Redbooks
ibm.com/redbooks

IBM Copy Services Manager Implementation Guide

William Rooney
Tariq Hanif
Gabriel Dragan
Dinu Stefan
Catalin Dumitriu
Robert Haimowitz
Octavian Lascu

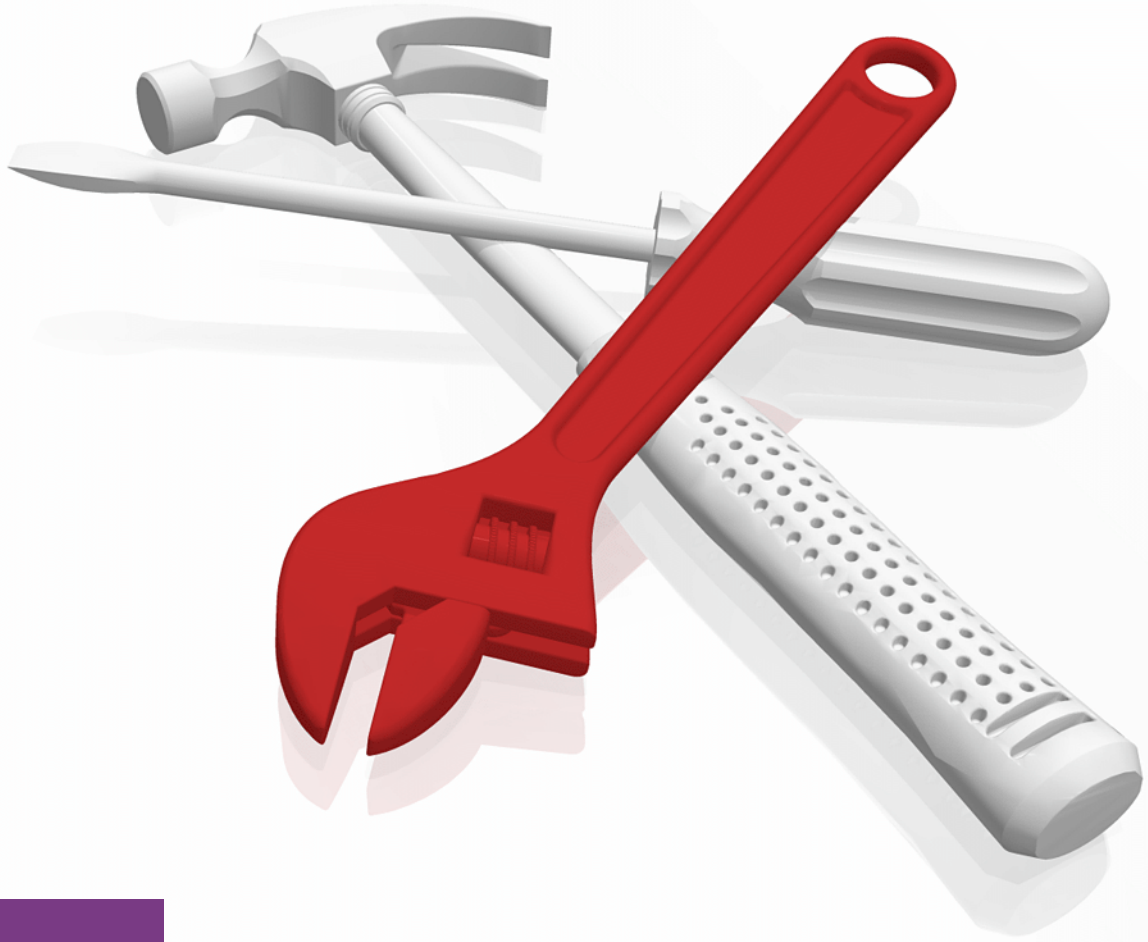

Storage

IBM

Redbooks

International Technical Support Organization

IBM Copy Services Manager Implementation Guide

September 2017

SG24-8375-00

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

First Edition (September 2017)

This edition applies to Version 6, Release 1, Modification 4 of IBM Copy Services Manager (product number 5725-Z54).

Contents

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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, MD-NC119, Armonk, NY 10504-1785, US

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 jurisdictions 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 provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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.

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

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 actual people or business enterprises 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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

AIX®
DS6000™
DS8000®
Enterprise Storage Server®
FlashCopy®
GDPS®
HACMP™
HyperSwap®
IBM®
IBM FlashSystem®
IBM SmartCloud®
IBM Spectrum™
IBM Spectrum Accelerate™
IBM Spectrum Virtualize™
IBM z®
IBM z Systems®
Parallel Sysplex®
Passport Advantage®
Power Systems™
PowerHA®
PowerPC®
RACF®
Redbooks®
Redpaper™
Redbooks (logo)  ®
Storwize®
System Storage®
SystemMirror®
Tivoli®
WebSphere®
XIV®
z Systems®
z/OS®
z/VM®

The following terms are trademarks of other companies:

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

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

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

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

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

Preface

This IBM® Redbooks® publication provides an overview of IBM Copy Services Manager (CSM) for IBM Z and open systems, and documents a set of scenarios for using IBM Copy Services manager to automate and manage replication tasks based on IBM Storage.

This book reviews and explains the usage of copy services functions and describes how these functions are implemented in IBM Copy Services Manager. IBM Copy Services Manager key concepts, architecture, session types and usage, and new functionality as of IBM Copy Services Manager version 6.1 are also described.

Authors

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

William Rooney is a Senior Technical Staff Member in the IBM Z Operating Systems Development organization in Poughkeepsie, New York. He has over 35 years of experience with IBM Z and z/OS®. He is the Software Architect for IBM HyperSwap® for z/OS. His areas of expertise include I/O configuration, I/O qualities of service, and storage replication.

Tariq Hanif is a Senior Software Engineer in the z/OS System Verification Test group. He received his PhD degree from Kings College, University of London, UK in 1995. During his research at Kings, he worked on algorithm-based architecture, signal processing, and parallel processing architecture for image-processing algorithms. After his PhD, he moved to Canada where he worked for Royal Bank of Canada as a Software Analyst. In 2000, he joined Nortel Networks Inc. in the US as a Software Engineer, where his focus was to develop efficient embedded signal processing software for VoIP gateways. He has several publications related to Specialized Computer Architecture for Computer Vision, and Parallel Processing. In 2003, he joined the IBM z/OS System Verification group where he tests the z/OS operating system, with a focus on z/OS HyperSwap and IBM Tivoli® Storage Productivity Center for Replication for IBM Z. In z/OS, his areas of interest are disaster recovery, IBM GDPS®, and IBM Parallel Sysplex®.

Gabriel Dragan is a Storage Consultant in IBM System Lab Services in Europe. He has more than 15 years of experience with IBM storage solutions, specialized in the Enterprise Storage area. He holds a degree in IT&C. His areas of expertise also include storage area network (SAN) solutions, storage replication, and disaster recovery solutions.

Dinu Stefan is a Senior IBM i Administrator from Romania. He has more than 20 years of experience in IBM i and the IBM Power Systems™ field. He holds a degree in Electronic Engineering from the Polytechnical Institute in Bucharest. His areas of expertise also include IBM i replication and disaster recovery, IBM Storage, and IBM i Operating System performance tuning.

Catalin Dumitriu is a Senior IBM AIX® System Administrator from Romania. He has more than 10 years of experience in AIX and IBM Power Systems. He holds a degree in Electronic Engineering from the Polytechnical Institute in Bucharest. He has worked extensively with IBM Power Systems, AIX, LPAR, and IBM Storage Systems. His areas of expertise also include IBM HACMP™, IBM Tivoli Storage Manager, and AIX performance tuning and disaster recovery.

Robert Haimowitz is a system programmer with the ITSO with over 38 years of large systems experience with IBM in IBM Z, z/OS, and z/VM®.

Octavian Lascu is a Senior IT Consultant for IBM Romania with over 25 years of experience. He specializes in designing, implementing, and supporting complex IT infrastructure environments (systems, storage, and networking), including high availability and disaster recovery solutions and high-performance computing deployments. He has developed materials for and taught workshops for technical audiences around the world for the past 17 years. He has written several IBM publications.

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

William G White
International Technical Support Organization, Poughkeepsie Center

Randy Blea
IBM Tucson

Bertrand Dufrasne
IBM Mountain View

Andrei Socoliuc
IBM Romania

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 2 - 6 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:

  ibm.com/redbooks

- Send your comments in an email:

  redbooks@us.ibm.com

- Mail your comments:

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

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

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

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

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

- Stay current on recent Redbooks publications with RSS Feeds:

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

1

IBM Copy Services Manager introduction

This chapter provides an overview of IBM Copy Services Manager for IBM Z and open systems. It reviews copy services functions and describes how these functions are implemented in IBM Copy Services Manager (CSM).

We provide an overview for IBM Copy Services Manager key concepts, architecture, session types and usage, and new functionality as of IBM Copy Services Manager version 6.1.

We also introduce IBM Storage Systems that are supported by IBM Copy Services Manager.

The following topics are described in this chapter:

- Overview
- New features in CSM 6.1
- Terminology
- Architecture
- Managed storage systems and session types

1.1 Overview

IBM Copy Services Manager (formerly *IBM Tivoli Storage Productivity Center for Replication*, a component of IBM Tivoli Storage Productivity Center and IBM SmartCloud® Virtual Storage Center) manages copy services in IBM storage environments. Copy services are features that are used by storage systems to configure, manage, and monitor data replication functions. These copy services include IBM FlashCopy®, Metro Mirror, Global Mirror, and Metro Global Mirror data replication.

Figure 1-1 on page 3 depicts an overview of the IBM Copy Services Manager environment.

IBM Copy Services Manager manages copy services for the following storage systems:[1,2]

- IBM System Storage® DS8000® series
- IBM Spectrum Storage™:
  - IBM Spectrum Virtualize™:
    - IBM SAN Volume Controller
    - IBM Storwize® Family:[3]

      - IBM Storwize V5000 and V5000F
      - IBM Storwize V7000 and V7000F
      - IBM Storwize V7000 Unified

    - IBM FlashSystem® V9000
  - IBM Spectrum Accelerate™:
    - IBM FlashSystem A9000/A9000R[4]
    - IBM XIV® System Storage[2]

IBM Copy Services Manager automates key replication management tasks to help you improve the efficiency of your storage replication. You can use a simple GUI to configure, automate, manage, and monitor all important data replication tasks in your environment, including the following tasks:

- Manage and monitor multisite environments to meet disaster recovery (DR) requirements
- Automate the administration and configuration of data replication features
- Keep data on multiple related volumes consistent across storage systems in a planned or unplanned outage
- Recover to a remote site to reduce downtime of critical applications
- Provide high availability (HA) for applications by using IBM HyperSwap technology
- Practice recovery processes while disaster recovery capabilities are maintained

[1] IBM Copy Services Manager might work with older IBM Storage Systems also. However, we do not list these products anymore because they are at end of service from IBM.
[2] Always check the latest information for end of service information for your IBM Storage (hardware and software).
[3] See the Storwize data sheet.
[4] IBM Copy Services Manager 6.1.4 or later.

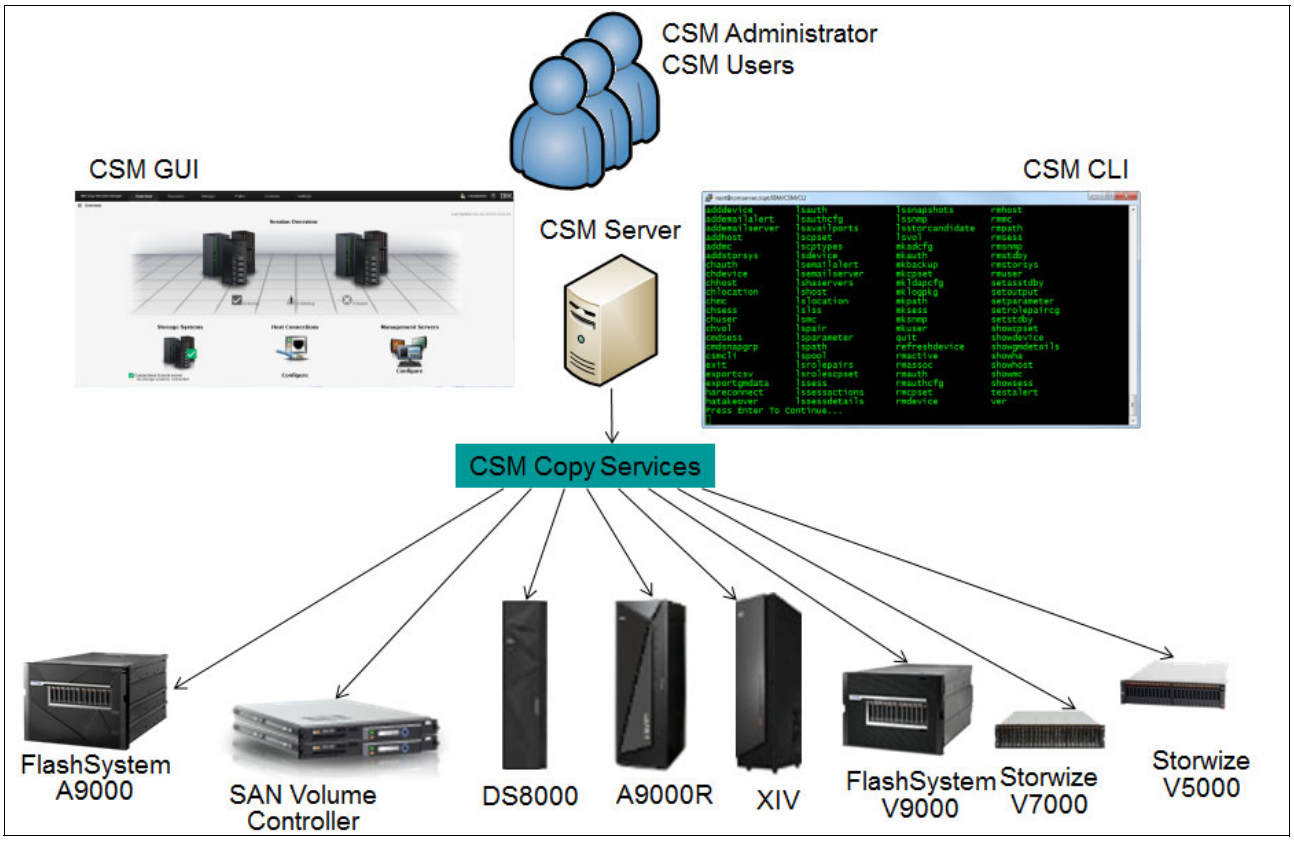

Figure 1-1 IBM Copy Services Manager: Managed IBM Storage Systems overview

1.2 New features in CSM 6.1

IBM Copy Services Manager 6.1 new features are provided in releases. At the time of this document, the release available was 6.1.4.

See the following product documentation for a list of new products and features.

Release Notes can be found at the Fix Central website.

Important: To access Release Notes information from IBM Support, you might need to authenticate using your IBM ID. If you do not have one, register.

1.3 Terminology

In this section, we describe the following key terms to help you understand and effectively use IBM Copy Services Manager:

- Management server

  The management server is a system that has IBM Copy Services Manager Server code installed. The management server provides a central point of control for managing data replication.

You can create a high availability environment by setting up a standby management server. A standby management server is a second instance of IBM Copy Services Manager server that runs on a different physical system, but is continuously synchronized with the primary (or active) IBM Copy Services Manager server.

The active management server issues commands and processes events, while the standby management server records the changes to the active server. As a result, the standby management server contains identical data to the active management server and can take over and run the environment without any loss of data.

- Storage system

  A storage system is a hardware device that contains data storage. Copy Services Manager (CSM) can control data replication within and between various storage systems. To replicate data among storage systems by using CSM, you must manually add a connection to each storage system.

- Host system

  A host system is an AIX or IBM z/OS system that connects to storage systems to enable certain replication features for those systems.

  A connection to a z/OS host system is required if you want to enable z/OS features, such as HyperSwap and hardened freeze in CSM sessions.

  A connection to an AIX host system is required if you want to use the Open HyperSwap[5] feature. This feature enables the automatic swap of input/output (I/O) to the volumes on a secondary site when a failure occurs when I/O is written to the primary site.

- Users and groups

  Copy Services Manager does maintain a directory of user names and passwords. CSM can also be configured to use a Lightweight Directory Access Protocol (LDAP) repository for user authentication.

  You can use the CSM graphical user interface (GUI) or command-line interface (CLI) to assign the users and groups that are defined in the user repository to a user role.

- User roles

  A user role determines the tasks and sessions that a user or group can manage. Copy Services Manager provides a set of predefined user roles: Monitor, Operator, and Administrator.

  - Administrators have unrestricted access to all features and functions in CSM.
  - Operators can manage specific sessions.
  - Monitors can view information in the CSM; however, they cannot modify or perform any commands or actions.

- Global Copy

  For DS8000 storage systems, Global Copy is an asynchronous long-distance copy option for data migration and backup.

- Session

  A session completes a specific type of data replication for a specific set of volumes. During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a copy set. A session can contain one or more copy sets. The type of data replication that is associated with the session determines the actions that you can perform against all copy sets in the session, the number of volumes that a copy set can contain, and the role that each volume plays.

[5] Open HyperSwap is not supported for AIX host servers that are in a clustered environment such as IBM PowerHA® SystemMirror® for AIX.

- Copy set

  A copy set is a set of volumes that represent copies of the same data. During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a copy set.

  Each volume in a copy set must be of the same size and volume type. For example, SAN Volume Controller volumes must be used with other SAN Volume Controller volumes. The number of volumes in the copy set and the role that each volume plays is determined by the session type that is associated with the session to which the copy set belongs.

- Volume roles

  Volume roles are given to every volume in the copy set.

  The volume role types are host volume, journal volume, intermediate volume, target volume, and change volume (SAN Volume Controller or Storwize storage systems only).

  The role defines how the volume is used in the copy set and the site location of the volume. For example, a host volume at the primary site has the role of Host1 (H1), while a journal volume at the secondary site has the role of Journal2 (J2).

- Role pair

  A role pair is the association of two roles in a session that take part in a copy relationship. For example, in a Metro Mirror session, the role pair can be the association between the volume roles of Host1 and Host2.

- Site

  The site determines the location of the volumes. The number of sites in a copy set is determined by the session type. IBM Copy Services Manager supports up three sites:

  - Site 1: The location of the primary storage system that contains the source data. Upon initial configuration, this site contains the host volumes with updates that are copied to the target volumes.
  - Site 2: The location of the secondary storage system that receives the copy updates from the primary storage system.
  - Site 3: The location of the tertiary storage system that receives the copy updates from either the primary or the secondary storage system (Multi-target sessions).

- Host volume

  A host volume is a volume that is connected to a server that reads and writes input/output (I/O). A host volume can be the source of updated tracks when the server that is connected to the host volume is actively issuing read and write I/O. A host volume can also be the target of the replication. When the host volume is the target, writes are inhibited.

  Host volumes are abbreviated as Hx, where *x* identifies the site.

- Journal volume

  A journal volume stores data that changed since the last consistent copy was created. This volume functions like a journal and holds the required data to reconstruct consistent data at the Global Mirror remote site. When a session must be recovered at the remote site, the journal volume is used to restore data to the last consistency point. A FlashCopy replication session can be created between the host or intermediate volume and the corresponding journal volume after a recover request is started to create another consistent version of the data.

  Journal volumes are abbreviated as Jx, where *x* identifies the site.

- Intermediate volume

  An intermediate volume receives data from the primary host volume during a replication with practice session. During a practice, data on the intermediate volumes is flash copied to the practice host volumes.

  Depending on the replication method that is used, data on intermediate volumes might not be consistent.

  Intermediate volumes are abbreviated as Ix, where *x* identifies the site.

- Target volume

  A target volume receives data from a host or intermediate volume. Depending on the replication type, that data might or might not be consistent. A target volume can also function as a source volume. For example, a common use of the target volume is as a source volume to allow practicing for a disaster, such as data mining at the recovery site while still maintaining disaster recovery capability at the production site.

- Change volume

  A change volume contains point-in-time images that are copied from the host or target volume.

  Change volumes are abbreviated as Cx, where *x* identifies the site.

1.4 Architecture

IBM Copy Services Manager high-level architecture is shown in Figure 1-2.

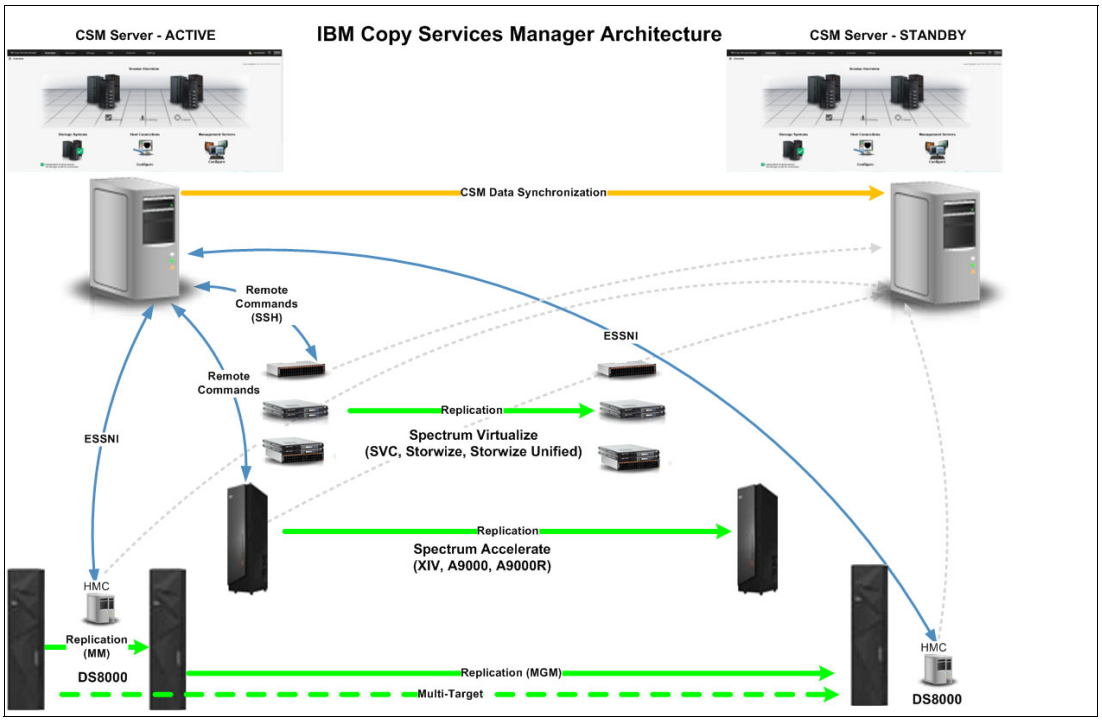

Figure 1-2 IBM Copy Services Manager high level architecture

Connectivity to storage systems is provided as follows:

- IBM Enterprise Storage Server® Network Interface (ESSNI) for the DS8000 series

- Remote commands on Secure Shell for IBM Spectrum Virtualize and IBM Spectrum Accelerate storage systems.

The IBM Copy Services Manager is implemented as an application using the IBM WebSphere® Liberty framework and Apache Derby database.

1.5 Managed storage systems and session types

Copy Services Manager controls data replication within and between various storage systems. All storage systems are grouped into site locations. The relationships between replicated volumes are managed by sessions. A session is used to complete a specific type of data replication against a specific set of volumes.

The type of copy service that is associated with the session determines the replication actions that are available for the session. For example, the options for FlashCopy sessions are different from the options for Metro Mirror sessions.

So, one session has the following components:

- One or more *sites*
- One or more *storage systems*
- One or more *copy sets*

When copy sets are added to the session, only the storage systems whose location matches the location of the site are allowed for selection. This ensures that a session relationship is not established in the wrong direction.

Table 1-1 shows the session types that are available in CSM 6.1.x and the storage systems that are supported.

Table 1-1 Sessions types and supported storage systems

Session name	Spectrum Virtualize[a]	Spectrum Accelerate[b]	ESS/DS8K	HyperSwap Function
Single site				
FlashCopy	Yes	N/A	yes	N/A
Snapshot	N/A	Yes	N/A	N/A
Dual site synchronous				
Basic HyperSwap	N/A	N/A	Yes[c]	N/A
Metro Mirror Single Direction	Yes	N/A	N/A	N/A
Metro Mirror Failover/Failback	Yes	Yes	Yes	Yes
Metro Mirror Failover/Failback with Practice	Yes	N/A	Yes	Yes
Dual site asynchronous				
Global Mirror Single Direction	Yes	N/A	Yes	N/A
Global Mirror Failover/Failback	Yes	Yes	Yes	N/A
Global Mirror Failover/Failback with Change Volumes	Yes	N/A	N/A	N/A

Session name	Spectrum Virtualize[a]	Spectrum Accelerate[b]	ESS/DS8K	HyperSwap Function
Global Mirror Failover/Failback with Practice	Yes	N/A	Yes	N/A
Global Mirror Either Direction with Two-Site Practice	N/A	N/A	Yes	N/A
Multitarget sessions				
Metro Mirror - Metro Mirror	N/A	N/A	Yes[d]	Yes
Metro Mirror - Global Mirror	N/A	N/A	Yes[d]	Yes
Metro Mirror - Global Mirror with practice	N/A	N/A	Yes[d]	Yes
Metro Mirror - Global Mirror with Site 3 Global Mirror[e]	N/A	N/A	Yes[d]	Yes

a. SVC, V5000, V7000, and V9000
b. XIV and A9000
c. Available only for z/OS
d. Available only for DS8000 series
e. Available in 6.1.4

1.5.1 Practice sessions

By using practice sessions, you can test disaster-recovery actions while maintaining disaster-recovery capability. In addition, the target volumes can be used for other purposes (for example, patch testing and so on).

Practice sessions include intermediate volumes on the remote site that contains the target volumes. A FlashCopy operation is completed from the intermediate volumes to the target volumes. The target volumes contain point-in-time data that you can use to test data-recovery actions.

For example, you can run scripts that map the target volumes to host systems in a remote site, or complete an initial program load (IPL) on the host in an isolated environment (for example, using different IPs). Because data replication continues from the source volume to the intermediate volume in a normal manner, your data is recoverable while you are testing the practice volume.

To use practice volumes, the session must be in the *Prepared* state.

Practice sessions:

- When practice sessions are used, Copy Services Manager assumes that the set of volumes that are used for practicing is also used in case of actual recovery. For this reason, the IBM Copy Services Manager operations and storage resources that are used for practicing are the same with a real recovery or site switch. This approach relieves the unpredictably of untested procedures during the real recovery operations.
- You can test disaster-recovery actions without the use of practice volumes. However, if you do not use practice volumes, data replication between sites is interrupted while you are recovering data to the remote site.
- Any existent session can be “upgraded” to use practices. However, on the new session, the old target volume becomes Intermediary (I), and the new target is a newly created FlashCopy volume.

1.5.2 Monitoring sessions icons and symbols

This section presents the icons and symbols used in the IBM Copy Services Manager GUI.

Session status icons

The Copy Services Manager GUI uses icons to represent the status of each session. Table 1-2 describes each session status icon.

Table 1-2 Session status icons

Icon	Meaning	Description
	Inactive	The session is in a defined state, with no activity on the hardware.
	Normal	A consistent copy of the data either exists or is being maintained.
	Warning	For Metro Mirror, Global Mirror, and Metro Global Mirror, the session might have volumes that are being synchronized or are about to be synchronized, with no suspended volumes. For FlashCopy, the warning status is valid only after the start command is issued and before the flash. This warning status means that the session is either preparing or is ready for a flash command, but targets do not yet have a consistent copy. If a HyperSwap session is degraded, which means it is enabled on one or more sysplex members and disabled on at least one sysplex member, then the session is in a warning state.
	Severe	One or more errors must be dealt with immediately. Possible causes include the following: ► One or more volumes are suspended ► A session is suspended ► A volume is not copying correctly

Volume role symbols

The Copy Services Manager GUI also provides a visual aid to help you create and manage your sessions. The visual aid shows the number of volume roles in the session and how the roles are distributed between the sites. The volume role symbols represent the replication status on the volumes.

Table 1-3 presents the meaning of each volume role symbol.

Table 1-3 Volume role symbols

Symbol	Description	Meaning
	Active host volumes	This symbol represents volumes that contain the source of updated tracks to which the application is actively issuing read and write input/output (I/O).

Symbol	Description	Meaning
	Active host volumes with change volumes	This symbol represents volumes that contain the source of updated tracks to which the application is actively issuing read and write I/O and change volumes.
	Recoverable volumes	This symbol represents volumes that contain a consistent copy of the data.
	Recoverable volumes with change volumes	This symbol represents volumes and change volumes that contain a consistent copy of the data.
	Inconsistent volumes	This symbol represents the volumes that do not contain a consistent copy of the data.
	Inconsistent volumes with change volumes	This symbol represents the volumes and change volumes that do not contain a consistent copy of the data.

Data copying symbols

The Copy Services Manager GUI also shows the copy method and direction. The data copying symbols indicate the type of copy that occurs between the volume roles. The direction that the symbol is displayed in the Copy Services Manager GUI depends on the direction of the copy.

Table 1-4 presents the meaning of each data copying symbol.

Table 1-4 Data copying symbols

Symbol	Description	Meaning
	FlashCopy	This symbol represents a FlashCopy operation.

Symbol	Description	Meaning
	FlashCopy with errors	This symbol represents a FlashCopy operation with errors on one or more pair.
	FlashCopy inactive	This symbol represents an inactive FlashCopy operation.
	FlashCopy inactive with errors	This symbol represents an inactive FlashCopy operation with errors on one or more pair.
	Synchronous	This symbol represents a synchronous copy.
	Synchronous with errors	This symbol represents a synchronous copy with errors on one or more pairs.
	Synchronous inactive	This symbol represents an inactive synchronous copy.
	Synchronous inactive with errors	This symbol represents an inactive synchronous copy with errors on one or more pair.
	Asynchronous	This symbol represents an asynchronous copy.
	Asynchronous with errors	This symbol represents an asynchronous copy with errors on one or more pair.
	Asynchronous inactive	This symbol represents an inactive asynchronous copy.
	Asynchronous inactive with errors	This symbol represents an inactive asynchronous copy with errors on one or more pair.
	HyperSwap or Open HyperSwap	This symbol indicates that the HyperSwap or Open HyperSwap feature is enabled for the relationship. If a failure occurs when I/O is being written to the primary storage system, these features automatically swap the I/O to the secondary site with no user interaction and little or no effect on the application.
	Suspended	This symbol represents a suspended copy relationship.
	Failed over	This symbol represents a failover copy relationship.

1.5.3 Session types

This section provides a brief description of the copy services session types and their associated icons in the CSM GUI.

FlashCopy

FlashCopy (see Table 1-1 on page 7) sessions copy the data that is on the source volume to the target volume on the same site (Site 1). The target volume contains the same data as the source volume at the point in time when the copy was established. Any subsequent write operations to the source volume are not reflected on the target volume. Figure 1-3 shows the volumes and data flow for the session.

Figure 1-3 FlashCopy session pictogram

Snapshot

Snapshot sessions (see Table 1-1 on page 7) create a point-in-time copy of a volume or set of volumes on the same site (Site 1) without having to define a specific target volume. The target volumes of a Snapshot session are automatically created when the snapshot is created. Figure 1-4 shows the volumes for the session.

Figure 1-4 Snapshot session pictogram

Metro Mirror Single Direction

Metro Mirror Single Direction sessions copy data in a single direction from the source volume on the local site (Site 1) to the target volume on the remote site (Site 2). Figure 1-5 shows the volumes and data flow for the session when data is copied from Site1 to Site 2.

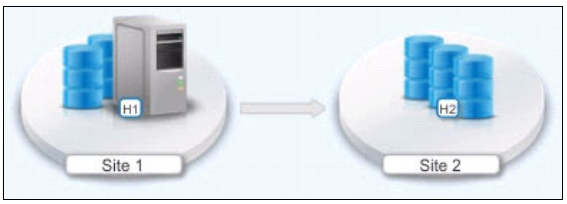

Figure 1-5 Metro Mirror single direction session pictogram

Metro Mirror Failover/Failback

Metro Mirror Failover/Failback sessions provide the same capabilities as Metro Mirror Single Directions sessions. The difference is that data replication for Metro Mirror Failover/Failback sessions is bidirectional.

Figure 1-6 shows the volumes and data flow for the session when data is copied from Site 1 to Site 2.

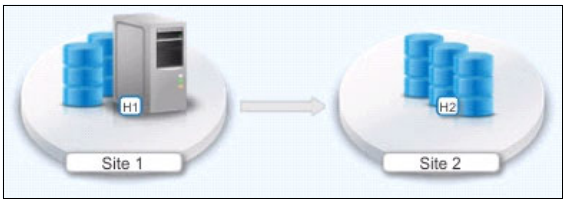

Figure 1-6 Metro Mirror failover and failback session pictogram

Basic HyperSwap

Basic HyperSwap sessions provide the same capabilities as Metro Mirror Failover/Failback session. A new HyperSwap function is added which commands the Z HyperSwap to switch I/O operations to the other site. Figure 1-7 shows the volumes and data flow for the session and also the HyperSwap sign.

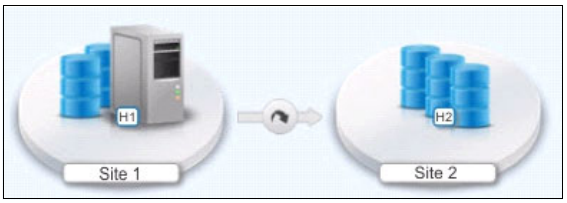

Figure 1-7 Basic HyperSwap session pictogram

Metro Mirror Failover/Failback with Practice

Metro Mirror Failover/Failback with Practice sessions combine Metro Mirror and FlashCopy replication to provide a point-in-time copy of the data on the remote site.

For this session type, a synchronous copy occurs from a source volume on the local site (Site 1) to an intermediate volume on the remote site (Site 2). A FlashCopy then occurs from the intermediate volume to a target volume on the remote site.

Figure 1-8 shows the volumes and data flow for the session when data is copied from Site 1 to Site 2.

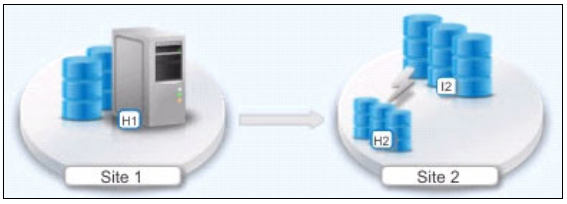

Figure 1-8 Metro Mirror failover/failback with practices session pictogram

Global Mirror Single Direction

Global Mirror Single Direction sessions copy data in a single direction from a source volume on the local site (Site 1) to a target volume on the remote site (Site 2).

DS8000 storage systems

For DS8000 storage systems[6], an asynchronous copy occurs from the source volume to the target volume. A FlashCopy then occurs from the target volume to a journal volume on the remote site.

Figure 1-9 shows the volumes and data flow for an IBM Enterprise Storage Server, IBM DS6000, or IBM DS8000 Global Mirror Single Direction session when data is copied from Site 1 to Site 2.

Figure 1-9 Global Mirror single direction session pictogram (DS8000 series)

IBM Spectrum Virtualize storage systems

For all other IBM Spectrum Virtualize storage system sessions, the data flow is the same. However, there is no FlashCopy to a journal volume.

Figure 1-10 shows the volumes and data flow for an IBM Spectrum Virtualize Global Mirror Single Direction session when data is copied from Site 1 to Site 2.

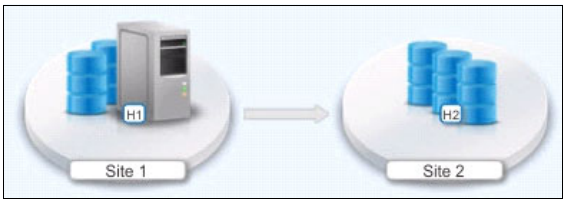

Figure 1-10 Global Mirror single direction session pictogram (IBM Spectrum Virtualize)

Global Mirror Failover/Failback

Global Mirror Failover/Failback sessions provide the same capabilities as Global Mirror Single Direction sessions. The difference is that data replication for Global Mirror Failover/Failback sessions is bidirectional.

DS8000 storage systems

Figure 1-11 shows the volumes and data flow for DS8000 Storage Systems[7] Global Mirror Single Direction session when data is copied from Site 1 to Site 2.

Figure 1-11 Global Mirror failover/failback session pictogram (DS8000 series)

[6] Also supported on ESS800 and IBM DS6000™
[7] Also supported on ESS800 and DS6000

IBM Spectrum Virtualize storage systems

For all other IBM Spectrum Virtualize storage system sessions, the data flow is the same. However, there is no FlashCopy to a journal volume, as shown in Figure 1-12.

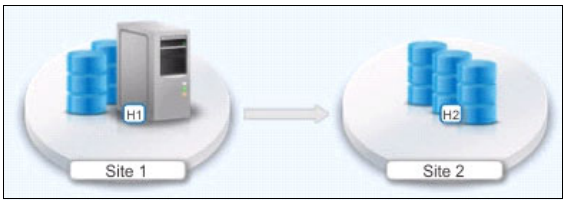

Figure 1-12 Global Mirror failover/failback session pictogram (IBM Spectrum Virtualize)

Global Mirror Failover/Failback with Change Volumes

Global Mirror Failover/Failback with Change Volumes (GMCV) sessions provide the same capabilities as Global Mirror Failover/Failback sessions but designed for lower bandwidth replication links. The difference is that the GMCV session also provides the option of enabling or disabling the use of change volumes.

Note: GMCV sessions are available only for IBM Spectrum Virtualize storage systems.

Figure 1-13 shows the volumes and data flow for the GMCV session when data is copied from Site 1 to Site 2. C1 and C2 are the change volumes for both sites.

Figure 1-13 GMCV session pictogram

Global Mirror Failover/Failback with Practice

Global Mirror Failover/Failback with Practice sessions combine Global Mirror and FlashCopy replication to provide a point-in-time copy of the data on the remote site (Site 2) for practice purpose.

DS8000 storage systems

For DS8000 storage systems[8], an asynchronous copy occurs from the source volume on the local site (Site 1) to an intermediate volume on the remote site (Site 2). A FlashCopy then occurs from the intermediate volume to both the practice volume and the journal volume on the remote site. Figure 1-14 represents the volumes and data flow for DS8000 Global Mirror Failover/Failback with Practice session when data is copied from Site 1 to Site 2.

Figure 1-14 Global Mirror failover/failback session pictogram (DS8000 series)

[8] Also supported on ESS800 and DS6000

IBM Spectrum Virtualize storage systems

For all other Spectrum Virtualize storage system sessions, the data flow is the same. However, there is no FlashCopy to a journal volume, as shown in Figure 1-15.

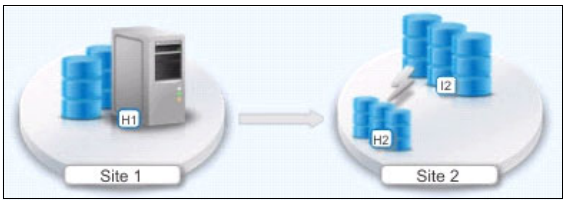

Figure 1-15 Global Mirror failover/failback w/ practice session pictogram (Spectrum Virtualize)

Global Mirror Either Direction with Two-Site Practice

Global Mirror Either Direction with Two-Site Practice sessions combines Global Mirror and FlashCopy replication to provide a point-in-time copy of the data on the local (Site 1) and remote site (Site 2). This session type includes one intermediate volume in each site, the local and remote one, so that you can practice disaster recovery from either site. Figure 1-16 shows the volumes and data flow for the session when data is copied from Site 1 to Site 2. I1 and I2 are the intermediate volumes in both sites.

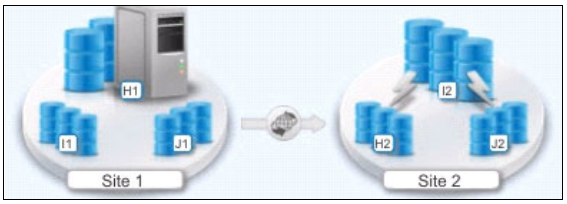

Figure 1-16 Global Mirror either direction with two-site practice pictogram

Metro Global Mirror (three sites)

Note: Metro Global Mirror sessions are available only on DS8000 and ESS storage systems.

Metro Global Mirror sessions combine Metro Mirror, Global Mirror, and FlashCopy replication into a single session. Metro Global Mirror sessions support three sites that are varying distances apart.

For this session type, a synchronous copy occurs from the source volume on the local site (Site 1) to the target volume on the second site (Site 2). An asynchronous copy then occurs from the volume in the second site to the target volume on the third site (Site 3) followed by a FlashCopy from the target to the journal volume on the same Site 3.

Figure 1-17 shows the volumes and data flow for the session when data is copied from Site 1 to Site 2 and then to Site 3.

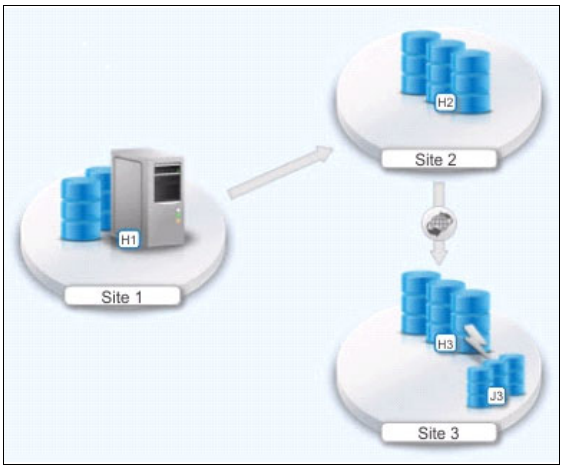

Figure 1-17 Metro Global Mirror session pictogram

Metro Global Mirror with Practice

Note: Metro Global Mirror with Practice sessions are available only on DS8000 and ESS storage systems.

Metro Global Mirror with Practice sessions combine Metro Mirror, Global Mirror, and FlashCopy replication across three sites to provide a point-in-time copy of the data on the third site.

For this session type, a synchronous copy occurs from the source volume on the local site (Site 1) to the target volume on the second site (Site 2). An asynchronous copy then occurs from the second site to the intermediate volume on the third site (Site 3) followed by a FlashCopy from the intermediate volume to the target and journal volumes on Site 3. Figure 1-18 shows the volumes and data flow for the session when data is copied from Site 1 through Site 3.

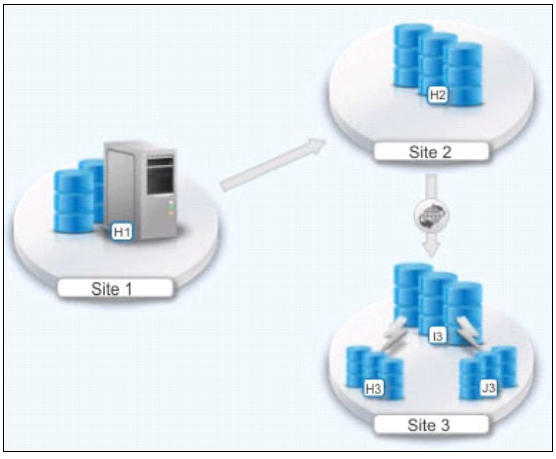

Figure 1-18 Metro Global Mirror with Practice session pictogram

Metro Mirror - Metro Mirror (multi-target)

Note: Metro Mirror - Metro Mirror sessions are available only for IBM DS8000 storage systems with a microcode level that supports single source to multi-target relationships. To determine whether you can use this session type, refer to the IBM DS8000 documentation for the microcode level that you are using.

Metro Mirror - Metro Mirror session is a multi-target session over three sites. One site (Site 1) is defined as source for two different Metro Mirror replications to other sites (Site 2 and Site 3).

Figure 1-19 shows the volumes and data flow for the session when data is copied from Site 1 to Site 2 and Site 3.

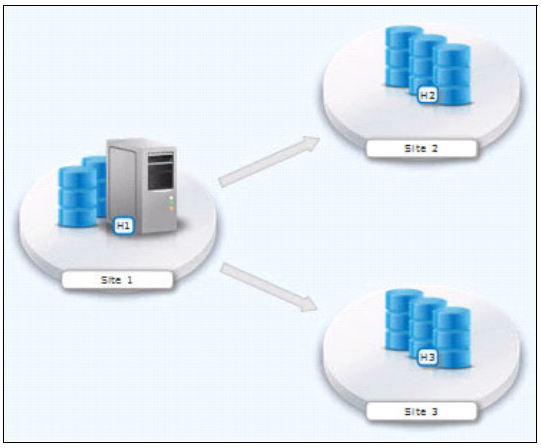

Figure 1-19 Metro Mirror - Metro Mirror session pictogram

Metro Mirror - Global Mirror

Note: Metro Mirror - Global Mirror sessions are available only for IBM DS8000 storage systems with a microcode level that supports single source to multi-target relationships. To determine whether you can use this session type, refer to the IBM DS8000 documentation for the microcode level that you are using.

Metro Mirror - Global Mirror session is a multi-target session against three sites. One site (Site 1) is defined as source for a Metro Mirror replication to second site (Site 2) and for another Global Mirror replication to the third site (Site 3).

Figure 1-20 on page 19 shows the volumes and data flow for the session when data is copied from Site 1 to Site 2 and Site 3.

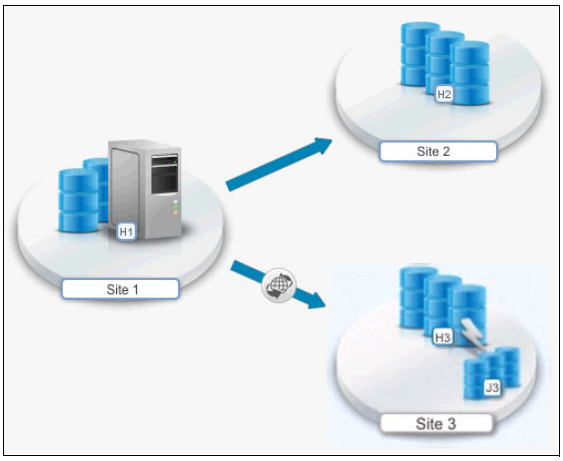

Figure 1-20 Metro Mirror Global Mirror (multi-target) session pictogram

Metro Mirror - Global Mirror with Practice

Note: Metro Mirror - Global Mirror with Practice sessions are available only for IBM DS8000 storage systems with a microcode level that supports single source to multi-target relationships. To determine whether you can use this session type, refer to the IBM DS8000 documentation for the microcode level that you are using.

Metro Mirror - Global Mirror with Practice session provide the same capabilities as Metro Mirror - Global Mirror session. In addition, this session uses FlashCopy volumes to practicing disaster recovery without losing your disaster recovery capability. Figure 1-21 shows the volumes and data flow for the session when data is copied from Site 1 to Site 2 and Site 3 and also the practice volumes.

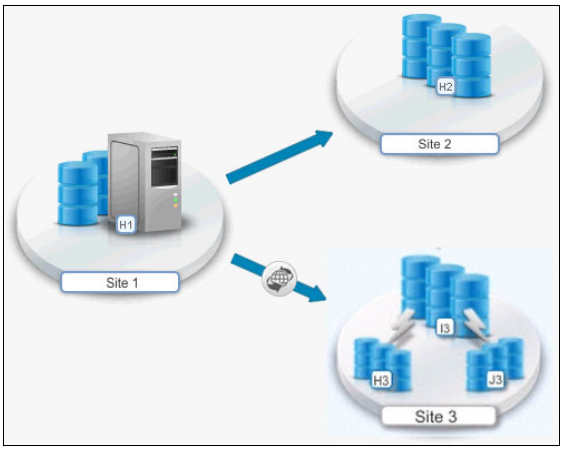

Figure 1-21 Metro Mirror Global Mirror (multi-target) w/ practice session pictogram

Metro Mirror - Global Mirror with Site 3 Global Mirror

Note: Metro Mirror - Global Mirror with Site 3 Global Mirror sessions are available only for IBM DS8000 storage systems with a microcode level that supports single source to multi-target relationships. To determine whether you can use this session type, refer to the IBM DS8000 documentation for the microcode level that you are using.

Metro Mirror - Global Mirror with Site 3 Global Mirror session provide the same capabilities as Metro Mirror - Global Mirror. In addition, after recovery in the third site (Site 3), this session allows a Global Mirror replication back to either primary site (Site 1) or secondary site (Site 2) and a cascaded Global Copy to the third site. This action provides disaster recover capabilities while production runs at Site 3. Figure 1-22 shows the volumes and data flow for the session when data is copied from Site 1 to Site 2 and Site 3 and also the journal volumes on each site.

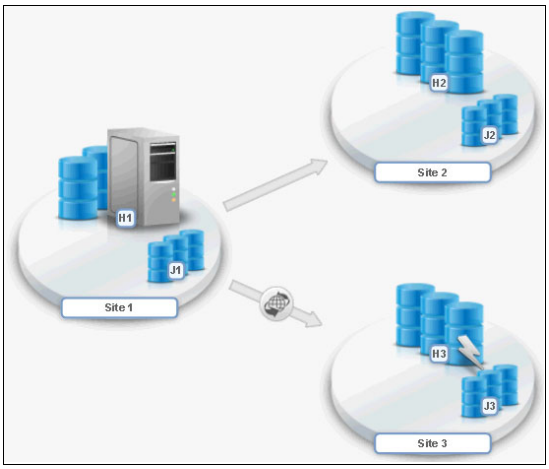

Figure 1-22 Metro Mirror - Global Mirror with Site 3 Global Mirror session pictogram

1.5.4 Session commands

The commands (actions) that are available for a session depend on the session type. Commands are issued synchronously to Copy Services Manager sessions. Any subsequent command that is issued to an individual session is not processed until the first command completes.

Some commands, such as the Start command, can take an extended amount of time to complete. By using the GUI, you can still issue commands to other sessions and not hold up functionality. When a command completes, the GUI console displays the results of the command.

The complete list of session commands with GUI details can be found in IBM Knowledge Center for CSM.

Note: The CSM CLI commands use specific syntax. For an example of using the CSM CLI see 2.2, “CSM command-line interface (CLI)” on page 23.

2

IBM Copy Services Manager deployment considerations

In this chapter, we provide considerations for deploying IBM Copy Services Manager within your environment.

We briefly describe supported platforms for CSM server deployment and considerations for CSM server replication options for deploying CSM in a disaster recovery environment, security considerations for typical deployments, and some considerations for using the CSM command-line interface (**`csmcli`**).

The following topics are described in this chapter:

- Planning considerations
  - High availability considerations
  - Security and users
- CSM command-line interface (CLI)

2.1 Planning considerations

IBM Copy Services Manager can be deployed on a range of platforms:

- IBM z/OS
- Linux for IBM z® Systems
- Linux on x86_64 bit
- Microsoft Windows Server
- IBM AIX
- Linux on IBM Power Systems (Big Endian - ppc64)

See the following link for detailed requirements (including virtualized environment).

For installation steps, see the IBM Copy Services Manager documentation IBM Copy Services Manager Documentation.

2.1.1 High availability considerations

IBM Copy Services Manager has built-in functionality for business continuity of the CSM management environment. The CSM server data can be replicated from the *Active* CSM server to a *Standby* CSM server (see Figure 2-1). Default communications port for CSM data replication is TCP:9561. The port can be changed in the CSM server configuration file for replication, named **`rmserver.properties`**.[1]

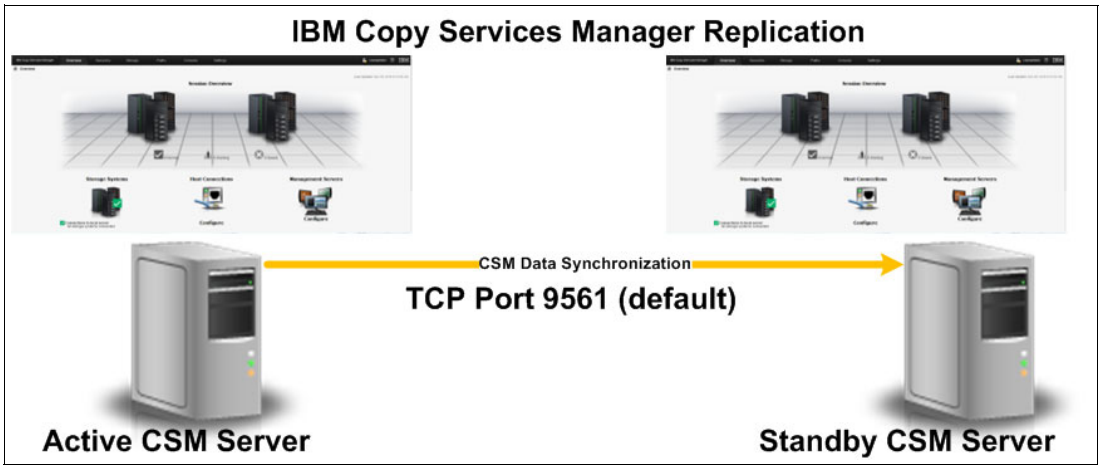

Figure 2-1 CSM Server replication

Failover and failback

Although there is no automated failover from the Active to the Standby server, nor automated failback (after a failover and initial CSM server recovery). These procedures are straight forward.

Important: When a CSM Server is defined as `Standby`, its configuration is overwritten completely. If you have an active configuration you want to use later, you need to save this configuration before defining the server as `Standby`. You can save the server from either GUI or *csmcli* (**`mkbackup`** command in *csmcli*).

Note: Both Active and Standby CSM servers must be at the same CSM software level.

[1] The path to the file differs depending on the platform on which the CSM server is deployed and your installation preferences.

In case the primary (Active) CSM Server is not available, manual intervention is required to "promote" the *Standby* CSM server to *Active* (failover) and continue the storage replication management. Upon primary CSM server recovery, the administrator must decide which configuration to use for further operations, and take appropriate actions.

An example of CSM server recovery is described in 6.2, "Configuring and testing high availability for the CSM Server" on page 129.

Disaster recovery and HyperSwap considerations

More security configuration is required if the managing CSM server instance is running outside the z/OS image or Parallel Sysplex instance configured for z/OS HyperSwap. See Chapter 7, "Securing IBM z/OS HyperSwap communication with IBM Copy Services Manager" on page 153.

2.1.2 Security and users

For IBM Copy Services Manager 6.1.x, users' definitions are stored by default in the CSM server's basic user registry. Copy Services Manager provides a set of predefined user roles: *Monitor*, *Operator*, and *Administrator*.

In addition to basic user registry, for CSM server deployment on z/OS, an operating system repository, such as IBM RACF®, can be used. LDAP can be used for both z/OS systems and other supported platforms.

For more information, see the IBM Copy Services Manager documentation IBM Copy Services Manager documentation.

2.2 CSM command-line interface (CLI)

During the disaster scenarios, failover and failback procedures involve multiple actions in the CSM-managed environment. For the handling of multiple operations, scripting would be the effective choice. For scripting commands, the **`csmcli`** (Copy Services Manager command-line interface) is provided.

Tip: For running the **`csmcli`**, proper Java Runtime Environment (JRE) configuration is required.

The CSM CLI can be used either for submitting batch commands (scripting, using the associated options and arguments), or interactively by starting the **`csmcli`** (shell) with no parameters.

2.2.1 Getting started with csmcli

For more information about how to install, customize, and start **`csmcli`**, and also for detailed help for all commands, go to IBM Knowledge Center.

2.2.2 Hints and tips

One of the benefits of using the command-line interface is that the output of the commands run can be used for other purposes (for example, documenting configurations, use as input for other commands, and so on). This section provides examples of how to use the CSM command-line interface (csmcli).

Creating a report for later use (documenting configuration)

When a human-readable report is needed for later use (for example, the output is too large to fit on a single display, or the output results need to be filtered, and so on), one option is to format the **csmcli** output with the `delim` option and redirect to a file. Later, this file can be imported in tabular program, for further analysis.

Example 2-1 shows how to use the **csmcli** in batch mode to retrieve information about the replication sessions managed by CSM.

Example 2-1 Displaying replication sessions using csmcli

```
[root@myLinux CLI]# ./csmcli.sh lssess -fmt delim -delim : >lssession.txt
[root@myLinux CLI]# cat lssession.txt
Name:Status:State:Copy Type
==========================
PROD:Inactive:Defined:Metro Mirror Failover/Failback
MMHF_MT:Inactive:Defined:Metro Mirror - Metro Mirror
TEST:Inactive:Defined:Metro Mirror Failover/Failback
A-A_Test:Inactive:Defined:Metro Mirror - Metro Mirror
HUK_load_issue:Inactive:Defined:Metro Mirror Failover/Failback
AD_Test:Warning:Prepared:Metro Mirror Failover/Failback
```

Imported in a tabular program, the output of the command in Example 2-1 looks similar to Figure 2-2.

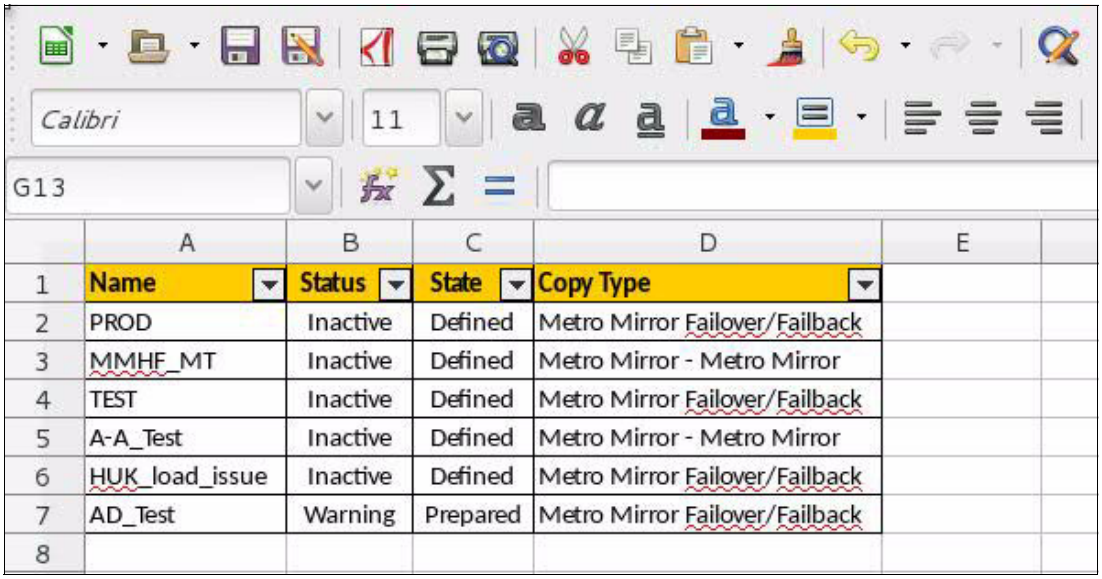

Figure 2-2 Using tabular program to display CSM data

Caveat: using the semicolon (“;”) character as separator

If the output of the **csmcli** needs to be delimited by the semicolon (“;”) character, running the **csmcli** command in batch mode might not return the wanted results (error), because the parent shell ($0) interprets the “;” as a shell command separator, as shown in Example 2-2.

Example 2-2 Wrong usage of semicolon delimiter

```
[root@myLinux CLI]# ./csmcli.sh -noinfo lsdevice -hdr off -devtype ds8000 -fmt
delim -delim ;
CMMCI9018E Parameter -delim is missing a required value.
Usage: lsdevice [ { -help|-h|-? } ] [ { -l|-s } ] [-fmt default|xml|delim|stanza]
[-p on|off] [-delim char] [-hdr on|off] [-r #] [-v on|off] -devtype
ds|ds6000|ds8000|ess|storwize-v3500|storwize-v3700|storwize-v5000|storwize-v7000|f
lashsystem-v9000|flashsystem-v840|svc|xiv [-connectionID connection_id] [id ... |
-]
Tip: Enter "help lsdevice" for more information.
[root@myLinux CLI]#
```

In this situation, you need to mask the semicolon (;) by the backslash character (\), as shown in Example 2-3.

Example 2-3 Correct usage of the semicolon (;) delimiter

```
[root@omyLinux CLI]# ./csmcli.sh -noinfo lsdevice -hdr off -devtype ds8000 -fmt
delim -delim \;
DS8000:BOX:2107.FAW31;HMC;DS8000;Connected;Connected
DS8000:BOX:2107.FCL91;HMC;DS8000;Connected;Connected
DS8000:BOX:2107.FCM21;HMC;DS8000;Connected;Connected
[root@omyLinux CLI]#
```

2.2.3 Advanced csmcli scripting examples

CSM command-line interface also can be used for advanced scripts: The command outputs can be used as input for other commands. This section provides some examples. Note that the examples are run on a system using a UNIX-like shell. Assume that you need to list all volumes that belong to LSS A3 on all storage systems registered in CSM. This listing can be accomplished completing the following steps:

1. List all storage systems registered in CSM:

   ```
   ./csmcli.sh -noinfo lsdevice -devtype ds8000
   ```

2. For each storage system, you need to list all volumes:

   ```
   ./csmcli.sh -noinfo lsvol -devtype ds -dev [device_name]
   ```

3. From the list of volumes, select only those of which the ID starts with `A3`.

The three steps can be combined in one single script. Let’s run them in reverse order:

1. To list all volumes with ID starting with `A3`, list all volumes in a **delim** formatted mode:

   ```
   ./csmcli.sh -noinfo lsvol -devtype ds -dev DS8000:BOX:2107.FAW31 -hdr off -fmt
   delim -delim \;
   ```

2. You can select only the volume ID in CSM format:

   ```
   ./csmcli.sh -noinfo lsvol -devtype ds -dev DS8000:BOX:2107.FAW31 -hdr off -fmt
   delim -delim \; |cut -d ';' -f 2
   ```

3. Applying another **cut**, but this time against the semicolon (**;**) separator, select only the real volume ID:

```
./csmcli.sh -noinfo lsvol -devtype ds -dev DS8000:BOX:2107.FAW31 -hdr off -fmt
delim -delim \; |cut -d ';' -f 2 |cut -d ':' -f 4
```

4. Finally, **grep** the output by ^A3 (use the carat (^) character to **grep** only on the beginning of the line):

```
./csmcli.sh -noinfo lsvol -devtype ds -dev DS8000:BOX:2107.FAW31 -hdr off -fmt
delim -delim \; |cut -d ';' -f 2 |cut -d ':' -f 4 |grep "^A3"
```

5. Now, use the **lsdevice** command as input for your customized command:

```
for i in `./csmcli.sh -noinfo lsdevice -hdr off -devtype ds8000 -fmt delim
-delim \; |cut -d';' -f 1`
do
echo "Volumes from system $i which belongs to LSS $1 : "
./csmcli.sh -noinfo lsvol -devtype ds -dev $i -p off -fmt delim -delim \; |cut
-d ';' -f 2 |cut -d ':' -f 4 |grep "^A3"
done
```

6. Optionally, modify the script to use the LSS ID from a parameter list. See Example 2-4.

Example 2-4 Sample script

```
[root@myLinux CLI]# cat list_lss.sh
for i in `./csmcli.sh -noinfo lsdevice -hdr off -devtype ds8000 -fmt delim
-delim \; |cut -d';' -f 1`
do
echo "Volumes from system $i which belongs to LSS $1 : "
./csmcli.sh -noinfo lsvol -devtype ds -dev $i -p off -fmt delim -delim \; |cut
-d ';' -f 2 |cut -d ':' -f 4 |grep "^$1"
done
[root@myLinux CLI]# ./list_lss.sh A3
Volumes from system DS8000:BOX:2107.FAW31 which belongs to LSS A3 :
A31A
A31B
A31C
A31D
A31E
A31F
Volumes from system DS8000:BOX:2107.FCL91 which belongs to LSS A3 :
A31A
A31B
A31C
A31D
A31E
A31F
Volumes from system DS8000:BOX:2107.FCM21 which belongs to LSS A3 :
A31A
A31B
A31C
A31D
A31E
A31F
[root@myLinux CLI]#
```

2.2.4 Remote CLI

The **csmcli** is especially helpful for developing replication session failover/failback procedures implemented as command scripts. The CSM command-line interface can be run directly from the CSM server, or it can be installed and run on a workstation that has network connectivity to the CSM server.

Also, starting with release 8.1 of the DS8000 series Hardware Management Console (HMC), the CSM server is preinstalled. When running the CSM server on the DS8000 HMC, only remote **csmcli** is available for handling scripts or interactive commands (shell).

When running the CSM CLI on a system other than the CSM server, the **csmcli** libraries are installed on that system during the installation process, and a secure connection to the CSM server (using an SSL certificate) is established. Figure 2-3 shows the secure communication port for remote CLI.

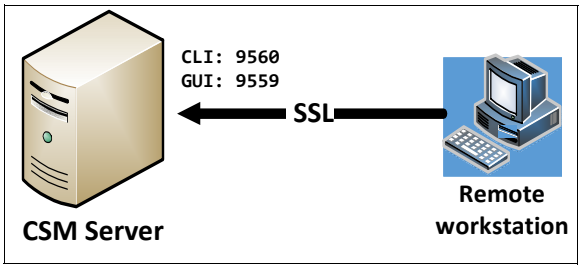

Figure 2-3 CSM server secure communication

Setup

To install the Copy Services Manager CLI on a remote system, you must download and extract the appropriate files for your software level and operating system. The following steps show how to install and start the **csmcli**:

1. Download the Copy Services Manager CLI installation package (**csmcli**) to the (remote) system where you want to run the CSM commands. The package can be downloaded from the same website as the Copy Services Manager server installation package. Optionally you can download the remote CLI from the IBM Fix Central site.

   The package types and their associated file names are listed in Table 2-1.

Table 2-1 Copy Services Manager CLI installation package file names by operating system

Operating system	Copy Services Manager CLI installation package file name
Microsoft Windows	csm-CLI-6.1.x-win.zip
AIX	csm-CLI-6.1.x-aix.tar.gz
Linux on IBM PowerPC®	csm-CLI-6.1.x-linux-ppc.tar.gz
Linux on IBM z Systems®	csm-CLI-6.1.x-linux-s390x.tar.gz
RedHat or SuSE Linux on x86	csm-CLI-6.1.x-linux-x86_64.tar.gz
z/OS	csm-CLI-6.1.x-zos.pax

2. Unpack the **csmcli** package to a directory on the remote system, in the following steps referred to as the CSM_CLI_DIR directory:
   - For Windows, extract the files by right-clicking and selecting **Extract All**.
   - For AIX or Linux, use the **gunzip**, and then the **tar** command to extract the files.
   - For z/OS, extract the files by using the **pax** command.
3. **For z/OS only:** After extracting files using the **pax** command, there is an extra step to set file attributes on Java libraries that the CSM CLI uses, as shown in Example 2-5.

Example 2-5 Setting the file attributes in z/OS

```
extattr +a CSM_CLI_DIR/csm/Java/lib/s390x/compressedrefs/libj9ifa27.so
extattr +a CSM_CLI_DIR/csm/Java/lib/s390x/default/libj9ifa27.so
```

4. Open the CSM_CLI_DIR/csm/CLI directory.
5. Edit the repcli.properties file by typing the *server name* (or IP address) and port of your Copy Services Manager server. The default settings are shown in Example 2-6.

Example 2-6 Updating repcli.properties

```
server=localhost /* Primary Copy Service Manager server name (IP Address is OK)
port=9560        /* Default port used for remote CLI
```

6. Optionally, you can update the csmcli-auth.properties with the *username* and *password* used for accessing the server. Remove the number sign (#) character in front of the *username* and *password* and update accordingly. See Example 2-7.

Example 2-7 Updating csmcli-auth.properties file

```
[root@myLinux CLI]# cat csmcli-auth.properties
# This is the properties file for CSM CLI authentication.
# You can place your CSM username and password in this file for automatic CLI
authentication.
# If the password was entered in the properties file, when the CLI shell is
entered or a CLI command is issued, the CLI
# encrypts your password and rewrites this properties file with the encrypted
password for security reasons.
# If a password is not entered, the CLI prompts for one and does not update
this file with an encrypted password.
# Place this properties file into the same directory that the csmcli.bat or
csmcli.sh file is located,
# or under the users home directory under csm-cli.
# For example in Windows:
C:\\Users\\Administrator\\csm-cli\\csmcli-auth.properties
password=XXXXXXXX
username=csmUser
[root@myLinux CLI]#
```

7. Start the Copy Services Manager CLI running **csmcli.bat** for Windows, as in Example 2-8.

Example 2-8 Starting csmcli on remote Windows workstation

```
C:\temp\csm-CLI-6.1-win\csm\CLI>csmcli.bat
IBM Copy Services Manager Command Line Interface (CLI)
 Copyright 2007, 2015 IBM Corporation
 CLI Client Version: 6.1.3, Build: a20160721-2100
 Authentication file: csmcli-auth.properties
```

```
Connected to:
 Server: PKSTN64.pok.stglabs.ibm.com     Port: 9560   UseREST: false
 Server Version: 6.1.4, Build: a20160920-1045

csmcli>
```

Alternatively, use **csmcli.sh** on AIX, Linux, or z/OS, as in Example 2-9.

Example 2-9 Starting csmcli on remote system using csmcli.sh

```
[root@myLinux CSM_kit]# ls
csm  csm-CLI-6.1.3-linux-x86_64.tar
[root@myLinux linux]# cd csm/CLI/
[root@myLinux CLI]# ls
csmcli-auth.properties  csmcli.bat  csmcli.sh  etc  lib  repcli.properties
rmserver.properties
[root@myLinux CLI]#
[root@myLinux CLI]# ./csmcli.sh
IBM Copy Services Manager Command Line Interface (CLI)
 Copyright 2007, 2015 IBM Corporation
 CLI Client Version: 6.1.3, Build: a20160721-2100
 Authentication file: csmcli-auth.properties

Connected to:
 Server: PKSTN64.pok.stglabs.ibm.com     Port: 9560   UseREST: false
 Server Version: 6.1.4, Build: a20160920-1045

csmcli>
```

3

Managing IBM FlashCopy and Metro Mirror sessions with IBM Spectrum Virtualize-based storage

In this chapter, we provide an overview of the tasks that are necessary to successfully deploy IBM Copy Services Manager (CSM) for managing IBM FlashCopy and Metro Mirror sessions with IBM Spectrum Virtualize-based storage. We have IBM SAN Volume Controller (SVC) in our environment.

We briefly describe supported platforms for CSM server deployment and considerations for CSM server replication options for deploying CSM in a disaster recovery (DR) environment, as well as security considerations for typical deployments.

The following topics are described in this chapter:

- Overview
- Configuring and managing FlashCopy sessions
- Configuring Metro Mirror sessions
- Metro Mirror failover and failback testing
- Enabling practice for Metro Mirror session

3.1 Overview

In this section, we describe a set of scenarios that we have tested in our environment. For testing, we used a configuration that consists of two IBM SAN Volume Controller clusters located in two distinct sites. A schematic diagram of our test environment is shown in Figure 3-1.

Note: In our test environment, we have used IP-based replication. Design and configuration of the transport infrastructure between sites is out of the scope of this document.

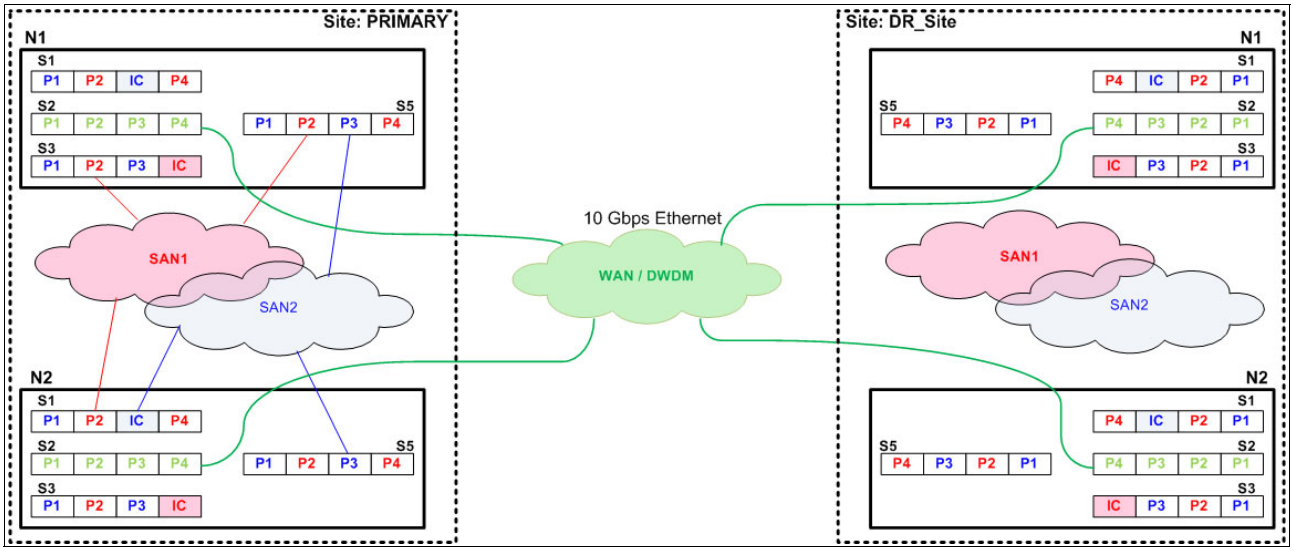

Figure 3-1 Test environment diagram

The scenarios described in the following chapters cover a set of what the authors of this document consider the most common tasks relevant for using IBM CSM. For a list of storage-based copy services types managed by IBM Copy Services Manager version, we have used (CSM 6.1.3) see Chapter 1, “IBM Copy Services Manager introduction” on page 1.

