set

Example 9-33 shows the DSCLI command used to complete this step. This command is run on the B volume set.

Example 9-33 The DSCLI command to remove a Remote Mirror and Copy volume pair relationship

```
rmpprc -dev IBM.2107-75BBBBB -at tgt -unconditional -quiet 1000
```

After this is completed, the GDPS/PPRC Master K-Sys NetView should reflect the status shown in Example 9-34.

Example 9-34 The copy and volume set status after the A to B relationship is removed at B

```
Copy Status
        A to L  Metro Mirror    Duplex
        B to L  Global Copy     Frozen
DASD Status
        A       PPRC Pri        DUPLEX          Metro Mirror to L
        B       Simplex         SMPLEX.CAS
                PPRC Pri        SUSP(A).X.L     Global Copy to L
        L       PPRC Sec        DUPLEX.CAS      Metro Mirror from A
                PPRC Pri        SUSP(3).CAS     Failed Over (B)
```

Note: This is performed in case you need to back out later. You must remove the secondary status from the B disk to allow you to FAILBACK to them. This step can be performed in parallel with the I/O configuration changes in later steps (hence the use of CLI rather than TSO commands).

Establish replication paths from the L to B volume set

Example 9-35 shows the DSCLI command used to complete this step. This command is run on the L volume set.

Example 9-35 The DSCLI command to establish the replications paths from L to B

```
mkpprcpath -dev IBM.2107-75LLLLL -remotedev IBM.2107-75BBBBB -remotewwnn BWWNN
-srclss 10 -srcss 8300 -tgtlss 10 -tgtss 8200 Illll:Ibbbb
```

After this is completed, the GDPS/PPRC Master K-Sys NetView should reflect the status shown in Example 9-36.

Example 9-36 This activity does not affect the state of the copy relationships

```
Copy Status
        A to L  Metro Mirror    Duplex
        B to L  Global Copy     Frozen
DASD Status
        A       PPRC Pri        DUPLEX          Metro Mirror to L
        B       Simplex         SMPLEX.CAS
                PPRC Pri        SUSP(A).X.L     Global Copy to L
        L       PPRC Sec        DUPLEX.CAS      Metro Mirror from A
                PPRC Pri        SUSP(3).CAS     Failed Over (B)
```

Note: This is performed in case you need to back out later. This step can be performed in parallel with the I/O configuration changes in later steps (hence the use of CLI rather than TSO commands).

Failback the B to L relationship at the L volume set

Example 9-37 shows the DSCLI command used to complete this step. This command is run on the L volume set.

Example 9-37 The DSCLI Failback command

```
failbackpprc -dev IBM.2107-75LLLLL -remotedev IBM.2107-75BBBBB -type gcp -cascade
1000:1000
```

After this process is complete, the GDPS/PPRC Master K-Sys NetView should reflect the status shown in Example 9-38.

Example 9-38 The copy and volume set status after the B to L relationship is failed back

```
Copy Status
        A to L  Metro Mirror    Duplex
        L to B  Global Copy     Pending.XD
DASD Status
        A       PPRC Pri        DUPLEX          Metro Mirror to L
        B       Simplex         SMPLEX.CAS
```

```
                PPRC Sec        PNDING.X.L      Global Copy from L
        L       PPRC Sec        DUPLEX.CAS      Metro Mirror from A
                PPRC Pri        PNDING.X.L      Global Copy to B
```

Note: This is performed in case we need to back out later. This step can be performed in parallel with the I/O configuration changes in later steps (hence the use of CLI rather than TSO commands).

Load the new I/O Configuration

Example 9-39 shows the commands required to load the new I/O configuration files.

Example 9-39 Load the new IOCDS and IODF files

```
Load IOCDS on all processor
Load IODF on all SYSPLEX
```

After this process is complete, the GDPS/PPRC Master K-Sys NetView should reflect the status shown in Example 9-40.

Example 9-40 These activities do not affect the state of the copy relationships

```
Copy Status
        A to L  Metro Mirror    Duplex
        L to B  Global Copy     Pending.XD
DASD Status
        A       PPRC Pri        DUPLEX          Metro Mirror to L
        B       Simplex         SMPLEX.CAS
                PPRC Sec        PNDING.X.L      Global Copy from L
        L       PPRC Sec        DUPLEX.CAS      Metro Mirror from A
                PPRC Pri        PNDING.X.L      Global Copy to B
```

Note: This step can be performed in parallel with the previous two steps.

Load the new GDPS/PPRC Configuration

Example 9-41 shows the commands that are required to load the new GDPS configuration files. These commands are run at the GDPS/PPRC Master K-Sys NetView.

Example 9-41 Load the new GDPS/PPRC Configuration

```
Use GDPS Panels
Use GDPS Option C, D – DASD Configuration
```

After this process is complete, the GDPS/PPRC Master K-Sys NetView should reflect the status shown in Example 9-42.

Example 9-42 This activity did not affect the state of the copy relationships

```
Copy Status
        A to L  Metro Mirror    Duplex
        L to B  Global Copy     Pending.XD
DASD Status
        A       PPRC Pri        DUPLEX          Metro Mirror to L
        B       Simplex         SMPLEX.CAS
                PPRC Sec        PNDING.X.L      Global Copy from L
```

```
        L       PPRC Sec        DUPLEX.CAS      Metro Mirror from A
                PPRC Pri        PNDING.X.L      Global Copy to B
```

Note: Update the GEOPARM to specify the A and L disks before loading the new configuration. After this GEOPARM is loaded, HyperSwap is re-enabled and full DR protection restored after the mirrors are all duplex.

Restore XRC Incremental Resync

Example 9-43 shows the commands that are required to restore the XRC Incremental Resync capability. These commands are run at the GDPS/PPRC and GDPS/XRC Master K-Sys NetViews.

Example 9-43 Steps to restore XRC Incremental Resync

```
Restore original THREESITE option in the GDPS/PPRC GEOPLEX OPTIONS
Suspend PPRC between A and C disk
Load new SA policy for GDPS/PPRC
Resync PPRC
Restore original THREESITE option in the GDPS/XRC GEOPLEX OPTIONS
Suspend XRC
Load new SA policy for GDPS/XRC
Resync XRC
```

After this process is complete, the GDPS/PPRC Master K-Sys NetView should reflect the status shown in Example 9-44.

Example 9-44 This activity did not affect the state of the copy relationships

```
Copy Status
        A to B  Metro Mirror    Duplex
        B to L  Global Copy     Pending.XD
DASD Status
        A       PPRC Pri        DUPLEX          Metro Mirror to B
        B       PPRC Sec        DUPLEX.CAS      Metro Mirror from A
                PPRC Pri        PNDING.X.L      Global Copy to L
        L       PPRC Sec        PENDING.XD      Global Copy from B
```

Note: This step is only required if the A and B disks are the source for an XRC mirror with Incremental Resynchronization enabled, that is, both the A and B disk have an XRC session.

You can use F ANTAS000,LISTSESS 2000 console command to check XRC sessions are in place on the PPRC secondary disk.

Remove the L to B relationship at the L volume set

This step can be completed using either one of commands below.

Example 9-45 shows the TSO command used to complete this step. This command is run on the GDPS/PPRC Master K-Sys TSO.

Example 9-45 TSO - Remove the L to B Global Copy relationship

```
CDELPAIR DEVN(X'2000') +
PRIM(X'8300' LLLLL X'00' X'00') +
SEC(X'8200' BBBBB X'10' X'00')
```

Example 9-46 shows the DSCLI command used to complete this step. This command is run at the L volume set.

Example 9-46 DSCLI - Remove the L to B Global Copy relationship

```
rmpprc -dev IBM.2107-75LLLLL -remotedev IBM.2107-75BBBBB -quiet 1000:1000
```

After this process is complete, the GDPS/PPRC Master K-Sys NetView should reflect the status shown in Example 9-47.

Example 9-47 The copy and volume set status after the B to L relationship is failed back

```
Copy Status
        A to L  Metro Mirror    Duplex
        L to B  n/a             n/a
DASD Status
        A       PPRC Pri        DUPLEX          Metro Mirror to L
        B       Simplex
        L       PPRC Sec        DUPLEX          Metro Mirror from A
```

Note: After the decision is made to stay with L disk as PPRC secondary, remove its cascaded relationship to B.

This completes the migration scenario.

9.4.3 Replacing the current Metro Mirror primary

This section describes the process to migrate the Metro Mirror primary volume set to a new DS8000. As mentioned in the 9.3, “Migration process” on page 188, use HyperSwap to eliminate the disruption and risk of replacing a current MzGM primary system. Doing so allows host processing to continue at the primary site, and minimizes interruptions to the Disaster Recovery environment.

The process to complete this migration is as follows:

1. HyperSwap and resynchronize Metro Mirror to make the current primary volume set (A) the secondary volume set (B).
2. Perform a Metro Mirror secondary migration as described in “Replacing the current Metro Mirror secondary” on page 193.
3. HyperSwap and resynchronize Metro Mirror to move production to the new volume set (L).

Part 2

z/OS based Migration Scenarios

This part examines typical migration scenarios based on several initial copy services configurations.

This part contains the following chapters:

- Migration using TDMF
- Migration using zDMF

© Copyright IBM Corp. 2015. All rights reserved.

10

Using TDMF z/OS

This chapter describes how to migrate disk data using the TDMF z/OS product. It addresses the steps that are required to use the product in either a typical local shared storage or in a remote migration scenario.

This chapter contains the following sections:

- General information about TDMF z/OS
- Using TDMF z/OS in a local shared storage environment
- Using TDMF z/OS in a remote environment

© Copyright IBM Corp. 2015. All rights reserved.

10.1 General information about TDMF z/OS

Transparent Data Migration Facility (TDMF) for z/OS is a host-based software solution for data migration in a z/OS environment.

10.1.1 Provided Functions

TDMF basically provides two distinct functions:

- Data migration with I/O redirection

  This function provides nondisruptive data transport from source to target DASD volumes with subsequent redirecting of I/Os to the new target devices. This function is described in 10.2, “Using TDMF z/OS in a local shared storage environment” on page 213.

- Point-in-Time migration

  This function provides nondisruptive data transport from source to target DASD volumes *without* redirecting I/Os to the new target devices. One special use of this function is data replication to remote target devices, which is described in 10.3, “Using TDMF z/OS in a remote environment” on page 246.

TDMF is also especially useful when migrating from non-IBM source volumes to IBM target volumes.

10.1.2 Terminology

The following terms are used when describing data migration processes with TDMF z/OS:

- Master system: The z/OS image where the TDMF Master batch job or started task is running and that is responsible for the data copy function. There is only one Master system in a local TDMF session.
- Agent system: A z/OS image in a storage environment shared with the Master system running a TDMF Agent batch job or started task. To ensure data integrity, any z/OS image with access to any migration volume must either run the TDMF Master or a TDMF Agent.
- Communications data set (COMMDS): The Master and associated Agent systems communicate through a shared system communications data set (COMMDS). See “Communications data set” on page 234 for details about this important component of a TDMF Migration session.
- Source volume: A DASD volume containing data to be migrated.
- Target volume: A DASD volume receiving the migrated data.
- Migration/Volume pair: The source and it associated target volume.
- Migration group: A group of volume pairs with an associated group name.
- Migration session: A Master system plus any number of Agent systems and the Volume pairs and Migration groups to be processed in a single TDMF execution.
- Swap migration: A migration session in which the source and target volumes are logically swapped and I/O is redirected to the new source after the migration completes. The original source device is usually marked offline and put into a Soft Fence state (if supported by the storage subsystem). Swap migrations are started by the MIGRATE control statement.
- Synchronization: The period when the following steps occur:

  a. I/O is inhibited to the source devices (Quiesce phase).

b. The last of the updates are collected and reflected to the target devices (Synchronize phase).

c. The following steps only apply to Swap migrations:

d. The source and target volumes are logically swapped (Redirect phase).

e. The I/O is resumed to the new source (original target) device (Resume phase)

- Point in Time (PIT) migration: A migration session where source and target volumes are not swapped after the migration completes. PIT migrations are started by the REPLICATE control statement.
- Perpetual Point in Time (PPIT) migration: A migration session where source and target volumes are not swapped after the migration completes. Other than a normal Point-in-Time migration, multiple synchronization points can be taken by replying to the outstanding PPIT prompt. The migration session does not complete before requested by the user. PPIT migrations are started by the REPLICATE control statement and specification of the PPIT common option.

10.2 Using TDMF z/OS in a local shared storage environment

A local shared DASD environment is when the TDMF z/OS product is an installation where all source and target volumes are locally connected to any number of systems. This can be a SYSPLEX or other configuration in which volumes to be migrated are shared between z/OS systems.

10.2.1 Overview of TDMF z/OS

The following sections provide you with an overview of the functions of TDMF z/OS in a shared DASD environment:

- Data migration and TDMF definitions and characteristics
- Architecture and Components
- Migration process flow
- Hardware compatibility
- Installation
- Customer requirements

Data migration and TDMF definitions and characteristics

Working with TDMF z/OS includes the following considerations:

- Multiple volume migrations can be performed within a single session (batch job).
- Multiple migration sessions can be performed simultaneously
- TDMF is dynamically activated and terminated.
- Applications remain unaware that migration is underway. The data on the source device is continuously and fully accessible for read and write activity.
- After the synchronization of source and target volumes is complete, the takeover of the target device is nondisruptive.
- TDMF supports a multiple-system shared data environment.
- TDMF guarantees complete physical data integrity.
- The use of TDMF is not restricted to any control unit model type or device type. Except as noted, all devices that are online in the data center can participate in a migration session as required.

- All volumes (source and target) of a migration must be online to the TDMF Master system.
- No data set can be allocated to a target volume during migration.

> **Restriction:** TDMF z/OS does not allow access to the target volume during the migration process.

- A source volume cannot contain an active local page data set, swap data sets, or active Sysplex Couple data sets (XCF).
- Some source volumes require special consideration like JES spool data sets and TDMF communication data sets (COMMDS)
- The source and target volumes must be of the same track geometry.

Architecture and Components

Figure 10-1 depicts the TDMF z/OS architecture and components. Volumes can be migrated from any storage vendor (or within the same storage subsystem).

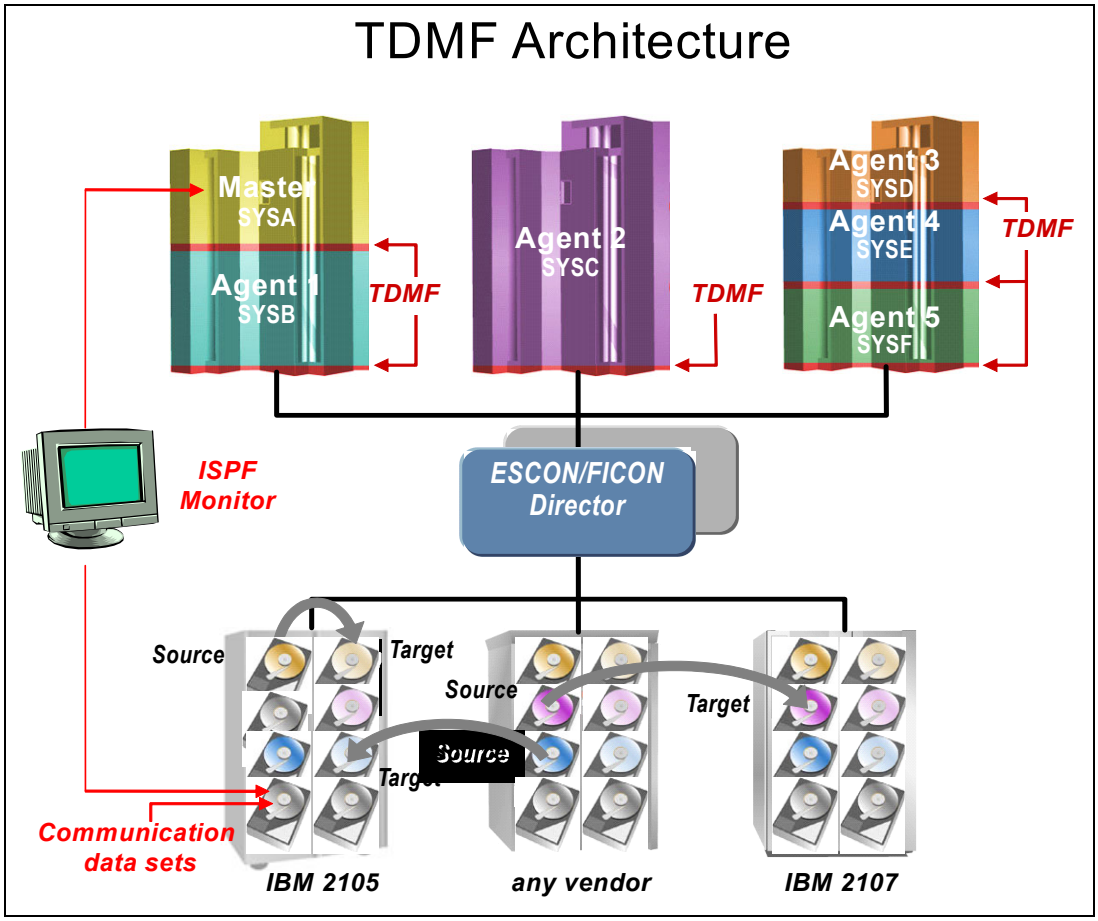

Figure 10-1 TDMF architecture (example with 6 z/OS images)

The TDMF z/OS architecture consists of the following components:

- TDMF Monitor: The TDMF Monitor is an ISPF application that is used to view and control an active or past migration session.
- TDMF Batch Monitor: TDMF sessions can also be controlled and monitored by a batch program called the TDMF Batch Monitor.

- Communications data set: The communications data set (COMMDS) is used to pass information between systems participating in a TDMF session. This data set contains the status and messages related to a specific session.
- Master and Agents: Each z/OS image that has access to the storage subsystems that contain the source and the target volumes must either run as a TDMF Master or a TDMF Agent. The Master and each Agent within a session must have access to the same COMMDS.

Master system responsibilities

The TDMF Master system is responsible for the following tasks:

- Initialize the TDMF Master system environment and the COMMDS
- Start and control each session for all participating systems
- Monitor source volume user I/O activity to detect updates
- Monitor target volume user I/O activity to prevent updates
- Copy data from the source volume to the target volume
- Process detected source volume updates from all systems
- Perform refresh operations to the target volume to reflect the update activity on the source volume
- Internal health checking

Agent system responsibilities

The TDMF Agent system is responsible for the following tasks:

- Initialize the TDMF Agent environment and establish communications to the Master system through the COMMDS
- Acknowledge and process migration requests (volumes to be moved) from the Master system
- Monitor source volume user I/O activity and detect updates
- Monitor target volume user I/O activity to prevent updates
- Notify the Master system of source volume update activity through the COMMDS
- Check the internal health of the Agent environments and the health of the Master system

Migration process flow

The TDMF migration process is a phased approach to make sure that data integrity is maintained in each phase. Data migration sessions can be interrupted at any time with no data loss. During the migration progress, all updates are written to the original source volume and are replicated by TDMF to the target volume. If an I/O violation occurs, the session is terminated and the source volumes remain intact.

System initialization process

All systems (Master and Agents) in a TDMF session must successfully start before the migration proceeds. If any violation occurs during system initialization on any system defined in the session, the migration is terminated.

If not all Agents specified by the Master control card input are started within 15 minutes after the Master was started, TDMF will terminate. If you start a TDMF Agent that was not specified in the TDMF Master JCL (but is using the same COMMDS) the TDMF migration session will be terminated.

If you use System Authorization Facility (SAF) and any migration volume fails SAF, the TDMF migration session is terminated.

SAF has these requirements:

- Swap migrations require ALTER authority for the source and target volumes.
- Point-in-Time migrations require READ authority for the source volume and UPDATE authority for the target volume.

Volumes in a session can be terminated by using the TDMF Monitor or TDMF Batch Monitor on the Master system before initialization of the Agent systems. If you select the History option to automatically record information about the migration session, the recording requires UPDATE authority for the data set. For more information, see the System Authorization Facility (SAF) and *TDMF Installation and Reference Guide*, TDM-Z55IR-001.

The TDMF Master system initiates and controls all migrations. The Agents must confirm completion of each phase to proceed. If any system detects a violation, that specific migration terminates. Depending on the state of the current migration, you must perform back-out processing to establish the original status before the migration session started.

Migration phases

The TDMF migration process flow can be separated into several phases as shown in Figure 10-2.

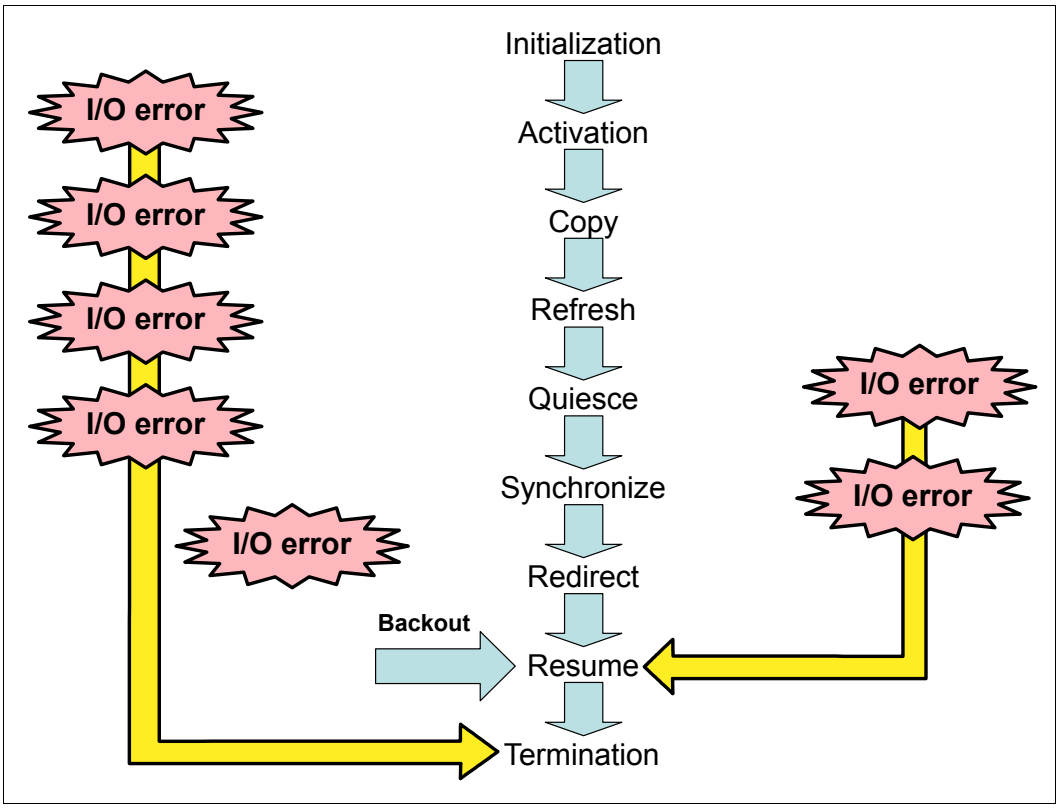

Figure 10-2 Phases of a TDMF session

- Initialization phase: During this phase, all systems participating in the migration session confirm the validity of the source and target volumes. The *Initialization phase* includes the following activities:

- Volume acknowledgment: Volumes or groups that require an acknowledgment are not eligible for volume confirmation and volume selection until an acknowledgment is received through the TDMF Monitor, a TDMF Batch Monitor, or the system console.

  This acknowledgment is required for migration volumes whose disaster recovery or mirror techniques are not compatible unless the **ALLOWmirrorchange(NOACK)** option is specified. TDMF recognizes volumes that use Remote Mirror/Copy functions like Metro Mirror (synchronous PPRC), Global Mirror (asynchronous PPRC, XRC), TrueCopy from Hitachi, and Symmetrix Remote Data Facility (SRDF) from EMC (these volumes are recognized with or without Consistency Groups).

  The disaster recovery or mirroring type must be compatible between the source and target volumes. Compatibility includes the vendor command interface structure and parameters. If they are not compatible, the **ALLOWmirrorchange** option must be specified and an acknowledgment received to continue the migration.

- Volume confirmation: Volumes that require confirmation are not eligible for volume or group selection until a confirmation is received. Eligibility can be confirmed by the TDMF Monitor, a TDMF Batch Monitor, or by using the system console. The order of confirmations determines the order of volume selection. Volumes or groups that do not require confirmation are immediately available for volume or group selection.

- Volume selection: By default all volume pairs that are defined in a session are automatically selected during the *Initialization phase*. However, volume selection is affected when certain user options are specified. The options that affect volume selection are volume confirmation (**CONFIRM|NOCONFIRM**), number of concurrent volumes and active in copy (**CONCurrent**), group options, and allow mirroring changes (**ALLOW|NOALLOW**).

- Volume initialization: Initialization of all volume level control blocks and page fixing of all real storage frames that are necessary for a volume migration.

- Activation phase: In this phase, I/O monitoring is activated by TDMF for the source and target devices. All systems in the session then attempt to allocate each source and target volume, preventing them from being inadvertently varied offline.

- Copy phase: In this phase, the Master system begins a COPY volume task to copy data from the source volume to the target volume. Each migration volume pair has an independent and asynchronous COPY volume task in the migration session.

  Updates (write I/Os) to source volumes are monitored and collected to be processed by the subsequent *Refresh phase*. Changes to the source volume are kept in a bitmap in the communications data set (COMMDS).

- Refresh phase: When the initial volume *Copy phase* completes, all source volume updates collected in the meantime are now copied to the target volume.

  Multiple *Refresh phases* can occur until TDMF determines that synchronization of the target volume can be achieved, at which time the Master system quiesces the source volume.

- Quiesce phase: In this phase, all TDMF Agents are requested to quiesce I/O to the source volume and return the final group of detected updates to the TDMF Master. The TDMF Master then also stops all I/O activity to the source volume.

- Synchronize phase: When all the systems comply with the quiesce request, all outstanding source volume updates are mirrored to the target volume by the TDMF Master.

- Volume I/O Redirect phase: In this phase, the TDMF Master requests that all TDMF Agents start redirecting I/Os and confirm back that redirection has been successfully started. The TDMF Master itself then redirects I/O activity from the source volume to the

target device. When this process is complete, the TDMF Master logically swaps the source with the target volume.

> **Tip:** If a system participating in the TDMF session fails an I/O redirection, TDMF resets the redirection on all systems (Master and Agents). TDMF proceeds to the *Resume phase* with an I/O error.

- Resume phase: Immediately after successful I/O redirect processing, the Master system runs **RESUME** processing and initiates the **RESUME** request for the Agent systems. Resuming allows user I/O to continue to the volume on its new device. After all systems process the **RESUME** request, the original (source) device is marked offline and put in a Soft Fenced state (if supported by the storage controller).

  If an I/O error occurs during the *Quiesce*, *Synchronize*, or I/O *Redirection phases*, the source device is re-enabled for I/O and the UCB is removed from the quiesce state.
- Terminate phase: When a volume completes a migration, the fixed storage of that volume is released for possible reuse in the current session. All dynamic allocations are also removed.

Hardware compatibility

TDMF is designed to support any Count Key Data/Extended (CKD/E) capable storage control units as a source or a target storage subsystem. These functions include Hyper-Volumes, Flexible-Volumes, and all user-defined volume sizes. Handling of extended address volumes (EAVs) depends on the z/OS level. See the *TDMF Installation and Reference Guide*, TDM-Z55IR-001, for more details.

Because TDMF z/OS is a host-based software migration tool, there are no specific hardware prerequisites.

Installation

There are two methods to install TDMF on a system:

- With SMP/E for z/OS
- Without SMP/E (preferred for migration service)

For installation details, see the *TDMF Installation and Reference Guide*, TDM-Z55IR-001.

Customer requirements

TDMF supports all z/OS-based operating systems that are currently supported by IBM. This product is not suitable for use with IBM z/VM®, Linux on System z, z/VSE®, or z/TPF.

If volumes to be migrated are accessed by any other non-z/OS operating systems, they can potentially be migrated by using an available z/OS system with physical access to those volumes. However, any non-z/OS system must be suspended (shut down) during the migration. To do so, the source and target volumes must be defined to the z/OS system where TDMF is supposed to perform the PIT migration that is specified with the REPLICATE control statement. After the volume migration (replication) completes, use ICKDSF REFORMAT to rename the target volumes to their original source volume serial numbers. Non-z/OS systems can now be restarted and are using the data on the new devices.

Special considerations must be taken when allocating the COMMDS in case z/OS runs as a guest system under z/VM. For more information, see the *TDMF Installation and Reference Guide*, TDM-Z55IR-001.

Tip: Perform periodic checks of the Required IBM Maintenance, IBM Technical Support (PTFs), and Technical Information Bulletins (TIBs). These requirements must be implemented to ensure successful TDMF operation.

Security

If the installation has a security package on the z/OS system on which TDMF is installed, you need to modify the security package. Otherwise, TDMF does not run properly. Check the profiles and command tables.

Limiting access to the TDMF authorized library to prevent unauthorized use of the TDMF system can be accomplished through security packages. The user submitting the SYSOPTN batch job (provided by the TDMF installation) must have UPDATE authority for the library indicated by the GTDKEY DD statement. The Master migration job also needs UPDATE authority for the GTDKEY file if the authorization key specifies volume or terabyte limits. To update authorization keys, use Option 10 of the TDMF Monitor as shown in Figure 10-3. The TSO user ID must have UPDATE authority for the GTDLLIB library.

```
 Softek TDMF Enterprise               Row 1 to 1 of 1
Command ===>                                          Scroll ===>

Company : IBM DEUTSCHLAND GMBH
Site  . : IBM MAINZ
Site ID : 11660
Enterprise Number  . . . . . . . . . . . : 50071
Enterprise Effective Date  . . . . . . . : September 24, 2014
Enterprise Expiration Date . . . . . . . December 31, 2014
Return Code from Security  . . . . . . . : 00000000
Features Enabled . . . . . . . . . . . . : MIGRATE  REPLICATE
                                           PPIT  OVA  TCP/IP

Commit Changes? (yes/no) . . . . . . . . ___
Softek TDMF for z/OS Version 5.5.0 Base Level.

   No    Security Key    Key Description        Added Date/Time    By

   01 A037858BE1C02F15  ENTERPRISE # 50071    09/24/2014 20:59 SYSOPTN
```

Figure 10-3 TDMF license keys

If the History option is selected, UPDATE authority is required for the data set specified. When viewing the history file (and any COMMDS) using the TDMF Monitor, the TSO user ID must have READ authority.

System Authorization Facility

If TDMF z/OS should use the installations SAF, select VOLUME SECURITY = YES when using the SYSOPTN batch job to set TDMF system default values. For more information about the TDMF system defaults, see TDMF System Defaults in the *TDMF Installation and Reference Guide*, TDM-Z55IR-001.