The scenarios cover the following actions for IBM SVC:

- Creating a FlashCopy session (single site)
- Configuring and managing Metro Mirror sessions
- Testing Metro Mirror failover and failback
- Enabling practice for a Metro Mirror session (Metro Mirror with Practice)

3.2 Configuring and managing FlashCopy sessions

In this section we describe how to configure and manage FlashCopy Sessions with IBM SAN Volume Controller:

Note: IBM Copy Services Manager does *not* provision (create, delete, or modify) volumes, pools, or other storage entities. IBM CSM only manages (copy services) session and consistency groups.

1. In the Copy Services Manager GUI, select **Create Session**, as shown in Figure 3-2.

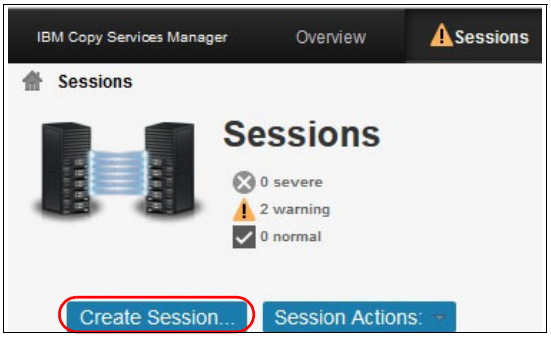

Figure 3-2 Create Session

2. For our tests, we used the IBM SVC cluster located in the secondary location, as shown in Figure 3-3. Complete the *Session name*, *Session type*, and chose the *location* (site). Click **OK** to proceed.

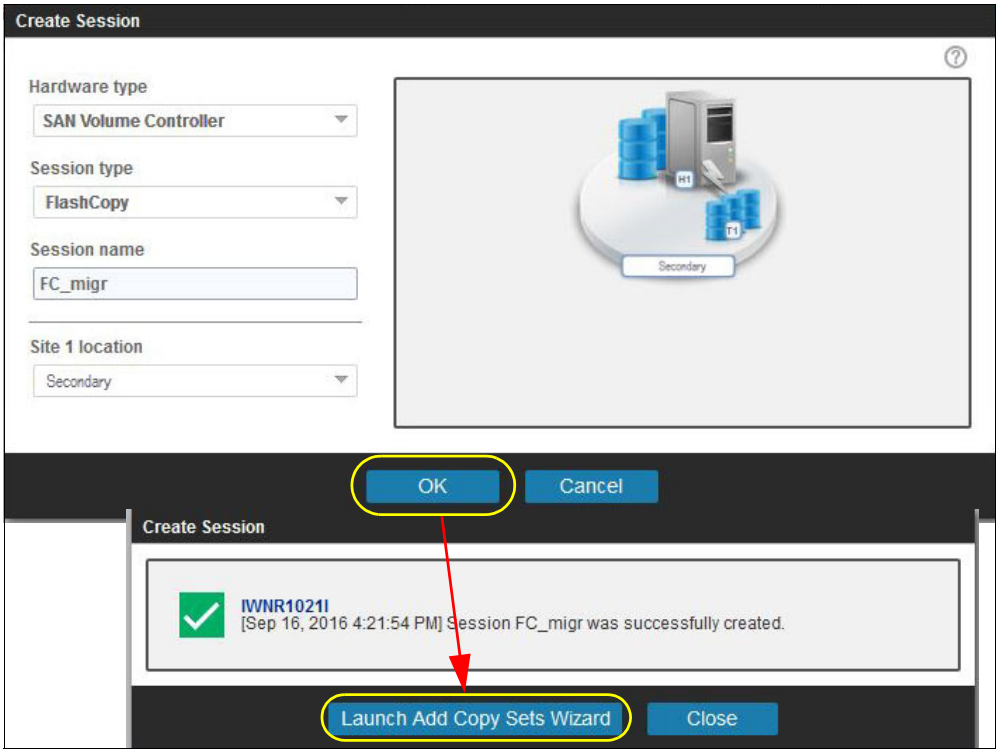

Figure 3-3 Defining the session

3. Select **Launch Add Copy Sets Wizard** to add copy sets to the newly defined sessions, and then click **Next**, as shown in Figure 3-4.

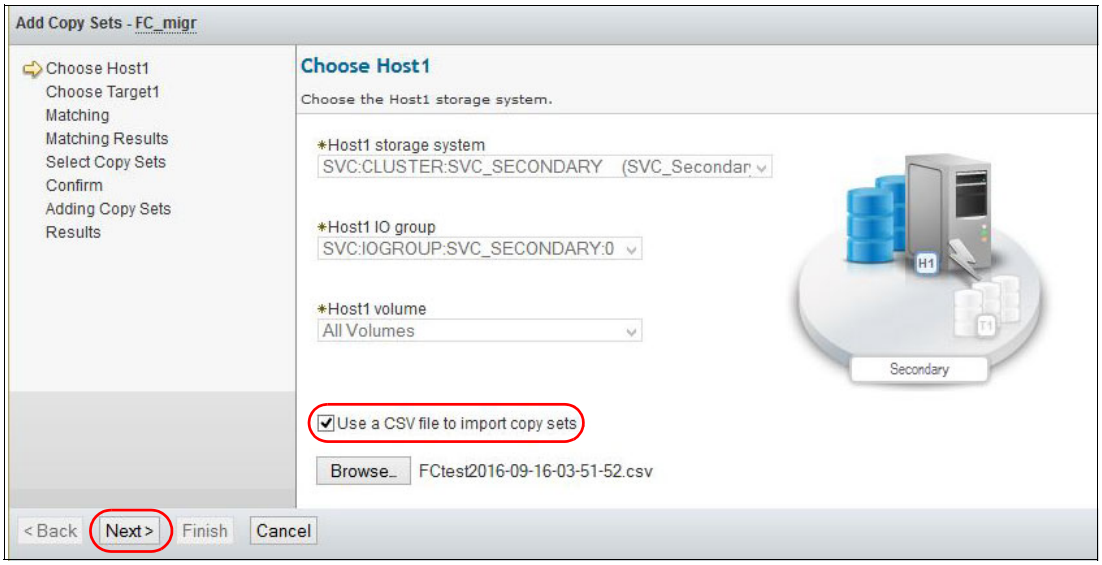

Figure 3-4 Adding copy sets

4. For defining the copy sets we have chosen to use a file that contains the list of volumes as comma-separated values (CSV) data file. Example 3-1 shows[1] an excerpt of the `.csv` file used to define the copy sets.

Example 3-1 Format of the .csv file

```
#FCtest,
#FlashCopy,
#Sep 16, 2016 3:51:52 PM
,
H1,T1

# Field significance
#<Stg_type>:VOL:<Stg_name>:<VOL_name>

SVC:VOL:SVC_SECONDARY:iOS_test_1,SVC:VOL:SVC_SECONDARY:iOS_test_1_fc01
SVC:VOL:SVC_SECONDARY:iOS_test_2,SVC:VOL:SVC_SECONDARY:iOS_test_2_fc01
SVC:VOL:SVC_SECONDARY:iOS_test_3,SVC:VOL:SVC_SECONDARY:iOS_test_3_fc01
...<< Snippet >>...
```

[1] The target volumes in SVC_Secondary have been created before defining the session in CSM.

5. When the `.csv` file has been loaded into CSM, it is matched with existing storage information. Figure 3-5 shows that the matching has been done successfully. Click **Next**.

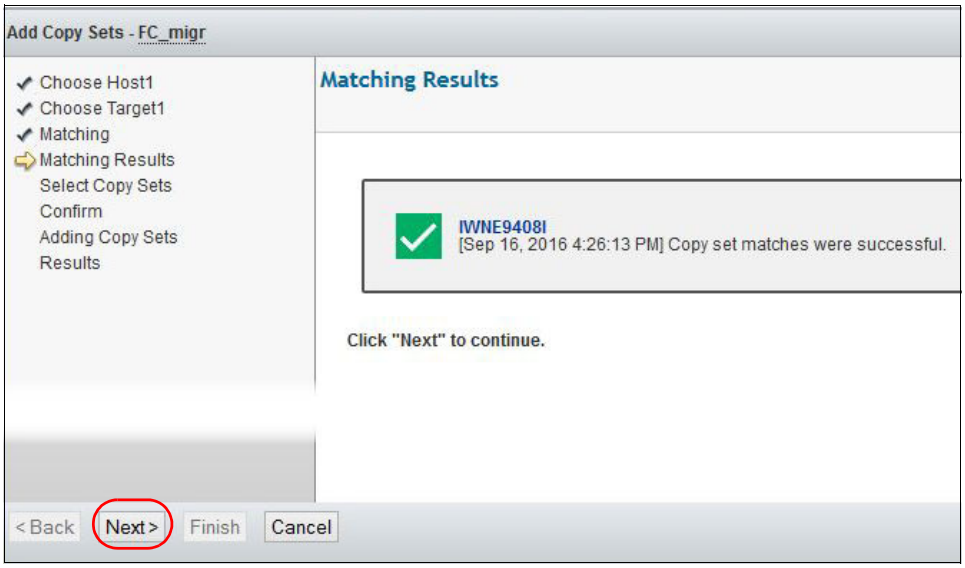

Figure 3-5 Data from the .csv file loaded into CSM

6. Check the copy sets in the following screen, shown in Figure 3-6. Click **Next**.

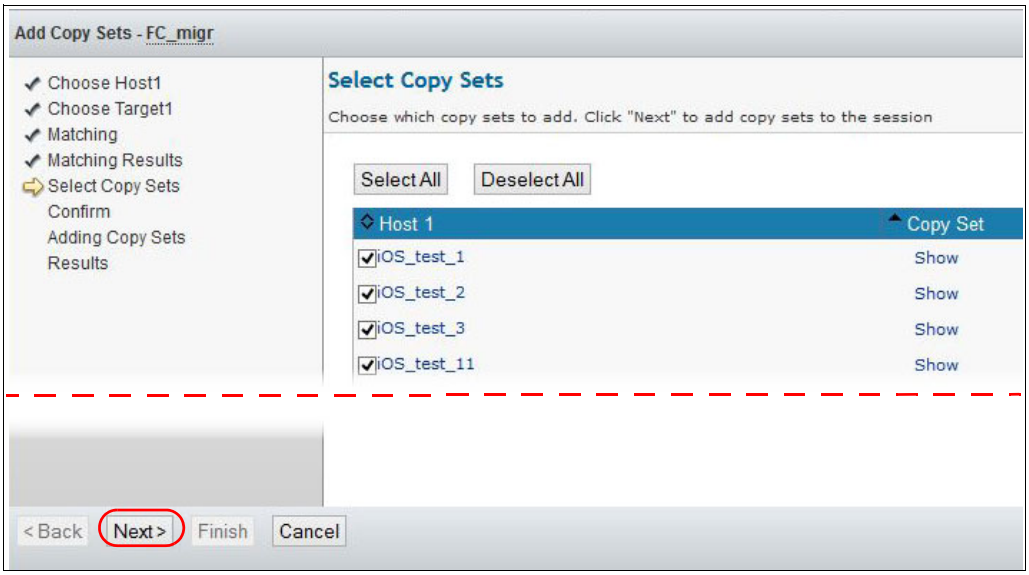

Figure 3-6 Checking selected copy sets

A summary of the actions that will be taken is shown in Figure 3-7.

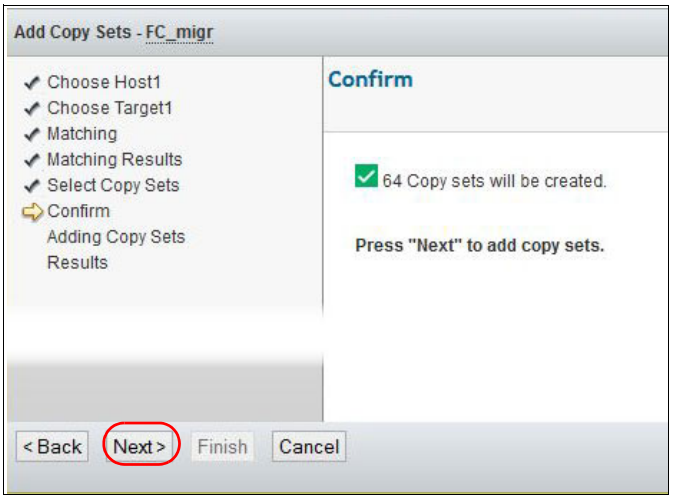

Figure 3-7 Actions summary

7. Click **Next** to confirm. The copy sets are added to the session and the results are displayed (see Figure 3-8).

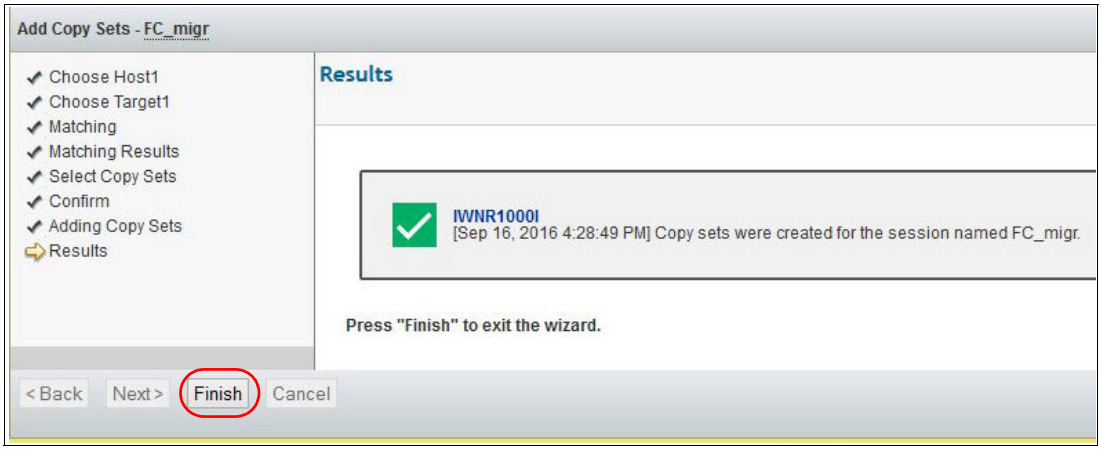

Figure 3-8 Add copy sets complete

The newly defined session is inactive at this point, as shown in Figure 3-9.

Note: At this time, no actions are performed at storage level (SVC).

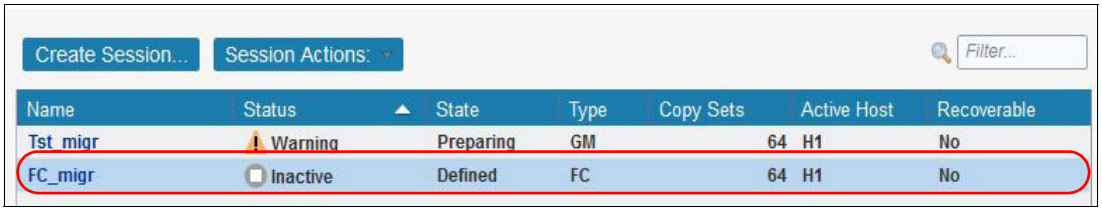

Figure 3-9 Defined session

8. Before initiating any action, you can check the settings for the FlashCopy (for example, incremental, background copy, and so on). Use **Session Actions** → **View/Modify Properties**. Figure 3-10 on page 37 shows the View / Modify Properties dialog.

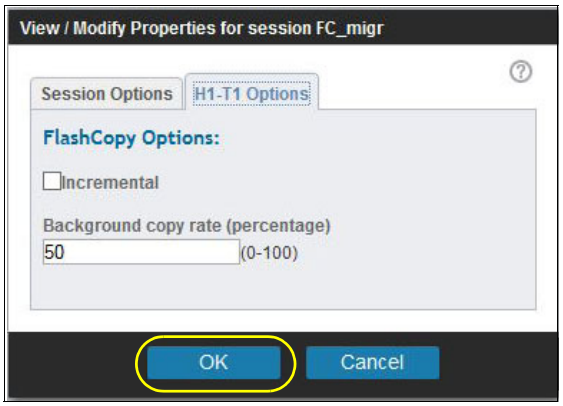

Figure 3-10 View/Modify session properties

9. After checking and adjusting parameters, we start the session (**Session Actions** → **Start** menu) and confirm the action, as shown in Figure 3-11. The first time the session is started (right after creating it), both options (**Start** or **Flash**) have the same result (FlashCopy).

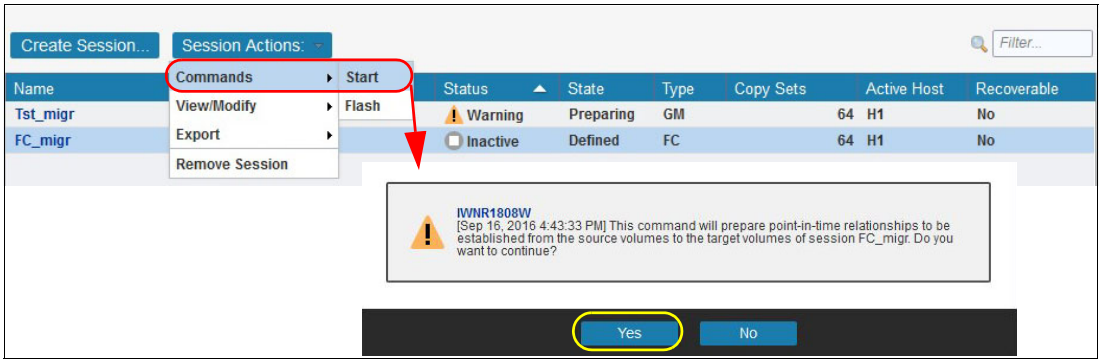

Figure 3-11 Starting the session

10. After the session has been started for the first time, a point-in-time copy is placed on the target volumes (T1). For subsequent actions, you can use the **Flash** action (see Figure 3-11) to create a new point-in-time copy (T1') of the source volumes. Note that creating a new point-in-time copy overwrites the previous copy (T1' overwrites T1).

3.3 Configuring Metro Mirror sessions

Metro Mirror provides a real-time copy of a storage volume (or multiple volumes) in the remote site that contains another SVC Cluster. The copy process is *synchronous* (primary site I/O is considered complete only when write to secondary site has been confirmed and acknowledgment has been received by primary site.

Application performance is strongly influenced by the distance between sites.

3.3.1 Scenario overview

In this scenario, in each site we use systems connected to the storage with similar configurations. After creating the Metro Mirror relationship, we test both failover and failback operations.

3.3.2 Creating the Metro Mirror relationship

In this section, we describe how we have configured a Metro Mirror relationship using the CSM GUI.

We assume that all required volumes are configured in both sites.

1. We define the session details and click **OK**, as shown in Figure 3-12 (Session type, name, and locations).

> **Note:** IBM Copy Services Manager does *not* provision (create, delete, or modify) volumes and pools or other storage entities. IBM CSM only manages (copy services) sessions (represented in HW as consistency groups).

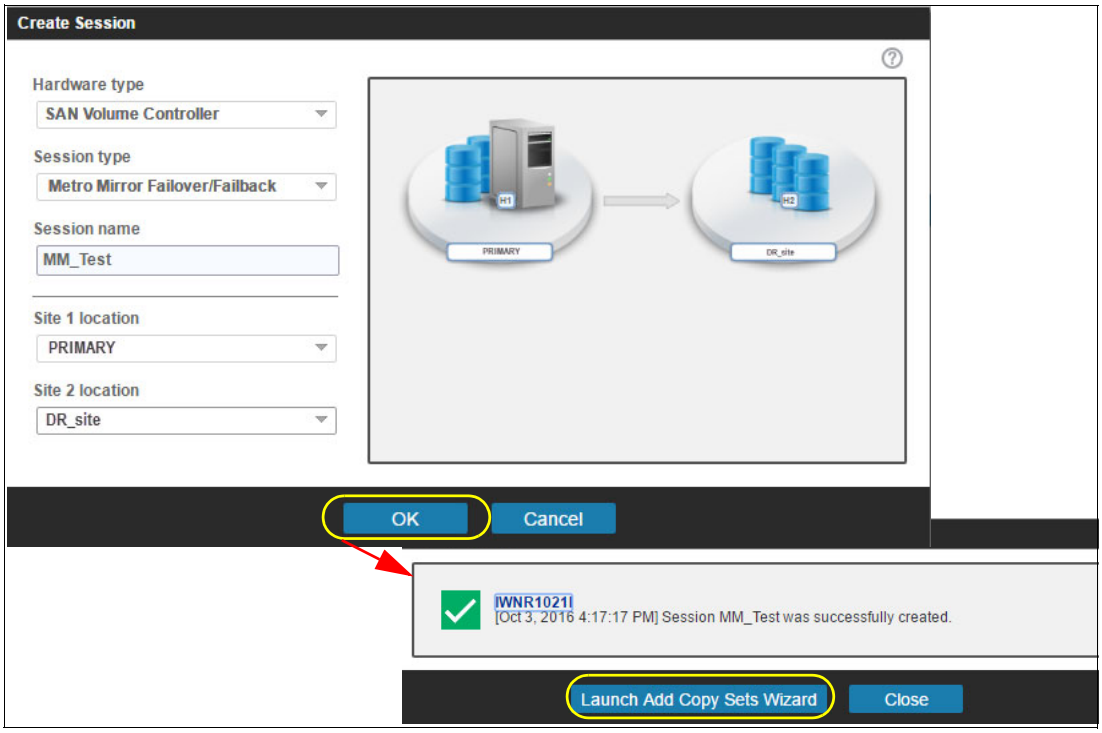

Figure 3-12 Defining the session

2. After we define the session, we launch the **Add Copy Sets** wizard and choose volumes from both sites (PRIMARY and DR_site), as shown in Figure 3-13. Click **Next**.

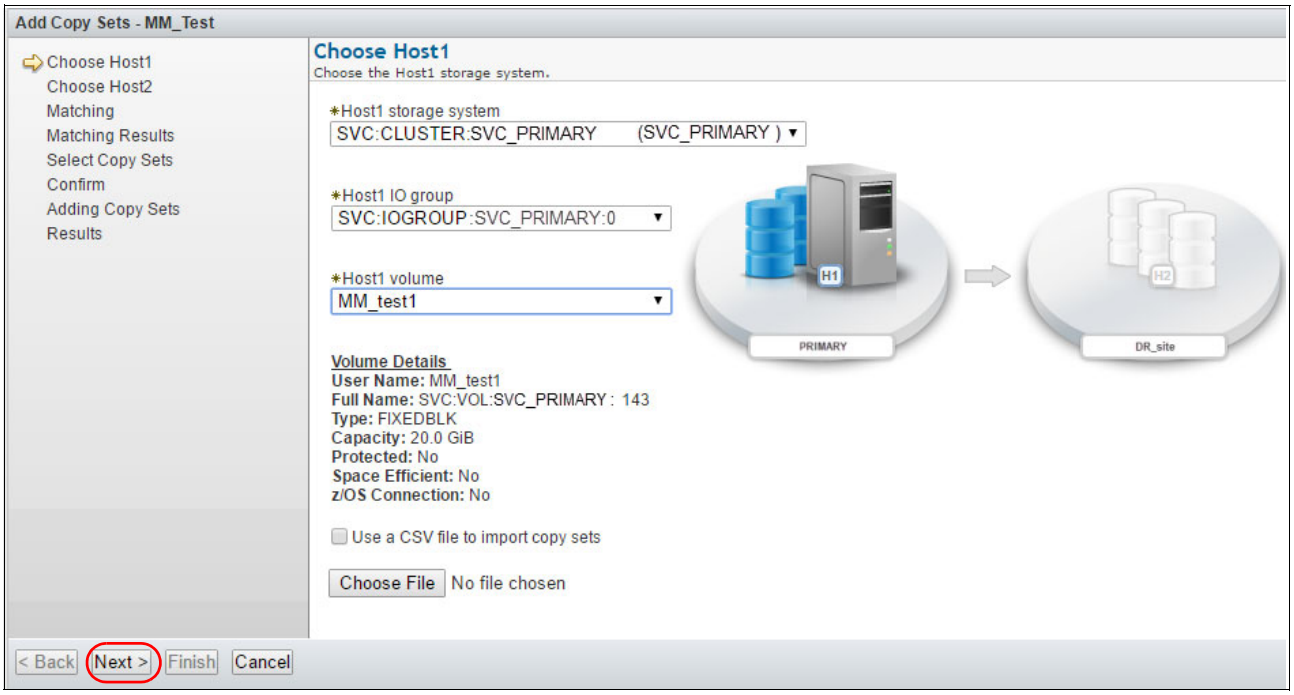

Figure 3-13 Selecting volumes from PRIMARY site

3. Next, we select the volume to be paired (DR_site), as shown in Figure 3-14. Click **Next**.

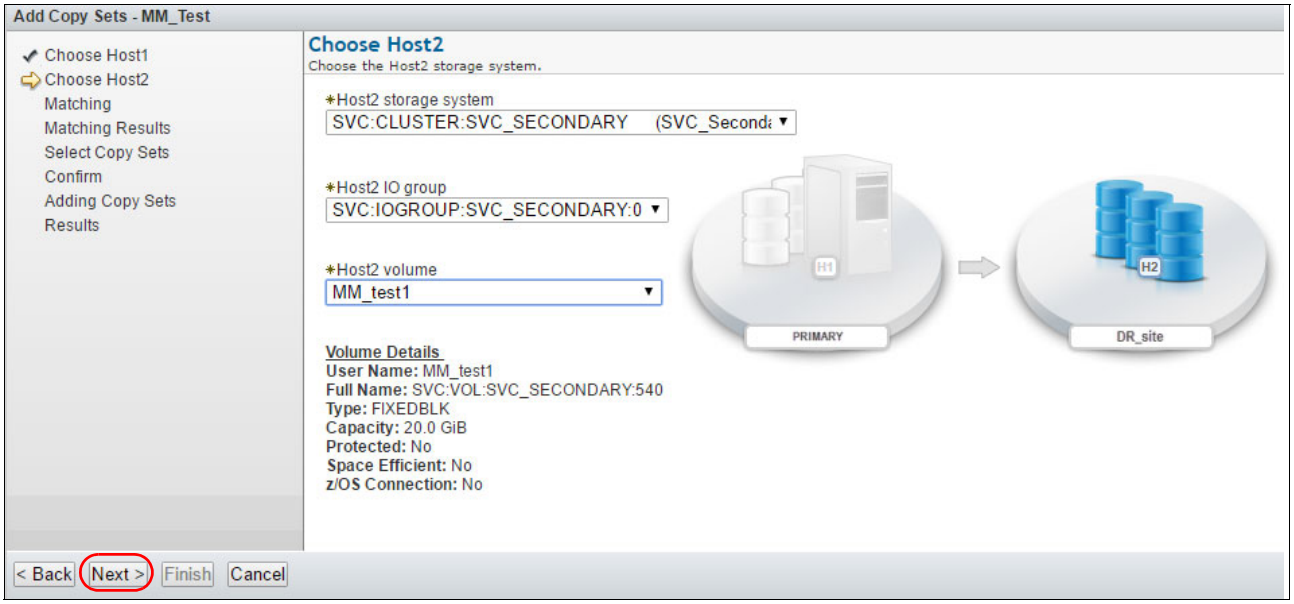

Figure 3-14 Selecting volume from DR_site (Secondary)

4. Select **Next** to proceed, as shown in Figure 3-15.

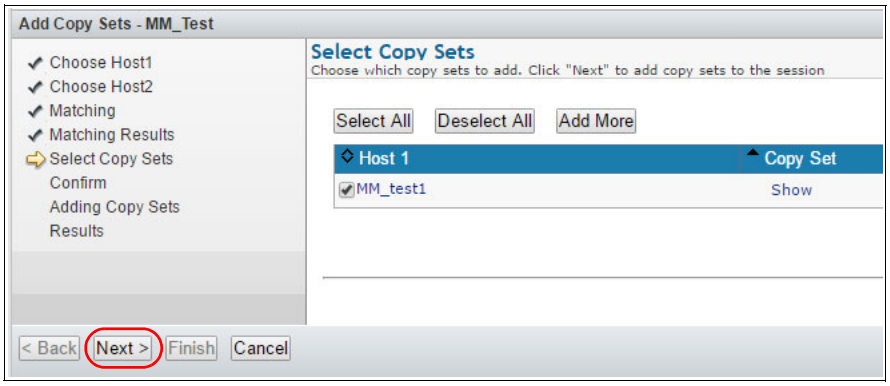

Figure 3-15 Select Copy Sets

5. Click **Select Copy Sets** (from both PRIMARY and DR_site). For this configuration, we have chosen to add the copy sets one by one, so, after adding the first copy set (shown in Figure 3-16) we add one more, repeating steps 2, 3, and 4.

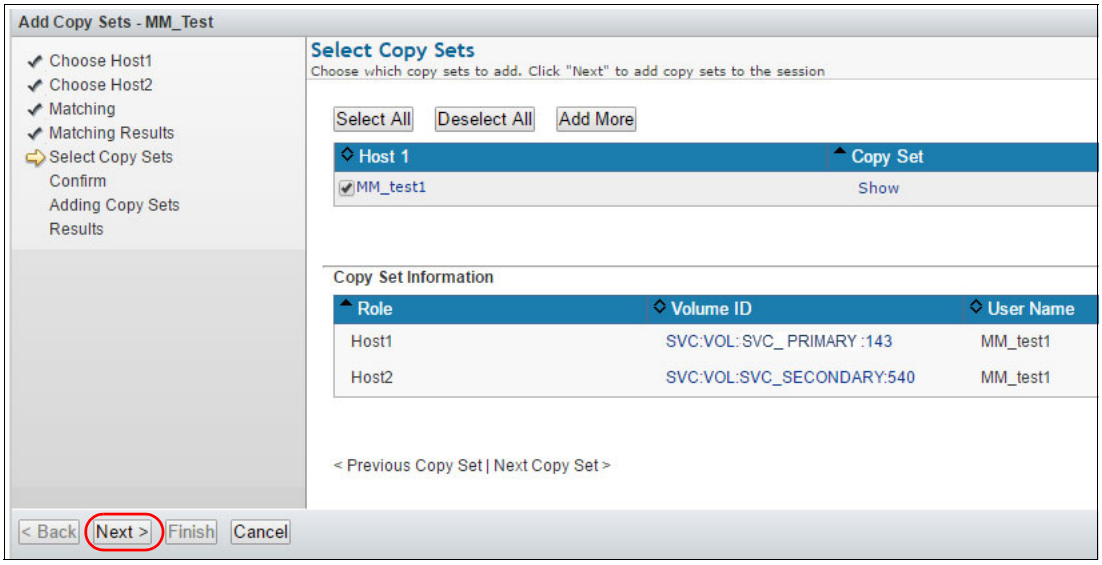

Figure 3-16 Copy sets for MM_Test session

6. When finished adding copy sets (H1 and H2), we finalize the process (**Confirm**), as shown in Figure 3-17 on page 41.

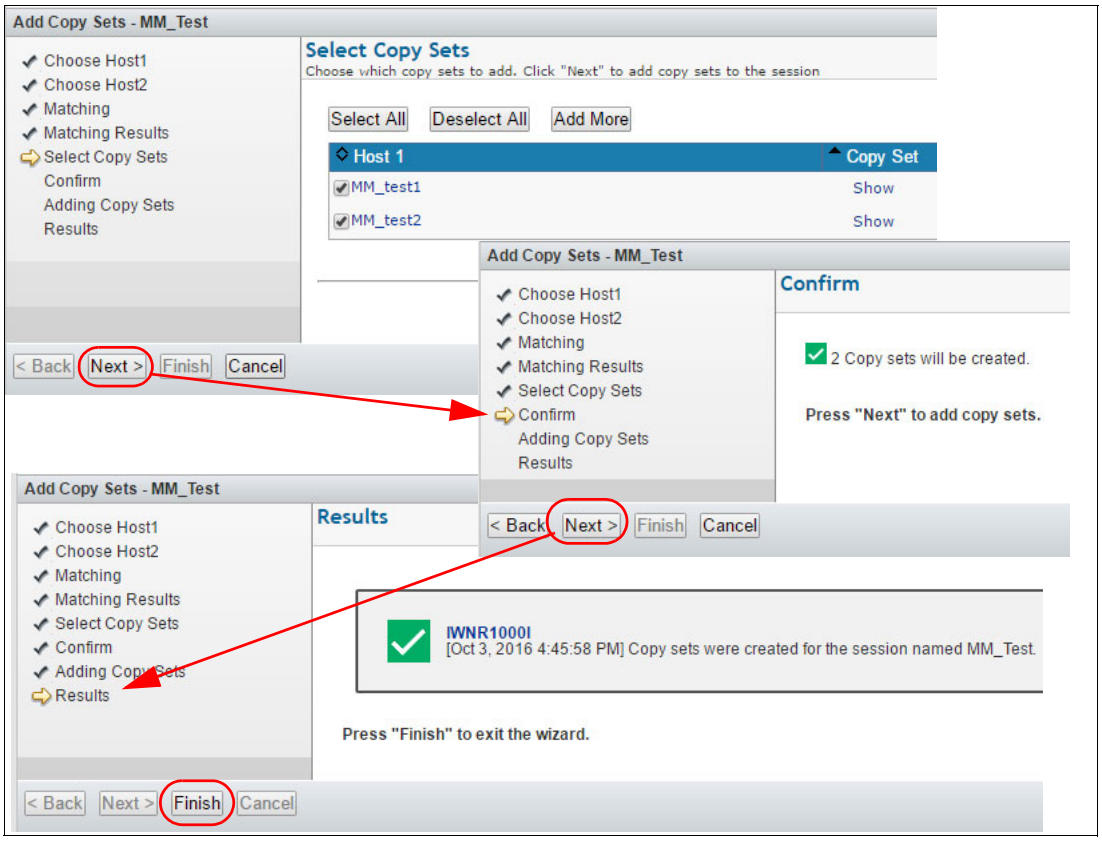

Figure 3-17 Finalizing the "Add Copy Sets" process

7. After adding the copy sets, the session is `Inactive` and no actions are performed at SVC level at this time. We visualize the session, as shown in Figure 3-18.

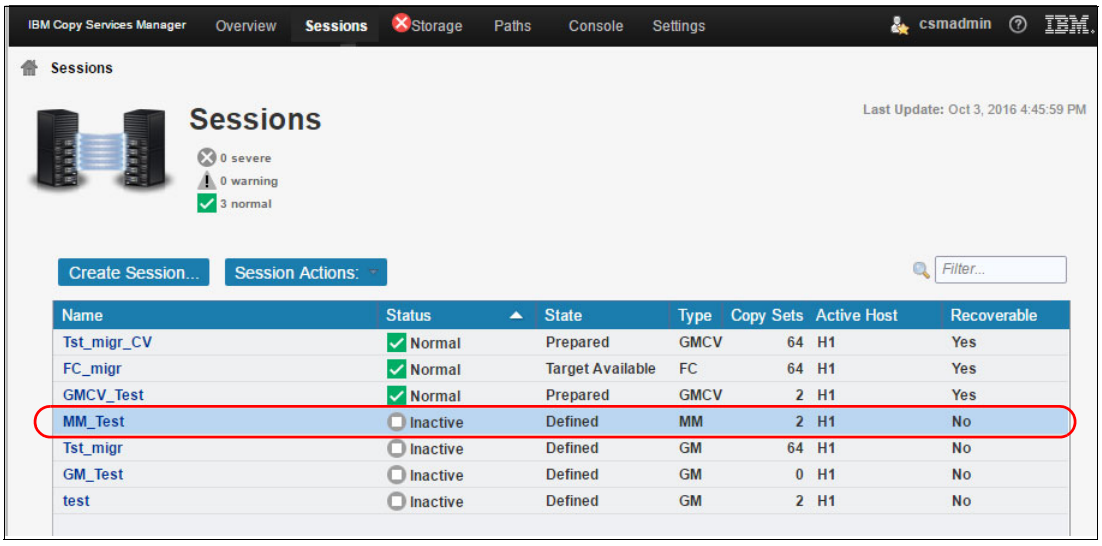

Figure 3-18 Session MM_Test defined (Inactive)

8. To start the session, we use **Session Actions** → **Commands** → **Start H1 -> H2** (Figure 3-19 on page 42). This action creates the *rcrelationship* in the SVC inside of a consistency group, as shown in Example 3-2 on page 43.

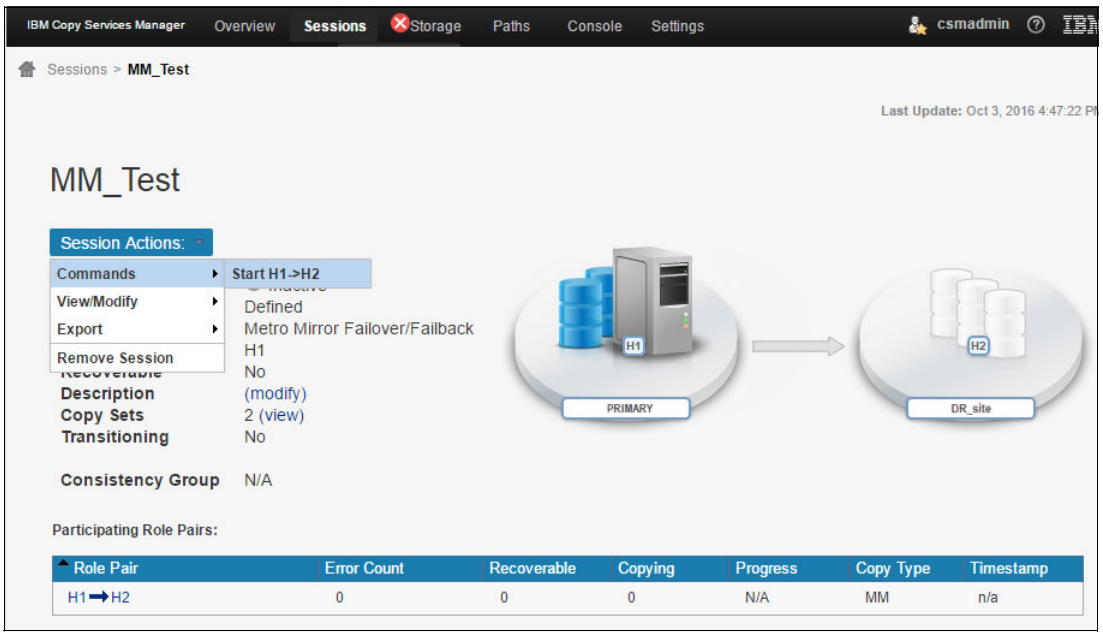

Figure 3-19 Start the session

9. Next, we confirm and monitor progress, as shown in Figure 3-20. When the relationship is in `Normal / Prepared` state, the copies are in consistent synchronized and H2 can be used for recovery actions (if necessary).

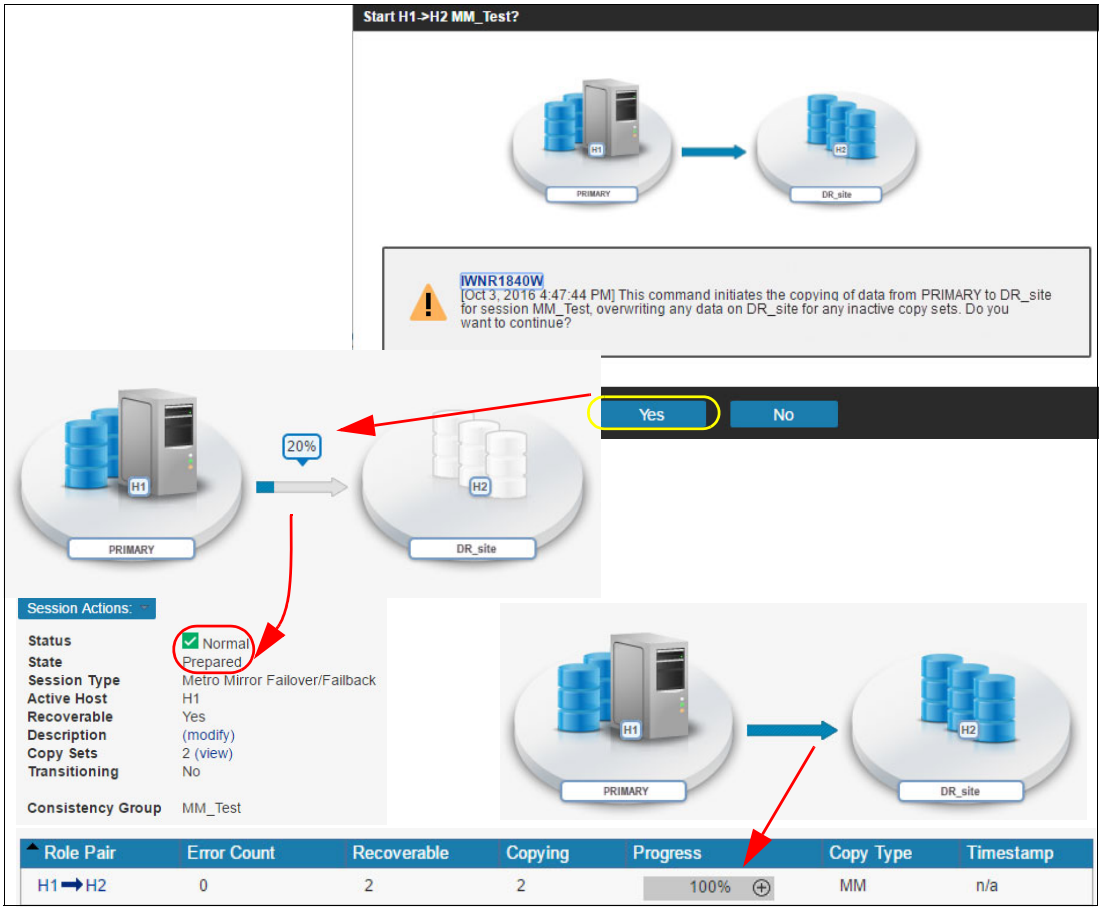

Figure 3-20 Confirm session start and monitor progress

10. We check the status of the relationships in the SVC interface, as depicted in Figure 3-21.

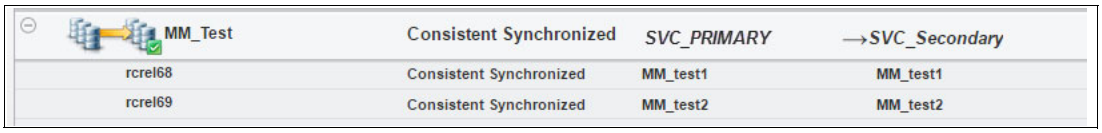

Figure 3-21 Metro Mirror relationship status in SVC GUI

For exemplification purposes, in Example 3-2 we also show the SVC commands that are issued by the IBM Copy Services Manager.

Example 3-2 Excerpt from SVC (PRIMARY) Audit Log

```
## Create MM relationship

1378         161003165318 csmuser                10.100.10.10   0      3
svctask mkrcconsistgrp -name MM_Test -cluster SVC_Secondary
1379         161003165318 csmuser                10.100.10.10   0      143
svctask mkrcrelationship -master 143 -aux 540 -cluster SVC_Secondary -consistgrp
MM_Test
1380         161003165318 csmuser                10.100.10.10   0      144
svctask mkrcrelationship -master 144 -aux 541 -cluster SVC_Secondary -consistgrp
MM_Test
1381         161003165319 csmuser                10.100.10.10   0
svctask startrcconsistgrp -primary master -force MM_Test
```

The session MM_Test is now available for data protection.

3.4 Metro Mirror failover and failback testing

In this section, we test the Metro Mirror (MM) configuration for failover and failback. The attached hosts configuration is not the subject of this document. However, for testing purposes we list the summary actions to be performed:

- Initial state:

  Hosts running connected to the SVC in PRIMARY site; MM replication from H1 -> H2 (copies are in consistent synchronized state)

- Failover:
  - Stop I/O on hosts in PRIMARY site
  - Suspend session (IBM CSM)
  - Recover volumes in secondary (DR_site) site (IBM CSM)
  - Access volumes in the SVC in DR_site (hosts in DR_site perform I/O)
- Failback:
  - Enable copy to PRIMARY site and start replication (H2 -> H1) (IBM CSM)
  - Stop I/O on hosts in DR_site site
  - Suspend session (IBM CSM)
  - Recover volumes in PRIMARY site (IBM CSM)
  - Access volumes in the SVC in PRIMARY site (hosts in PRIMARY site perform I/O)
  - Enable copy to DR_site and start replication (H1 -> H2) (IBM CSM)
- End of test: Hosts and MM relationship in initial state

3.4.1 Failover (controlled)

Note: We stop I/O on the host systems accessing volumes in the PRIMARY site. Use platform-specific tools or utilities for this action.

To perform a controlled failover,[2] complete the following steps:

1. When the I/O on the host has been stopped, we can suspend the MM session, as shown in Figure 3-22.

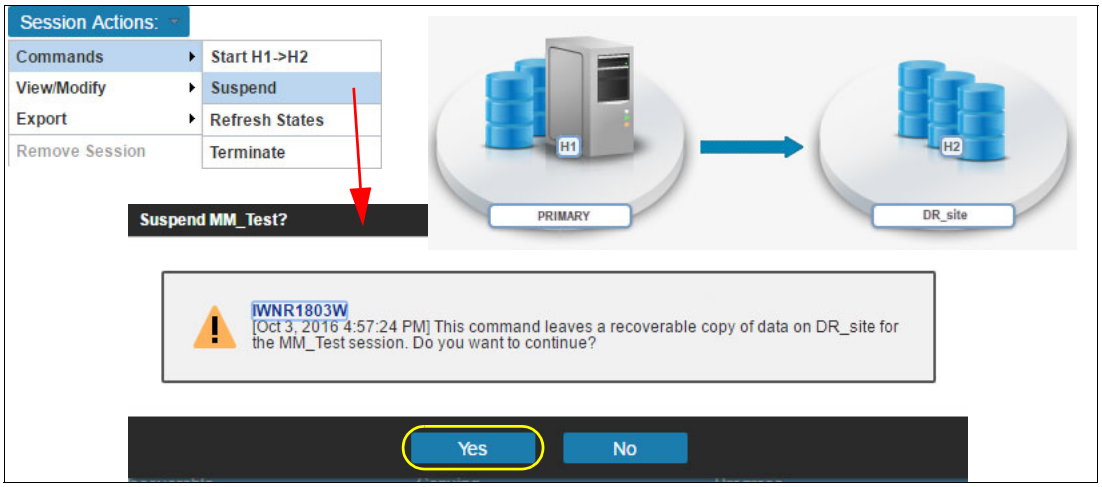

Figure 3-22 Suspending MM_Test session

Figure 3-23 shows the result of the Suspend action in CSM.

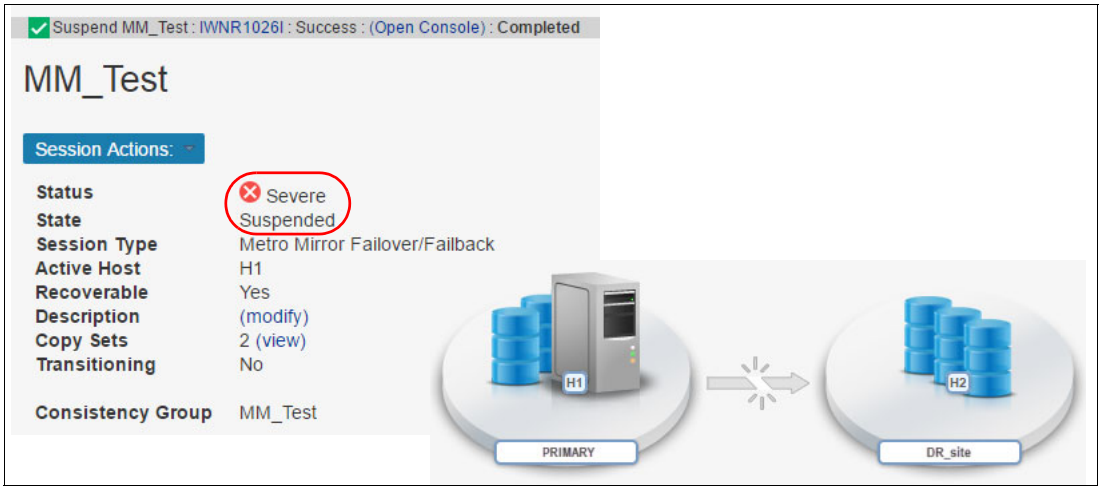

Figure 3-23 Session MM_Test suspended

2. We recover the volumes in DR_site for host access (H2), as shown in Figure 3-24 on page 45.

[2] Metro Mirror (synchronous replication) is block-level replication. Applications are unaware of the moment of suspending a relationship. In case of primary storage failure, the target can be made available for I/O. However, application-level recovery procedures might be required.

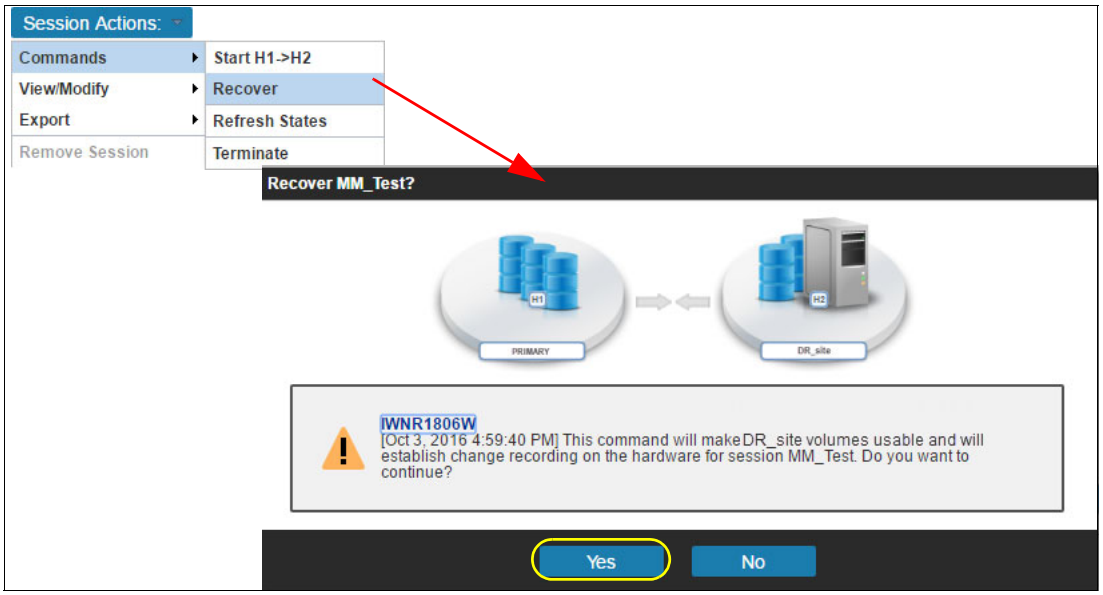

Figure 3-24 Recovering volumes in DR_site for host access

3. When the session is in Normal / Target Available state, we can activate host systems in DR_site to perform I/O (H2), as shown in Figure 3-25.

> **Note:** It is out of scope for this material to describe details about how to activate host systems and access volumes from these. Each operating system has platform-specific methods and tools or utilities for this purpose.

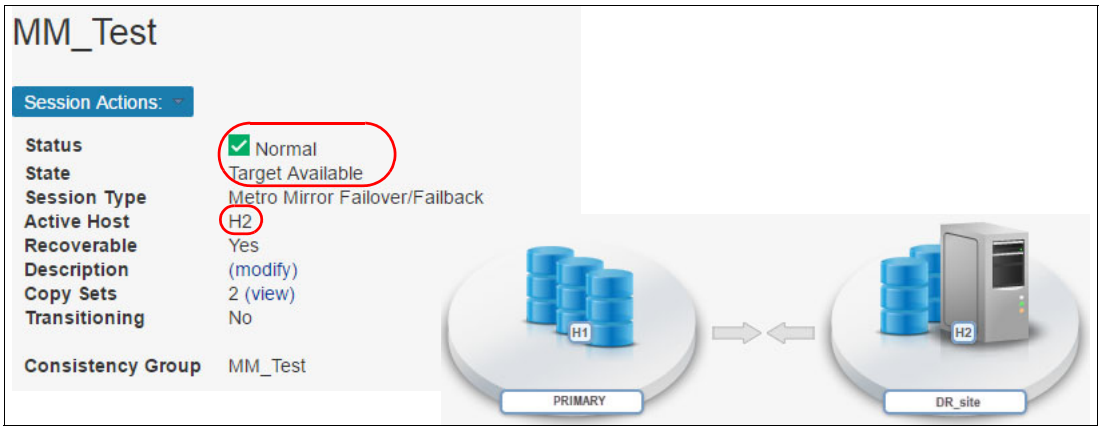

Figure 3-25 Session MM_Test prepared for I/O in DR_site

4. After I/O has been performed on volumes in DR_site, we can enable (if wanted) reverse replication, as shown in Figure 3-26 on page 46.

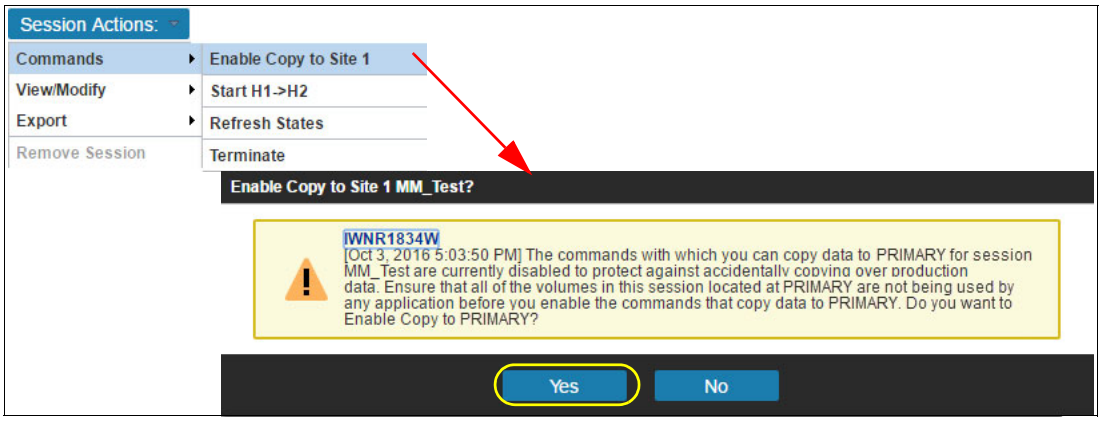

Figure 3-26 Enabling copy to PRIMARY site

At this stage, reverse replication has been enabled, but not yet started, as shown in Figure 3-27.

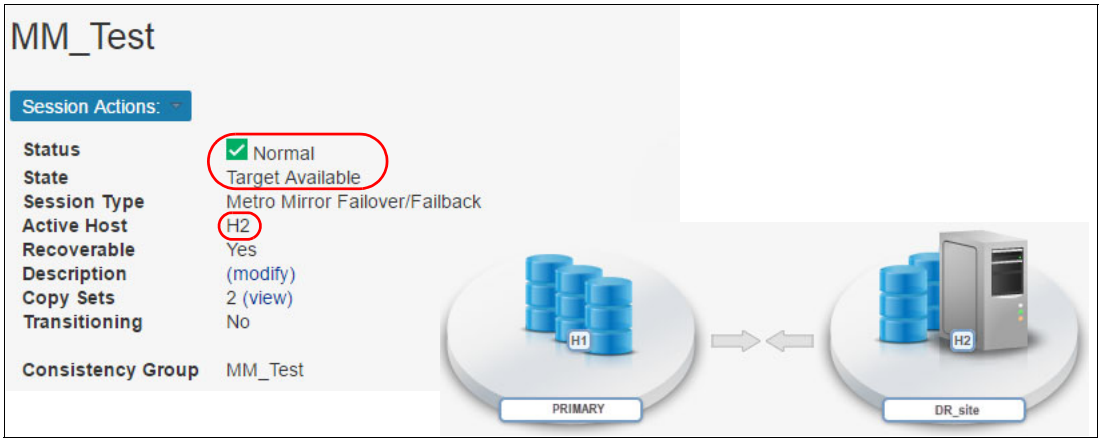

Figure 3-27 Reverse replication prepared

5. We start replication to the PRIMARY site, as shown in Figure 3-28.

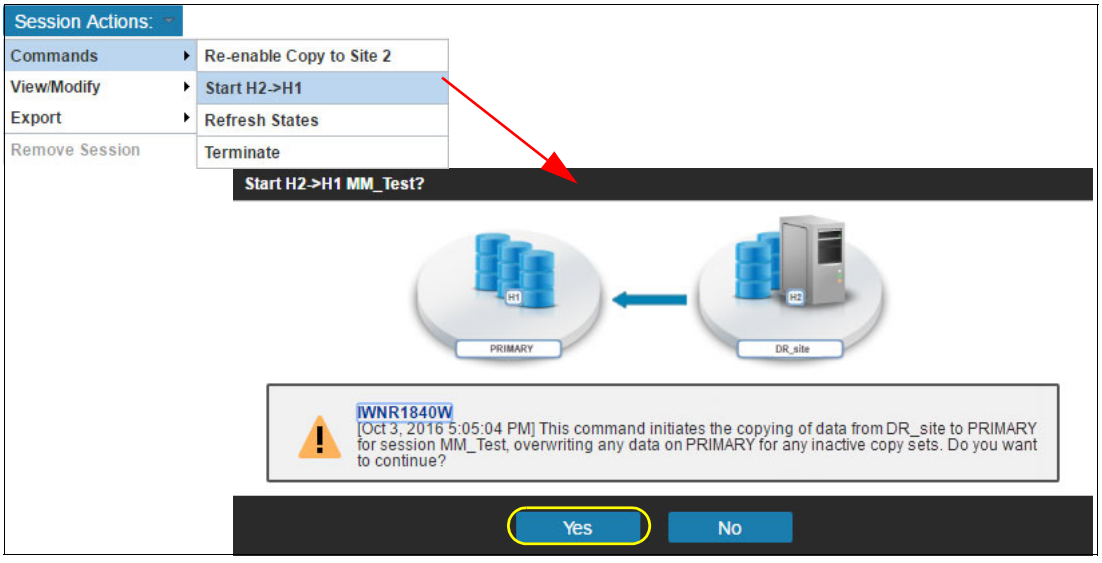

Figure 3-28 Enabling H2 -> H1 replication

3.4.2 Failback (controlled)

When the volumes are synchronized (relationship in Normal / Prepared state, as shown in Figure 3-29), we can start the failback to PRIMARY site operation. At this time, the host system in DR_site is active and performs I/O.

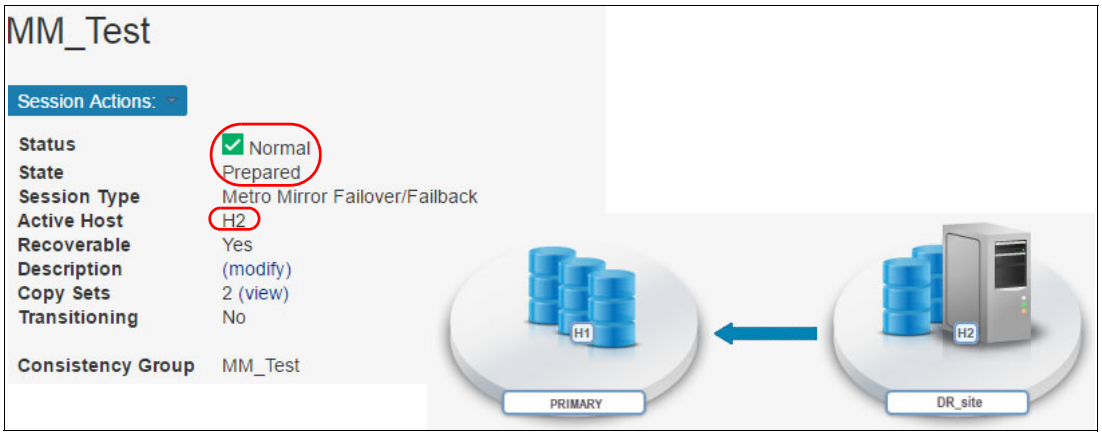

Figure 3-29 H2 -> H1 replication active

6. For data consistency,[3] we stop the I/O to the volumes in DR_site in preparation for the failback operation. After the host I/O has been stopped, we can suspend the H2 -> H1 relationship, as shown in Figure 3-30.

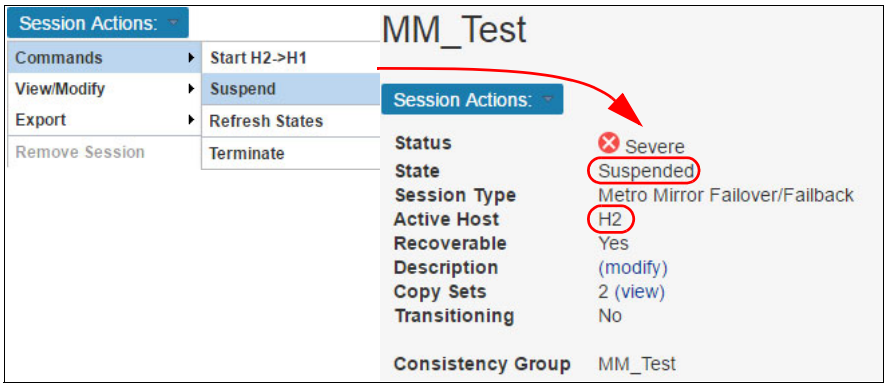

Figure 3-30 Suspending H2 -> H1 replication

[3] Metro Mirror (synchronous replication) is block-level replication. Applications are unaware of the moment of suspending a relationship. In case of active storage failure, the target volumes can be made available for I/O. However, application-level recovery procedures might be required.

7. To render the volumes in PRIMARY site accessible for the host systems, we perform recovery, as shown in Figure 3-31.

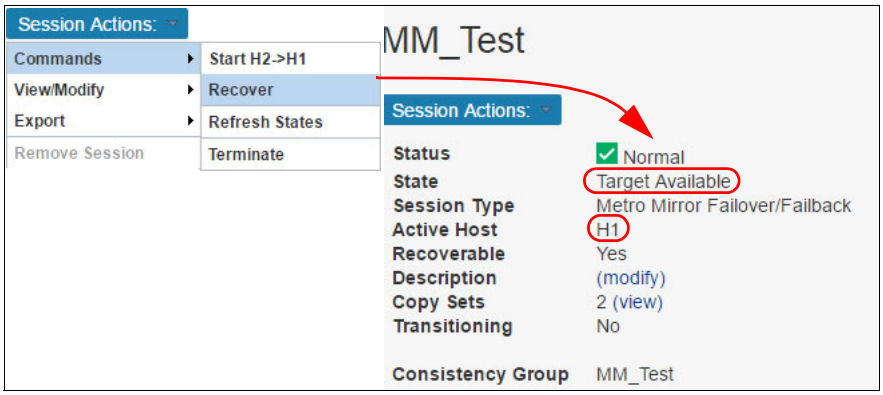

Figure 3-31 Recover H1 volumes for host access

8. At this time we can activate host systems in the PRIMARY site and access the volumes for I/O, and enable copy to DR_site (for continuous protection in case of storage failure), as shown in Figure 3-32.

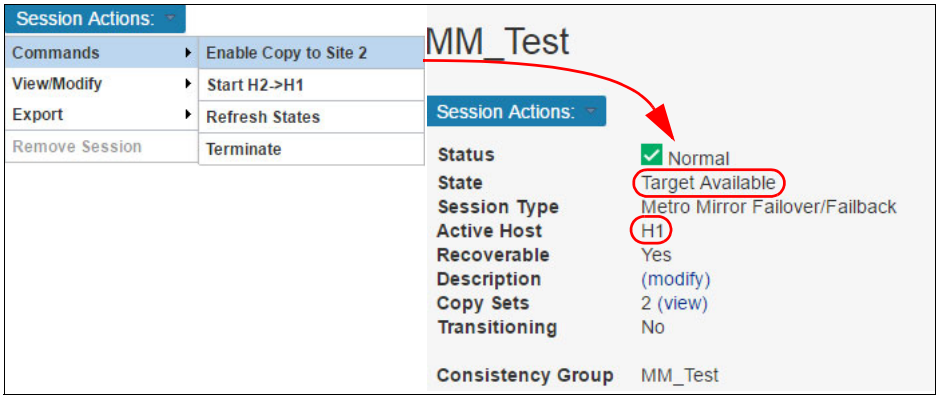

Figure 3-32 Enabling H1 -> H2 replication

9. We start the H1 -> H2 replication, as shown in Figure 3-33.

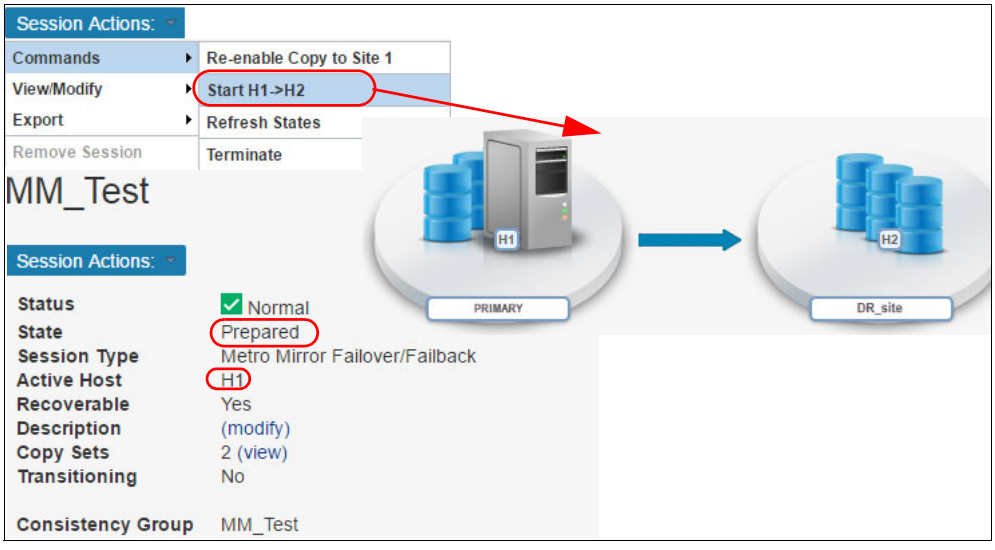

Figure 3-33 Failback complete: MM_Test session "Normal / Prepared" for H1 -> H2

The status has been reverted to the original status (H1 -> H2 replication).

You can also check the SVC Activity log (for both PRIMARY and DR_site) for the commands that have been issued by IBM Copy Services Manager. Example 3-3 shows the commands issued to the SVC cluster in PRIMARY site.

Example 3-3 SVC in PRIMARY site: Audit log excerpt for Failover / Failback

```
## FailOver MM

1382         161003170315 csmuser                10.100.10.10   0
svctask stoprcconsistgrp MM_Test
1383         161003171029 csmuser                10.100.10.10   0
svctask startrcconsistgrp -primary aux -force MM_Test
```

3.5 Enabling practice for Metro Mirror session

In this scenario, we use an existing (active) Metro Mirror (MM) relationship and convert it to Metro Mirror with Practice (MM with Practice).

Tip: It is important to follow the steps in the right order to avoid application downtime and maximize data availability.

Note: During this change, the copy sets part of the initial MM session continues to be replicated by the underlying HW (IBM SVC), and is in Consistent Synchronized state.

Overview

The following is an overview of the steps we have followed to change an existing MM relationship to *MM with Practice*:

1. Based on existing session (identified as a consistency group in the SVC), we create in the DR_site FlashCopy volumes (to be used for practice) for *each* target volume in the MM session. This step is done on the SVC directly (GUI or CLI), no CSM involvement.
2. While the MM relationship is in the Normal / Prepared state, we remove the MM relationship from Copy Services Manager (H1 -> H2) *without removing the base hardware relationship* on the storage system. This leaves the copy sets in the base hardware (storage) untouched (in all copy sets in *Consistent Synchronized* state).
3. In CSM, we define a new session of type *Metro Mirror Failover/Failback with Practice* and “populate” it with the designated volumes (H1, H2, and I2).
4. We start the new (MM with Practice) session and monitor it.

3.5.1 Scenario

The starting point is the existing MM_Test session, as shown in the SVC GUI depicted in Figure 3-34.

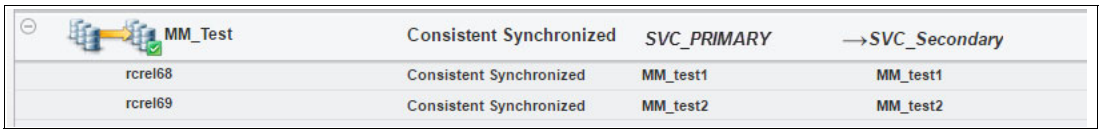

Figure 3-34 Metro Mirror relationship status (Consistency Group MM_Test) in SVC GUI

Complete the following steps:

1. We create the FlashCopy volumes in DR_site based on the existing relationship (represented as a Consistency Group in the SVC).

   Note: It is important to create the volumes before removing the session from CSM, because this action can be scripted based on the SVC CG information. This is especially useful for sessions that consist of many copy sets.

   If you remove the session from the CSM (without removing the relationships from the underlying hardware), the Consistency Group information will not be available in SVC (CLI or GUI), making it difficult to identify the copy sets.

   The FlashCopy volumes (designated as H2) in the Secondary site (DR_site) can be created from the GUI, but this can be cumbersome for sessions with a large number of volumes.

   We have used a script and SVC CLI. The script shown in Example 3-4 can be used for similar actions. This script creates the following volumes in the Secondary site (DR_site):

   ```
   MM_Test1_prct
   MM_Test2_prct
   ```

Example 3-4 Script for creating volumes (to be used for FlashCopy) for each H2 volume in an MM session

```
CG=MM_Test
POOL=IT_test
lsrcrelationship -filtervalue consistency_group_name=$CG -nohdr |while read -a rc
do
lsvdisk -nohdr -bytes -filtervalue id=${rc[8]} | while read -a v
do
echo Creating volume ${v[1]}"_prct" with virtual size ${v[7]} bytes
mkvdisk -autoexpand -iogrp ${v[3]} -mdiskgrp $POOL -name ${v[1]}"_prct" -rsize 0 -size ${v[7]} -unit b
done
done
```

Hint: The script can be used for other sessions by adjusting the `CG` and `POOL` variables.

Attention: SAN Volume Controller scripting and CLI skills are required to perform this action. Make sure that you understand the actions in the script before running it in your environment.

2. At this point, we proceed to defining the new session, as shown in Figure 3-35 on page 51.

   Note: Defining a new session does not impact the existing MM_Test session. By defining the new session into CSM, we minimize the transition time from MM to MM with practice session.

   Also, data replication is not affected during the transition of the copy sets from MM to MM with practice sessions (copy sets remain in a *Consistent Synchronized* state).