ALTER authority for the source and target volumes is required for a swap migration[1]. Error messages are issued for all volumes not meeting these requirements in a session.

[1] A swap migration is a TDMF z/OS migration where I/O is redirected from the source to the target volume after the migration completes. It is specified with the **MIGRATE** control statement in the TDMF Master control card input.

TDMF checks two types of classes:

- The DATASET class for the COMMDS history data set and the TDMF library defined on the GTDKEY DD statement
- The DASDVOL class for volumes allowed in a pairing

If the authorization keys specify volume or terabyte limits, the TSO user ID submitting the jobs must have update authority for the data set specified by the GTDKEY DD statement.

Protecting Data Integrity

As TDMF z/OS migrates data at the z/OS operating system level, all write I/Os to migration source volumes must be monitored before any I/O is finally redirected to the new target volumes. TDMF Agents must be run on all z/OS systems that access migration source and target volumes. It also must be prevented that any non-z/OS systems like z/VM, zLinux, Linux access source, and target volumes during the migration by either shutting down appropriate applications or systems.

TDMF z/OS takes extensive measures to protect data integrity during volume migrations:

- Before the actual migration (*Copy phase*) is started, TDMF z/OS analyzes physical access paths to the migration volumes and storage controllers. It then determines which systems have access and whether required TDMF z/OS Agents are running on those systems.
  - DS8000 storage systems can provide the *Query Host Access (QHA)* function used by TDMF z/OS to determine system access at the volume level
  - For storage systems not providing the QHA function logical path data is examined to determine system access at the storage control unit level
  - The default for the TDMF z/OS session option **`UNIDENT`** is defined as **`Terminate`**. This option prevents a migration session to start if not all systems accessing migration volumes have TDMF Agents running (or are manually excluded by the **`IGNORESYS`** session option)
- As soon as TDMF z/OS has determined that all required TDMF Agents are running and before entering the *Copy phase*, source and target volumes are put into the Set Path group ID (SPID) Fence state if the function is available on z/OS and the storage subsystem. This prevents any system that did not already have access at the time the migration was started to establish an access path to any of the migration volumes (SPID fenced offline devices cannot be brought online). The SPID Fence function can be available on DS8700, DS8800, and DS8870 disk subsystems. The SPID Fence state is removed as soon as the migration for a volume pair completes.
- After the migration of a volume pair completed and I/O has been redirected to the new target volume, the original source volume will be put into a Soft Fence state (if provided by the disk subsystem). This effectively prevents any I/O to the original source volume after I/O has been redirected by TDMF. The Soft Fence function can be provided by DS8870 storage systems. TDMF also takes target volumes out of a possible Soft Fence state before the migration starts.

> **Note:** By default, TDMF z/OS puts the original source volume into a Soft Fence state after a volume migration completes. To make these volumes available again, the Soft Fence state must be manually removed. For more information, see “Clearing a fence status using ICKDSF” on page 223. TDMF provides various parameters and options to control Soft Fencing.

Query Host Access (QHA)

DS8000 series storage subsystems can provide the Query Host Access function that provides system access information about the volume level. This facilitates TDMF during identification of required TDMF Agents.

To determine whether a device supports QHA, use the **DEVSERV QDASD** command with the **QHA** parameter. Example 10-1 lists 18 distinct path-group-ids for a device.

Example 10-1 Query Host Access (QHA)

```
DS QD,7802,QHA
IEE459I 11.57.31 DEVSERV QDASD 314
 UNIT VOLSER SCUTYPE DEVTYPE       CYL  SSID SCU-SERIAL DEV-SERIAL EFC
07802 JZ7802 2107961 2107900     30051  0078 0175-W2201 0175-W2201 *OK
  QUERY HOST ACCESS TO VOLUME
 PATH-GROUP-ID          FL STATUS  SYSPLEX   MAX-CYLS
880116AD6E2094CDF683E7  50 ON               268434453
88011EE15E2097CDF686A8  50 ON               268434453
880117E15E2097CDF6726D  50 ON               268434453
880005D5B62818CDF688D7  50 ON               268434453
800008D5B62818CDF667F2  50 ON      TDMFPLEX   1182006
80000B6D352817CDF6682B  50 ON      TDMFPLEX   1182006
80000A9B672827CDF66833* 50 ON      TDMFPLEX   1182006
80000A6D352817CDF66833  50 ON      TDMFPLEX   1182006
880334D46D2827CDF6818E  50 ON               268434453
8800079B672827CDF6773C  50 ON               268434453
88000D7B452817CDF686E0  50 ON               268434453
8800066D352817CDF6894E  50 ON               268434453
8800016D352817CDF67666  50 ON               268434453
8800099B672827CDF67C90  50 ON               268434453
8800026D352817CDF682A8  50 ON               268434453
8800047B452817CDF67DE7  50 ON               268434453
8800056D352817CDF687CA  50 ON               268434453
80000E6D352817CDF6927E  00 OFF     RALNS57    1182006
****     18 PATH GROUP ID(S) MET THE SELECTION CRITERIA
****      1 DEVICE(S) MET THE SELECTION CRITERIA
****      0 DEVICE(S) FAILED EXTENDED FUNCTION CHECKING
```

Each of the systems accessing the device must either run a TDMF Agent (in case it is a z/OS system), disconnect access to the device (**VARY *devn*,OFFLINE**), completely be shut down, or explicitly specified by using the TDMF session option **IGNORESYS** (in case the customer can safely assume that no data integrity exposure will be created by tolerating access).

Note: TDMF does not prevent uncontrolled I/Os to the source and target devices for any system specified by the IGNORESYS option. It is the user's responsibility to ensure data integrity by ignoring systems with an active physical connection to migration volumes.

Soft Fence

The **DEVSERV QDASD** command can be used to determine whether a device is in the Soft Fence state. In Example 10-2, the value SOF indicates that the device is in a Soft Fence state.

Example 10-2 Determining Soft Fence state by using DEVSERV QDASD

```
DS QD,7802
IEE459I 12.07.06 DEVSERV QDASD 064
 UNIT VOLSER SCUTYPE DEVTYPE       CYL  SSID SCU-SERIAL DEV-SERIAL EFC
07802 JZ7802 2107961 2107900     30051  0078 0175-W2201 0175-W2201 SOF
****      1 DEVICE(S) MET THE SELECTION CRITERIA
****      1 DEVICE(S) FAILED EXTENDED FUNCTION CHECKING
```

In Example 10-3, the **DEVSERV PATH** command is used to indicate whether the device is in a Soft Fence state.

Example 10-3 Determining Soft Fence state by using DEVSERV PATH

```
DS PATH,7802
IEE459I 12.07.39 DEVSERV PATHS 066
 UNIT DTYPE  M CNT VOLSER  CHPID=PATH STATUS
      RTYPE  SSID CFW TC   DFW PIN DC-STATE CCA DDC       CYL CU-TYPE
07802,33909 ,0,000,JZ7802,63=+ 71=+
      2107   0078  Y  YY.  YY.  N  SIMPLEX   02  02     30051 2107
  ** FENCED DEVICE            D8FFFF00 D0FFFF00 D0FFFF00 D0FFFF00
                              00000000 00000000
************************ SYMBOL DEFINITIONS ************************
O = ONLINE                         + = PATH AVAILABLE
```

The **DEVSERV PATH** command displays "FENCED DEVICE" if the device is soft fenced.

Set Path Group ID (SPID) Fence

DS8000 series storage subsystems can provide the SPID Fence function. This function is used by TDMF to prevent unintentional access to volumes currently being migrated. A z/OS system cannot bring a volume online that is in a SPID Fence state. While a TDMF Master system migrates data, all z/OS images that access those volumes must run a TDMF agent to prevent data integrity issues. The SPID Fence function prevents any system that did not access any migration volume at the time a TDMF migration is started[2] to establish an access path to those volumes.

The **DEVSERV QDASD** command can be used to determine whether a device is in the SPID Fence state. In Example 10-4, the value SPF indicates that the device is in a SPID Fence state.

Example 10-4 Determining SPID Fence state by using DEVSERV QDASD

```
DS QD,7803
IEE459I 09.22.18 DEVSERV QDASD 136
 UNIT VOLSER SCUTYPE DEVTYPE       CYL  SSID SCU-SERIAL DEV-SERIAL EFC
07803 JZ7803 2107961 2107900     30051  0078 0175-W2201 0175-W2201 SPF
****      1 DEVICE(S) MET THE SELECTION CRITERIA
****      1 DEVICE(S) FAILED EXTENDED FUNCTION CHECKING
```

[2] For this reason, it did not require a TDMF Agent

In Example 10-5, the **DEVSERV PATH** command is used to indicate whether the device is in a SPID Fence state.

Example 10-5 Determining SPID Fence state by using DEVSERV PATH

```
DS PATH,7803
IEE459I 09.23.12 DEVSERV PATHS 138
 UNIT DTYPE  M CNT VOLSER  CHPID=PATH STATUS
      RTYPE  SSID CFW TC   DFW PIN DC-STATE CCA DDC       CYL CU-TYPE
07803,33909 ,A,000,JZ7803,63=+ 71=+
      2107   0078  Y  YY.  YY.  N  SIMPLEX   03  03     30051 2107
  ** FENCED SPID               CCFFFF00 C4FFFF00 C4FFFF00 C4FFFF00
                               00000000 00000000
************************ SYMBOL DEFINITIONS ************************
A = ALLOCATED                      + = PATH AVAILABLE
```

Clearing a fence status using ICKDSF

A fenced status that is caused by TDMF can be reset for a device using ICKDSF. Example 10-6 depicts how the Soft Fence and SPID Fence state can be reset for unit 7802.

Example 10-6 Resetting Soft or SPID Fence state using ICKDSF

```
//STEP     EXEC PGM=ICKDSF,PARM='NOREPLYU'
//SYSPRINT DD SYSOUT=*
//SYSIN    DD *
 CONTROL -
 UNIT(7802) -
 CLEARFENCE -
 SERIAL(W2201) -
 SCOPE(DEV)
```

10.2.2 Example migration scenario in a shared storage environment

This section illustrates a simple example of migrating data from an ESS 800 to a DS8000. In the example four volumes are moved, but the process would be the same for any number of volumes.

Experimental environment

The experimental environment has the following characteristics:

- Hardware and software setup as shown in Figure 10-4

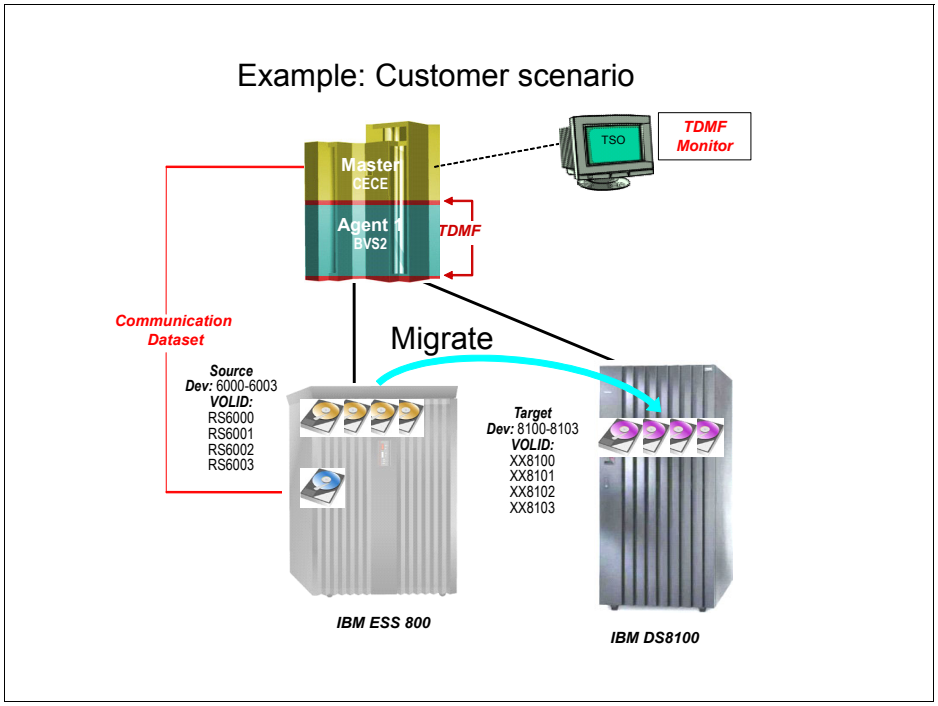

Figure 10-4 Example: Graphical overview (customer scenario)

- z/OS level 01.12.00 (Figure 10-5).

```
IEE421I RO *ALL,D IPLINFO 955
MCECEBC  RESPONSES ------------------------------------
IEE254I  18.27.15 IPLINFO DISPLAY 954
 SYSTEM IPLED AT 12.41.01 ON 01/21/2014
 RELEASE z/OS 01.12.00    LICENSE = z/OS
 USED LOADB9 IN SYS1.IPLPARM ON CA13
 ARCHLVL = 2   MTLSHARE = N
 IEASYM LIST = BC
 IEASYS LIST = 00
 IODF DEVICE: ORIGINAL(CA13) CURRENT(CA13)
 IPL DEVICE: ORIGINAL(CA10) CURRENT(CA10) VOLUME(CEBCR1)
MZBCVS2  RESPONSES ------------------------------------
IEE254I  18.27.15 IPLINFO DISPLAY 855
 SYSTEM IPLED AT 13.35.54 ON 01/27/2014
 RELEASE z/OS 01.12.00    LICENSE = z/OS
 USED LOADS9 IN SYS1.IPLPARM ON CA13
 ARCHLVL = 2   MTLSHARE = N
 IEASYM LIST = S2
 IEASYS LIST = 00
 IODF DEVICE: ORIGINAL(CA13) CURRENT(CA13)
 IPL DEVICE: ORIGINAL(CA10) CURRENT(CA10) VOLUME(CEBCR1)
```

Figure 10-5 Display IPL information including the z/OS level

- Two LPARs (LPAR 4 and 5) as seen in Figure 10-6

```
IEE421I RO *ALL,D M=CPU 973
MCECEBC  RESPONSES ----------------------------------------
IEE174I 18.31.15 DISPLAY M 972
PROCESSOR STATUS
ID  CPU                  SERIAL
00  +                     04423A2064
01  +                     14423A2064
02  +                     24423A2064
03  +                     34423A2064
04  +                     44423A2064
05  +                     54423A2064
06  +                     64423A2064

CPC ND = 002064.116.IBM.51.00000002423A
CPC SI = 2064.116.IBM.51.000000000002423A
CPC ID = 00
CPC NAME = P002423A
LP NAME = BVS1       LP ID = 4
MIF ID  = 4

MZBCVS2  RESPONSES -------------------------
IEE174I 18.31.15 DISPLAY M 857
PROCESSOR STATUS
ID  CPU                  SERIAL
00  +                     05423A2064
01  +                     15423A2064
02  +                     25423A2064
03  +                     35423A2064
04  +                     45423A2064
05  +                     55423A2064
06  +                     65423A2064

CPC ND = 002064.116.IBM.51.00000002423A
CPC SI = 2064.116.IBM.51.000000000002423A
CPC ID = 00
CPC NAME = P002423A
LP NAME = BVS2       LP ID = 5
MIF ID  = 5
```

Figure 10-6 Displaying processor information

- Storage subsystems: IBM TotalStorage ESS 800 and IBM System Storage DS8000
- Channel attachment (ESCON / FICON) for storage

Migration steps

Complete the following steps to migrate data with TDMF z/OS:

1. Verify that all volumes to be migrated are online to the Master TDMF system as shown in Figure 10-7. In the example, one volume is offline on the TDMF Master.

```
D U,,,6000,4
IEE457I 22.56.27 UNIT STATUS 111
UNIT TYPE STATUS        VOLSER     VOLSTATE
6000 3390 A             RS6000     PRIV/RSDNT
6001 3390 O             RS6001     PRIV/RSDNT
6002 3390 A             RS6002     PRIV/RSDNT
6003 3390 OFFLINE       RS6003     PRIV/RSDNT
```

Figure 10-7 Display source device address on the TDMF Master system

2. Bring the volume online by issuing a **VARY ONLINE** command for the offline devices as shown in Figure 10-8.

```
V 6003,ONLINE
IEE302I 6003 ONLINE
```

Figure 10-8 Bring source device online

3. Check that the target devices (8100-8103) are online to the TDMF Master system. Figure 10-9 shows that all the devices are offline.

```
D U,,,8100,4
IEE457I 13.23.16 UNIT STATUS 583
UNIT TYPE STATUS        VOLSER     VOLSTATE
8100 3390 OFFLINE                      /RSDNT
8101 3390 OFFLINE                      /RSDNT
8102 3390 OFFLINE                      /RSDNT
8103 3390 OFFLINE                      /RSDNT
```

Figure 10-9 Check target device address

4. Bring the target devices online on the TDMF Master system by issuing the **VARY ONLINE** command to the devices as seen in Figure 10-10.

```
V 8100-8103,ONLINE
IEE457I 23.05.42 UNIT STATUS 238
UNIT TYPE STATUS        VOLSER     VOLSTATE
8100 3390 O             XX8100     PRIV/RSDNT
8101 3390 O             XX8101     PRIV/RSDNT
8102 3390 O             XX8102     PRIV/RSDNT
8103 3390 O             XX8103     PRIV/RSDNT
```

Figure 10-10 Vary target devices online

5. After all devices (source and target) used in the data migration are online to the TDMF Master system, start the TDMF Monitor. The customized member to start the monitor in the experimental environment is stored in library `SYS1.IBM.HGTD550.GTDELIB`. Start the TDMF Monitor by issuing the **`TSO EXEC`** command as shown in Figure 10-11.

```
DSLIST - Data Sets Matching SYS1.IBM.HGTD550.*                    Row 1 of 7
Command ===> TSO EX 'SYS1.IBM.HGTD550.GTDELIB(TDMF)'          Scroll ===> CSR

Command - Enter "/" to select action                 Message          Volume
-------------------------------------------------------------------------------
         SYS1.IBM.HGTD550.GTDELIB                                      VSLNSM
         SYS1.IBM.HGTD550.GTDLLIB                                      VSLNSM
         SYS1.IBM.HGTD550.GTDMLIB                                      VSLNSM
         SYS1.IBM.HGTD550.GTDPLIB                                      VSLNSM
         SYS1.IBM.HGTD550.GTDTLIB                                      VSLNSM
         SYS1.IBM.HGTD550.SAMPLIB                                      VSLNSM
         SYS1.IBM.HGTD550.TIB                                          VSLNSM
***************************** End of Data Set list ***************************
```

Figure 10-11 Start the TDMF Monitor

6. The TDMF start window is displayed as seen in Figure 10-12. Press Enter to continue.

```
 Command ===>

              TTTTTTTTTT    DDDDDDDDD     MM      MM    FFFFFFFFFF
              TTTTTTTTTT    DDDDDDDDDD    MMM    MMM    FFFFFFFFFF
                  TT        DD      DD    MMMM  MMMM    FF
                  TT        DD      DD    MM MMMM MM    FF
                  TT        DD      DD    MM  MM  MM    FFFFFFFF
                  TT        DD      DD    MM      MM    FFFFFFFF
                  TT        DD      DD    MM      MM    FF
                  TT        DD      DD    MM      MM    FF
                  TT        DDDDDDDDDD    MM      MM    FF
                  TT        DDDDDDDDD     MM      MM    FF

            Transparent Data Migration Facility for z OS Version 5.5.0
                       LICENSED MATERIALS - PROPERTY OF IBM

                                Press Enter to continue
```

Figure 10-12 TDMF start window

7. The TDMF Monitor Menu is displayed as seen in Figure 10-13. Each option available is covered in detail in the *TDMF Installation and Reference Guide*, TDM-Z55IR-001. Select option 0 to open `SYS1.IBM.HGTD550.SAMPLIB`, where you can find example migration jobs.

```
 Softek TDMF Monitor Menu
Option ===>

   0  Change or Submit Replication or Migration Jobs
   1  Current Sessions: Monitor Progress
   2  Current Sessions: User Interaction and Status
   3  Current Sessions: Display Messages
   4  Current Sessions: Display Offline Volume Access (OVA)
   5  Current/Past Sessions: Display Performance Data
   6  Past Sessions: Display Summary
   7  Past Sessions: Display Details
   8  Past Sessions: Display Communication Dataset History
   9  Display Installation Options and Environment
   10 Display/Modify Installation Security Environment
   11 Display/Modify User's TSO Monitor Options
   12 Build Replication or Migration Jobs
   U  Support Utilities
   H  Help and Message Detail Facility
```

Figure 10-13 TDMF Monitor Menu

8. In the experimental environment, TDMF migration jobs have been created from the sample members and stored in `STXXXXX.TDMF.ITSO.TDMF.CNTL`. The customized JCL to define a new COMMDS is shown in Figure 10-14.

```
EDIT       STXXXXX.TDMF.ITSO.TDMF.CNTL(CDSL$G01) - 01.00   Columns 00001 00072
Command ===>                                                  Scroll ===> CSR
****** **************************** Top of Data ****************************
000001 //CDSL$G01 JOB A185,'IBM',MSGCLASS=I,REGION=0M,
000002 //      CLASS=A,TIME=1440,NOTIFY=&SYSUID
000003 /*JOBPARM S=CEBC
000004 //*****************************************************
000005 //* Create SYSCOM DATA SETS AT LOCAL SITE
000006 //*****************************************************
000007 //STEP01  EXEC PGM=IEFBR14
000008 //SYSPRINT DD SYSOUT=*
000009 //SYSCOM   DD DSN=SYS1.TDMF.ITSO.LOCAL.SYSCOM01,
000010 //            DISP=(NEW,CATLG,DELETE),
000011 //            SPACE=(CYL,70,,CONTIG),UNIT=3390,
000012 //            DCB=(LRECL=4096,BLKSIZE=4096,RECFM=F,DSORG=PS),
000013 //            VOL=SER=RS6100
000014 /*
****** ***************************** Bottom of Data **************************
```

Figure 10-14 Define new COMMDS

Figure 10-15 shows customized JCL for a MASTER job.

```
EDIT       STXXXXX.TDMF.ITSO.TDMF.CNTL(TDMLBC01) - 01.01   Columns 00001 00072
Command ===>                                                  Scroll ===> CSR
****** **************************** Top of Data ****************************
000001 //TDMLBC01 JOB A185,'IBM',MSGCLASS=I,REGION=0M,
000002 //      CLASS=A,TIME=1440,NOTIFY=&SYSUID
000003 /*JOBPARM S=CEBC
000004 //*****************************************************
000005 //* Master JOB TDMF running: LOCAL
000006 //*****************************************************
000007 //LOCAL01  EXEC PGM=GTDMAIN,PARM=MASTER,REGION=0M
000008 //STEPLIB  DD DISP=SHR,DSN=SYS1.IBM.HGTD550.GTDLLIB
000009 //GTDKEY   DD DISP=SHR,DSN=SYS1.IBM.HGTD550.GTDLLIB
000010 //SYSCOM   DD DISP=SHR,DSN=SYS1.TDMF.ITSO.LOCAL.SYSCOM01
000011 //SYSPRINT DD SYSOUT=*
000012 //SYSSNAP  DD SYSOUT=*
000013 //SYSIN    DD *
000014 SESSION M(CEBC) AGENT(CVS2)
000015 OPT(FAST UNIDENT(T) NOPROMPT CONC(2)
000016 PAC CHECKT)
000017 MIGRATE RS6000 XX8100 YY6000
000018 MIGRATE RS6001 XX8101 YY6001
000019 MIGRATE RS6002 XX8102 YY6002
000020 MIGRATE RS6003 XX8103 YY6003
****** ***************************** Bottom of Data **************************
```

Figure 10-15 Master job

Figure 10-16 shows customized JCL for an AGENT job

```
EDIT       STXXXXX.TDMF.ITSO.TDMF.CNTL(TDMLBC01) - 01.01   Columns 00001 00072
Command ===>                                                  Scroll ===> CSR
****** **************************** Top of Data ****************************
000001 //TDALCV01 JOB A185,'IBM',MSGCLASS=I,REGION=0M,
000002 //      CLASS=A,TIME=1440,NOTIFY=&SYSUID
000003 /*JOBPARM S=CVS2
000004 //*****************************************************
000005 //* AGENT JOB TDMF running: CVS2
000006 //*****************************************************
000007 //STEP01   EXEC PGM=GTDMAIN,PARM=AGENT,REGION=0M
000008 //STEPLIB  DD DISP=SHR,DSN=SYS1.IBM.HGTD550.GTDLLIB
000009 //GTDKEY   DD DISP=SHR,DSN=SYS1.IBM.HGTD550.GTDLLIB
000010 //SYSCOM   DD DISP=SHR,DSN=SYS1.TDMF.ITSO.LOCAL.SYSCOM01
000011 //SYSPRINT DD SYSOUT=*
000012 //SYSUDUMP DD SYSOUT=*
000013 //SYSSNAP  DD SYSOUT=*
****** ***************************** Bottom of Data **************************
```

Figure 10-16 Agent job

9. Before starting the actual migration, a communications data set must be allocated by submitting the appropriate job as shown in Figure 10-14 on page 228. Also, see “Sizing a COMMDS” on page 235.
10. After verifying that the new COMMDS exists, submit the TDMF Master and Agent jobs as shown in Figure 10-15 on page 229 and Figure 10-16 on page 229.
11. Change to the TDMF Monitor using option 1 (`Current Sessions: Monitor Progress`) in the Monitor menu as seen in Figure 10-13 on page 228. Option 1 allows you to monitor the progress of the current session (Figure 10-17). In this scenario, four volumes are being migrated.

```
                                Session Monitor                Row 1 to 8 of 8
Command ===>                                                  Scroll ===> PAGE

Softek TDMF Master V5.5.0 Session Active.
ComDataSet : SYS1.TDMF.ITSO.LOCAL.SYSCOM01
         Source  Migration    Percent Complete --------->
Group    VolSer    Phase      ..10...20...30...40...50...60...70...80...90..100
         RS6000   Inactive
         RS6001   Inactive
         RS6002   Inactive
         RS6003   Inactive
```

Figure 10-17 Monitor the session (migration still inactive)

After repeatedly pressing Enter to refresh the screen, you can see two of the volumes copying data simultaneously. The progress monitor is updated to show the migration phases and the percentage of completion for each phase after pressing the Enter key. You also get an estimated end time for the initial *Copy phase* for each volume as shown in Figure 10-18.

```
                                Session Monitor                Row 1 to 8 of 8
Command ===>                                                  Scroll ===> PAGE

Softek TDMF Master V5.5.0 Session Active.
ComDataSet : SYS1.TDMF.ITSO.LOCAL.SYSCOM01
         Source  Migration    Percent Complete --------->
Group    VolSer    Phase      ..10...20...30...40...50...60...70...80...90..100
         RS6000   Copy        --------------->                          0:01:04
         RS6001   Copy        --------------->                          0:01:04
         RS6002   Inactive
         RS6003   Inactive
```

Figure 10-18 Monitor the session (2 volumes in Copy phase)

Figure 10-19 displays the output of the session monitor with the first two volumes in the *Refresh phase* at 100% complete. The copy has not yet started for the next two volumes.

```
 Session Monitor                Row 1 to 8 of 8
Command ===>                                              Scroll ===> PAGE

Softek TDMF Master V5.5.0 Session Active.
ComDataSet : SYS1.TDMF.ITSO.LOCAL.SYSCOM01
         Source  Migration    Percent Complete --------->
Group    VolSer    Phase      ..10...20...30...40...50...60...70...80...90..100
         RS6000   Refresh 1   ------------------------------------------------>
         RS6001   Refresh 1   ------------------------------------------------>
         RS6002   Inactive
         RS6003   Inactive
```

Figure 10-19 Monitor the session (2 volumes reached Refresh phase)

Figure 10-20 shows that the first two volumes are completed and the next two volumes are in the *Copy phase*. Based on the parameter (**CONCurrant=2**), only two volumes are processed at a time. However, this parameter can be dynamically changed in the monitor.

```
 Session Monitor                Row 1 to 8 of 8
Command ===>                                              Scroll ===> PAGE

Softek TDMF Master V5.5.0 Session Active.
ComDataSet : SYS1.TDMF.ITSO.LOCAL.SYSCOM01
         Source  Migration    Percent Complete --------->
Group    VolSer    Phase      ..10...20...30...40...50...60...70...80...90..100
         RS6000   Complete
         RS6001   Complete
         RS6002   Copy        -------------->                          0:01:04
         RS6003   Copy        --------------->                         0:01:03
```

Figure 10-20 Monitor the session (2 volumes complete, 2 in Copy phases)

Figure 10-21 to Figure 10-23 on page 232 illustrates the transition from the *Refresh phase* to the *Synchronize phase* for the second set of volumes.

```
 Session Monitor                Row 1 to 8 of 8
Command ===>                                              Scroll ===> PAGE

Softek TDMF Master V5.5.0 Session Active.
ComDataSet : SYS1.TDMF.ITSO.LOCAL.SYSCOM01
         Source  Migration    Percent Complete --------->
Group    VolSer    Phase      ..10...20...30...40...50...60...70...80...90..100
         RS6000   Complete
         RS6001   Complete
         RS6002   Refresh 1   ------------------------------------------------>
         RS6003   Refresh 1   ------------------------------------------------>
```

Figure 10-21 Monitor the session (two volumes still in Refresh phase)

```
 Session Monitor                Row 1 to 8 of 8
Command ===>                                                Scroll ===> PAGE

Softek TDMF Master V5.5.0 Session Active.
ComDataSet : SYS1.TDMF.ITSO.LOCAL.SYSCOM01
         Source  Migration    Percent Complete --------->
Group    VolSer    Phase      ..10...20...30...40...50...60...70...80...90..100
         RS6000   Complete
         RS6001   Complete
         RS6002   Quiescing   ------------------------------------------------>
         RS6003   Quiescing   ------------------------------------------------>
```

Figure 10-22 Monitor the session (two volumes quiescing)

```
 Session Monitor                Row 1 to 8 of 8
Command ===>                                                Scroll ===> PAGE

Softek TDMF Master V5.5.0 Session Active.
ComDataSet : SYS1.TDMF.ITSO.LOCAL.SYSCOM01
         Source  Migration    Percent Complete --------->
Group    VolSer    Phase      ..10...20...30...40...50...60...70...80...90..100
         RS6000   Complete
         RS6001   Complete
         RS6002   Synchronize ------------------------------------------------>
         RS6003   Synchronize ------------------------------------------------>
```

Figure 10-23 Monitor the session (two volumes synchronizing)

Figure 10-24 shows that the migration for this session is now complete. This was just a simple migration with only four volume pairs. A typical migration scenario usually involves many more volumes.

```
 Session Monitor       ** no sessions active **
Command ===>                                                Scroll ===> PAGE

No Softek TDMF sessions active.
```

Figure 10-24 TDMF session complete

12. Check the online state of the volumes. As you can see in Figure 10-25, the old source volumes 6100-6103 are now offline to all systems participating in the migration. The volume serial numbers are changed according to the **MIGRATE** statements used in the TDMF Master JCL.

```
RO *ALL,D U,,,6000,4
IEE421I RO *ALL,D U,,,6000,4 978
MCECEBC  RESPONSES ---------------------------------------------------
IEE457I 23.05.35 UNIT STATUS 977
UNIT TYPE STATUS        VOLSER     VOLSTATE
6000 3390 OFFLINE       YY6000         /RSDNT
6001 3390 OFFLINE       YY6001         /RSDNT
6002 3390 OFFLINE       YY6002         /RSDNT
6003 3390 OFFLINE       YY6003         /RSDNT
MZBCVS2  RESPONSES ---------------------------------------------------
IEE457I 23.05.35 UNIT STATUS 884
UNIT TYPE STATUS        VOLSER     VOLSTATE
6000 3390 OFFLINE       YY6000         /RSDNT
6001 3390 OFFLINE       YY6001         /RSDNT
6002 3390 OFFLINE       YY6002         /RSDNT
6003 3390 OFFLINE       YY6003         /RSDNT
******************************** BOTTOM OF DATA ********************************
```

Figure 10-25 Query old source volumes after migration

Figure 10-26 shows the status of the old target volumes (8100-8103). These volumes are now online with the volume serial numbers of the original source volumes. They are now transparently being used by applications without interruption by the physical data movement. The data migration is complete.

```
IEE421I RO *ALL,D U,,,8100,4 007
MCECEBC  RESPONSES ---------------------------------------------
IEE457I 14.26.09 UNIT STATUS 006
UNIT TYPE STATUS        VOLSER     VOLSTATE
8100 3390 O             RS6100     PRIV/RSDNT
8101 3390 O             RS6101     PRIV/RSDNT
8102 3390 O             RS6102     PRIV/RSDNT
8103 3390 O             RS6103     PRIV/RSDNT
MZBCVS2  RESPONSES ---------------------------------------------
IEE457I 14.26.09 UNIT STATUS 589
UNIT TYPE STATUS        VOLSER     VOLSTATE
8100 3390 O             RS6100     PRIV/RSDNT
8101 3390 O             RS6101     PRIV/RSDNT
8102 3390 O             RS6102     PRIV/RSDNT
8103 3390 O             RS6103     PRIV/RSDNT
******************************** BOTTOM OF DATA ***************
```

Figure 10-26 Verify that the new source volumes are online

10.2.3 Tips and preferred practices

This section offers general advice and preferred practices for migrating disk volumes using TDMF z/OS.

Keeping current

Before you start the migration, make sure that your installed TDMF z/OS software is up-to-date by using these steps:

- Install each program temporary fix (PTF) as it becomes available.
- Manuals are periodically updated. Make sure that you have the most current information available by checking for TIBs. TIBs are typically issued between documentation releases.

Setting default options

As part of the postinstallation tailoring of TDMF, option defaults are set by the SYSOPTN batch job. For example, if you want to use pacing in every migration session, set this value to YES in the SYSTOPTN batch job during first activation. Those values will be the default for all subsequent sessions.

TDMF system default values set by the SYSOPTN batch job can be overridden within a session. The following is the order of precedence:

- Session statement
- Group statements
- Migrate statements

For more information, see the *TDMF Installation and Reference Guide*, TDM-Z55IR-001, *Softek TDMF Control Statements*.

Storage requirements

The storage required for running TDMF Master and Agents depends on the number of volumes, the number of Agent systems, and the values of various TDMF options.

For example, a migration of 16 volumes using one TDMF Agent (two systems in total) without specifying the **OVA** (*offline volume access*) option has the following storage requirements for the Master and Agent systems:

- Master system:
  - Fixed common storage: 612 K HVCOMMON, 9 K ESQA, 256-bytes CSA
  - Pageable common storage: 28 K HVCOMMON
  - Fixed extended private area storage: 16,600 K
- Agent system:
  - Fixed common storage: 617 K HVCOMMON, 5 K ESQA, 256-bytes CSA
  - Pageable common storage: 28 K HVCOMMON
  - Fixed extended private area storage: 656 K

If the **COMPare** common option or *Full Speed Copy* (**PAcing(FULLSPEED)**) volume option is requested, an extra 900 K buffer for each volume pair is allocated. For more information about storage requirements, see *Storage Requirements* in chapter 2 of the *TDMF Installation and Reference Guide*, TDM-Z55IR-001.

Communications data set

Each TDMF session requires a unique communications data set (COMMDS) shared between the Master and Agent systems. The COMMDS must not be on any source or target migration volume. Best practice is to allocate the COMMDS for each concurrently run TDMF session on

a separate volume that is not managed by SMS. The COMMDS must be placed on a device that supports CKD/E.

Important: Unpredictable results occur when multiple sessions use the same COMMDS.

If a problem occurs, the COMMDS is the primary tool for problem determination and resolution. Although reusing a COMMDS is possible, be aware that the previous session data is no longer available. Do not reuse the COMMDS to avoid overwriting it so that it is available for problem determination and resolution.

Tip: Use a new COMMDS for each session, so that if a problem occurs, all information is available for analysis.

Member ALLOCCM in the SAMPLIB can be used to create a customized COMMDS allocation JCL. The communications data set must be physically on a cylinder boundary with contiguous space. For an example, see Figure 10-14 on page 228.

COMMDS reserve

TDMF periodically issues a RESERVE for the COMMDS to serialize communication between the Master and Agent systems. For more information, see *Unicenter CA-MIM Resource Sharing* and *Global Resource Serialization* in chapter 5 of the *TDMF Installation and Reference Guide*, TDM-Z55IR-001.

Sizing a COMMDS

To calculate the size of a COMMDS, use the following formula:

$$\text{COMMDS CYLS} = \text{MIN}\{256, V * (S + K)\}$$

where:

V is a number that depends on the number of migration volume pairs as follows:

- 64 volumes = 2.5
- 128 volumes = 5.0
- 256 volumes = 7.5
- 512 volumes = 10.0

S is the number of participating systems

K is the size of the source volumes that are involved as follows:

- For 3390-3, K = 4
- For 3390-9, K = 6
- For 3390-27, K = 15
- For 3390-54, K = 47
- For 3390-A, K = 184

For example, consider a TDMF migration session with 128 3390-27 and 128 3390-54 volumes with eight participating systems. The largest device type in the session determines the K value (*K* = 47) and the size is calculated as follows:

$$\text{CYLS} = 7.5 * (8 + 47) = 412.5$$

For this example, a communication data set with a size of 413 3390 cylinders would be required.

Tip: Although the COMMDS can be defined as a generation data group (GDG), this is not recommended because they can be overwritten during the migration process. Instead, use non-GDG data sets, which are the default.

For more information, see *Placement of the Communications data set* in Chapter 4 (*TDMF Performance Considerations*) of the *TDMF Installation and Reference Guide*, TDM-Z55IR-001.

Participation of Agent systems

In a shared DASD environment, any z/OS system that is not running the TDMF Master and has access to a source or target migration volume must run a TDMF Agent. Each Master or Agent system only monitors I/O to the source or target device from its perspective. If a TDMF migration session is run without including all systems connected to any source or target device, data integrity might be compromised.

The **`UNIDENTifiedsystems`** option can be used to verify that all required TDMF Agents are being run. The default action can be set by using the SYSOPTN batch job, and can be overridden using the options within the SESSION control statement. The following options are available to specify the action taken if a system accessing any source or target migration volume is detected without running a TDMF Agent:

- Terminate: Issues an error message and terminates the migration/replication (RC08).
- Warning: Issues a warning message, but the migration/replication continues (RC04).
- Informational: Issues an informational message and the migration/replication continues (RC00).

TDMF only detects systems accessing migration volumes for storage systems that support either logical path data (3990-6 or later) or the QHA function (DS8K).

Important: You are responsible for running TDMF Agents on all z/OS systems accessing migration volumes. Non-z/OS systems must be prevented from accessing source or target volumes during the migration.

For more information, see Chapter 3, *Session Options* in the *TDMF Installation and Reference Guide*, TDM-Z55IR-001.

Protection of target volume data

If a target volume contains data, TDMF overwrites that data if no files are open at the time the session starts. If a file is open on a target volume at volume initiation, TDMF terminates that migration session. The system displays return code 12 (RC 12) and any error messages pertinent to the situation.

The CHECK TARGET option ensures that TDMF does not overwrite data on a target device. Selection of this option informs TDMF that only the VTOC, VTOCIX, and VVDS entries are allowed on the target volume.

Tip: Use the **`CHECKTarget`** option as default to protect your data. Set the **`CHECKTarget`** option in each session control statement.

For more information, see the *Session Options* in Chapter 3 (*TDMF Control Statements*) of the *TDMF Installation and Reference Guide*, TDM-Z55IR-001.

Pacing

Always use pacing so that application performance is not affected by a TDMF migration session. Some volumes regularly experience high I/O rates and the normal pacing feature of TDMF reduces the amount of system resources consumed until the application impact is tolerable.

The following options are available to control pacing:

- Reverse pacing
- Full speed copy
- FastCopy
- Active in Copy

Pacing is active during the *Copy* and *Refresh phases* only. After the volumes or session enter the *Quiesce phase* and continue through the *Termination phase*, pacing is no longer used. For more information, see *Phases of Migration* in Chapter 1 (*TDMF for z/OS Overview*) of the *TDMF Installation and Reference Guide*, TDM-Z55IR-001.

Tip: Generally, use the default pacing options. Move frequently updated volumes during a low activity time period.

Reverse pacing

Reverse pacing starts at one track per I/O operation, and scales upwards depending on the activity on the source volume. For more information, see *Dynamic Volume Pacing* in Chapter 4 (*TDMF Performance Considerations*) of the *TDMF Installation and Reference Guide*, TDM-Z55IR-001. This setting is a more conservative approach compared to using normal pacing because it slowly increases the amount of system resources TDMF consumes.

Full speed copy

This option causes TDMF to increase the number of buffers from one to two to allow interleaving of the I/O. Using two buffers shortens the migration time. Use this option judiciously because the additional load to the source and target volumes can negatively affect applications and the migration itself. Field statistics show that the decrease in migration time is 40% on average. For more information, see *Full Speed Copy Impact* in Chapter 3 (*TDMF Control Statements*) of the *TDMF Installation and Reference Guide*, TDM-Z55IR-001.

FastCopy

FastCopy instructs TDMF to copy only allocated tracks and cylinders, ignoring empty ones. This option can decrease the amount of time that is required to migrate a volume with no performance impact. Always use the FastCopy option unless there are reasons that prevent its use like migrating volumes of a non-z/OS system while those systems are down (for example, z/VM or Linux for System z disks). For more information, see *Common Options for the SESSION, GROUP, MIGRATE, and REPLICATE Statements* in Chapter 3 (*TDMF Control Statements*) of the *TDMF Installation and Reference Guide*, TDM-Z55IR-001.

Active in Copy

The volume pairs in a migration are considered active from volume initialization to volume termination. Active in Copy instructs TDMF to treat a volume as being active from volume initialization through the completion of the first refresh pass. Use of this option limits the number of active copy tasks, but speeds the normal copy process. For more information, see the *TDMF Installation and Reference Guide*, TDM-Z55IR-001, Chapter 3 *Session Options*.

Avoid COMPARE option for production volumes

TDMF performs data comparison between source and target volumes if requested by the **COMPARE** option. This comparison takes place in the *Quiesce phase* and can take a considerable amount of time during that no I/O to the production source volume is performed. Usually this is not tolerable, and so the **COMPARE** option should not be specified.

Rank contention and storage subsystem performance

When large numbers of volumes are to be copied, it is important to understand the following about the system:

- The architecture of the subsystem
- The logical to physical mapping of the volumes within the array

Use this information during the planning, coding, and execution of TDMF sessions. This information helps ensure that these sessions do not adversely affect application, array, and subsystem performance.

It is advised to perform some experimental migrations. By using a performance monitor like RMFMON, check if there is an impact to system performance. Also, check channel utilization to make sure that it is no more than 70-75% per channel path. Change the number of concurrent running volumes based on the results using the TDMF Monitor.

For more information, see the *TDMF Installation and Reference Guide*, TDM-Z55IR-001, Chapter 4, *RAID Subsystems and Rank Contention*.

TDMF interaction with other programs

Before using TDMF, see the *TDMF Installation and Reference Guide*, TDM-Z55IR-001, *Softek TDMF Planning Considerations*. Ensure that all requirements are met for running TDMF sessions.

If you use an array-based mirroring solution in your environment, it must be protected. Protection is especially important if you need to keep that environment active during a migration. In this case, review the following topics in chapter *Softek TDMF Planning Considerations*:

- Migrating from SRDF to PPRC Volumes
- Migrating to SRDF Volumes
- Peer-to-Peer Remote Copy

If GDPS/PPRC (HyperSwap) is active, be aware that HyperSwap is turned off by TDMF at the time of swap. See the *TDMF Installation and Reference Guide*, TDM-Z55IR-001 for details.

Identification of volumes requiring special handling

Some volumes might require special handling before, during, or after a migration. These volumes need to be identified and potentially put into one or more separate sessions.

In addition, some volume migrations require an application recycle to complete the migration. For more information, see the *TDMF Installation and Reference Guide*, TDM-Z55IR-001, chapter *Softek TDMF Planning Considerations*.

Important: Do not migrate active Local Page data sets or Sysplex Couple data sets.

Migrating Local Page data sets

Local Page data sets cannot be moved while active. To migrate them, complete the following steps:

1. Analyze location of these data sets by using the **D ASM** command at each system accessing the source/target storage (Example 10-7).

   Example 10-7 Display Local Page data sets

   ```
   COMMAND INPUT ===> /D ASM                                    SCROLL ===> CSR
   RESPONSE=MCECEBC
   IEE200I 11.32.10 DISPLAY ASM 536
   TYPE     FULL STAT   DEV DATASET NAME
   PLPA      82%   OK  5249 PAGE.MCECEBC.PLPA
   COMMON     0%   OK  5249 PAGE.MCECEBC.COMMON
   LOCAL      0%   OK  5249 PAGE.MCECEBC.LOCAL1
   LOCAL      0%   OK  524A PAGE.MCECEBC.LOCAL2
   LOCAL      0%   OK  5249 PAGE.MCECEBC.LOCAL3
   LOCAL      0%   OK  524A PAGE.MCECEBC.LOCAL4
   PAGEDEL COMMAND IS NOT ACTIVE
   ```

2. Check whether the Local Page data sets are on separate volumes. If not, create temporary Local Page data sets on different volumes (old storage subsystem).
3. Release all active Local Page data sets on the same volume using the PAGE DELETE command. It takes time to release them from the system. Use the DISPLAY ASM command to check on the progress.
4. When all Local Page data sets on the same volume are released from the system Page Pool, migrate that volume.
5. After migrating the volume, use the PAGE ADD command to add Local Page data sets to the System Page Pool again.
6. Repeat the proceeding steps for all volumes where Local Page data sets are located.

PLPA and COMMON Page data sets can be migrated by TDMF, but the volume where it resides must be the only volume in a migration session.

Migrating JES spool and check point data sets

JES2 spool and checkpoint data sets should be migrated separate from other application data volumes. Mark all volumes that contain these data sets in your volume migration list and migrate them in the same session. Migrate one volume at a time during a time when there is not a heavy I/O load on the system. Preferably stop JES activity to spool and checkpoint files during migration. For more information about systems running JES3, see the *TDMF Installation and Reference Guide*, TDM-Z55IR-001.

Migrating Couple data sets

Determine the volumes on which couple data sets reside by using the command **D XCF, COUPLE,TYPE=SYSPLEX** as shown in Example 10-8.

Example 10-8 Display sysplex couple data sets

```
COMMAND INPUT ===> /D XCF,COUPLE,TYPE=SYSPLEX                 SCROLL ===> CSR
RESPONSE=MCECEBC
 IXC358I  22.08.02  DISPLAY XCF 886
 SYSPLEX COUPLE DATA SETS
 PRIMARY    DSN: SYS1.CEBCPLEX.XCF.CDS01
            VOLSER: COUPL1     DEVN: B23D
            FORMAT TOD         MAXSYSTEM MAXGROUP(PEAK) MAXMEMBER(PEAK)
            12/09/2010 10:43:02       32     200   (39)      303    (8)
            ADDITIONAL INFORMATION:
             ALL TYPES OF COUPLE DATA SETS ARE SUPPORTED
             GRS STAR MODE IS SUPPORTED
             CLUSTER RESOURCE MANAGEMENT IS SUPPORTED
             SYSTEM STATUS DETECTION PROTOCOL IS SUPPORTED
 ALTERNATE  DSN: SYS1.CEBCPLEX.XCF.CDS02
            VOLSER: COUPL2     DEVN: B23E
            FORMAT TOD         MAXSYSTEM MAXGROUP       MAXMEMBER
            12/09/2010 10:46:44       32     200             303
            ADDITIONAL INFORMATION:
             ALL TYPES OF COUPLE DATA SETS ARE SUPPORTED
             GRS STAR MODE IS SUPPORTED
             CLUSTER RESOURCE MANAGEMENT IS SUPPORTED
             SYSTEM STATUS DETECTION PROTOCOL IS SUPPORTED
```

Check whether these data sets are on separate volumes.

Note: Do not try to migrate volumes with an active couple data set (CDS). A system wait state 0A2 might be caused when XCF is unable to update the active CDS during the TDMF *Quiesce* and *Synchronize phase*.

You also need to check the content of the couple data set volumes and ensure that they do not contain other data that requires special consideration when migrating with TDMF.

To migrate couple data sets, complete the following steps:

1. Switch the couple data set from primary to alternate:

   setxcf couple,type=sysplex,pswitch

   This also removes the former primary CDS from service.

2. Migrate the volume (COUPL1) with the alternate (now inactive old primary) couple data set. Make sure to run TDMF Agents on all systems that access the couple data set volumes.

3. Add the migrated alternate (old primary) to the system again as the new alternate couple data set:

   setxcf couple,type=sysplex,acouple=(SYS1.CEBCPLEX.XCF.CDS01,COUPL1)

4. Switch the couple data set again:

   setxcf couple,type=sysplex,pswitch

5. Migrate the volume (COUPL2) with the alternate (now inactive old alternate) couple data set.
6. Add the migrated alternate (old alternate) to the system as the new alternate again:

```
setxcf couple,type=sysplex,acouple=(SYS1.CEBCPLEX.XCF.CDS02,COUPL2)
```

Migration considerations

Consider the following tips when planning your migration:

- Consider using a naming convention for the TDMF Master and Agent jobs. If multiple sessions are required for the migration, use a naming convention that groups the Master, Agents, and COMMDS in an easy-to-find format. For more information, see "JCL naming convention" on page 242.
- The keyword **TIME=1440** or **TIME=(***mm***,***ss***)** can be specified on job cards to avoid system ABEND 322 (S322). **TIME=1439** is preferable because **TIME=1440** disables SMF time recording for that job. **TIME=(***mm***,***ss***)** allows recording of SMF times, but limits the amount of processor time used by the TDMF session. Ensure that a reasonable amount of time is allowed for each TDMF session, especially if the **COMPARE** option is used. See members MASTER and AGENT in the SAMPLIB for examples. For more information about the **TIME** parameter, see the MVS JCL manuals in the z/OS Internet Library at:

  http://www-03.ibm.com/systems/z/os/zos/library/bkserv/index.html
- You might consider migrating as many volumes as possible within a TDMF session. The overall real storage requirements for a specific number of migration volumes are identical for many volumes within a single (or few) session compared to only a few volumes using many sessions. By coding the maximum number of volumes per session, the management of multiple sessions is easier.

  Note: The current release of TDMF allows up to 512 volumes to be coded per session. However, the number of volumes multiplied by the number of systems (Master plus Agents) must not exceed 2048.

- Set *Active in Copy* and the *Number of Concurrent Volumes*[3] to the number of channels the system can access in the source subsystem. In a typical technology refresh, the target subsystem has newer storage devices and newer channels, such as FICON versus ESCON connections. Because ESCON channels are slower than FICON channels, limit the number of active volumes based on the slower channel paths. A general rule of thumb for ESCON channel paths is two active volumes per path.
- Employ *FastCopy* whenever possible. This option causes that only allocated tracks and cylinders are copied from source to target volumes. Empty tracks are not copied, and thus reduce the overall migration time.
- Use *Full Speed Copy* if you want to employ double buffering. When this option is selected, two 900-K buffers are allocated and pacing is automatically started. Overall time reduction depends on a number of factors, but an average of 45 percent can be expected.
- Select the *Dynamic ICKDSF REFVTOC* option. ICKDSF is dynamically called at the end of the migration if reformatting the VTOC is required.
- Select *Toleration of Invalid Count Fields* to ensure that volumes with invalid count fields (CCHH mismatch) are copied without error. These count fields are not corrected.

[3] Either as a system default by using the SYSOPTNS batch job or with the **CONCurrent(***nn***,ACTIVE)** session option.

JCL naming convention

When creating JCL for TDMF Master and Agents, establish a naming convention to avoid confusion during the migration process. The following is an example from a service project where one Master and six Agents were required per session. It was decided to perform the migration using three sessions. In total, 42 jobs (3 (sessions) * 7 (1 Master + 6 Agents) need to be created.

The naming convention used is shown in Example 10-9.

Example 10-9 Job naming convention (example)

```
//TDMLCP01 JOB ....
    !!! !
    !!! +--> session number (2 digits)
    !!+----> 2 char out of the LPAR name (first or last 2 char)
    !!       to identify AGENTs at different LPARs, ensure to
    !!       use unique jobnames
    !+-----> L=Local site, R=Remote site
    +------> M=Master, A=AGENT
```

Estimating how long it takes to move the data

How long it takes to migrate the data depends on a number of factors. Some of the factors are the options selected for a specific volume migration or for the entire session. Others are the activity of the processor, channel paths, and devices. This section addresses the following factors:

- Pacing
- Number of sessions
- Factors external to TDMF
- Channel paths
- Storage subsystem cache
- Device size and rotational speed
- Logical to physical mapping

Pacing

The term pacing with regard to TDMF specifies the number of tracks read/written in a single I/O operation. Standard I/O pacing is 15 tracks per I/O operation. TDMF samples device activity every 30 seconds during the *Copy* and *Refresh phases*. If device or channel path negatively activity affects the production environment, TDMF dynamically decreases the number of tracks read/written in a single I/O operation. You can also specify that reverse pacing is used, which is used when migrating high activity volumes. This option starts at one track per I/O operation and, if activity allows, increases the number of tracks per I/O dynamically.

User-specified pacing using the **MAXTracks** volume option allows you to determine the number of tracks read/written in a single I/O operation. Values of 1, 3, 5, or 15 can be specified for the number of tracks per I/O. This value might or might not be static depending on whether pacing is specified.

If an operating system is heavily used, the number of available real storage frames can be limited. TDMF tracks this value so that it does not affect production from that standpoint. If real storage pacing is started, I/O pacing is automatically affected.

To speed up the migration, you can select *Full Speed Copy* with the pacing volume option as shown in Example 10-10.

Example 10-10 Selecting Full Speed Copy with the pacing volume option

```
MIGRATE EM9010 EM9030 OPTIONS(PACING(FULLSPEED))
```

This option allows TDMF to use two buffers and can increase speed by around 45% on normal systems. However, additional real storage requirements and performance degradation for other users or applications using the same storage subsystem can advise against using this *Full Speed Copy* option.

With the *Number of concurrent volumes* and *Active in Copy* options, you can control the effect that the migration has on your production environment.

The number of concurrent running volumes can be changed dynamically by using the TDMF Monitor option 2 as shown in Figure 10-27. Select the *Number of concurrent volumes* field and change it, then press Enter to activate the change. If you reduce the number, the running volumes continue copying until the end of the volume. New volumes start when the number of volumes in the *Copy phase* drop below the new defined value.

```
 Session Status                 Row 1 to 4 of 4
Command ===>                                                Scroll ===> PAGE

System: Master               Softek TDMF for z/OS               Version 5.5.0
ComDataSet SYS1.TDMF.ITSO.LOCAL.SYSCOM01                        Sessions 0001
Number of volumes migrating 004              Number of concurrent volumes 002
Number of volumes  complete 000              Number of  volumes   waiting 002

Requested   Volume Device  Group     --- Migration --- - Error Info - Sync
 Action     Serial Number  Name      Status      Type  System Message Goal

   __       RS6000  6000             Copy        Swap A                005
            XX8100  8100

   __       RS6001  6001             Copy        Swap A                005
            XX8101  8101

   __       RS6002  6002             Inactive    Swap A                005
            XX8102  8102

   __       RS6003  6003             Inactive    Swap A                005
            XX8103  8103
```

Figure 10-27 Changing the number of concurrent volumes dynamically

You can manage a single volume pair from this TDMF monitor panel. Use one of the listed line commands[4] as shown in Figure 10-28.

```
 Transparent Data Migration Facility

                  PF3=End      PF7=Page Up       PF8=Page Down


               Following valid on all Systems                  Command
Display        Migration Messages                            : M
Display        Migration Performance Data                    : D
Display        Volume's Offline Volume Access (OVA) Job(s)   : X

               Following valid only on MASTER System           Command
Acknowledge    Mirror Change on Volume Swap                  : A
Approval       Approve this migration                        : B
Continue       Volume Migration or its Group                 : C or CG
Confirm        PPiT Recycle for Volume or its Group          : R or RG
Confirm        Start Volume Migration or its Group           : F or FG
Confirm        Synchronization for Volume or its Group       : P or PG
Disallow       OVA  registrations for Volume or its Group    : H or HG
Discontinue    PPiT migrations at "Recycle" prompt           : E or EG
Reinitialize   Migration of terminated Volume                : I
Set Goal       Synchronization Goal for Volume or its Group  : Z or ZG
Suspend        Volume Migration or its Group                 : S or SG
Terminate      Volume Migration or its Group                 : T or TG
               Terminate not allowed at APPROVAL PROMPT


PF4            Display the next active session
PF6            Toggle key for displaying the session options
PF10           Immediately display the session progress monitor
PF11           Toggle key for filtering volumes
```

Figure 10-28 TDMF line commands

Number of sessions

There is no limit to the number of TDMF sessions that can be started. The determining factor is the amount of available system resources.

Factors external to TDMF

When TDMF is moving data, it uses processor resources and peripheral resources such as channel paths and storage subsystem resources. These resources also determine how fast data is moved and the effect on other applications.

Channel paths

TDMF interacts with the Input/Output Supervisor, which manages transport mechanism and channel paths. Various types of channel paths are in use today, ranging from older technologies such as parallel paths to newer ones such as FICON. Each type of channel has a specific limit on the amount of data that can be transported. See the *IBM System z Connectivity Handbook*, SG24-5444 for throughput information:

http://www.redbooks.ibm.com/abstracts/sg245444.html

[4] The help panel is displayed when pressing PF1 with the cursor in a volume pairs selection field.

Storage subsystem cache

The size of cache in a storage subsystem has increased over the years. However, when a migration or migration service is being performed, it frequently involves a technology refresh. Older technology has smaller cache sizes compared to that of new technology. Although TDMF issues the channel command **`Inhibit Cache Load`** during the read operation, some storage subsystems do not honor this command. Therefore, while cache is being used by the production environment, there is the additional load of the data being moved by TDMF.

Device size and rotational speed

Over the years, the size of devices has increased along with the rotational speed of the devices. However, you often migrate off older technology, which is generally slower. Know the source subsystem and all the intricacies that go with it.

Logical to physical mapping

In an RAID-1 environment, multiple logical devices are mapped to a single physical device. In an RAID-5 environment, multiple logical devices are mapped to multiple physical devices. Although these new devices are much faster than older technology, there is still only one head-disk assembly per physical device. Moving too many logical devices at once can have a negative impact on a production environment. However, if solid-state disks (SSDs) are installed, the internal interfaces can be a bottleneck instead of the disks themselves.

In summary

The best migration practice is to plan properly. Understanding the production environment and when the heaviest loads occur is a part of planning, along with the effect of old and new technology. Remember that data can be moved only as fast as the slowest device or slowest channel allows.

Terminating a TDMF session

With the TDMF Monitor or TDMF Batch Monitor, you can terminate a specific volume pairing, volume groups, or all volumes in a session dynamically. A termination request for a volume pairing might take up to a minute to be processed by TDMF. If the Master and all the Agent systems are active in the same IBM Parallel Sysplex, the interval is less.

Under extreme conditions, you can use the MVS **`CANCEL`** command. If the Master system fails leaving an Agent system active, allow a 15-minute interval to allow the Agent system to shut down automatically. If it does not, cancel the Agent system job manually.

If a TDMF session still hangs after issuing an MVS **`CANCEL`** command, use **`Display GRS,C`** to determine whether canceling another job might break the deadlock. When a TDMF job is canceled, the program signals the other systems in the session to terminate the active volumes and clean up resources. Canceling for a second time, or using the MVS **`FORCE`** command, bypasses this recovery and causes the TDMF RTM Resource Manager routine to take control. Canceling this way might result in migration volumes being varied offline and boxed[5] as part of cleanup. If you issue a **`FORCE`** or second **`CANCEL`** command, run the TDMFCLUP program as described in “Batch Utilities” of the *TDMF Installation and Reference Guide*, TDM-Z55IR-001.

> **Important:** During an error, avoid issuing a CANCEL command against TDMF Master or Agents. Instead, use the TDMF Monitor option 2 (interaction with TDMF) to terminate a migration session.

[5] And possibly left in a SPID Fence or Soft Fence state.

10.3 Using TDMF z/OS in a remote environment

This section describes the remote replication using TDMF z/OS product. For remote replication, TDMF uses TCP/IP connections between local source and remote target environments. Therefore, to use the product in a typical remote migration scenario, there are additional steps in relation to a TDMF migration in a local shared storage environment.

This section contains the following topics:

- Overview of Remote Replication with TDMF z/OS
- General guidelines for using TDMF in a remote environment
- Example Remote Replication scenario with TDMF z/OS

10.3.1 Overview of Remote Replication with TDMF z/OS

TDMF z/OS uses TCP/IP to create Point-in-Time copies of local source volumes on remote target volumes. It is a host-based software solution for data replication. It allows migration in a z/OS environment to new locations over long distances. TDMF z/OS is user-initiated and controlled, and ensures full access to the local data at any point during a replication operation. However, it does not redirect I/O to the target devices as is possible in a local shared storage environment.

Long-distance data migration software solution

TDMF z/OS has the following advantages:

- It is hardware independent, and therefore especially useful when migrating from non-IBM source volumes.
- It can be used when other migration technologies like Remote Mirror/Copy, z/OS Global Mirror are not applicable.
- It is easy to manage, and no additional automation or management tools are needed.
- It is easy to install, with no additional hardware or software that needs to be installed.