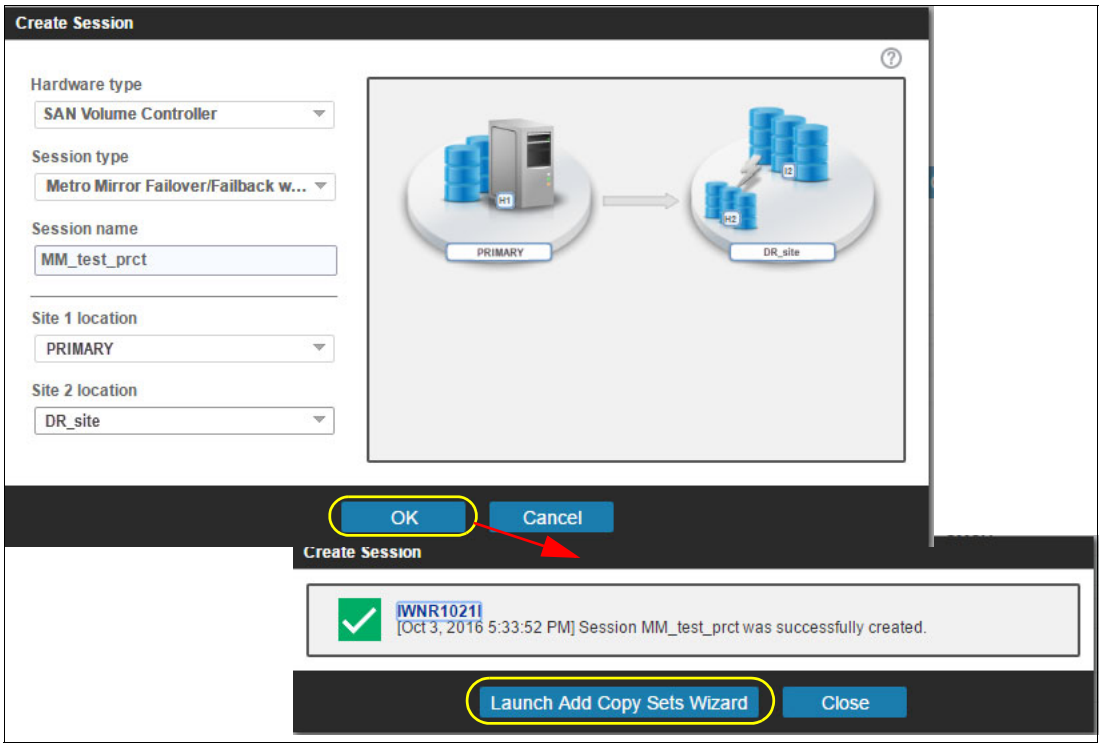

Figure 3-35 Creating a new session for MM w/ Practice

3. In the Add Copy Sets wizard, we define H1, as shown in Figure 3-36.

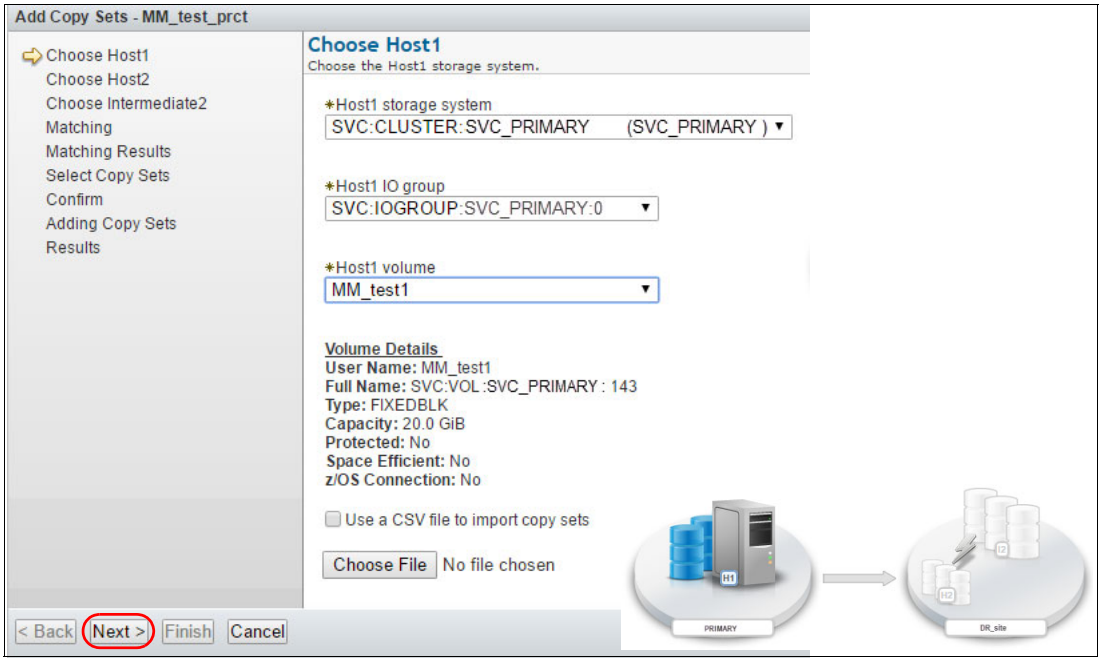

Figure 3-36 Defining H1 for first copy set

4. Next, we define H2, as shown Figure 3-37.

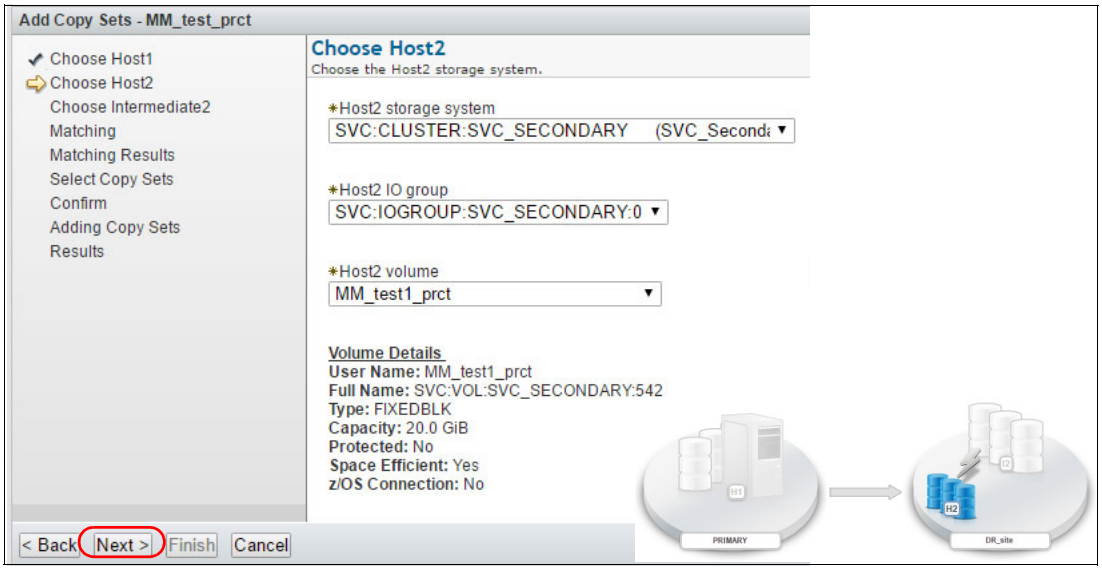

Figure 3-37 Defining H2 for the first copy set

5. Next, we define I2 (intermediate), as shown in Figure 3-38.

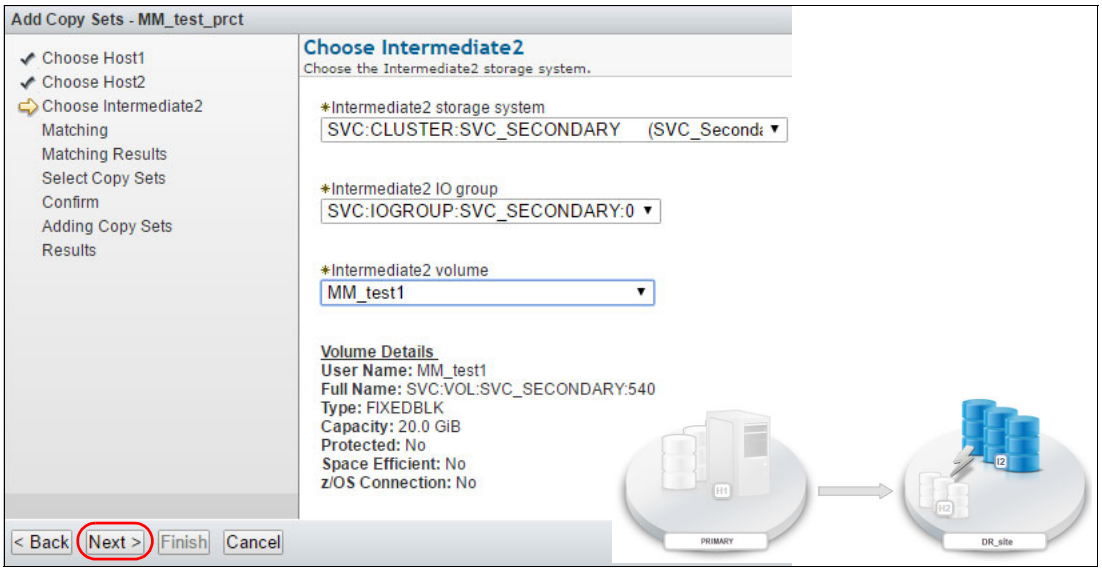

Figure 3-38 Defining I2 for first copy set

6. Repeat the procedure for additional copy sets. When all wanted copy sets ($H1_{2..n}$, $H2_{2..n}$, $I2_{2..n}$) have been defined, you select all copy sets and confirm the action, as shown in Figure 3-39.

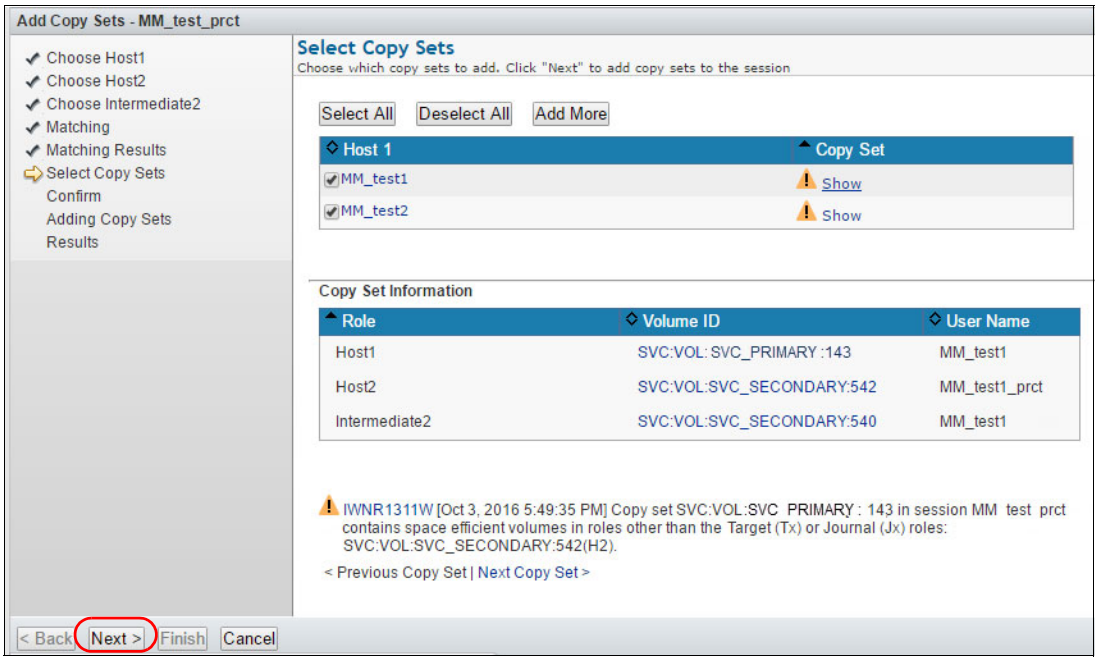

Figure 3-39 Confirming copy sets to be added to MM_Test_prct session

7. When copy sets have been added to the session, the results are shown. Click **Finish** to exit the Add Copy Sets wizard, as shown in Figure 3-40.

> **Note:** No commands are issued to the underlying hardware (SVC) at this time. The session has been defined into CSM, but we do not activate it before removing the previous MM session (MM_Test).

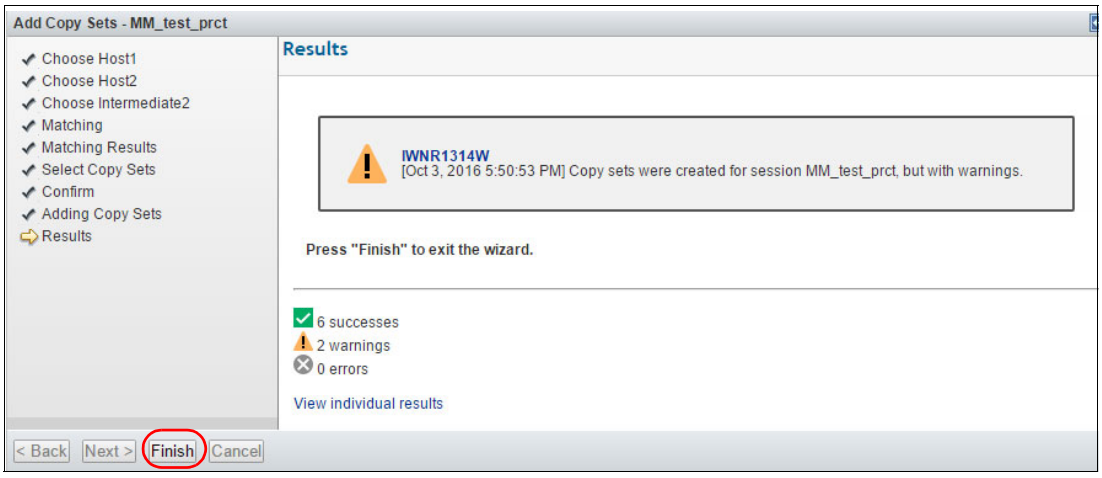

Figure 3-40 Add Copy Sets wizard results

8. To start removing the current session select **Session Action** → **View / Modify** → **Remove Copy Sets**, as shown in Figure 3-41.

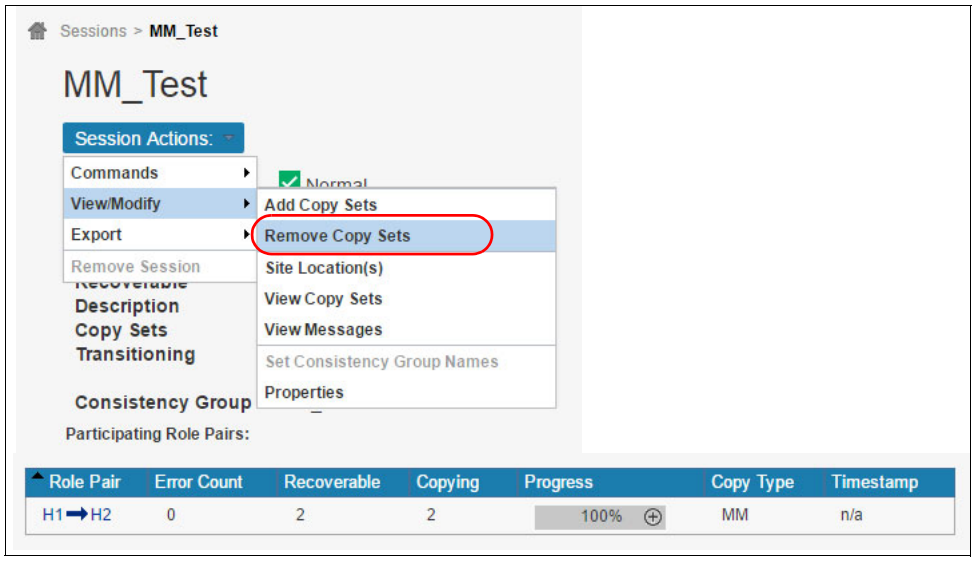

Figure 3-41 Initiating the Remove Copy Sets action

9. We select the copy sets to be removed (MM_Test1 and MM_Test2 in this case), as shown in Figure 3-42.

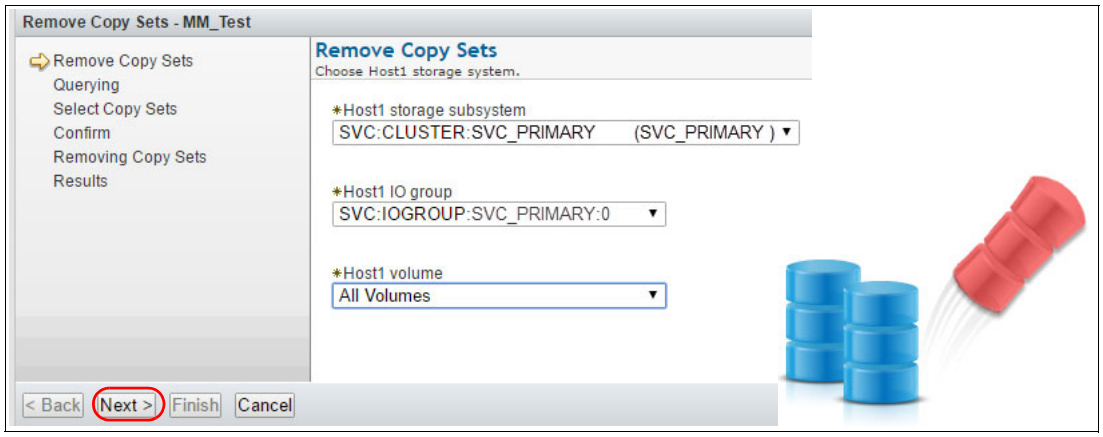

Figure 3-42 Remove Copy Sets wizard

10. Select the copy sets (in this case, all: MM_Test1 and MM_Test2), as shown in Figure 3-43.

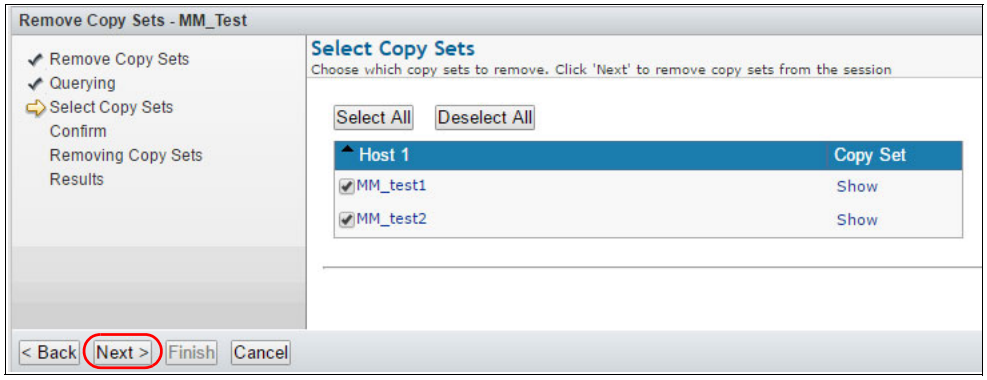

Figure 3-43 Selecting copy sets to be removed

11. Make sure that you select **Yes, keep base hardware relationships on the storage system** and then confirm. Check the results and click **Finish** to exit the wizard, as shown in Figure 3-44.

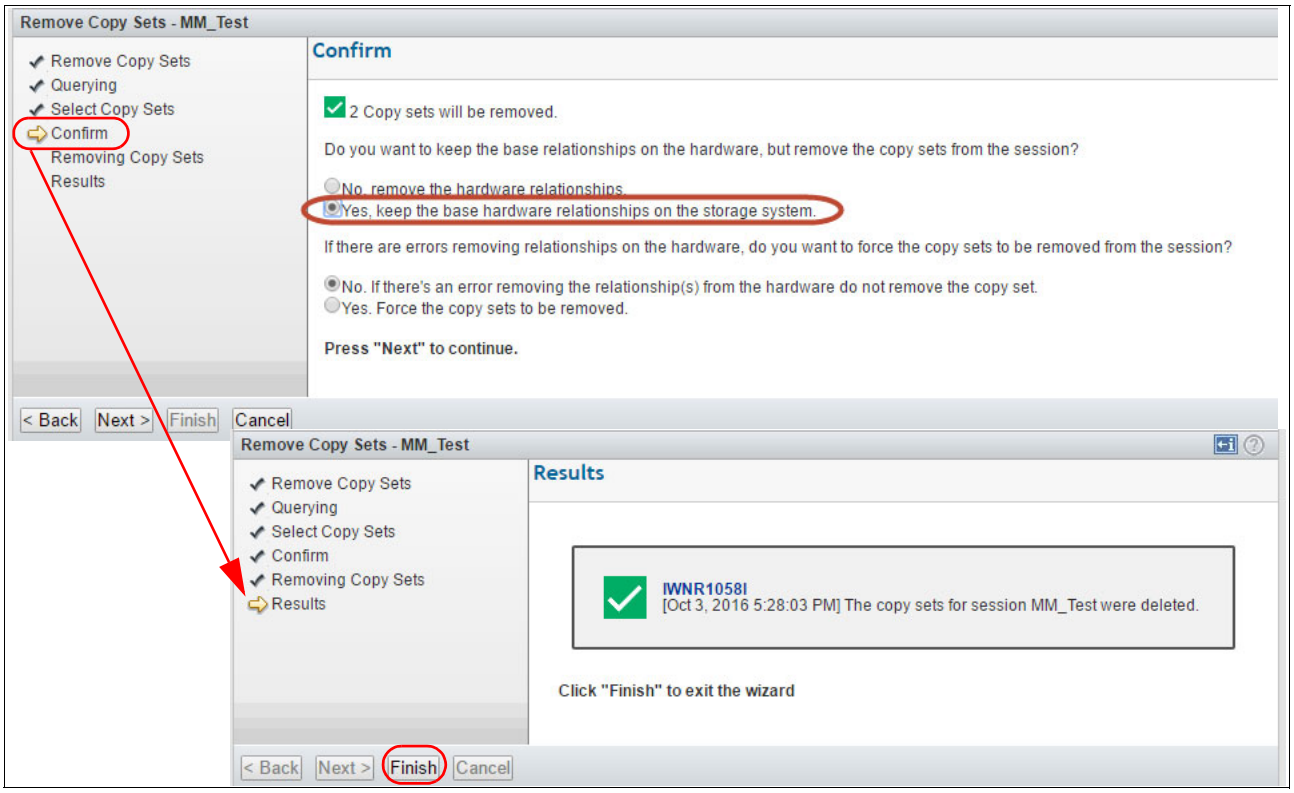

Figure 3-44 Finalizing the copy sets removal

12. You can check the underlying hardware to verify that the MM relationships are available and in the Consistent Synchronized state, but are not part of a consistency group anymore, as shown in Figure 3-45.

Figure 3-45 SVC view of MM relationship removed from CSM

13. At this point, we can start the previously prepared relationship (MM_Test_prtc), as shown in Figure 3-46.

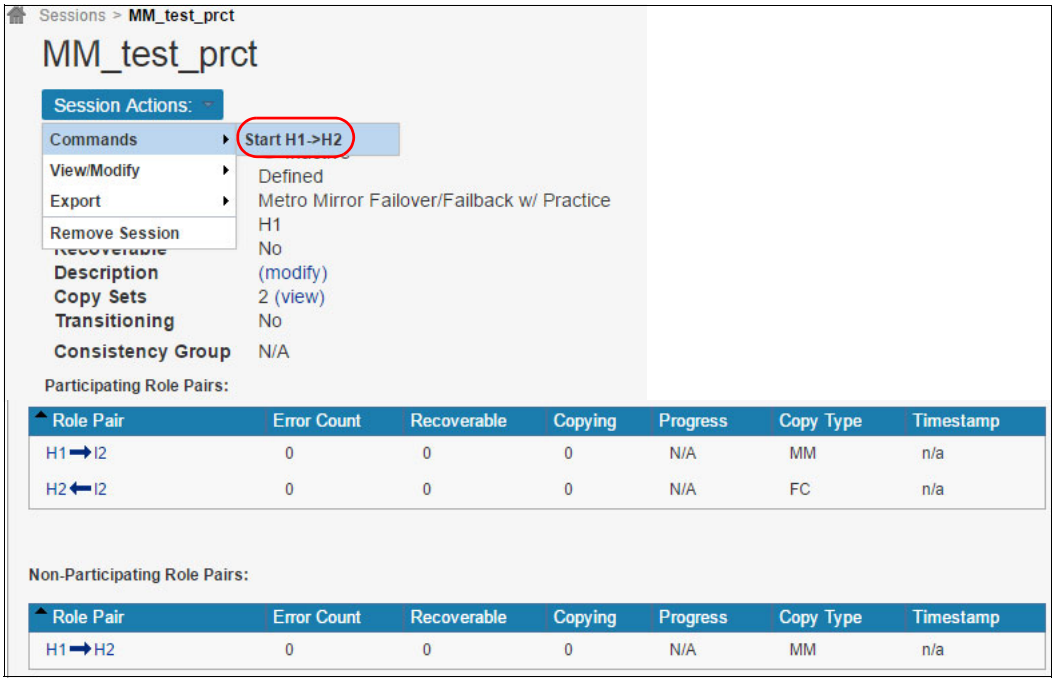

Figure 3-46 Starting the session

14. Confirm by clicking **Yes**, as shown in Figure 3-47.

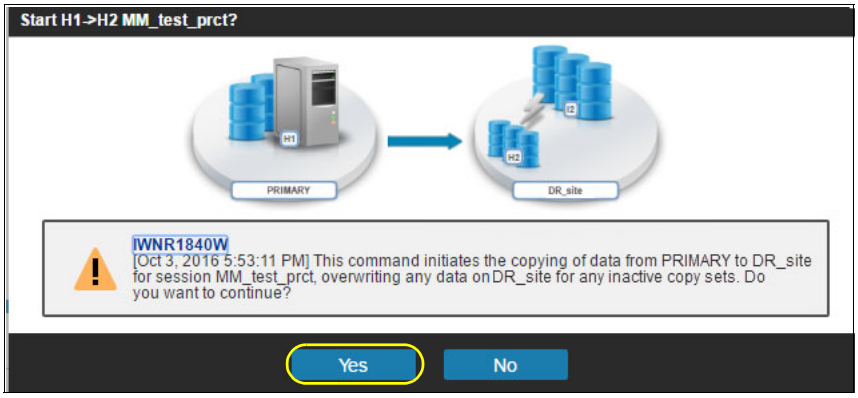

Figure 3-47 Confirming session start

15. Check progress, as shown in Figure 3-48. When the H1 -> I2 has reached 100%, the session is ready for practice DR configuration.

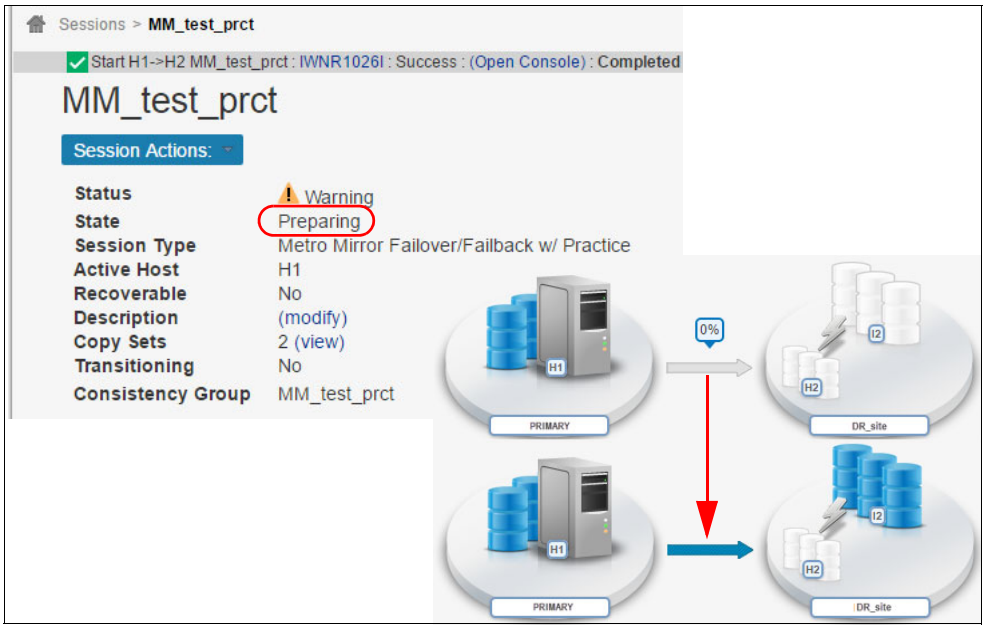

Figure 3-48 Session status and progress

16. To prepare the practice volumes (H2) for host access (in secondary site), select **Session Actions** → **Commands** → **Flash**, as shown in Figure 3-49.

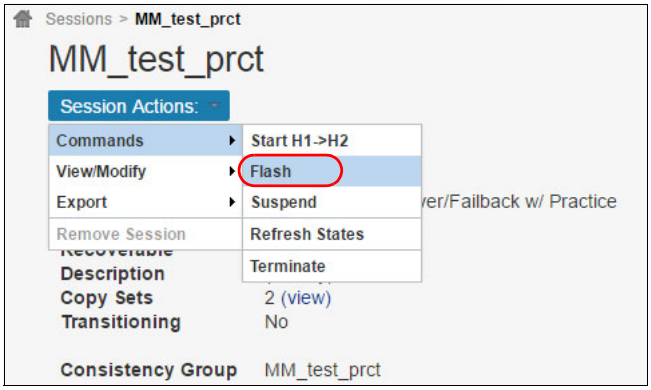

Figure 3-49 Preparing H2 for practice (host access in DR_site)

17. Confirm as shown in Figure 3-50, and H2 can be accessed.

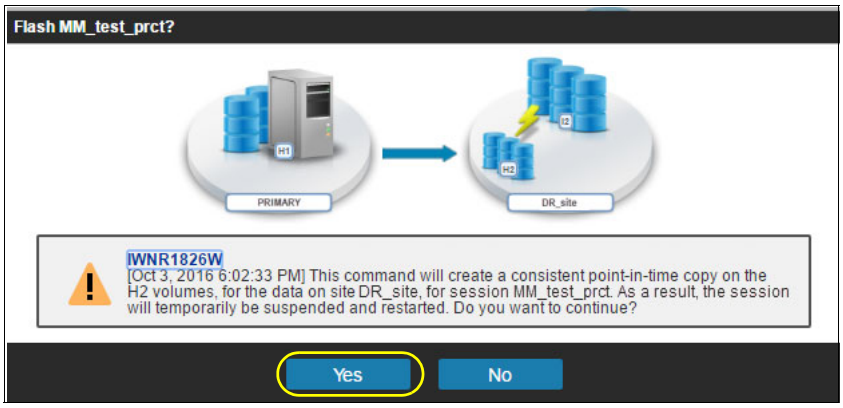

Figure 3-50 Confirming FlashCopy

Note: It is out of scope for this material to describe details about how to activate host systems and access the volumes from these. Each operating system has platform-specific methods and tools or utilities for this purpose.

4

Global Mirror Sessions with IBM Spectrum Virtualize-based storage

In this chapter, we provide an overview of the tasks that are necessary to successfully deploy IBM Copy Services Manager within your environment.

We briefly describe supported platforms for CSM server deployment and considerations for CSM server replication options for deploying CSM in a disaster recovery environment, and security considerations for typical deployments. In addition, the chapter describes exporting and importing sessions and copy sets.

In this chapter, we describe the following scenarios:

- Configuring and managing Global Mirror sessions
- Global Mirror failover and failback
- Converting a Global Mirror session to a Global Mirror with Change Volumes

4.1 Configuring and managing Global Mirror sessions

In this scenario, we create a Global Mirror (asynchronous replication) session between the two sites in our test environment (PRIMARY and DR_site: See Figure 3-1 on page 32). The focus in this document is on IBM Copy Services Manager capabilities.

Note: IBM Copy Services Manager does *not* provision (create, delete, or modify) volumes, pools, or other storage entities. All required volumes must be configured before defining and activating the IBM CSM (copy services) sessions and consistency groups.

4.1.1 Creating a Global Mirror session

The scenario assumes that the necessary volumes and the replication links on the underlying storage hardware (we used IBM SVC for this scenario) are already configured, and we use IBM CSM to create and manage a Global Mirror session.

In the initial phase, the volumes are accessed by one or more hosts in PRIMARY site (H1).

1. We create a session named `GM_Test`, as shown in Figure 4-1.

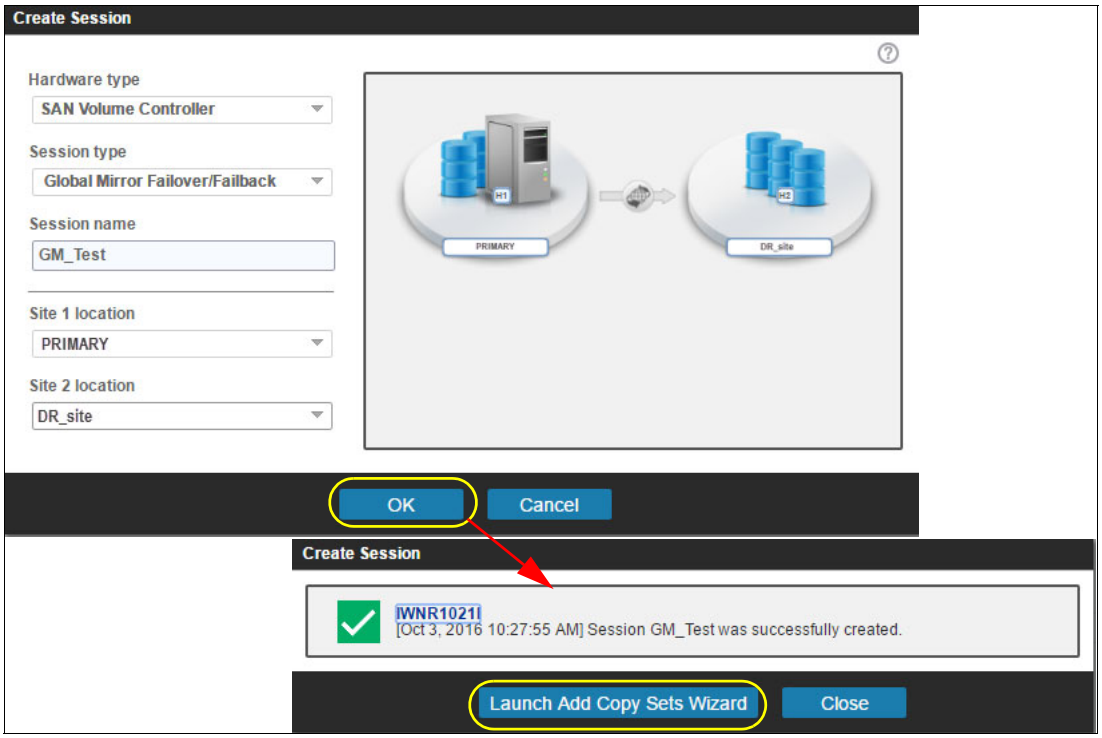

Figure 4-1 Creating the GM session and launching the Add Copy Sets wizard

2. The Add Copy Set wizard guides us to select H1 volumes, as shown in Figure 4-2.

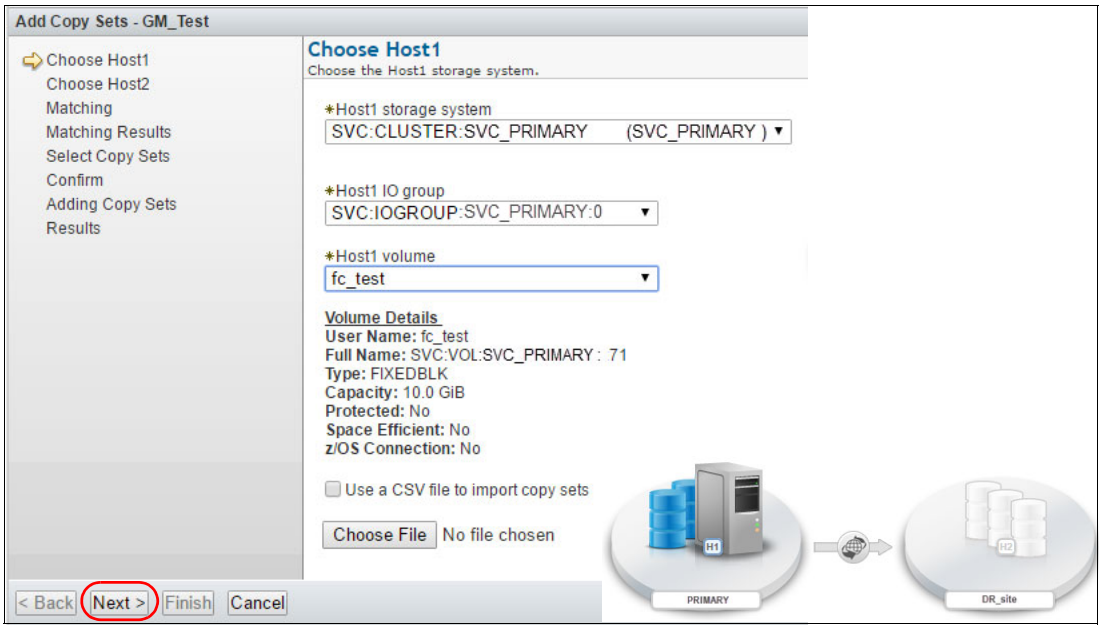

Figure 4-2 Adding H1 volumes

3. Next, we add the H2 volumes, as shown in Figure 4-3.

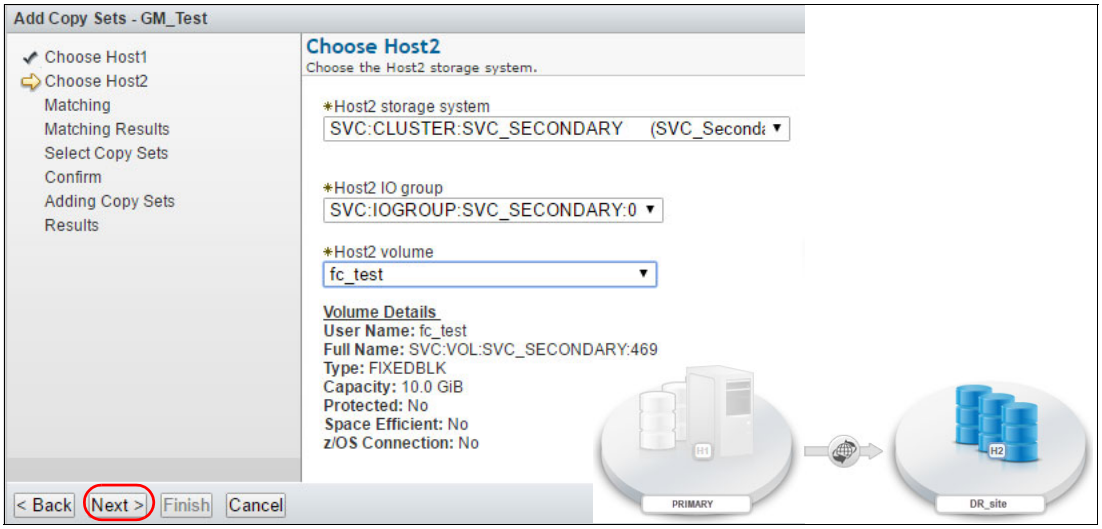

Figure 4-3 Adding H2 volumes

4. When all copy sets (volumes) have been added, double-check and finalize, as shown in Figure 4-4.

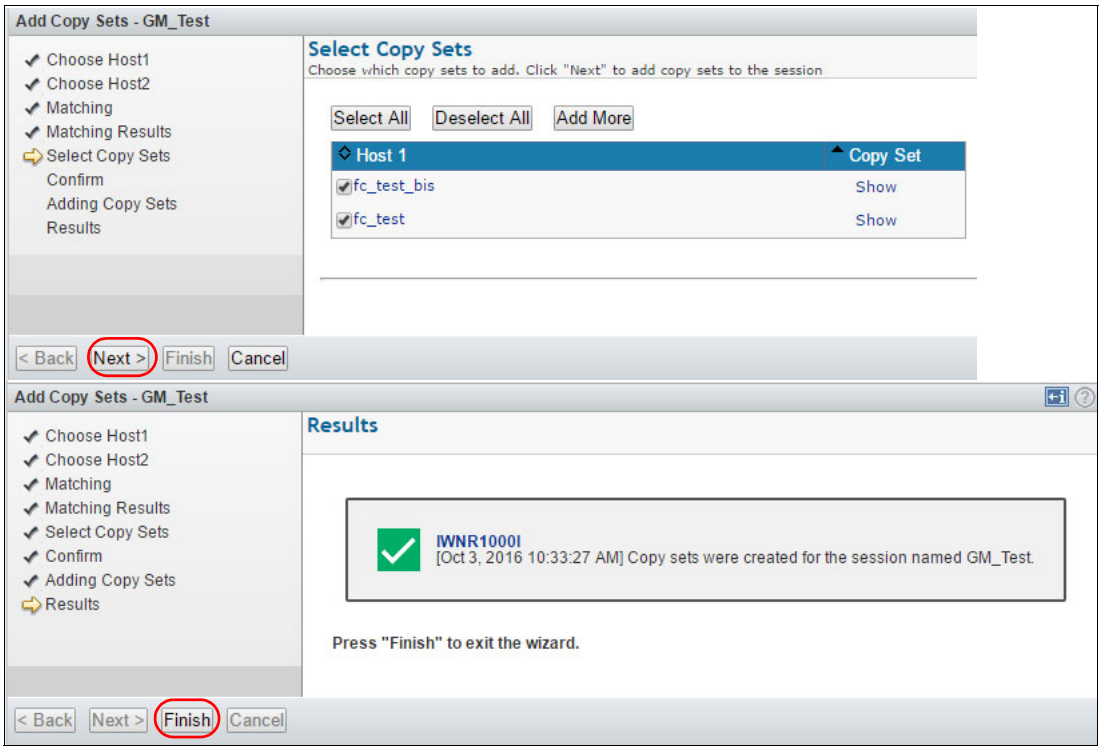

Figure 4-4 Finalizing session

5. The new session (GM_Test) has been created, but is inactive at this time, as shown in Figure 4-5.

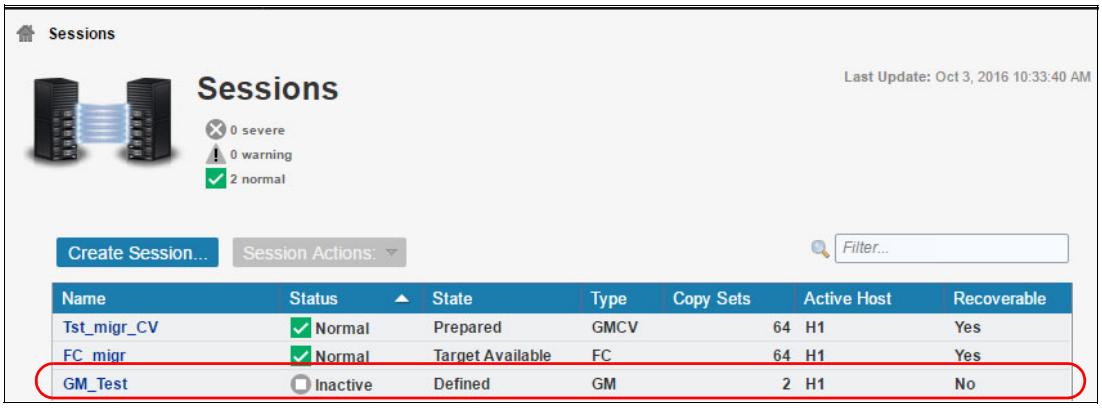

Figure 4-5 Session GM_Test defined

6. When the copy sets have been added to the GM_Test session, we start the replication to DR_site (H1 -> H2), as shown in Figure 4-6.

Figure 4-6 Starting replication

7. Confirm session start, as shown in Figure 4-7.

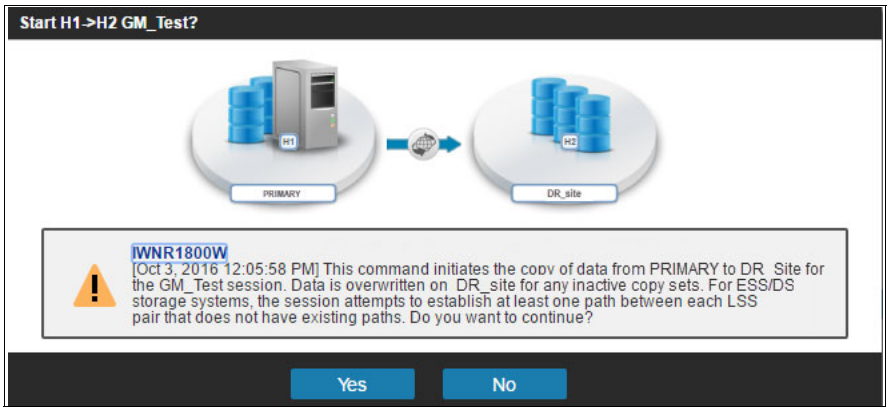

Figure 4-7 Confirming session start

8. The session is represented as a Consistency Group in the underlying hardware (IBM SVC), as shown in Figure 4-8.

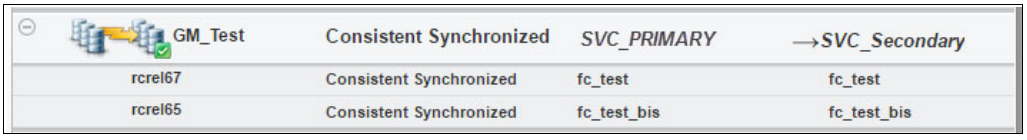

Figure 4-8 Session GM_Test represented in underlying HW

9. When the replication process has synchronized the copy sets, the session will be in Normal / Prepared state, as shown in Figure 4-9.

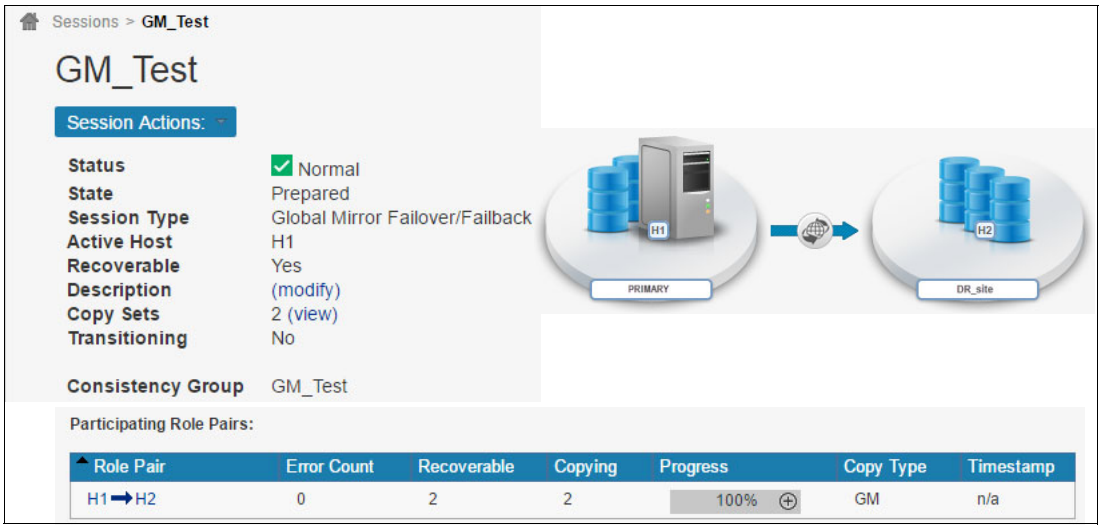

Figure 4-9 Session synchronized (Normal/Prepared)

4.2 Global Mirror failover and failback

In this section, we describe the tests we have performed in our environment for GM sessions.

4.2.1 Testing Global Mirror failover (controlled)

In this section, we test the Global Mirror failover and failback operations[1]. The starting point is with the session in “Normal / Prepared” state, as shown in Figure 4-9.

The objective of this test is to verify that the data can be made accessible to one or more hosts in secondary site in case disaster recovery is needed.

[1] Global Mirror (asynchronous replication) is storage block level replication. Applications are unaware of the moment of suspending a relationship. In case of primary storage failure, the target can be made available for I/O. However, application level recovery procedures might be required.

We follow these steps:

1. We suspend the GM session, as shown in Figure 4-10.

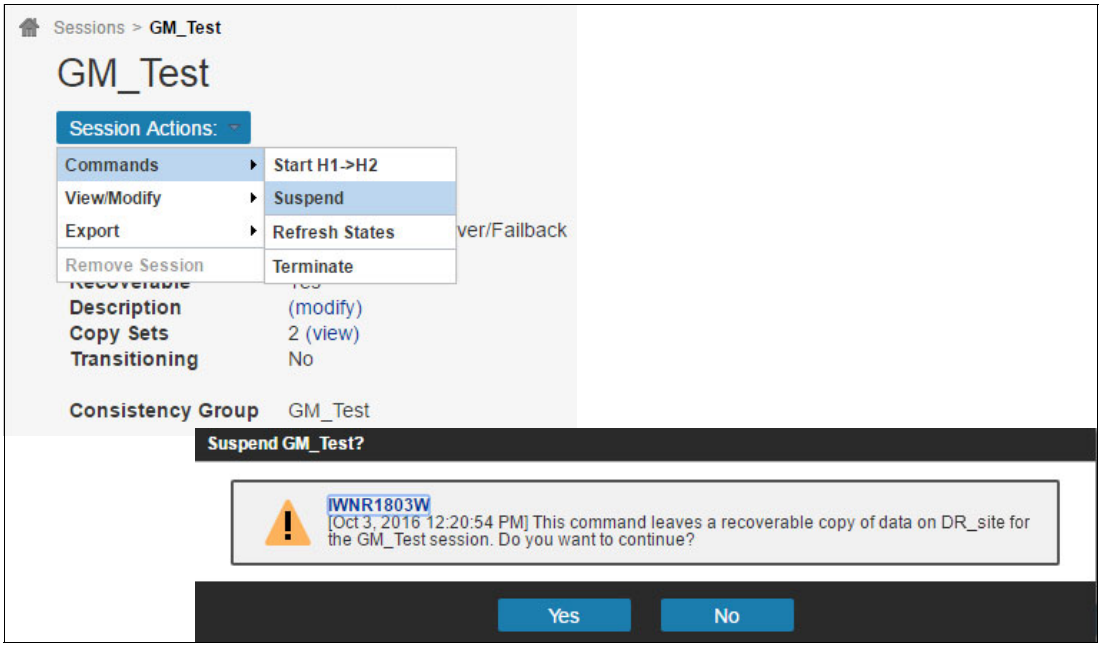

Figure 4-10 Suspending GM session

2. Before suspending the session, we stop I/O from hosts in PRIMARY site (I/O to H1). The suspended session is shown in Figure 4-11.

> **Note:** We stop I/O on the host systems accessing H1 volumes using platform-specific tools or utilities for this action.

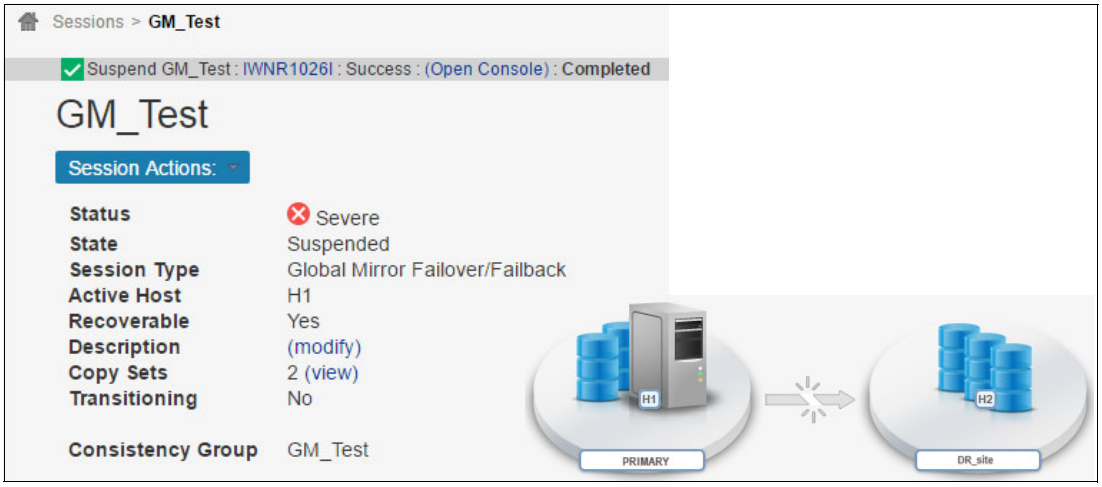

Figure 4-11 Session suspended

3. To make the volumes in DR_site available for host access, we perform session recovery, as shown in Figure 4-12.

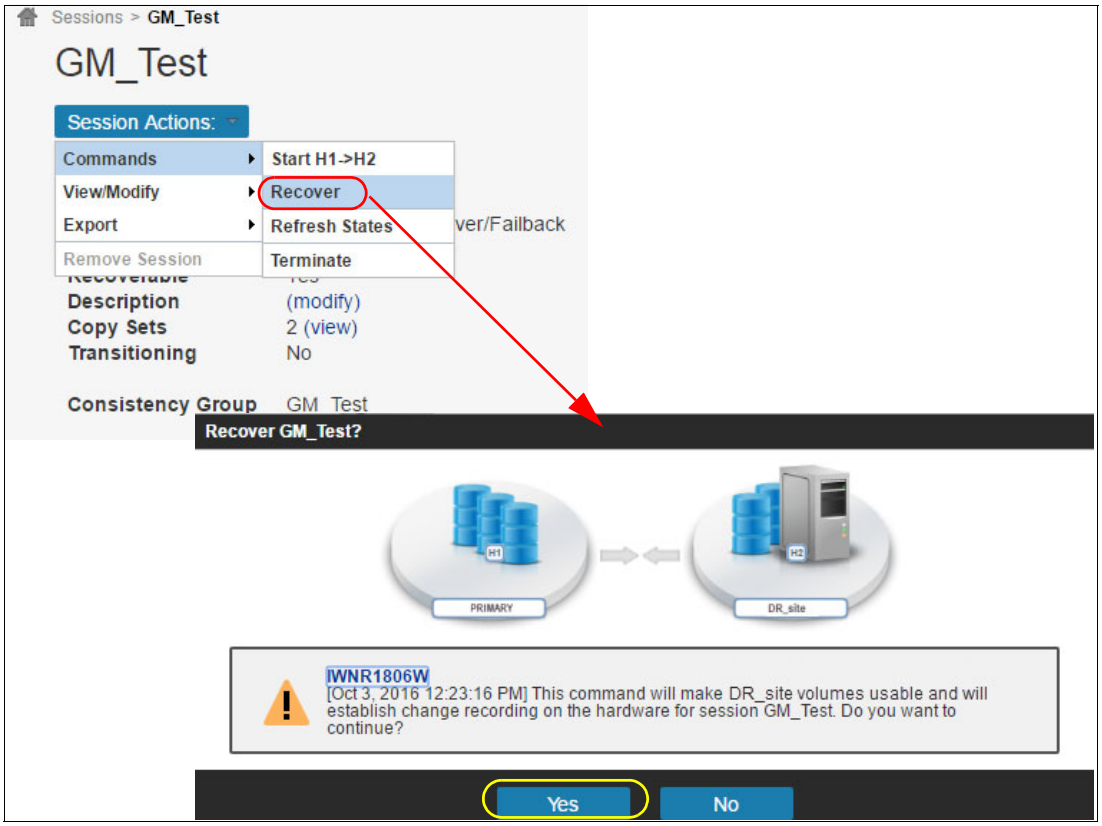

Figure 4-12 Making target (H2) available to hosts in DR_site

4. When the recovery command has been performed, we start I/O from hosts in the Secondary site. The H2 volumes are available for I/O, as shown in Figure 4-13.

Note: It is out of scope for this material to describe how to activate host systems and access volumes from these. Each operating system has platform-specific methods and tools or utilities for this purpose.

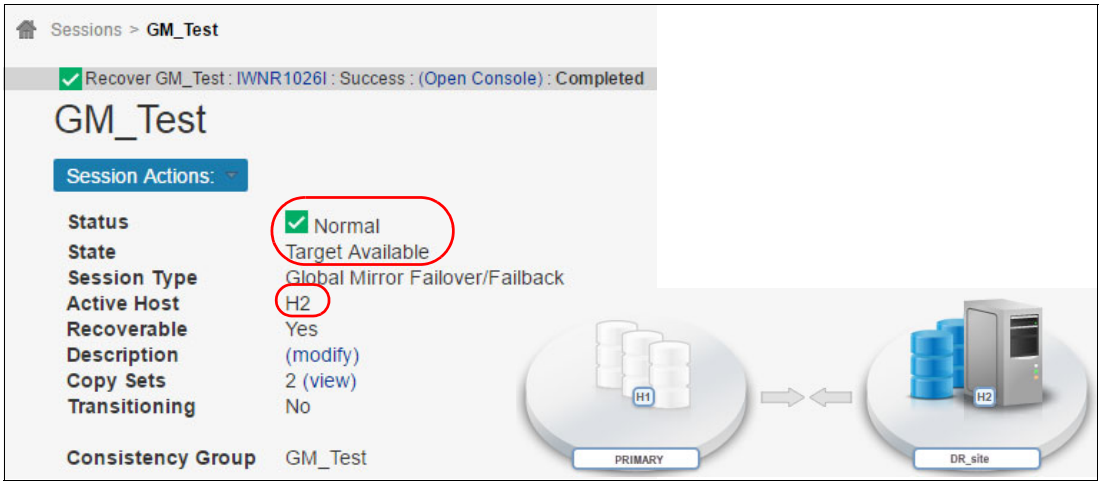

Figure 4-13 H2 volumes available for host access

At this time, I/O is performed on volumes in DR_site (H2). Depending on the test scenario, you can either failback and replicate the changes performed on H2 while host systems were active in DR_site, or you can resume operation in PRIMARY site, discarding changes in DR_site (overwrite changes).

The failback scenario we tested assumes that failback is performed with replicating changes (H2 -> H1) while I/O has been performed on volumes in DR_site.

4.2.2 Testing Global Mirror failback (controlled)

We perform the following steps:

1. To replicate changes while H2 volumes are active (I/O in DR_site), we enable copy to Site 1 (PRIMARY), as shown in Figure 4-14.

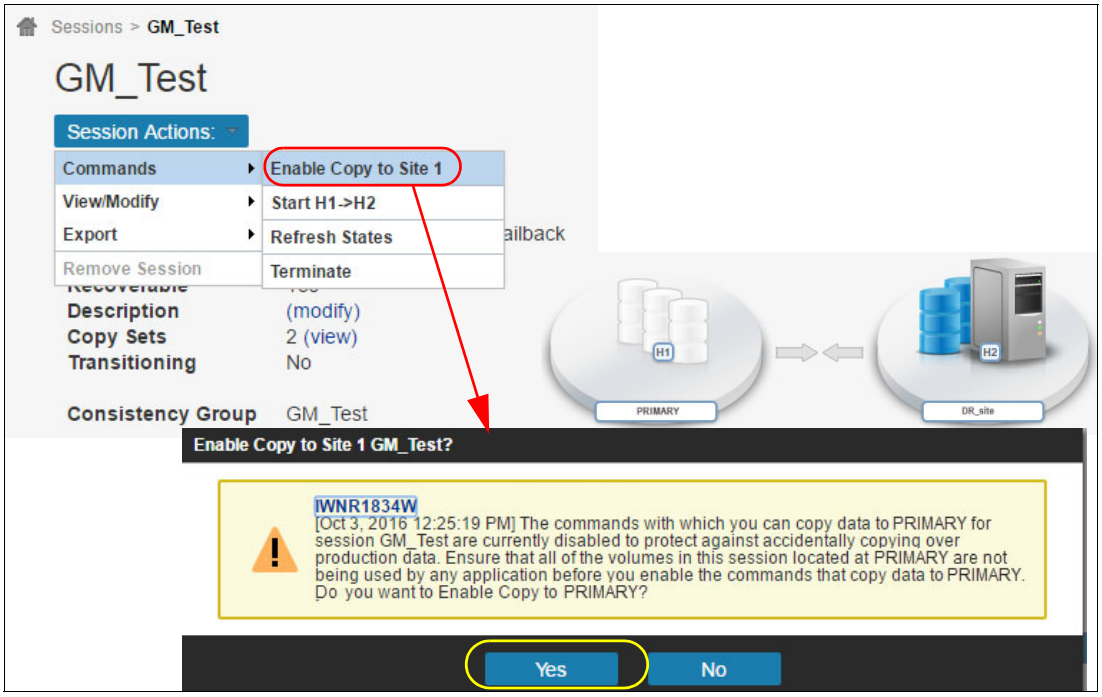

Figure 4-14 Enabling copy to PRIMARY site

2. At this time, we start replication to the PRIMARY site, as show in Figure 4-15.

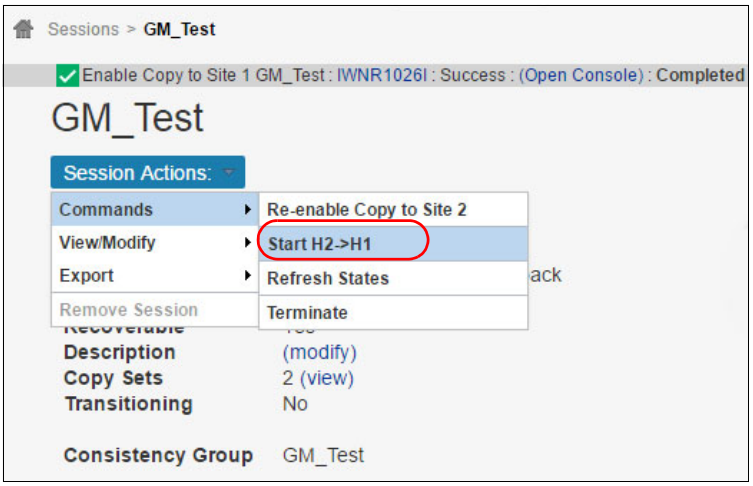

Figure 4-15 Starting copy to PRIMARY site

3. Note that the volumes in PRIMARY site will be overwritten. We confirm, as shown in Figure 4-16.

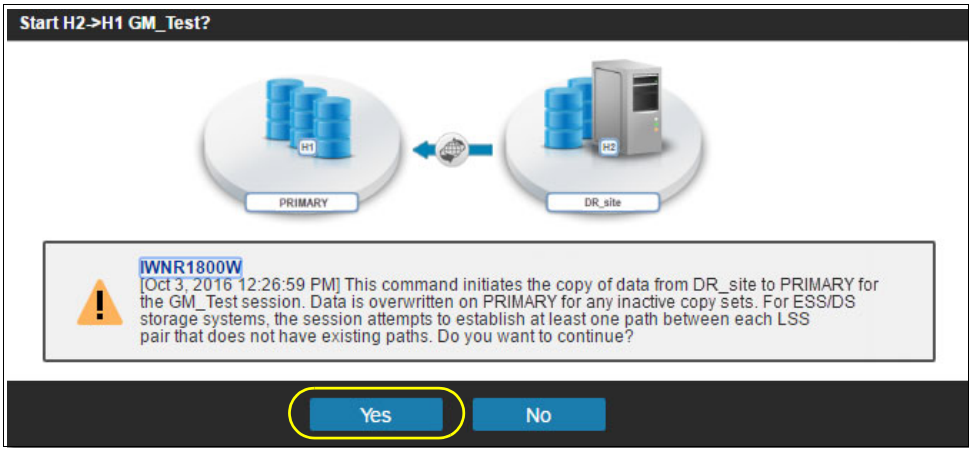

Figure 4-16 Confirm starting H2 -> H1

4. When the session has reached Normal / Prepared state (see Figure 4-17), we proceed to moving host systems' access from DR_site to PRIMARY site.

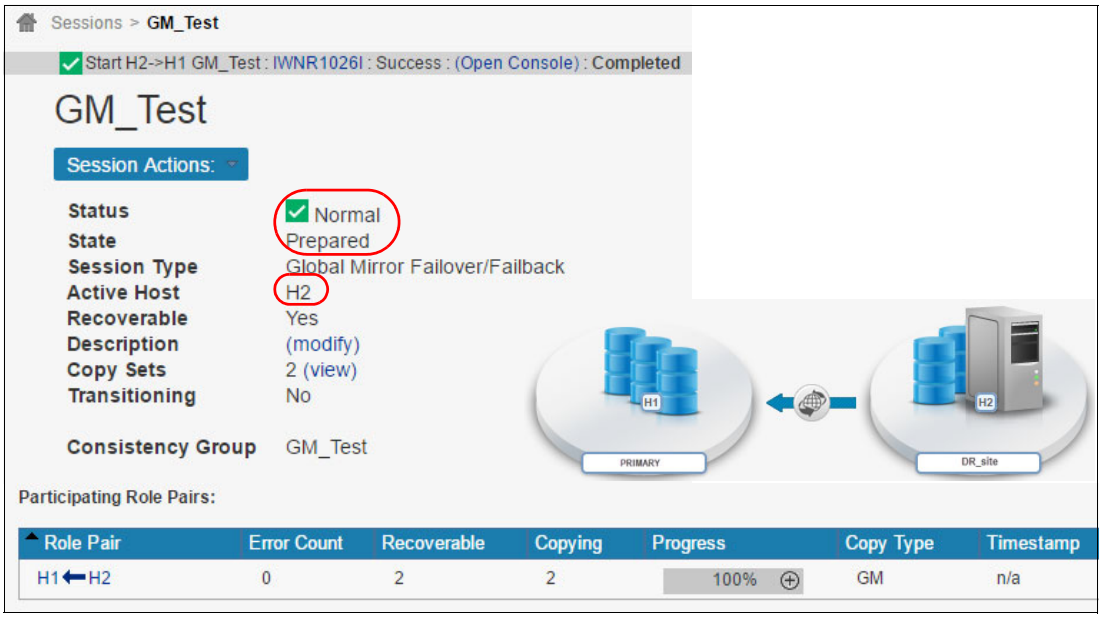

Figure 4-17 H1 -> H2 in consistent copying state

5. Because this is a controlled failback, we stop host systems I/O on volumes in DR_site before suspending the session (see Figure 4-18).

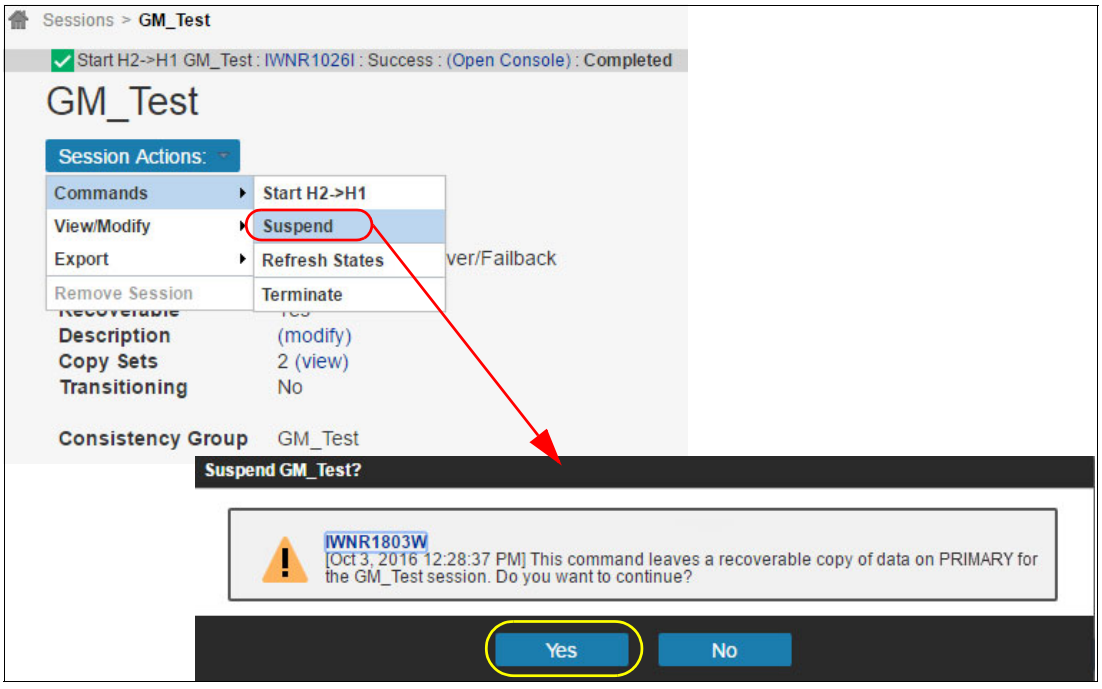

Figure 4-18 Suspending H2 -> H1

6. Before starting I/O for host systems in PRIMARY site, we perform session recovery, as shown in Figure 4-19.

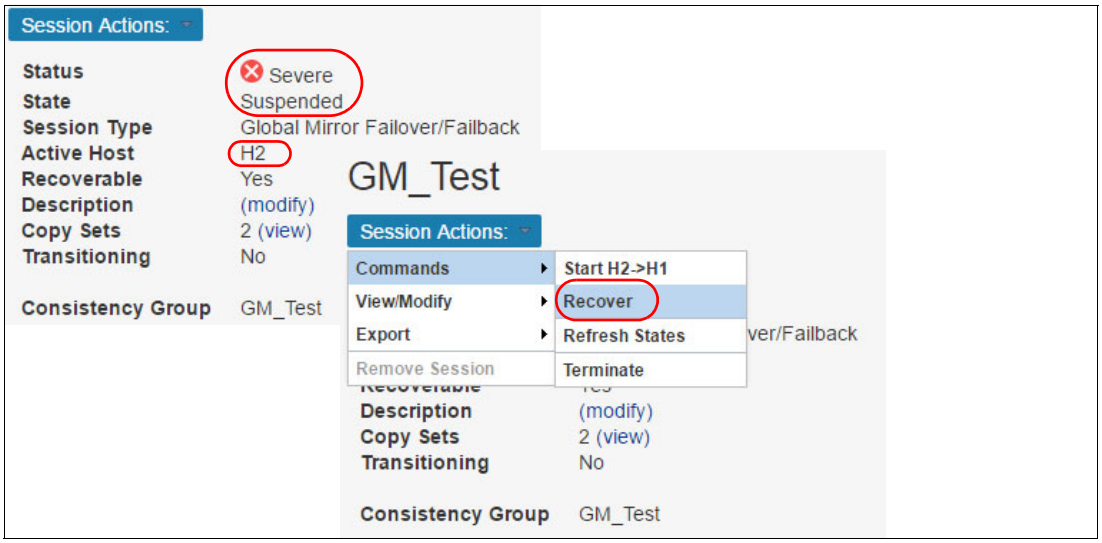

Figure 4-19 Session Recovery

7. We confirm (see Figure 4-20). At this time, we can start I/O on host systems in PRIMARY site (H1).

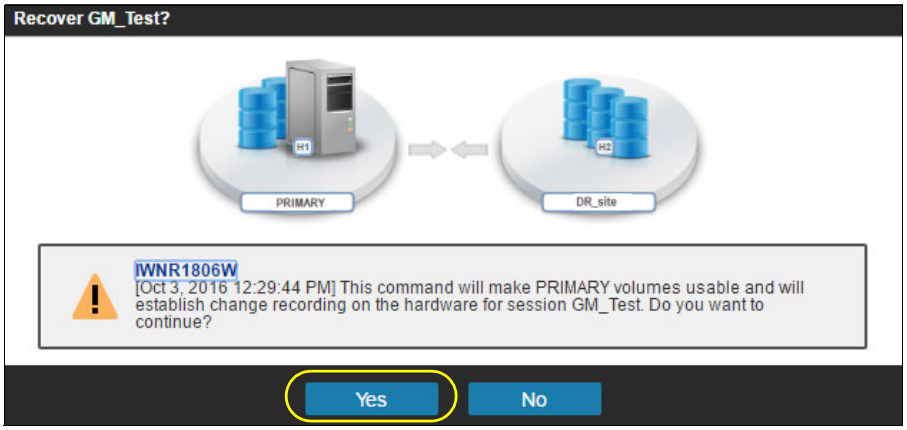

Figure 4-20 Confirm recovery action

8. To prepare for further activities, we enable copy from DR_site to PRIMARY, as shown in Figure 4-21.

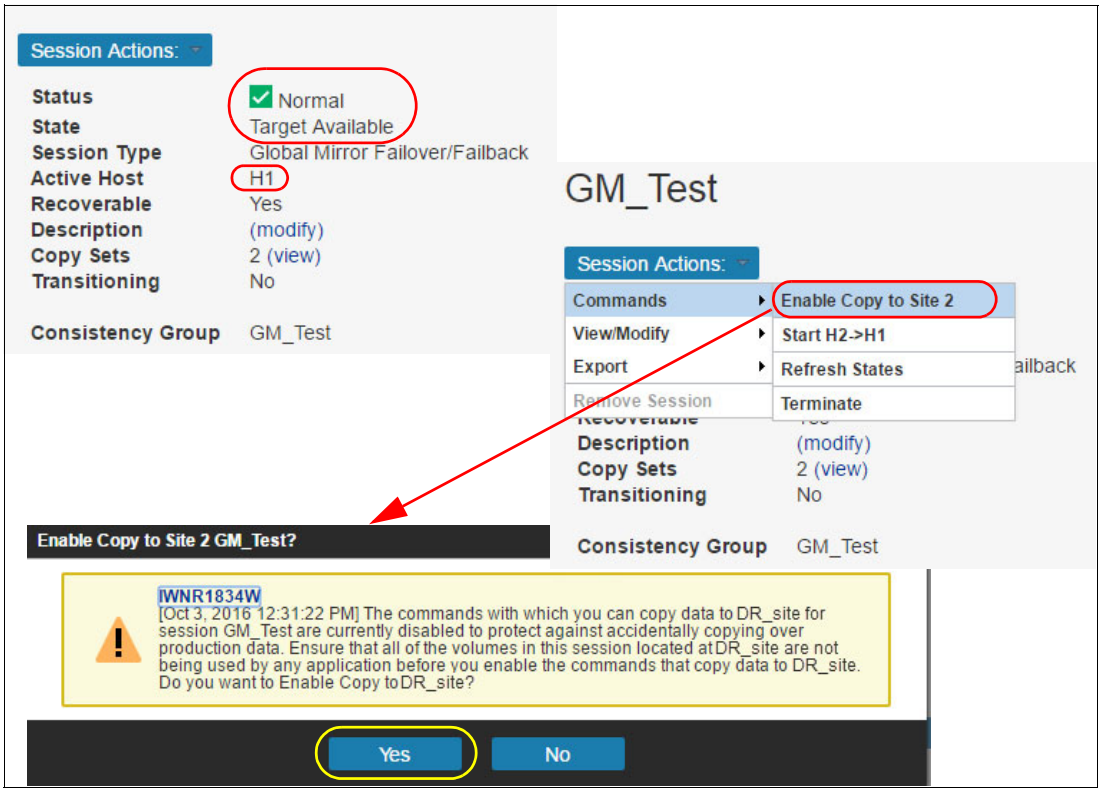

Figure 4-21 Re-enabling H1 -> H2 copy

9. We start H1 -> H2 session, as shown in Figure 4-22.

Figure 4-22 Starting H1 -> H2

10. Confirm, as shown in Figure 4-23.

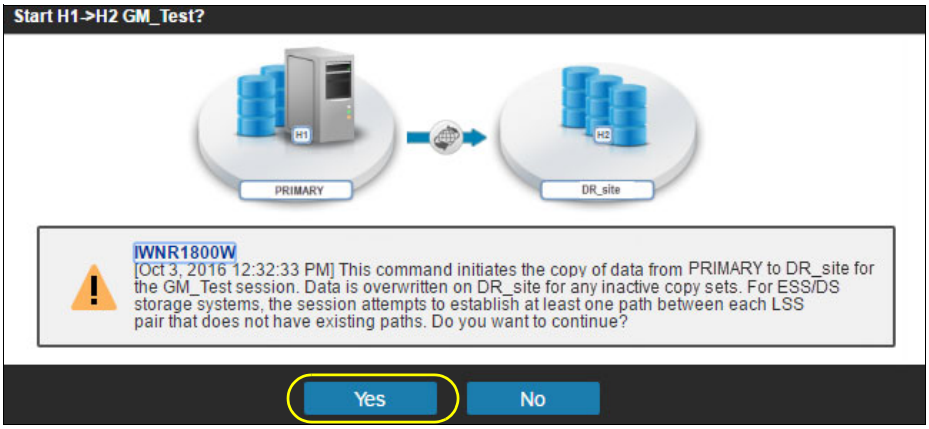

Figure 4-23 Confirm starting H1 -> H2

11. When the session reaches Normal / Prepared state (Figure 4-24), the configuration is back in the initial state and is prepared for disaster recovery.