Long-distance data migration characteristics

The TCP/IP replication function of TDMF z/OS has the following characteristics:

- It allows long-distance data migration. Data is copied from source devices to target devices over any available TCP/IP connection.
- A replication is the logical relationship between a source and a target device.
- The user initiates and controls all migrations. The user identifies the “from” (source) volumes and the “to” (target) volumes.
- Multiple volume replication can be established within a single session.
- Multiple sessions can be run in parallel.
- The replication tool is dynamically activated and manageable.
- Applications are unaware that a replication is performed. The local data is continuously and fully accessible for read and write activity.
- For synchronization, all applications that access the source volumes must be stopped at a certain point to reach a data consistency state at the remote site.
- TDMF supports a multiple-system shared data environment.
- TDMF ensures complete physical data integrity.

- The use of TDMF is not restricted to any control unit model type or device type. Except as noted, all devices in the data center can participate in replication sessions as required.
- A source volume cannot contain an active local page data set.
- The source and target volumes must be of the same track geometry.

The proceeding characteristics represent the ideal of a nondisruptive data migration.

Restriction: TDMF z/OS does not allow access to the target volume during the migration process.

Remote Replication Terminology

In addition to the terms already listed in 10.1.2, “Terminology” on page 212 these additional terms are used with regard to a remote replication using TDMF over TCP/IP:

Master system (local) The local z/OS system running a TDMF Master (a z/OS batch job or started task) that is responsible for the data copy function. Unlike a TDMF migration in a shared local storage environment, two Master systems are involved in a remote TDMF replication session: A local and a remote Master system.

Agent system (local) A local z/OS system in a shared storage environment running a TDMF Agent. To ensure data integrity, any system besides the local Master system that accesses any source volume must run a TDMF Agent. The TDMF Master system and associated Agent systems communicate through a shared system COMMDS.

Master system (remote) The remote z/OS system running a TDMF Master that is responsible for receiving and storing the data at the remote site.

Agent systems (remote) A remote z/OS system running a TDMF Agent. Agent systems are required at the remote site if target volumes are shared. Running TDMF Agents prevent access to the remote target volumes while a replication process is active.

Migration session The local and remote TDMF Masters plus any local and remote TDMF Agents and the volume pairs and groups to be processed in a single TDMF execution.

Data consistency The procedures to ensure that data remains intact and correct throughout the migration. Applications accessing local source volumes must be shut down to allow TDMF synchronization and to reach a state of data consistency.

Synchronization When I/O is inhibited to the source devices (*Quiesce phase*) and access to all source volumes is prevented (all applications stopped). The last set of source volume updates are collected and reflected to the target devices (*synchronization phase*). Volume serial numbers of target volumes are *not* changed during this process. They can be changed separately if required. When this process completes, the TDMF session ends and the target device is taken offline (depending on the replication parameter).

TDMF z/OS architecture in a remote replication process

Figure 10-29 depicts the TDMF z/OS architecture and components in a remote replication process. In this example, the volumes are being moved from a local to a remote site. All storage vendors are supported (any to any).

The following components are depicted:

- TDMF Monitor: The TDMF Monitor is used to manage the active replication jobs and view active or past replications. The TDMF Monitor component consists of REXX execs, which require ISPF version 3.6 and TSO/E Version 2.4 at a minimum.
- Communications data set (local / remote): The local COMMDS is used to pass information between local systems participating in a TDMF TCP/IP session (Figure 10-29). This data set contains the status and messages related to a specific session. The COMMDS also serves as the input file to the TDMF Monitor. A COMMDS is also needed at the remote site, and needs to have the same size (number of cylinders).

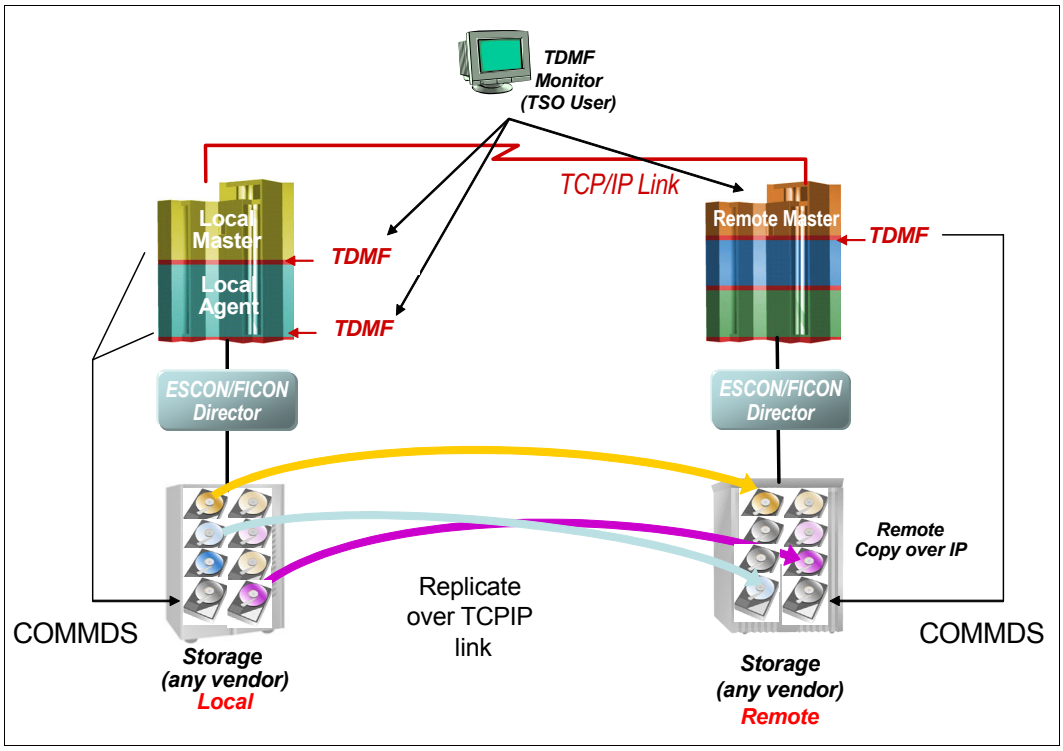

Figure 10-29 TDMF TCP/IP Basic overview

- Master and Agents: Each system that has access to the local source or remote target volumes must run a TDMF Agent. Each side (that is, local and remote) of the TDMF Master/Agent session must have access to its own communication data set.

Important: To avoid a possible data integrity problems, all systems that access replicated volumes must be identified to the Master system. TDMF includes various controls and checks that ensure that the user does not make the following errors:

- Assign or direct conflicting migrations to the same devices
- Attempt migrations to nonexistent devices
- Use the same COMMDS for two simultaneous or overlapping migration sessions

Master system responsibilities

The local TDMF Master system is responsible for the following tasks:

- Initialize the TDMF Master system environment and the COMMDS.
- Start and control each session for all participating systems.
- Check the communication with the remote site (remote Master).
- Monitor source volume I/O activity to detect updates.
- Copy data from the source volume to the target volume through TCP/IP.
- Process detected source volume updates from all local systems.
- Perform refresh operations to the target site to reflect the update activity on the source volumes.
- Internal health checking.
- The remote TDMF Master system receives data from the local site and writes those data to the target volumes. It also monitors target volume system I/O activity to prevent updates from applications to the target volumes.

Agent system responsibilities

The TDMF Agent system is responsible for the following items:

- Initialize the TDMF Agent environment and establish communications to the Master system through the COMMDS.
- Acknowledge and process replication requests (volumes to be copied) from the Master system.
- Monitor source volume I/O activity and detect updates.
- Notify the TDMF Master system about source volume update activity through the COMMDS.
- Internal health checking.
- Monitor target volume I/O activity to prevent updates from applications to the target volumes.

TDMF z/OS hardware compatibility

TDMF is designed to support CKD/E-capable control units of any vendor as a source or a target storage subsystem. This functionality includes Hyper-Volume, Flexible-Volumes, and all user-defined volume sizes. Extended address volumes (EAVs) might be supported depending on the z/OS version. For more information, see the *TDMF Installation and Reference Guide*, TDM-Z55IR-001.

Because TDMF is a host software migration tool, there are no particular hardware prerequisites.

Install TDMF z/OS

For the installation procedure, see “Installation” on page 218.

You need to install TDMF at the remote site (a license key to run the TCP/IP replication function of TDMF z/OS on the remote site is not required).

TDMF z/OS customer requirements

See “Customer requirements” on page 218 for any customer requirements.

10.3.2 General guidelines for using TDMF in a remote environment

The following sections provide general guidelines when using TDMF z/OS to replicate local source volumes to remote volumes:

- Keeping current
- Preparation and planning
- Creating a consistent data structure
- Summary of replication tasks and steps
- Link performance information

Keeping current

It is important to make sure before any migration that your TDMF z/OS software is current and up-to-date:

- Check whether there are any new PTFs available and update the TDMF software as needed
- Manuals are periodically updated. To ensure that you have the most current information, check for new TIBs, which are typically issued between manual releases.

Preparation and planning

Make sure that you cover the following considerations when planning:

- Make sure that your TSO user ID has authority to run TDMF and ICKDSF.
- Create a volume list. Each target volume should be of the same or larger size than its corresponding source volume. Make sure that you do not have duplicates or overlaps (volume serial number, source volume unit address, target volume unit address) in your list. Mark your special consideration volumes in the list (IPL, IODF, XCF DS, local page data sets, PLPA and COMMON page data set, and so on) so they are easy to find at the remote site.
- Group volumes in relation to application data (storage groups) into the same TDMF migration session. You might also need to group volumes by control units for performance reasons. Performance problems might be caused by the ESCON connection or HW with less internal bandwidth. Make sure that the same volume is not defined to multiple groups.
- Consult your network support group to reserve a range of port numbers that can be used for replication. The number of ports that are needed depends on the total number of migration sessions. However, remember that one TDMF session needs two adjacent port numbers. Make sure that ports are not blocked by firewalls between the two sites. Also, review the BPXPRMxx parmlib member for parameters that affect a TDMF migration session over TCP/IP. If your installation uses multiple TCP/IP stacks, the TDMF migration jobs require a SYSTCPD DD statement allocating the appropriate profile data set.
- Reserve a number of volumes for local COMMDSs. Allocate no more than 3-4 COMMDSs on each volume to prevent performance issues when sessions are run in parallel. The total number of COMMDSs must be the same as the number of sessions. Allocate the same number of COMMDSs at the remote site.
- Define available and active initiators (one initiator for each session) for both sites. The service ratio should be SYSSTC or similar for all TDMF jobs
- Create temporary local page data sets if needed and do a page delete to the original local page data sets. TDMF does not copy volumes containing active local page data sets.

- Establish naming conventions for the job names to enable concurrent execution of many migration jobs. Naming conventions can help to avoid confusion. For more information, see “JCL naming convention” on page 242.
- Schedule the migration for a time when heavy system I/O load such as database reorganization and production batch processing can be suspended.

Normally replication of storage volumes is combined with a tape/library system replication. It is also sometimes done in combination with open storage systems and servers. Therefore, it is important to consolidate the TDMF z/OS replication project with the migration team.

Creating a consistent data structure

To move an entire data center, you must have a consistent data structure. This section highlights the components of a consistent data structure. At the local site, you have one or more processors, and one or more storage systems. At the remote site, you have a minimum of one z/OS system to receive the data.

The data structure has the following components. Figure 10-30 on page 252 shows how these components relate to each other.

- Local system: The local system consists of all volumes in relation to the system. Examples include IPL, IODF, local page data sets, and system-related volumes.
- TDMF volumes: These volumes are TDMF libraries and COMMDS volumes. These volumes are not moved.
- Temporary local page data sets: These data sets are on volumes that are not replicated. TDMF cannot copy active local page data sets, so the original data sets are removed from the page pool by a page delete command. They can then be replicated.
- Data volumes: This term includes all volumes, including SMS-managed and non-SMS managed volumes with applications, application data, and temporary work data sets.
- Remote system: This system includes all volumes at the remote site in relation to the remote Master system. This system receives the data from the local TDMF Master system over TCP/IP and writes the data to the corresponding target volumes.
- TDMF Monitor: This TDMF ISPF application is run by a TSO user. It is useful to run the TDMF Monitor function at the local and remote sites of the replication process.
- TCP/IP Link: This link is the physical connection from the local to the remote site using TCP/IP as the protocol. The normal bandwidth for a service engagement (2-12 TB) is a 1 GBit link speed. With this connection, you can expect up to120 MBps at the link. The throughput of the TCP/IP links is a major factor and must be considered when estimating the size of the overall replication time frame.

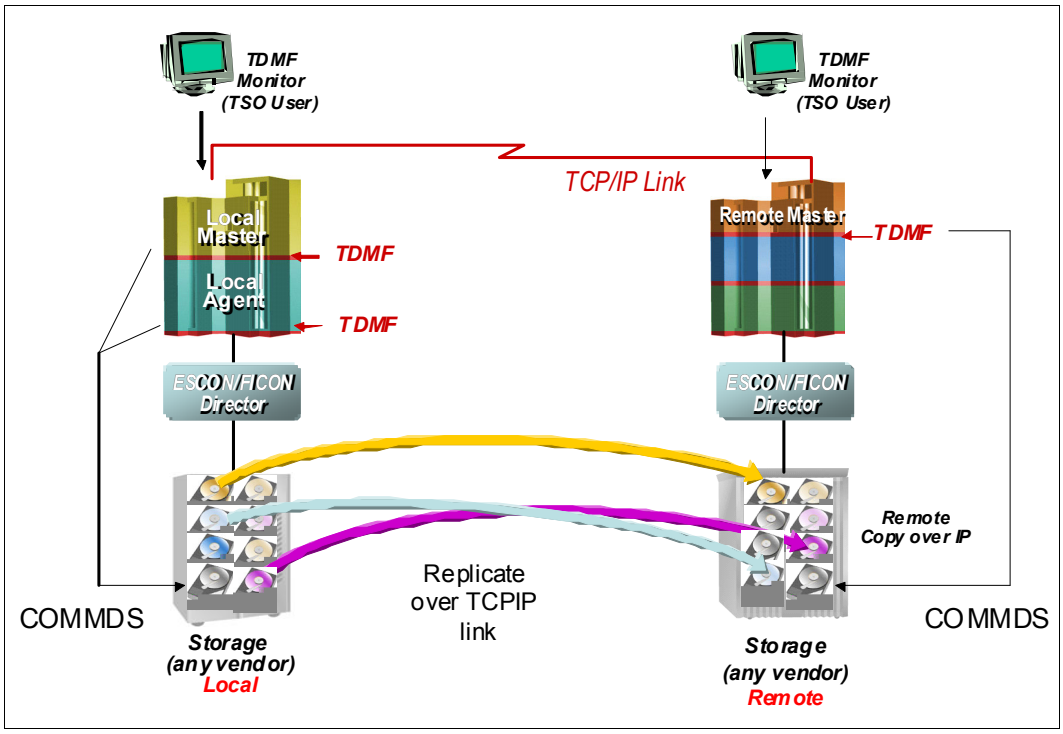

Figure 10-30 TDMF overview for a data center move to a new location over TCP/IP

Summary of replication tasks and steps

Before you begin the replication, complete the following steps:

- Do system assurance for both sites, including the TCP/IP connection. For more information, see the *TDMF Installation and Reference Guide*, TDM-Z55IR-001.
- Check that all target volumes are available and online on the target system (new storage).
- Install TDMF on both sites.
- Define the volume pairs (source/targets) with the aid of the storage administrator.
- Create TDMF JCL members for both sites.
- Create ICKDSF INIT/Rename Jobs for the remote site.
- If necessary, prepare temporary local page data sets on volumes that are not part of the replication. These data sets are necessary because active local pages cannot be handled by TDMF.
- Perform an initial network test to check the link, connection, and bandwidth.

The replication is normally done in three stages. Along with the storage administrator, make sure the volume list (relation between source and target volumes) is up to date before each stage.

1. Perform a test without actually synchronizing the data as shown in Figure 10-31. This test is normally done with system/network volumes only, and is also a network test. Application data are copied to test the bandwidth.

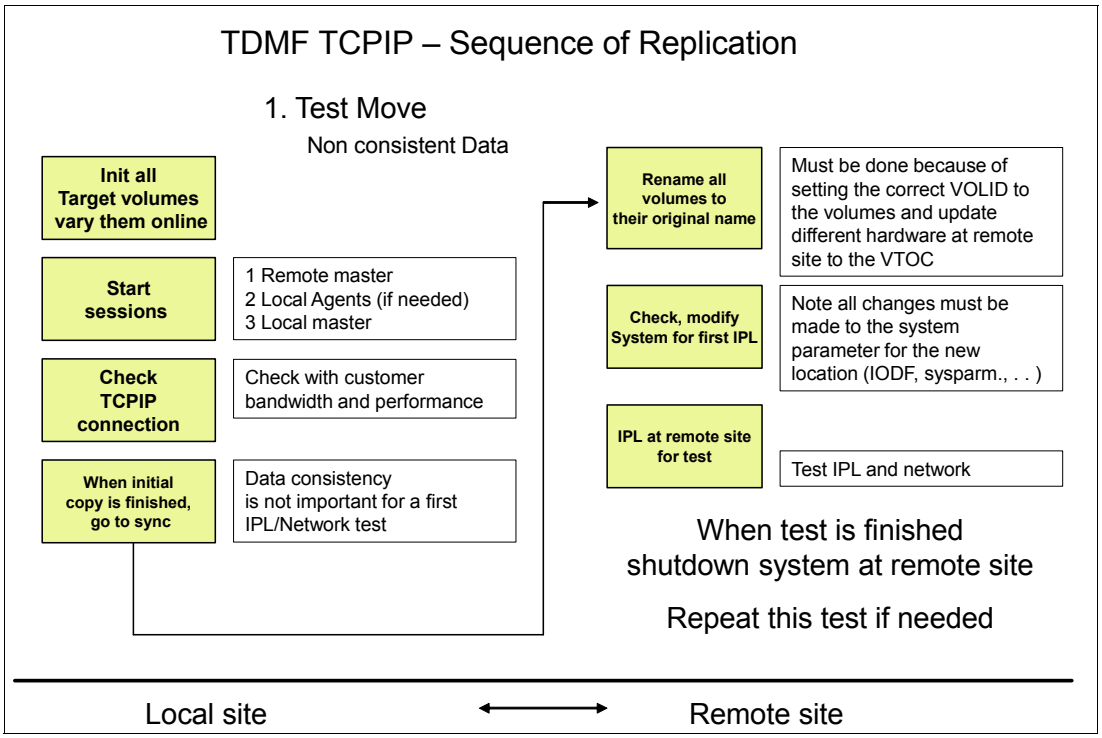

Figure 10-31 Test without sync data

2. Run a test with a complete data synchronization step to test the system, network, and application data at the remote site (Figure 10-32).

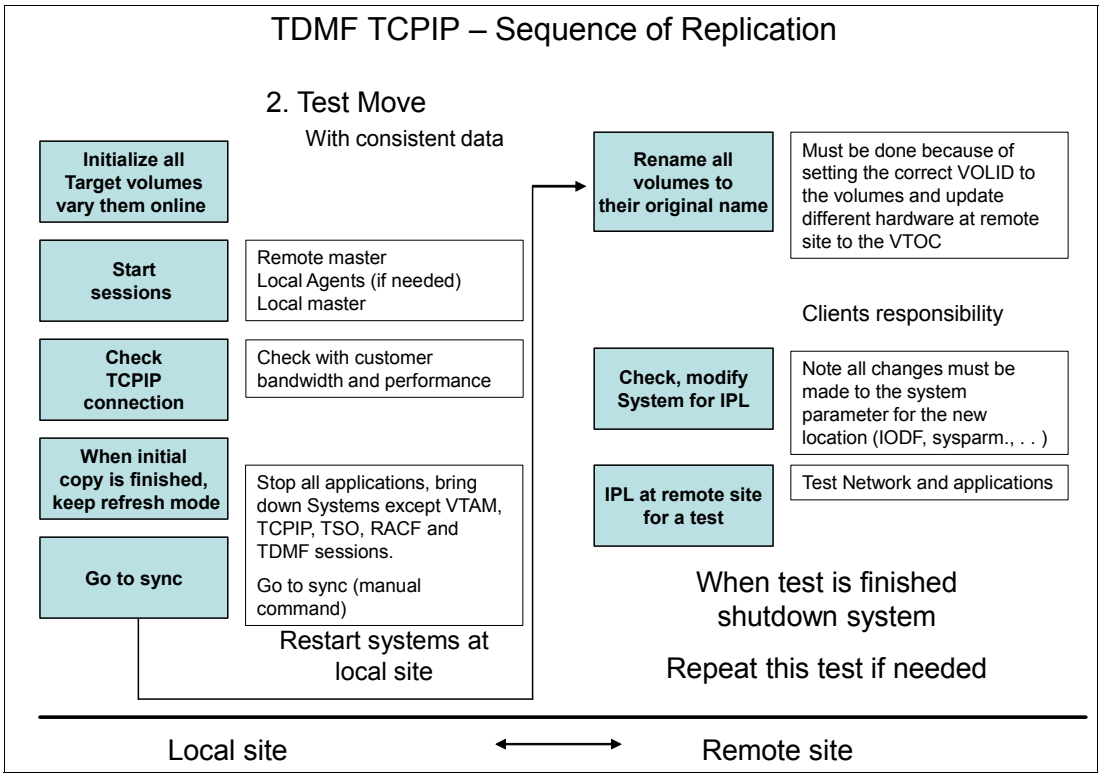

Figure 10-32 Test with synchronization data, Production Test Environment (PTE)

3. After completing the tests, run the actual production move (Figure 10-33).

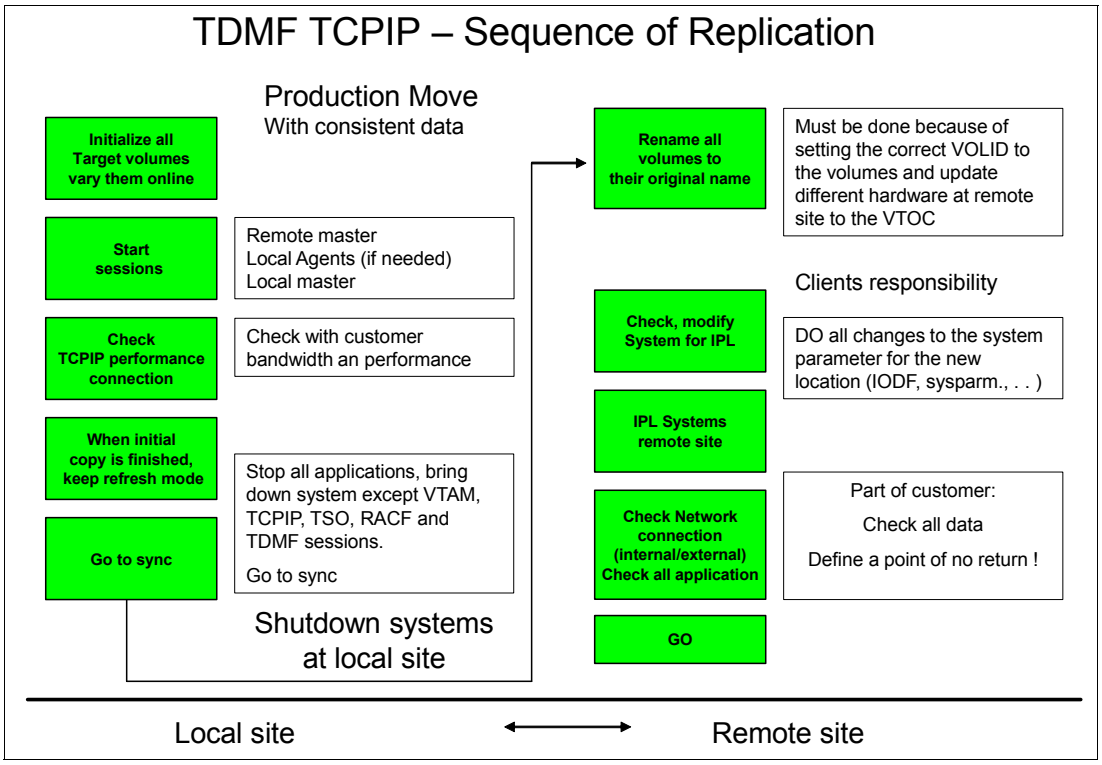

Figure 10-33 Production move, sequence of replication

Link performance information

Figure 10-34 shows the performance found in test environments for various connection types, and the expected throughput for a TDMF replication over TCP/IP.

Connection Type	Bytes Per Minute	MB Per Minute	MB per Hour	GB per Hour	TDMF Speed GB per Hour
384K DSL	2,949,120	2.8125	168.75	.1647	.1318
T1	11,520,000	10.98	658.18	.6437	.5149
10 base T Ethernet	75,000,000	71.52	4,291.53	4.19	3.35
T3	345,600,000	329.589	19,775	19.31	15.45
100 Base T Ethernet	750,000,000	715.255	42,915	41.90	33.52
OC3	1,162,500,000	1,108.64	66,518	65.95	51.96
OC12	4,500,000,000	4,291.53	257,492	251.45	201.16
Gigabit Ethernet	7,500,000,000	7,152.55	429,153	419.09	335.27
OC48	18,000,000,000	17,166.13	1,029,968	1,005.82	804.66
OC192	72,000,000,000	68,664.52	4,119,872	4,023.28	3,218.64

Figure 10-34 Link speed examples

Figure 10-35 shows the impact of TDMF software compression. At a TCP/IP link speed around 120 MBps, the FICON channel throughput rate increased up to 160 MBps. However, the CPU load also increased significantly. In the customer scenario, software compression cannot be used during online business because of the high processor load.

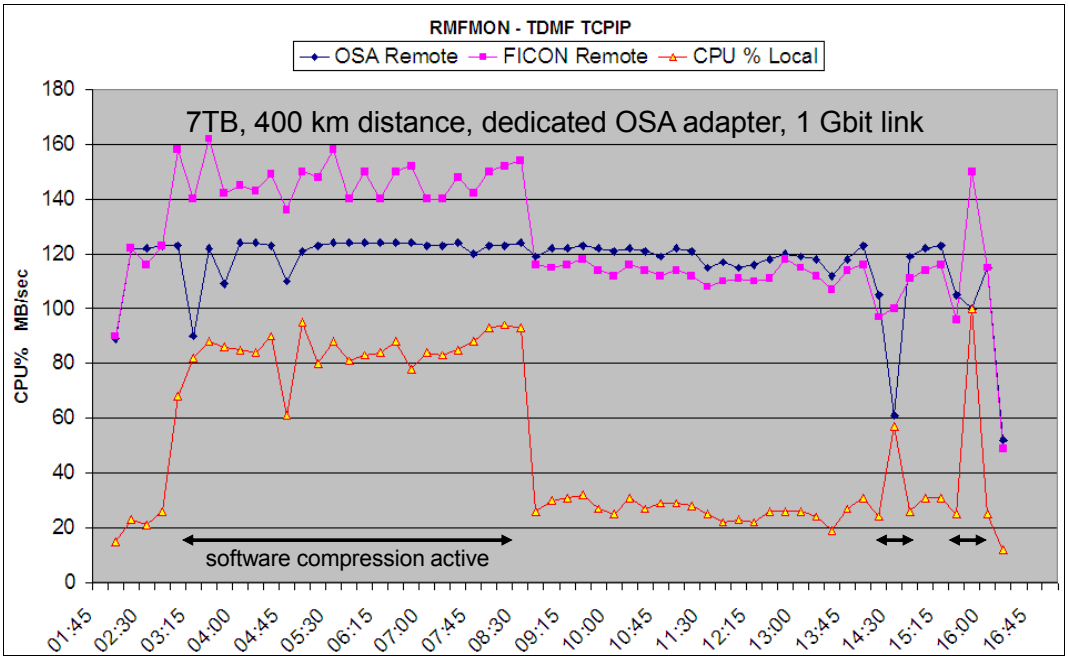

Figure 10-35 Example of a performance information client production move

10.3.3 Example Remote Replication scenario with TDMF z/OS

This scenario described here demonstrates the required TDMF, ICKDSF rename, and ICKDSF initialization jobs. The network test is included with screen captures of the JCL, the TDMF Monitor, and some performance information.

The distance in this scenario was around 100 km between the local and the remote site. The client wanted to test the connectivity and bandwidth for the two new defined 1GBit links. To retain the performance information, the test was run twice and two sets of COMMDSs were used. Sessions 01 - 03 were in the first test, and sessions 04 - 06 were in the second.

Both tests used the same 30 volumes (10 out of each DISK Storage unit). Therefore, after the first test the RENAME and INIT jobs were run, and all 30 volumes were set ONLINE again. All devices with 2xxx and 3xxx addresses were attached with ESCON Channel paths. The rest were FICON attached. Example 10-11 shows the volumes that were used.

Example 10-11 Volume list for test

```
Source volumes at source site
        UC2010 - UC2019         (ESCON connection)
        UC3010 - UC3019         (ESCON connection)
        UC4010 - UC4019         (FICON connection)
Target volumes at remote site
        UC8400 – UC8409         (FICON connection)
        UC8500 – UC8509         (FICON connection)
        UC8600 - UC8609         (FICON connection)
```

The TDMF Master at the local site was run on system SYCP (local TDMF Master). TDMF Agents were run on all other local z/OS systems. The TDMF Master at the remote site was run on system SYX1 (remote TDMF Master). No TDMF Agents were required at the remote site because the target volumes were not shared with other remote z/OS images.

Overview of a remote migration/replication

Figure 10-36 shows a simplified view of the physical setup in the described client scenario. The Channel Path IDs (CHPIDs) are listed to check the performance using the z/OS performance monitor (RMFMON).