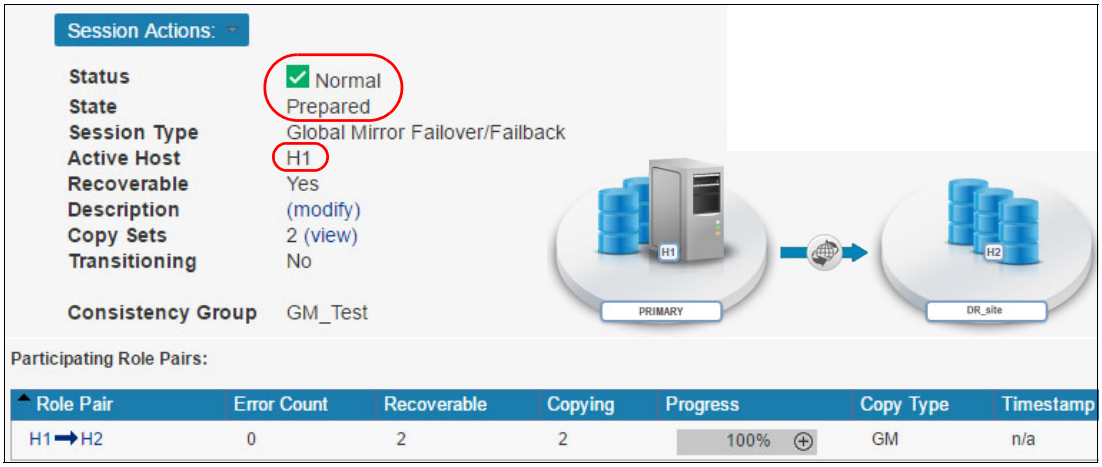

Figure 4-24 Session recovered to initial state

4.3 Converting a Global Mirror session to a Global Mirror with Change Volumes

The Global Mirror with Change Volumes (GM w/ CVs) has been designed to overcome and mask communication issues between sites. Possible communications issues might the following influences:

- Temporary replication line fluctuations caused by network infrastructure.
- Sudden increases in workload (spikes) which might temporarily increase the replication volume, and for which the replication line bandwidth has not been designed for.

Tip: For IBM SAN Volume Controller sessions, IBM Copy Services Manager server can be configured to automatically restart the session that has been suspended due to reason code 1720 or 1920.[a] SAN Volume Controller automatically suspends replication for Global Mirror relationships if it cannot keep up with forming consistency groups to avoid application impact. This CSM feature adds automation for the session restart.

Global Mirror with Change Volumes (GM w/ CVs) can be implemented to also mask line fluctuations. In addition, this option can provide protection for workload spikes.

The drawback of using GM w/ CVs is that the delta between the primary and secondary copies can be significant (increased Recovery Point Objective - RPO) if line throughput is not properly sized to accommodate the spikes in a timely manner.

a. The IBM CSM feature is named “Basic automatic restart for SAN Volume Controller Global Mirror suspend operations with reason code 1720 or 1920”.

Scenario overview

In this scenario, we use an existing (active) Global Mirror (GM) relationship and convert it to Global Mirror with Change Volumes (GM w/ CVs).

Tip: It is important to perform the conversion steps in the right sequence to avoid any downtime and maximize data availability.

Note: During this change, the copy sets managed by the initial GM session continue to be replicated by the underlying HW (SVC), and are in “Consistent Synchronized” state.

An overview of the steps we have followed to change an existing GM relationship to GM w/ CVs is presented in the following list:

1. Based on the existing GM session (represented as a consistency group in the IBM SVC), we create in each site the volumes (to be used as change volumes) for each H1 and H2 volume in the GM session. This step is done on the SVC directly (GUI or CLI), no CSM involvement.
2. In CSM, we define a new session of type `Global Mirror with Change Volumes` and “populate” it with the designated volumes (H1, CV1, H2, and CV2).
3. We remove the GM relationship from Copy Services Manager (H1 -> H2) *without* removing the base hardware relationship on the storage system. This leaves the copy sets in the base hardware (storage) untouched (in `Consistent` state).
4. In the IBM CSM GUI, we start the GM w/ CVs session and monitor it.

4.3.1 GM w/ CVs session preparation

Note: IBM Copy Services Manager does *not* provision (create, delete, or modify) volumes and pools or other storage entities. IBM CSM only manages (copy services) sessions (represented in the HW as consistency groups).

The existing GM session (GM_Test) in the underlying HW (IBM SVC) is shown in Figure 4-25.

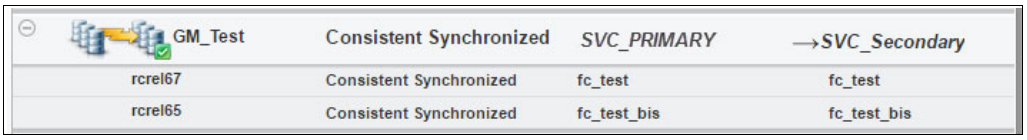

Figure 4-25 Session GM_Test represented in underlying HW

Hint: It is important to create the change volumes before removing the GM session from CSM because this action can be scripted based on the SVC Consistency Group (CG) information. This is especially useful for sessions that consist of many copy sets.

If you remove the session from the CSM (without removing the relationships from the underlying hardware), the Consistency Group information will not be available in SVC (CLI or GUI), making it difficult to identify the copy sets.

We complete the following steps:

1. We start in the secondary SVC by creating the change volumes (CVs) for all volumes in the GM_Test consistency group. Example 4-1 shows the script we have used.

   Hint: The script can be used for other sessions by adjusting the `CG` and `POOL` variables.

   Attention: SAN Volume Controller scripting and CLI skills are required to perform this action. Make sure that you understand the actions in the script before running it in your environment.

   Note: Creating the volumes to be used as change volumes can be done from SVC GUI or CLI. We have chosen the script method because it can be easily adapted and repurposed for other relationships.

   Example 4-1 Script to create CVs for Secondary site CVs for H2 volumes)

```
CG=GM_Test
pool=IT_test
lsrcrelationship -filtervalue consistency_group_name=$CG -nohdr |while read -a
rc
do
lsvdisk -nohdr -bytes -filtervalue id=${rc[8]} | while read -a v
do
echo Creating volume ${v[1]}"_cv" with virtual size ${v[7]} bytes
mkvdisk -autoexpand -iogrp ${v[3]} -mdiskgrp $pool -name ${v[1]}"_cv" -rsize 0
-size ${v[7]} -unit b
done
done
```

   This creates in Secondary site the following volumes and associates them as change volumes with auxiliary volumes:

   fc_test_cv (with virtual size 10737418240 bytes)
   fc_test_bis_cv (with virtual size 10737418240 bytes)

2. We run a similar script in the PRIMARY site storage. Example 4-2 shows the script we ran in the PRIMARY site for creation of the change volumes.

Example 4-2 Script to create CVs for PRIMARY site (CVs for H1 volumes)

```
CG=GM_Test
pool=IT_test
lsrcrelationship -filtervalue consistency_group_name=$CG -nohdr |while read -a
rc
do
lsvdisk -nohdr -bytes -filtervalue id=${rc[4]} | while read -a v
do
echo Creating volume ${v[1]}"_cv" with virtual size ${v[7]} bytes
mkvdisk -autoexpand -iogrp ${v[3]} -mdiskgrp $pool -name ${v[1]}"_cv" -rsize 0
-size ${v[7]} -unit b
done
done
```

This creates in the PRIMARY site the following volumes that will be associates as change volumes to the master volumes:

fc_test_cv (with virtual size 10737418240 bytes)
fc_test_bis_cv (with virtual size 10737418240 bytes)

3. In the CSM, we create another session of type `Global Mirror Failover/Failback w/ Change Volumes`, as shown in Figure 4-26.

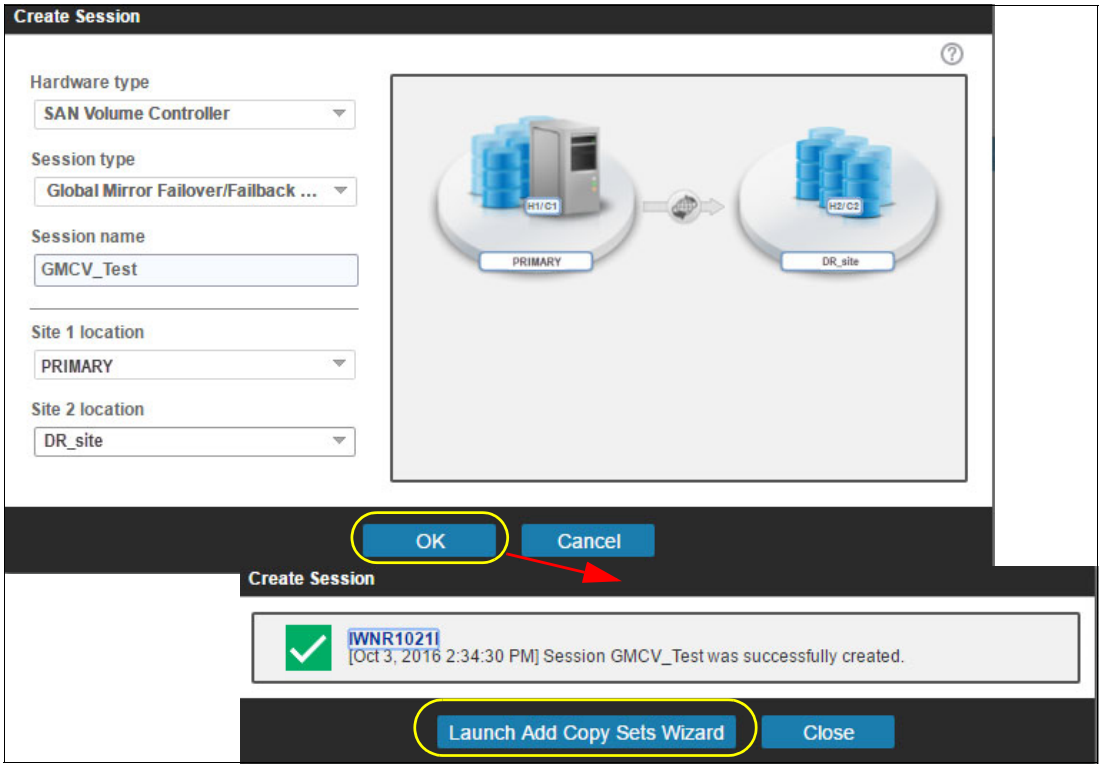

Figure 4-26 Create Session and Launch Add Copy Sets Wizard

4. We select volumes for H1, as shown in Figure 4-27.

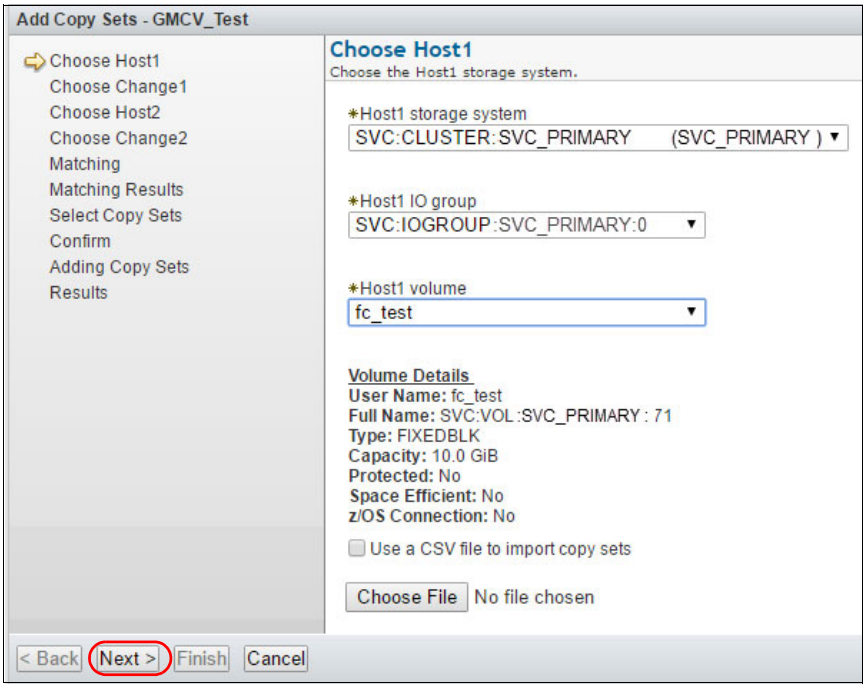

Figure 4-27 Selecting H1 volumes

5. Select Change1 (C1) volumes (see Figure 4-28).

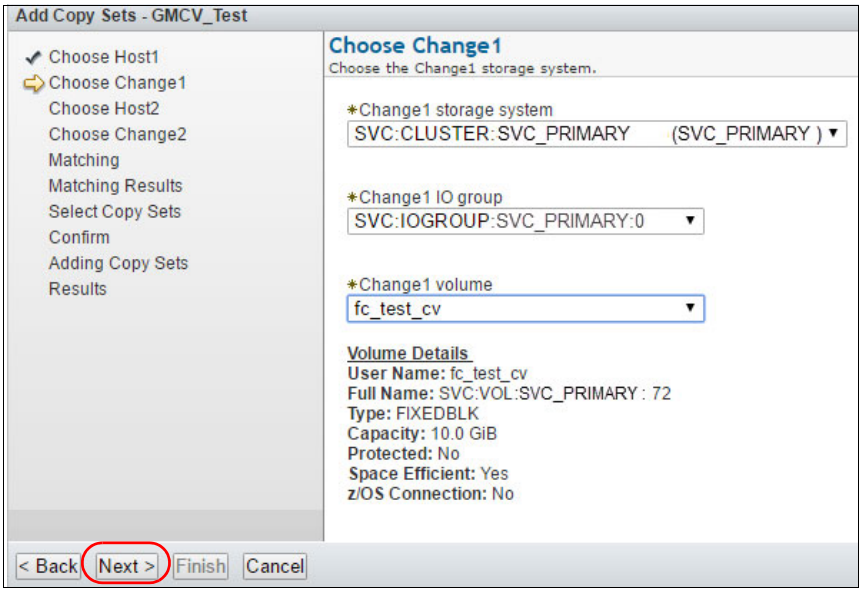

Figure 4-28 Change1 (C1) selection

6. Select Host2 (H2) volumes, as shown in Figure 4-29.

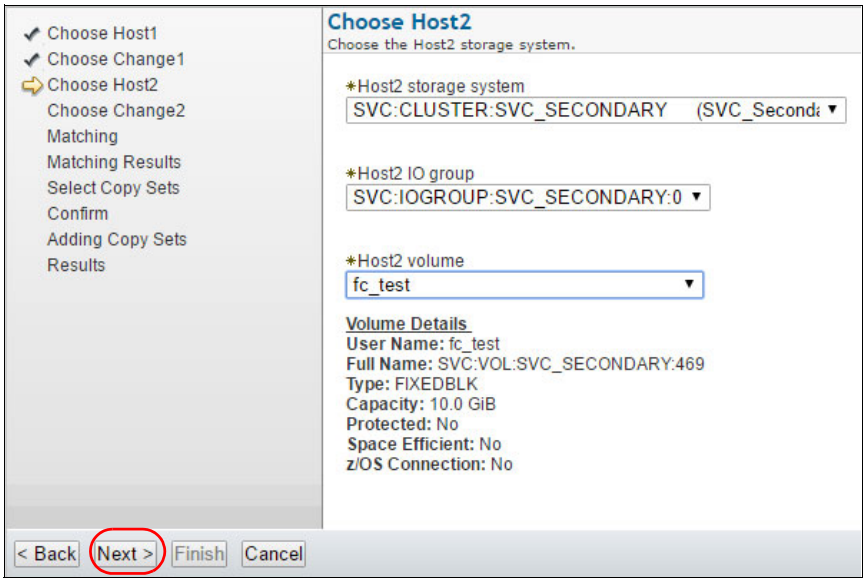

Figure 4-29 Select H2

7. We select Change2 (C2) volumes, as shown in Figure 4-30.

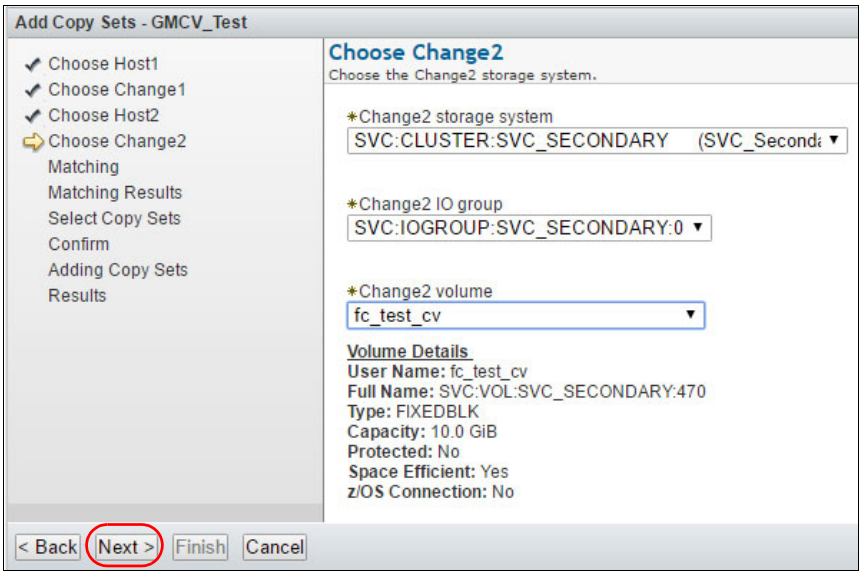

Figure 4-30 Selecting C2

8. After the wizard matches the results, we select the copy sets to create, as shown in Figure 4-31.

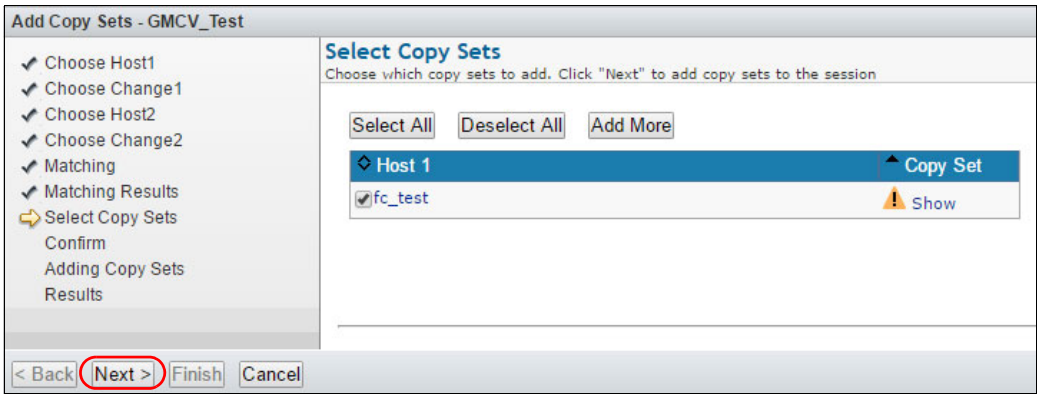

Figure 4-31 Selecting copy sets to be added to the session

9. We verify copy set information, as shown in Figure 4-32.

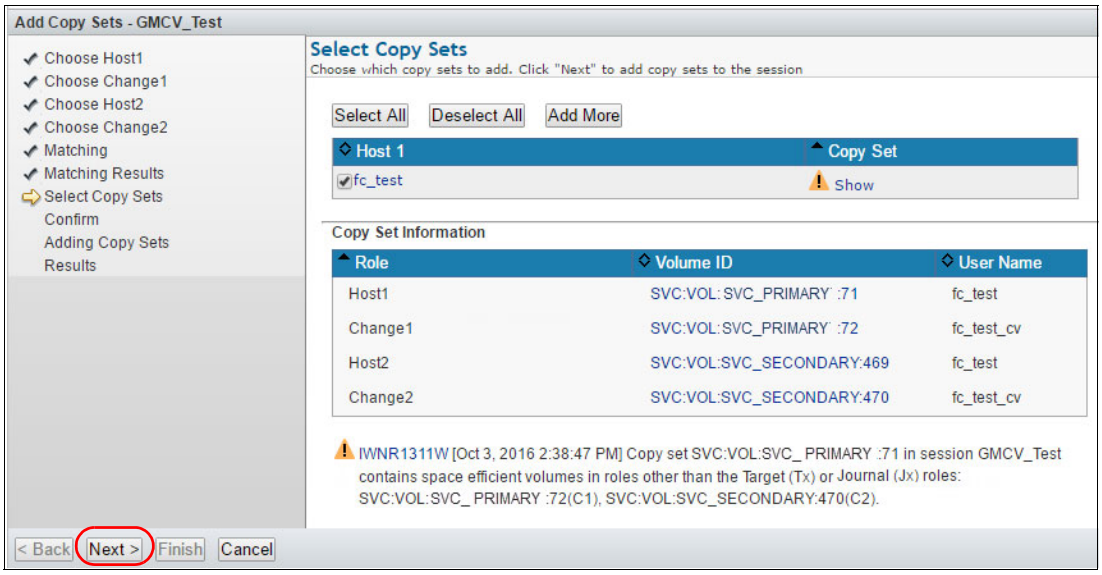

Figure 4-32 Copy set details

10. We confirm and finalize adding the copy sets, as shown in Figure 4-33.

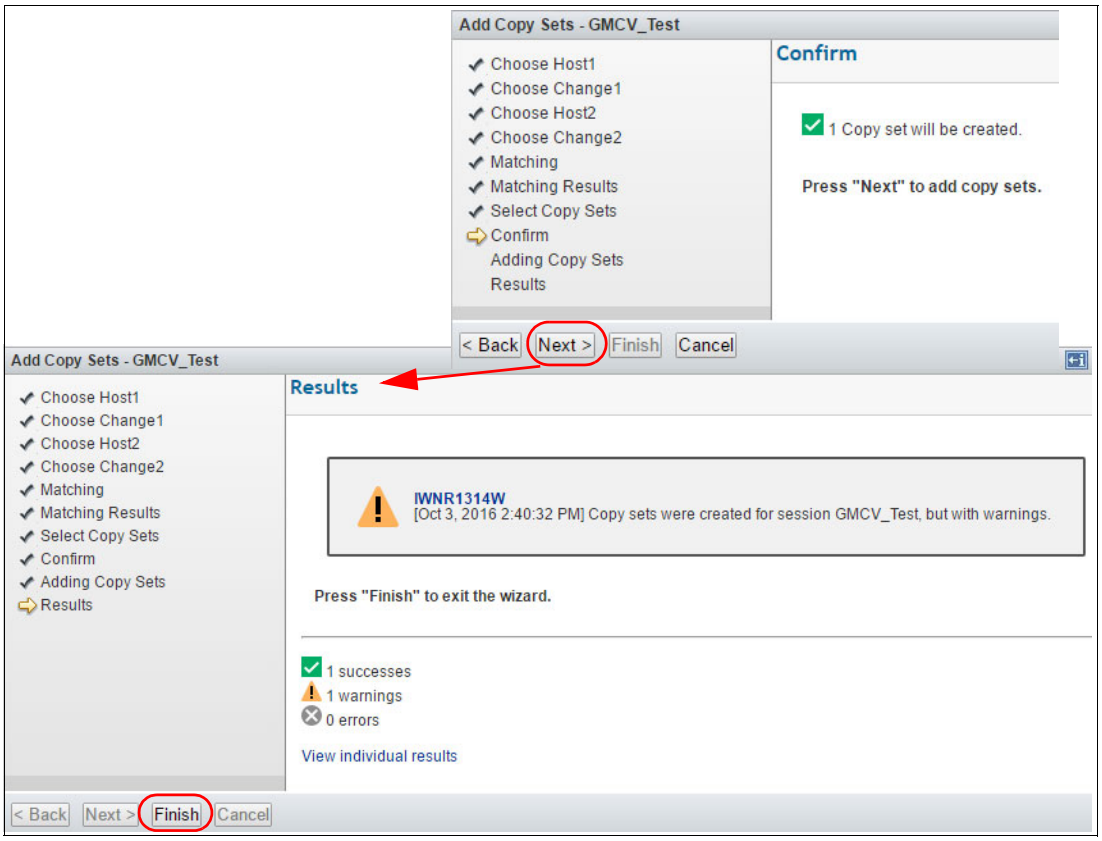

Figure 4-33 Finalizing the Add Copy Sets

11. We have added the second set of volumes (copy set) using a `.csv` file shown in Example 4-3.

Example 4-3 The .csv file used for adding the second copy set

```
#GMCV_Test                          ,,,
#Global Mirror Failover/Failback w/ Change Volumes                    ,,,
#Oct 3   2016 2:42:54 PM                ,,,
                          ,,,
H1,C1,H2,C2
SVC:VOL:SVC_PRIMARY:fc_test_bis,SVC:VOL:SVC_PRIMARY:fc_test_bis_cv,SVC:VOL:SVC_
SECONDARY:fc_test_bis,SVC:VOL:SVC_SECONDARY:fc_test_bis_cv
```

12. The NEW session (GM w/ CVs) is not started at this time. We check the copy sets, as shown in Figure 4-34.

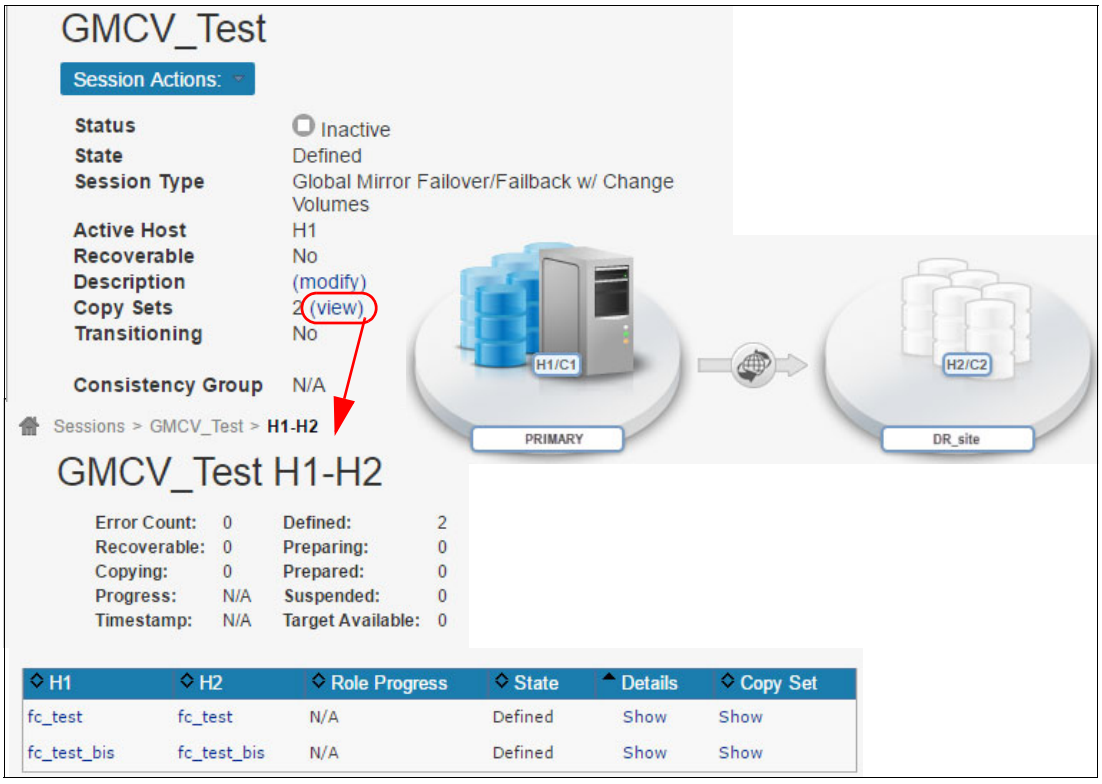

Figure 4-34 Checking the copy sets

13. We remove the copy sets from the existing GM_Test session as shown in Figure 4-35.

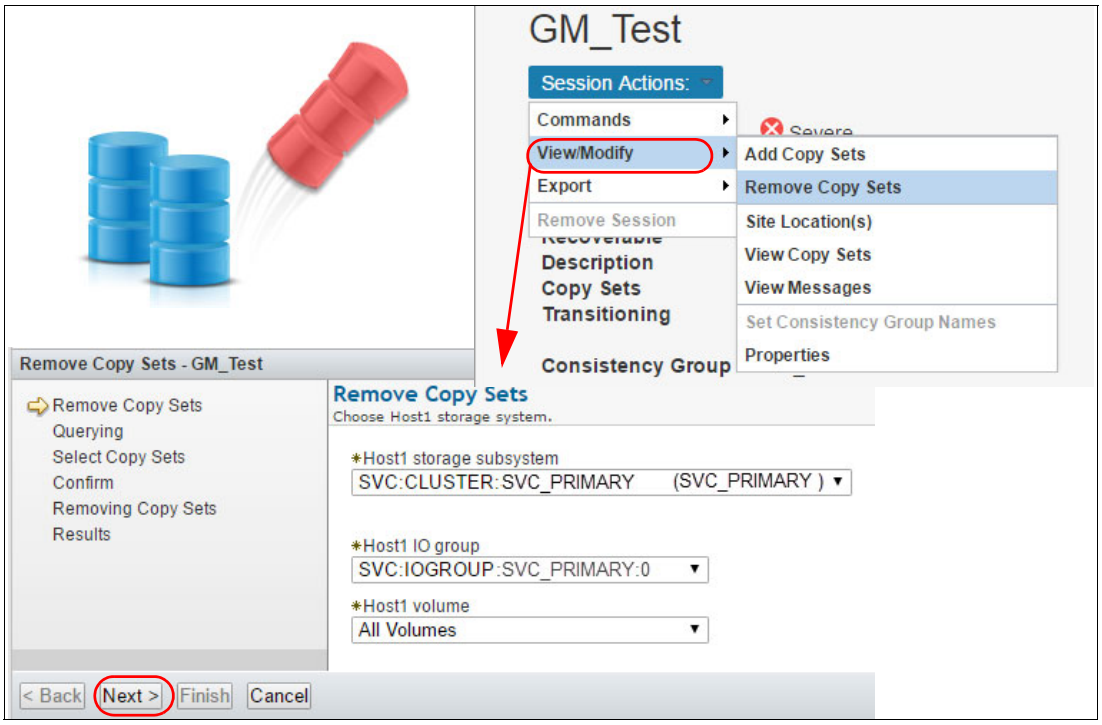

Figure 4-35 Initiating the Remove Copy Set operation

14. We make sure that the underlying HW relationship is preserved, as shown in Figure 4-36.

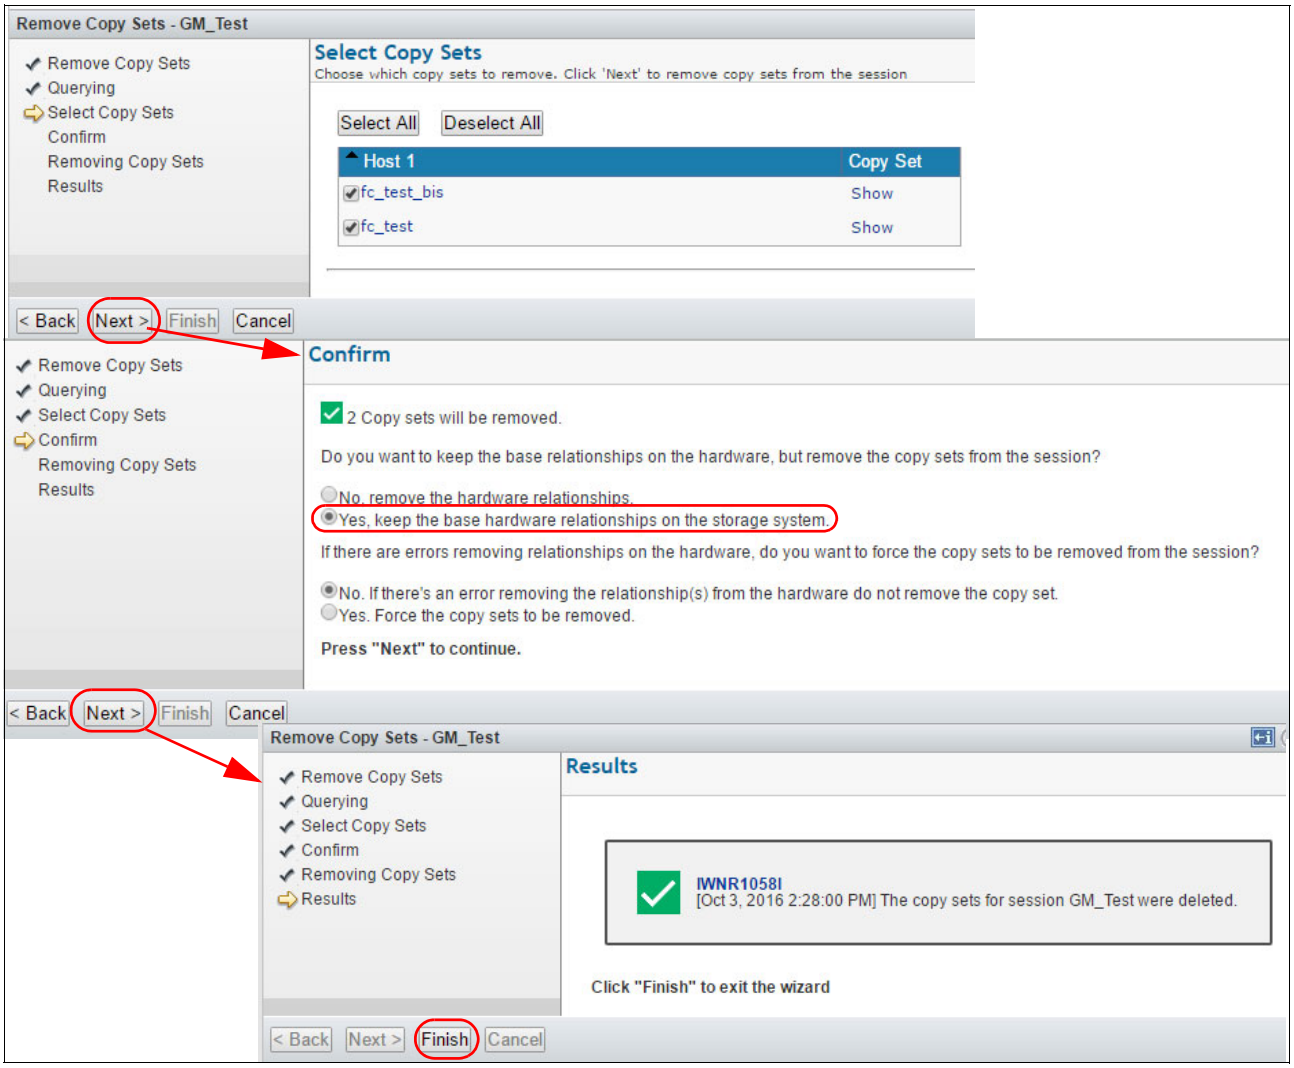

Figure 4-36 Finalizing copy set removal

15. The consistency group in the underlying hardware is removed, as shown in Figure 4-37, but the relationship between the volumes in PRIMARY and DR_site storage is preserved.

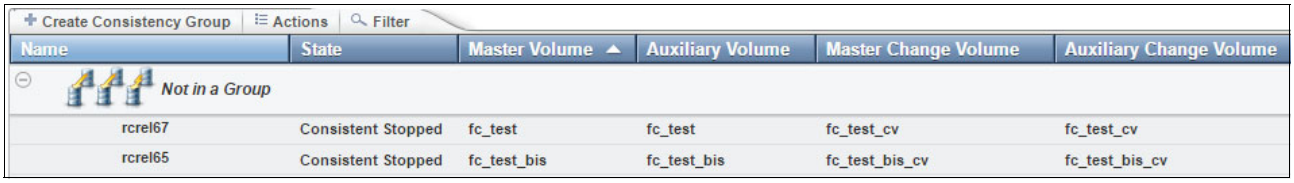

Figure 4-37 SVC relationships are preserved

16. We start the newly created session GMCV_Test (Global Mirror Failover/Failback w/ Change Volumes), as shown in Figure 4-38.

Figure 4-38 Starting the GM w/ CV session

17. We confirm action, as shown in Figure 4-39.

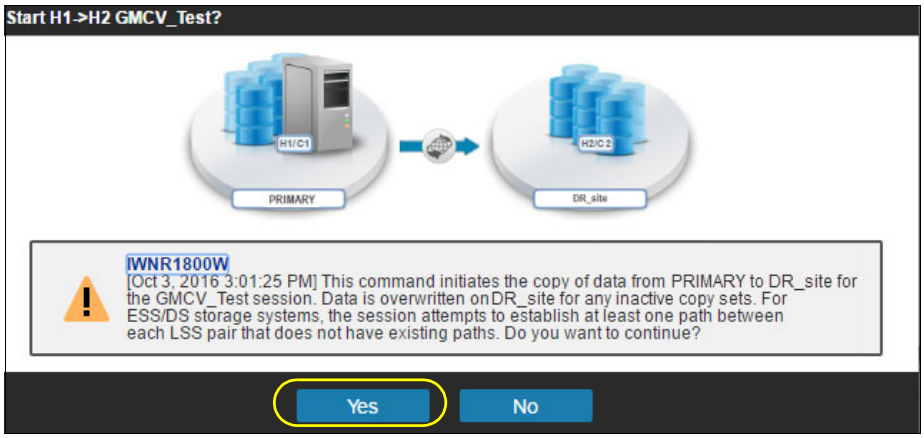

Figure 4-39 Confirm session start

18. We check the progress and monitor session, as shown in Figure 4-40.

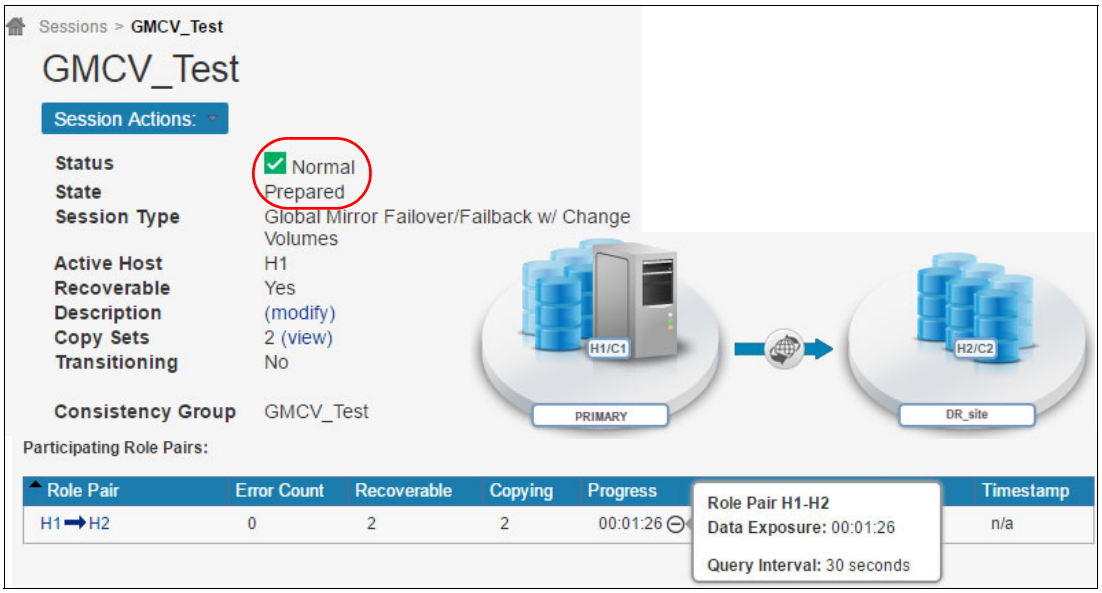

Figure 4-40 Session in Normal/Prepared state

19. Figure 4-41 shows how the session is displayed in the underlying hardware (SVC).

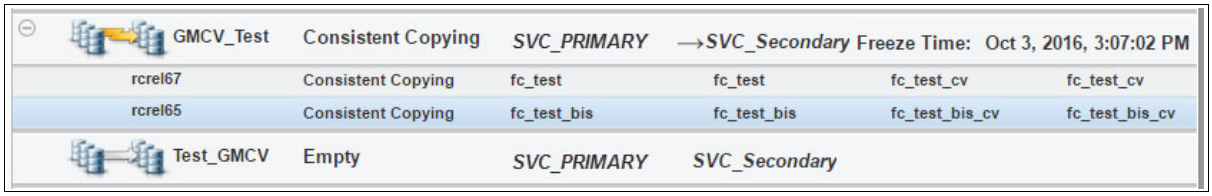

Figure 4-41 Session view from SVC

This is the end of the scenario.

4.3.2 Hint: exporting session copy sets data to a file

To keep a record of the session, you can export session copy set data to a comma-separated values file (`.csv`). This file can be further used as a base for adding multiple copy sets to the existing session, or to other sessions. For sessions with many volumes, it is practical to supply copy set data to CSM in a file.

Complete the following steps:

1. To export session (copy sets) data to a `.csv` file select **Session Actions** → **Export Copy Set**, as shown in Figure 4-42.

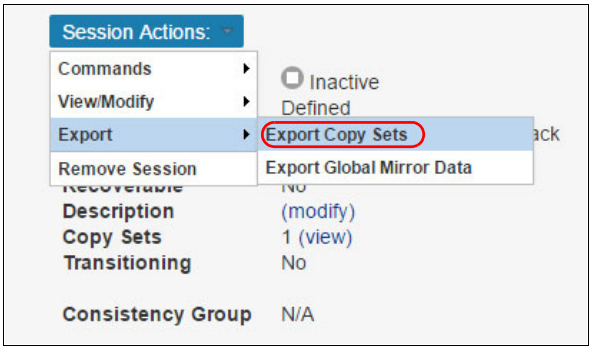

Figure 4-42 Exporting Copy Sets data

2. The copy sets are exported to a file generated by CSM, as shown in Figure 4-43.

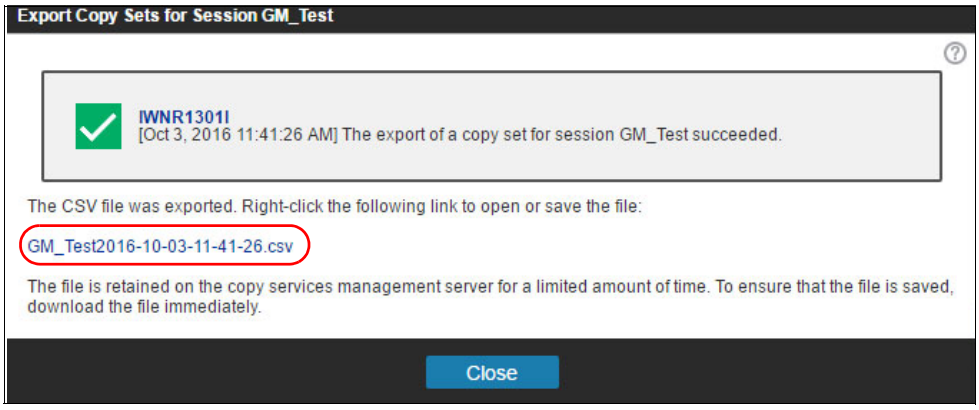

Figure 4-43 Export Copy Sets result

The content of the file is shown in Example 4-4.

Example 4-4 Export data for session GM_Test

```
#GM_Test,
#Global Mirror Failover/Failback,
#Oct 3, 2016 11:41:26 AM
,
H1,H2
SVC:VOL:SVC_PRIMARY:fc_test,SVC:VOL:SVC_SECONDARY:fc_test
SVC:VOL:SVC_PRIMARY:fc_test_bis,SVC:VOL:SVC_SECONDARY:fc_test_bis
```

5

Managing IBM DS8000 replication

This chapter describes replication scenarios implemented with IBM DS8000 series storage subsystems and managed with IBM Copy Services Manager Version 6 Release.

Replication using IBM DS8000 series provides multi-target capabilities. The scenarios presented cover the creation of such configurations and how to manage the environments using IBM CSM for achieving optimal resiliency and ease of use.

This chapter describes the following scenarios:

- Migrating from Global Mirror replication to Multi-Target Metro Mirror - Global Mirror
  - Creating the GM session
  - Configuring a Multi-Target MM-GM session
- Migrating from a Basic z/OS HyperSwap session to Multi-Target Metro Mirror - Metro Mirror
  - Creating the basic HyperSwap session
  - Testing the HyperSwap functionality (z/OS)
  - Migrating to a Multi-Target MM-MM session

5.1 Migrating from Global Mirror replication to Multi-Target Metro Mirror - Global Mirror

Business case

This scenario describes the procedure we have used to improve the resiliency of a DR solution based on a two-site Global Mirror (asynchronous) replication to an improved, three-site configuration without having to reconfigure the GM (initial) replication.

Although a GM-based solution provides an RPO > 0, due to the business requirements, high availability and a third site must be integrated in the current DR solution. The target of this deployment is a Multi-Target Metro Mirror - Global Mirror configuration.

Scenario overview

This scenario consists of the following steps:

1. Creating the GM session (an existing GM session can be used as well).
2. Preparing the migration: Creating the new Multi-Target MM-GM session.
3. Managing the replication under the new session (three site, Multi-Target MM-GM)

5.1.1 Creating the GM session

In a GM session, the source volumes (H1) are copied asynchronously to the target volumes (H2). The writes made by the host in site 1 are not chronologically replicated in the target (site 2, “NJ”), therefore journal volumes (J2) are used to provide data consistency in the DR site (site 2).

Complete the following steps:

1. To create the new GM session, click the **Create Session** button in the CSM GUI. We define the session name (`GM_initial`) and the two sites, as shown in Figure 5-1. In our case, the production site is **Manhattan** and the DR site is **NJ**.

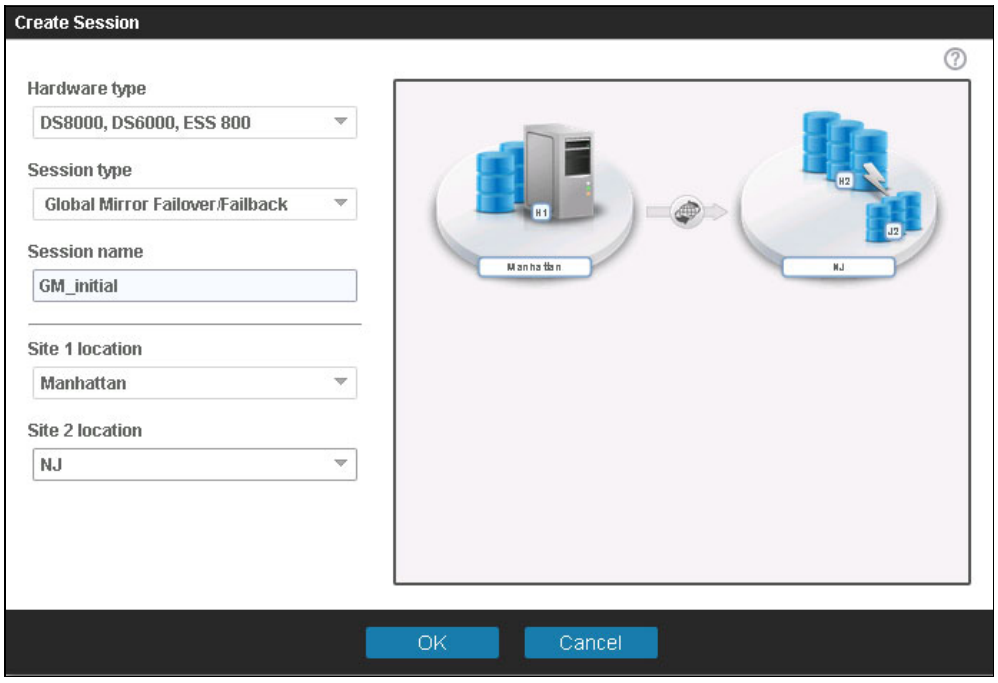

Figure 5-1 Creating the GM session

2. In the following step, we add copy sets to the session. Click the **Launch Add Copy Sets Wizard** button, as shown in Figure 5-2.

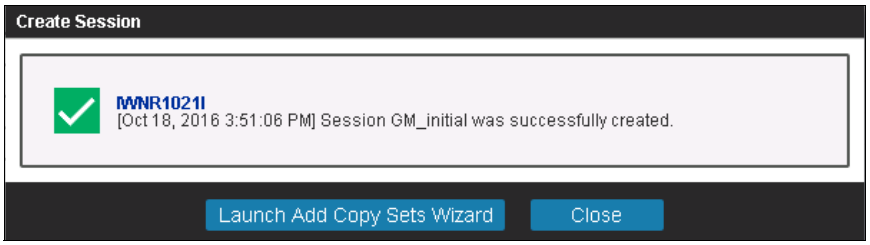

Figure 5-2 Launch Add Copy Sets Wizard

3. We enter, one by one, the volumes for H1 (Figure 5-3), H2 (Figure 5-4), and J2 (Figure 5-5 on page 88).

Figure 5-3 Choose H1 volumes

In our case, we replicate all volumes from LSS 05 in Manhattan to volumes from LSS 05 in NJ using journal volumes from LSS 06 in NJ.

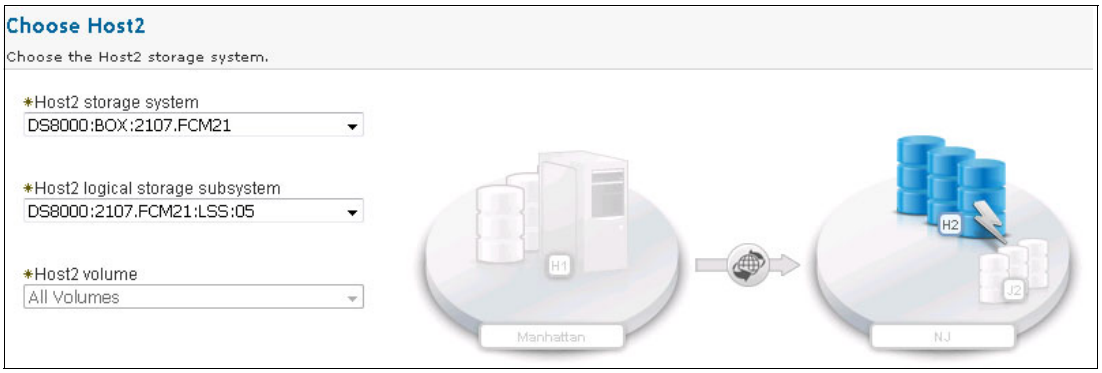

Figure 5-4 Choose H2 volumes

Optionally you can use a `.csv`[1] format file to import all copy sets. For details about how to import copy sets using a `.csv` file, see the CSM documentation.

[1] The data format in the .csv file depends on the session type and storage stem. You can check the format by exporting an existing session data.

Figure 5-5 shows entering the J2 volume.

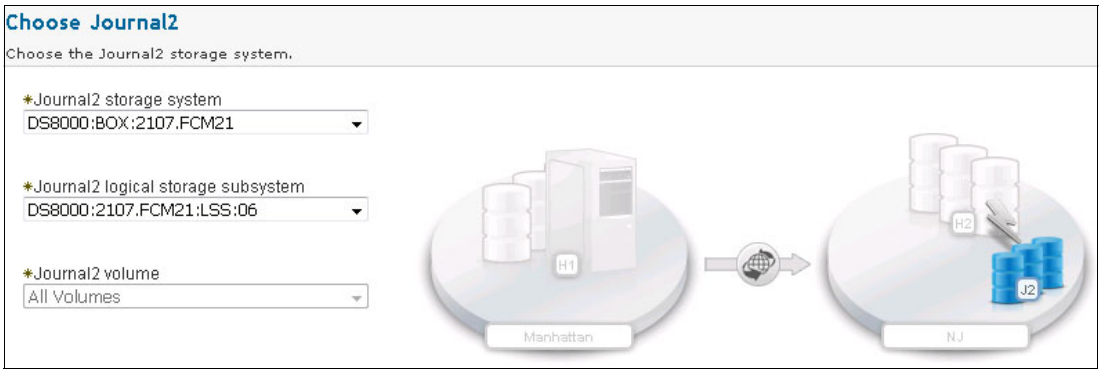

Figure 5-5 Choose J2 volumes

4. When the wizard has finished and all copy sets have been defined, the new session becomes Inactive / Defined, as shown in Figure 5-6. At this point in time, the CSM server has not yet sent any commands to the storage, so there is no active relationship set up on storage level.

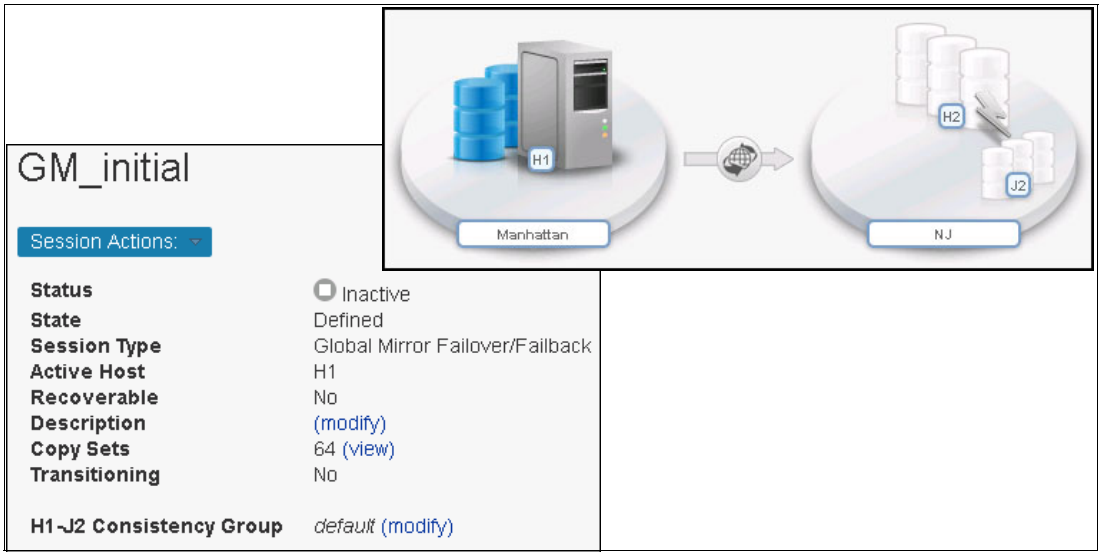

Figure 5-6 Session in "Inactive/Defined" state

5. To activate and use the session, we first start the (asynchronous) copy of data from H1 to H2. If the volumes in H2 were used before, all data will be overwritten by the Global Copy operation. Note that the Global Copy will not create a consistent copy of the data in the DR site (*NJ*).

Figure 5-7 shows the warning that explains what the **Start GC H1->H2** action will do.

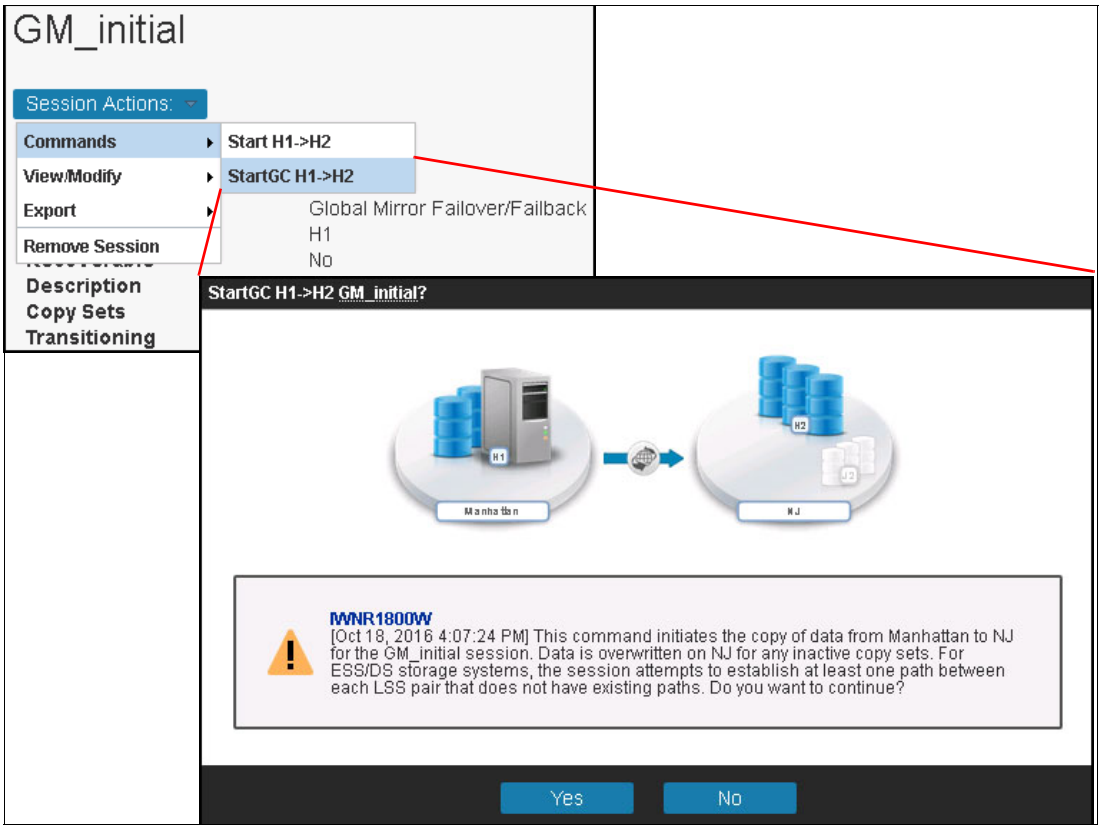

Figure 5-7 Starting GC H1->H2

6. The session enters the *Preparing* state while the data is copied to NJ, as shown in Figure 5-8. The *Warning* indicates that the volumes are being synchronized (see the session status definition in 1.5.2, “Monitoring sessions icons and symbols” on page 9).

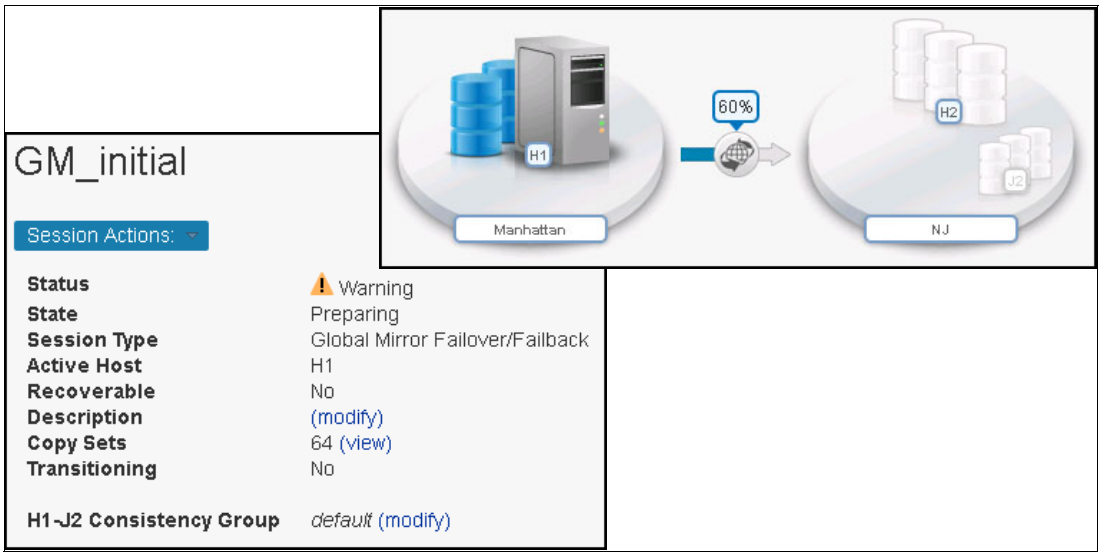

Figure 5-8 Session in Preparing state

7. While the operation is ongoing, you can check the newly created Global Copy relations using **dscli**, as shown in Example 5-1.

Example 5-1 The lssession command output (dscli)

```
dscli> lssession 05
Date/Time: 19 octombrie 2016 20:52:44 EEST IBM DSCLI Version: 7.7.21.39 DS:
IBM.2107-13FCL91
LSS ID Session Status  Volume  VolumeStatus  PrimaryStatus  SecondaryStatus
FirstPassComplete
================================================================================
05  03 CG In Progress  0500  Active  PrimaryCopy Pending   Secondary Simplex  True
05  03 CG In Progress  0501  Active  PrimaryCopy Pending   Secondary Simplex  True
05  03 CG In Progress  0502  Active  PrimaryCopy Pending   Secondary Simplex  True
05  03 CG In Progress  0503  Active  PrimaryCopy Pending   Secondary Simplex  True
05  03 CG In Progress  0504  Active  PrimaryCopy Pending   Secondary Simplex  True
05  03 CG In Progress  0505  Active  PrimaryCopy Pending   Secondary Simplex  True
05  03 CG In Progress  0506  Active  PrimaryCopy Pending   Secondary Simplex  True
05  03 CG In Progress  0507  Active  PrimaryCopy Pending   Secondary Simplex  True
05  03 CG In Progress  0508  Active  PrimaryCopy Pending   Secondary Simplex  True
05  03 CG In Progress  0509  Active  PrimaryCopy Pending   Secondary Simplex  True
..........
..........
05  03 CG In Progress  053D  Active  PrimaryCopy Pending   Secondary Simplex  True
05  03 CG In Progress  053E  Active  PrimaryCopy Pending   Secondary Simplex  True
05  03 CG In Progress  053F  Active  PrimaryCopy Pending   Secondary Simplex  True
dscli>
```

8. After the copy from primary to DR has completed and the out-of-sync tracks result is zero (or almost zero), CSM will show 100% progress for the session (see Figure 5-9). The session type is still Global Copy, but it is now ready for creating the consistency groups required for the Global Mirror session.

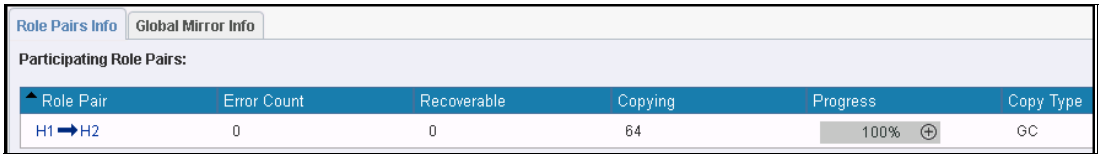

Figure 5-9 Global Copy session progress 100%

Note: You can directly start the GM session, but the behavior is the same: a Global Copy session starts and it waits for out-of-sync tracks to become zero (or near zero). Only then is a consistency group created together with the GM relation.

9. Next, we create the GM consistency groups. From the **Session Actions** menu, select **Start H1->H2**, as shown in Figure 5-10. Note the *Warning* message.

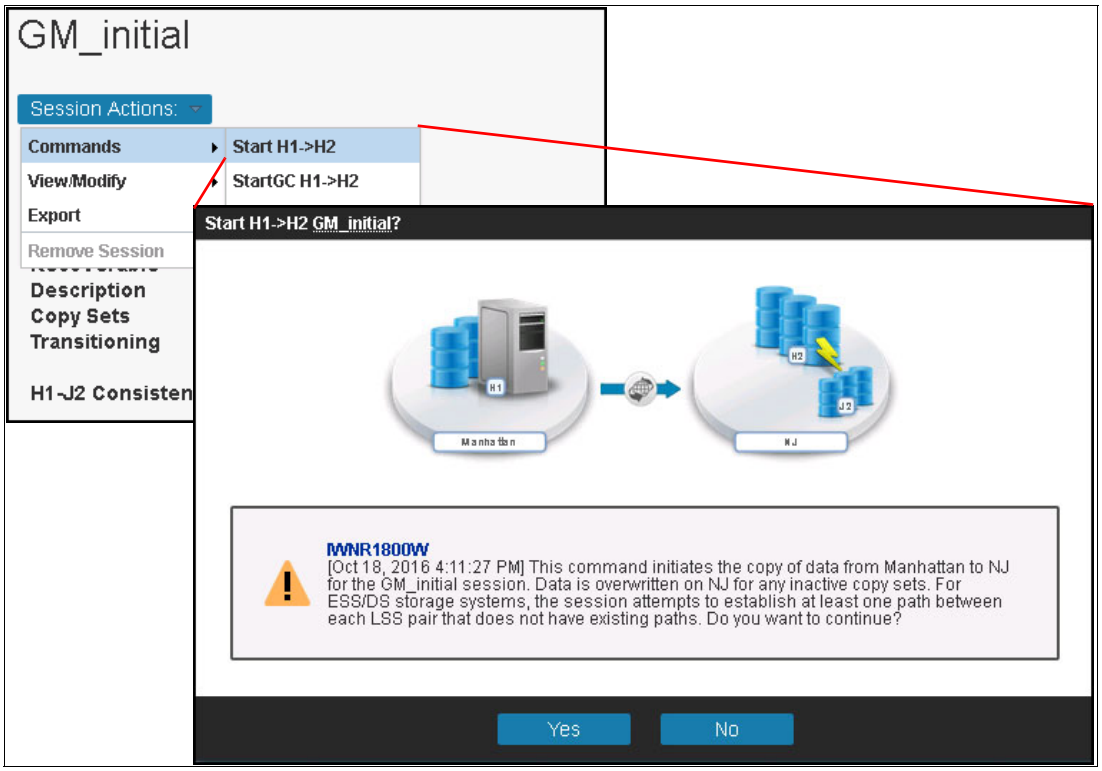

Figure 5-10 Start H1->H2 (Global Mirror)

10. When the first consistency group is established, the session state transitions to *Normal/Prepared* and J2 has a consistent copy of the data, as shown in Figure 5-11.

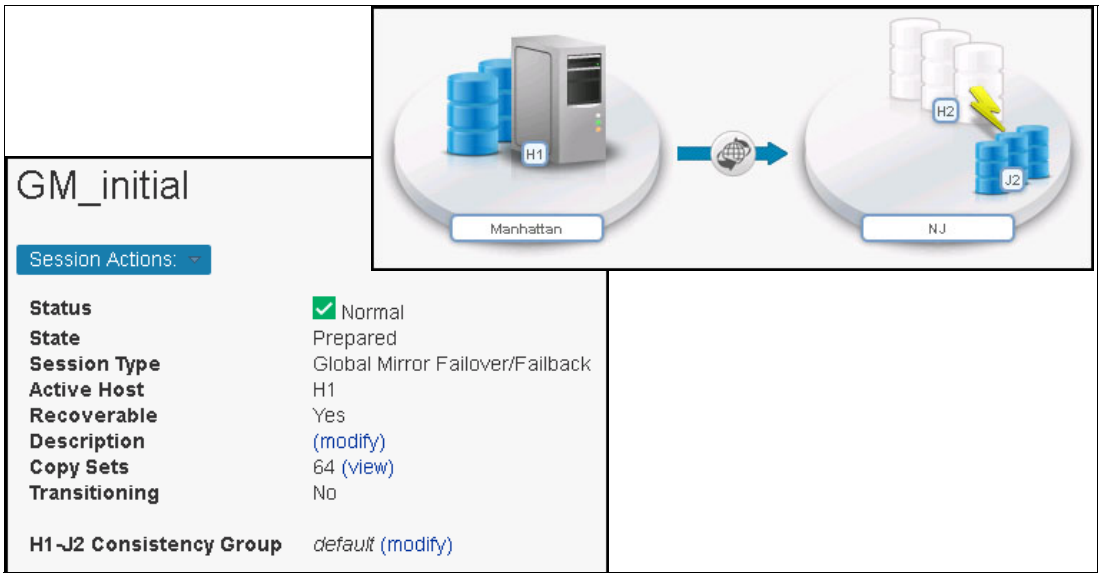

Figure 5-11 GM Session in Normal/Prepared state

Hints: CSM provides additional information

Copy Services Manager provides useful information about the session in a form that is easy to understand and in a single place.

Figure 5-12 shows an example of detailed information provided in the CSM GUI for each relation by hovering the mouse cursor over it.

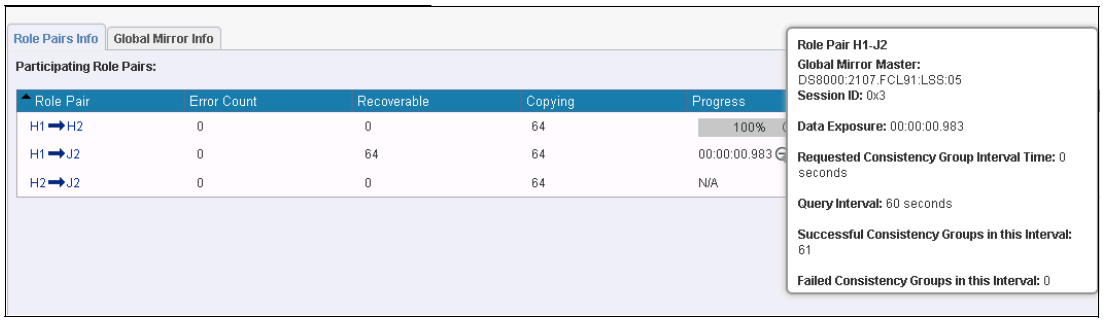

Figure 5-12 Relation details

Figure 5-13 provides a sample of additional information provided by CSM, such as *Coordination time*, *drain time*, or *Data exposure history* (last 15 minutes).

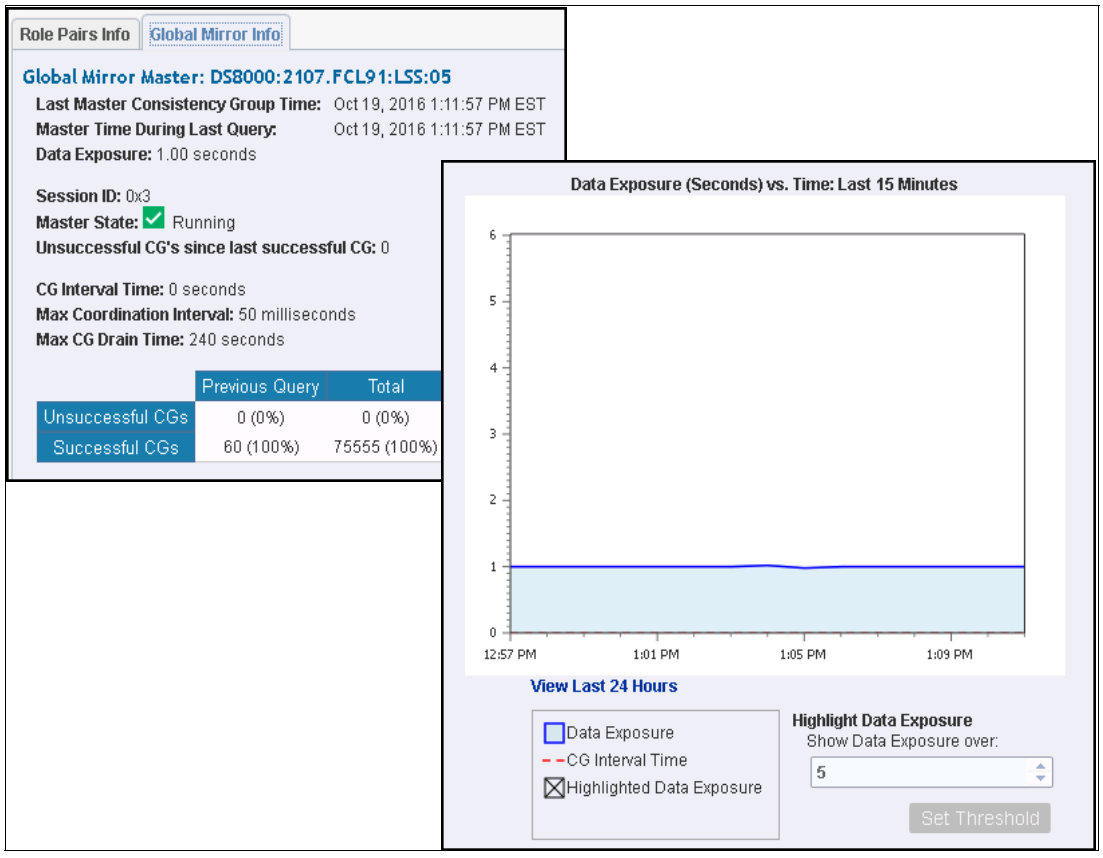

Figure 5-13 Additional information provided by CSM

5.1.2 Configuring a Multi-Target MM-GM session

The starting point of this task is an existing (active) GM session (asynchronous) from *Manhattan* to *NJ*. The objective of this configuration is to add a synchronous copy of the data in *Manhattan* site (H1) in a third site, named *PoK (H2)* while maintaining the asynchronous copy in the *NJ (H3)* site.

The multi-target MM-GM solution replicates data synchronously (Metro Mirror) from *Manhattan* to *PoK* and asynchronously (Global Mirror) from *Manhattan* to *NJ*, adding the high availability dimension to the initial DR solution (GM only).

Tip: When a session is first created in the Copy Services Manager, no action is performed on the storage systems before the session is explicitly started. As such, it is safe to create the new Multi-Target MM-GM session with the same set of volumes that are active in the current GM session.

Overview

This configuration has the following steps:

1. Creating (defining) a new session (we named it `MT_MM-GM`) and "populate" it with the existing copy sets (even though these already belong to an existing GM session), then add the new volumes from the PoK site used for the synchronous replication (*NJ* becomes *H3* in this configuration).
2. Removing the copy sets from the existing GM session. This removes the consistency groups (as logical entities) without stopping the Global Copy pairs at storage level.
3. Starting GC and GM from Manhattan (H1) to NJ (H3, J3).
4. Starting GC from Manhattan (H1) to PoK (H2).
5. Starting MM from Manhattan (H1) to PoK (H2).

Scenario

Complete the following steps:

1. We start by creating (defining) the new Multi-Target Metro Mirror - Global Mirror session. We specify the same site information as for the previous GM session, plus the data for the new site (PoK, which becomes H2), as shown in Figure 5-14.

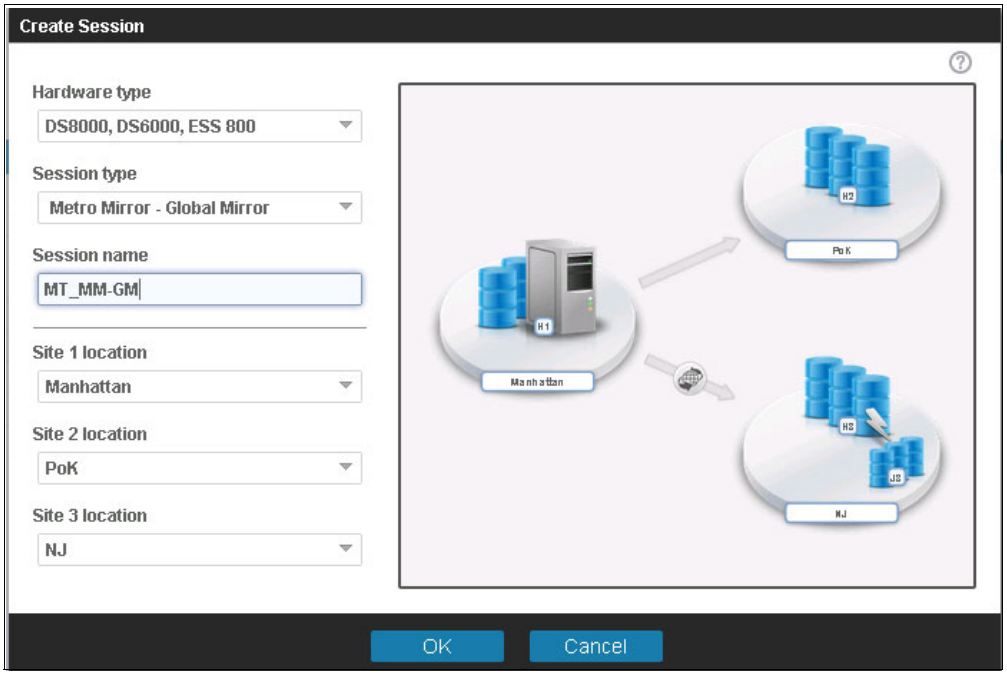

Figure 5-14 Creating the new MT_MM-GM session

2. Next, we add the copy sets. Click **Launch Add Copy Sets Wizard**, as shown in Figure 5-15.

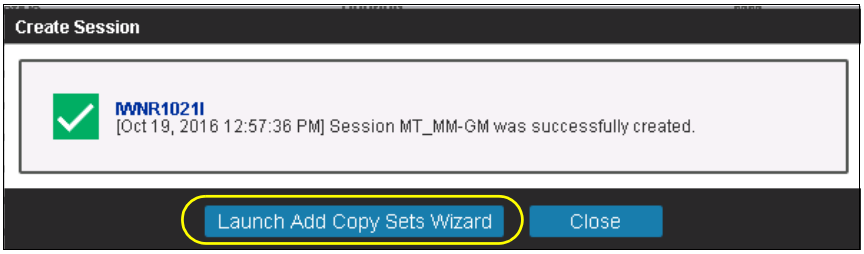

Figure 5-15 Launching the Add Copy Sets Wizard

3. One by one, we specify the H1 (Figure 5-16), H2 (Figure 5-17), H3 (Figure 5-18 on page 96) and J3 (Figure 5-19 on page 96) volumes.

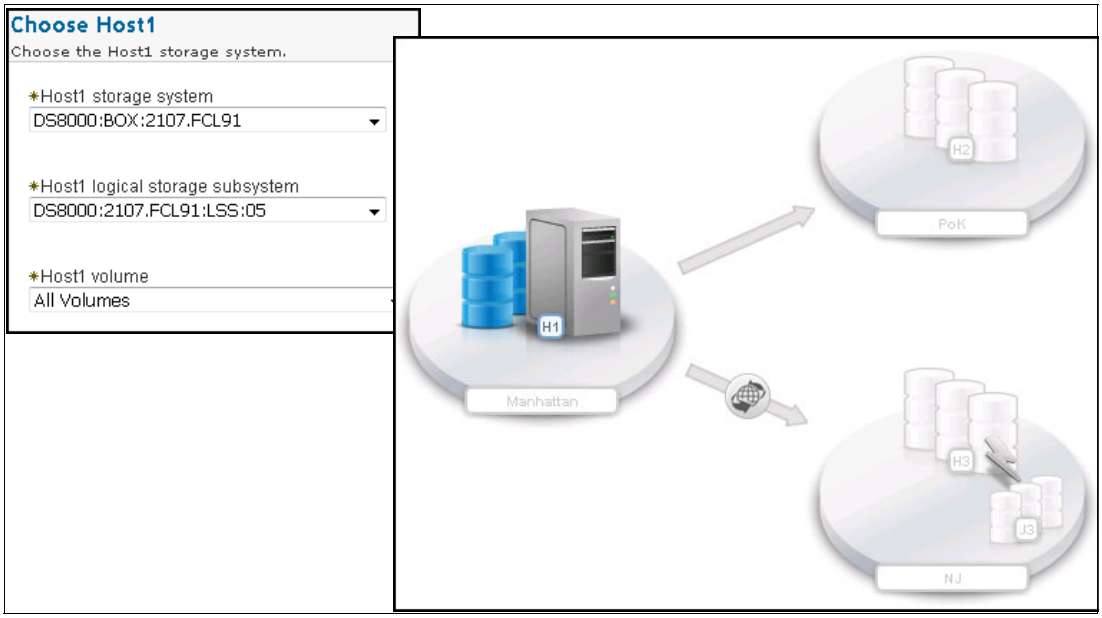

Figure 5-16 Host1 (Manhattan) volumes

4. For the Global Mirror relation, we replicate all volumes from LSS 05 in Manhattan to volumes from LSS 05 in NJ using journal volumes from LSS 06 in NJ.

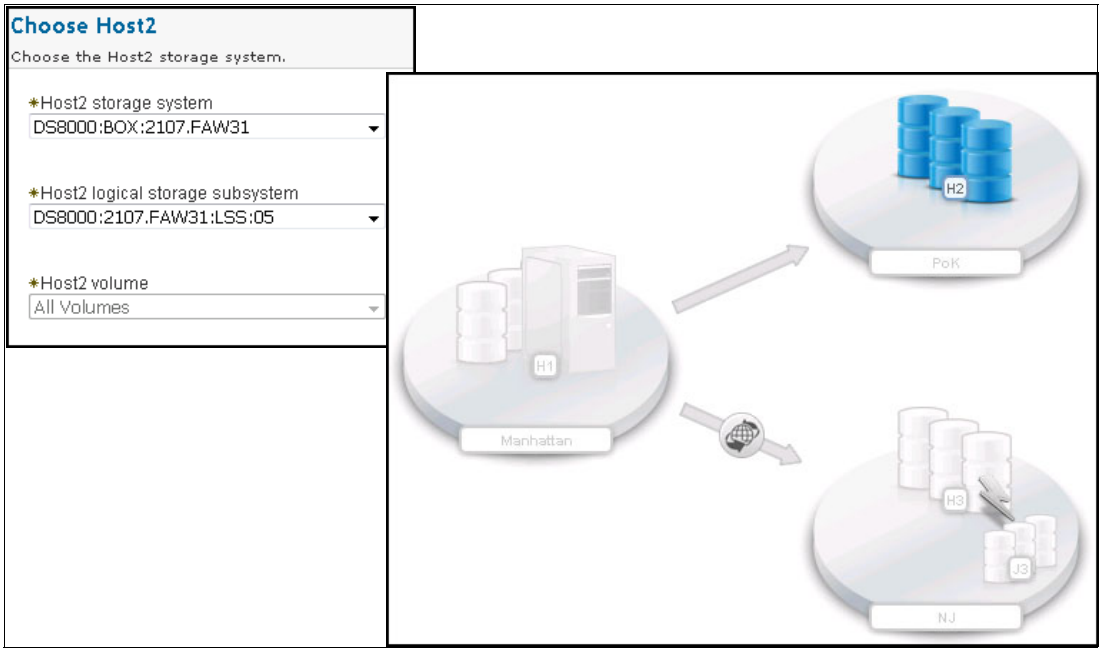

Figure 5-17 Host 2 (PoK) volumes

Optionally you can use a `.csv` format file to import the copy sets (in one operation).

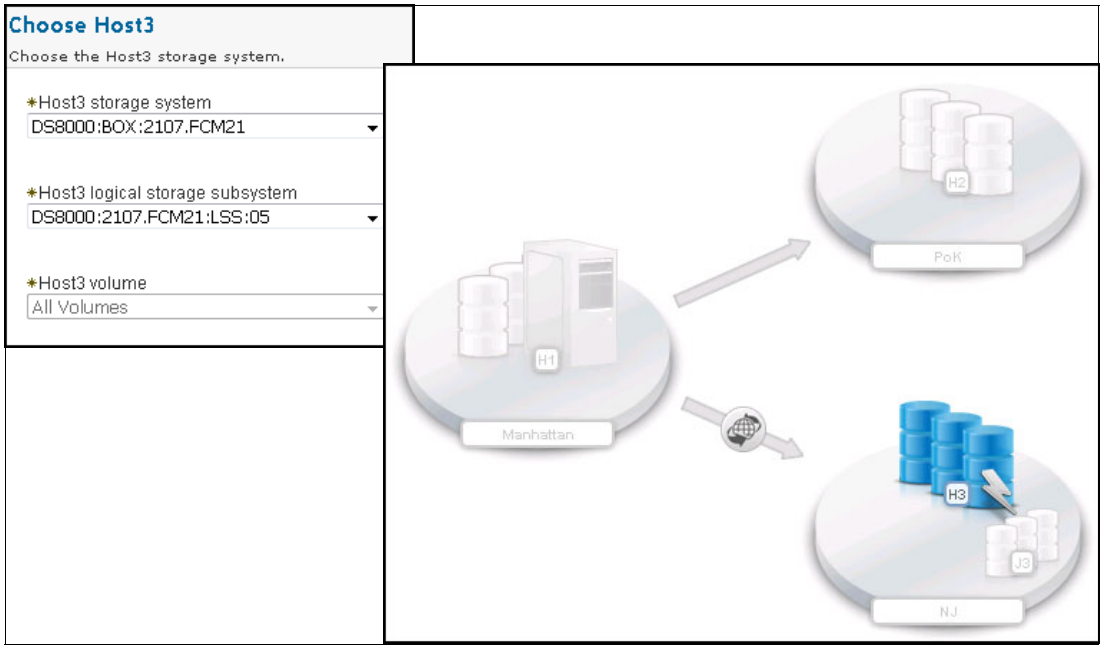

Figure 5-18 Host3 (NJ) volumes

Tip: During the matching process (for volumes and sites) the Copy Services Manager detects that the defined volumes are already part of an existing session and issues a Warning message.

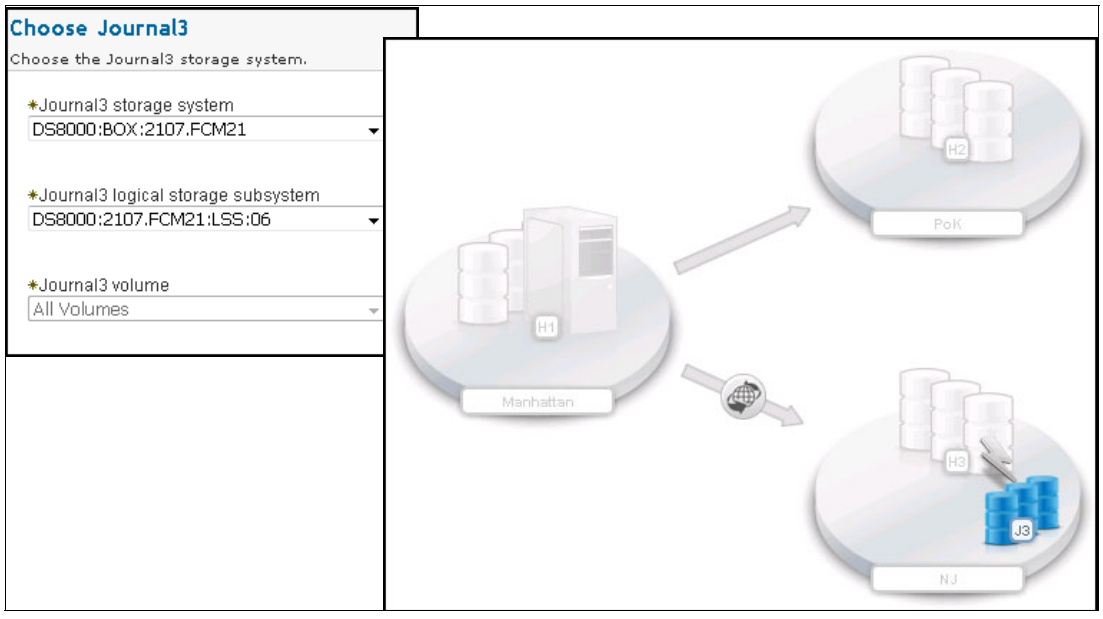

Figure 5-19 Journal3 (NJ) volumes

5. In this case, we chose to ignore the Warning message because we will not start this session while the old (GM) one is active. An example of the warning message in shown in Figure 5-20.

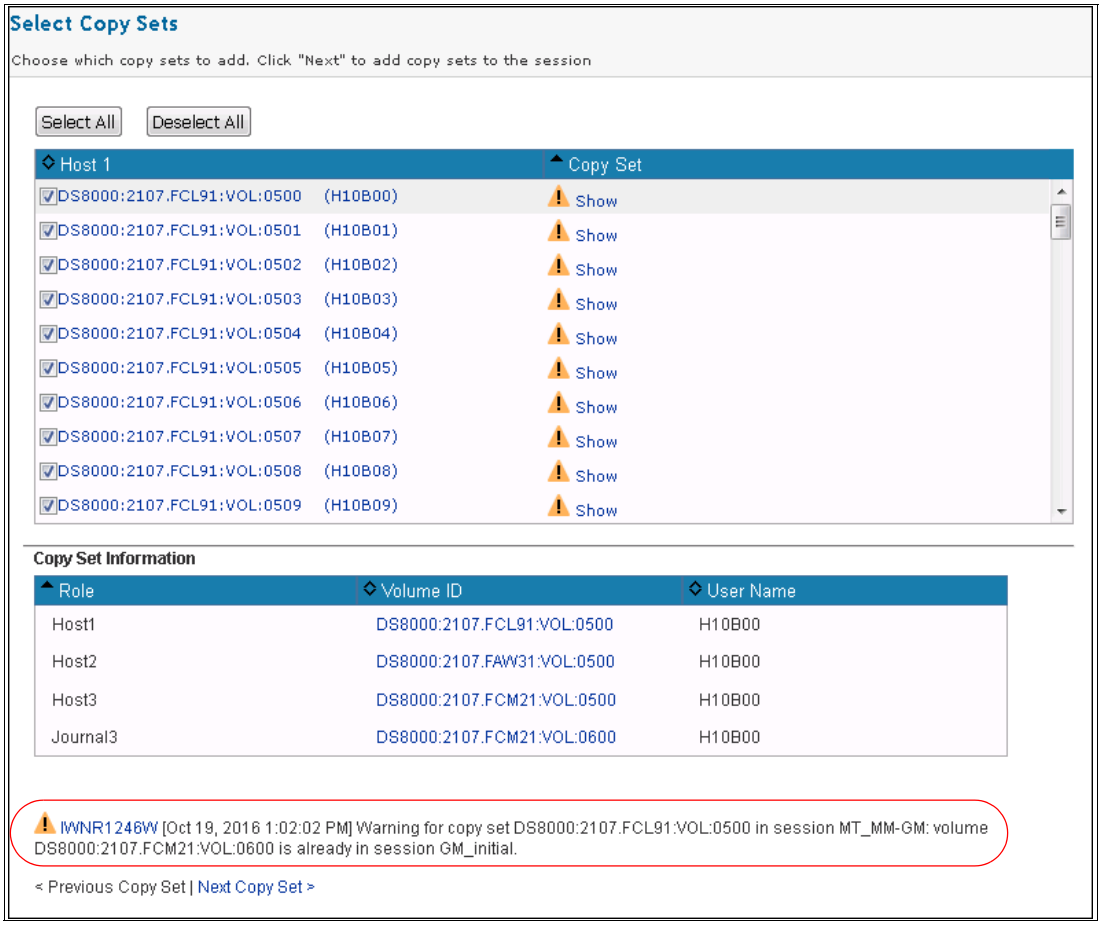

Figure 5-20 Warning message

6. When the copy sets have been added, the new Multi-Target Metro Mirror - Global Mirror is ready to be used (*Inactive/Defined* state). Figure 5-21 shows the Inactive/Defined status for the new session.

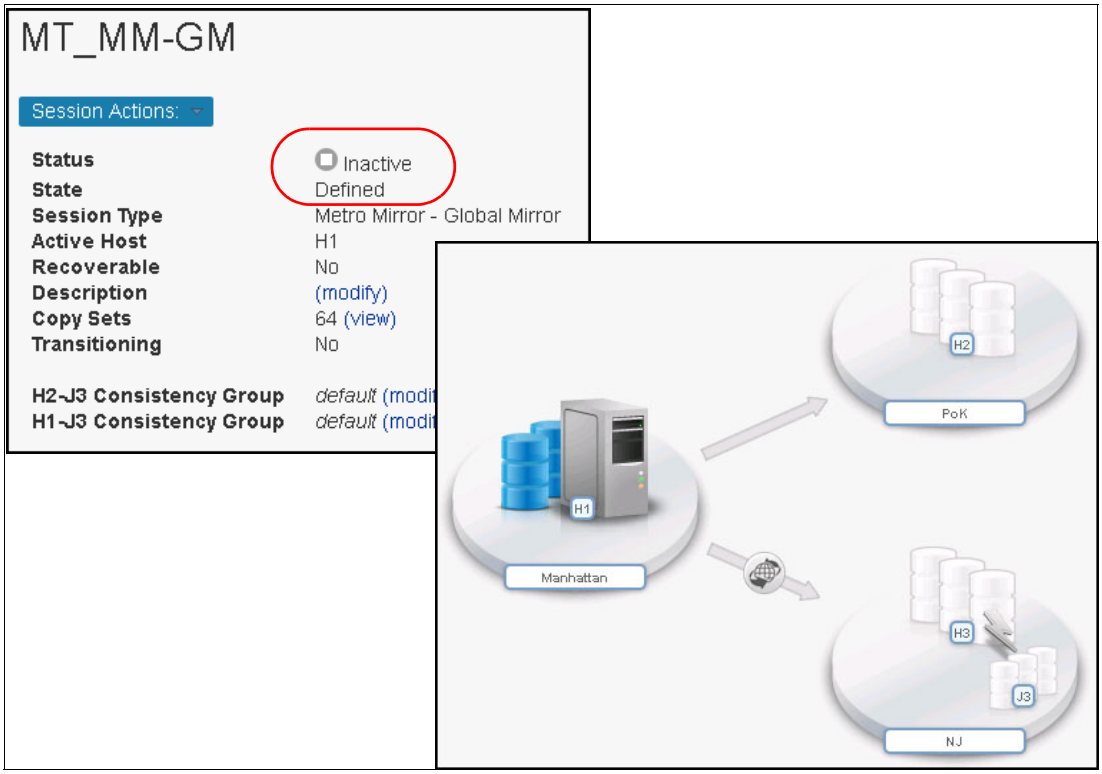

Figure 5-21 New MT_MM-GM session

7. In this step, we remove the copy sets from the active GM session. This does not stop the copy process, but there is no consistency in NJ because the FlashCopy relations used for consistency groups are removed. The Global Copy relation is active at the storage level.

   To remove the copy sets, we use the **Session Action** menu and select **View/Modify** → **Remove Copy Sets**, as shown in Figure 5-22.

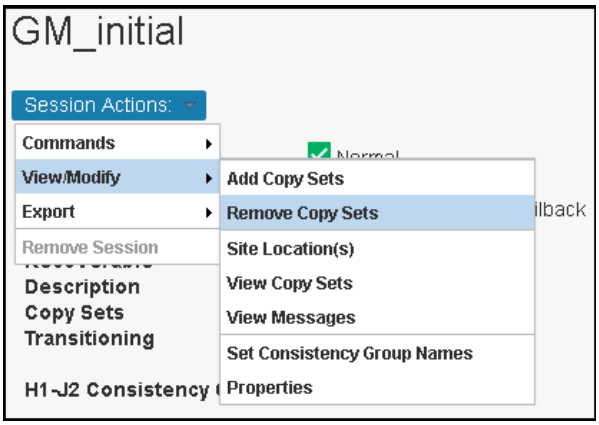

Figure 5-22 Remove Copy Sets

8. When prompted, only the source volumes (H1) must be specified, as shown in Figure 5-23.

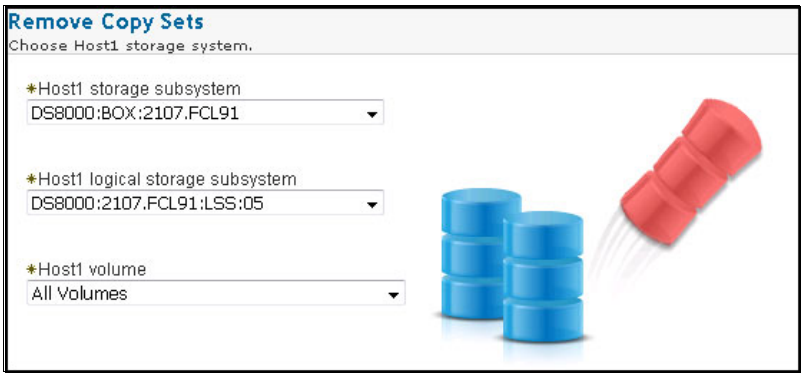

Figure 5-23 Removing Copy Sets from GM session

Important: During the confirmation of the removal action, we chose **Yes, keep the base relationships on the hardware**, as shown in Figure 5-24. This setting makes CSM leave the Global Copy active on the storage systems.

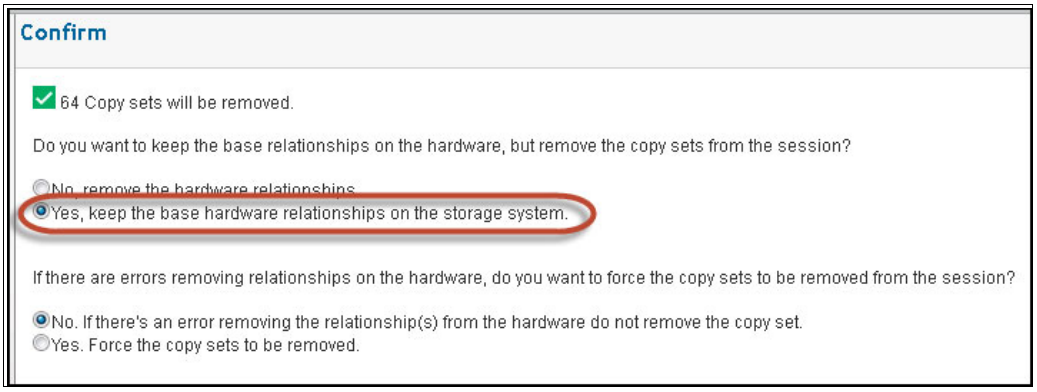

Figure 5-24 Maintaining HW relations on the storage system

9. While the copy sets are removed from the session we select the **Global Mirror Info** tab for the **Data Exposure** time, as shown in Figure 5-25.

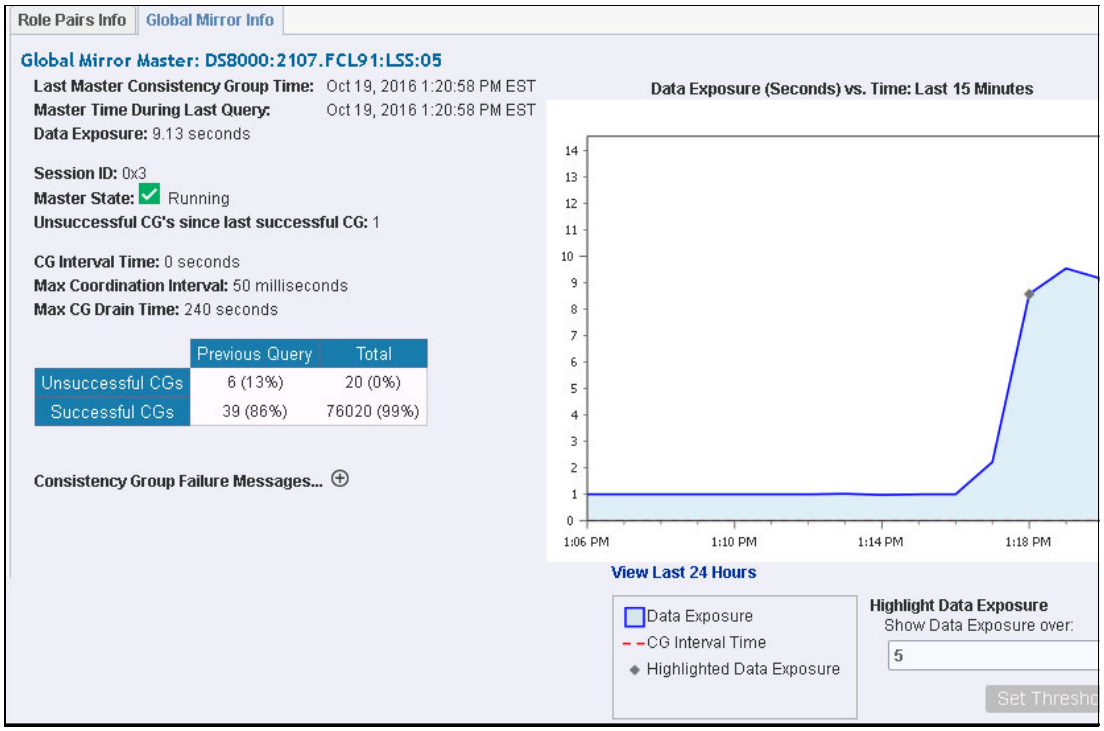

Figure 5-25 Global Mirror info

The old GM session becomes Inactive/Defined without any copy sets, as shown in Figure 5-26.

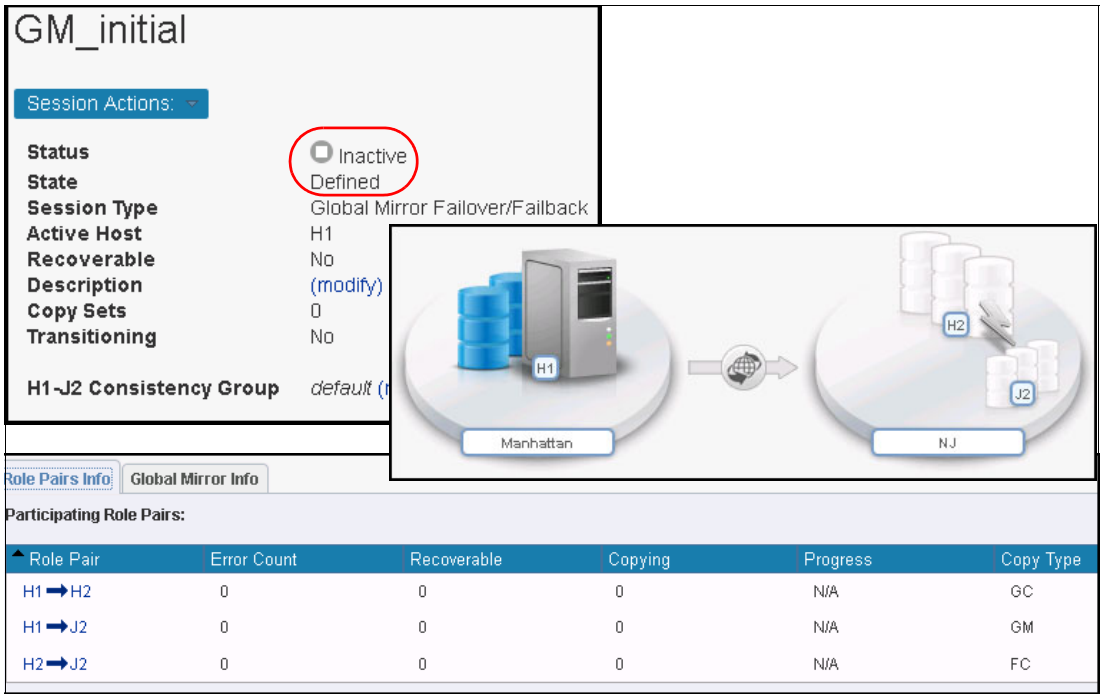

Figure 5-26 Empty GM session

10. Also, we can verify on storage level that the Global Copy relations are active (but there are no more Global Mirror relations because there is no FlashCopy, J3, in the NJ site), as shown in Example 5-2.

Example 5-2 Verifying relations using dscli

```
dscli> lssession 05
Date/Time: 19 octombrie 2016 21:30:46 EEST IBM DSCLI Version: 7.7.21.39 DS:
IBM.2107-13FCL91
CMUC00234I lssession: No Session found.
dscli>
dscli>
dscli> lspprc 0500-05FF
Date/Time: 19 octombrie 2016 21:29:31 EEST IBM DSCLI Version: 7.7.21.39 DS:
IBM.2107-13FCL91
ID         State        Reason  Type     Out Of Sync  SourceLSS Timeout (secs)
Critical Mode  First Pass Status
================================================================================
=============
0500:0500  Copy Pending  -  Global Copy  0  05  120  Disabled  False
0501:0501  Copy Pending  -  Global Copy  0  05  120  Disabled  False
0502:0502  Copy Pending  -  Global Copy  0  05  120  Disabled  False
0503:0503  Copy Pending  -  Global Copy  0  05  120  Disabled  False
0504:0504  Copy Pending  -  Global Copy  0  05  120  Disabled  False
0505:0505  Copy Pending  -  Global Copy  0  05  120  Disabled  False
0506:0506  Copy Pending  -  Global Copy  0  05  120  Disabled  False
0507:0507  Copy Pending  -  Global Copy  0  05  120  Disabled  False
0508:0508  Copy Pending  -  Global Copy  0  05  120  Disabled  False
0509:0509  Copy Pending  -  Global Copy  0  05  120  Disabled  False
...........
...........
053E:053E Copy Pending  -  Global Copy  0  05  120  Disabled  False
053F:053F Copy Pending  -  Global Copy  0  05  120  Disabled  False
```

11. On the new Multi Target Metro Mirror - Global Mirror session, first we start the GM from H1 to H3 in order to start the consistency groups for existing Global Copy relations.

    If the session was created from scratch (no existing GM session), you could start the GC first. In our case, we can go directly and start the GM relations.

12. From the **Session Action** menu, choose **Commands** → **Start**, as shown in Figure 5-27.

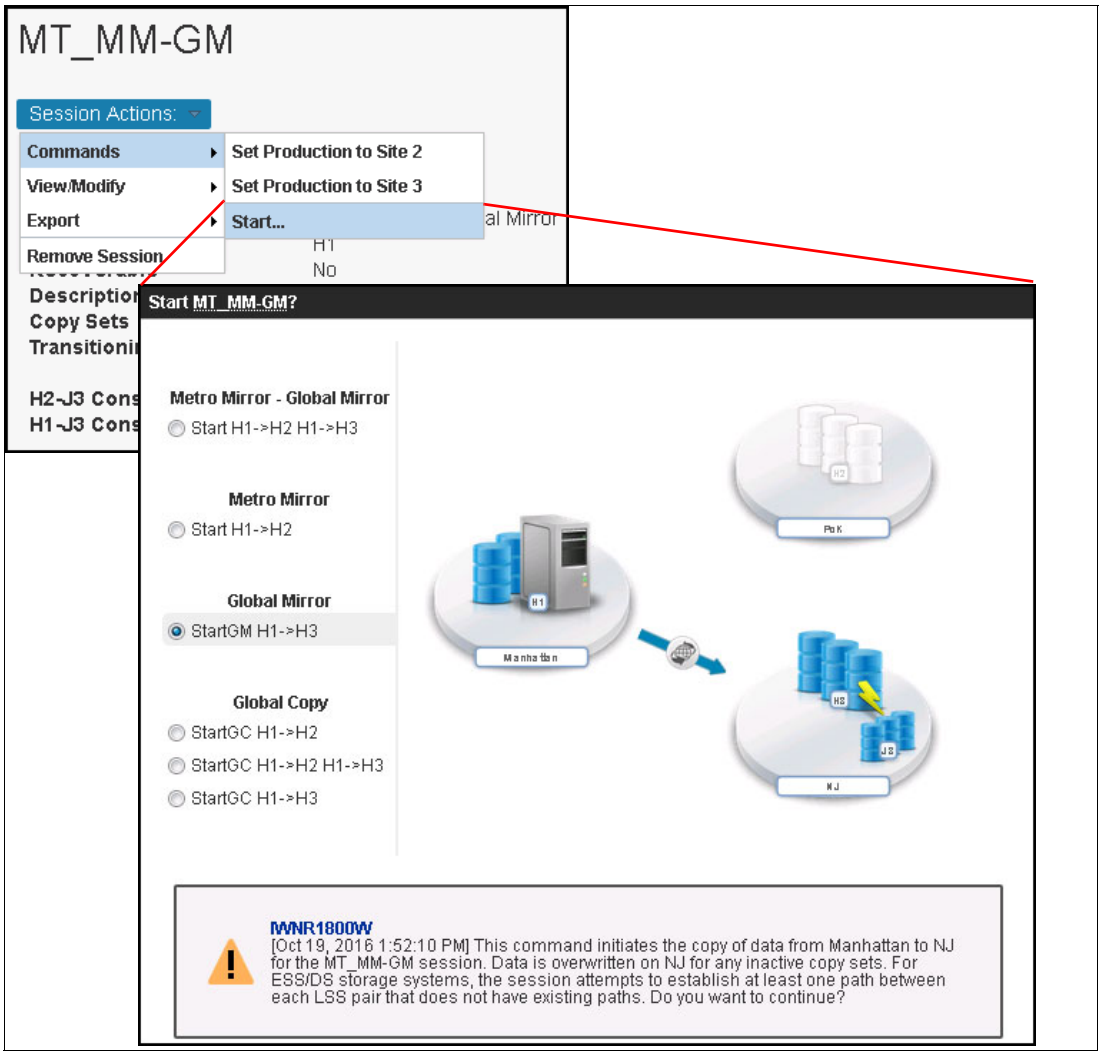

Figure 5-27 Starting the GM relations

13. When all copy sets have been reactivated, we are again ready for DR when J3 volumes become consistent, as shown in Figure 5-28.

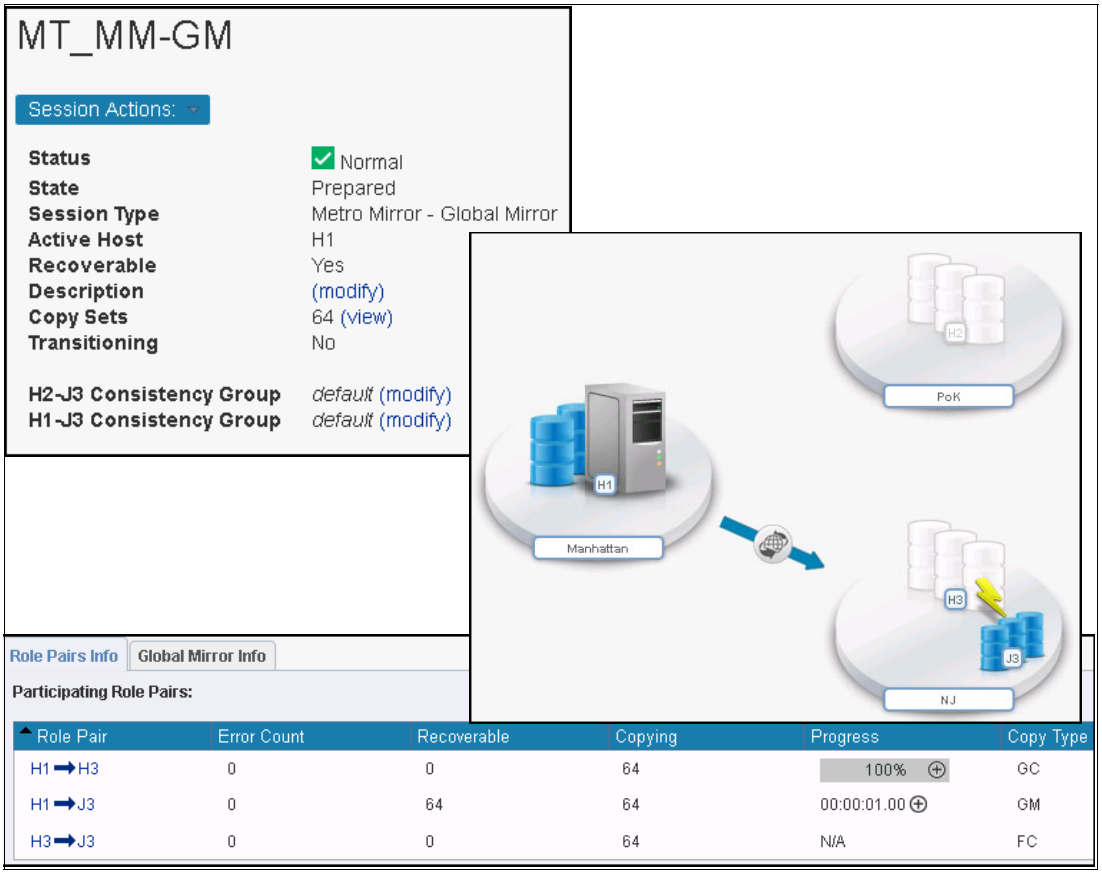

Figure 5-28 GM active

14. Due to the large amount of data that needs to be synchronized for the newly created MM relations, before starting the Metro Mirror relations we have to perform an asynchronous copy to H2 (to avoid any performance impact on production site - H1). Therefore, we start the GC H1->H2, as shown in Figure 5-29.

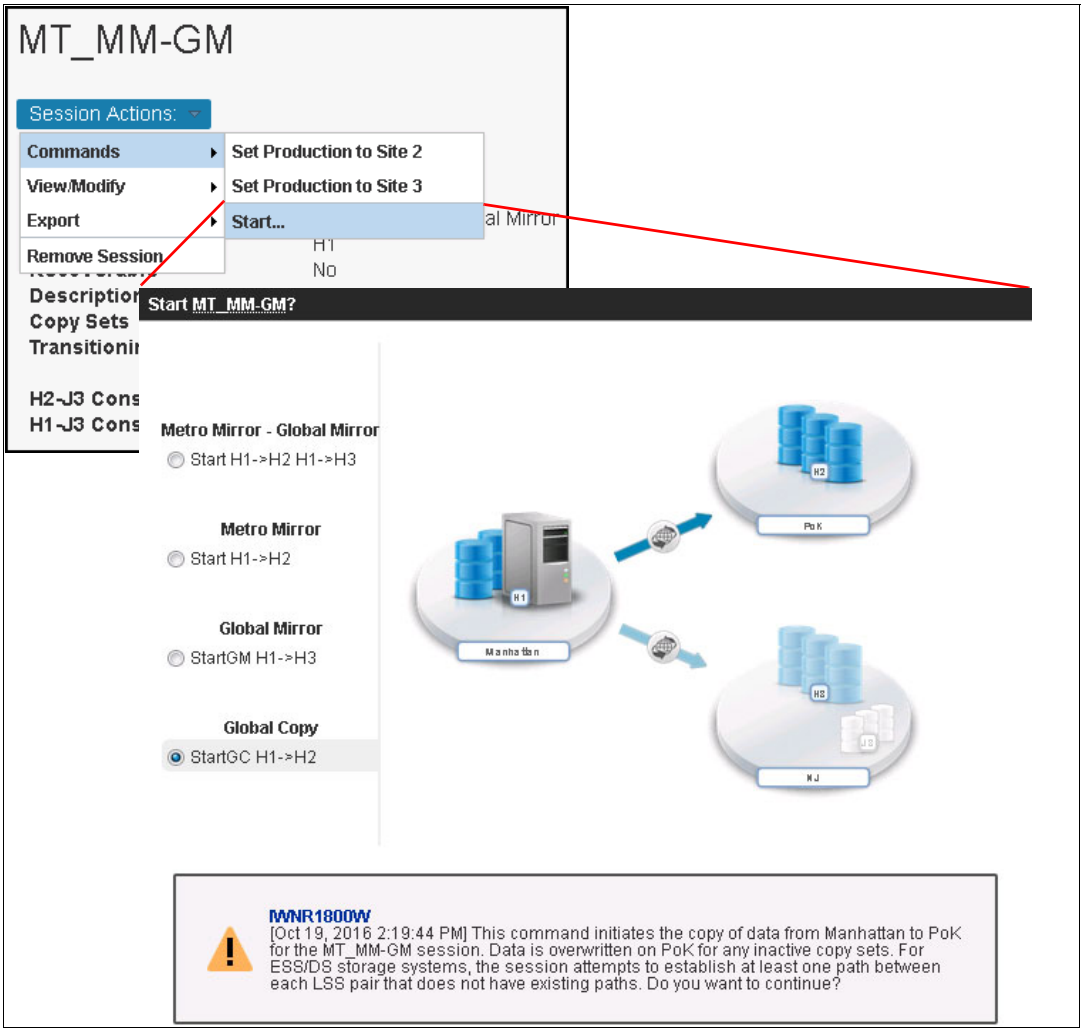

Figure 5-29 Global Copy H1->H2

15. By using GC, the data is copied from H1 to H2 without performance impact on the production site. GC operation progress is shown in Figure 5-30.

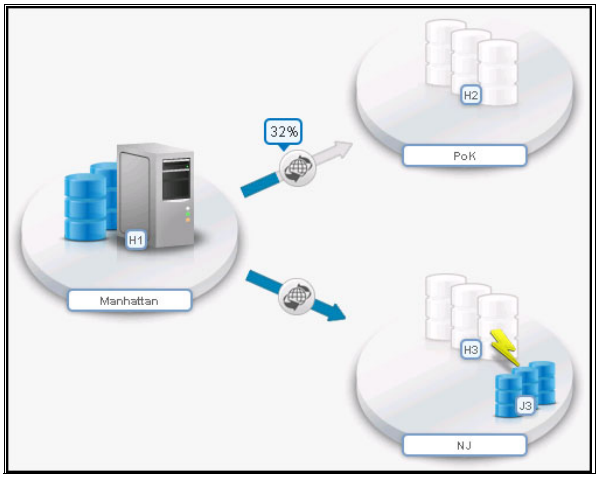

Figure 5-30 GC progress

16. When the GC progress has reached 100% (Out-of-Sync tracks is zero or almost zero) we can start the Metro Mirror (synchronous) relation (H1->H2), as shown in Figure 5-31.

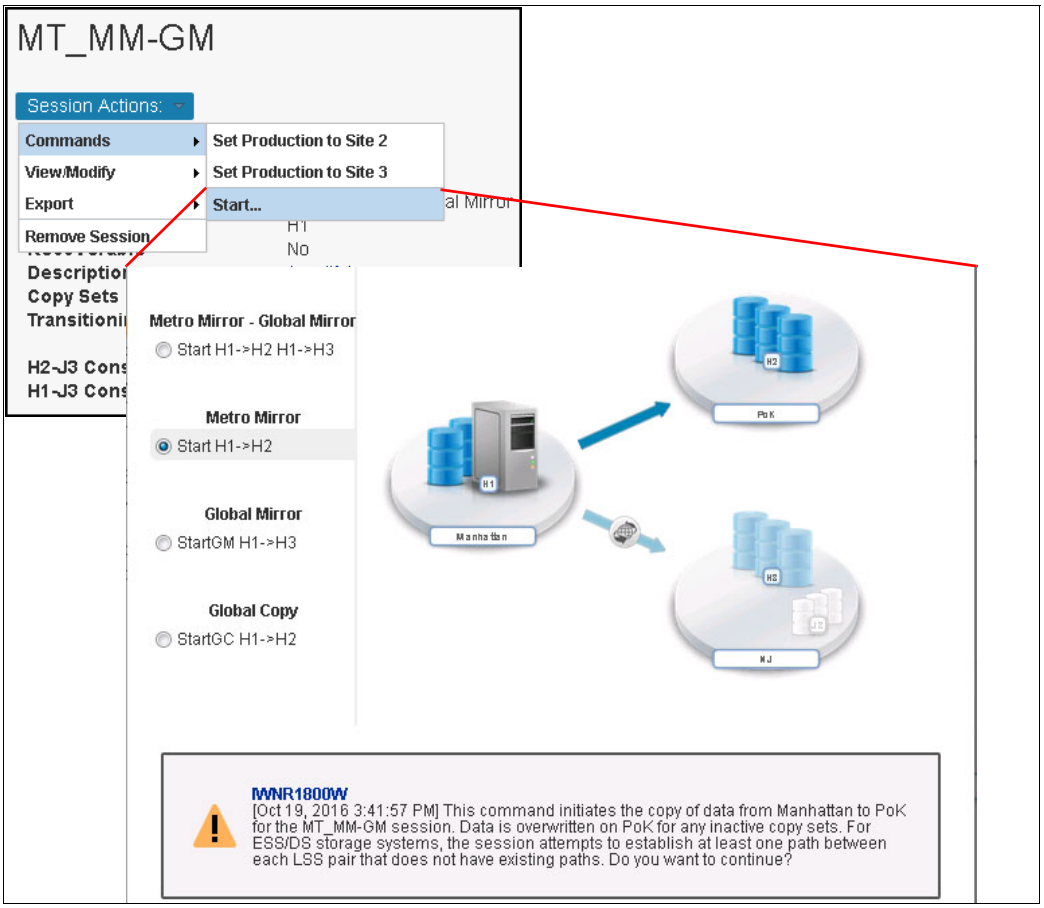

Figure 5-31 Starting Metro Mirror H1->H2

- Finally, we have a DR (Manhattan to NJ) solution with high availability (*Manhattan* to *PoK*) - Multi-Target Metro Mirror - Global Mirror, as shown in Figure 5-32.

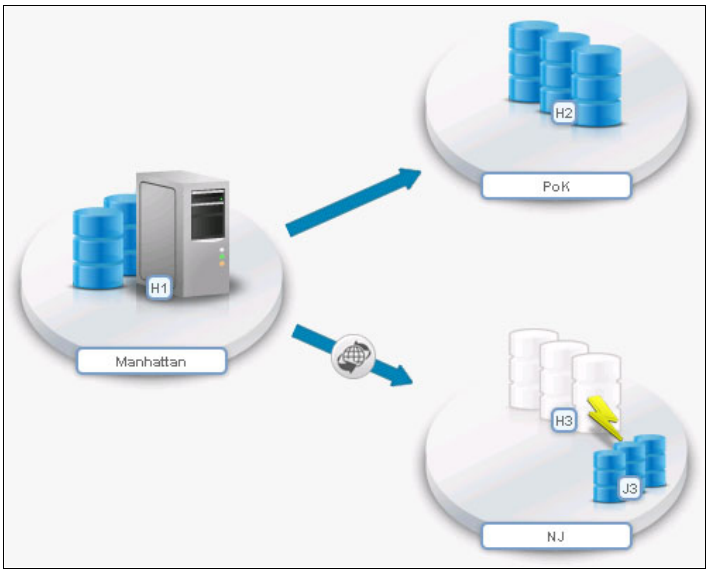

Figure 5-32 MT_MM-GM session

5.2 Migrating from a Basic z/OS HyperSwap session to Multi-Target Metro Mirror - Metro Mirror

Business case

In this scenario, we upgrade a synchronous DR solution to a multi-target, enhanced high availability replication solution without business impact (downtime).

The initial configuration (z/OS Basic HyperSwap) with two sites in synchronous replication distance (*PoK* and *Manhattan*) is later upgraded to an enhanced high availability solution stretched across three sites: *PoK*, *Manhattan*, and *NJ*. (Multi-Target Metro Mirror - Metro Mirror).

Scenario overview

The scenario consists of the following steps:

1. Creating the Basic HyperSwap session
2. Testing HyperSwap functionality
3. Preparing for migration: Creating the new Multi-Target MM-MM session
4. Managing the replication under the new session

5.2.1 Creating the Basic HyperSwap session

Attention: The current publication covers IBM Copy Services Manager (CSM) usage. z/OS HyperSwap (AIX HyperSwap) is discussed from a CSM perspective. We do not describe how to configure and test HyperSwap at an operating system level.

IBM z/OS Basic HyperSwap session requires that storage systems (DS8k) are configured for Metro Mirror. From CSM's perspective, the Basic HyperSwap session is similar to a Metro Mirror session: source volumes (H1) are copied synchronously to the target volumes (H2).

The z/OS HyperSwap can be configured in a z/OS image or in a parallel sysplex. The z/OS HyperSwap component is responsible for redirecting the I/O operations to a specific site (storage system).

In our example, we use a parallel sysplex that has been configured in the Copy Services Manager server. The communication configuration between the CSM server and the z/OS HyperSwap is already configured and uses the z/OS Java Native Interface (JNI), as the CSM server is installed in one z/OS image that is part of the parallel sysplex for this scenario:

1. To create a new Basic HyperSwap session, we click **Create Session** in the Session section in the GUI, as shown in Figure 5-33. We enter the session name and the two locations, Site 1, production (*PoK*), and Site 2, DR (*Manhattan*).

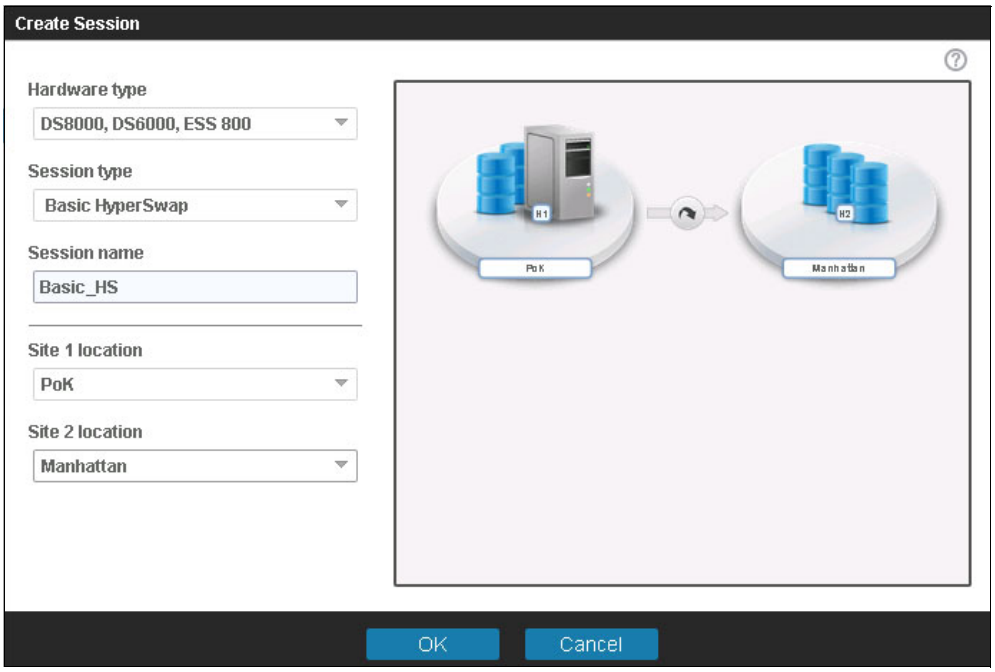

Figure 5-33 New Session for HyperSwap

2. Next, we add the copy sets: click **Launch Add Copy Sets Wizard**, as shown in Figure 5-34.

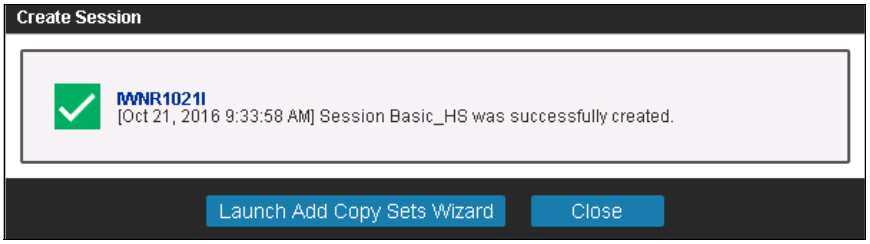

Figure 5-34 Launch Add Copy Sets Wizard

3. One by one, we specify the H1 (Figure 5-35 on page 108) and H2 (Figure 5-36 on page 108) volumes. Optionally you can use a `.csv` format file to import all the copy sets.

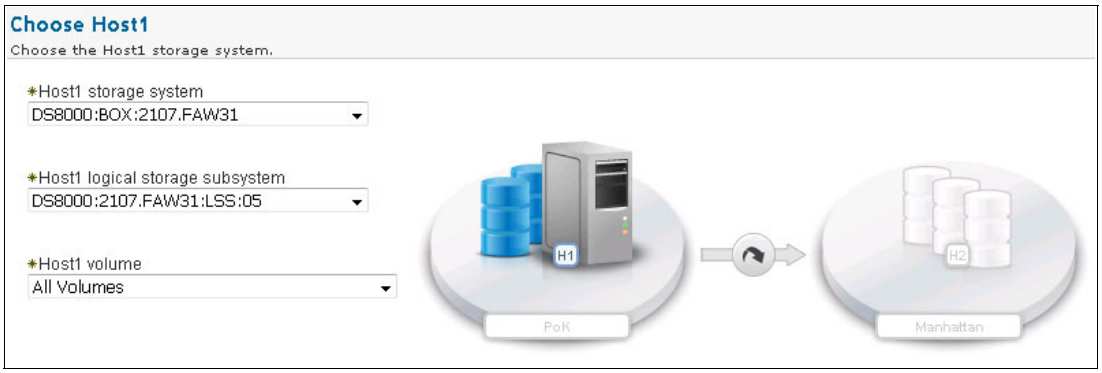

Figure 5-35 Choosing H1 volumes

4. In our case, we replicate all volumes from LSS 05 in *PoK* to volumes from LSS 05 in *Manhattan*.

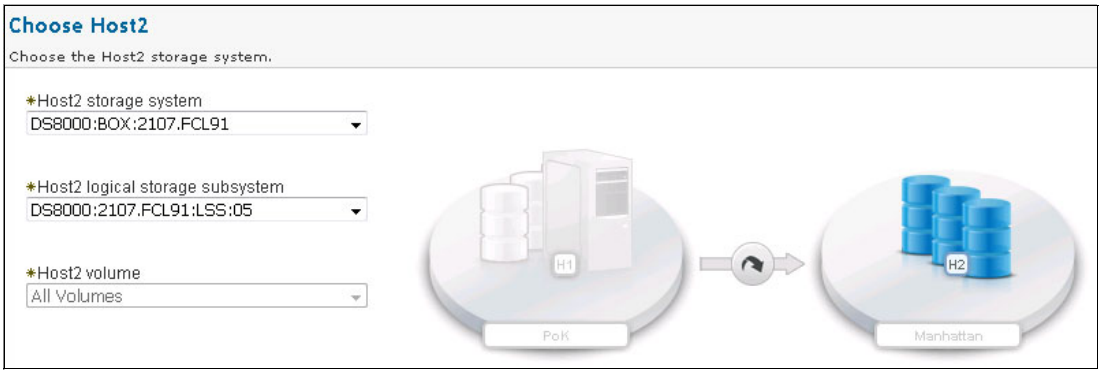

Figure 5-36 Choose H2 volumes

5. When all copy sets have been entered and the wizard has completed, the session becomes *Inactive/Defined*, as shown in Figure 5-37.

   At this time, the CSM server has not yet sent any commands to the storage system, so there are no Metro Mirror (PPRC) relations set up on storage level.

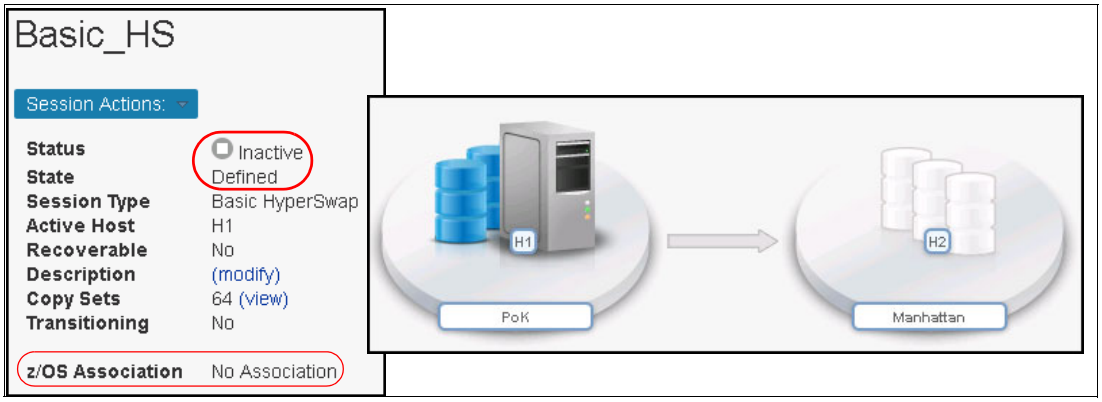

Figure 5-37 Initial state

6. Also, there is no *z/OS Association* yet. Before starting any action, go to **Session Actions** → **View/Modify** → **Properties** first, as shown in Figure 5-38.

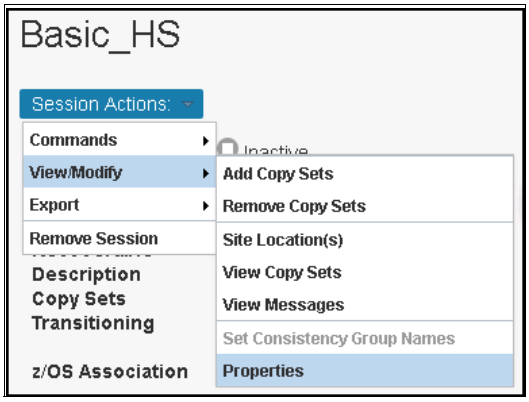

Figure 5-38 Session properties

In the *H1-H2 Options* tab we can see that there is no z/OS association yet, and also an informational message, in red, which explains that a *system or a sysplex association is required for hardened freeze and HyperSwap management.*

7. Here we select the required z/OS system,[2] or sysplex, as shown in Figure 5-39.

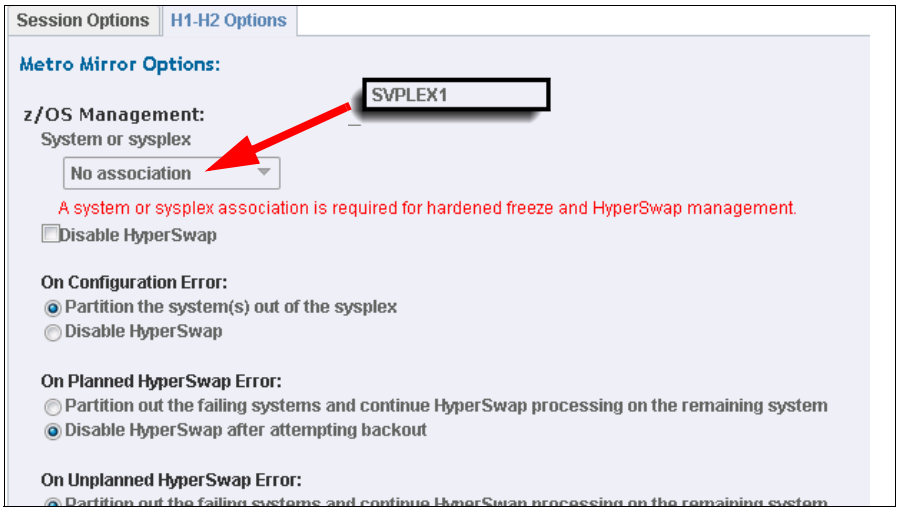

Figure 5-39 H1-H2 Options

[2] The z/OS systems (Host and parallel sysplex information) have been previously configured in CSM.

Figure 5-40 shows the session *z/OS Association*.

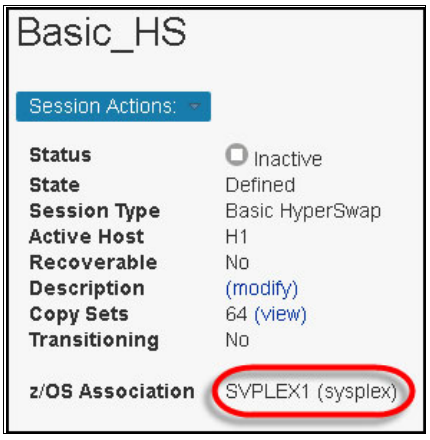

Figure 5-40 z/OS association

8. Next, we start the synchronous copy from H1 to H2. If volumes in H2 were used before, all data is overwritten by the new Metro Mirror relations. Figure 5-41 shows the warning that explains the actions for **Start H1->H2**.

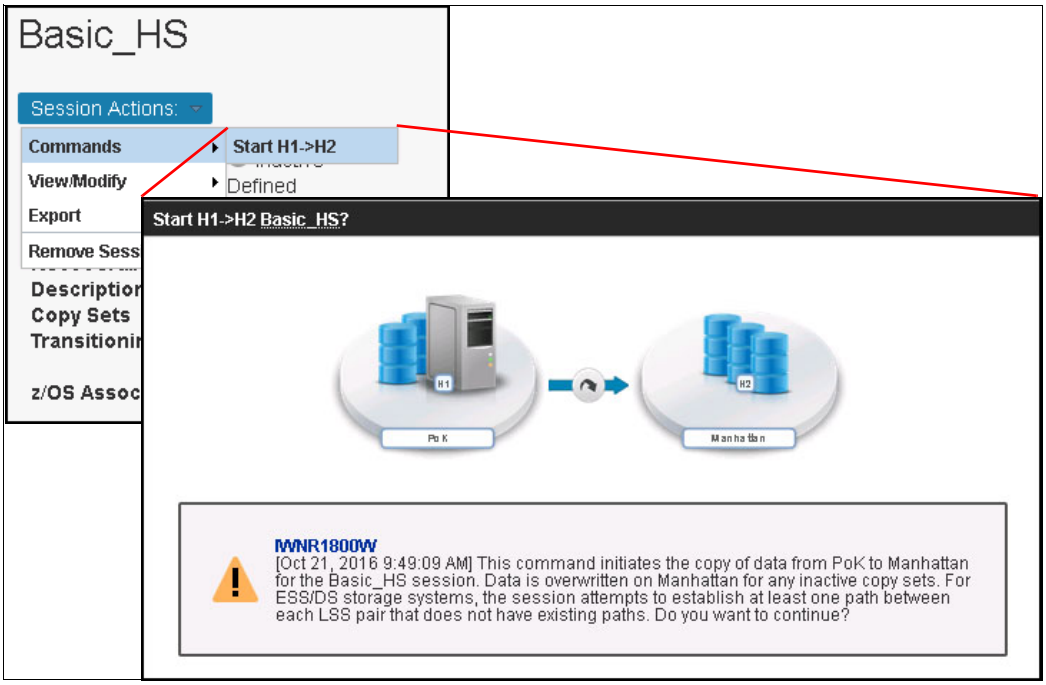

Figure 5-41 Start H1->H2

9. The session changes to the *Preparing* state while the data is copied to the *Manhattan* site, as shown in Figure 5-42.

   The warning status means that the volumes are being synchronized (see 1.5.2, “Monitoring sessions icons and symbols” on page 9).

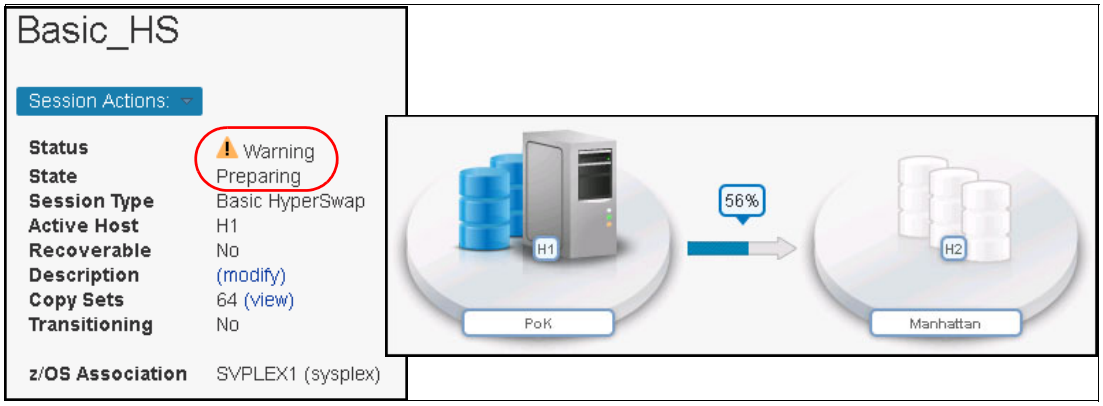

Figure 5-42 New Session in Preparing state

10. You can check the newly created Metro Mirror relations and the copy process, as shown in Example 5-3.

Example 5-3 Checking status at storage system level (dscli)

```
dscli> lspprc -type mmir 0500-05FF
Date/Time: 21 octombrie 2016 17:50:42 EEST IBM DSCLI Version: 7.7.21.39 DS:
IBM.2107-75FAW31
ID State Reason Type    SourceLSS Timeout (secs) Critical Mode First Pass Status
=======================================================================================
0500:0500 Copy Pending -   Metro Mirror 05   120   Disabled   Invalid
0501:0501 Copy Pending -   Metro Mirror 05   120   Disabled   Invalid
0502:0502 Copy Pending -   Metro Mirror 05   120   Disabled   Invalid
0503:0503 Full Duplex  -   Metro Mirror 05   120   Disabled   Invalid
0504:0504 Full Duplex  -   Metro Mirror 05   120   Disabled   Invalid
0505:0505 Full Duplex  -   Metro Mirror 05   120   Disabled   Invalid
0506:0506 Copy Pending -   Metro Mirror 05   120   Disabled   Invalid
0507:0507 Full Duplex  -   Metro Mirror 05   120   Disabled   Invalid
0508:0508 Full Duplex  -   Metro Mirror 05   120   Disabled   Invalid
0509:0509 Copy Pending -   Metro Mirror 05   120   Disabled   Invalid
050A:050A Full Duplex  -   Metro Mirror 05   120   Disabled   Invalid
050B:050B Copy Pending -   Metro Mirror 05   120   Disabled   Invalid
050C:050C Full Duplex  -   Metro Mirror 05   120   Disabled   Invalid
050D:050D Copy Pending -   Metro Mirror 05   120   Disabled   Invalid
...........
...........
053B:053B Copy Pending -   Metro Mirror 05   120   Disabled   Invalid
053C:053C Full Duplex  -   Metro Mirror 05   120   Disabled   Invalid
053D:053D Full Duplex  -   Metro Mirror 05   120   Disabled   Invalid
053E:053E Full Duplex  -   Metro Mirror 05   120   Disabled   Invalid
053F:053F Full Duplex  -   Metro Mirror 05   120   Disabled   Invalid
dscli>
```

11. After the initial copy process is complete, the session changes to *Normal/Prepared*, as shown in Figure 5-43.

    This means that all Metro Mirror relations are in *Full Duplex* state and that the systems are ready for recovery in case of a disaster.

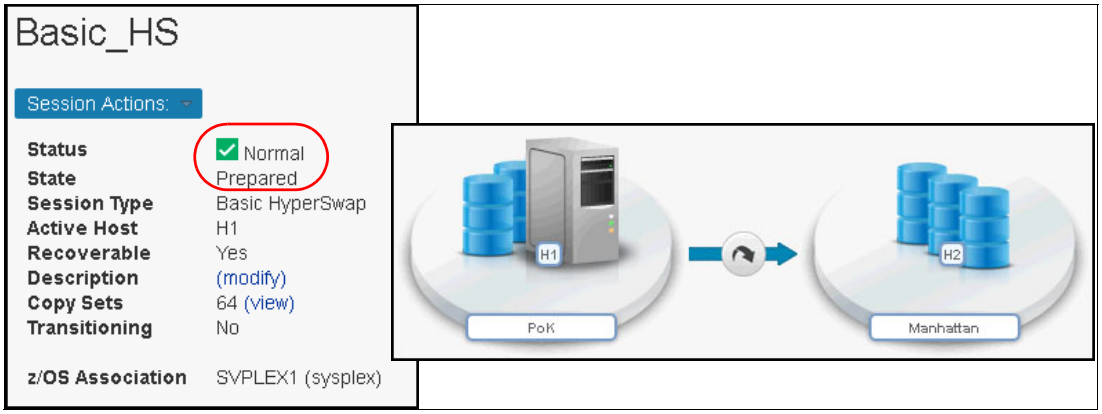

Figure 5-43 Prepared state

5.2.2 Testing the HyperSwap functionality (z/OS)

Starting with CSM 6.1, in addition to the **`Start`** and **`Stop`** session commands (specific to any Metro Mirror sessions), there is a new command, **`HyperSwap`** (see Figure 5-44).