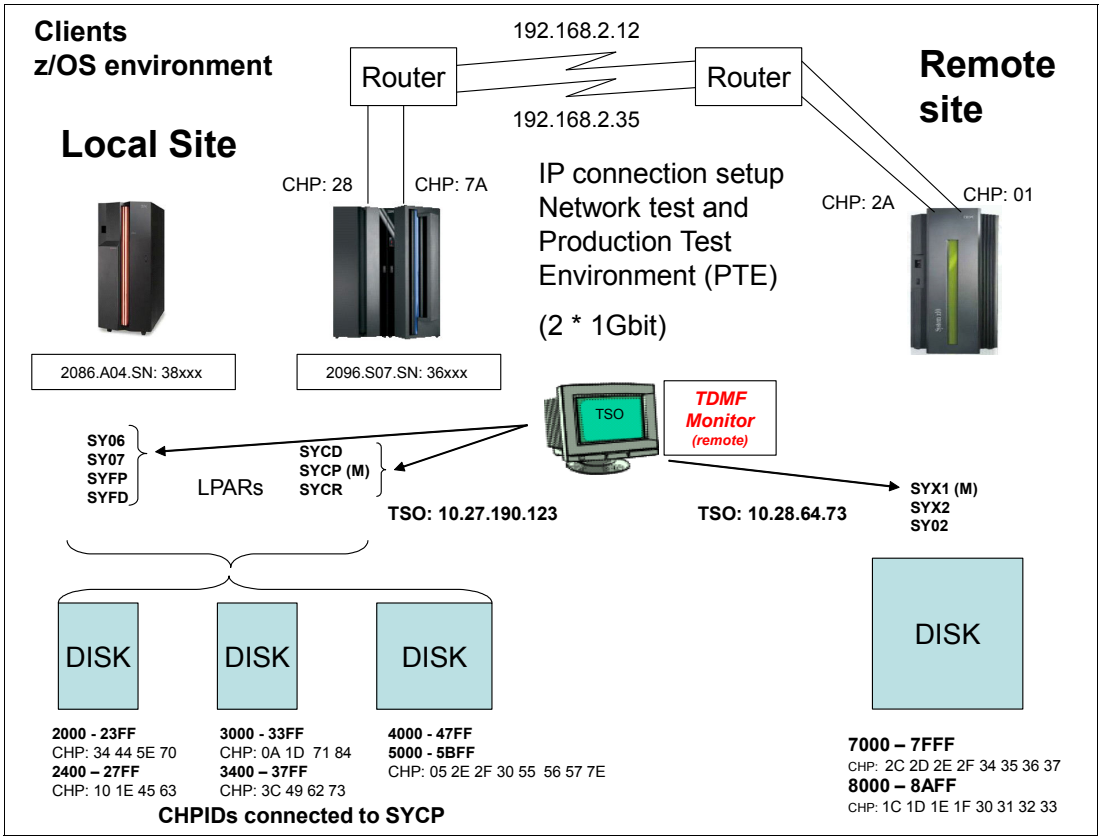

Figure 10-36 TDMF replication over TCP/IP overview

JCL at the local (source) site

The job names used at the local site are shown in Figure 10-37.

```
EDIT              STXXXXX.TDMF.NWTS1.TDMF.CNTL              Row 00001 of 00044
Command ===>                                                  Scroll ===> CSR
           Name     Prompt       Size   Created          Changed          ID
_________ $OVR00                           --> Overview
_________ TDALCD01       Agents for SYCD
_________ TDALCD02                    "
_________ TDALCD03                    "
_________ TDALCD04                    "
_________ TDALCD05                    "
_________ TDALCD06                    "
_________ TDALCR01       Agents for SYCR
_________ TDALCR02                    "
_________ TDALCR03                    "
_________ TDALCR04                    "
_________ TDALCR05                    "
_________ TDALCR06                    "
_________ TDALFD01       Agents for SYFD
_________ TDALFD02                    "
_________ TDALFD03                    "
_________ TDALFD04                    "
_________ TDALFD05                    "
_________ TDALFD06                    "
_________ TDALFP01       Agents for SYFD
_________ TDALFP02                    "
_________ TDALFP03                    "
_________ TDALFP04                    "
_________ TDALFP05                    "
_________ TDALFP06                    "
_________ TDAL0601       Agents for SY06
_________ TDAL0602                    "
_________ TDAL0603                    "
_________ TDAL0604                    "
_________ TDAL0605                    "
_________ TDAL0606                    "
_________ TDAL0701       Agents for SY07
_________ TDAL0702                    "
_________ TDAL0703                    "
_________ TDAL0704                    "
_________ TDAL0705                    "
_________ TDAL0706                    "
_________ TDCL$G00                              <-- create COMMDS
_________ TDMLCP01       Master for SYCP
_________ TDMLCP02                    "
_________ TDMLCP03                    "
_________ TDMLCP04                    "
_________ TDMLCP05                    "
_________ TDMLCP06                    "
          **End**
```

Figure 10-37 Job names used in replication (local site)

For a long-distance data migration using TDMF over TCP/IP, the IP address of the remote TDMF Master system and the assigned port number are specified using the REMOTE control statement as shown in Figure 10-38 (line 20).

```
EDIT       MBSM.TDMF.NWTS1.TDMF.CNTL(TDMLCP01) - 01.00     Columns 00001 00072
Command ===>                                                  Scroll ===> CSR
****** **************************** Top of Data ****************************
000001 //TDMLCP01 JOB (,DASD),'IBM DATA MIGRATION',CLASS=T,
000002 //      REGION=0M,MSGCLASS=X,NOTIFY=&SYSUID
000003 //* TDMF DATA MIGRATION - DO NOT CANCEL THIS JOB !!!
000004 /*XEQ TXRPRD
000005 /*JOBPARM SYSAFF=SYCP
000006 //*****************************************************
000007 //* Master JOB TDMF running: SYCP local site
000008 //*****************************************************
000009 //LOCAL01  EXEC PGM=GTDMAIN,PARM=MASTER,REGION=0M
000010 //STEPLIB  DD DISP=SHR,DSN=SYS1.GTD550.GTDLLIB
000011 //GTDKEY   DD DISP=SHR,DSN=SYS1.GTD550.GTDLLIB
000012 //SYSCOM   DD DISP=SHR,DSN=SYS1.TDMF.SDNT.LOCAL.SYSCOM01
000013 //SYSPRINT DD SYSOUT=*
000014 //SYSSNAP  DD SYSOUT=*
000015 //SYSIN    DD *
000016 SESSION M(SYCP)
000017 AGENTS( SY06 SY07 SYFP SYFD SYCD SYCR)
000018 OPT(FAST UNIDENT(W) PROMPT CONC(4 ACTIVE)
000019 NOCOMPRESS PAC CHECKT OFFLINET SINGLE)
000020 REMOTE SYX1 ADDR(192.168.2.12) PORT(8200)
000021 REPLICATE UC2010 UC8400  OPT(MAXTR(1))
000022 REPLICATE UC2011 UC8401  OPT(MAXTR(1))
000023 REPLICATE UC2012 UC8402  OPT(MAXTR(1))
000024 REPLICATE UC2013 UC8403  OPT(MAXTR(1))
000025 REPLICATE UC2014 UC8404  OPT(MAXTR(1))
000026 REPLICATE UC2015 UC8405  OPT(MAXTR(1))
000027 REPLICATE UC2016 UC8406  OPT(MAXTR(1))
000028 REPLICATE UC2017 UC8407  OPT(MAXTR(1))
000029 REPLICATE UC2018 UC8408  OPT(MAXTR(1))
000030 REPLICATE UC2019 UC8409  OPT(MAXTR(1))
****** ************************** Bottom of Data **************************
```

Figure 10-38 TDMF Master JCL at the local site (TDMLCP01)

The following other TDMF session options are used in this TDMF Master JCL:

FAST Copy only space allocated to data sets.

UNIDENT(W) Issue a warning message in case systems are detected without a TDMF Agent that are accessing any source or target volume. For anything other than testing purposes, use UNIDENT(T) to terminate the session if such an agentless system is detected[6].

PROMPT Suspends the migration in the *Refresh phase* and waits for confirmation before continuing into the *Synchronize phase*.

CONC(4 ACTIVE) Copy up to four volumes concurrently. This value can be changed dynamically during the migration by using the TDMF Monitor.

[6] Any system without an active TDMF Agent accessing any of the source or target volumes during a migration might cause data integrity issues.

NOCOMPRESS	Do not use software compression.
PAC	Use pacing to automatically reduce the number of tracks per I/O in case the migration activity degrades overall system performance.
CHECKT	Stop the migration if data sets are allocated on target volumes. This setting ensures that only empty volumes are accepted as target volumes.
OFFLINET	Take target volumes offline after migration completes.
SINGLE	Combines all volumes into a single session (with an ID of SINGLE) that you can end (go to sync) using a manual command.
MAXTRACKS(15)	This parameter adjusts the number of tracks that are read per I/O and the memory usage in the common area. If you run out of memory, use MAXTR(1) to reduce the performance impact. This parameter is also useful for old devices attached using ESCON.

For more information about the parameters, see the *TDMF Installation and Reference Guide*, TDM-Z55IR-001.

Figure 10-39 shows one of the local Agent jobs. The COMMDS specified using the SYSCOM DD statement must be the same as the one specified in the corresponding local Master JCL for this TDMF session.

```
EDIT       MBSM.TDMF.NWTS1.TDMF.CNTL(TDALCD01) - 01.00     Columns 00001 00072
Command ===>                                                  Scroll ===> CSR
****** ***************************** Top of Data ******************************
000001 //TDALCD01 JOB (,DASD),'IBM DATA MIGRATION',CLASS=T,
000002 //      REGION=0M,MSGCLASS=X,NOTIFY=&SYSUID
000003 //* TDMF DATA MIGRATION - DO NOT CANCEL THIS JOB !!!
000004 /*XEQ TORCSD
000005 /*JOBPARM SYSAFF=SYCD
000006 //*****************************************************
000007 //* AGENT JOB TDMF running: SYCD
000008 //*****************************************************
000009 //STEP01   EXEC PGM=GTDMAIN,PARM=AGENT,REGION=0M
000010 //STEPLIB  DD DISP=SHR,DSN=SYS1.GTD550.GTDLLIB
000011 //GTDKEY   DD DISP=SHR,DSN=SYS1.GTD550.GTDLLIB
000012 //SYSCOM   DD DISP=SHR,DSN=SYS1.TDMF.SDNT.LOCAL.SYSCOM01
000013 //SYSPRINT DD SYSOUT=*
000014 //SYSUDUMP DD SYSOUT=*
000015 //SYSSNAP  DD SYSOUT=*
****** *************************** Bottom of Data **************************
```

Figure 10-39 TDMF Agent JCL local site (TDALCD01)

JCL at the remote (target) site

The remote site requires JCL for the remote TDMF Master sessions. For each local TDMF Master session, a remote counterpart is required. Because target volumes are not shared at the remote site, TDMF Agents are not required.

The job names used at the local site are shown in Figure 10-40.

```
EDIT              MBSM.TDMF.NWTS1.TDMF.CNTL                 Row 00001 of 00009
Command ===>                                                  Scroll ===> CSR
           Name     Prompt       Size   Created          Changed          ID
_________ $OVR00                                    <-- Overview
_________ TDCR$G00                                  <-- create COMMDS
_________ TDMRX101       Master for SYX1
_________ TDMRX102                    "
_________ TDMRX103                    "
_________ TDMRX104                    "
_________ TDMRX105                    "
_________ TDMRX106                    "
_________ VDASIN01                                  <-- device list (UCBs)
          **End**
```

Figure 10-40 TDMF Master jobs (remote site)

Tip: The port numbers specified in the local TDMF Master and their corresponding remote TDMF Master JCL must match. The port number defines a pair of two consecutive ports. The first port number is used to establish a control session between the local and remote TDMF Master systems. The subsequent second port number is used to establish a separate TCP session for each source/target volume pair. See Figure 10-54 on page 268.

Figure 10-41 shows the remote TDMF Master JCL. The TCP port number is specified with the **PORT** parameter on the **EXEC** job control statement.

```
EDIT       MBSM.TDMF.NWTS1.TDMF.CNTL(TDMR0601) - 01.00     Columns 00001 00072
Command ===>                                                  Scroll ===> CSR
****** ***************************** Top of Data ******************************
000001 //TDMRX101 JOB (,DASD),'IBM DATA MIGRATION',CLASS=T,
000002 //      REGION=0M,MSGCLASS=X,NOTIFY=&SYSUID
000003 //* TDMF DATA MIGRATION - DO NOT CANCEL THIS JOB !!!
000004 //*XEQ SYX1
000005 /*JOBPARM SYSAFF=SYX1
000006 //*******************************************************
000007 //* MASTER JOB TDMF  SYX1 remote site
000008 //*******************************************************
000009 //STEP01   EXEC PGM=GTDMAIN,PARM='MASTER,PORT=8200'
000010 //STEPLIB  DD DISP=SHR,DSN=SYS1.GTD550.GTDLLIB
000011 //GTDKEY   DD DISP=SHR,DSN=SYS1.GTD550.GTDLLIB
000012 //SYSCOM   DD DISP=SHR,DSN=SYS1.TDMF.SDNT.REMOTE.SYSCOM01
000013 //SYSPRINT DD SYSOUT=*
000014 //SYSUDUMP DD SYSOUT=*
000015 //SYSSNAP  DD SYSOUT=*
000016 //SYSIN    DD *
000017 SESSION M(SYX1)
000018 OPT(FAST UNIDENT(T))
****** **************************** Bottom of Data ****************************
```

Figure 10-41 TDMF Master JCL remote site

4. After the TDMF replication completes, you can rename the target volumes to their original volume serial numbers using ICKDSF REFORMAT. Figure 10-42 shows the JCL used in the experimental environment.

```
EDIT       MBSM.TDMF.NWTS1.RENAME.CNTL(TDRR$G01) - 01.0    Columns 00001 00072
Command ===>                                                  Scroll ===> CSR
****** **************************** Top of Data ****************************
000001 //TDRR$G01 JOB (,DASD),'IBM DATA MIGRATION',CLASS=T,
000002 //      REGION=0M,MSGCLASS=X,NOTIFY=&SYSUID
000003 //* TDMF DATA MIGRATION - DO NOT CANCEL THIS JOB !!!
000004 //STEP0001 EXEC  PGM=ICKDSF,REGION=3000K,PARM='NOREPLYU'
000005 //SYSPRINT DD  SYSOUT=(*,,1TAB)
000006 //SYSIN    DD  *
000007 REFORMAT UNIT(8400) VFY(UC8400) VOLID(UC2010)
000008 REFORMAT UNIT(8401) VFY(UC8401) VOLID(UC2011)
000009 REFORMAT UNIT(8402) VFY(UC8402) VOLID(UC2012)
000010 REFORMAT UNIT(8403) VFY(UC8403) VOLID(UC2013)
000011 REFORMAT UNIT(8404) VFY(UC8404) VOLID(UC2014)
000012 REFORMAT UNIT(8405) VFY(UC8405) VOLID(UC2015)
000013 REFORMAT UNIT(8406) VFY(UC8406) VOLID(UC2016)
000014 REFORMAT UNIT(8407) VFY(UC8407) VOLID(UC2017)
000015 REFORMAT UNIT(8408) VFY(UC8408) VOLID(UC2018)
000016 REFORMAT UNIT(8409) VFY(UC8409) VOLID(UC2019)
****** ***************************** Bottom of Data ***************************
```

Figure 10-42 ICKDSF RENAME job for renaming target volumes

Network testing: TDMF migration over TCP/IP

The following steps show the first session and what the TDMF Monitor and IBM RMF™ performance monitor display during the replication:

1. The local TDMF Master and Agents for the first TDMF session are started. The TDMF Monitor option 1 display (*Current Sessions: Monitor Progress*) is shown in Figure 10-43.

```
                              Session Monitor             Row 1 to 14 of 14
Command ===>                                                Scroll ===> CSR

Softek TDMF Master V5.5.0 Session Active.
ComDataSet : SYS1.TDMF.SDNT.LOCAL.SYSCOM01
         Source  Migration    Percent Complete --------->
Group    VolSer    Phase      ..10...20...30...40...50...60...70...80...90..100
SINGLE   UC2010   Inactive
SINGLE   UC2011   Inactive
SINGLE   UC2012   Inactive
SINGLE   UC2013   Inactive
SINGLE   UC2014   Inactive
SINGLE   UC2015   Inactive
SINGLE   UC2016   Inactive
SINGLE   UC2017   Inactive
SINGLE   UC2018   Inactive
SINGLE   UC2019   Inactive
```

Figure 10-43 TDMF Monitor Progress (local session started)

TDMF is now waiting for the remote Master system. As soon as the TDMF remote Master is started, TDMF starts the copy process for the defined concurrent number of volumes.

After a few seconds, you can see a progress bar and the estimated remaining time for the *Copy phase* of each volume as shown in Figure 10-44.

```
                                Session Monitor              Row 1 to 14 of 14
Command ===>                                                   Scroll ===> CSR

Softek TDMF Master V5.5.0 Session Active.
ComDataSet : SYS1.TDMF.SDNT.LOCAL.SYSCOM01
         Source  Migration    Percent Complete --------->
Group    VolSer    Phase      ..10...20...30...40...50...60...70...80...90..100
SINGLE   UC2010   Copy        ------------------->               0:01:12
SINGLE   UC2011   Copy        ------------------->               0:01:11
SINGLE   UC2012   Copy        -------------------->              0:01:05
SINGLE   UC2013   Copy        ------------------->               0:01:13
SINGLE   UC2014   Inactive
SINGLE   UC2015   Inactive
SINGLE   UC2016   Inactive
SINGLE   UC2017   Inactive
SINGLE   UC2018   Inactive
SINGLE   UC2019   Inactive
```

Figure 10-44 TDMF Monitor Progress (four volumes copying)

2. If TDMF does not negatively affect your systems, you can increase the number of concurrent volumes using the TDMF Monitor option 2 (*Current Sessions: User Interaction and Status*). In the experimental environment, we switched to *User Interaction and Status* by pressing PF10 and changed the number of concurrent volumes to 8 as shown in Figure 10-45.

```
                                Session Status               Row 1 to 10 of 10
Command ===>                                                   Scroll ===> CSR

System: Master               Softek TDMF for z/OS               Version 5.5.0
ComDataSet SYS1.TDMF.NWTS1.LOCAL.SYSCOM01                       Sessions   01
Number of volumes migrating 010              Number of concurrent volumes 008
Number of volumes  complete 000              Number of  volumes   waiting 002

Requested   Volume Device  Group     --- Migration --- - Error Info - Sync
 Action     Serial Number  Name      Status      Type  System Message Goal

   __       UC2010  2010  SINGLE     Copy        PIT  P                005

            UC8400       192.168.002.012 Port=08200
   __       UC2011  2011  SINGLE     Copy        PIT  P                005

            UC8401       192.168.002.012 Port=08200
                              :
```

Figure 10-45 TDMF Monitor User Interaction and Status (changing number of concurrent volumes)

3. Switching back to TDMF Monitor *Progress* by pressing F10, eight volumes are now displayed as being in the *Copy phase* (Figure 10-46).

```
                              Session Monitor              Row 1 to 14 of 14
Command ===>                                                 Scroll ===> CSR

Softek TDMF Master V5.5.0 Session Active.
ComDataSet : SYS1.TDMF.SDNT.LOCAL.SYSCOM01
         Source  Migration    Percent Complete --------->
Group    VolSer    Phase      ..10...20...30...40...50...60...70...80...90..100
SINGLE   UC2010   Copy        -------------------------->               0:01:24
SINGLE   UC2011   Copy        -------------------------->               0:01:23
SINGLE   UC2012   Copy        ---------------------------->              0:01:17
SINGLE   UC2013   Copy        -------------------------->               0:01:25
SINGLE   UC2014   Copy        ----->                                    0:05:04
SINGLE   UC2015   Copy        ----->                                    0:04:58
SINGLE   UC2016   Copy        ----->                                    0:05:07
SINGLE   UC2017   Copy        ----->                                    0:05:03
SINGLE   UC2018   Inactive
SINGLE   UC2019   Inactive
```

Figure 10-46 TDMF Monitor Progress (eight volumes copying)

After a few minutes, the refreshed TDMF Monitor *Progress* display shows that some volumes have reached the *Refresh 1 phase* (Figure 10-47).

```
                              Session Monitor              Row 1 to 14 of 14
Command ===>                                                 Scroll ===> CSR

Softek TDMF Master V5.5.0 Session Active.
ComDataSet : SYS1.TDMF.SDNT.LOCAL.SYSCOM01
         Source  Migration    Percent Complete --------->
Group    VolSer    Phase      ..10...20...30...40...50...60...70...80...90..100
SINGLE   UC2010 Refresh 1     ------------------------------------------------>
SINGLE   UC2011 Refresh 1     ------------------------------------------------>
SINGLE   UC2012 Refresh 1     ------------------------------------------------>
SINGLE   UC2013 Refresh 1     ------------------------------------------------>
SINGLE   UC2014   Copy        ------------------------------>           0:01:59
SINGLE   UC2015   Copy        ------------------------------>           0:02:01
SINGLE   UC2016   Copy        ------------------------------->          0:01:57
SINGLE   UC2017   Copy        ------------------------------->          0:01:58
SINGLE   UC2018   Copy        --->                                      0:04:33
SINGLE   UC2019   Copy        --->                                      0:04:40
```

Figure 10-47 TDMF Monitor Progress (four volumes have reached Refresh 1)

4. When the *Copy phase* completes, the volumes enter the *Refresh phase*. The message `SYNC Volume Needed` indicates that an outstanding prompt needs to be replied before the replication process for volume group continues to the *Synchronize phase* (Figure 10-48).

```
                                Session Monitor            SYNC Volume Needed.
Command ===>                                                   Scroll ===> CSR

Softek TDMF Master V5.5.0 Session Active.
ComDataSet : SYS1.TDMF.SDNT.LOCAL.SYSCOM01
         Source  Migration    Percent Complete --------->
Group    VolSer    Phase      ..10...20...30...40...50...60...70...80...90..100
SINGLE   UC2010 Refresh 1     ------------------------------------------------>
SINGLE   UC2011 Refresh 1     ------------------------------------------------>
SINGLE   UC2012 Refresh 1     ------------------------------------------------>
SINGLE   UC2013 Refresh 1     ------------------------------------------------>
SINGLE   UC2014 Refresh 1     ------------------------------------------------>
SINGLE   UC2015 Refresh 2     --------------------------->
SINGLE   UC2016 Refresh 1     ------------------------------------------------>
SINGLE   UC2017 Refresh 1     ------------------------------------------------>
SINGLE   UC2018 Refresh 1     ------------------------------------------------>
SINGLE   UC2019 Refresh 2     ------------------------------------------------>
```

Figure 10-48 TDMF Monitor Progress (group in Refresh phase)

In the experimental environment, updates to the source volumes have been stopped (emulating an application suspension). The Synchronize phase can be initiated by issuing the PG line command for the first volume on the *User Interaction and Status* panel (press the F10 key to switch between the *User Interaction and Status* and *Monitor Progress* screens) as shown in Figure 10-49. The command must be confirmed in the request panel.

```
                                Session Status             SYNC Volume Needed.
Command ===>                                                   Scroll ===> CSR

System: Master               Softek TDMF for z/OS               Version 5.5.0
ComDataSet SYS1.TDMF.SDNT.LOCAL.SYSCOM01                        Sessions   01
Number of volumes migrating 010              Number of concurrent volumes 008
Number of volumes  complete 000              Number of  volumes   waiting 000

Requested   Volume Device  Group     --- Migration --- - Error Info - Sync
 Action     Serial Number  Name      Status      Type  System Message Goal

   pg      *UC2010  2010  SINGLE   Waiting Sync  PIT  P                005

            UC8400       192.168.002.012 Port=08200
   __       UC2011  2011  SINGLE   Refresh 1     PIT  P                005

            UC8401       192.168.002.012 Port=08200
   __       UC2012  2012  SINGLE   Refresh 1     PIT  P                005
```

Figure 10-49 TDMF Monitor User Interaction and Status (initiate Synchronize phase for group)

By repeatedly pressing Enter and refreshing the panel, you can see volumes entering the Synchronization phase as shown in Figure 10-50. The upper left area displays the total numbers of volumes in this session and a dynamically updated counter of completed volumes.

```
System: Master               Softek TDMF for z/OS              Version 5.5.0
ComDataSet SYS1.TDMF.SDNT.LOCAL.SYSCOM01                       Sessions   01
Number of volumes migrating 010             Number of concurrent volumes 008
Number of volumes  complete 007             Number of  volumes   waiting 000

Requested   Volume Device  Group     --- Migration --- - Error Info - Sync
 Action     Serial Number  Name      Status      Type  System Message Goal

  __        UC2010  2010  SINGLE     Synchronize PIT  P                005

            UC8400       192.168.002.012 Port=08200
  __        UC2011  2011  SINGLE     Complete    PIT  P                005

            UC8401       192.168.002.012 Port=08200
  __        UC2012  2012  SINGLE     Complete    PIT  P                005

            UC8402       192.168.002.012 Port=08200
  __        UC2013  2013  SINGLE     Synchronize PIT  P                005

            UC8403       192.168.002.012 Port=08200
                      :
```

Figure 10-50 TDMF Monitor User Interaction and Status (start SYNC phase)

As soon as all volumes within the session are completed, the session disappears from the TDMF Monitor screens. All TDMF session jobs terminate. Check all job logs to see whether MAXCC is greater than zero (Figure 10-51).

```
06.21.19 JOB02611 $HASP165 TDMRX101 ENDED AT TXRTEST  MAXCC=0 CN(INTERNAL)
```

Figure 10-51 TDMF job information at the remote site

A MACCCC=4 indicates that warning messages have been issued. Review the job logs and determine whether the warning messages can be tolerated.

In case of MAXCC>4, check the job logs for error messages. You will possibly need to restart the replication after taking any measure required to successfully re-execute the TDMF session.

Performance of a TDMF migration over TCP/IP

In the following example, RMFMON was used to display the channel load of the Open System adapter (OSA) CHPID 28 (Figure 10-52).

With a utilization of less than 50 percent, the transfer rate was not higher than 64 MBps. This rate is an indication that the OSA on this side of the migration is not causing delays. During this experiment, the source volumes were ESCON attached with four channel paths. Because those paths only provided 15 MBps each, the maximum migration throughput was limited to approximately 60 MBps plus some additional protocol-related processor burden.

```
F                                CPU= 11/ 11 UIC= 65K PR=   0  SYCP CHANNEL  T
06:07:15  CHANNEL UTILIZATION(%)   READ(B/S)  WRITE(B/S)  FICON OPS  ZHPF OPS
ID NO  G  TYPE  S PART  TOT  BUS   PART  TOT  PART  TOT   RATE ACTV  RATE ACTV
28        OSD   Y 46.5 47.2 35.6   227K  229K  64M  64M
29        OSE   Y  0.0  0.3  7.3      0    0     0    0
```

Figure 10-52 RMFMON Channel statistics

With faster channels like FICON, the limiting resource in a remote replication session is usually the bandwidth that is provided by the network connection. There can be many factors contributing to reduced throughput rates, and sophisticated analysis of network problems is beyond the scope of this publication.

TCP/IP information

This section covers information about network aspects related to migrations using TDMF z/OS over the TCP/IP connections.

TCP/IP settings for TDMF replication

You can specify the size of the send and receive buffers used by the TCP/IP communication by specifying the **SNDBuf** and **RCVBuf** options on the **REMOTE** control statement for a TDMF migration session. If these are not specified, the system defaults are used. These options can be verified by issuing the TSO **netstat config** command. The example output in Figure 10-53 shows that both the default send and receive buffer sizes are set to 256 K.

```
EZZ2350I MVS TCP/IP NETSTAT CS V1R11       TCPIP Name: TCPIPBC         04:52:02
EZZ2700I Home address list:
EZZ2701I Address          Link             Flg
EZZ2702I -------          ----             ---
EZZ2703I 192.168.2.04     LOSAEG09         P
EZZ2703I 127.0.0.1        LOOPBACK
EZZ2350I MVS TCP/IP NETSTAT CS V1R11       TCPIP Name: TCPIPBC         04:52:02
EZZ2720I TCP Configuration Table:
EZZ2721I DefaultRcvBufSize:  00262144  DefaultSndBufSize: 00262144
EZZ2727I DefltMaxRcvBufSize: 00262144  SoMaxConn:         0000000010
EZZ2722I MaxReTransmitTime:  120.000   MinReTransmitTime: 0.500
EZZ2723I RoundTripGain:      0.125     VarianceGain:      0.250
EZZ2724I VarianceMultiplier: 2.000     MaxSegLifeTime:    30.000
EZZ2725I DefaultKeepALive:   00000120  DelayAck:          No
EZZ2726I RestrictLowPort:    Yes       SendGarbage:       No
EZZ2728I TcpTimeStamp:       Yes       FinWait2Time:      300
EZZ2729I TTLS:               No
                      :
```

Figure 10-53 Result of TSO netstat config command

TCP port mapping of a TDMF session

For each TDMF session, a pair of consecutive TCP port numbers is required at the remote site. The first of these port numbers is specified in the TDMF Master JCL (local and remote site) and used to establish the communication between the two Master jobs. The second port number is used to establish data transfer sessions for each volume pair (source and associated target volume).

In the experimental environment, a TSO **`netstat conn`** command shows port numbers 8200 and 8201 being used (Figure 10-54). At the local site, a dynamically allocated ephemeral port number is mapped to the remote port number 8201. In the example, 10 volumes are being replicated and you can see 10 different connections to remote port number 8201. The TCP session connected to the remote port number 8200 is used as a control communication path between the two Master sessions.

```
EZZ2350I MVS TCP/IP NETSTAT CS V1R11       TCPIP Name: TCPIPCP         04:52:02
EZZ2585I User Id  Conn     Local Socket           Foreign Socket         State
EZZ2586I -------  ----     ------------           --------------         -----
EZZ2587I BPXOINIT 00000013 0.0.0.0..10007         0.0.0.0..0             Listen
EZZ2587I DB1ADIST 0000349D 0.0.0.0..5051          0.0.0.0..0             Listen
EZZ2587I DB1ADIST 0000349E 0.0.0.0..5054          0.0.0.0..0             Listen
EZZ2587I DB1ADIST 000034A2 0.0.0.0..5052          0.0.0.0..0             Listen
EZZ2587I FTPDB1   00013FE7 0.0.0.0..21            0.0.0.0..0             Listen
                  :
EZZ2587I TDMLBC01 000587A7 192.168.2.04..2230     192.168.2.12..8201     Establsh
EZZ2587I TDMLBC01 000587A8 192.168.2.04..2231     192.168.2.12..8201     Establsh
EZZ2587I TDMLBC01 000587AF 192.168.2.04..2238     192.168.2.12..8201     Establsh
EZZ2587I TDMLBC01 000587AE 192.168.2.04..2237     192.168.2.12..8201     Establsh
EZZ2587I TDMLBC01 000587A9 192.168.2.04..2232     192.168.2.12..8201     Establsh
EZZ2587I TDMLBC01 000587AC 192.168.2.04..2235     192.168.2.12..8201     Establsh
EZZ2587I TDMLBC01 000587AA 192.168.2.04..2233     192.168.2.12..8201     Establsh
EZZ2587I TDMLBC01 00058787 192.168.2.04..2198     192.168.2.12..8200     Establsh
EZZ2587I TDMLBC01 000587AD 192.168.2.04..2236     192.168.2.12..8201     Establsh
EZZ2587I TDMLBC01 000587AB 192.168.2.04..2234     192.168.2.12..8201     Establsh
EZZ2587I TDMLBC01 000587A6 192.168.2.04..2229     192.168.2.12..8201     Establsh
                  :
```

Figure 10-54 Viewing the TCP/IP port mapping for a TDMF session

Note: Each TDMF session running in parallel requires its own communication data set and unique pair of TCP port numbers.