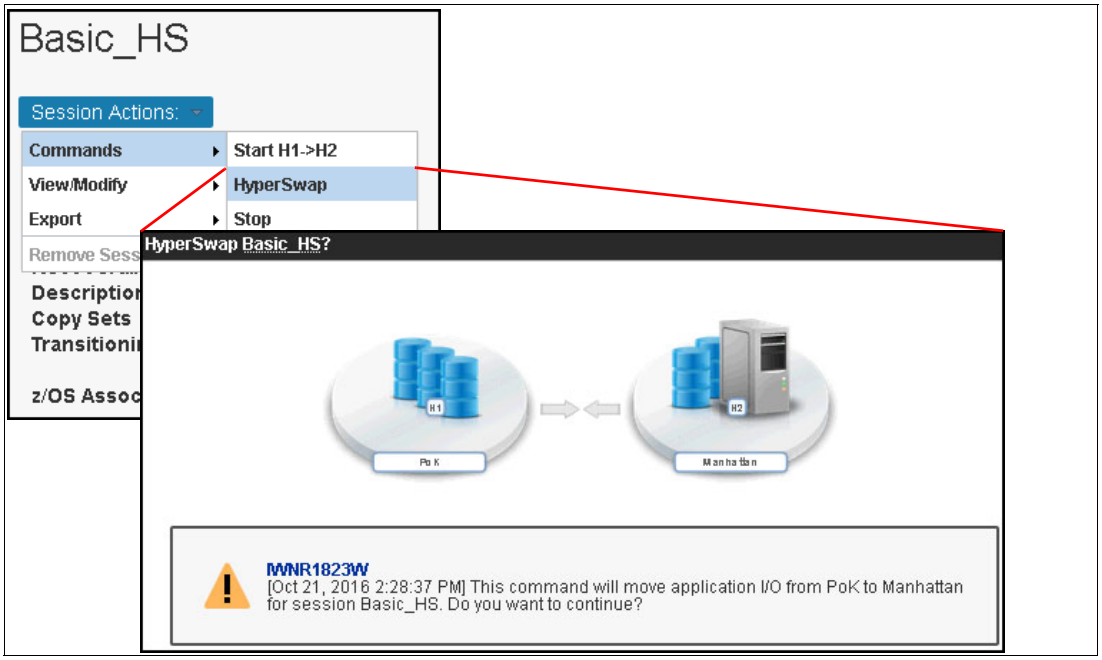

Figure 5-44 HyperSwap session command

Important: When activating the HyperSwap action on a session, IBM Copy Services Manager sends the commands to the z/OS HyperSwap in the Sysplex (or z/OS image) to move the I/O on the other site. CSM does not issue the commands directly to the storage for performing the HyperSwap action. HyperSwap takes over and performs the necessary actions at storage system level. IBM Copy Services Manager *receives and displays updated information* for the session status while z/OS HyperSwap *performs the actions*.

When the HyperSwap has been initiated, the Metro Mirror relations are suspended and the I/O is redirected to H2 (*Manhattan*). Also, the session state changes to *Target Available*, as shown in Figure 5-45.

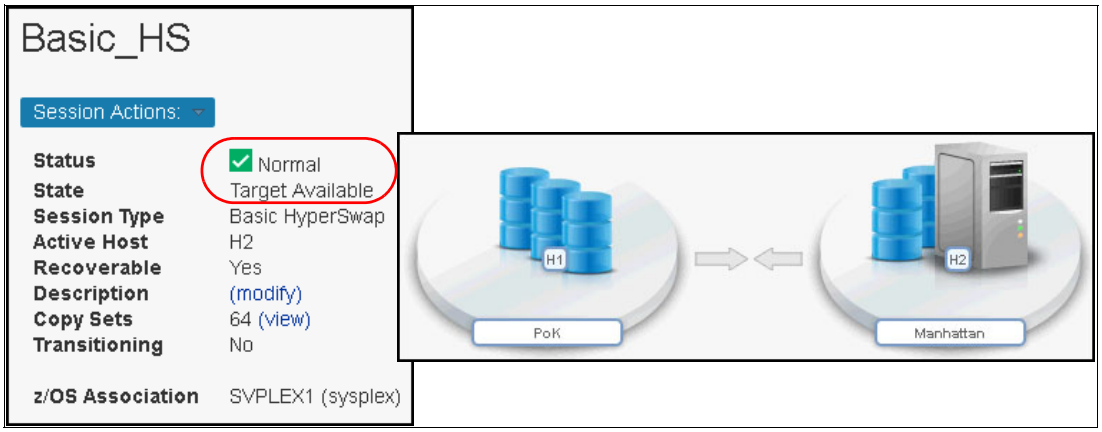

Figure 5-45 Target Available

Example 5-4 shows the **dscli** output, and we can observe the *Suspended* status of the Metro Mirror relationships.

Example 5-4 Checking HyperSwap results (dscli)

```
dscli>
dscli> lspprc 0500-05FF
Date/Time: 21 octombrie 2016 22:34:06 EEST IBM DSCLI Version: 7.7.21.39 DS:
IBM.2107-75FAW31
ID        State     Reason      Type         SourceLSS Timeout (secs) Critical
Mode First Pass Status
==========================================================================================
0500:0500 Suspended Host Source Metro Mirror 05   120   Disabled   Invalid
0501:0501 Suspended Host Source Metro Mirror 05   120   Disabled   Invalid
0502:0502 Suspended Host Source Metro Mirror 05   120   Disabled   Invalid
0503:0503 Suspended Host Source Metro Mirror 05   120   Disabled   Invalid
0504:0504 Suspended Host Source Metro Mirror 05   120   Disabled   Invalid
0505:0505 Suspended Host Source Metro Mirror 05   120   Disabled   Invalid
0506:0506 Suspended Host Source Metro Mirror 05   120   Disabled   Invalid
0507:0507 Suspended Host Source Metro Mirror 05   120   Disabled   Invalid
0508:0508 Suspended Host Source Metro Mirror 05   120   Disabled   Invalid
0509:0509 Suspended Host Source Metro Mirror 05   120   Disabled   Invalid
..........
..........
053E:053E Suspended Host Source Metro Mirror 05   120   Disabled   Invalid
053F:053F Suspended Host Source Metro Mirror 05   120   Disabled   Invalid
dscli>
```

At this time, all I/Os are performed to H2 (*Manhattan*) and the replication is not active. To bring back the configuration ready for DR, we activate the replication from *Manhattan* to *PoK* (*Start H2->H1*), as shown in Figure 5-46.

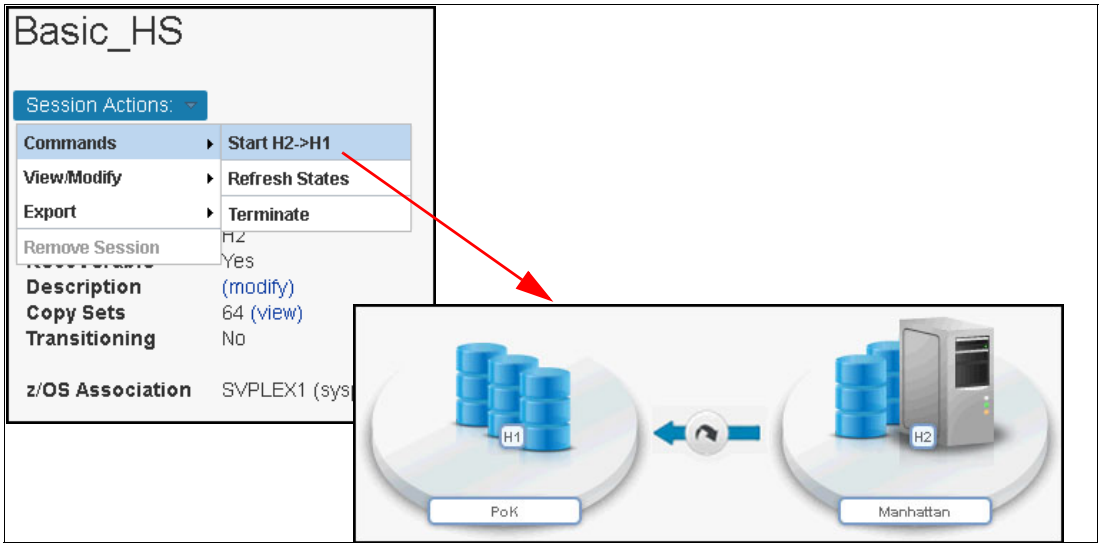

Figure 5-46 Start H2-H1

To revert, you can switch back the I/O to the initial direction (*PoK* (H1) -> *Manhattan* (H2)).

5.2.3 Migrating to a Multi-Target MM-MM session

At this point, we have tested the high availability solution between sites *PoK* and *Manhattan*. The new solution should improve configuration availability by extending (migrating) the current HyperSwap configuration to cover a third site (named *NJ*).

Tip: When a session is created in Copy Services Manager, no implied (automatic) action is performed at the storage system level. Therefore, it is safe to create the new Multi -Target MM-MM session with the same set of volumes that are active on the current MM session.

To extend the configuration to a Multi-Target Metro Mirror - Metro Mirror configuration we perform the following high-level actions.:

1. Define/Create a new MT_MM-MM session and add the existing copy sets and the new volumes from NJ site (volumes in NJ site are also used for synchronous replication).
2. Remove the copy sets from existing Basic HyperSwap session.
   The Metro Mirror relations remain active.
3. Start MM from *PoK* (H1) to *Manhattan* (H2)
4. Start GC from *PoK* (H1) to *NJ* (H3)
5. Start MM from *PoK* (H1) to *NJ* (H3)

Specifically, complete the following steps:

1. To create the new Multi-Target Metro Mirror - Metro Mirror session we need to specify the same locations as for the previous MM session, and add the location of the new site (*NJ*), which becomes H3, as shown in Figure 5-14.

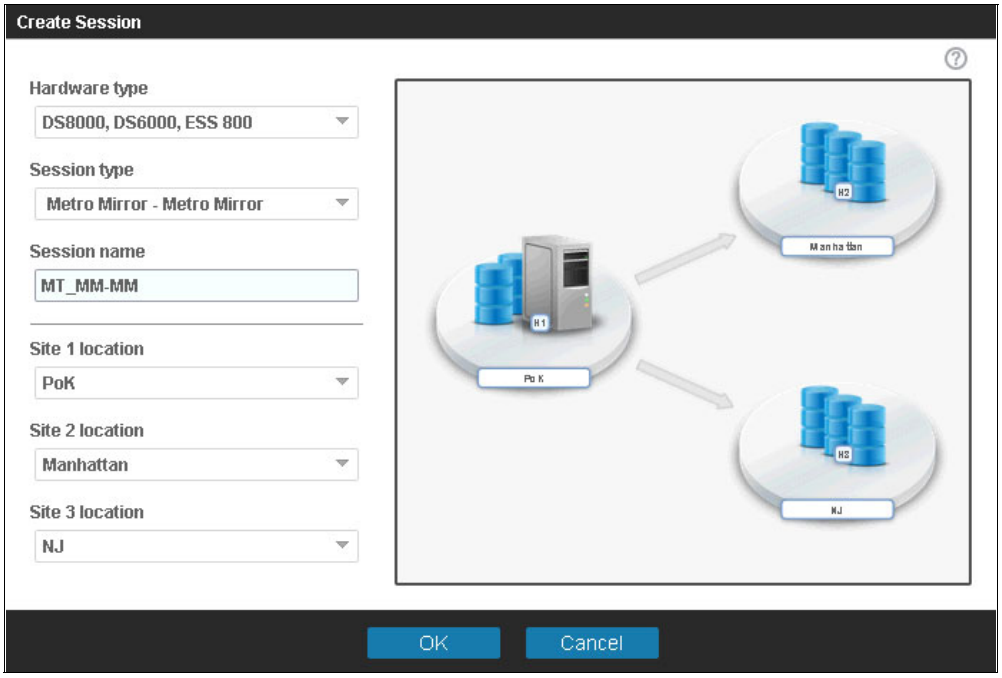

Figure 5-47 Create MT_MM-GM session

2. Next, we add copy sets to the session definition. We click **Launch Add Copy Sets Wizard**, as shown in Figure 5-15.

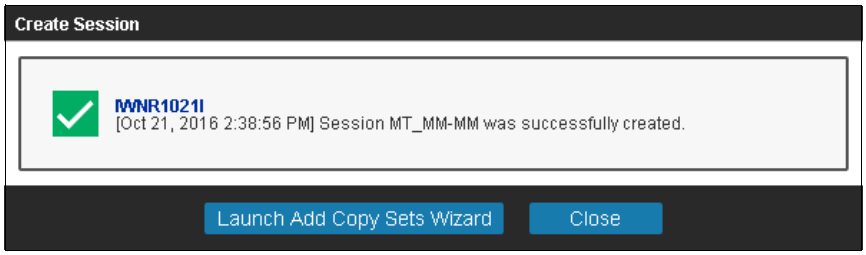

Figure 5-48 Add Copy Sets Wizard

3. We specify (one by one) the H1 (Figure 5-49), H2 (Figure 5-50) and H3 (Figure 5-51 on page 117) volumes.

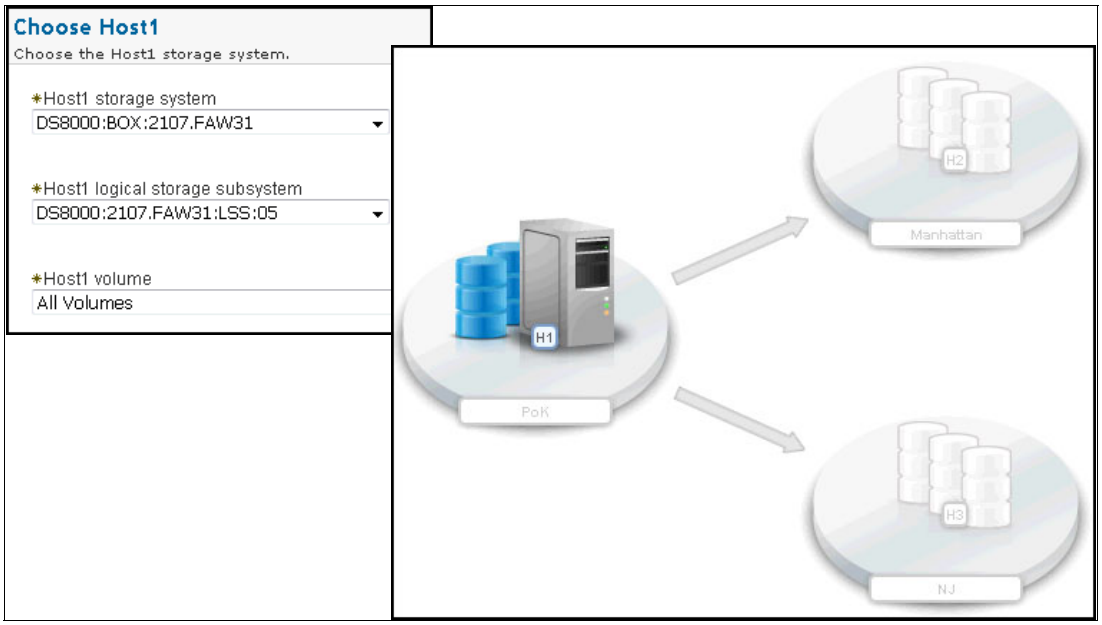

Figure 5-49 H1 (PoK) volumes

In our case, we replicate all volumes from LSS 05 in all three sites.

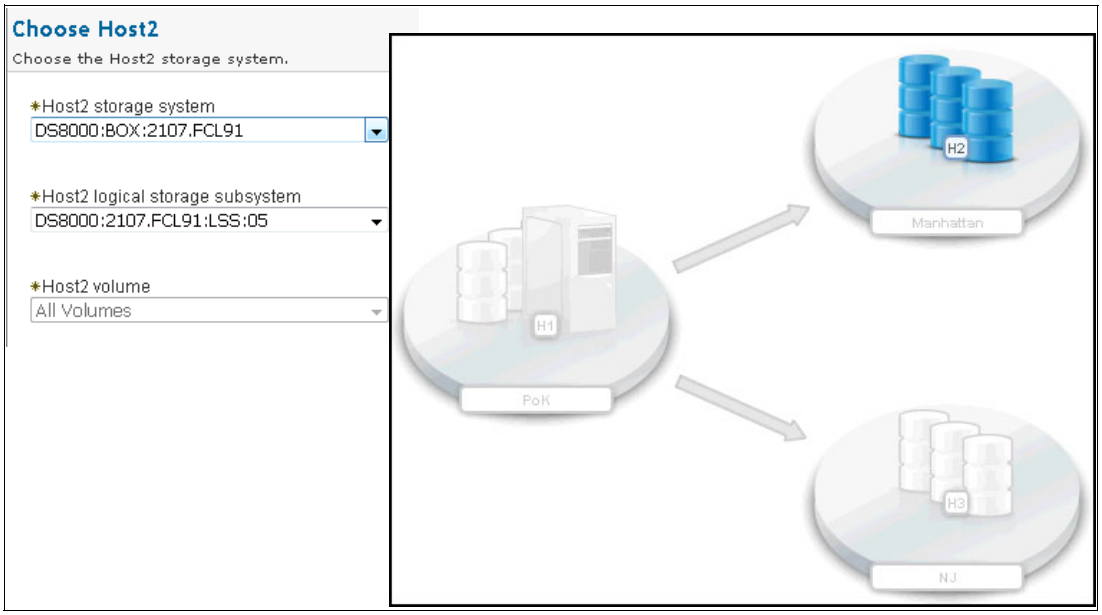

Figure 5-50 H2 (Manhattan) volumes

Optionally you can use a `.csv` format file to import all of the copy sets.

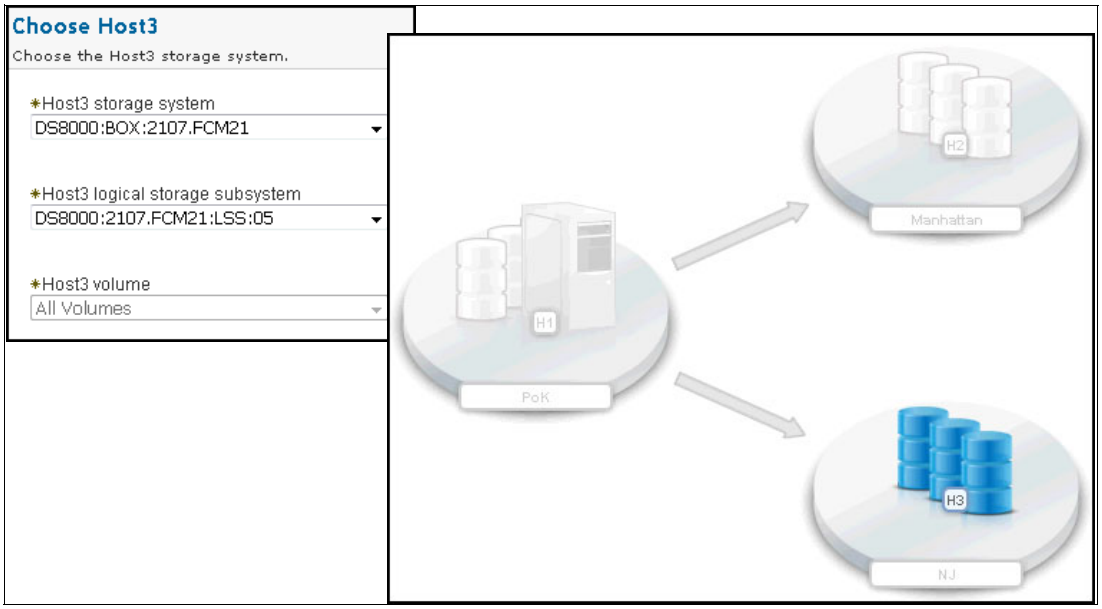

Figure 5-51 H3 (NJ) volumes

4. During the matching process, Copy Services Manager detects that the defined volumes are already part of an existing session.

   In this case, we chose to ignore the warning message (see Figure 5-52) because we do not start this session while the old one is still active.

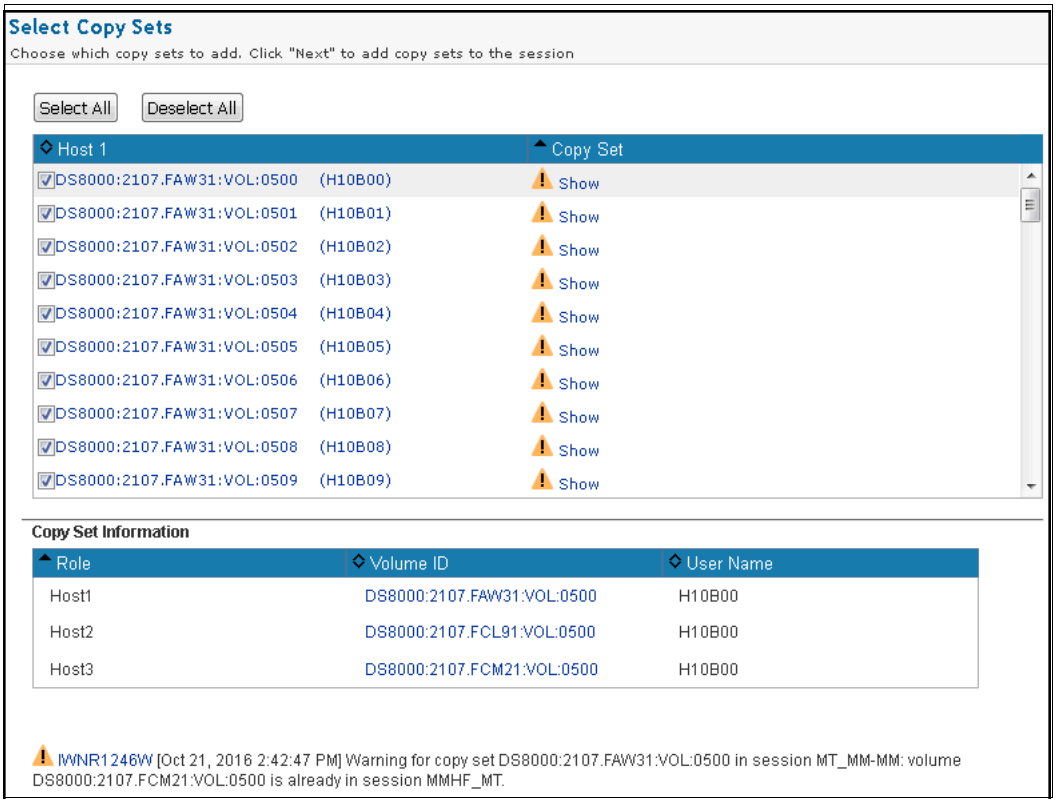

Figure 5-52 Warning message

5. After the copy sets have been added, the new Multi-Target Metro Mirror - Metro Mirror session is ready to be used. Figure 5-53 shows the new session as *Inactive/Defined*.

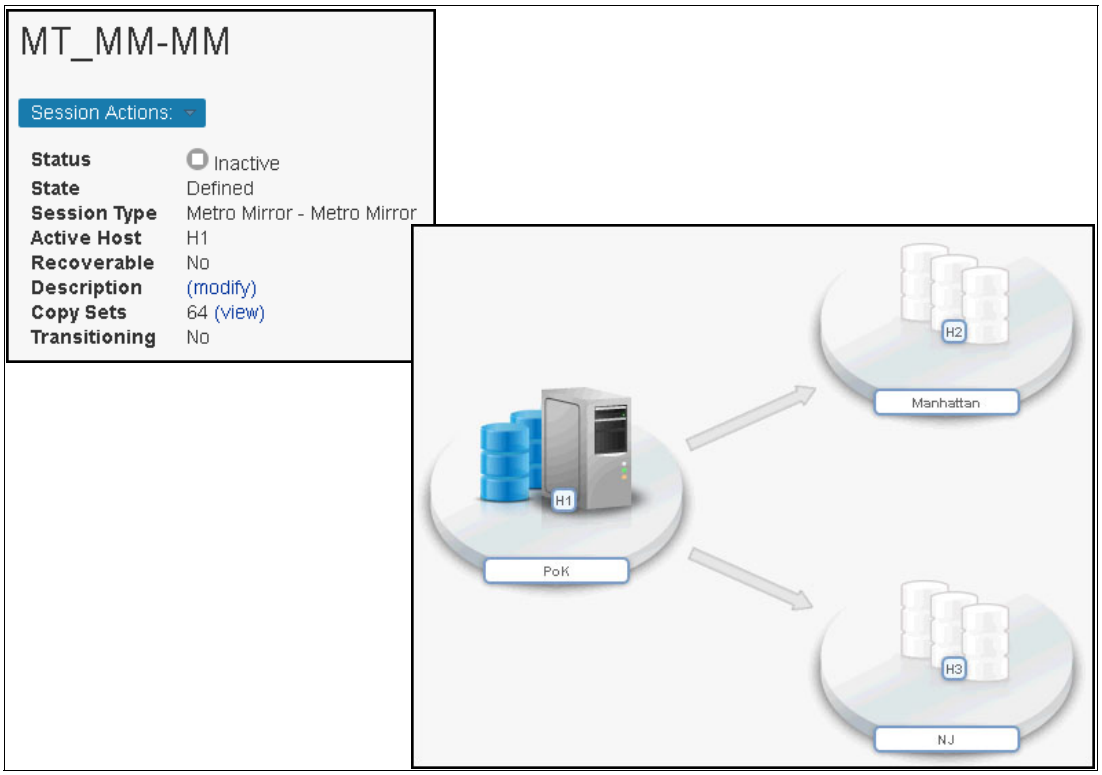

Figure 5-53 New MT_MM-MM session

6. Next, we associate the sysplex information to the session. We use the **Session Action** menu and select **Properties**, as shown in Figure 5-54.

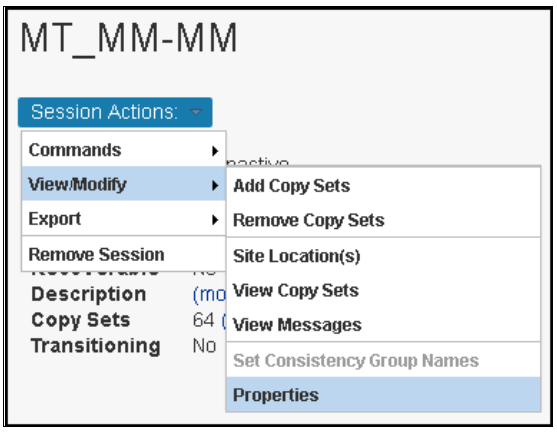

Figure 5-54 Session properties

7. In the **Session Options** tab, we select the z/OS Management System or Sysplex, as shown in Figure 5-55.

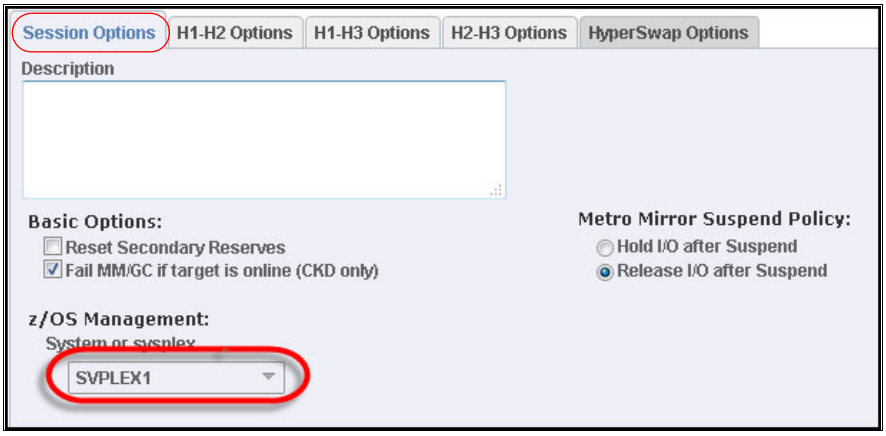

Figure 5-55 Session Options

8. For each relation, H1-H2, H1-H3, and H2-H3, we also specify that Metro Mirror management is handled by HyperSwap (see Figure 5-56).

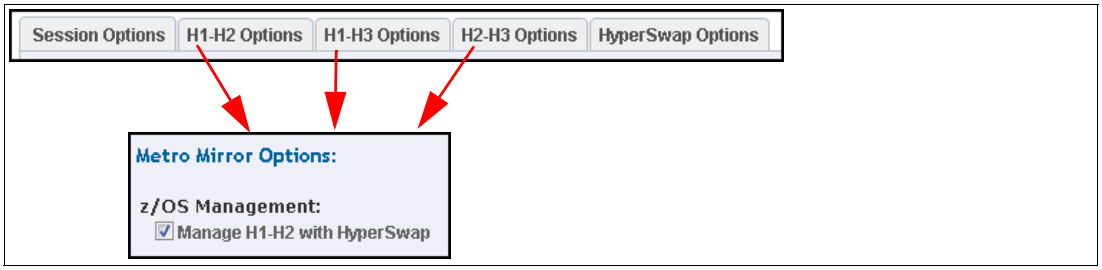

Figure 5-56 Metro Mirror Options

9. In the final configuration step on the **HyperSwap Options** tab, we have the option to let the HyperSwap Manager determine the active site, or we can specify the priorities, as shown in Figure 5-57.

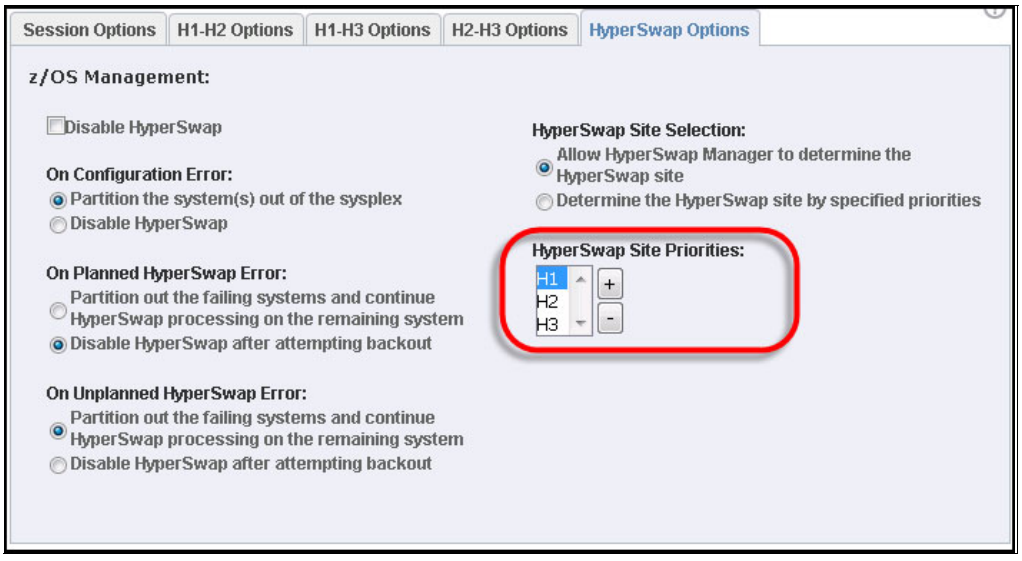

Figure 5-57 HyperSwap Site Priority

10. Now, we remove the copy sets from the active MM session. This does not stop the copy process because all Metro Mirror relations remain active on storage level.

    To remove the copy sets, we use **Session Actions** menu in the Basic HyperSwap session and select **View/Modify** → **Remove Copy Sets**, as shown in Figure 5-58.

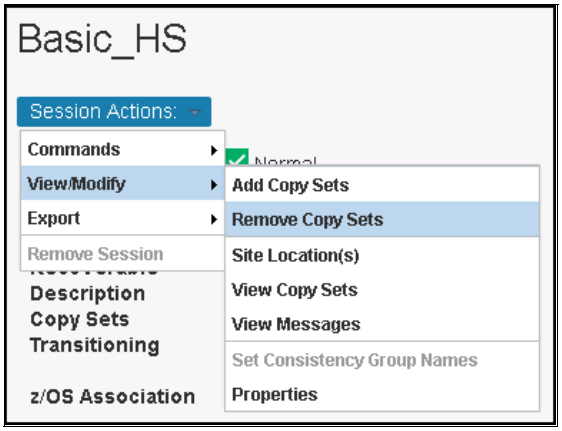

Figure 5-58 Remove Copy Sets

11. When asked, only the source volumes (H1) must be specified, as shown in Figure 5-59.

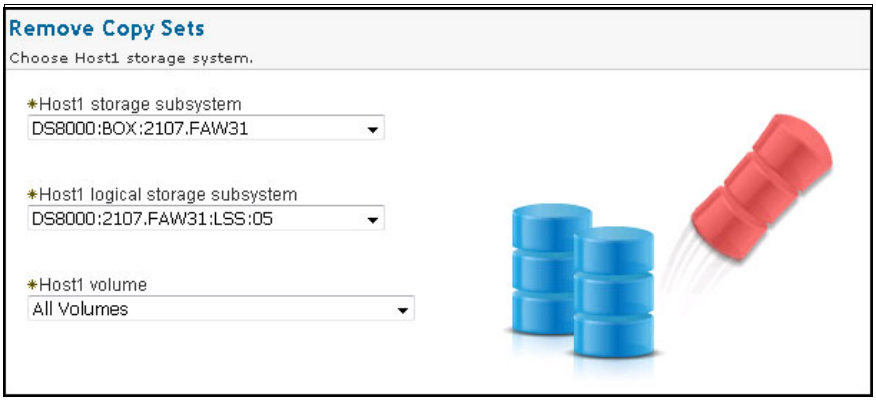

Figure 5-59 Remove Copy Sets

Important: You must choose **Yes, keep the base relationships on the storage system**. This tells CSM to leave the Global Copy active on the hardware (Figure 5-60).

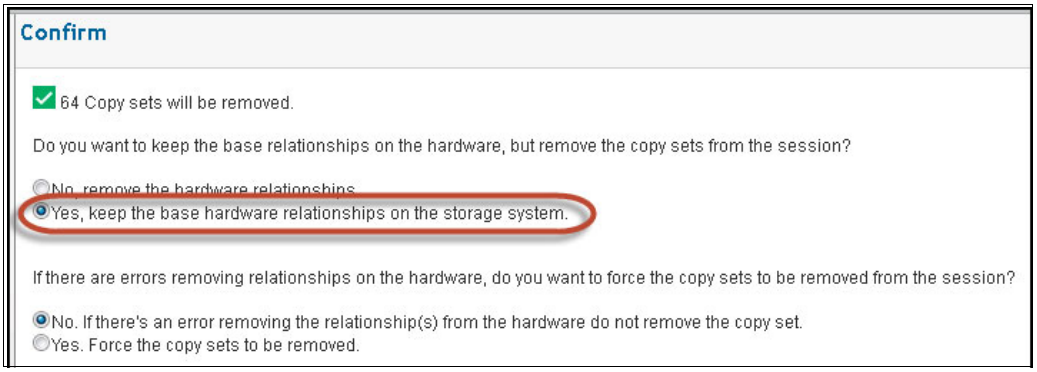

Figure 5-60 Confirm keeping base hardware relations

12. After the copy sets have been removed from the Basic_HS session, the status changes to *Inactive/Defined*, as shown in Figure 5-61.

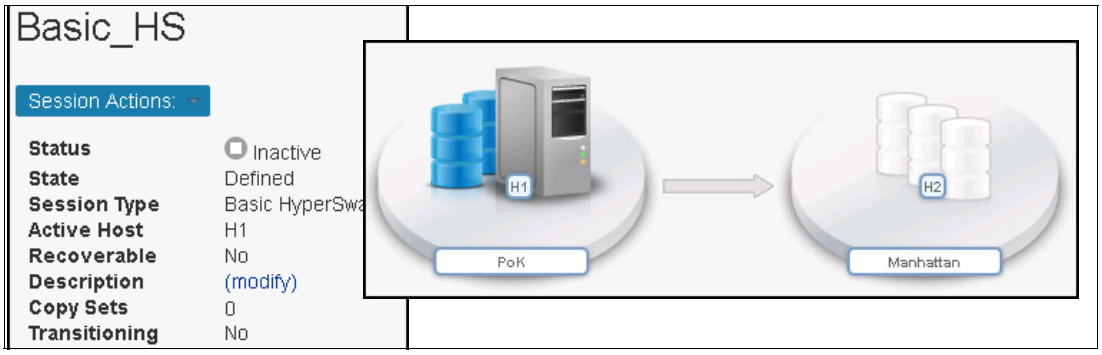

Figure 5-61 Basic HyperSwap session after removing copy sets

13. Now, we integrate the Metro Mirror relations with the new session by selecting **Start H1->H2**, as shown in Figure 5-62.

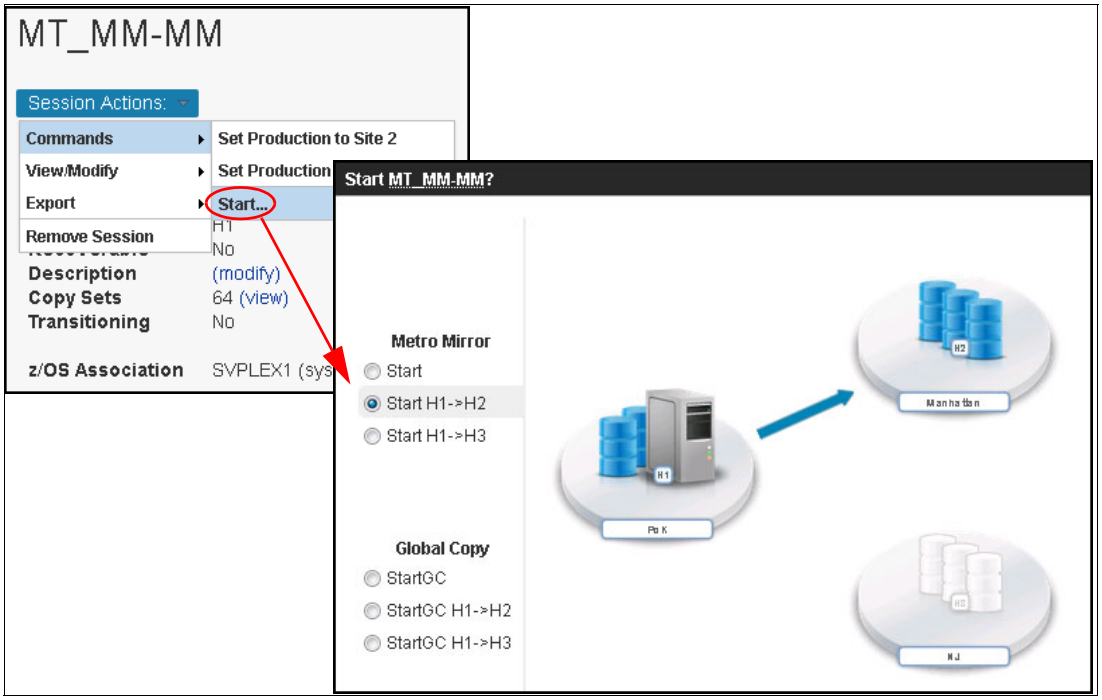

Figure 5-62 Start H1->H2

The warning message, that *data will be overwritten on Manhattan for all inactive copy sets* can be ignored this time because all copy sets were active before. Because the underlaying hardware relations (storage system) are active, the task should complete almost immediately.

In addition, the session is configured to manage the Metro Mirror relation using HyperSwap, as shown in Figure 5-63.

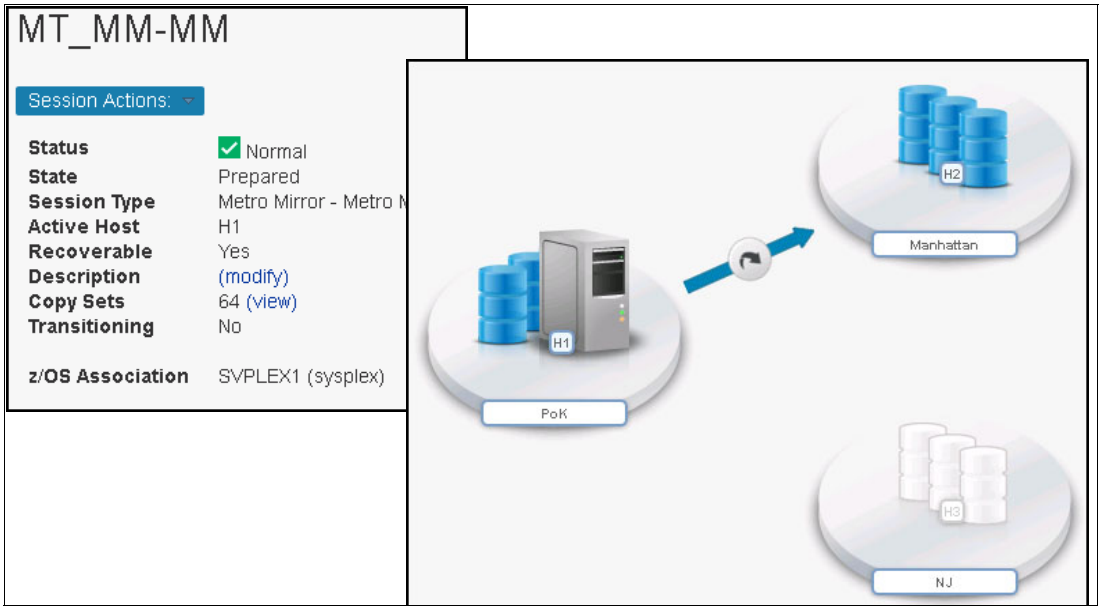

Figure 5-63 HyperSwap H1-H2

14. The next step is to replicate data to the third site, *NJ* (H3). Because we do not want any performance impact on the production site, we start an asynchronous Global Copy from H1 to H3, as shown in Figure 5-64.

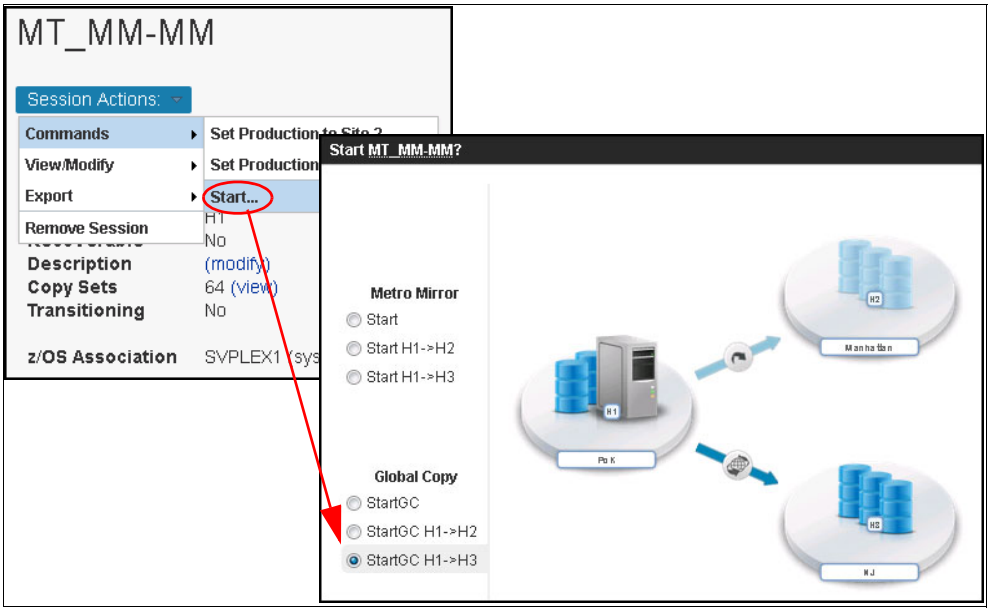

Figure 5-64 Start GC H1-H3

15. This action starts Global Copy relations on storage level. You can check in the **dscli** the new Global Copy in *Copy pending* state together with the existing Metro Mirror relations, as shown in Example 5-5.

Example 5-5 Checking relations status

```
dscli> lspprc 0500-05FF
Date/Time: 21 octombrie 2016 23:01:23 EEST IBM DSCLI Version: 7.7.21.39 DS:
IBM.2107-75FAW31
ID        State        Reason Type         SourceLSS Timeout (secs) Critical Mode
First Pass Status
==========================================================================================
0500:0500 Full Duplex  -   Metro Mirror 05   120   Disabled   Invalid
0500:0500 Copy Pending -   Global Copy  05   120   Disabled   True
0501:0501 Full Duplex  -   Metro Mirror 05   120   Disabled   Invalid
0501:0501 Copy Pending -   Global Copy  05   120   Disabled   True
0502:0502 Full Duplex  -   Metro Mirror 05   120   Disabled   Invalid
0502:0502 Copy Pending -   Global Copy  05   120   Disabled   True
0503:0503 Full Duplex  -   Metro Mirror 05   120   Disabled   Invalid
0503:0503 Copy Pending -   Global Copy  05   120   Disabled   True
0504:0504 Full Duplex  -   Metro Mirror 05   120   Disabled   Invalid
0504:0504 Copy Pending -   Global Copy  05   120   Disabled   True
0505:0505 Full Duplex  -   Metro Mirror 05   120   Disabled   Invalid
0505:0505 Copy Pending -   Global Copy  05   120   Disabled   True
0506:0506 Full Duplex  -   Metro Mirror 05   120   Disabled   Invalid
0506:0506 Copy Pending -   Global Copy  05   120   Disabled   True
........
........
053E:053E Full Duplex  -   Metro Mirror 05   120   Disabled   Invalid
053E:053E Copy Pending -   Global Copy  05   120   Disabled   True
053F:053F Full Duplex  -   Metro Mirror 05   120   Disabled   Invalid
053F:053F Copy Pending -   Global Copy  05   120   Disabled   True
dscli>
```

During the asynchronous copy, the MT_MM-MM session remains in the *Preparing* state. The warning status means that the copy sets are in *resync* state. See Figure 5-65.

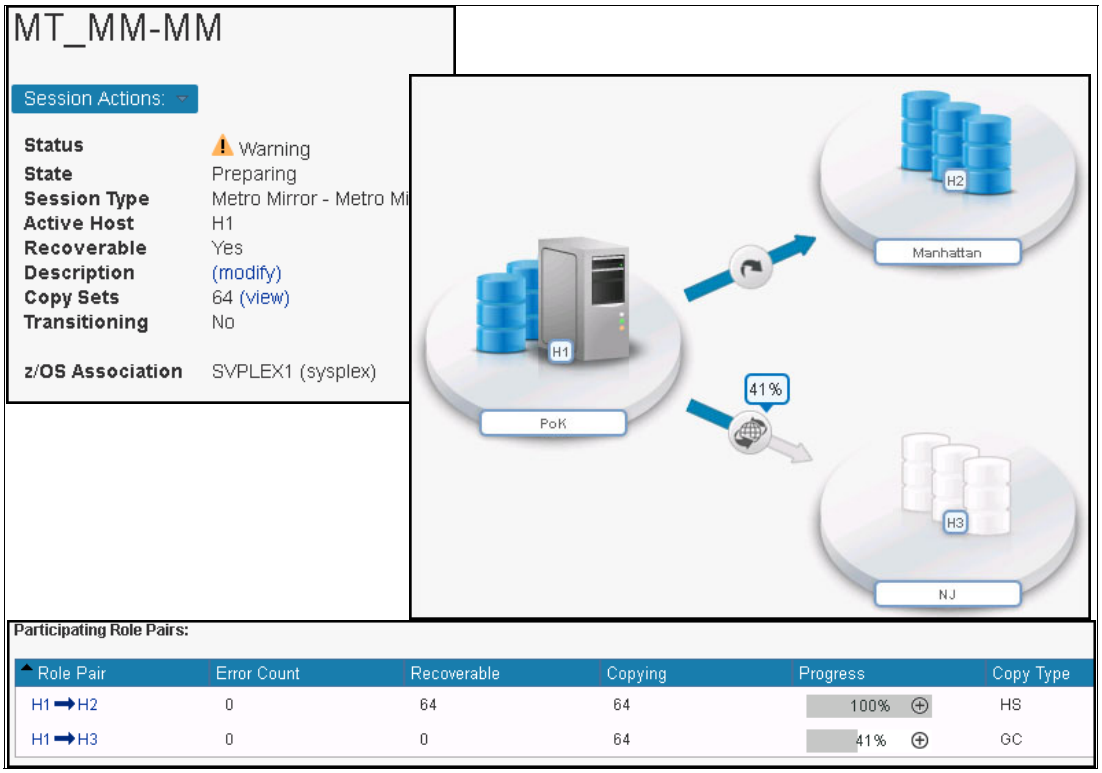

Figure 5-65 Session status during Global Copy H1-H3

16. When the GC operation reaches 100% (Out-of-Sync tracks is zero or almost zero), we can select **Start H1->H3** (Metro Mirror) with minimum performance impact, as shown in Figure 5-66.

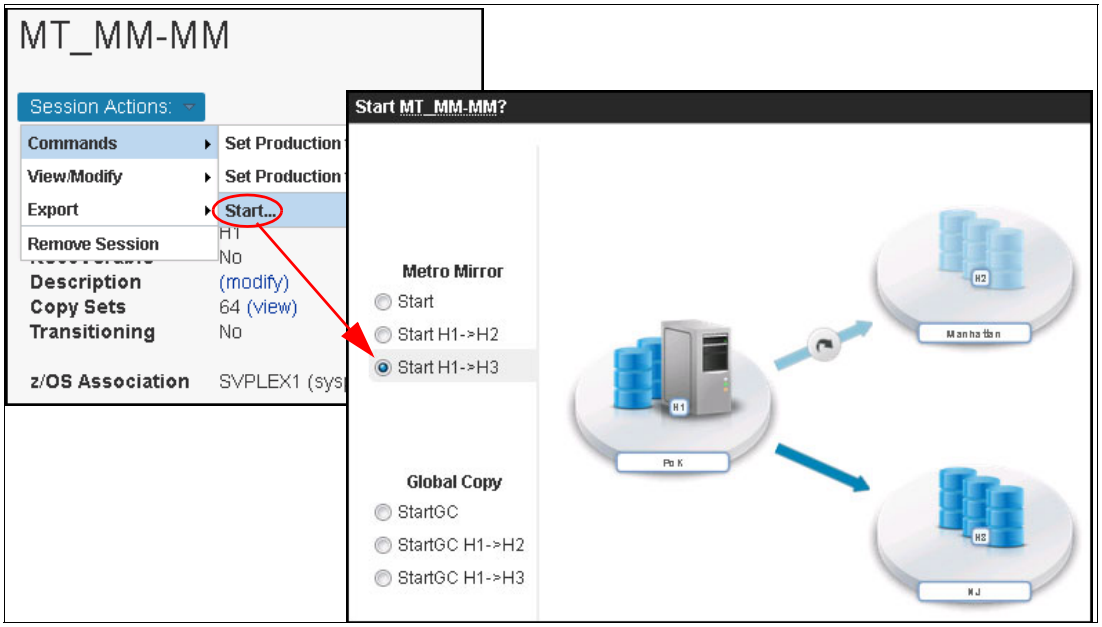

Figure 5-66 Start H1-H3

This activates the Metro Mirror on H1-H3 while also enabling it to be managed by HyperSwap manager from the parallel sysplex. See Figure 5-67.

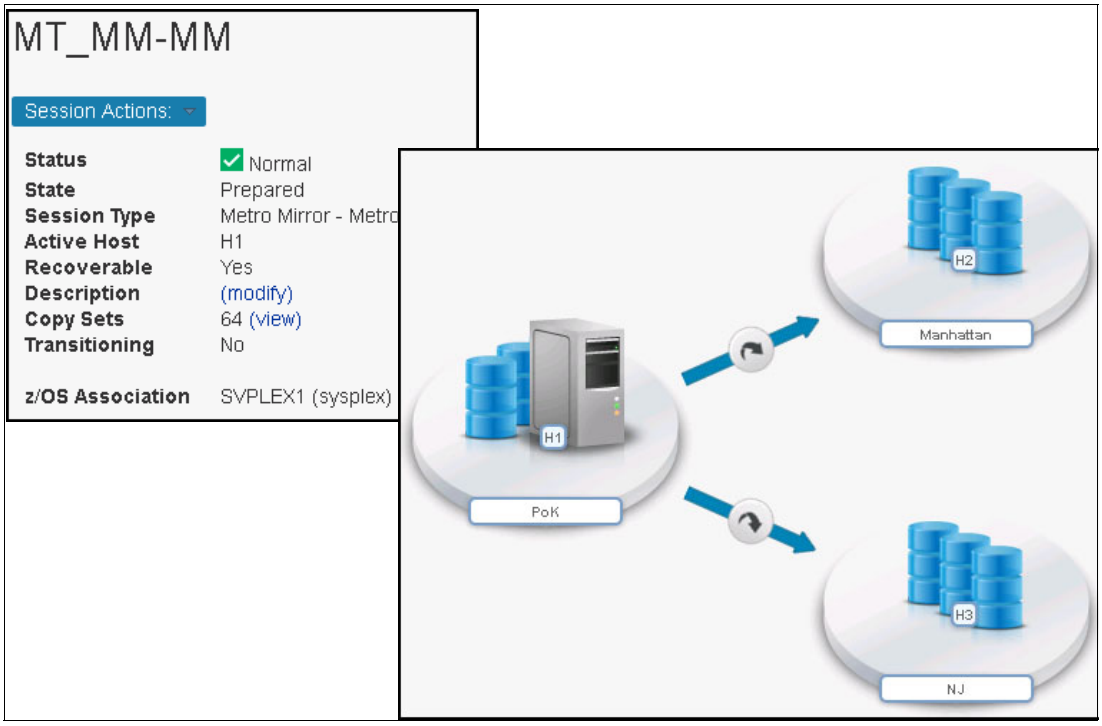

Figure 5-67 Final Status MT_MM-MM

Conclusion

The three site Multi-Target MM-MM solution provides enhanced high availability with Disaster Recovery. The three-site HyperSwap configuration enables you to chose where to move the I/O, as shown in Figure 5-68.

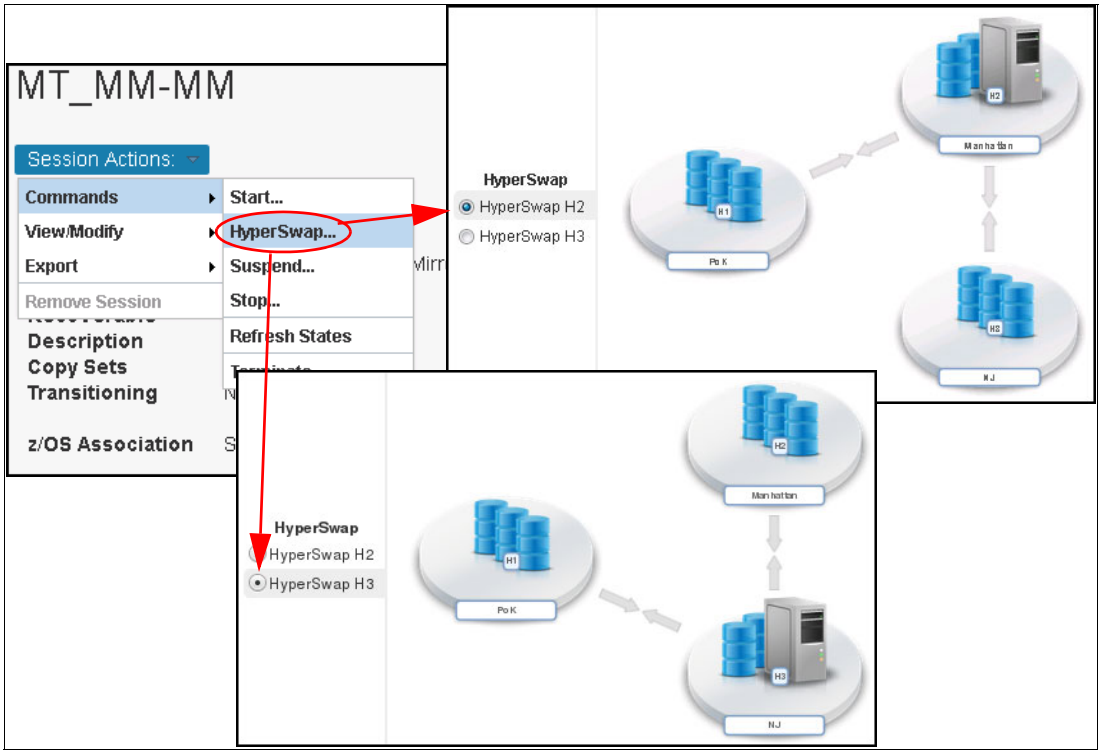

Figure 5-68 HyperSwap Options

6

Implementing IBM DS8000 Multi-Target Metro Mirror - Global Mirror environment with high availability for the CSM Server

In this chapter, we describe how to configure and test the resiliency of a Multi-Target Metro Mirror - Global Mirror solution in case of double failure that includes the primary (Active) CSM server.

The following tasks are covered in this chapter:

- Introduction
  - Scenario overview
- Configuring and testing high availability for the CSM Server
  - Configuring the Standby CSM Server on H3 HMC
  - Simulating (primary) CSM server and sysplex failure
  - Takeover on Standby CSM server
  - Recovering H1 and H2 sites and failback to H1
  - Recovering the initial CSM server

6.1 Introduction

Many businesses require their applications to be continuously available and cannot tolerate any service interruption. Loss of a disaster recovery capability is considered a severe impact to the business.

For example, if the production site fails, swapping to a Metro Mirror target allows applications to continue running. However, without another target to act as a safeguard in case of a disaster, many business applications are left unprotected.

6.1.1 Scenario overview

In this scenario, we test the resiliency of a Multi-Target Metro Mirror - Global Mirror solution in case of double failure.

The current environment has a Metro Mirror relation between *PoK* (H1) and *Manhattan* (H2) managed in a parallel sysplex configuration. Also, the *PoK* (H1) site has another target in *NJ* (H3). The replication to *NJ* (H1-H3) is asynchronous (GM).

Due to business considerations (for example, the Z administrator should administer the CSM server), the Copy Services Manager is installed on z/OS so the Host Connection for the Sysplex is ZOS_NATIVE (Java Native Interface - JNI).

In case of a failure on H1-H2, the HyperSwap manages the disk I/O operations, but a question arises: What happens in case of double failure?

Let's assume we have a planned outage on *Manhattan* (H2) and the H1-H2 is suspended. An unexpected failure occurs in *PoK* (H1), and the parallel sysplex becomes unavailable. In this situation, the data is still available in *NJ* (H3) and we can fail over to this site and activate the applications. The main issue here is that the Copy Services Manager server is also unavailable because the parallel sysplex is down and we don't have a tool to issue the storage failover.

There are two options in this situation:

a. Use storage interface (**`dscli`**) to perform the failover actions:

   This requires complex operations and, in case of multiple sessions and relations, failover could be time-consuming, making recovery time unacceptable.

b. Alternatively, we can use a Standby CSM server on the third site:

   This is easy to manage because CSM can take over and resume replication management on the Standby CSM server. This gives access to all CSM sessions, and you can issue the failover actions.

   Moreover, the standby CSM server can be activated on the DS8000 HMC[1] in *NJ* (H3).

For the scenario described in this section, we have performed the following steps:

1. Configuring the Standby CSM Server on H3 (DS8000 HMC)
2. Double failure occurs (simulation)
3. Standby CSM server takes over
4. Use the CSM (now active) to issue failover on H3
5. Recover H1 and H2 sites and failback production in H1

[1] Check for DS8880 HMC version (firmware Release 8.1 or later) for CSM support.

6.2 Configuring and testing high availability for the CSM Server

Note: Starting with DS8000 Version 8.1, Copy Services Manager comes preinstalled on the Hardware Management Console (HMC). Therefore, you can enable the Copy Services Manager software that is already on the hardware system. Doing so results in shorter CSM environment setup time, and eliminates the need to maintain a separate server for Copy Services functions.

The HMC embedded version is the Basic Edition of the CSM server, which allows you to connect only to an LDAP repository for remote authentication. You can obtain (pending valid CSM entitlement) the full product license key from either IBM Passport Advantage® or IBM Data storage feature activation.

6.2.1 Configuring the Standby CSM Server on H3 HMC

The following steps must be performed to prepare the embedded CSM in DS8000 HMC:

1. Connect by browser using HMC management IP or fully qualified domain name (FQDN), which are case-sensitive:

   `https://<HMC_IP>/CSM/`

   Attention:

   - There is no need to specify the 9559 port: The HMC uses a proxy so that CSM uses the same certificate. This is because certain deployments may require changing the DS8000 HMC certificate.
   - If you do *not* use the "/" character after "CSM" in the URL, the browser returns a Console Internal Error (HTTP status code: 404).

2. Log in using the default user and password (`csmadmin`/`passw0rd`). After first login, the password for the `csmadmin` user can be changed. Move the mouse over the `csmadmin` user on the upper-right side of the GUI, and then click **Edit Password**, as shown in Figure 6-1.

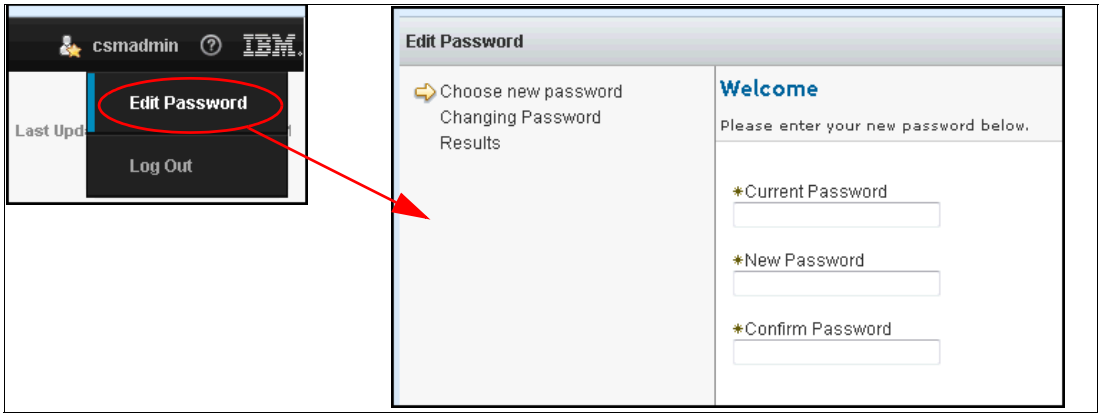

Figure 6-1 Edit the password

In this step, we upgrade the CSM server, applying the license file:

1. From the CSM GUI, click the **Update Licenses** menu and browse for the License file. Use your license key (see “Configuring the Standby CSM Server on H3 HMC” on page 129). Optionally you can use a Try-and-Buy key, as shown in Figure 6-2.

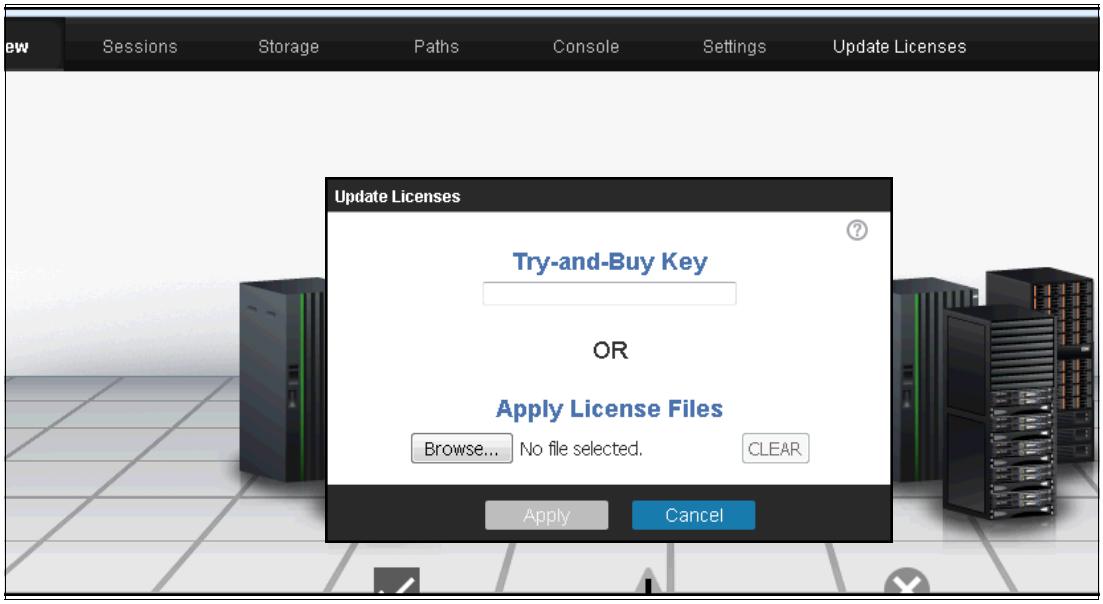

Figure 6-2 Updating CSM Licenses

2. Inform the two CSM servers of each other and their roles:
   a. Connect to the active CSM server and define the Standby CSM server. From the GUI, go to the **Settings** menu and click **Management Servers**. In the management section, click **Define Standby**.
   b. Connect to the designated Standby CSM server, and define it as a standby server:
      i. From the GUI go to **Settings** menu and click **Management Servers**.
      ii. In management section, click **Set this Server as Standby**.

      In both cases, the new window asks for the partner CSM server. See Figure 6-3 on page 131.

> **Important:** Defining a server as *Standby* wipes any existing configuration on the designated standby CSM server. As a precaution, when you run this action from the Active CSM server, you are asked for the *username* and *password* of the standby CSM server.

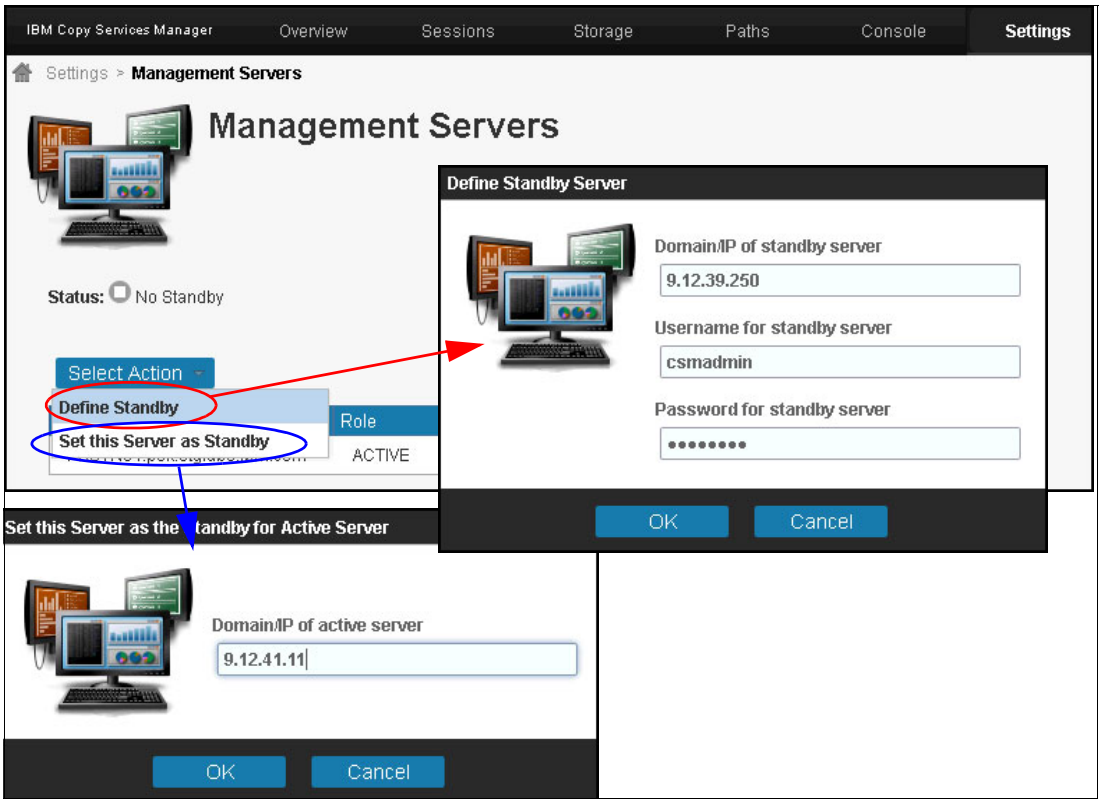

Figure 6-3 Defining the standby CSM Server

3. The warning message informs about the overwrite of any existing configuration on the designated standby CSM server. See Figure 6-4.

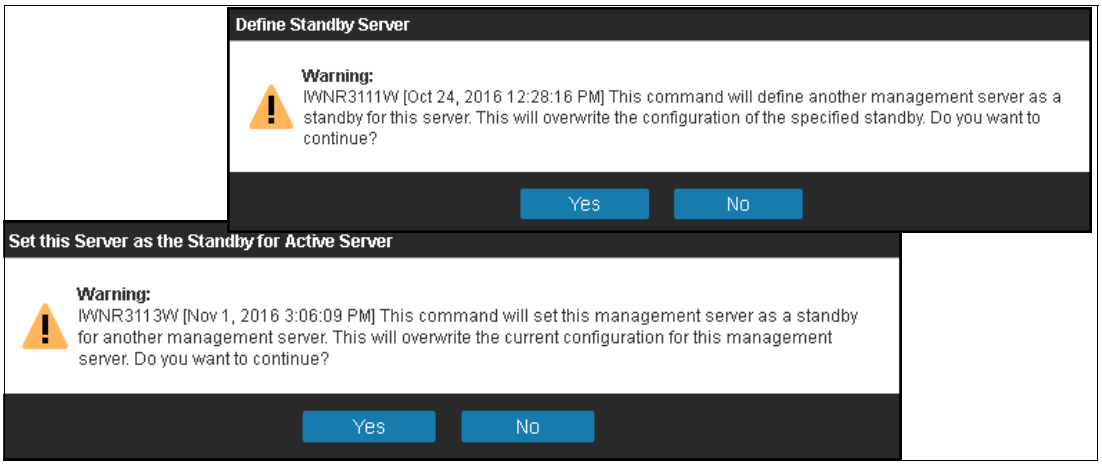

Figure 6-4 Overwrite warning for the designated Standby CSM server

4. After clicking **Yes**, the two servers initiate the database synchronization from the active CSM server to the Standby CSM server. The *Management Server* status reflects the process as `Synchronization Pending` (see Figure 6-5 on page 132).

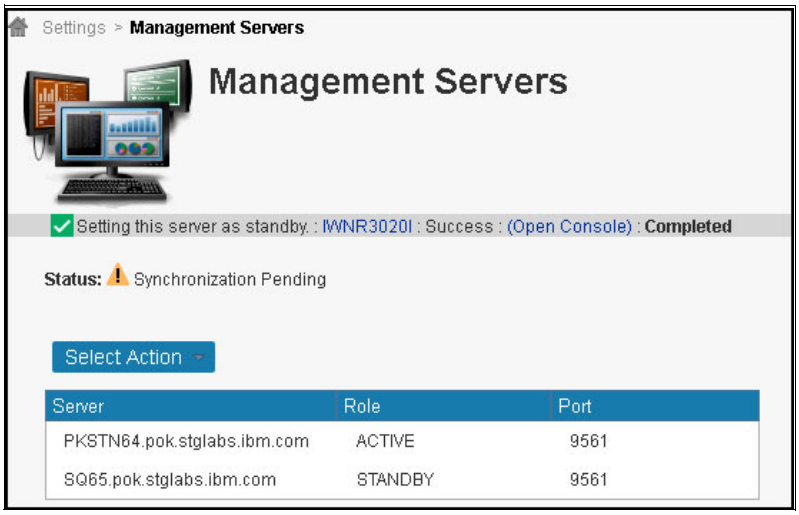

Figure 6-5 Synchronization Pending

5. When the synchronization is complete, the Standby server is ready to take over.

6.2.2 Simulating (primary) CSM server and sysplex failure

The DR solution used for testing is a Multi-Target Metro Mirror - Global Mirror managed partially (on the synchronous replication) by a sysplex. The production is running in *PoK* (H1) is in a Metro Mirror relation with *Manhattan* (H2) for high availability.

Additionally, the data is asynchronously replicated from *PoK* (H1) to the third site, *NJ* (H3). Copy Services Manager (MT_MM-GM) session configured for this solution is shown in Figure 6-6.

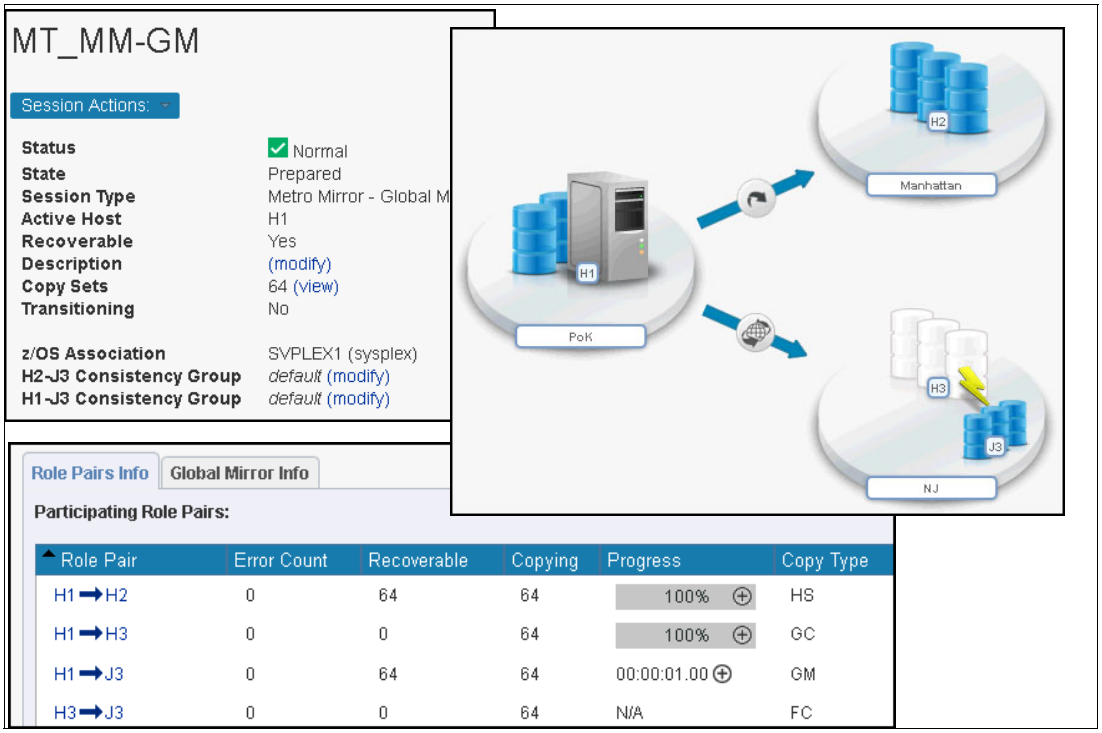

Figure 6-6 Multi-Target Metro Mirror - Global Mirror session in CSM

Complete the following steps:

1. The first (simulated) failure is a planed outage for DS8000 maintenance in *Manhattan* (H2) site. This disables the HyperSwap, but the production continues without impact in *PoK* (H1). In our case, we bring offline all volumes in H2. We can see the message from z/OS after the first failure in Example 6-1.

Example 6-1 HyperSwap disabled

```
RESPONSE=N64
IOSHM0303I HyperSwap Status 623
Number of configurations: 1
Replication Session: MT_MM-GM____H1H2
Socket Port: 14000
HyperSwap disabled:
One or more members unable to verify PPRC secondary device connectivity
N64:
Member unable to verify PPRC secondary device connectivity
N65:
Member unable to verify PPRC secondary device connectivity
Swap Highest Priority: No
Disallow Non-MultiTarget System: No
New member configuration load failed: Partition
Planned swap recovery: Disable
Unplanned swap recovery: Partition
FreezeAll: Yes
Stop: No
```

2. Copy Services Manager detects this status and marks the H1-H2 MM relation with error status *Severe*, as shown in Figure 6-7.

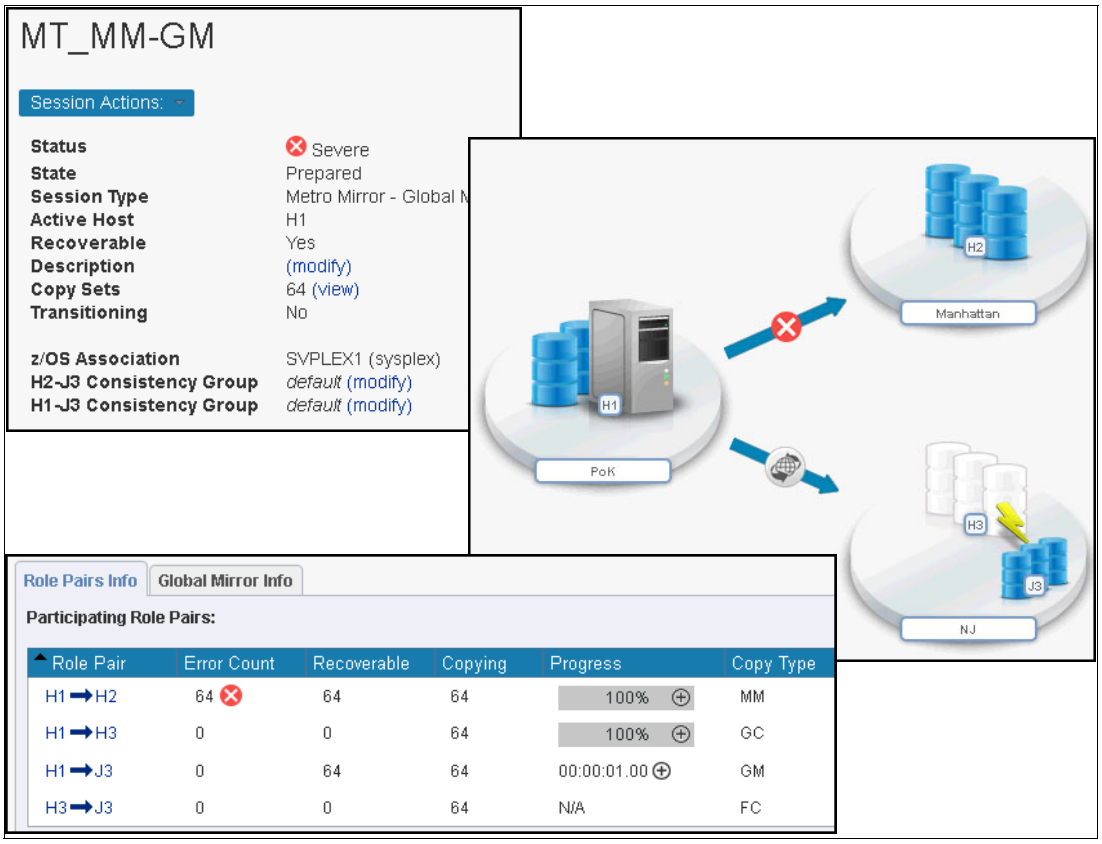

Figure 6-7 HyperSwap session error

3. The “planned outage” means that the H2 volumes are unavailable, so we suspend the H1-H2 relation (see Figure 6-8).

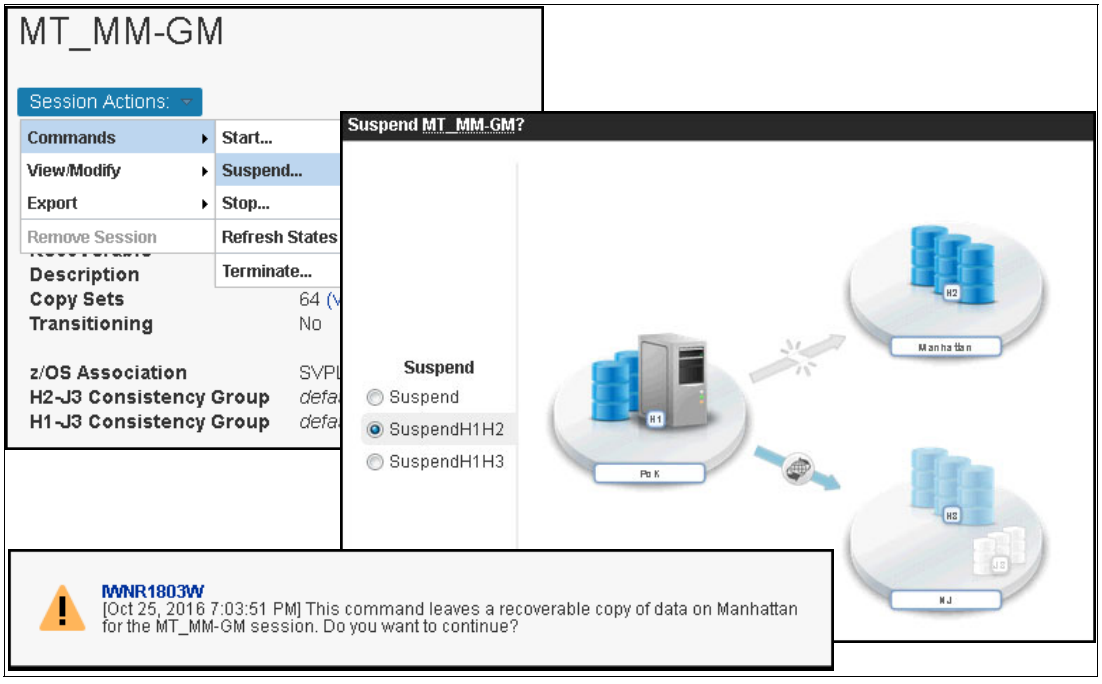

Figure 6-8 Suspend H1-H2

This renders the Session **State** in `Suspended (Partial)`, as shown in Figure 6-9.

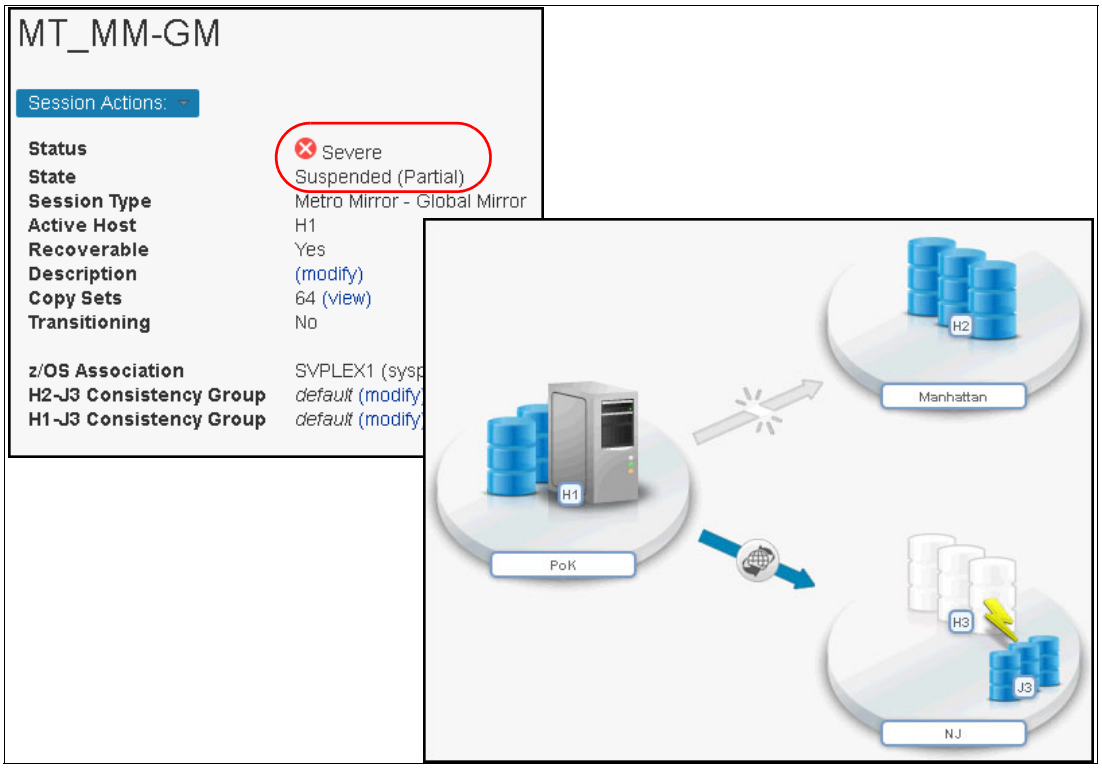

Figure 6-9 Session Suspended (Partial)

4. The second failure brings down the *PoK* (H1) site. To simulate this failure, the volumes in H1 are set offline in z/OS. This causes the parallel sysplex to go down, and, as a direct consequence, the CSM server becomes unavailable (Active CSM server running in z/OS on one of the parallel sysplex members). The CSM GUI is unreachable, as shown in Figure 6-10.

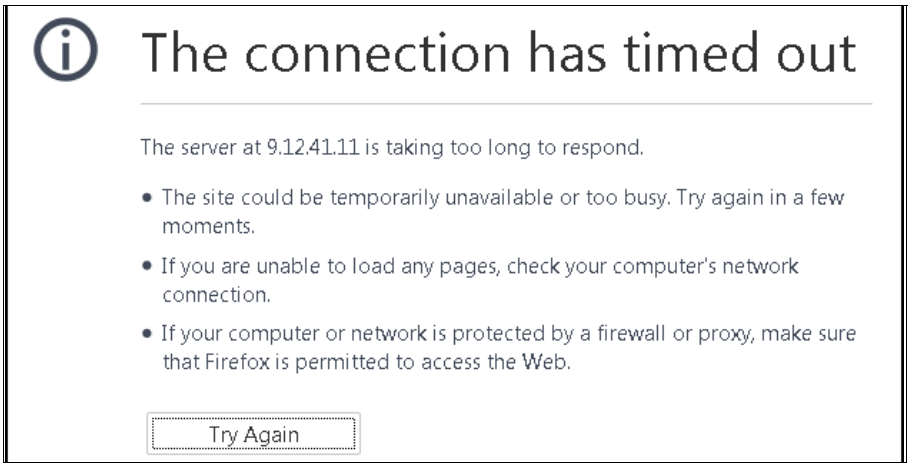

Figure 6-10 CSM GUI time-out

Also, on the storage level, the Global Mirror relations become suspended with consistency in NJ site. This can be verified with **dscli**, as shown in Example 6-2.

Example 6-2 Checking relations' status (dscli)

```
dscli> lspprc -type gcp 0500-05FF
Date/Time: 26 octombrie 2016 03:44:41 EEST IBM DSCLI Version: 7.7.21.39 DS:
IBM.2107-75FAW31
ID        State     Reason      Type        SourceLSS Timeout (secs) Critical Mode
First Pass Status
================================================================================
0500:0500 Suspended  Host Source  Global Copy 05  120  Disabled  True
0500:0500 Suspended  Host Source  Global Copy 05  120  Disabled  True
0501:0501 Suspended  Host Source  Global Copy 05  120  Disabled  True
0501:0501 Suspended  Host Source  Global Copy 05  120  Disabled  True
0502:0502 Suspended  Host Source  Global Copy 05  120  Disabled  True
0502:0502 Suspended  Host Source  Global Copy 05  120  Disabled  True
0503:0503 Suspended  Host Source  Global Copy 05  120  Disabled  True
0503:0503 Suspended  Host Source  Global Copy 05  120  Disabled  True
0504:0504 Suspended  Host Source  Global Copy 05  120  Disabled  True
0504:0504 Suspended  Host Source  Global Copy 05  120  Disabled  True
0505:0505 Suspended  Host Source  Global Copy 05  120  Disabled  True
0505:0505 Suspended  Host Source  Global Copy 05  120  Disabled  True
0506:0506 Suspended  Host Source  Global Copy 05  120  Disabled  True
0506:0506 Suspended  Host Source  Global Copy 05  120  Disabled  True
.........
.........
053E:053E Suspended  Host Source  Global Copy 05  120  Disabled  True
053E:053E Suspended  Host Source  Global Copy 05  120  Disabled  True
053F:053F Suspended  Host Source  Global Copy 05  120  Disabled  True
053F:053F Suspended  Host Source  Global Copy 05  120  Disabled  True
dscli>
```

Note: The output in Example 6-2 on page 136 displays two lines for each volume pair because of *two* suspended relations: H1-H2 and H1-H3.

6.2.3 Takeover on Standby CSM server

In an actual situation, we still have consistent data in the third site, NJ (H3). Because the CSM server in PoK (H1) is unavailable, we have to take over on the standby CSM server, which is located in NJ (H3):

1. From the GUI of the standby CSM server, go to **Settings** → **Management Servers** and click the **Takeover** option from **Select Action**, as shown in Figure 6-11 on page 137.

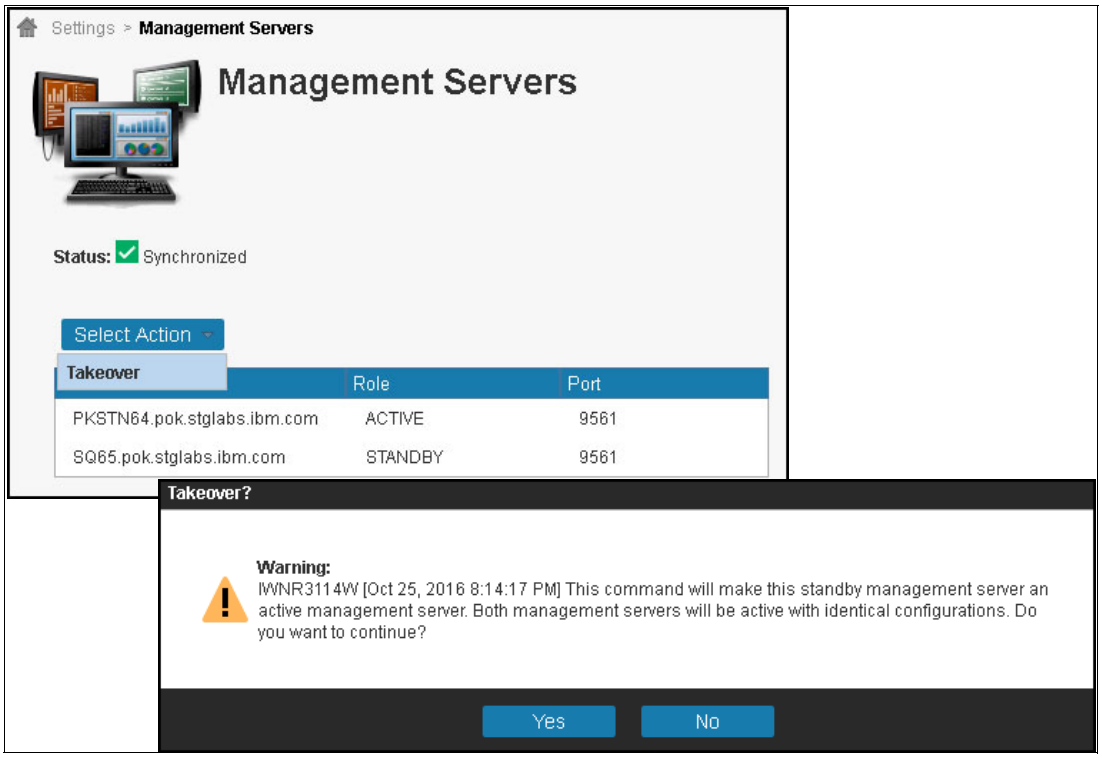

Figure 6-11 CSM server takeover

2. When the CSM server in NJ (H3) takes over, it becomes the active server (stand-alone) as shown in Figure 6-12.

Figure 6-12 CSM server in NJ site after takeover

3. We can now perform sessions operations. We can see the suspended status but also the consistent data in *NJ* (J3), as shown in Figure 6-13.

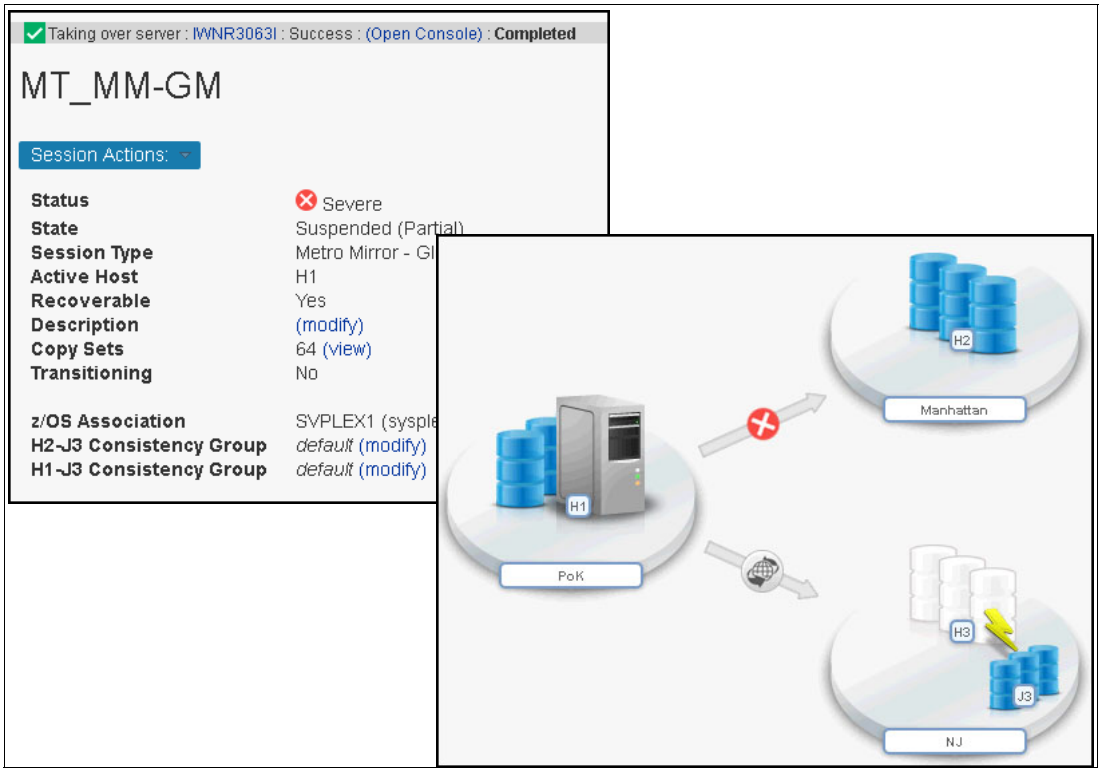

Figure 6-13 Inactive session viewed from CSM server in NJ

The gray arrow between H1 and H3 represents an inactive asynchronous copy. Because we want to move the production to H3, we have to completely suspend all relations and recover in the *NJ* site.

4. Next, we suspend H1-H3, as shown in Figure 6-14.

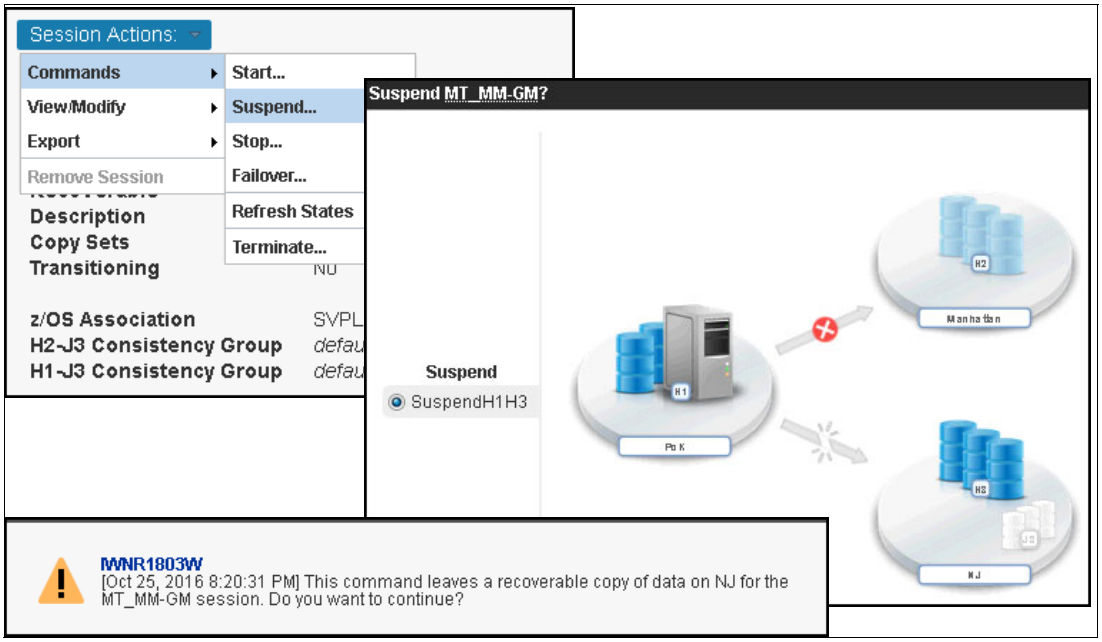

Figure 6-14 Suspend H1-H3

This creates a consistent copy of the data (volumes in H3) which is available for restarting production, as shown in Figure 6-15.

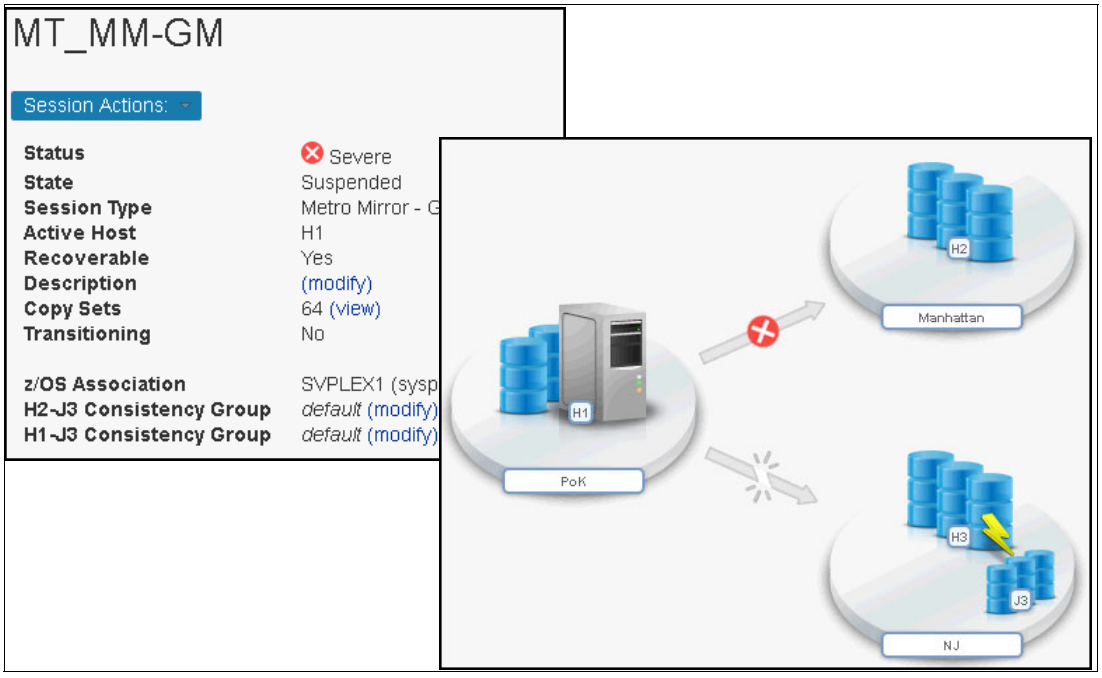

Figure 6-15 Consistent volumes H3

5. We recover in H3 by choosing **Recover** from **Session Actions** → **Commands**, as shown in Figure 6-16.

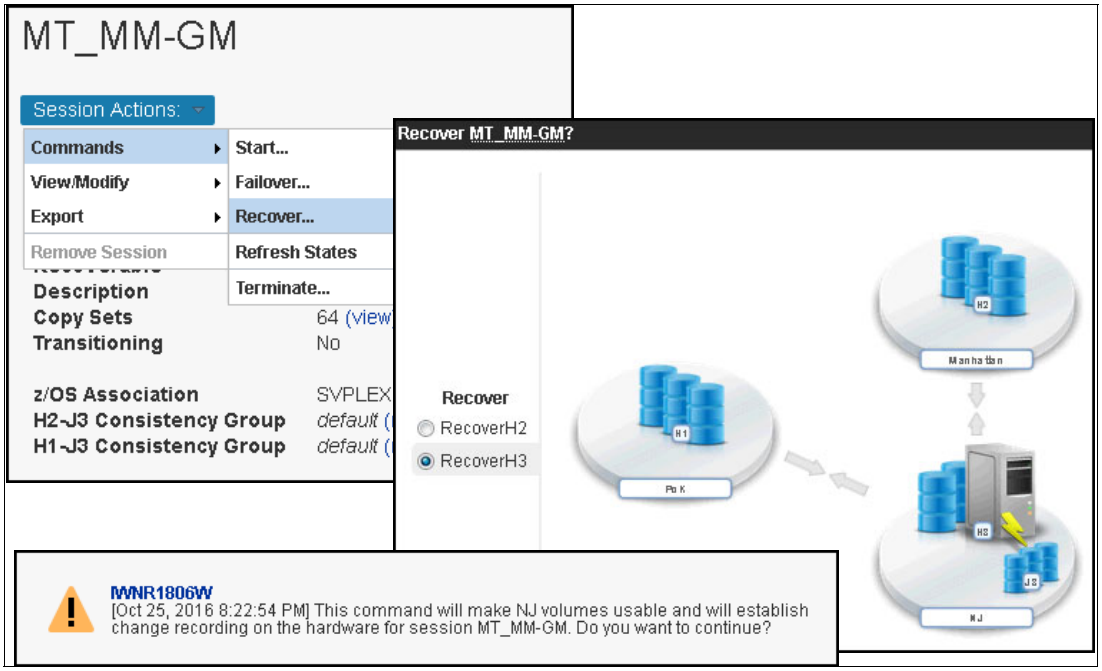

Figure 6-16 Recover

6. The H3 volumes in NJ site are now ready to start the production, but the data copy option (to H1 or H2) is disabled for the moment. CSM considers the possibility to revert to original data in H1 or H2 and prevent accidentally overwriting the data. In our case, we confirm that the production is in NJ (H3), as shown in Figure 6-17.

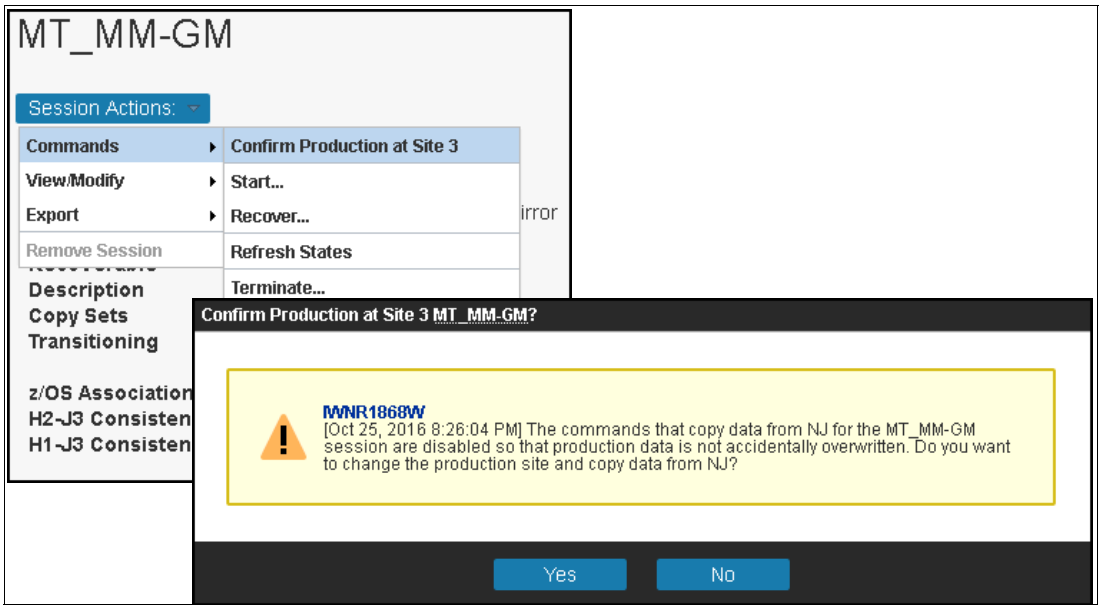

Figure 6-17 Confirm Production in Site 3 (NJ)

6.2.4 Recovering H1 and H2 sites and failback to H1

The data in *PoK* (H1) and *Manhattan* (H2) sites is now obsolete due to the production running in *NJ* (H3). To failback the production under the sysplex, we must first start the replication to these sites:

1. From **Session Actions** click **Start** and chose the option to copy data asynchronously from *NJ* (H3) to *PoK* (H1) and then to *Manhattan* (H2). This initiates the Global Copy relations, as shown in Figure 6-18.

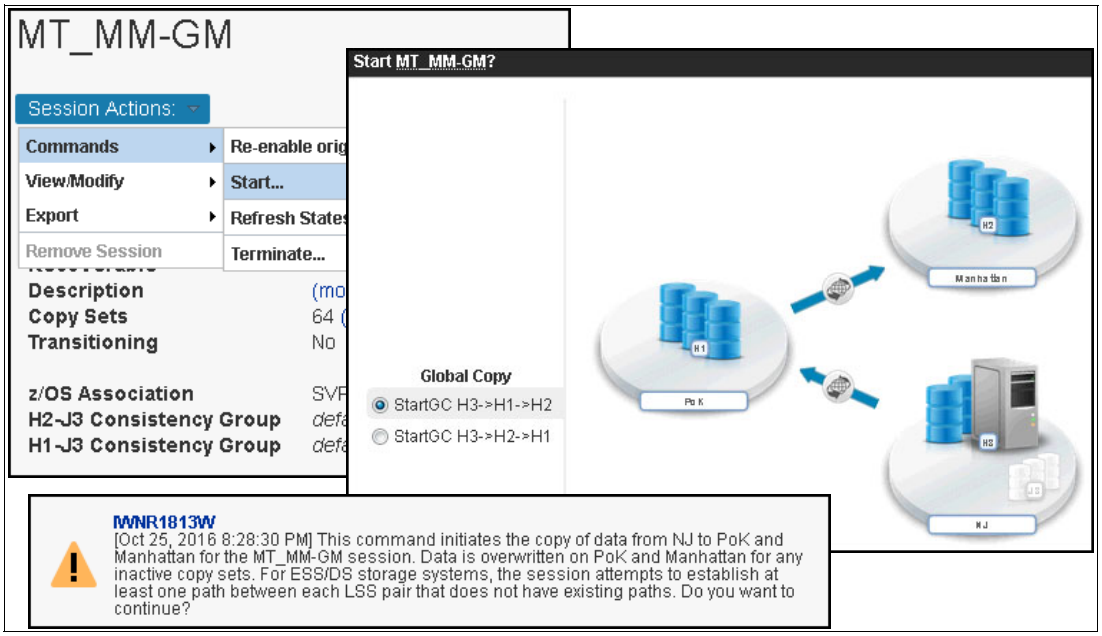

Figure 6-18 Starting GC H3-H1-H2

2. While the data is copied to H1 and to H2, the *Session state* changes to *Preparing*. We can monitor the copy progress for each site, as shown in Figure 6-19.

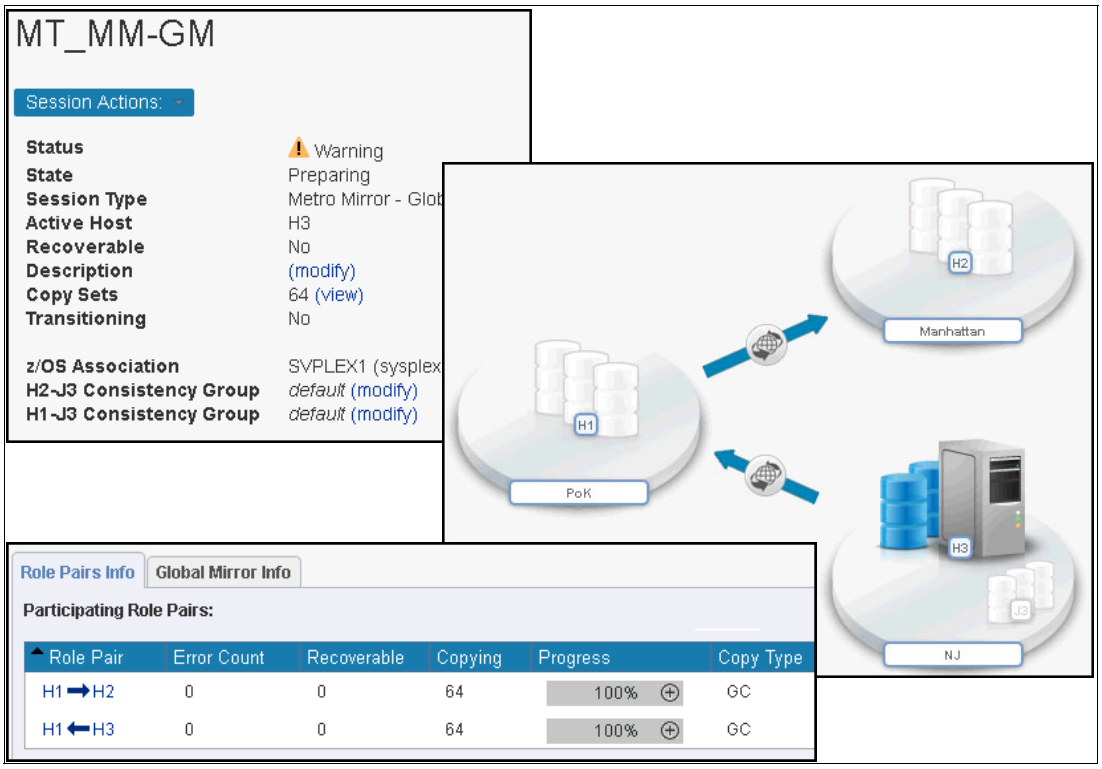

Figure 6-19 Preparing session for recovery

In certain situations, especially after a real failure of the entire site, before moving the production back you might need to perform more tests to ensure that everything is configured properly. For example, we want to have consistent data in *PoK* (H1) for tests, but also we want production to run in parallel in *NJ* (H3).

3. Because we have initiated a Global Copy to H1, the data is not consistent in H1. CSM provides two options to achieve data consistency in H1. The session has to be suspended with one of the following options:
   – *Suspend* while Global Copy is active. This changes the replication in synchronous MM and the production affects performance because of the (long) distance between H1-H3 (see Figure 6-20).

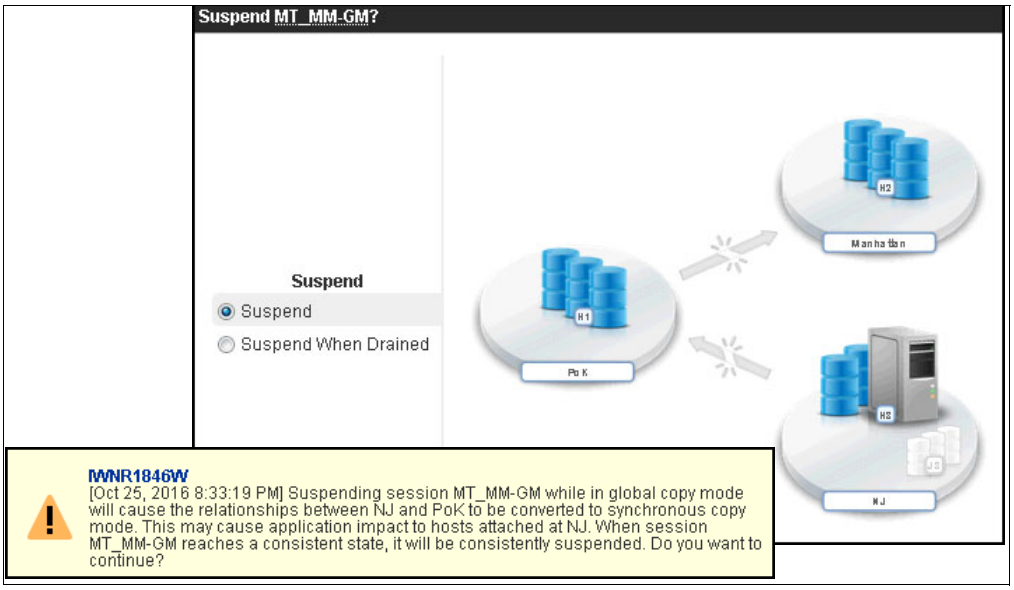

Figure 6-20 Suspend session

   – *Suspend When Drained* (Figure 6-21). This option is asking the administrator to pause the I/O on H3 and wait for synchronization. In this case H1 becomes consistent.

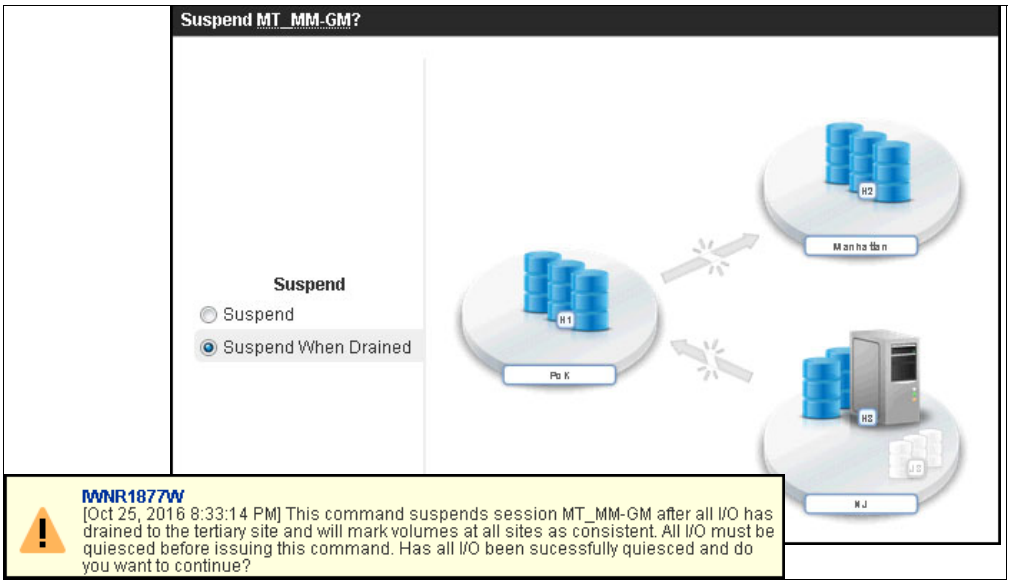

Figure 6-21 Suspend when drained

Note: Optionally, to avoid performance impact on production, Metro Mirror - Global Mirror with Site 3 Global Mirror session can be used instead.

This provides additional journal volumes in *PoK* (J1) and *Manhattan* (J2), and the Goal Copy process is followed by Global Mirror, which provides consistent data in H1 without any effect on performance.

Whichever option is chosen, consistent data in H1 becomes available and ready for testing (for example, start application with isolate IPs). After all tests are performed and we want to recover in H1, we have to start GC H3-H1-H2 again to update all of the changes performed in H3 since the relation was suspended.

When the Out-of-Sync tracks are again zero or almost zero, and the relations' progress is 100%, we can suspend the relations with option *when Drained*. Now, applications must be stopped in H3 and restarted in H1.

4. We recover in H1, as shown in Figure 6-22.

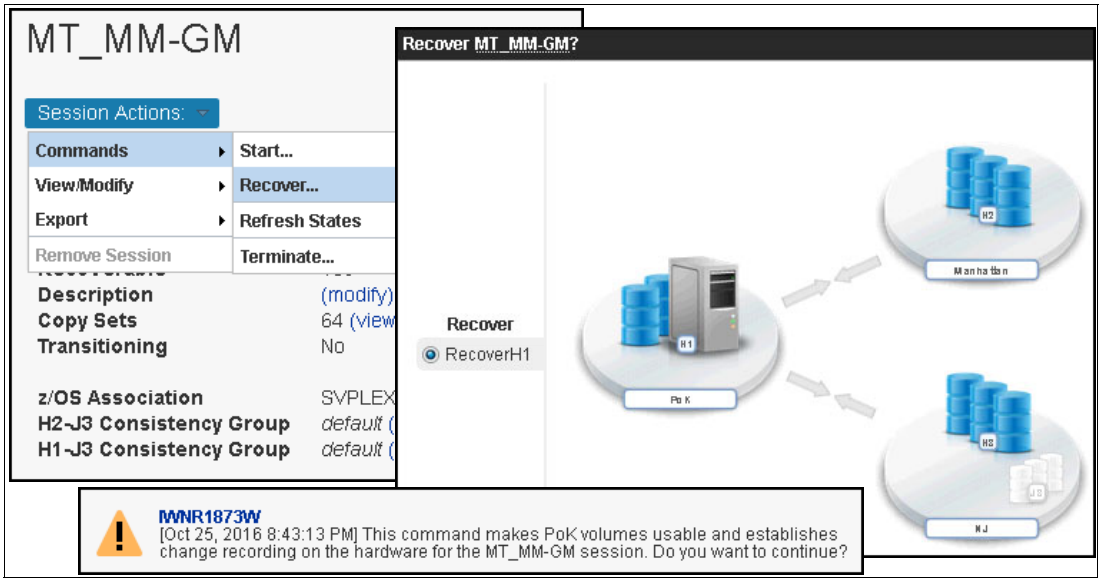

Figure 6-22 Recovering H1

5. Now the *Session status* is *Normal* and the Targets are available. This means that the volumes in *PoK* (H1) are back in production, but the data is not replicated to other sites yet, as shown in Figure 6-23.

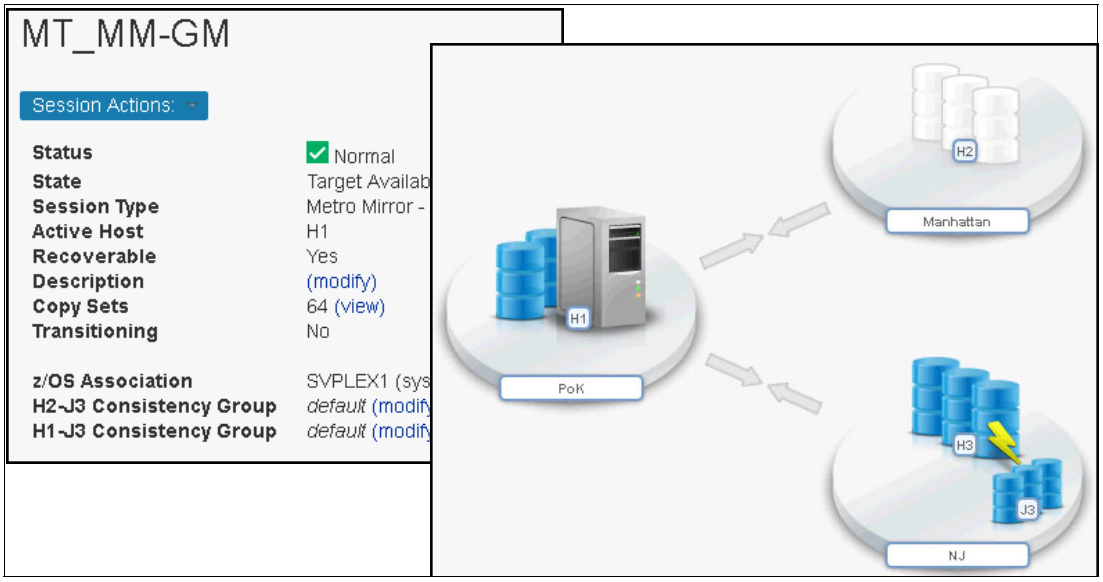

Figure 6-23 Session back to normal

6. The first step after switching back to *PoK* and the applications are up and running, is to confirm the production at Site 1 and enable the possibility to start replication to other sites, as shown in Figure 6-24.

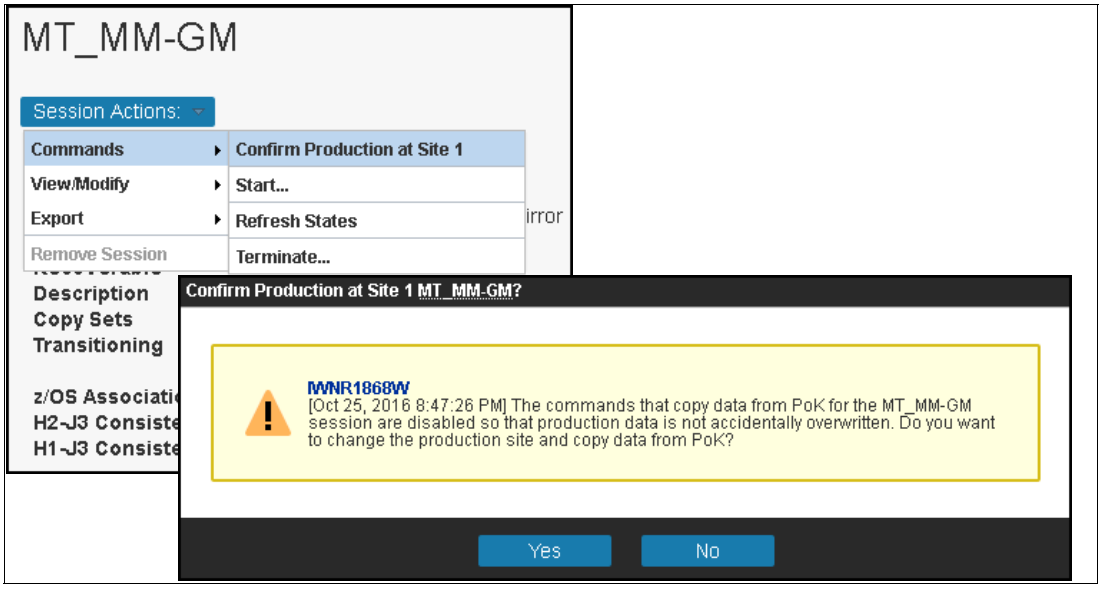

Figure 6-24 Confirm Production at Site 1

7. Next, we start replication to *NJ* (H3). We initiate asynchronous copy (Global Copy) to both H2 and H3 to replicate all changed tracks/blocks, as shown in Figure 6-25.

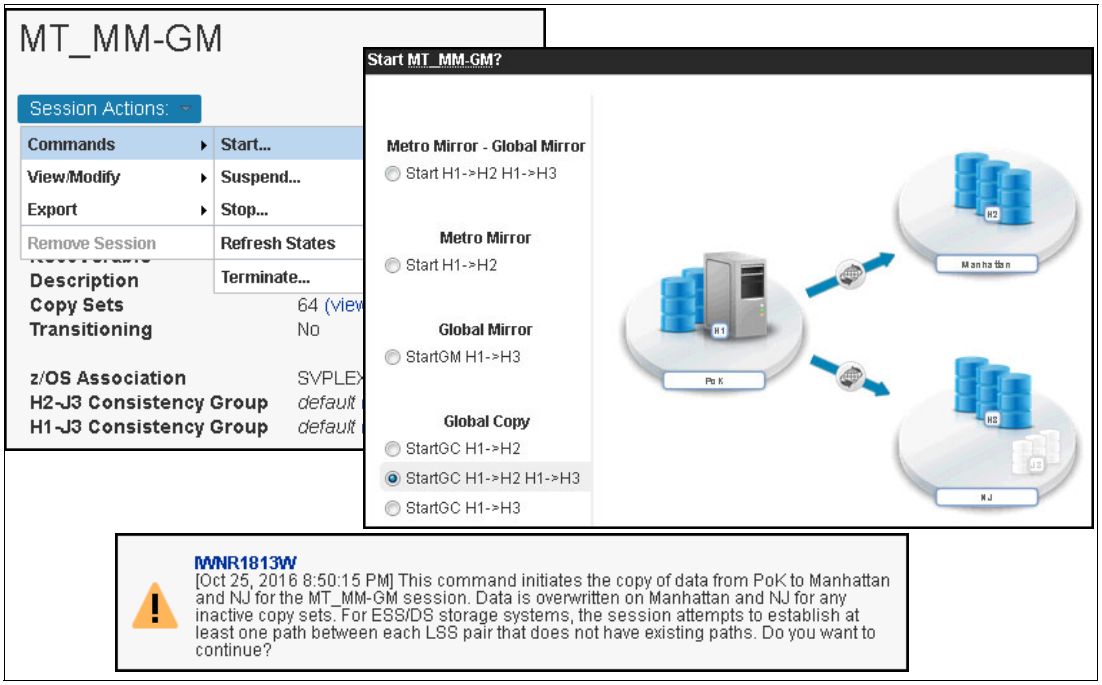

Figure 6-25 Start GC H1-H2 and H1-H3

8. We wait for copy progress to complete (100%) and we start GM on H1-H3 as shown in Figure 6-26.

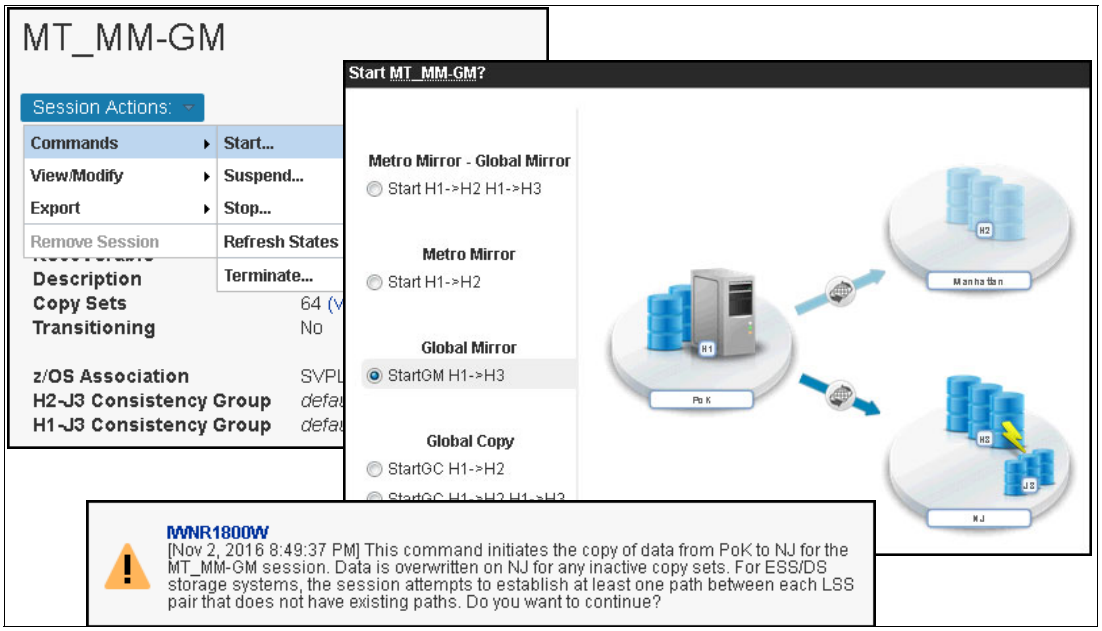

Figure 6-26 Start GM H1-H3

9. We observe the information about Global Mirror replication (such as *Data Exposure,* shown in Figure 6-27).

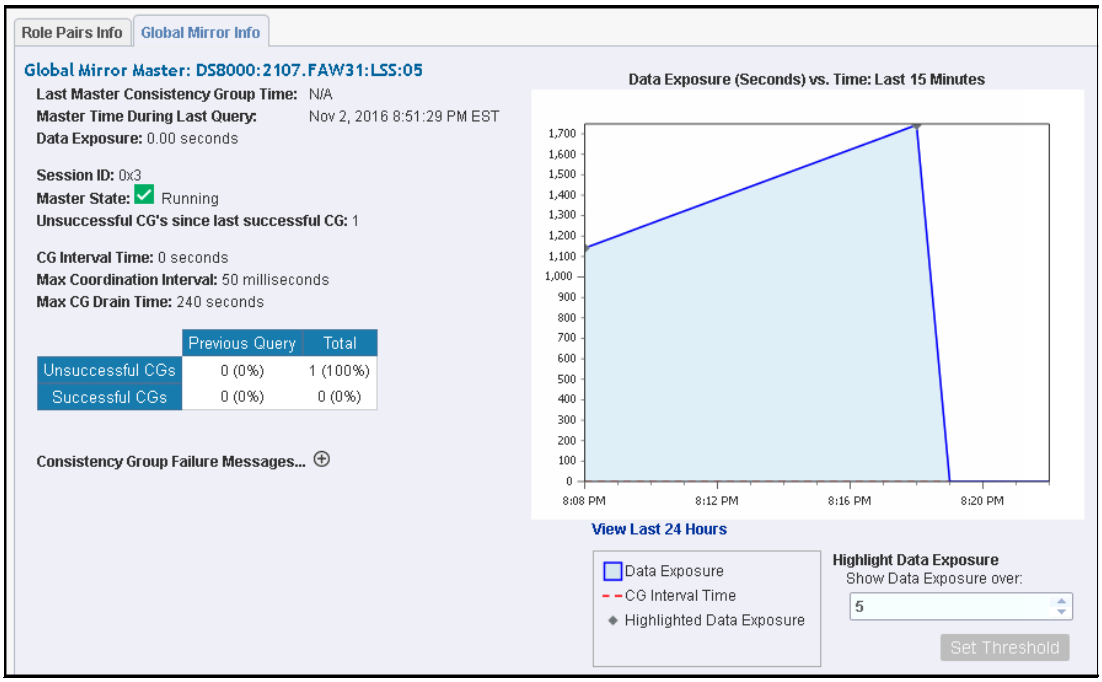

Figure 6-27 Global Mirror information

10. We start also the Metro Mirror (H1-H2) as shown in Figure 6-28.

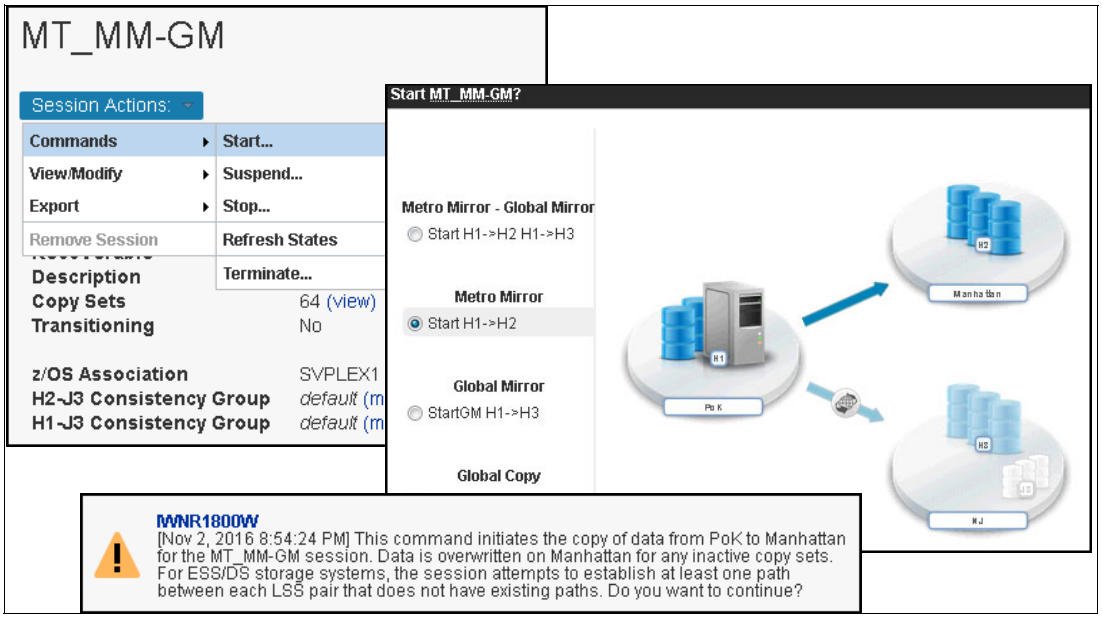

Figure 6-28 Start MM H1-H2

11. Because the session is configured to be managed by parallel sysplex, the MM will also activate, and, after a while, HyperSwap and the configuration is back in the initial phase, as shown in Figure 6-29.

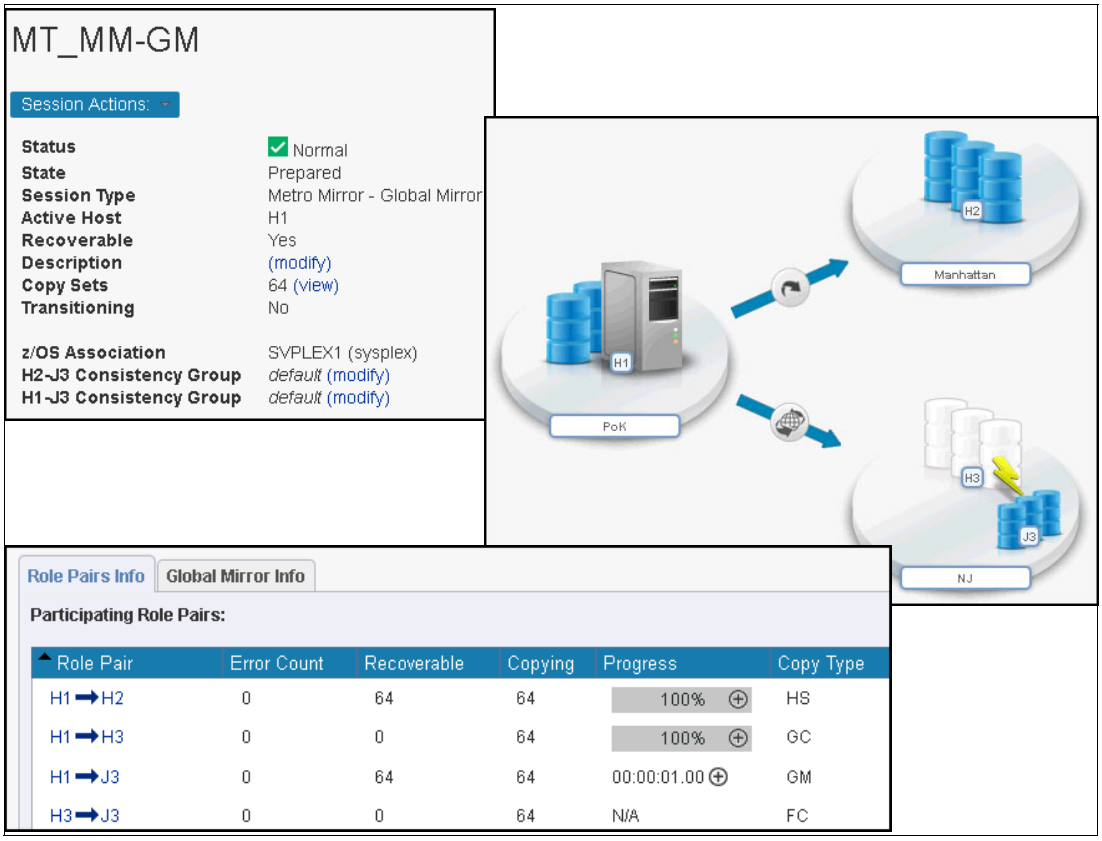

Figure 6-29 Initial state (after recovery)

6.2.5 Recovering the initial CSM server

We can now activate the initial Copy Services Manager server. The database of the initial CSM server is obsolete because the CSM server has been unavailable and it needs to update its configuration.

Reverting to the initial CSM server in *PoK* can be done assigning it as the standby server to the active CSM server that is running in *NJ* (H3). The process is repeated in reverse. This overwrites the configuration of the CSM server in *PoK* and synchronizes it with the current one from the CSM server in *NJ*. After synchronization is complete, we repeat the takeover and make the Copy Services Manager server running on z/OS the Active CSM server:

1. After the z/OS system is running on *PoK* (H1), you can connect to the initial CSM server using the browser.

   The server configuration data is outdated (the configuration before failure still exists). The GUI shows that this is the active server, but status is *Disconnected Consistent*.

The currently active CSM server from NJ (H3) is known as STANDBY (see Figure 6-30). Also, you can have access to sessions but they are not actual.

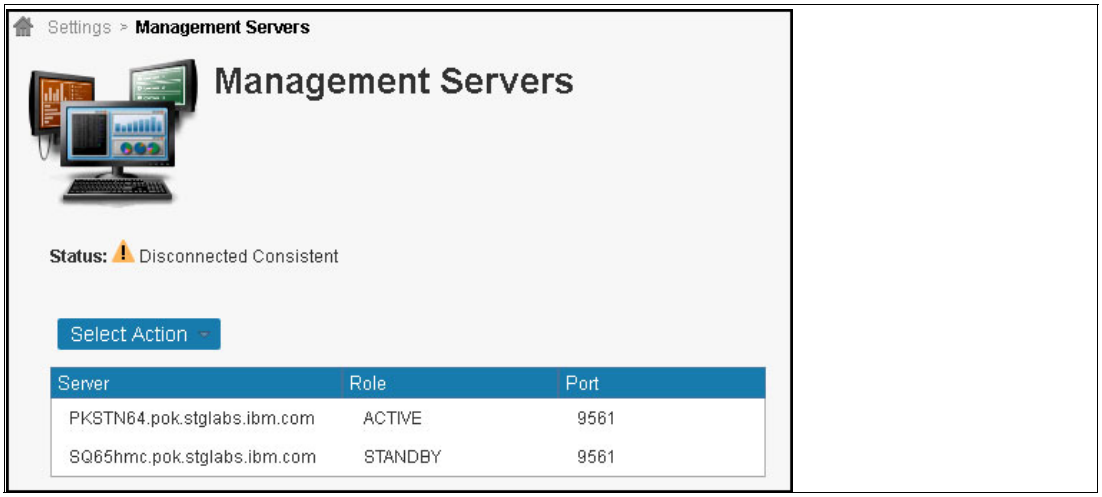

Figure 6-30 Disconnected Consistent status

2. Before assigning this CSM server as Standby we have to remove the old standby definition, as shown in Figure 6-31.

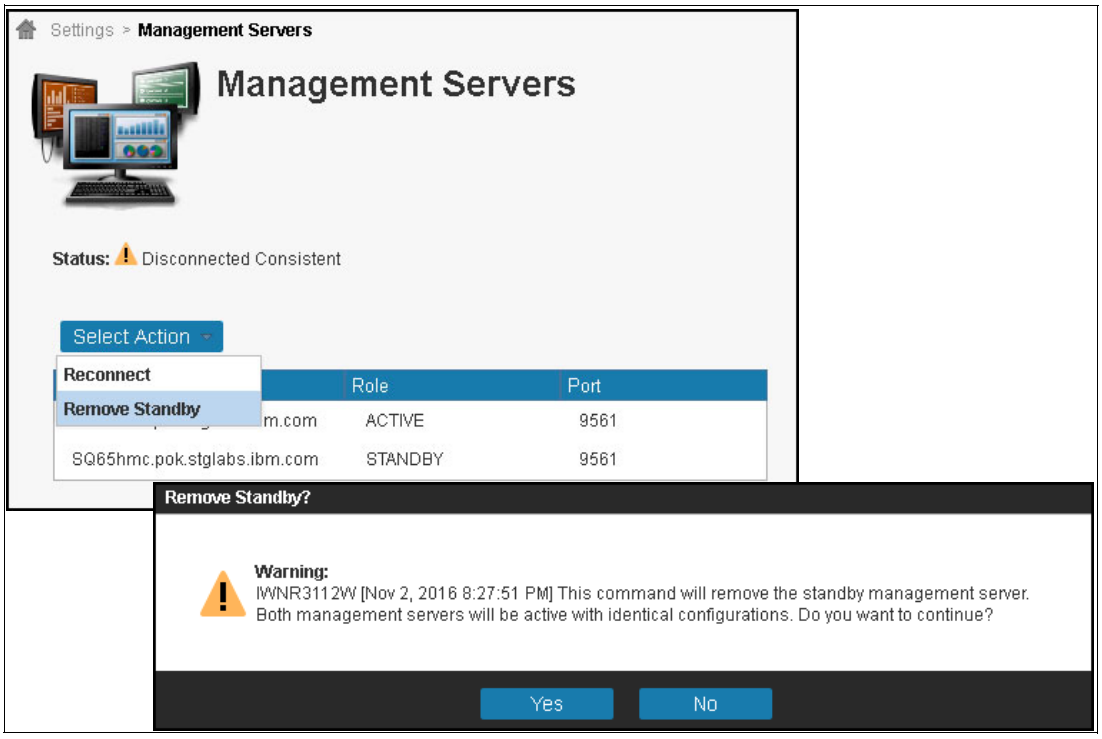

Figure 6-31 Remove Standby definition

3. After the standby definition is removed, we can "*Set this Server as Standby*" for the active CSM server in *NJ* (H3). This removes any old configurations and synchronizes with the active server, as shown in Figure 6-32.

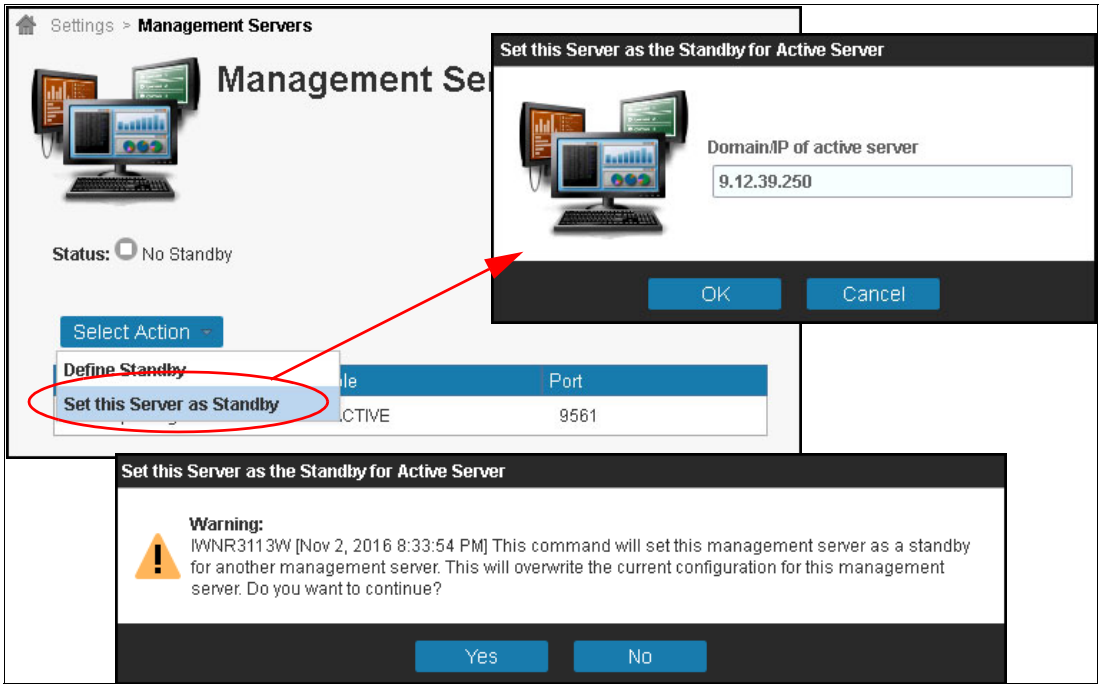

Figure 6-32 Set this Server as Standby

- When the configuration becomes synchronized, we issue the takeover, as shown in Figure 6-33.

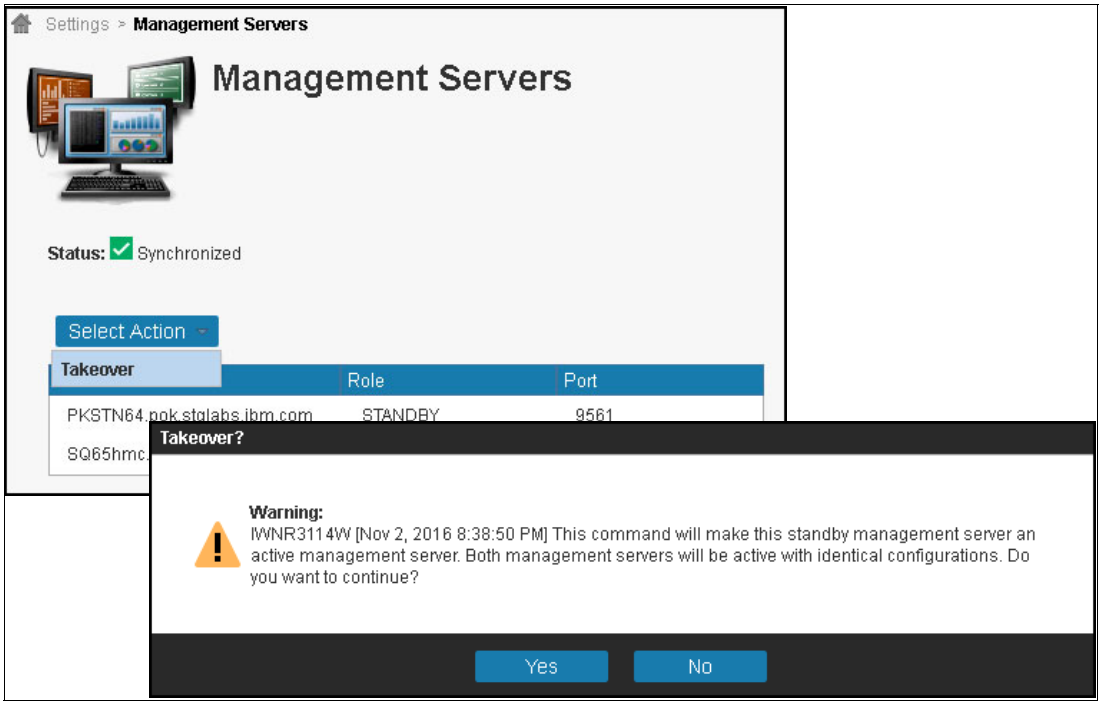

Figure 6-33 Initiating takeover to the CSM server in PoK

Attention: Taking over from the active CSM server doesn't mean that the roles switch and that the old active server becomes the standby.