TCP/IP information for an active remote TDMF replication session

Issuing the TSO **netstat all** with the **CLIent** filter option displays useful information about each TCP session established between the local and remote TDMF Master systems. Figure 10-55 shows an excerpt from the output of the TSO **netstat all (cli tdmlbc01** command for an active local TDMF Master job.

```
----
Client Name: TDMLBC01                 Client Id: 000587A7
  Local Socket: 192.168.2.04..2230
  Foreign Socket: 192.168.2.12..8201
    BytesIn:            00000000000000000460
    BytesOut:           00000000000038915276
    SegmentsIn:         00000000000000001625
    SegmentsOut:        00000000000000027598
    StartDate:          11/06/2014       StartTime:          04:57:57
    Last Touched:       04:58:02         State:              Establsh
    RcvNxt:             2190992725       SndNxt:             2675598415
    ClientRcvNxt:       2190992725       ClientSndNxt:       2675598415
    InitRcvSeqNum:      2190992264       InitSndSeqNum:      2636683138
    CongestionWindow:   0000101951       SlowStartThreshold: 0000049504
    IncomingWindowNum:  2191000457       OutgoingWindowNum:  2675679847
    SndWl1:             2190992725       SndWl2:             2675499751
    SndWnd:             0000180096       MaxSndWnd:          0000229312
    SndUna:             2675499751       rtt_seq:            2675559219
    MaximumSegmentSize: 0000001448       DSField:            00
    Round-trip information:
      Smooth trip time: 3.000            SmoothTripVariance: 2.000
    ReXmt:              0000000001       ReXmtCount:         0000000000
    DupACKs:            0000000002       RcvWnd:             0000007732
    SockOpt:            0C00             TcpTimer:           C2
    TcpSig:             04               TcpSel:             C0
    TcpDet:             E0               TcpPol:             02
    TcpPrf:             00
    QOSPolicy:          No
    RoutingPolicy:      No
    ReceiveBufferSize:  0000004096       SendBufferSize:     0000114688
    ReceiveDataQueued:  0000000000
    SendDataQueued:     0000000000
    SendStalled:        Yes
    Ancillary Input Queue: N/A
----
```

Figure 10-55 Output from the netstat all command

11

z/OS Dataset Mobility Facility

z/OS Dataset Mobility Facility (zDMF) is a host-based (mainframe) data migration facility that migrates data at the logical data set (extent) level between sets of volumes. All data movement is accomplished without application downtime.

This chapter outlines the steps needed to conduct a successful data migration project that uses zDMF. It also describes the standards and provides guidelines when using zDMF for a data migration project.

zDMF is a data set level migration tool that updates the information in the ICF catalog and interacts with SMS. Therefore, there are sections that address its role in the migration process, detailed explanations of the environmental requirements, and reference materials.

This chapter includes the following sections:

- z/OS Dataset Mobility Facility overview
- Terminology
- zDMF migration process
- Important zDMF-related topics
- Migration in an experimental shared DB2 test environment

© Copyright IBM Corp. 2015. All rights reserved.

11.1 z/OS Dataset Mobility Facility overview

The task of migrating small volumes onto new, large volume capacity storage subsystems is difficult and time-consuming. Historically, merging volumes required outages of revenue-generating and often mission-critical applications for the entire duration of a data migration. zDMF allows you to migrate and consolidate data sets while minimizing or eliminating application downtime.

zDMF moves data sets without disrupting access to applications during the migration. zDMF also provides the following advantages:

- Facilitates combining smaller capacity volumes onto larger volumes
- Reduces consumption of available system resources
- Supports the implementation of tiered storage
- Improves application performance for service level compliance
- Maintains application availability during data set level migrations
- Allows continued growth of important or new applications
- Reduces storage platform total cost of ownership (TCO)

In addition, allocated data sets can be migrated with a dynamic swap of metadata. Table 11-1 shows the capabilities of zDMF.

Table 11-1 zDMF capabilities

Key factors	Description	zDMF capability
Dynamic swap	Automatically redirects I/O to a new location after diversion completes.	Allows data migration without disruption.
Flexible migration options	Provides control over the I/O rate for data migration reads/writes.	Allows online migration activity while maintaining optimal application performance and service levels.
Data set grouping	Allows a group of data sets to be collectively migrated.	Provides easy management of large migrations.
Robust interface	ISPF application for managing and monitoring.	Allows easy configuration, monitoring, and operations.
Resilient architecture	Current state of migration is always maintained, regardless of controlled or non-controlled shutdown, to ensure continued/recoverable migration.	Migration process is “checkpoint” restartable during server shutdown and restart.

Although zDMF provides a sophisticated data migration tool, it is a sensitive mechanism observing every data set activity on volumes containing identified migration data sets. Although it is a stable tool, it causes some burden to the I/O processing and it makes sense to reduce that burden to a minimum during a migration. Overhead can be reduced by avoiding high I/O activity during a migration, reducing the overall migration time by considerate planning, and the correct choice of migration tools.

Significant burden, for example, is caused by starting a migration using zDMF for data sets that frequently get deleted and reallocated during the migration process. High I/O activity can also cause some burden, especially during the *Mirror phase* where each write I/O is run to the source and target volume.

After all, even when fully migrated with zDMF, some data sets require a short application outage to fully complete the migration process. Considerate planning of required application outages reduces the overall migration burden.

11.2 Terminology

The following terms are used when describing data migration processes using zDMF:

- zDMF Server: The zDMF Server is a software component of zDMF running as a started task on a z/OS image.
- Source volume: A volume where data sets to be migrated are located.
- Target volume: A volume where data sets are migrated to.
- Source data set: A data set to be migrated that is on its source volume.
- Target data set: A data set on a target volume that receives an associated source data set.
- Migration data set: A data set that should be migrated with zDMF. It is also used to refer to the pair of original source data set and its associated target data set.
- Migration group: An entity that assigns the data sets to be migrated with certain migration attributes.
- zDMF data base: A data set used by the zDMF Servers running on all z/OS images participating in a migration to share migration-related information.
- Mirroring: The process where write I/Os are performed to both source and target data set.
- Synchronization: The process where identity between source and target data set is achieved.
- Diversion: The process when I/Os to the original source data set are diverted to its corresponding target data set.

11.3 zDMF migration process

This section outlines the steps that are needed to conduct a successful data migration project using zDMF. zDMF is designed to move allocated data sets that require nondisruptive data migration.

11.3.1 zDMF migration steps

The zDMF Monitor application is usually used to define a process that performs all steps of a data set migration. The sequence of steps involved in a migration using zDMF is described in Table 11-2.

Table 11-2 zDMF migration steps

Migration steps	Description
Group definition	The data sets to be migrated and other migration specifics are defined in a migration group using the zDMF Monitor, or by using the zDMF batch utility. When using the zDMF monitor, the group definition is stored in a partitioned data set.
Promoting a migration group definition	Storing a migration group definition in the zDMF database.

Migration steps	Description
Activating a migration group (Activation phase)	Activating a migration group initiates the data migration process for that group.
Copy phase	Data is asynchronously copied from the source data sets to the target data sets defined in the migration group.
Synchronization phase	All final differences between source and target data sets in a migration group are synchronized and the group is prepared for mirroring.
Mirror phase	The migration group is put into a state of synchronous mirroring. I/Os to the source data set are mirrored to the target data set.
Diversion phase	The actual logical relocation of data sets occurs. Source and target data set metadata (that is, catalog entries, VTOC, VVDS) is modified and all I/O activity is redirected to the new data sets location.
Completion phase (post migration)	Although the metadata has been modified, applications that were active before diversion continue to have their I/O redirected until they release (de-allocate) the data set. A short planned application outage is required to fully complete the data set migration.
Post-Completion phase	After the migration is complete, the original source data sets (which have been renamed to *hlq.group.date.time.sequence_number*[a]) and other used storage resources can be cleaned up.

a. *hlq* has been specified in the group definition as the target data set high-level qualifier

11.3.2 Planning the migration

Important: The planning phase is the most vital phase for a successful migration.

This section outlines the generic steps, personnel functions, and tasks that pertain to the planning phase. The high-level process consists of the following steps:

1. Establish a migration management team that consists of:
   - Primary Migration Manager
   - Alternate Migration Manager
   - Account Coordinator
   - Security Coordinator (if required for sensitive data)
   - Technical Lead Coordinator
2. Announce the migration with a notification that aligns with the application requirements.
3. Size the target configuration to accommodate the source configuration plus predicted future expansion.
4. Create a migration team to perform the pre-migration tasks listed in the following section. The migration team should include a technical data migration team trained to use zDMF and any other necessary utilities (such as DFSMShsm, DFSMSdss, SMS Redirection). This technical team can also act as migration mentors.

11.3.3 Pre-migration tasks

Complete the following tasks before starting a migration that includes zDMF:

1. Make sure that the account coordinator identifies the owners of the data. Inform the owners about the migration and update them as needed.
2. Inform the security and compliance groups about the migration if you will be migrating sensitive data.
3. Use the Pre-Migration Planning Checklist to ensure that all of the pre-migration planning steps are run.
4. Examine any involved source catalog using IDCAMS EXAMINE with the INDEXTEST and DATATEST parameters. Verify that the source catalogs have appropriate free space and structural integrity to accommodate growth outside the scope of the migration.
5. Create a dedicated zDMF user catalog with appropriate free space to accommodate all target data sets. Connect this catalog to the same systems as the source catalogs. Place the catalog on a non-SMS management volume that does not contain customer data.
6. Create and relate some unique aliases to the previously created zDMF user catalog. The aliases are used as the high-level qualifier to create temporary target data set entries according to the following pattern:

   `hlq.groupname.date.time.sequence_number`
7. Create or add the new target volume SMS Storage Group Configuration. The new storage capacity must be the same or greater than the source storage. Consideration must also be given to any multivolume data sets.
8. Add the new target volumes to the appropriate storage group and change their status to ENABLE before the migration is started.
9. Create a list of the source volumes that contain data to be moved and the associated target volumes or storage group.
10. Change the status of the source volumes to DISNEW to accommodate SMS redirection and prevent allocation of new data sets on source volumes during the migration.
11. Identify suitable time frames for the different migration phases. Consider data set and Group activity, and fallback requirements when scheduling. In particular, schedule around heavy I/O loads like batch work or IBM DB2® reorganizations during the migration phase.
12. Identify tools that help reducing the migration processor requirements.

Remember: Some of these steps might not be applicable because of your internal policies on data migration.

11.3.4 Tools for migrating data sets

There are a number of tools available to facilitate a z/OS based data migration at the data set level from source volumes onto new target volumes. Each of them has its own characteristics and capabilities. A successful migration is best achieved by combining the most suitable tools.

The following tools are described in this section:

- DFSMS
- DFSMShsm
- DFSMSdss
- TDMF
- zDMF

How DFSMS can help during a migration

Data Facility Storage Management Subsystem (DFSMS), or SMS in short, can be used to facilitate a data migration project. SMS manages data set creation. If you want to migrate entire volumes (all data sets on a particular volume), you can prevent allocation of new data sets on these volumes by changing their volume status to DISNEW and enabling a new target volume (change status to ENABLE).

In an environment that has a high delete and reallocation activity, you can move that data off the source to new target volumes by using this feature of SMS.

Verify that the storage group that contains those volumes has enough spare volumes to satisfy allocation requests for new data sets that will not be migrated.

However, use this technique only for data sets that are frequently deleted and reallocated.

Note: SMS can be used to automatically migrate data sets with high delete and reallocation activity.

DFSMShsm as a convenient migration tool

For data sets that are not deleted and reallocated regularly and are not continuously allocated, a window of opportunity[1] can be determined that accommodates a migration. During this time, you can use Hierarchical Storage Manager (DFSMShsm) or another standard utility like DFSMSdss to migrate those data sets.

By disabling original source level 0 volumes and enabling new target volumes for allocation, HSM can be used to migrate data sets that are *not in use* (not allocated) from level 0 volumes onto migration level 1 volumes. It can then recall them onto a new target volume using a simple HSM command.

For example, to replace volume SG2002 with a new volume, you can complete these steps:

1. Modify or create an appropriate ACS routine to select the new volume when recalling data sets.
2. Prevent new allocations to volume SG2002.

   To prevent new allocations to the volume, change the status of the volume in the storage group. Usually the status for that volume can be set to DISNEW using ISMF.
3. Move the non-allocated data sets off the volume using the following HSM command:

   ```
   MIGRATE VOLUME(SG2002 MIGRATE(0)) CONVERT
   ```

   DFSMShsm migrates all of the non-allocated (*not in use*) data sets that it can process off volume SG2002 and recalls them to another level 0 volume. The automatic class selection (ACS) routine selects the volume to which the data set is recalled.

   Although the actual migration process requires only a simple HSM command, sufficient space must be available on a migration level 1 volume that will be used as an intermediate storage location.

Note: Using DFSMShsm requires some definition changes and intermediate disk storage. It can only be used to migrate data sets that are not in use.

[1] In which the data set is not allocated by any application and is long enough to perform a migration

DFSMSdss is a powerful migration tool

A successful migration requires a combination of various utilities that complement each other and enhance the migration process. For data sets that are not permanently in use (allocated to applications), there might be a window of opportunity to accommodate a migration using a standard data set migration utility such as DFSMSdss.

Although DFSMSdss can be used to copy allocated and opened data sets, these copies will not be in a consistent state because applications might have updated the existing source data sets after the copy process completed.

For data sets that are not in use (not allocated) it is a powerful tool that requires JCL to be created, but does not require intermediate storage areas.

Note: DFSMSdss can be used to migrate data sets that are not in use. If not using DFSMSdss, a Job Control Language (JCL) script must be created to migrate the data set.

Using TDMF can be part of a migration

Although TDMFs main purpose is the transparent, nondisruptive migration of entire disk volumes, it can play a role when a merge migration from smaller disk volumes onto larger new volumes is planned.

TDMF can replace a small source volume with a new large disk volume without requiring any application downtime. However, TDMF can only swap a smaller volume with exactly one larger volume by moving the original source data at the beginning of the new large target volume.

Note: TDMF cannot merge multiple small source volumes onto one large target volume.

Why use zDMF at all?

zDMF is useful for migrating all permanently allocated data sets or those for which it proves difficult to find an application downtime window.

zDMF migrates data sets while they are in use and maintains a mirror image of the source data set on a new target volume. Only a short application outage is required to finalize the migration process, and the application can reuse the data set on its new location.

11.3.5 zDMF migration planning

Using zDMF as a tool in a data migration project helps reduce application outages to a minimum. zDMF migrates data in a phased process. In one phase, data are copied from the source volumes to the target volumes, and in a later phase any updates to the source are mirrored to the target device. This causes some I/O processor burden to the affected applications. This burden can somehow be reduced by taking data sets or an entire group of data sets into the *Divert phase*, in which I/Os are redirected from the original source to the new target volume. However, this phase can be considered vulnerable as it usually does not provide a way back to using the source data set in case of serious errors or other harmful events. For this reason, the time that data sets remain in the *Divert phase* while in use by applications should be planned carefully if it cannot be avoided.

zDMF data set selection planning

zDMF includes a multitude of data set selection criteria, enabling you to select from a single data set to thousands of data sets in one or multiple groups.

Consider the following when specifying data set selection criteria for data sets to be migrated using zDMF:

- Prevent new allocations to source volumes in an SMS configuration by setting the volume status to DISNEW.
- Data sets that will probably be deleted and re-created during the zDMF migration cycle should be handled by DFSMS ACS Routines and excluded from being migrated with zDMF.
- Review the delete activity cycle of GDG. If the cycle is shorter than the zDMF migration cycle, allow them to move themselves by disabling source volumes for new GDG allocations.
- Change non-SMS-managed volumes that are mounted with a STORAGE status to PRIVATE.
- The storage subsystem must have sufficient space to accommodate normal processing. This might not be the case for data set relocation or workload balancing.
- Identify data sets with primary space 0 (SPACE=(CYL,0,100)) because they should not be migrated.
- The security profile that is associated with a particular data set might prevent you from moving the data set.
- When moving a data set, for performance reasons make sure that the new target volume does not already have any performance constraints.
- When doing volume consolidation, make sure not to select multiple source volumes with performance constraints within the same zDMF migration.
- Try to schedule the *Completion phase* and if possible the *Divert phase* to coincide with a planned application outage.
- Identify data sets that have a specific volser dependency. Moving data sets that have a volser dependency without including the new volser can cause a problem.
- Identify all the volumes that are associated with multi-volume data sets. The same number of target volumes are required.
- Use the Migration Planning Checklist to establish tasks, assignments, and status.

Important: Data sets on volumes that contain catalogs should be migrated separately. Make sure that catalog maintenance is not scheduled during the migration.

Restrictions

The following data sets are not supported by zDMF:

- Data sets cataloged in the master catalog or on any system resident volumes should not be included in a migration.
- Data sets used by the operating system such as LINKLIST APF-authorized, page, JES2, and JES3. These data sets are not supported because they might not divert properly, causing serious system outages.
- Virtual Storage Access Method (VSAM) with IMBED, KEYRANGE, and Replicate parameters. For more information, see the *zDMF Installation & Reference Guide*.
- Catalogs.
- ISAM.
- Individual PDS members.
- HFS data sets.

- Page and swap data sets.
- Temporary (&&) data sets.
- Undefined data sets.
- Data sets allocated with no *primary* space or *utilized* space.
- Volume-specific data sets, including volume table of contents (VTOC), VSAM volume data set (VVDS), and volume table of contents index (VTOCIX).

Both control unit types must be of an equal or higher type.

```
Example:
2105 -> 3990 will not be moved
3990 -> 2105/2107 will be moved
2105 -> 2105/2107 will be moved
```

The following data sets are flagged and do not go into the *Divert phase* until all associated applications are stopped on all participating systems. This Extended Mirroring is enforced in environments that do not convert hardware reserves. For more information, see the "Migration scheduling techniques" section in the *zDMF Installation & Reference Guide*.

- VSAM record level sharing data sets
- BDAM data sets that require absolute track
- PDSE data sets in use by the SMSPDSE address space

Data migration classifications

The following are the classifications used with data migration:

- Specify data set selection criteria carefully so that the following limitations are met for any data set migration group:
  - Less than 1,500 data sets
  - Less than 40,000 extents in total
- Appropriate sizing of migration groups reduces exposures related to ECSA, CSA, and extended private storage consumption. It also mitigates the ENQ impact of migration processing on catalogs and volumes.
- An EXCLUDE LIST to separate data sets into various groups. Create a MASTER EXCLUDE LIST. This list includes identifying the following types of data sets:
  - RESTRICTED (ISAM, HFS, and so on)
  - VOL SER DEPENDANT data sets
  - SOFTWARE/SYSTEM data sets that are sensitive to movement

Migration group standards

Follow these guidelines when creating migration groups:

- Create group naming standards that provide easy identification. zDMF prefixes a member name with 'ZD', limiting your choice to six characters.
- Create groups that you can manage based on the number of data sets.
- Retain group definition PDS members for reference.

11.3.6 Typical migration scenario

This section outlines the steps of a typical storage migration. In this scenario, a technology upgrade is required because the lease of the current storage subsystem is ending. Therefore,

data needs to be migrated from old to new storage. It was also decided to install new higher capacity drives.

The migration metrics are shown in Table 11-3.

Table 11-3 Migration metrics

Old storage	New storage	Capacity
800 MOD-3 to	800 MOD-3 to	2270 GB
120 MOD-3 to	40 MOD-9	340 GB
200 MOD-9 to	200 MOD-9	1702 GB
240 MOD-9 to	80 MOD-27	2043 GB
1360 volumes to	1120 volumes	TOTAL=6355 GB

In this scenario, the old subsystem had to be removed within 30 days after delivery of the new subsystem to avoid costs. A conventional disk-to-disk copy was not feasible because long application outages would have been required. The old hardware migration utility could not be used because of compatibility issues going from ESCON to FICON. Ultimately, a combination of host-based utilities and products was chosen to run the migration.

Because 80-90% of data to be migrated were on SMS-managed volumes, the following methods were used during the migration:

- DFSMS to reallocate data sets with high delete and reallocation activity
- DFSMShsm to copy non-allocated files
- TDMF for volume level migrations
- zDMF for allocated files

Selection logic

TDMF was chosen to move volume data and facilitate the activation of new disk hardware. However, a nondisruptive logical data set migration tool was needed to complement DFSMS for moving allocated data sets and use new larger disks.

Performing the migration

The following steps were performed:

- Labeling new target volumes
- Migrate smaller to larger volumes with TDMF
- Migrate using DFSMS
- Migrate same capacity volumes with TDMF
- Migrate with zDMF
- Completing the migration
- Verifying the migration

Labeling new target volumes

An ICKDSF Minimal Init was performed for the new volumes (the migration target volumes) labeling them with a XXucb# volume serial number. This function was used rather than initializing the new volumes with volume serial numbers that conformed to the customer's naming convention. The new volumes were not defined to DFSMS.

Migrate smaller to larger volumes with TDMF

TDMF was used to migrate smaller volumes onto larger volumes. In this case, MOD-9 were migrated onto MOD-27. This completed one-third of the migration without any disruption to

production. It took less than one hour to migrate 40 MOD-3 to MOD-9 volumes, and about four hours to migrate 80 MOD-9 to MOD-27 volumes.

The TDMF migration was set up to dynamically start ICKDSF to reformat and expand the VTOC of the target volumes to accommodate their larger capacity.

Migrate using DFSMS

The remaining two-thirds of the source volumes needed to be consolidated to larger capacity volumes. These volumes were placed in the DISNEW status to prevent any new data sets from being allocated on the old source volumes. This also caused data sets that were deleted and reallocated during daily and weekly production runs to be migrated to the new devices.

This task was monitored to determine which files did not get migrated by SMS.

DFSMShsm was used to migrate data sets that were not allocated off old source volumes by using the following command:

```
MIGRATE VOLUME(srcvol MIGRATE(0)) CONVERT
```

Migrate same capacity volumes with TDMF

TDMF was used to migrate same capacity volumes that also contained data sets contained in the restriction list of zDMF. 800 MOD-3 were migrated to MOD-3 and 200 MOD-9 were migrated to MOD-9 DASD.

Migrate with zDMF

Data sets that were not reallocated by DFSMS or moved by DFSMShsm onto the new devices were migrated by zDMF.

The remaining data sets were primarily those that are permanently allocated like DB2 or IBM CICS® files.

Consideration was given to using DFSMSdss to migrate remaining files that are not permanently allocated. To preserve data consistency, DFSMSdss should only be used for data sets that are not allocated. However, application downtimes might not be sufficient for DFSMSdss to accommodate complete data set migrations. zDMF can migrate data sets independent of their current allocation status, and can tolerate application access during a migration. Therefore, zDMF was used to migrate all remaining data sets.

Completing the migration

All data sets except the ones with persistent allocations were migrated by zDMF to completion. The files that remained in the zDMF *Divert phase* needed an application outage to complete migration. Corresponding applications were scheduled to be restarted over the weekend.

After zDMF completed all data set migrations, the Storage Migration project was complete.

Verifying the migration

An extra day was spent to verify that all data had migrated as planned and post-migration tasks had been completed.

Migration timeline estimate

Figure 11-1 shows an estimated timeline for the migration described previously:

- Pre-migration: Install new storage the first week.
- Migration Step 1: Label new disk volumes

- Migration Step 2: TDMF small volumes to large volumes
- Migration Step 3: SMS DISNEW redirection
- Migration Step 4: TDMF same capacity volumes
- Migration Step 5: zDMF migration
- Migration Step 6: Scheduled application restart (if required)
- Migration Step 7: Clean-up
- Post-migration: Two weeks to remove old storage (installation, migration, and de-installation complete in one month)

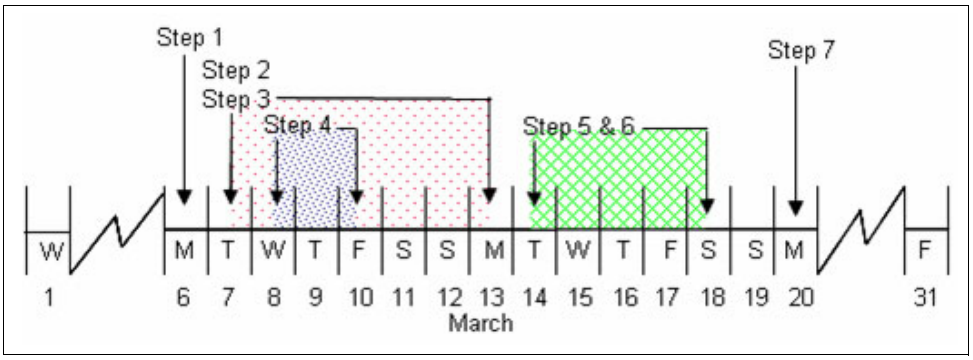

Figure 11-1 Sample migration timeline

11.3.7 Migration performance and scalability

This section addresses the zDMF performance parameters that establish the user needs with specific migration requirements and system capabilities. Guidelines for sizing an zDMF migration group are also included.

zDMF performance considerations

zDMF uses both asynchronous and synchronous techniques to copy data from source locations to target locations, minimizing performance impact.

A zDMF Server must always be active on all z/OS images sharing any data set to be migrated. zDMF requires a high execution priority and should be put in a response-oriented performance group.

zDMF parameters

The following general zDMF parameters can be specified by a configuration member:

- MAXIO: Specifies the maximum overall number of I/O requests that can be active at a time on the server. Because I/O buffers and control areas are allocated based on the MAXIO value, select a number appropriate to the resources available.
- MAX_CHANNEL_IO: Specifies the maximum number of concurrent I/O requests that can be issued to a channel during the *Copy phase* or *Mirror/Synchronization phase*.
- MAX_DEVICE_IO: Specifies the maximum number of concurrent I/O requests that can be issued to a device during the *Copy* and *Mirror/Synchronization phases*.
- MAXTRK: This optional parameter specifies the size of I/O operations in tracks that zDMF I/O copy operations use to transfer less than a full cylinder of data. The MAXTRK value is used to reduce the application response time impact of zDMF Copy operations immediately following activation. For example, MAXTRK=5 causes zDMF to move one

extent in three I/O operations. Splitting the extent reduces the time that the device is unavailable to application I/O operations into three short windows.

- **U**NIDENTIFIED_SYSTEMS_ACTION: Tells zDMF to check that a zDMF Server is active on all z/OS images that can access a shared migration volume on a standard 3990-6 compatible control unit. It has these possible values:
  - IGNORE causes no action
  - WARNing results in warning messages (GZD4192W or GZD4194W)
  - TERMinate causes the group activation to fail if not all z/OS images with access to the source volume have an active zDMF Server. The failure message is GZD4191E or GZD4193E. If a target volume fails the check, it is removed from candidacy, message GZD4195I is issued, and the activation continues if other target volumes are available.

  Tip: Initially run in TERMinate mode to verify correct operations and build an IGNORE_SYSTEMS list if needed.

  This parameter works with 3990-6 control units or later (2105/2107) only.

  Important: You must ensure that every z/OS image that has access to any migration data set runs a zDMF Server. Any system without a zDMF Server (z/OS and non-z/OS) must be prevented from accessing any migration data set during the migration process.

- IGNORE_SYSTEMS: Allows the specification of system CPU IDs that will be tolerated without a zDMF Server being run. This is often required if TERMinate is specified as the UNIDENTIFIED_SYSTEMS_ACTION because storage controllers can report systems back to zDMF that are known to not access migration data sets.

Sizing a zDMF migration group

Serialize state changes when multiple groups are used to migrate a population of data sets. To avoid possible application performance impact and deadlock exposure, ensure that groups progress serially through the *Activation* and *Pending-Divert phases*. When activating migration groups, make sure that each group is showing some progress in the *Copy/Synchronization phase* before activating the next group. Likewise, ensure that each group is in a state of Mirrored before entering the **`divert`** command to put successive groups in diversion.

Application performance impact can be reduced by specifying the MAXTRK zDMF server parameter. The MAXTRK parameter causes zDMF I/O operations to be performed using less than one cylinder. For more information, see “Configuring the zDMF Server Parameters” in the *zDMF Installation & Reference Guide*.

zDMF storage requirements

zDMF checks the available free space in ECSA while the group is activated. In the case of ECSA shortages, groups are not activated. The following are the basic ECSA, CSA, and extended private storage requirements for zDMF:

- A global area is acquired that needs X'CE0' (3,296) bytes of fixed ECSA.
- Operating system interface modules are loaded into fixed ECSA, and require at least X'DAA08' bytes (874.5 KB).
- The PC routine is loaded into pageable ECSA, which requires approximately 320 KB.

- The storage pools (fixed ECSA) are initialized with X'480000' bytes (4.5 MB). When groups are mirroring or diverting, the storage pools expand depending on the complexity of the channel programs.
- When a group is promoted and activated, the parsing routine uses ECSA to store images.

  For each active group, the following considerations apply (in fixed ECSA):

  – Control blocks for which space is not freed:
    - One DWGRP block, X'1000' (4096) bytes, until the next IPL or the zDMF Server is shut down.
    - One DWGDEV block for each source and target device, X'F8' (248) bytes.
    - A sparse array for extent look-up from each device block. Minimum size is X'400' bytes (1 KB).

  – Control blocks for which space is freed when the group completes:
    - One QCDSN block for every data set and each individual VSAM component or AIX, X'2D8' (728) bytes.
    - One DWEXT block for each data set extent, X'A6' (166) bytes.

Note: At least 0.85 MB will be resident in fixed ECSA until the next IPL for all components, except for the PC routine and possibly the storage pools.

Security

Protect zDMF using IBM RACF®, ACF2, or other security utilities to prevent use by unauthorized personnel. Use of zDMF by unauthorized personnel might result in the inappropriate transfer of a data set out of a specified isolated environment.

Volume checklist

Be sure that the new target volumes are added to the appropriate storage group and all the volumes remain in DISALL status until you are ready to start the migration. The volume status then can be changed to ENABLE.

Be sure that all of the target volumes meet the following requirements:

- Are initialized for SMS (ICKDSF STORAGEGROUP)
- Contain a VVDS large enough to accommodate the target configuration
- Contain an INDEX VTOC of sufficient size
- Have additional 10 cylinders available for the zDMF extent map file

Because the INDEX VTOC is required on an SMS-managed volume, ISMF can be used to identify a disabled or broken INDEX VTOC.

When starting the migration, make sure that all source volumes are set to the DISNEW status.

Note: If the DISNEW status was enabled using an SMS VARY command, the next ACS TRANSLATE and VALIDATE can reverse that status.

Migrating software vendor data sets

The zDMF data set selection criteria and other restrictions are based on specific access methods and data set types. These restrictions are clearly outlined in the zDMF product documentation.

However, you must first consider additional software vendor restrictions on some product data sets. These restrictions are essential to preserve data integrity and product-related functionality. Additionally, before running the zDMF migration, make sure that you have documentation of user exits, product customization, and cross-application dependencies unique to your environment.

zDMF can move any data sets that conform to supported data set lists and restrictions. However, consider additional application restrictions and consult with the person responsible for the application. Also, keep in mind that each installation, although apparently similar, can have subtle differences that might jeopardize the migration.

11.3.8 Post-migration

The creation of a separate user catalog for the zDMF target data sets expedites the post-migration cleanup process. The following are the steps for various post-migration scenarios:

- For a full volume migration (without post migration verification):
  - Delete the aliases that are related to the zDMF user catalog.
  - Delete the zDMF user catalog.
  - Initialize the old source volumes.
- For a selective data set migration:
  - Delete the work data sets using IDCAMS:

    `DELETE (hlq.*)`
  - Delete the data sets using ISPF 3.4 or ISMF, or use the zDMF Monitor function `Generate Source DSN Cleanup JCL` (issue the Z line command on a completed group entry) to delete all work data sets.
  - Delete the aliases related to the zDMF user catalog.
  - Delete the zDMF user catalog.

Tip: After the migration group completes and the verification process is satisfied, delete the migration group from the zDMF database. Use the zDMF Monitor Option 2 (`Interact With Promoted Groups`) to delete the group.

When a migration group is deleted, the zDMF server performs the following tasks:

- Ensures that the migration group is in a state that allows it to be deleted.
- Removes the migration group from internal z/OS storage (memory). Removing the group frees valuable system memory.
- Deletes the migration group from the zDMF database.

As part of the post-migration activities, address the following concerns:

- Create a clean-up schedule to determine how long you want to keep the zDMF target data sets.
- Determine how you want to treat the zDMF data sets:
  - Back up before delete
  - Delete without backup
  - Delete the BCS and Init the volumes
  - Content verification for post-migration audit

11.4 Important zDMF-related topics

This section includes other important concepts that you need to understand in the context of an zDMF migration.

11.4.1 Understanding the catalog structure

Because zDMF is a data-set-level migration tool that includes updating the information in the ICF catalog, you should have a basic understanding of the ICF catalog structure and be able to complete these tasks:

- Perform diagnostic tests for the pre-migration process
- Resolve any catalog structure integrity issues
- Identify the cause of potential migration problems that result from erroneous entries

The ICF catalog records are stored in two components:

- The basic catalog structure (BCS): BCS is considered as the catalog. The BCS is a VSAM key-sequenced data set (KSDS). Its primary function is to point to the volumes on which a data set is located. The BCS is created when either a user or master catalog is defined using access method services (AMS). A BCS does not have to be on the same volume as the data set it references. There can be more than one BCS on a volume.
- The VSAM volume data set (VVDS): The VVDS is considered an extension of the VTOC. The VVDS is a VSAM entry-sequenced data set (ESDS). It contains the information that is required to process VSAM data sets. In a storage management subsystem (SMS) environment, it contains the volume-related information for non-VSAM SMS-managed data sets. There is one VVDS on each DASD volume that contains a VSAM or SMS-managed data set cataloged in an ICF catalog. The VVDS is always on the same volume as the data set it references.

11.4.2 Catalog diagnostic recommendations

BCS, VVDS, and VTOC can be considered metadata required to process a data set. Damage to any piece or a mismatch between them can prevent you from accessing your data even if there is no problem with the data set itself.

Because the zDMF data movement process includes catalog alteration, it is essential to eliminate any structural integrity problem with any involved catalog. Verify that the data sets to be migrated are correctly cataloged on all systems involved in the migration. Use IDCAMS to perform the required diagnostic tests.

Table 11-4 lists IDCAMS commands that can be useful to fix any existing catalog problems on the source volume environment.

Table 11-4 Catalog command and action summary table

Command	Action
DELETE NOSCRATCH	Removes only BCS entries. VVDS and VTOC are not affected.
DELETE TRUENAME	Removes unrelated truename record if the associated cluster record does not exist.
DELETE VVR or DELETE NVR	Removes unrelated VVR or NVR. Also removes VTOC entry, if present.

Command	Action
DEFINE RECATALOG	Creates a BCS entry pointing to existing VVRs or NVR.
DEFINE NONVSAM	Creates a BCS entry pointing to non-system-managed non-VSAM data set.

IDCAMS EXAMINE

To ensure that a BCS and an associated backup are structurally sound, the IDCAMS EXAMINE command should be included in any scheduled catalog backup processes. This command traces failures that are caused by unsynchronized entries in the catalog and reports about its structural integrity. Test both the index and data components of a BCS and identify the amount of free space remaining. Example 11-1 shows a sample JCL to analyze and report on structural errors of a catalogs index component.

Example 11-1 IDCAMS EXAMINE sample JCL

```
//EXAMEX1  JOB
//STEP1    EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//SYSIN    DD *
 EXAMINE NAME(ICFCAT.V338001/MASTRPW) -
   ERRORLIMIT(200)INDEXTEST DATATEST
```

The sample JCL in Example 11-1 produces output similar to what is shown in Figure 11-2.

```
EXAMINE NAME(ICFT.APPL.UCAT1) -
INDEXTEST DATATEST ERRORLIMIT(200)
IDC01700I INDEXTEST BEGINS
1DC01724I INDEXTEST COMPLETE - NO ERRORS DETECTED
IDC01701I DATATEST BEGINS
IDC01709I DATATEST COMPLETE - NO ERRORS DETECTED
IDC01708I 3593 CONTROL INTERVALS ENCOUNTERED
IDC01710I DATA COMPONENT CONTAINS 49216 RECORDS
IDC01711I DATA COMPONENT CONTAINS 8 DELETED CONTROL INTERVALS
IDC01712I MAXIMUM LENGTH DATA RECORD CONTAINS 15290 BYTES
IDC01722I 73 PERCENT FREE SPACE
IDC0001I FUNCTION COMPLETED, HIGHEST CONDITION CODE WAS 0
```

Figure 11-2 IDCAMS EXAMINE sample output

IDCAMS DIAGNOSE

The linkage between VVDS and BCS entries can be checked using the IDCAMS DIAGNOSE command. See Example 11-2 for checking forward pointers of BCS entries to a specified VVDS.

Example 11-2 Checking forward pointers of BCS entries

```
DIAGNOSE ICFCAT INFILE(INCAT) -
   COMPAREDS(SYS1.VVDS.VXXXXXX) ERRORLIMIT(100) NODUMP LIST
```

Example 11-3 shows the command to check the backward pointers of VVDS entries to a specified catalog.

Example 11-3 Check the backward pointers of VVDS entries

```
DIAGNOSE VVDS IDS(SYS1.VVDS.VXXXXXX) -
   COMPAREDD(INCAT) ERRORLIMIT(100) NODUMP LIST
```

11.4.3 Global resource serialization (GRS)

In a GRS complex, programs can serialize access to data sets on shared DASD volumes at the data set level. A program on one system can access a data set on a shared volume while other programs on other systems can still access other data sets on the same volume. Therefore, serialization can reduce contention for these resources and minimize the chance of an interlock occurring between systems.

Reserve handling requirements

zDMF has specific requirements in the area of reserve handling. More specifically, zDMF requires that all hardware reserves for volumes or migration data sets issued by z/OS address spaces must be converted to globally propagated ENQ requests.

When using zDMF to migrate data, the following protections are provided while converting hardware reserves to globally propagated ENQ requests:

- Ensures the data integrity of data for data sets being migrated
- Avoids deadly embraces in multisystem environments

Ensuring data integrity for data sets being migrated

To ensure the data integrity of data sets being migrated, hardware reserves must be converted to globally propagated ENQ requests. In the following examples, data set DS1 is being migrated from source Volume A to target Volume B. The applications are running on the two z/OS images MVSA and MVSB, and both require exclusive serialization of the resource.

- zDMF impact on serialization before diversion

  zDMF phases before the *Divert phase* have no impact on resource serialization. These phases include Group Definition, Activation, Synchronization, and Mirroring. Because all applications are continuing to read and write data from the source location, all serialization is targeted to the source location. For example, data set DS1 on Volume A has all reserves issued to that location across all applications and systems for serialization of the resource. Therefore, all I/O activity to the Mirrored data set or target on Volume B is also serialized.

- zDMF impact on serialization after diversion

  During the zDMF *Divert phase*, hardware reserves are converted to ENQ requests to ensure data integrity of data sets being migrated.

  After entering the *Divert phase*, applications that remain active after diversion will continue to serialize usage of the resource as though it still is on the source volume. Applications starting after diversion are allocated directly to the new target location and do not require any I/O diversion activity by zDMF. These applications do not require alteration because all catalog and volume level metadata are modified to reflect the new location of the data set.

  If you do not convert hardware reserves to ENQ requests, the following situation can cause a data integrity problem:

  – MVSA reserves volume A for data set DS1 that is in diversion, reads data, and prepares to write it back. The data is actually read from volume B because I/Os are diverted to the target data set.

- A new application starts on MVSB and reserves volume B. The application reads data from data set DS1, modifies it, and writes it back to data set DS1.
- During this time, the zDMF channel program for MVSA must wait because Volume B is reserved to MVSB. After MVSB releases the reserve on Volume B, MVSA overwrites what MVSB wrote, and then releases the reserve on Volume A.

If the reserve is converted to a global enqueue, the resource is protected by name and remains secure. This protection remains while all applications protecting any associated resource use the same name and do not rely on the actual device reserve for serialization.

Avoiding deadly embraces in a multi-system environment

While zDMF is mirroring migration data sets, it sometimes must update VTOC and VVDS entries for the target data set. If hardware reserves are not converted for the VTOC and VVDS, deadlock situations can occur.

To avoid this situation, these resources must be converted from hardware reserves to ENQ requests. This conversion allows zDMF to prevent multiple systems from trying to exclusively acquire the same volume resources.

zDMF provides a stand-alone utility called Reserve Monitor. This utility detects hardware reserves and determines whether they are currently being converted to globally propagated ENQ requests. For more information, see “Reserve Monitor” in the *zDMF Installation & Reference Guide*.

11.5 Migration in an experimental shared DB2 test environment

In this example, active allocated data sets in a shared DB2 environment are moved from 3390-3 to 3390-9 volumes. The following steps were completed to prepare the experimental environment:

- DB2 table spaces were allocated on the source volumes and data records were added. All records were checked.
- The DFSMS storage group was prepared:
  - Source volumes were set to DISNEW.
  - New target volumes were set to ENABLE.
- Applications on two systems were started which performed updates to the table spaces during the migration.
- Jobs were created to define and activate zDMF migration groups.
- Migration was monitored using the zDMF Monitor.

After migration, all data were checked for consistency.

11.5.1 Overview of experimental environment

Figure 11-3 shows the experimental environment.

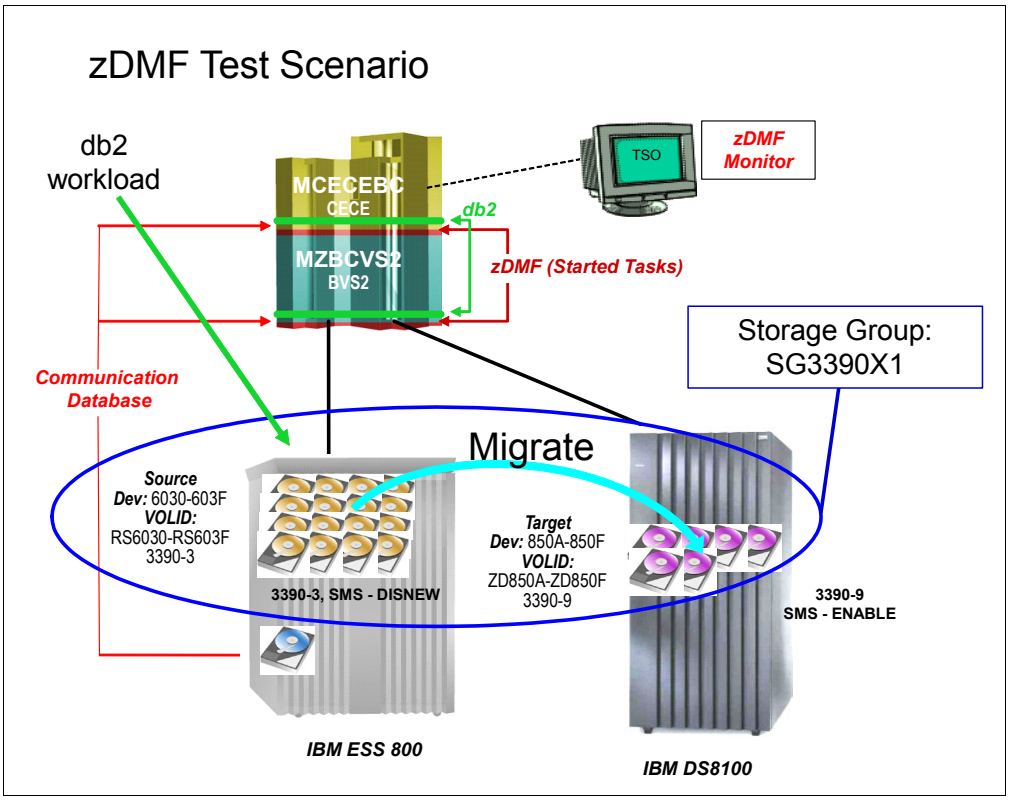

Figure 11-3 zDMF test environment overview

For performance reasons, the zDMF communication database and the user catalog for zDMF work data sets were allocated on volumes that are not managed by DFSMS.

The name and location of the zDMF database can be checked by using the zDMF Monitor as shown in Figure 11-4.

```
 Installation Options              Row 1 to 26 of 26
Command ===>                                                Scroll ===> CSR

Company : IBM Deutschland GMBH
Site  . : IBM MAINZ
Site ID : 11660  Date Initialized  . . . : January 10, 2011
Local Time on Messages . . : Yes

                          Operating Environment

IBM zDMF
zDMF VERSION                  : 3.2.5
zDMF COMMAND PREFIX (CPFX)    : ZD
zDMF DATA BASE DATA SET NAME  : SYS1.HGZD325.DB
zDMF DATA BASE VOLUME         : ZDMFDB
zDMF DB SYSTEM AUTHORIZATION  : ALTER
zDMF SIMULATE DATA SET NAME   : SYS1.HGZD325.SIMULATE.REPORTS
zDMF SIMULATE VOLUME          : VSLNSM
zDMF SIM SYSTEM AUTHORIZATION : ALTER
SMF INTERVAL (# of minutes)   : 5
SMF TYPE                      :
SMF SUBTYPES INCLUDED         :
UNIDENTIFED SYSTEMS ACTION    : IGNORE
    IGNORED SYSTEMS           : NONE DEFINED
USER ID                       : STXXXXX
CPU ID                        : 0004423A 2064
SCP NAME                      : SP7.1.2
SCP FMID                      : HBB7770
ETR ID                        : 00
LOCAL TIME                    : 07/09/2011 00:10:45.09
GMT TIME                      : 07/08/2011 21:59:40.67
LOCAL OFFSET                  : +02:11:04
LEAP SECONDS                  : +000
```

Figure 11-4 zDMF Monitor Display Installation Options

An alias for the work data sets used (with high-level qualifier XXX) was added to the zDMF user catalog.

Source and target volumes were defined to storage group SG3390X1. The source volumes were labeled RS603x, and the target volumes were labeled ZD850x. Figure 11-5 shows the status of each volume in storage group SG3390X1.

```
SDSF ULOG  CONSOLE STXXXXX1                        LINE  COMMAND ISSUED
COMMAND INPUT ===> /D SMS,SG(SG3390X1),LISTVOL                SCROLL ===> CSR
RESPONSE=MCECEBC
 IGD002I 00:34:31 DISPLAY SMS 172

 STORGRP  TYPE    SYSTEM= 1 2
 SG3390X1 POOL            + +

 VOLUME   UNIT    SYSTEM= 1 2                             STORGRP NAME
 RS603A   603A            D D                               SG3390X1
 RS603B   603B            D D                               SG3390X1
 RS603C   603C            D D                               SG3390X1
 RS603D   603D            D D                               SG3390X1
 RS603E   603E            D D                               SG3390X1
 RS603F   603F            D D                               SG3390X1
 RS6030   6030            D D                               SG3390X1
 RS6031   6031            D D                               SG3390X1
 RS6032   6032            D D                               SG3390X1
 RS6033   6033            D D                               SG3390X1
 RS6034   6034            D D                               SG3390X1
 RS6035   6035            D D                               SG3390X1
 RS6036   6036            D D                               SG3390X1
 RS6037   6037            D D                               SG3390X1
 RS6038   6038            D D                               SG3390X1
 RS6039   6039            D D                               SG3390X1
 ZD850A   850A            + +                               SG3390X1
 ZD850B   850B            + +                               SG3390X1
 ZD850C   850C            + +                               SG3390X1
 ZD850D   850D            + +                               SG3390X1
 ZD850E   850E            + +                               SG3390X1
 ZD850F   850F            + +                               SG3390X1
 ***************************** LEGEND ******************************
 . THE STORAGE GROUP OR VOLUME IS NOT DEFINED TO THE SYSTEM
 + THE STORAGE GROUP OR VOLUME IS ENABLED
 - THE STORAGE GROUP OR VOLUME IS DISABLED
 * THE STORAGE GROUP OR VOLUME IS QUIESCED
 D THE STORAGE GROUP OR VOLUME IS DISABLED FOR NEW ALLOCATIONS ONLY
 Q THE STORAGE GROUP OR VOLUME IS QUIESCED FOR NEW ALLOCATIONS ONLY
 > THE VOLSER IN UCB IS DIFFERENT FROM THE VOLSER IN CONFIGURATION
 SYSTEM  1 = MCECEBC    SYSTEM  2 = MZBCVS2
```

Figure 11-5 Displaying all volumes within storage group SG3390X1 with SMS status

Statistical data for the volume allocation status showed that all new volumes were empty (Figure 11-6).

VOLUME SERIAL	DEV TYPE	DEV NUM	MOUNT CHAR	FREE SPACE CYLS	TRKS	EXTS	LARGEST FREE CYLS	TRKS	PCT	PCT FREE	FRAG INDEX
RS603A	33903	603A	SMS X	1020	8	1	1020	8	100%	30%	.000
RS603B	33903	603B	SMS X	1462	14	1	1462	14	100%	43%	.000
RS603C	33903	603C	SMS X	1616	11	1	1616	11	100%	48%	.000
RS603D	33903	603D	SMS X	568	7	1	568	7	100%	17%	.000
RS603E	33903	603E	SMS X	1324	10	1	1324	10	100%	39%	.000
RS603F	33903	603F	SMS X	996	6	1	996	6	100%	29%	.000
RS6030	33903	6030	SMS X	841	14	1	841	14	100%	25%	.000
RS6031	33903	6031	SMS X	696	6	2	696	0	99%	20%	.000
RS6032	33903	6032	SMS X	1764	10	1	1764	10	100%	52%	.000
RS6033	33903	6033	SMS X	1460	3	1	1460	3	100%	43%	.000
RS6034	33903	6034	SMS X	1002	4	1	1002	4	100%	30%	.000
RS6035	33903	6035	SMS X	697	12	1	697	12	100%	20%	.000
RS6036	33903	6036	SMS X	1762	14	1	1762	14	100%	52%	.000
RS6037	33903	6037	SMS X	1168	7	1	1168	7	100%	35%	.000
RS6038	33903	6038	SMS X	302	9	3	297	0	98%	9%	.013
RS6039	33903	6039	SMS X	704	5	1	704	5	100%	21%	.000
ZD850A	**33909**	**850A**	**SMS X**	**9974**	**5**	**3**	**9822**	**0**	**98%**	**99%**	**.007**
ZD850B	**33909**	**850B**	**SMS X**	**9974**	**5**	**3**	**9822**	**0**	**98%**	**99%**	**.007**
ZD850C	**33909**	**850C**	**SMS X**	**9974**	**5**	**3**	**9822**	**0**	**98%**	**99%**	**.007**
ZD850D	**33909**	**850D**	**SMS X**	**9974**	**5**	**3**	**9822**	**0**	**98%**	**99%**	**.007**
ZD850E	**33909**	**850E**	**SMS X**	**9974**	**5**	**3**	**9822**	**0**	**98%**	**99%**	**.007**
ZD850F	**33909**	**850F**	**SMS X**	**9974**	**5**	**3**	**9822**	**0**	**98%**	**99%**	**.007**

Figure 11-6 Space information for Storage Group SG3390X1

11.5.2 Preparing a list of data sets to be migrated

The migration in this example should be performed using the zDMF batch utility, which can be used to define and activate zDMF migration groups with JCL.

To have better control over the migration process, each data set to be migrated should be defined with its fully qualified data set name. Figure 11-7 shows the list of DB2 data sets that were created in preparation for this experimental environment.

```
DSLIST - Data Sets Matching DSNCAT.DSNDBD.DBX1ST01.*              Row 1 of 30
Command ===>                                                  Scroll ===> CSR

Command - Enter "/" to select action                 Message           Volume
-----------------------------------------------------------------------------
         DSNCAT.DSNDBD.DBX1ST01.IXA010A.I0001.A001                     RS603A+
         DSNCAT.DSNDBD.DBX1ST01.IXA010B.I0001.A001                     RS6037+
         DSNCAT.DSNDBD.DBX1ST01.IXA010C.I0001.A001                     RS6031+
         DSNCAT.DSNDBD.DBX1ST01.IXA010D.I0001.A001                     RS603E+
         DSNCAT.DSNDBD.DBX1ST01.IXA010E.I0001.A001                     RS6035+
         DSNCAT.DSNDBD.DBX1ST01.IXA010F.I0001.A001                     RS6030+
         DSNCAT.DSNDBD.DBX1ST01.IXA0101.I0001.A001                     RS603C+
         DSNCAT.DSNDBD.DBX1ST01.IXA0102.I0001.A001                     RS6031+
         DSNCAT.DSNDBD.DBX1ST01.IXA0103.I0001.A001                     RS6030+
         DSNCAT.DSNDBD.DBX1ST01.IXA0104.I0001.A001                     RS6033+
         DSNCAT.DSNDBD.DBX1ST01.IXA0105.I0001.A001                     RS6038+
         DSNCAT.DSNDBD.DBX1ST01.IXA0106.I0001.A001                     RS6037+
         DSNCAT.DSNDBD.DBX1ST01.IXA0107.I0001.A001                     RS603F+
         DSNCAT.DSNDBD.DBX1ST01.IXA0108.I0001.A001                     RS603A+
         DSNCAT.DSNDBD.DBX1ST01.IXA0109.I0001.A001                     RS6033+
         DSNCAT.DSNDBD.DBX1ST01.TSA010A.I0001.A001                     RS603B+
         DSNCAT.DSNDBD.DBX1ST01.TSA010B.I0001.A001                     RS6030+
         DSNCAT.DSNDBD.DBX1ST01.TSA010C.I0001.A001                     RS6032+
         DSNCAT.DSNDBD.DBX1ST01.TSA010D.I0001.A001                     RS6034+
         DSNCAT.DSNDBD.DBX1ST01.TSA010E.I0001.A001                     RS6030+
         DSNCAT.DSNDBD.DBX1ST01.TSA010F.I0001.A001                     RS6032+
         DSNCAT.DSNDBD.DBX1ST01.TSA0101.I0001.A001                     RS6038+
         DSNCAT.DSNDBD.DBX1ST01.TSA0102.I0001.A001                     RS603F+
         DSNCAT.DSNDBD.DBX1ST01.TSA0103.I0001.A001                     RS6037+
         DSNCAT.DSNDBD.DBX1ST01.TSA0104.I0001.A001                     RS6035+
         DSNCAT.DSNDBD.DBX1ST01.TSA0105.I0001.A001                     RS6034+
         DSNCAT.DSNDBD.DBX1ST01.TSA0106.I0001.A001                     RS6031+
         DSNCAT.DSNDBD.DBX1ST01.TSA0107.I0001.A001                     RS6031+
         DSNCAT.DSNDBD.DBX1ST01.TSA0108.I0001.A001                     RS603A+
         DSNCAT.DSNDBD.DBX1ST01.TSA0109.I0001.A001                     RS6035+
```

Figure 11-7 Data set name list for test migration

11.5.3 Creating JCL for the zDMF batch utility

Figure 11-8 on page 296 and Figure 11-9 on page 297 show an example for a job using the zDMF batch utility. Because the DSN subparameter only allows the specification of up to 16 fully qualified data set names (or masks), the data set name list must be split into groups[2] of up to 16 data sets, each starting with a **SET** statement. It might be helpful to use a REXX procedure to create this job out of an input file containing the list of data set names.

[2] Do not confuse these groups with a zDMF migration group. The sample job defines only one zDMF migration group with 2 sets of data set names

The following describes the structure of a job calling the zDMF batch utility and the zDMF parameters used in this test example:

- **`IF GROUP ITS999 EXIST`**: The action following that conditional statement is run in case a group of the name ITS999 exists in the zDMF database. If it does not, the subsequent statement will be skipped.
  - **`COMMAND DELETE ITS999`**: This is the zDMF command that is run if the condition is met.
- **`PROMOTE DLM=ZZ`**: Instructs the zDMF batch utility to interpret the following statements up until the delimiter characters ZZ as a group definition, which is promoted to the zDMF database.
- **`GROUP(ITS999)`**: Specifies the name assigned to the zDMF migration group defined by the subsequent group definition statements.
- **`DELETE_EXISTING_TARGET_DATASETS (NO)`**: Do not delete existing work data sets.
- **`EARLY_DATA_SET_COMPLETION (YES)`**: Divert and complete non-allocated data sets as soon as they have reached a status of MIRROR.
- **`ALLOCSEQ (NONE)`**: Migrate data sets as they appear in the list rather than sorting them by their size.
- **`SET`**: Starts the definition of a collection of source data sets to be migrated with their associated migration parameters.
- **`SPHERE (YES)`**: Move all parts of a VSAM Cluster.
- **`TOLERATE_SOURCE_VALIDATE_FAILURE(YES)`**: Indicates that instead of terminating the entire zDMF migration group processing, to only skip a data set in case an allocation error is encountered for a source data set.
- **`SOURCE(DSN(`*dsnlist*`))`**: In this case, is used to specify a comma-separated list of the fully qualified source data set names.
- **`TOLERATE_TARGET_CREATE_FAILURE(YES)`**: Indicates that instead of terminating the entire zDMF migration group processing, to only skip a data set in case an allocation error is encountered for a target data set.
- **`TARGET(DSN(XXX))`**: Specifies the high-level qualifier to be used when allocating target data sets.
- **`ZZ`**: Marks the end of the group definition.
- **`COMMAND ACTIVATE ITS999`**: Initiates activation of the zDMF migration group named ITS999 (the one that has just been defined).

```
//STRZ999  JOB (A185,SYS),'STXXXXXX IBM',
//           MSGLEVEL=(1,1),NOTIFY=&SYSUID,MSGCLASS=X,CLASS=A
//* IBM user name and telefon number
/*JOBPARM S=CEBC
//BATCH    EXEC PGM=GZDBAT,PARM='CPFX=ZD'
//STEPLIB  DD DISP=SHR,DSN=CPAC.HGZD325.GZDLLIB
//SYSPRINT DD SYSOUT=*
//LDMFBAT  DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN    DD *
IF GROUP ITS999 EXIST
COMMAND DELETE ITS999
PROMOTE DLM=ZZ
GROUP (ITS999) -
  DELETE_EXISTING_TARGET_DATASETS (NO) -
  EARLY_DATA_SET_COMPLETION (YES) -
  ALLOCSEQ (NONE)
SET -
  SPHERE (YES) -
  TOLERATE_SOURCE_VALIDATE_FAILURE (YES) -
  SOURCE ( DSN(                          -
   DSNCAT.DSNDBD.DBX1ST01.IXA010A.I0001.A001    -
   DSNCAT.DSNDBD.DBX1ST01.IXA010B.I0001.A001    -
   DSNCAT.DSNDBD.DBX1ST01.IXA010C.I0001.A001    -
   DSNCAT.DSNDBD.DBX1ST01.IXA010D.I0001.A001    -
   DSNCAT.DSNDBD.DBX1ST01.IXA010E.I0001.A001    -
   DSNCAT.DSNDBD.DBX1ST01.IXA010F.I0001.A001    -
   DSNCAT.DSNDBD.DBX1ST01.IXA0101.I0001.A001    -
   DSNCAT.DSNDBD.DBX1ST01.IXA0102.I0001.A001    -
   DSNCAT.DSNDBD.DBX1ST01.IXA0103.I0001.A001    -
   DSNCAT.DSNDBD.DBX1ST01.IXA0104.I0001.A001    -
   DSNCAT.DSNDBD.DBX1ST01.IXA0105.I0001.A001    -
   DSNCAT.DSNDBD.DBX1ST01.IXA0106.I0001.A001    -
   DSNCAT.DSNDBD.DBX1ST01.IXA0107.I0001.A001    -
   DSNCAT.DSNDBD.DBX1ST01.IXA0108.I0001.A001    -
   DSNCAT.DSNDBD.DBX1ST01.IXA0109.I0001.A001    -
   DSNCAT.DSNDBD.DBX1ST01.TSA010A.I0001.A001    ))-
TOLERATE_TARGET_CREATE_FAILURE (YES) -
  TARGET (DSN (XXX))
             :
```

Figure 11-8 zDMF batch utility job (part 1)

```
            :
SET -
  SPHERE (YES) -
  TOLERATE_SOURCE_VALIDATE_FAILURE (YES) -
  SOURCE ( DSN(                          -
   DSNCAT.DSNDBD.DBX1ST01.TSA010B.I0001.A001    -
   DSNCAT.DSNDBD.DBX1ST01.TSA010C.I0001.A001    -
   DSNCAT.DSNDBD.DBX1ST01.TSA010D.I0001.A001    -
   DSNCAT.DSNDBD.DBX1ST01.TSA010E.I0001.A001    -
   DSNCAT.DSNDBD.DBX1ST01.TSA010F.I0001.A001    -
   DSNCAT.DSNDBD.DBX1ST01.TSA0101.I0001.A001    -
   DSNCAT.DSNDBD.DBX1ST01.TSA0102.I0001.A001    -
   DSNCAT.DSNDBD.DBX1ST01.TSA0103.I0001.A001    -
   DSNCAT.DSNDBD.DBX1ST01.TSA0104.I0001.A001    -
   DSNCAT.DSNDBD.DBX1ST01.TSA0105.I0001.A001    -
   DSNCAT.DSNDBD.DBX1ST01.TSA0106.I0001.A001    -
   DSNCAT.DSNDBD.DBX1ST01.TSA0107.I0001.A001    -
   DSNCAT.DSNDBD.DBX1ST01.TSA0108.I0001.A001    -
   DSNCAT.DSNDBD.DBX1ST01.TSA0109.I0001.A001    ))-
TOLERATE_TARGET_CREATE_FAILURE (YES) -
  TARGET (DSN (XXX))
ZZ
COMMAND ACTIVATE ITS999
```

Figure 11-9 zDMF batch utility job (part 2)

11.5.4 Starting and Monitoring the migration

To start the migration process, the zDMF Servers need to be active on all systems sharing the migration data sets. The zDMF batch utility JCL will be submitted and monitored by using the zDMF Monitor option 2 (*Interact with Promoted Groups*). The migration group is promoted by the batch job and automatically activated. As soon as the group is successfully activated, zDMF starts migrating the data sets.

The batch job will end after sending the data to the zDMF Server and does not wait for the actual migration process.

Figure 11-10 shows the output of the submitted zDMF batch utility job:

```
            :
IEF373I STEP/BATCH   /START 2011193.0005
IEF032I STEP/BATCH   /STOP  2011193.0009
        CPU:     0 HR  00 MIN  00.03 SEC    SRB:     0 HR  00 MIN  00.00 SEC
        VIRT:   240K  SYS:   288K  EXT:    16468K  SYS:    17476K
IEF375I  JOB/STRZ999 /START 2011193.0005
IEF033I  JOB/STRZ999 /STOP  2011193.0009
        CPU:     0 HR  00 MIN  00.03 SEC    SRB:     0 HR  00 MIN  00.00 SEC
DWRQBUF§ 2E104120 LN=00FFFEE0 USED=00000000 REQ=00000000
==> IF GROUP ITS999 EXISTS
==> COMMAND DELETE ITS999
*** COMMAND SKIPPED DUE TO CONDITION TEST

==> PROMOTE DLM=ZZ
    GROUP (ITS999) -
      DELETE_EXISTING_TARGET_DATASETS (NO) -
      EARLY_DATA_SET_COMPLETION (YES) -
      ALLOCSEQ (NONE)
    SET -
      SPHERE (YES) -
      TOLERATE_SOURCE_VALIDATE_FAILURE (YES) -
      MIGRATE_ONLY_SPECIFIED_VOLUMES (NO)    -
      SOURCE ( DSN(                          -
       DSNCAT.DSNDBD.DBX1ST01.IXA010A.I0001.A001    -
       DSNCAT.DSNDBD.DBX1ST01.IXA010B.I0001.A001    -
                  :
       DSNCAT.DSNDBD.DBX1ST01.TSA0108.I0001.A001    -
       DSNCAT.DSNDBD.DBX1ST01.TSA0109.I0001.A001    -
                                                   )-
          )                                         -
TOLERATE_TARGET_CREATE_FAILURE (YES) -
   TARGET (DSN (XXX))
                 :
*GROUP PROMOTED, STOW LENGTH IS 2256
*** RC=0
==> COMMAND ACTIVATE ITS999
GROUP ITS999 ACTIVATED
*** RC=0
*** MAXIMUM RETURN CODE WAS 0
```

Figure 11-10 Output produced by a zDMF batch utility job

By running the zDMF Monitor option 2 (*Interact with Promoted Groups*), you can monitor the progress of the migration process. Repeatedly pressing the Enter key refreshes the screens. After the migration group is promoted and activated by the zDMF batch utility job, the zDMF Server checks each data set specified in the group definition. zDMF also verifies that required space on the new target volumes for each migration data set is available. After the data sets are validated, the groups status changes to `Active` as shown in Figure 11-11.

```
                            Subsystem ZD   Groups             Row 1 to 2 of 2
Command ===>                                                  Scroll ===> CSR

  Group                                                            CEBC
    Data Set                                                      Status
        Extents

_ ITS999    Owning System: CEBC                                   Active
        Group control blocks initializing (00040/00060)
```

Figure 11-11 zDMF Monitor option 2 (Interact with Promoted Groups): Active status

As soon as all data sets (source and work/target data sets) are allocated and required control blocks have been established, zDMF changes the group status to `Pending-Active` as shown in Figure 11-12.

```
                            Subsystem ZD   Groups             Row 1 to 2 of 2
Command ===>                                                  Scroll ===> CSR

  Group                                                            CEBC
    Data Set                                                      Status
        Extents

_ ITS999    Owning System: CEBC                            Moved=  0% P-Active
  Activate report available
```

Figure 11-12 zDMF Monitor option 2 (Interact with Promoted Groups): Pending Active status

The *Copy phase* starts automatically (Figure 11-13).

```
                            Subsystem ZD   Groups            Row 1 to 2 of 2
Command ===>                                                 Scroll ===> CSR

  Group                                                            CEBC
    Data Set                                                      Status
        Extents

_ ITS999    Owning System: CEBC                            Moved= 77% Copy
  Activate report available
```

Figure 11-13 zDMF Monitor Item 2 (Interact with Promoted Groups): Copy phase

When all data is copied to the target volumes, the group status changes to `Mirror` as shown in Figure 11-14.

```
                            Subsystem ZD   Groups            Row 1 to 3 of 3
Command ===>                                                 Scroll ===> CSR

  Group                                                            CEBC
    Data Set                                                      Status
        Extents

_ ITS999    Owning System: CEBC                            Moved=100% Mirror
  Activate report available
  Early Data Set Complete status: 000008 Data sets diverted.
```

Figure 11-14 zDMF Monitor option 2 (Interact with Promoted Groups): Mirror phase

When the group has reached `Mirror` status, zDMF modifies the metadata of currently not allocated (not in use) data sets[3]. Figure 11-15 shows that eight data sets were already migrated from the source volumes to new target volumes.

```
DSLIST - Data Sets Matching DSNCAT.DSNDBD.DBX1ST01                 Row 1 of 30
Command ===>                                                  Scroll ===> CSR

Command - Enter "/" to select action                 Message           Volume
-------------------------------------------------------------------------------
         DSNCAT.DSNDBD.DBX1ST01.IXA010A.I0001.A001                      RS603A+
         DSNCAT.DSNDBD.DBX1ST01.IXA010B.I0001.A001                      RS6037+
         DSNCAT.DSNDBD.DBX1ST01.IXA010C.I0001.A001                      RS6031+
         DSNCAT.DSNDBD.DBX1ST01.IXA010D.I0001.A001                      RS603E+
         DSNCAT.DSNDBD.DBX1ST01.IXA010E.I0001.A001                      RS6035+
         DSNCAT.DSNDBD.DBX1ST01.IXA010F.I0001.A001                      RS6030+
         DSNCAT.DSNDBD.DBX1ST01.IXA0101.I0001.A001                      RS603C+
         DSNCAT.DSNDBD.DBX1ST01.IXA0102.I0001.A001                      RS6031+
         DSNCAT.DSNDBD.DBX1ST01.IXA0103.I0001.A001                      RS6030+
         DSNCAT.DSNDBD.DBX1ST01.IXA0104.I0001.A001                      RS6033+
         DSNCAT.DSNDBD.DBX1ST01.IXA0105.I0001.A001                      RS6038+
         DSNCAT.DSNDBD.DBX1ST01.IXA0106.I0001.A001                      RS6037+
         DSNCAT.DSNDBD.DBX1ST01.IXA0107.I0001.A001                      RS603F+
         DSNCAT.DSNDBD.DBX1ST01.IXA0108.I0001.A001                      RS603A+
         DSNCAT.DSNDBD.DBX1ST01.IXA0109.I0001.A001                      RS6033+
         DSNCAT.DSNDBD.DBX1ST01.TSA010A.I0001.A001                      RS603B+
         DSNCAT.DSNDBD.DBX1ST01.TSA010B.I0001.A001                      RS6030+
         DSNCAT.DSNDBD.DBX1ST01.TSA010C.I0001.A001                      RS6032+
         DSNCAT.DSNDBD.DBX1ST01.TSA010D.I0001.A001                      RS6034+
         DSNCAT.DSNDBD.DBX1ST01.TSA010E.I0001.A001                      RS6030+
         DSNCAT.DSNDBD.DBX1ST01.TSA010F.I0001.A001                      RS6032+
         DSNCAT.DSNDBD.DBX1ST01.TSA0101.I0001.A001                      ZD850C+
         DSNCAT.DSNDBD.DBX1ST01.TSA0102.I0001.A001                      ZD850E+
         DSNCAT.DSNDBD.DBX1ST01.TSA0103.I0001.A001                      ZD850F+
         DSNCAT.DSNDBD.DBX1ST01.TSA0104.I0001.A001                      ZD850F+
         DSNCAT.DSNDBD.DBX1ST01.TSA0105.I0001.A001                      ZD850E+
         DSNCAT.DSNDBD.DBX1ST01.TSA0106.I0001.A001                      ZD850D+
         DSNCAT.DSNDBD.DBX1ST01.TSA0107.I0001.A001                      ZD850B+
         DSNCAT.DSNDBD.DBX1ST01.TSA0108.I0001.A001                      ZD850C+
         DSNCAT.DSNDBD.DBX1ST01.TSA0109.I0001.A001                      RS6035+
***************************** End of Data Set list *****************************
```

Figure 11-15 Data set status (source data sets)

[3] EARLY_DATA_SET_COMPLETION(YES) was specified in the group definition.

11.5.5 Diverting the Group

After some time, 16 data sets were completely migrated. The remaining data sets were allocated and used by DB2 and stay in a `Mirror` status. To divert I/Os to the target volumes instead of mirroring the **V** line, command is issued for the migration group as shown in Figure 11-16.

```
                            Subsystem ZD   Groups            Row 1 to 3 of 3
Command ===>                                                 Scroll ===> CSR

  Group                                                            CEBC
    Data Set                                                      Status
        Extents

v ITS999    Owning System: CEBC                            Moved=100% Mirror
  Activate report available
  Early Data Set Complete status: 000016 Data sets diverted.
```

Figure 11-16 zDMF Monitor option 2 (Interact with Promoted Groups): Start diversion

After confirming the command, the `Divert Group Options` panel shows up that enables you to schedule the diversion for a specific date and time (Figure 11-17) or by pressing Enter to start the diversion.

```
 ------------------------ Divert Group Options -----------------------

   _   Retry.
   _   Schedule - Date: YY/MM/DD Time: HH:MM
       (Type over date and time)

   Select 'S' divert types for this group and press Enter to process.
   If Schedule is not selected, divert processing will begin at once.
```

Figure 11-17 Scheduling the diversion

After repeatedly refreshing the panel by pressing Enter, you can see the status changing to `P-Divert` (pending diversion) with an indication of its progress as shown in Figure 11-18.

```
                            Subsystem ZD   Groups             Row 1 to 2 of 2
Command ===>                                                  Scroll ===> CSR

  Group                                                                CEBC
    Data Set                                                          Status
        Extents

_ ITS999    Owning System: CEBC                          Diverted= 25% P-Divert
  Activate report available
```

Figure 11-18 zDMF Monitor option 2 (Interact with Promoted Groups): Pending diversion

More details can be displayed by modifying the display options (press the F4 key and modify the first three entries as shown in Figure 11-19).

```
                                Display Options
Command ===>                                                  Scroll ===> CSR

Show Source Data Set Name  . . . . . . Y or N  . . . y       <--
Show Target Data Set Name  . . . . . . Y or N  . . . y       <--
Show Extent Information  . . . . . . . Y or N  . . . y       <--
Show Command Messages  . . . . . . . . Y or N  . . . N
Show Command Diagnostic Information  . Y or N  . . . N

                     F1=Help     F5=Save Settings     F3=Exit
```

Figure 11-19 zDMF Monitor (modify display options)

Save your new settings by pressing the F5 key.

After leaving the *Display Options* panel by pressing F3, the details as shown in Figure 11-20 are displayed. Cylinder and head information together with track number and progress percentage for each extent of each data set (source and target) is displayed.

```
                              Subsystem ZD   Groups          Row 1 to 35 of 372
Command ===>                                                     Scroll ===> CSR

  Group                                                                  CEBC
    Data Set                                                            Status
        Extents

_ ITS999    Owning System: CEBC                                   n/a  Divert
  Activate and Divert reports available
_   Src:DSNCAT.DSNDBC.DBX1ST01.IXA010A.I0001.A001
        VSAM Cluster
    Tar:XXX.ITS099.D1193.T0020546.S00001
        VSAM Cluster
_   Src:DSNCAT.DSNDBC.DBX1ST01.IXA010B.I0001.A001
        VSAM Cluster
    Tar:XXX.ITS099.D1193.T0020546.S00002
        VSAM Cluster
             :
_   Src:DSNCAT.DSNDBD.DBX1ST01.TSA010F.I0001.A001
        Unit=6032(RS6032) Cyl=000004D Hd=B                              Divert
        Unit=6032(RS6032) Cyl=00002D9 Hd=5                              Divert
        Unit=6032(RS6032) Cyl=00003D7 Hd=5                              Divert
        Unit=6032(RS6032) Cyl=00004FD Hd=5                              Divert
    Tar:XXX.ITS099.D1193.T0020546.S00021.DATA
        Unit=850C(ZD850C) Cyl=000015B Hd=1 #Trks=00000976   Moved=100%
        Unit=850C(ZD850C) Cyl=000019C Hd=2 #Trks=00000375   Moved=100%
        Unit=850C(ZD850C) Cyl=00001B5 Hd=2 #Trks=00000435   Moved=100%
        Unit=850C(ZD850C) Cyl=00001D2 Hd=2 #Trks=00000495   Moved=100%
             :
_   Src:DSNCAT.DSNDBD.DBX1ST01.TSA0101.I0001.A001
        Unit=6038(RS6038) Cyl=000000C Hd=0                              Complete
        Unit=6038(RS6038) Cyl=000000C Hd=B                              Complete
        Unit=6038(RS6038) Cyl=000038A Hd=0                              Complete
        Unit=6038(RS6038) Cyl=00008BC Hd=0                              Complete
        Unit=6038(RS6038) Cyl=0000A35 Hd=0                              Complete
    Tar:XXX.ITS099.D1193.T0020546.S00022.DATA
        Unit=850C(ZD850C) Cyl=00001F3 Hd=2 #Trks=00000001   Moved=100%
        Unit=850C(ZD850C) Cyl=00001F3 Hd=3 #Trks=00000975   Moved=100%
        Unit=850C(ZD850C) Cyl=0000234 Hd=3 #Trks=00000375   Moved=100%
        Unit=850C(ZD850C) Cyl=000024D Hd=3 #Trks=00000435   Moved=100%
        Unit=850C(ZD850C) Cyl=000026A Hd=3 #Trks=00000495   Moved=100%
_   Src:DSNCAT.DSNDBD.DBX1ST01.TSA0102.I0001.A001
        Unit=603F(RS603F) Cyl=000000C Hd=0                              Complete
        Unit=603F(RS603F) Cyl=0000017 Hd=0                              Complete
        Unit=603F(RS603F) Cyl=000040D Hd=9                              Complete
        Unit=603F(RS603F) Cyl=0000584 Hd=9                              Complete
        Unit=603F(RS603F) Cyl=0000737 Hd=9                              Complete
             :
```

Figure 11-20 zDMF Monitor Detailed extent list of an active group

To see data sets that are allocated by an application, issue the **J** line command for the group. The system SMFID and the application names allocating the data sets are displayed (Figure 11-21).

```
 Allocated Data Sets for Group ITSO99       Row 1 to 37 of 37
Command ===>                                                  Scroll ===> CSR
Data Set Name                              SMFID  -------- Job Names -------
DSNCAT.DSNDBC.DBX1ST01.IXA010A.I0001.A001   CVS2  DB2ADBM1
                                            CEBC  DB1ADBM1
DSNCAT.DSNDBC.DBX1ST01.IXA010B.I0001.A001   CVS2  DB2ADBM1
                                            CEBC  DB1ADBM1
DSNCAT.DSNDBC.DBX1ST01.IXA010C.I0001.A001   CVS2  DB2ADBM1
                                            CEBC  DB1ADBM1
DSNCAT.DSNDBC.DBX1ST01.IXA010D.I0001.A001   CVS2  DB2ADBM1
                                            CEBC  DB1ADBM1
DSNCAT.DSNDBC.DBX1ST01.IXA010E.I0001.A001   CVS2  DB2ADBM1
                                            CEBC  DB1ADBM1
DSNCAT.DSNDBC.DBX1ST01.IXA010F.I0001.A001   CVS2  DB2ADBM1
                                            CEBC  DB1ADBM1
DSNCAT.DSNDBC.DBX1ST01.IXA0101.I0001.A001   CVS2  DB2ADBM1
                                            CEBC  DB1ADBM1
DSNCAT.DSNDBC.DBX1ST01.IXA0102.I0001.A001   CVS2  DB2ADBM1
                                            CEBC  DB1ADBM1
DSNCAT.DSNDBC.DBX1ST01.IXA0103.I0001.A001   CVS2  DB2ADBM1
                                            CEBC  DB1ADBM1
DSNCAT.DSNDBC.DBX1ST01.IXA0104.I0001.A001   CVS2  DB2ADBM1
                                            CEBC  DB1ADBM1
DSNCAT.DSNDBC.DBX1ST01.IXA0105.I0001.A001   CVS2  DB2ADBM1
                                            CEBC  DB1ADBM1
DSNCAT.DSNDBC.DBX1ST01.IXA0106.I0001.A001   CVS2  DB2ADBM1
                                            CEBC  DB1ADBM1
DSNCAT.DSNDBC.DBX1ST01.IXA0107.I0001.A001   CVS2  DB2ADBM1
                                            CEBC  DB1ADBM1
DSNCAT.DSNDBC.DBX1ST01.IXA0108.I0001.A001   CVS2  DB2ADBM1
                                            CEBC  DB1ADBM1
DSNCAT.DSNDBC.DBX1ST01.IXA0109.I0001.A001   CVS2  DB2ADBM1
                                            CEBC  DB1ADBM1
DSNCAT.DSNDBC.DBX1ST01.TSA010A.I0001.A001   CEBC  DB1ADBM1
DSNCAT.DSNDBC.DBX1ST01.TSA010B.I0001.A001   CEBC  DB1ADBM1
DSNCAT.DSNDBC.DBX1ST01.TSA010C.I0001.A001   CEBC  DB1ADBM1
DSNCAT.DSNDBC.DBX1ST01.TSA010D.I0001.A001   CEBC  DB1ADBM1
DSNCAT.DSNDBC.DBX1ST01.TSA010E.I0001.A001   CEBC  DB1ADBM1
DSNCAT.DSNDBC.DBX1ST01.TSA010F.I0001.A001   CEBC  DB1ADBM1
DSNCAT.DSNDBC.DBX1ST02.IXA020C.I0001.A001   CEBC  DB1ADBM1
```

Figure 11-21 zDMF detailed data sets allocation list (zDMF Monitor)

The auto-completion feature of zDMF and DB2 interact with Media Manager services. The migration of many DB2 data sets, such as table spaces, can automatically complete within 24-25 hours. However, certain types of DB2 data sets like the BSDS and log data sets most likely require a shutdown of the DB2 system to complete. The zDMF migration group can also be completed earlier by closing individual DB2 table spaces or by shutting down appropriate DB2 systems. It only takes a short time for zDMF to complete the migration group and DB2 can be brought up again or reopen the affected table spaces.

Usually the DB2 and associated application outages are scheduled to allow the completion of the migration. This outage might be required to release all allocations to migration data sets. While a data set within a zDMF migration group is allocated (in use), the migration process cannot complete. Before restarting affected applications and corresponding DB2 systems, check with the **J** line command that data sets are no longer allocated for a migration group as shown in Figure 11-22.

```
                 Allocated Data Sets for Group ITS999         Row 1 to 4 of 4
Command ===>                                                  Scroll ===> CSR

Data Set Name                              SMFID  -------- Job Names -------

No jobs allocated to any data sets
included in this group.
```

Figure 11-22 zDMF Monitor data set allocation status after application/DB2 was stopped

In addition, make sure that the migration process is completed and the migration group has reached the `Complete` status as shown in Figure 11-23.

```
                            Subsystem ZD   Groups            Row 1 to 2 of 2
Command ===>                                                 Scroll ===> CSR

  Group                                                            CEBC
    Data Set                                                      Status
        Extents

_ ITS999                                                         Complete
  Activate, Divert, and Terminate reports available
```

Figure 11-23 zDMF Monitor: Status of group is Complete

After the group reaches the `Complete` status, check the new location of all migrated data sets as shown in Figure 11-24.

```
DSLIST - Data Sets Matching DSNCAT.DSNDBD.DBX1ST01.*              Row 1 of 30
Command ===>                                                 Scroll ===> CSR

Command - Enter "/" to select action                 Message          Volume
-------------------------------------------------------------------------------
         DSNCAT.DSNDBD.DBX1ST01.IXA0101.I0001.A001                     ZD850F+
         DSNCAT.DSNDBD.DBX1ST01.IXA0102.I0001.A001                     ZD850F+
                         :
         DSNCAT.DSNDBD.DBX1ST01.IXA010E.I0001.A001                     ZD850D+
         DSNCAT.DSNDBD.DBX1ST01.IXA010F.I0001.A001                     ZD850C+
         DSNCAT.DSNDBD.DBX1ST01.TSA0101.I0001.A001                     ZD850C+
         DSNCAT.DSNDBD.DBX1ST01.TSA0102.I0001.A001                     ZD850E+
                         :
         DSNCAT.DSNDBD.DBX1ST01.TSA010E.I0001.A001                     ZD850D+
         DSNCAT.DSNDBD.DBX1ST01.TSA010F.I0001.A001                     ZD850C+

***************************** End of Data Set list ****************************
```

Figure 11-24 Location of source data sets

Use available DB2 utilities to check data consistency of the migrated data sets (table spaces).

11.5.6 Performing post migration cleanup

To perform a post migration cleanup processing, issue the **Z** line command for the migration group as shown in Figure 11-25.

```
                              Subsystem ZD   Groups            Row 2 to 3 of 3
Command ===>                                                  Scroll ===> CSR

  Group                                                                CEBC
    Data Set                                                          Status
        Extents

z ITS999                                                             Complete
  Activate, Divert, and Terminate reports available
```

Figure 11-25 Start Cleanup procedure with Z line command

This creates a job that deletes the original source data sets that have now been assigned work data set names according to the high-level qualifier specified in the group definition.

As a security measure, it is also possible to generate an ICKDSF step to completely erase the content of the original source data sets. This step is generated if the user session option `Create ICKDSF TRKFMT Statements` is specified as **Y** as shown in Figure 11-26.

```
                              User Session Options
Command ===>                                                 07/17/11 01:25:09

zDMF Command Prefix  . . . . . . ZD
Group Definition Data Set  . . . STXXXXX.ZDMF.CONFIG
Use Log Data Set . . . . . . . . N (Y/N) Retain 07 generations of logs (01-30)
  Use Browse or View . . . . . . V (B/V)
  Log Data Set Name  . . . . . . ____________________________________________
Messages with Local or GMT time. L (L/G)
Show Command Messages  . . . . . N (Y/N)
Show Command Diagnostic Info . . N (Y/N)
Create ICKDSF TRKFMT Statements. Y (Y/N) <-- "Y"
Early Data Set Completion  . . . Y (Y/N)

zDMF Load Library  . . . . . . . CPAC.HGZD325.GZDLLIB
zDMF Rexx Library  . . . . . . . CPAC.HGZD325.GZDELIB
zDMF Panel Library . . . . . . . CPAC.HGZD325.GZDPLIB
zDMF Table Library . . . . . . . CPAC.HGZD325.GZDTLIB
zDMF Message Library . . . . . . CPAC.HGZD325.GZDMLIB
zDMF Security Library  . . . . . CPAC.HGZD325.GZDLLIB

            F1=Help     F3=Exit     F5=Save Settings     F12 = Cancel
```

Figure 11-26 zDMF User Session Options (zDMF Monitor option 6) for ICKDSF TRKFMT

Figure 11-27 on page 308 shows the first ICKDSF steps of a generated cleanup job when `Create ICKDSF TRKFMT Statements` has been specified as **Y**.

```
VIEW       SYS11198.T013218.RA000.STXXXXX.R0147420         Columns 00001 00072
Command ===>                                                  Scroll ===> PAGE
****** ***************************** Top of Data ****************************
000001 Jobcard1              Have sufficient region size
000002 Jobcard2                   to run ICKDSF
000003 //*
000004 //*  You may submit the jcl from this panel after you
000005 //*  update the job card(s).  To save a copy, you must use
000006 //*  the CREATE or REPLACE primary command to place the
000007 //*  data in a data set of your choosing.
000008 //*  Reference the ISPF User's Guide on the use of VIEW.
000009 //*
000010 //*  All statements should be reviewed prior to submitting.
000011 //*
000013 //*  User MUST have authority to run ICKDSF on volumes.
000014 //*
000015 //*  Run ALL the ICKDSF job steps before the data sets are
000016 //*  deleted, else the extent information may not be valid.
000017 //*
000018 //*  ICKDSF statements for Data Sets successfully migrated.
000019 //*
000020 //S00001   EXEC PGM=ICKDSF
000021 //D00001   DD   DISP=OLD,DSN=XXX.ITS999.D1193.T0020546.S00022
000022 //LDRS6038 DD   DISP=OLD,UNIT=3390,VOL=SER=RS6038
000023 //SYSPRINT DD   SYSOUT=*
000024 //SYSIN    DD   *
000025
000026 /* Data Set Name: XXX.ITS999.D1193.T0020546.S00022.DATA            */
000027 TRKFMT DDNAME(LDRS6038) VERIFY(RS6038) ERASEDATA CYCLES(3) -
000028        FROMRANGE(00000012,00) TORANGE(00000012,00)
000029 TRKFMT DDNAME(LDRS6038) VERIFY(RS6038) ERASEDATA CYCLES(3) -
000030        FROMRANGE(00000012,11) TORANGE(00000077,10)
000031 TRKFMT DDNAME(LDRS6038) VERIFY(RS6038) ERASEDATA CYCLES(3) -
000032        FROMRANGE(00000906,00) TORANGE(00000930,14)
000033 TRKFMT DDNAME(LDRS6038) VERIFY(RS6038) ERASEDATA CYCLES(3) -
000034        FROMRANGE(00002236,00) TORANGE(00002264,14)
000035 TRKFMT DDNAME(LDRS6038) VERIFY(RS6038) ERASEDATA CYCLES(3) -
000036        FROMRANGE(00002613,00) TORANGE(00002645,14)
000037 /*
000038 //S00002   EXEC PGM=ICKDSF,COND=(0,NE)
000039 //D00002   DD   DISP=OLD,DSN=XXX.ITS999.D1193.T0020546.S00023
000040 //LDRS603F DD   DISP=OLD,UNIT=3390,VOL=SER=RS603F
000041 //SYSPRINT DD   SYSOUT=*
000042 //SYSIN    DD   *
000043
000044 /* Data Set Name: XXX.ITS999.D1193.T0020546.S00023.DATA            */
000045 TRKFMT DDNAME(LDRS603F) VERIFY(RS603F) ERASEDATA CYCLES(3) -
000046        FROMRANGE(00000012,00) TORANGE(00000012,00)
000047 TRKFMT DDNAME(LDRS603F) VERIFY(RS603F) ERASEDATA CYCLES(3) -
000048        FROMRANGE(00000023,00) TORANGE(00000087,14)
                :
```

Figure 11-27 Generated cleanup JCL (including ICKDSF TRKFMT steps)

Figure 11-28 shows the generated steps to delete the original source data sets and extent map data sets. These steps are generated regardless of the setting of the `Create ICKDSF TRKFMT Statements` session option.

```
VIEW       SYS11198.T013954.RA000.STXXXXX.R0147421         Columns 00001 00072
Command ===>                                                  Scroll ===> PAGE
****** ***************************** Top of Data ******************************
                :
000115 //DELETES  EXEC PGM=IDCAMS
000116 //SYSPRINT DD  SYSOUT=*
000117 //SYSIN    DD  *
000118
000119   /* GROUP NAME: ITS999                                                -
000120      DELETE STATEMENTS FOR DATA SETS SUCCESSFULLY MIGRATED.          */
000121
000122   DELETE 'XXX.ITS999.D1193.T0020546.S00001'             CLUSTER
000123   DELETE 'XXX.ITS999.D1193.T0020546.S00002'             CLUSTER
000124   DELETE 'XXX.ITS999.D1193.T0020546.S00003'             CLUSTER
000146   DELETE 'XXX.ITS999.D1193.T0020546.S00004'             CLUSTER
000147   DELETE 'XXX.ITS999.D1193.T0020546.S00005'             CLUSTER
000148   DELETE 'XXX.ITS999.D1193.T0020546.S00006'             CLUSTER
                :
000159   DELETE 'XXX.ITS999.D1193.T0020546.S00027'             CLUSTER
000160   DELETE 'XXX.ITS999.D1193.T0020546.S00028'             CLUSTER
000161   DELETE 'XXX.ITS999.D1193.T0020546.S00029'             CLUSTER
000162   DELETE 'XXX.ITS999.D1193.T0020546.S00030'             CLUSTER
000163
000164   /* GROUP NAME: ITS999   END                                        */
000165
000166 /*
000167 //EXZD850C EXEC PGM=IDCAMS,DYNAMNBR=1
000168 //SYSPRINT DD   SYSOUT=*
000169 //SYSIN    DD   *
000170
000171   /* ATTEMPT TARGET EXTENT MAP DATA SET DELETE ONLY IF               */
000172   /* NOT ALLOCATED AND STILL PRESENT ON VOLUME.                      */
000173   SET LASTCC=0
000174   ALLOCATE                                                          -
000175      DSNAME('SYS1.HGZD325.EXTMAP.ZDZD850C.DB') OLD
000176   IF LASTCC EQ 0                                                    -
000177      THEN                                                           -
000178      DELETE 'SYS1.HGZD325.EXTMAP.ZDZD850C.DB'
000179 /*
```

Figure 11-28 Clean up with data set IDCAMS delete statements for the work data sets

Note: Only delete extent map data sets if all zDMF migration groups have reached the `Complete` status. If other groups are still active and you want to clean up completed groups, the steps deleting the extent map data sets should be removed from the generate job.

Do not forget to export delete the user catalog for the work data sets and delete all work data set aliases.

Related publications

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

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 DS8000 Copy Services for IBM System z*, SG24-6787
- *IBM DS8000 Copy Services for Open Systems*, SG24-6788
- *IBM DS8000 Multiple Target Peer-to-Peer Remote Copy*, REDP-5151
- *IBM System Storage DS8000: Remote Pair FlashCopy (Preserve Mirror)*, REDP-4504
- *IBM System Storage DS8000 Series: IBM FlashCopy SE*, REDP-4368

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:

- *DS8000 Command-Line Interface User's Guide,* GC27-4212
- *z/OS DFSMS Advanced Copy Services*, SC23-6847

Online resources

These websites are also relevant as further information sources:

- Documentation for the DS8000:

  https://www-01.ibm.com/support/knowledgecenter/HW213_7.4.0/com.ibm.storage.ssic.help.doc/f2c_ichomepage_v7.40.html

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

© Copyright IBM Corp. 2015. All rights reserved.

Redbooks

DS8870 Data Migration Techniques

SG24-8257-00

ISBN 0738440604

IBM

(0.5” spine)
0.475”<->0.873”
250 <-> 459 pages

SG24-8257-00

ISBN 0738440604

Printed in U.S.A.

Get connected

Redbooks
ibm.com/redbooks