4. Finally, we assign the CSM server in NJ as the standby server. We use the **Settings** → **Management Servers** section and chose **Define Standby** Server, as shown in Figure 6-34.

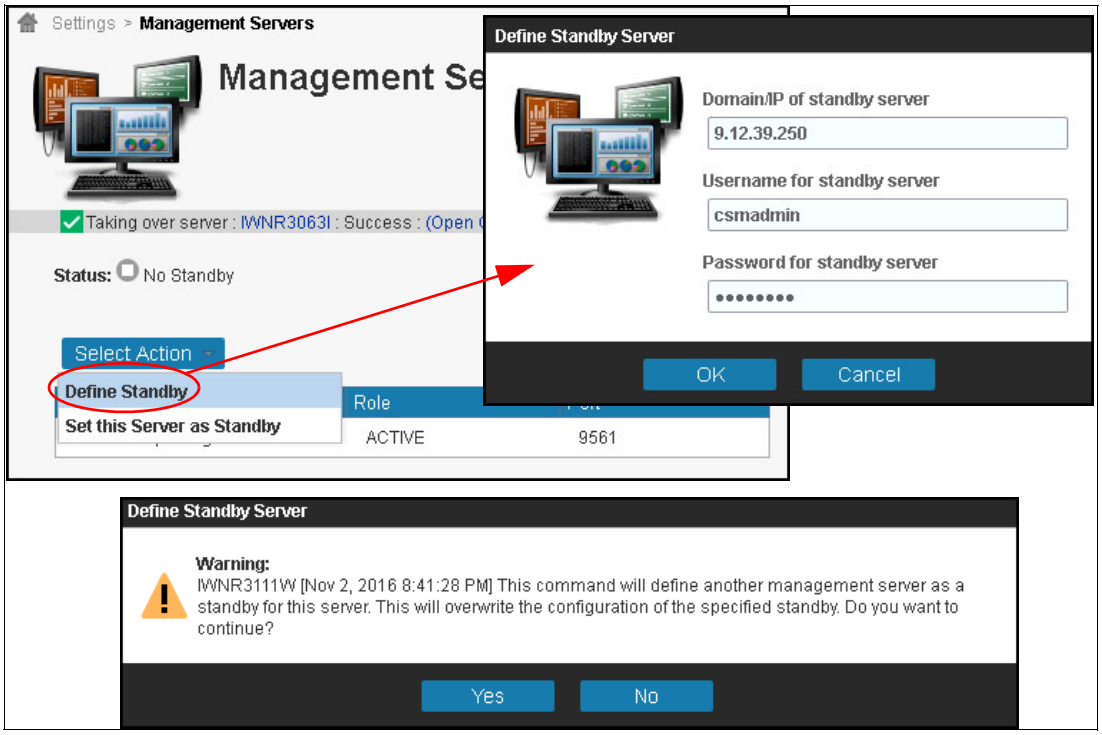

Figure 6-34 Defining Standby Server

7

Securing IBM z/OS HyperSwap communication with IBM Copy Services Manager

The scenario in this chapter describes how to configure secure TCP/IP communication between IBM z/OS HyperSwap function and an IBM Copy Services Manager (CSM) server instance.

In addition to managing storage replication, IBM Copy Services Manager can be used to initiate a planned HyperSwap operation. The configuration described in this chapter enables z/OS HyperSwap to communicate with CSM securely using Transport Layer Security (TLS). z/OS HyperSwap can be also triggered or initiated with no CSM involvement.

This chapter describes the following topic:

- IBM z/OS HyperSwap

7.1 IBM z/OS HyperSwap

IBM z/OS HyperSwap function has been designed to enhance system high availability by masking storage failures in a (IBM DS8000 series) Metro Mirror configuration. HyperSwap is a z/OS component which manages storage access paths on an IBM Z server. If there is an I/O error in the primary set of volumes (for example, a storage subsystem failure), HyperSwap automatically swaps I/O to a secondary set of volumes (Metro Mirror target) with little to no impact to the application.

7.1.1 Introduction

IBM Copy services Manager (CSM) manages z/OS HyperSwap using a *HyperSwap session*. For Copy Services Manager, the HyperSwap session is managed as a Metro Mirror session. CSM manages the z/OS HyperSwap function in the z/OS systems for controlled swaps.

IBM Copy Services Manager supports enabling the following session types with HyperSwap capabilities:

- Basic HyperSwap
- Metro Mirror
- Metro-Global Mirror
- Multi-target session types:
  - Metro Mirror - Metro Mirror
  - Metro Mirror - Global Mirror
  - Metro Mirror - Global Mirror with Practice

7.1.2 z/OS HyperSwap overview

This section provides a high-level overview of the z/OS HyperSwap functionality. For additional information and details about how to configure z/OS HyperSwap, see the IBM Redpaper™ publication *IBM DSS 8880 and z Systems Synergy*, REDP-5186.

A typical deployment scenario is shown in Figure 7-1.

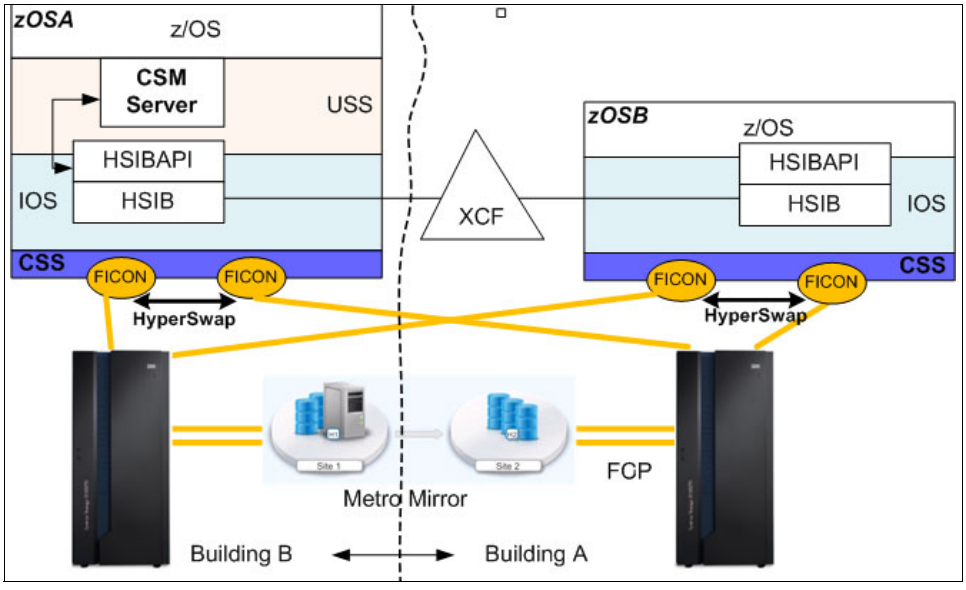

Figure 7-1 z/OS Basic HyperSwap in a Parallel Sysplex

In the course of the actual swap operation, HyperSwap is responsible to issue all necessary Copy Services commands to perform the complete failover to H2. This failover also leads to the Metro Mirror state change of the H2 volumes from secondary DUPLEX to primary SUSPENDED.

7.1.3 Securing TCP/IP communication for CSM: An overview

IBM Copy Services manager can be deployed on several platforms. For supported platforms check the current information at IBM Support.

CSM communicates with the z/OS HyperSwap using two methods:

- Java Native Interface (JNI)
- TCP/IP

Communication using JNI

For most common z/OS environments, IBM Copy Services Manager server code runs in a z/OS image (in an address space) and utilizes the underlying Unix System Services (USS).

When the CSM server instance (code) runs inside the same z/OS image that has z/OS HyperSwap configured (or in a z/OS image part of a Parallel Sysplex, as shown in Figure 7-1 on page 154), the CSM server communicates with the z/OS HyperSwap through the HyperSwap API (HSIBAPI z/OS address space) directly (no TCP/IP communication), using the Java Native Interface (JNI). In this case, no additional security is required for this type of communication.

Communication using TCP/IP

In a configuration where the CSM server runs outside the scope of the z/OS HyperSwap configuration, the communication between the CSM server and the z/OS HyperSwap is performed via TCP/IP. In this case, regardless of the relative location of the platform on which the CSM server runs, communication (via TCP/IP) with the z/OS HyperSwap address space *must* be secured.

The CSM server, outside the scope of the z/OS HyperSwap, can be deployed:

- In a different z/OS image, not part of a Parallel Sysplex
- In a z/OS image in a different Parallel Sysplex
- In a system not running z/OS (AIX, Linux, or Microsoft Windows, or the DS8000 HMC)

Scenario

Consider the following DR sample scenario, presented in Figure 7-2 on page 156:

- In the primary site, the z/OS LPAR is configured with HyperSwap over the MM relationship between two storage subsystems (H1 and H2). The z/OS LPAR also runs the active (primary) instance of the CSM server.
- The active (primary) CSM Server replicates to the standby (backup) CSM server hosted on a system in the DR site, outside the z/OS or the Parallel Sysplex. The Standby CSM server can be installed and configured on an “open” system (AIX, Linux, or Windows) located in a DR site (Site 3 in the example).

Tip: After a planned or unplanned HyperSwap, the active volumes are swapped from H1 to H2, and CSM remains passive and not involved. HyperSwap notifies CSM about the session state change after the HyperSwap operation is completed.

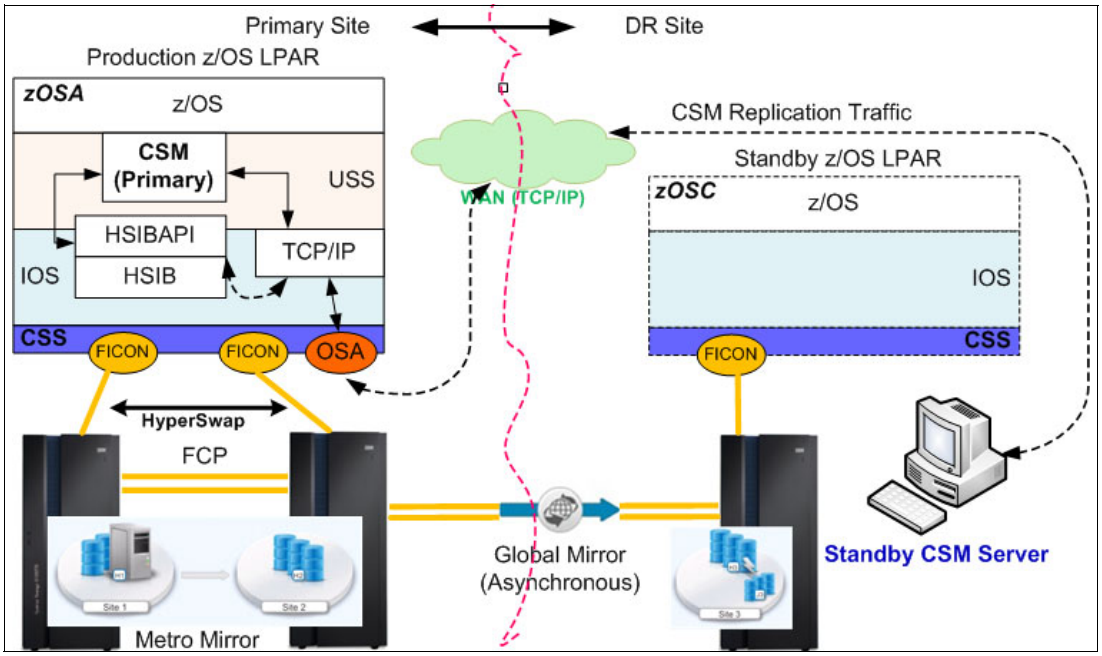

Figure 7-2 Basic MGM scenario using Copy Services Manager

Tip: The standby (backup) CSM server can also be configured to run on a DS888x Hardware Management Console (HMC)[a] that manages the storage subsystems located in the DR site (Site 3).

a. There may be certain limitations for setting up secure communication between the z/OS HyperSwap and a CSM server instance running on the DS888x HMC. Check the latest DS888x HMC and CSM documentation before you configure this HMC as a CSM server.

- Under normal circumstances, the active CSM server is the primary CSM server, and it is used to manage the z/OS HyperSwap configured on the production z/OS LPAR.

  Because the CSM (server) and the HyperSwap address run on the same z/OS image, CSM uses the Java Native Interface (JNI) to communicate to the z/OS HyperSwap address space.
- For high availability, the CSM configuration data is replicated over a TCP/IP connection to the standby (backup) CSM server running in the DR site (Site 3).
- The DS8000 in the DR site (Site 3) is configured for Global Mirror. The z/OS LPAR in the DR site is not activated when production runs in the primary z/OS LPAR.
- In case the primary (production) location becomes unavailable (including the primary CSM server), the standby (backup) CSM server must be activated to manage the replication and to make available the H3 copy (Global Mirror) in the DS8000 located in the DR site (Site 3).
- When the CSM server running in the DR site (Site 3) becomes active, it should be used to manage the replication even after the primary site becomes available (z/OS HyperSwap), and until proper recovery is performed. Recovery may require to resynchronize (Site 3 to Site 1) data between the CSM servers (active CSM server located in DR site, which has the latest CSM data, and the CSM server in the primary site).

  In this situation, the communication between the (now active) CSM instance in the DR location and the z/OS HyperSwap in the Primary site (Site 1) must be secured (encrypted).

Why use AT-TLS?

Because the z/OS HyperSwap has not been coded to support encrypted communication, the z/OS Application Transparent Transport Layer Security (AT-TLS) facility can be used to provide encrypted data transfer between the active CSM server and the z/OS HyperSwap.

7.1.4 Securing communication between z/OS HyperSwap and the CSM Server

This section describes the basic configuration steps for securing TCP/IP communication between the z/OS HyperSwap (identified as the SERVER role) and the IBM Copy Services Manager Server (identified as the CLIENT role).

Note: For our scenario we have deployed a standalone CSM server running a Linux (x86_64) platform.

Overview

The z/OS HyperSwap has been enhanced through the *HyperSwap Sockets Server* (BHIHSRV z/OS address space) to allow IBM Copy Services Manager to connect to the HyperSwap address space using sockets, in addition to the Java Native Interface (JNI).

This enables a CSM instance running on z/OS in one sysplex or a standalone z/OS image to manage HyperSwap sessions running one or more other sysplexes, and also allows HyperSwap management from a CSM instance running on a non-z/OS system (AIX, Windows, or Linux).

Scenario considerations

For the scenario presented in Figure 7-2 on page 156 we need to secure TCP/IP communication between the CSM instance running on a non-z/OS operating system located in Site 3 (activated after a CSM failover) and the z/OS HyperSwap. This configuration is shown in Figure 7-3.

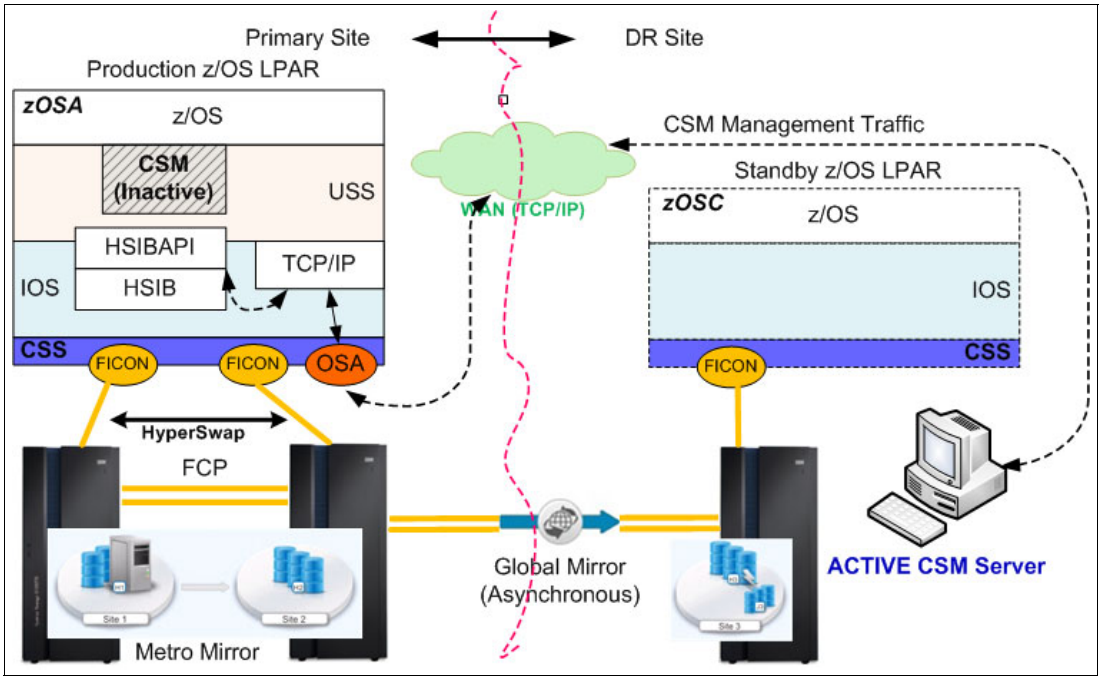

Figure 7-3 z/OS HyperSwap management from DR CSM instance

Secure TCP/IP communication configuration is needed if we want to manage the z/OS HyperSwap in Site 1 before the CSM instance is recovered to the initial server running on z/OS (CSM active/standby configuration is shown in Figure 7-2 on page 156).

The z/OS SSL/TLS infrastructure must be configured prior to using encrypted TCP/IP communication between the z/OS HyperSwap and the CSM instance in Site 3.

Terminology

This section introduces general terminology that is useful in providing understanding of the parties and entities involved in building and serving the secure communication context.

Client-server model

The client-server model has the following roles:

- Client role: A piece of a computer application that communicates with another entity (server role) that provides access to some function or feature (data) that is requested by the client. The server responds by providing the services and data requested. The client can be:
  - A user (person) that interacts with some kind of computer interface used to request information.
  - Another application in an environment with automated actions.
- Server role: A piece of computer application that listens to client requests and responds with the requested data.

Security: identity management

Identity management in terms of security includes the following elements:

- *Identity:* When two parties need to communicate (client and server roles), the first requirement is to identify with each other. The client must present some kind of identity form to the server, and the server must be able to identify itself to the client. The identity is presented in the form of a *Digital Certificate*.
- *Certificate Authority:* To establish the parties' identities, some kind of third party that can guarantee the identity presented by each party may be required. The third party is trusted by the parties that need to identify with each other. The Certificate Authority is used in the initial phase of the communication to establish trust.
- When a party presents its identity (*Digital Certificate*), the other party checks the Digital Certificate by referring to the *Certificate Authority.*
- *Digital Certificate:* A Digital Certificate conforms to a standard format (for example, x.509) and contains information about the identity that it represents. In addition to the public key of the identity that it represents, the certificate contains an expiration date, information about the identity (organization, division, and so on), issuer information, and key fingerprint and algorithm information.
- *Key Ring:* A collection of certificates (identities) available to a party for use during the secure communication process.

Security: Data privacy and data integrity

Data privacy is achieved by encrypting the data sent over a communication channel (between the client and the server). *Data integrity* is achieved by providing data digest mechanisms, for example by using hashing algorithms (one way hashing functions):

- *Cypher/decipher:* Cypher is a mathematical method used to encrypt the data (data -> code). Decipher is the reverse action used to retrieve the data from the code.

- *Encryption/decryption:* Encryption is the process converting some kind of data into a form (code) that cannot be (easily) understood by a third party that accidentally intercepts a copy of the encrypted data. Decryption is the process of extracting the data from the encrypted format (code). Over a communication channel, the data is encrypted at its source and decrypted at its destination.
- *Integrity:* To ensure that the data reaches its destination unaltered (unchanged in any form), additional information (data digest generated at source) must be provided at he destination. This information is compared at destination with the digest of the data (generated at destination) to ensure integrity.

Client and server roles (client-server modes) applied to our environment

Important: In this model of deployment (IBM z/OS HyperSwap managed by Copy Services Manager), for securing the client-server communication using TLS, the *HyperSwap Manager (part of z/OS) plays the SERVER role*, while the *CSM server instance plays the CLIENT role*.

Configuration FLOW

A local Certificate Authority is used to sign the HyperSwap Manager (server role) certificate:

- The RACF (Resource Access Control Facility) is used to generate the Certificate Authority (local to the z/OS image where the HyperSwap address space runs). The z/OS (local) Certificate Authority (CA) fits the purpose to provide the client (CSM instance) with the level of trust required to control the z/OS HyperSwap server.
- The Certificate Authority provides a self-signed certificate (*trusted root certificate*). The self-signed certificate is generated using RACF. The CA's self-signed certificate is provided to the client (the CSM server) that needs to verify the certificate (identity) of the party involved in communication (the HyperSwap Manager identity).

  Tip: A publicly available certificate authority can be used to sign the z/OS HyperSwap server certificate; however, this would generate unnecessary administrative burden for the limited function that must be served (authentication and securing communication between the CSM instance and the z/OS HyperSwap Manager).

- RACF is also used to generate the HyperSwap Manager's key pair (public and private keys for the server role). The HyperSwap Manager (server) certificate is signed using the CA's trusted root certificate (previously generated).
- Certificates (CA's certificate and the HyperSwap Manager certificate) are stored in a key ring that is made available to the AT-TLS policy manager. The key ring a is used by AT-TLS on behalf of the parties involved in the communication (for authentication and data encryption).
- AT-TLS policies must be configured and installed.

AT-TLS overview

Socket applications that have not been coded to support encrypted communication send and receive clear text over the socket.

- In z/OS, AT-TLS (Application Transparent Transport Layer Security) provides TLS support for existing clear-text (no encryption) socket applications without requiring any application changes (re-coding).
- The Transport Layer Security (TLS) protocol provides transport layer security (authentication, integrity, and confidentiality) for a secure connection between two applications.

- The TLS protocol begins with a handshake, in which the two applications (client and server) agree on a cipher suite. A cipher suite is composed of cryptographic algorithms to be used for authentication, session encryption, and hashing (digital fingerprinting).
- After the client and server applications have negotiated a cipher suite, they authenticate each other and generate a session key. The session key is used to encrypt and decrypt all data traffic sent between client and server.
- Applications that have not been coded to use the SSL/TLS libraries (such as z/OS HyperSwap) are unaware of the encrypted traffic carried on their behalf by AT-TLS.
- In addition, support is provided for applications that require awareness of AT-TLS for status or to control the negotiation of security.

AT-TLS is policy-driven, managed by a policy agent (PAGENT):

- AT-TLS provides application-to-application security using policies. The policies are defined and loaded into the stack by *policy agent* (PAGENT).
- When AT-TLS is enabled and a newly established connection is first used, the TCP layer of the stack searches for a matching AT-TLS policy.
- If no policy is found, the connection is made without AT-TLS involvement. If a policy is found, a sequence of activities is generated to catty out encrypted traffic, similar to the one illustrated in Figure 7-4.

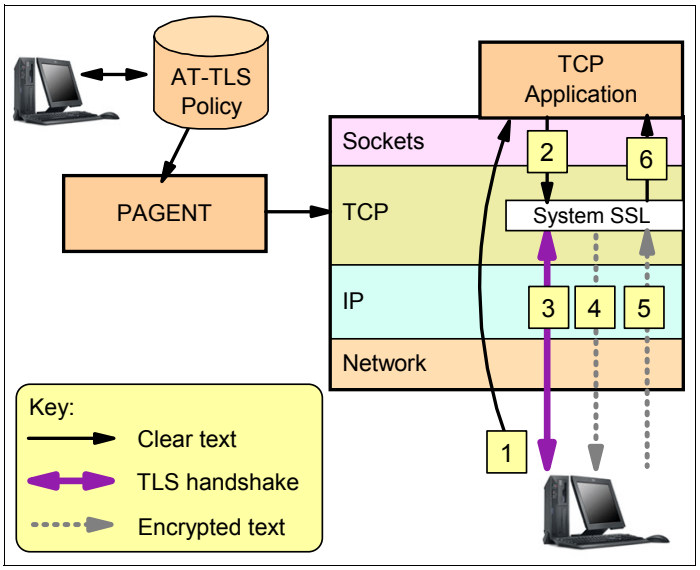

Figure 7-4 AT-TLS basic flow

This diagram illustrates the following flow:

1. The client connection to the server gets established in the clear (no security, TCP handshake only).
2. The server sends data in the clear and the TCP layer queues it.
3. The TCP layer invokes System SSL to perform a TLS handshake under the identity of the server, using policy information.
4. The TCP layer invokes System SSL to encrypt the queued data and sends it to the client.
5. The client then sends encrypted data and the TCP layer invokes System SSL to decrypt the data.
6. The server receives data in the clear.

For more information on z/OS AT-TLS, see the Redbook *IBM z/OS V2R1 CS TCP/IP Implementation Volume 4: Security and Policy-Based Networking*, SG24-8099.

7.1.5 Generating a self-signed (CA) certificate in z/OS

This certificate represents the local certificate authority (CA) and is used as a CA certificate.

To create a RACF key ring, you must first generate a RACF CA certificate and a personal certificate for IBM z/OS Basic HyperSwap, then connect the certificates to the key ring.

A RACF key ring is connected to a set of personal certificates and trusted certificates that are stored in the RACF database. The RACF command **RACDCERT** is used to create and delete key rings and to connect or disconnect certificates to the key rings

To create a RACF key ring to be used by AT-TLS on behalf of IBM z/OS Basic HyperSwap, complete the following steps:

1. We use RACF as a local Certification Authority. We generate a RACF certificate authority (CA) certificate as shown in Example 7-1.

   Example 7-1 Generating the CA certificate (self signed)

   ```
   RACDCERT CERTAUTH GENCERT -
     SUBJECTSDN (OU('ITSO CSM Certificate Authority') -
     O('ITSO') -
     C('US')) -
     KEYUSAGE(HANDSHAKE DATAENCRYPT DOCSIGN CERTSIGN) -
     WITHLABEL('CSM Local Certificate Authority')
   ```

2. Next, we export this certificate to a data set so that we can use it on the system where the IBM Copy Services Manager server is running, as shown in Example 7-2.

   Note: This data set (file) will be later transferred to the CSM server (which plays the client role for the z/OS HyperSwap server in this setup) to be used to verify the z/OS HyperSwap Manager (server) identity.

   Example 7-2 RACF commands to export the CA certificate to a data set

   ```
   RACDCERT EXPORT (LABEL('CSM Local Certificate Authority')) -
     CERTAUTH DSN('CSM.LOCCERTA.CERT') -
     FORMAT(CERTDER)
   ```

3. We run the RACF commands shown in Example 7-3 to refresh and ensure that the certificate is in storage.

   Example 7-3 Sample RACF SETR commands

   ```
   SETR CLASSACT(DIGTCERT)
   SETR RACLIST(DIGTCERT)
   SETR RACLIST(DIGTCERT) REFRESH
   ```

   RACF is also used to generate the *personal* certificate that will be used for the z/OS HyperSwap Manager and the CA certificate is used to sign the personal certificate.

4. We generate a *personal* certificate for the IBM z/OS Basic HyperSwap server.This certificate identifies the instance of z/OS HyperSwap Manager. This certificate is presented to the partner application during the SSL handshake (CSM instance in this case).

This certificate must be associated with the user ID under which the z/OS HyperSwap Manager is running.

Finding the ID for HyperSwap

To find the ID, complete the following steps:

1. To identify the user under which the z/OS HyperSwap is running, use the command shown in Example 7-4.

Example 7-4 Identifying the user for HyperSwap

```
RLIST STARTED HSIB STDATA
 CLASS      NAME
 -----      ----
 STARTED    HSIB*.** (G)

 LEVEL  OWNER      UNIVERSAL ACCESS  YOUR ACCESS  WARNING
 -----  --------   ----------------  -----------  -------
  00    HSIB            NONE               NONE    NO

 INSTALLATION DATA
 -----------------
 NONE

 APPLICATION DATA
 ----------------
 NONE

 AUDITING
 --------
 FAILURES(READ)

 NOTIFY
 ------
 NO USER TO BE NOTIFIED

 STDATA INFORMATION
 ------------------
 USER= HSIB
 GROUP= SYS1
 TRUSTED= NO
 PRIVILEGED= NO
 TRACE= NO
    READY
```

2. The following Example 7-5 shows how to use a RACF command to generate the *personal* certificate for z/OS HyperSwap that is running under user ID *HSIB*.

Example 7-5 Generating the personal certificate for the HyperSwap user ID

```
RACDCERT ID(HSIB) GENCERT -
  SUBJECTSDN(CN('CSM Client') -
  OU('Hyperswap Server') -
  O('CSM') -
  C('US')) -
  WITHLABEL('Hyperswap Manager') -
  SIGNWITH(CERTAUTH LABEL('CSM Local Certificate Authority')) -
  KEYUSAGE(HANDSHAKE DATAENCRYPT DOCSIGN)
```

3. We create a RACF key ring and connect the certificates to the key ring. The key ring is needed by AT-TLS. The RACF key ring must be associated with a user ID (in this case, the HSIB user ID). The key ring must have a name (in this case, *csmkeyring*, as shown in Example 7-6).

Example 7-6 Creating the key ring

```
RACDCERT ID(HSIB) ADDRING(csmkeyring)
READY
```

4. The z/OS HyperSwap Manager *personal* certificate must be connected to the key ring together with the CA certificate, as shown in Example 7-7.

Example 7-7 Adding the certificate to the key ring

```
RACDCERT ID(HSIB) CONNECT(LABEL('Hyperswap Manager') RING(csmkeyring) default)
READY

RACDCERT ID(HSIB) CONNECT(CERTAUTH LABEL('CSM Local Certificate Authority')
RING(csmkeyring))
READY
```

5. We must allow user ID HSIB permission to READ the key ring, as shown in Example 7-8.

Note: Before allowing user ID permission to read the key ring, you may need to define the IRR.DIGTCERT.LISTRING to class FACILIY:

```
RDEFINE FACILITY IRR.DIGTCERT.LISTRING UACC(NONE)
```

Example 7-8 Allowing user ID READ permission to the FACILITY class

```
PERMIT IRR.DIGTCERT.LISTRING CLASS(FACILITY) ID(HSIB) ACCESS(READ)
ICH06011I RACLISTED PROFILES FOR FACILITY WILL NOT REFLECT THE UPDATE(S) UNTIL
A SETROPTS REFRESH IS ISSUED
READY

SETR RACLIST(FACILITY DIGTCERT) REFRESH
ICH14063I SETROPTS command complete.
READY
```

6. Complete the following verifications:

   a. We verify the CA certificate, as shown in Example 7-9.

Example 7-9 Verifying the certificate authority

```
RACDCERT LIST(LABEL('CSM Local Certificate Authority')) CERTAUTH

 Digital certificate information for CERTAUTH:

   Label: CSM Local Certificate Authority
   Certificate ID: 2QiJmZmDhZmjgcPi1EDTloOBkODDhZmjiYaJg4GjhUDBpKOIlpmJo6hA
   Status: TRUST
   Start Date: 2016/11/02 23:00:00
   End Date:   2017/11/03 22:59:59
   Serial Number:
        >00<
   Issuer's Name:
```

```
        >OU=ITSO CSM Certificate Authority.O=ITSO.C=US<
   Subject's Name:
        >OU=ITSO CSM Certificate Authority.O=ITSO.C=US<
   Signing Algorithm: sha256RSA
   Key Usage: CERTSIGN
   Key Type: RSA
   Key Size: 2048
   Private Key: YES
   Ring Associations:
     Ring Owner: HSIB
     Ring:
        >csmkeyring<

 READY
```

b. We verify the personal certificate for the z/OS HyperSwap server as shown in Example 7-10.

Example 7-10 Personal certificate for z/OS HyperSwap server

```
RACDCERT LIST(LABEL('Hyperswap Manager')) ID(HSIB)

 Digital certificate information for user HSIB:

   Label: Hyperswap Manager
   Certificate ID: 2QTI4snCyKiXhZmipoGXQNSBlYGHhZlA
   Status: TRUST
   Start Date: 2016/11/09 00:00:00
   End Date:   2017/11/09 23:59:59
   Serial Number:
        >05<
   Issuer's Name:
        >OU=ITSO CSM Certificate Authority.O=ITSO.C=US<
   Subject's Name:
        >CN=CSM Client.OU=Hyperswap Server.O=CSM.C=US<
   Signing Algorithm: sha256RSA
   Key Usage: HANDSHAKE
   Key Type: RSA
   Key Size: 2048
   Private Key: YES
   Ring Associations:
     Ring Owner: HSIB
     Ring:
        >csmkeyring<

 READY
```

c. And we check the key ring, as shown in Example 7-11.

Example 7-11 Listing the key ring

```
RACDCERT LISTRING(csmkeyring) ID(HSIB)

 Digital ring information for user HSIB:

   Ring:
        >csmkeyring<
```

```
Certificate Label Name              Cert Owner     USAGE      DEFAULT
-------------------------------     ------------   --------   -------
Hyperswap Manager                   ID(HSIB)       PERSONAL     YES
CSM Local Certificate Authority     CERTAUTH       CERTAUTH     NO

READY
```

7.1.6 Preparing the CSM server for secure communication with HyperSwap

Tip: The CSM server in DR site has been deployed on a Linux (x86_64) platform. The CSM installation provides all the tooling required to perform key and certificate management for the CSM server.

Important: Starting with CSM 6.1.5 the z/OS HyperSwap server certificate can be imported directly into CSM GUI during the Host Connection setup procedure. If your CSM server is at 6.1.5 or later, the steps in this section can be skipped.

Complete the following steps to prepare the CSM server:

1. Transfer the CA certificate to the CSM server (in Site 3) and add it to the CSM server key store configuration:
   a. Place the certificate that was exported in Step 2 on page 161 in the z/OS UNIX System Services hierarchical file system (HFS) by using the **OPUT** command shown in Example 7-12. Use the BINARY option.

   Example 7-12 OUTPUT command example

   ```
   OPUT 'CSM.LOCCERTA.CERT' '/u/myuser/sc30.cert' BINARY
   ```

   b. Log on to the CSM instance (in Site 3) and transfer the CA certificate file on the local file system. Use binary FTP to download. We have transferred it into the CSM directory tree, as shown in Example 7-13.

   Example 7-13 CA certificate transferred to CSM server in Site 3

   ```
   [root@csmserver wlp]# pwd
   /opt/IBM/CSM/liberty/wlp
   [root@csmserver wlp]# ls -l *cert
   -rw-r--r-- 1 root root 910 Jan 23 18:31 sc30.cert
   ```

   c. Create a Java Key Store (JKS - `.jks` file type) file and import the CA certificate into the JKS. The JKS will be supplied to the CSM server for verifying the identity for the z/OS HyperSwap server it has to connect to.
   d. Identify the **ikeyman** utility delivered with the IBM Java (in our Linux installation there is one installed together with the IBM Copy Services Manager code), as shown in Example 7-14.

   Example 7-14 Identifying the ikeyman utility

   ```
   [root@csmserver ~]# find /opt -name ikeyman
   /opt/IBM/CSM/liberty/wlp/IBM/Java/jre/bin/ikeyman
   ```

e. Launch the **ikeyman** utility in a graphical environment, as shown in Figure 7-5.

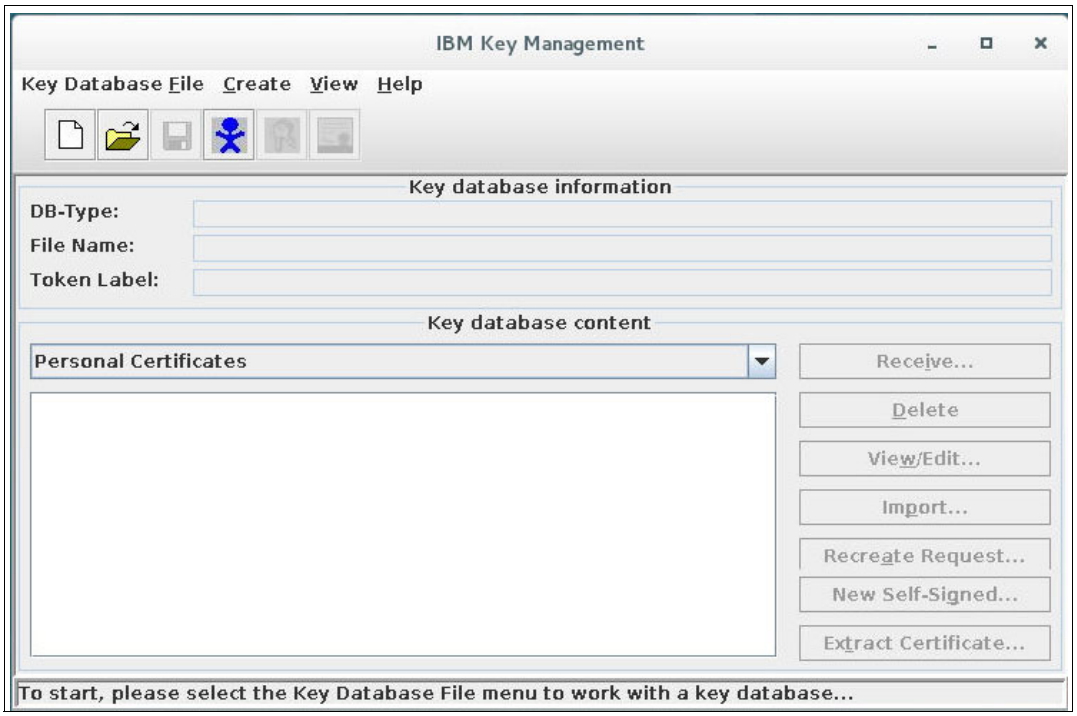

Figure 7-5 The ikeyman application GUI

f. Initiate the creation of a new Key Database file, as shown in Figure 7-6.

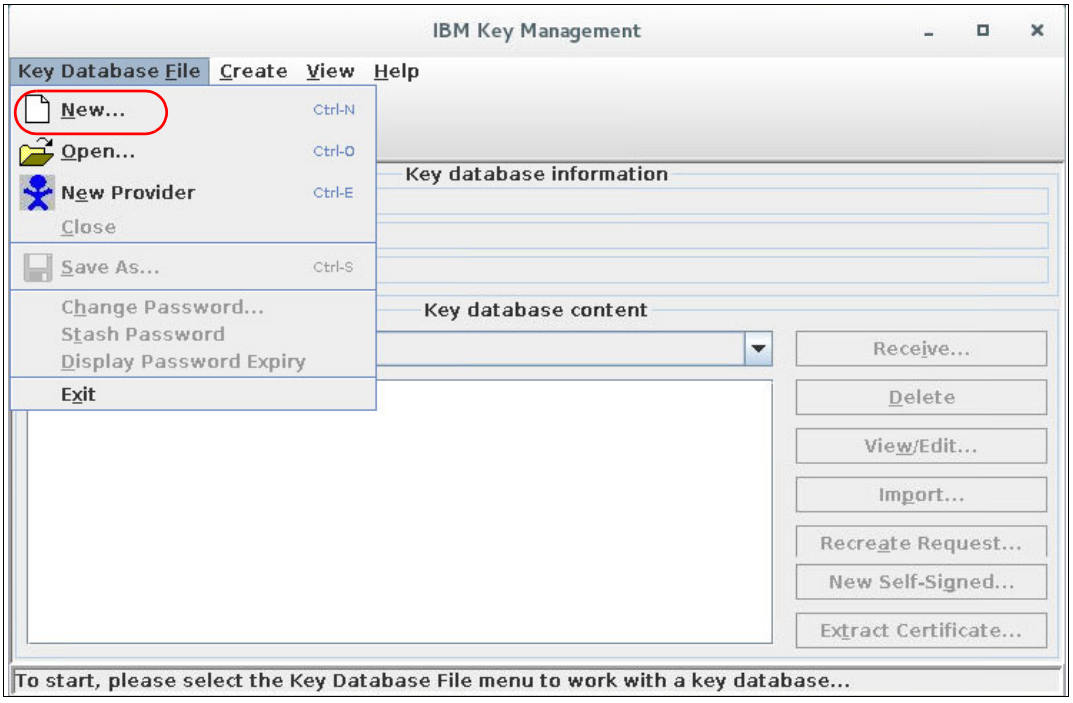

Figure 7-6 Creating a new Key Database file

g. Define the name and location of the Key Database file, as shown in Figure 7-7. The Key Database file will be later copied in a location suitable for the CSM server.

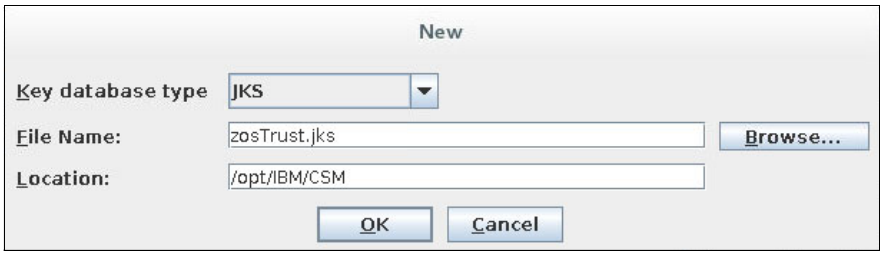

Figure 7-7 The key Database file

h. Supply a password for accessing the Key Database, as shown in Figure 7-8.

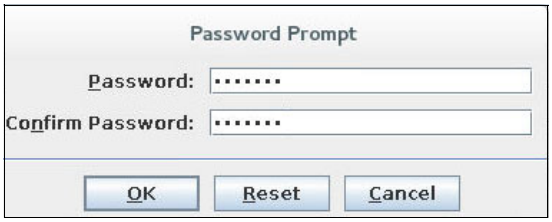

Figure 7-8 Protecting the Key Database with a password

i. Specify the type of certificate you want to add (import) to the newly created Key Database. We want to add a signer certificate (CA certificate), as shown in Figure 7-9.

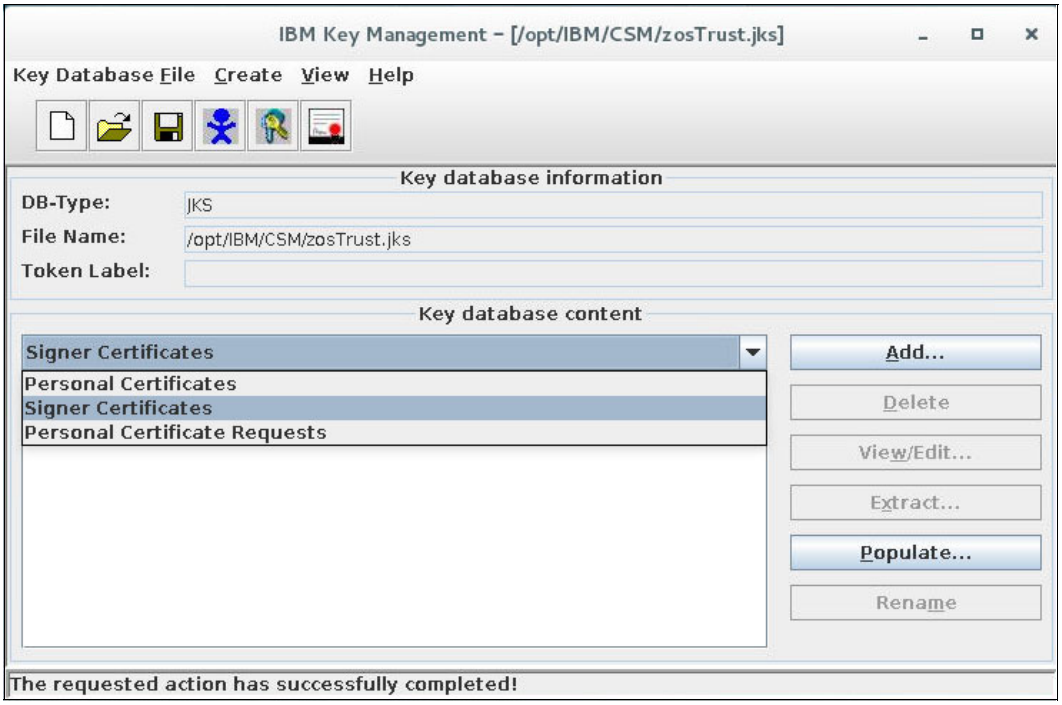

Figure 7-9 Selecting the type of certificate you want to add (import)

j. Import the CA certificate from the file transferred earlier (b on page 165) into the Key Database, as shown in Figure 7-10.

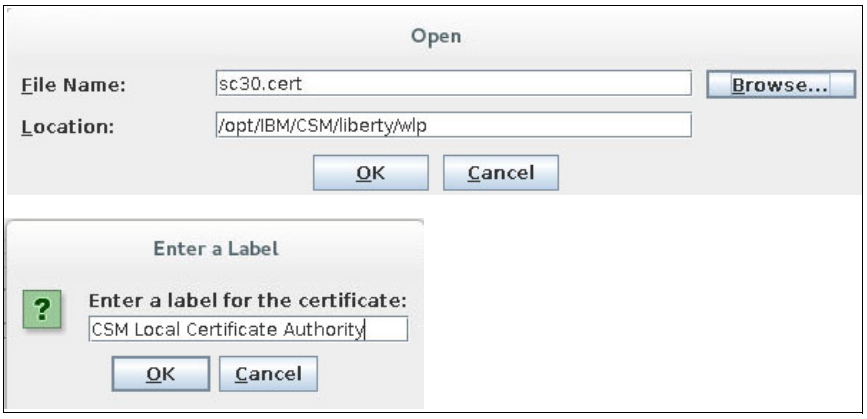

Figure 7-10 Importing the CA certificate

k. The imported certificate is shown in Figure 7-11.

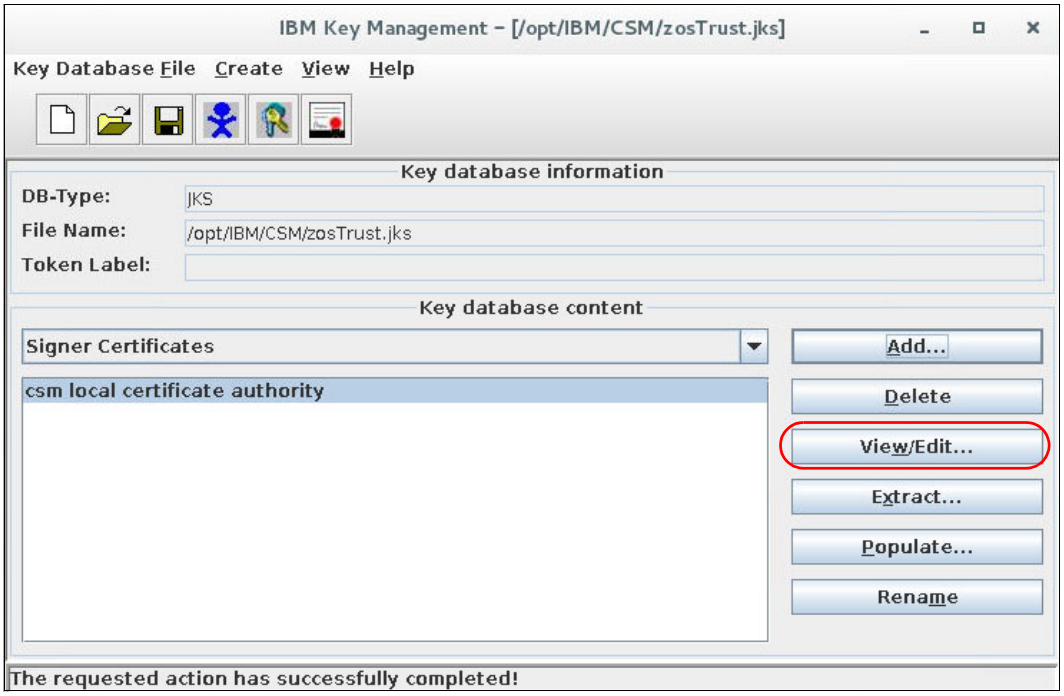

Figure 7-11 Imported CA certificate

l. Check the details by clicking the **View/Edit** button. The window shown in Figure 7-12 opens. Compare.

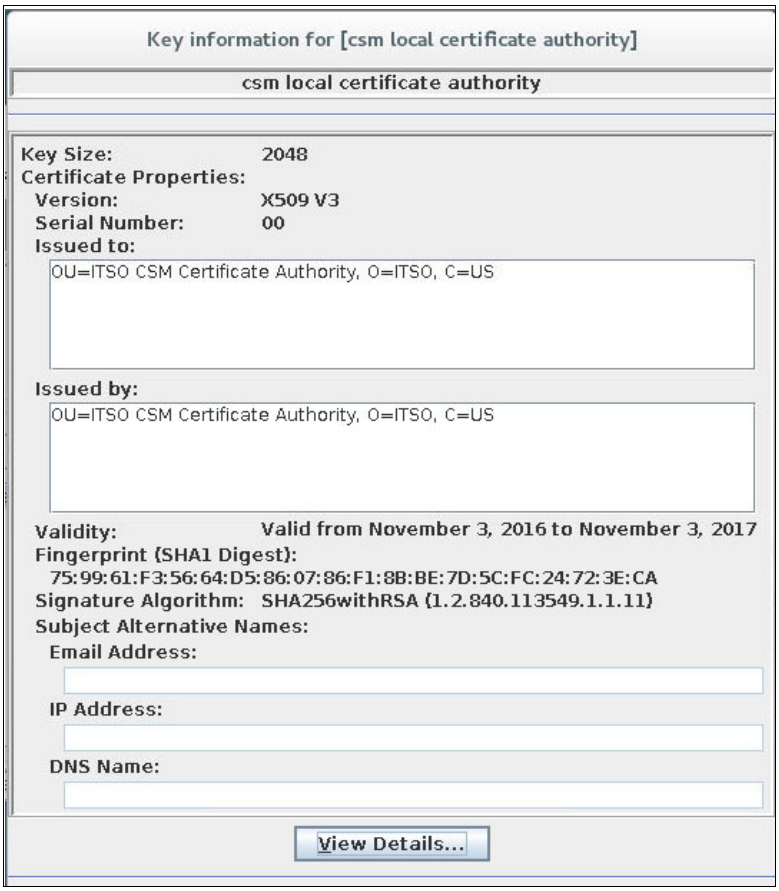

Figure 7-12 Certificate details

You can compare the certificate details with the information shown in z/OS (Example 7-9 on page 163).

2. After you close the **ikeyman** application, you can transfer (copy) the Key Database file (zosTrust.jks) to the following location on the CSM server (default install tree for CSM server for Linux on x86_64):

   ```
   /opt/IBM/CSM/liberty/wlp/usr/servers/csmServer/etc
   ```

   Check the file size and date:

   ```
   cd /opt/IBM/CSM/liberty/wlp/usr/servers/csmServer/etc/
   ls -l zosTrust.jks
   -rw-r--r-- 1 root root 978 Jan 23 19:51 zosTrust.jks
   ```

3. Restart the CSM server to pick up the configuration changes:

   ```
   /opt/IBM/CSM/stopAuth.sh
   /opt/IBM/CSM/stopCSM.sh

   /opt/IBM/CSM/startAuth.sh
   /opt/IBM/CSM/startCSM.sh
   ```

7.1.7 Configuring AT-TLS on z/OS

In preparation for enabling secure communications for the z/OS HyperSwap, we need to configure AT-TLS and the Policy Agent.

In this section, we use the z/OS Management Facility (z/OSMF) for configuring secure TCP/IP communication between the z/OS HyperSwap and the managing CSM server (CSM server running outside z/OS).

The present scenario describes how to create the Application Transparent Transport Layer Security (AT-TLS) policy file and upload it to the required z/OS systems.

For more information about AT-TLS, view the introduction to AT-TLS presentation.

Important: The assumption in this chapter is that AT-TLS is already installed on your system. Thus, the steps outlined here cover only the creation of the policy file:

- For any z/OS LPARs where you want to configure AT-TLS, you must ensure that the TCP/IP profile is updated to enable AT-TLS and that the policy agent (PAGENT) is started.
- To configure AT-TLS in the TCP/IP profile, issue the **`TCPCONFIG TTLS`** command. For more information about the `TCPCONFIG` statement or to start PAGENT, see the *z/OS Communications Server IP Configuration Reference* V2R2, SC27-3651.
- This policy was created using the IBM Configuration Assistant for z/OS Communications Server V2R3 (available as a task in z/OSMF).

Complete the following steps:

1. Log on to the z/OS Management Facility on our system, as shown in Figure 7-13 on page 171.

Note: Even though the z/OSMF runs on a particular z/OS image (and is accessed through an IP address active on that z/OS image), the AT-TLS configuration can be tailored to suit your needs by selecting the appropriate scope:

- A z/OS image (can be different from the one that runs the z/OSMF)
- A Parallel Sysplex (consisting of multiple z/OS images)
- A subset of z/OS images from a Parallel Sysplex

Understanding the scope will allow you to select the correct z/OS image to configure.

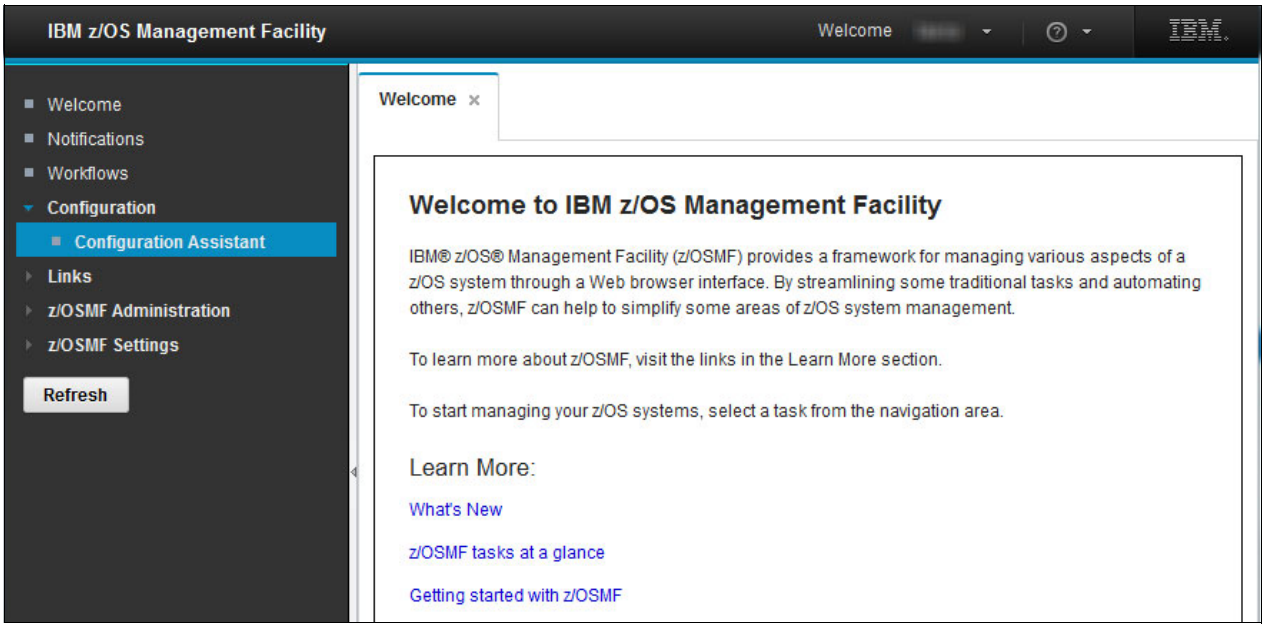

Figure 7-13 z/OSMF Welcome: z/OS Communications Server Configuration Assistant

2. Expand the “Configuration” menu and Select “Configuration Assistant”, and select an existing (or define a new backing store) as shown in Figure 7-14.

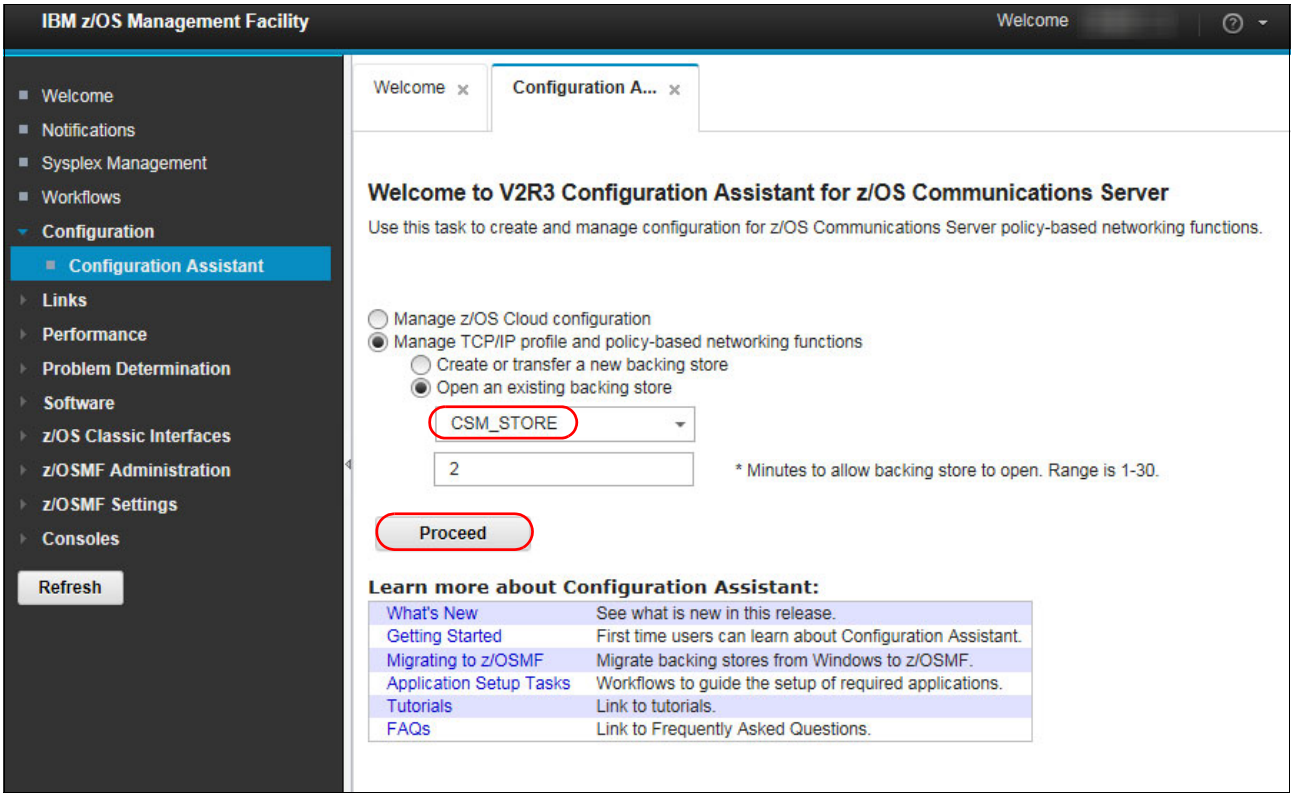

Figure 7-14 z/OSMF Configuration Assistant backing store

3. Select the Traffic Descriptors tab and create a new traffic descriptor as shown in Figure 7-15. The traffic descriptor tells AT-TLS which service to secure.

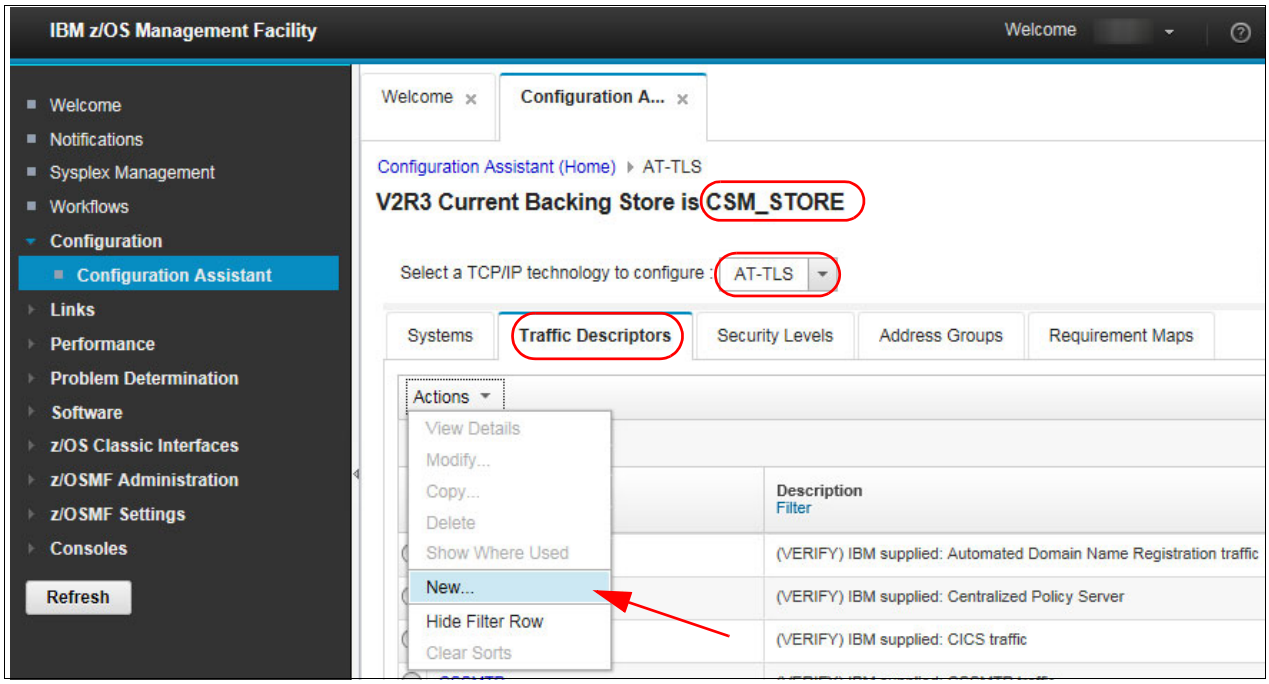

Figure 7-15 Creating a new Traffic Descriptor

4. Define a name for the Traffic Descriptor: We have chosen `HS_descriptor`, as shown in Figure 7-16.

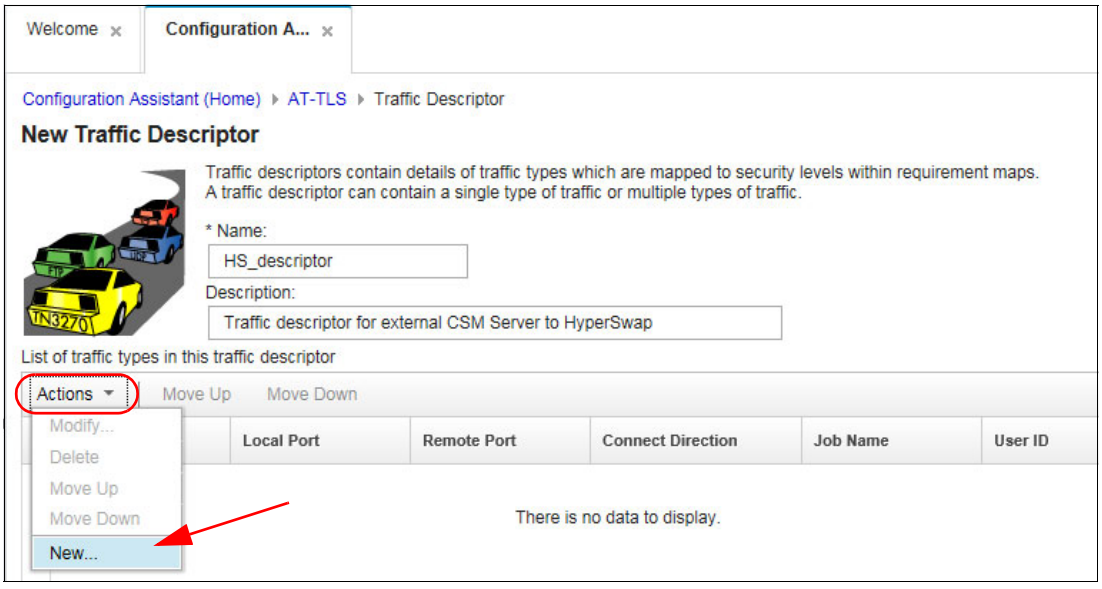

Figure 7-16 Defining the Traffic Descriptor name

5. We define the Traffic Descriptor Details (Figure 7-17 shows the TCP/IP application and protocol details). In our environment we have chosen to use port TCP 5858 for the z/OS HyperSwap Server (this is the default port the CSM server expects to talk to).

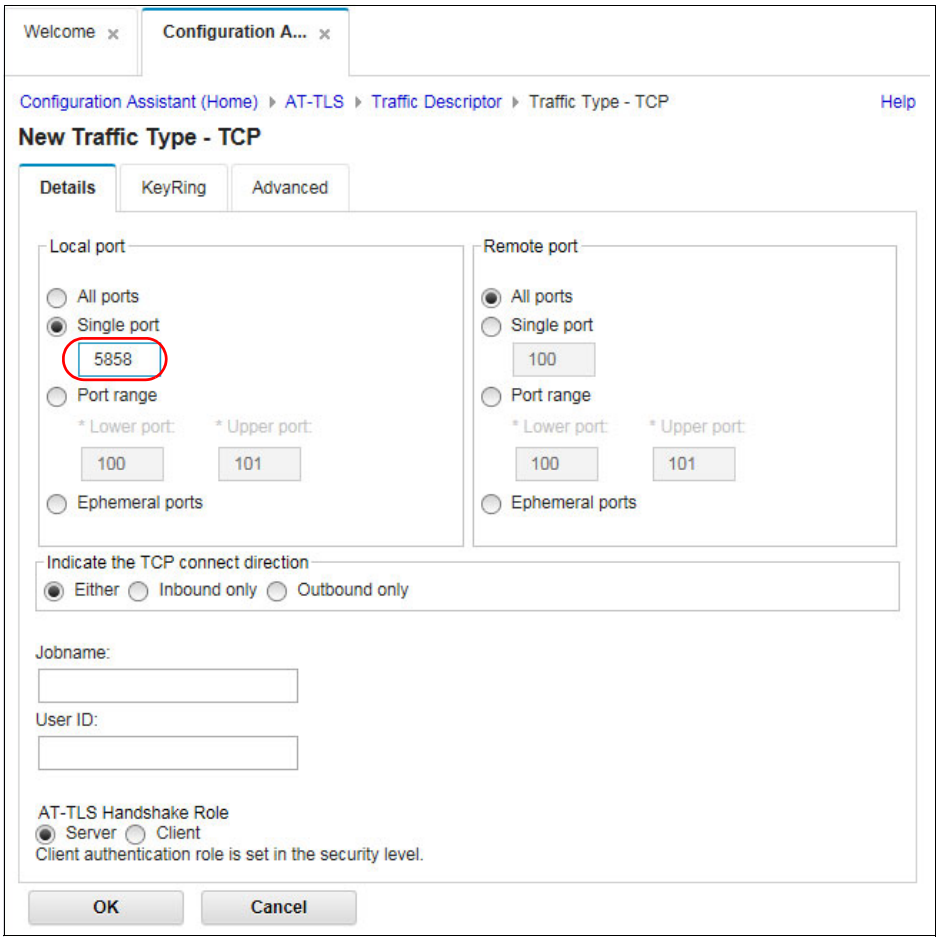

Figure 7-17 Traffic Descriptor TCP/IP details

6. We define the key ring that will be used for securing communication (AT-TLS on behalf of the application) as shown in Figure 7-18. We have chosen to define the key ring at system image level (described in a later step).

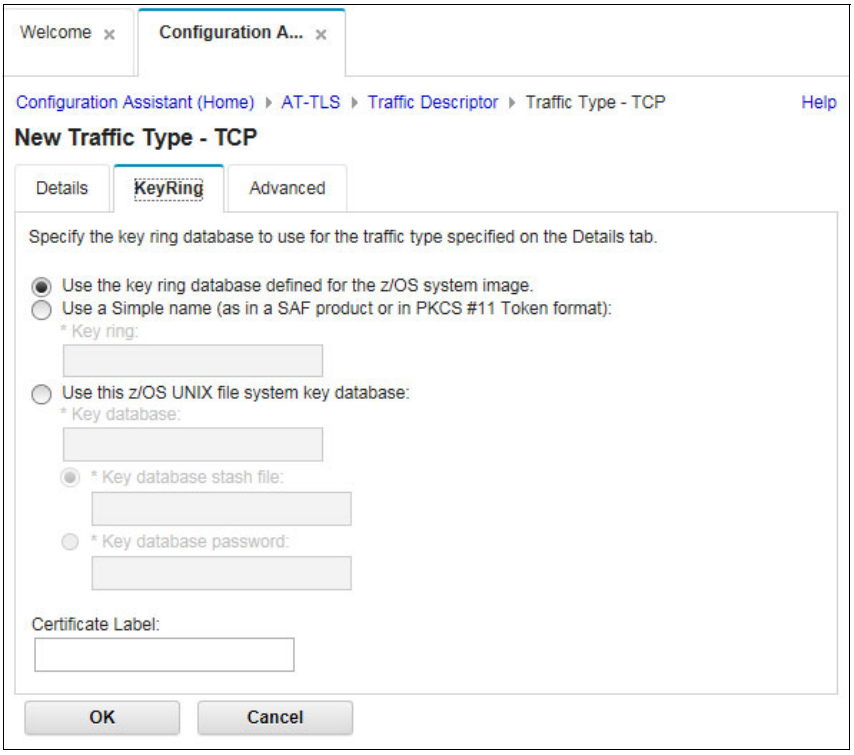

Figure 7-18 Defining the key ring for the traffic descriptor

7. The Traffic Descriptor summary is shown in Figure 7-19.

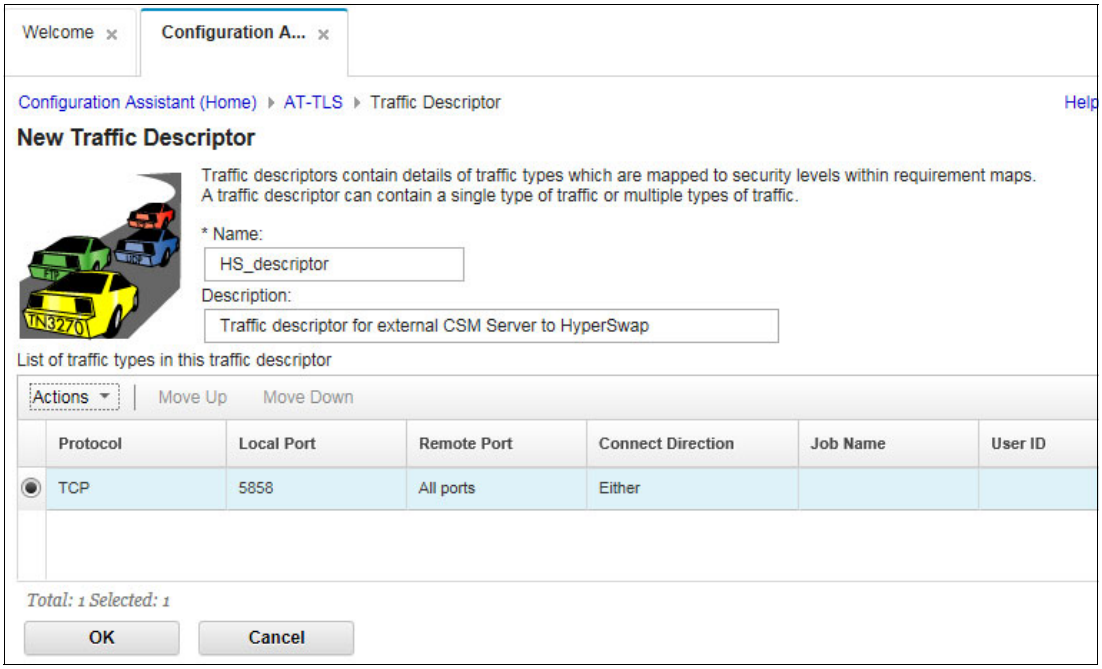

Figure 7-19 Traffic Descriptor summary

8. Figure 7-20 shows the new Traffic Descriptor listed in the CSM_STORE.

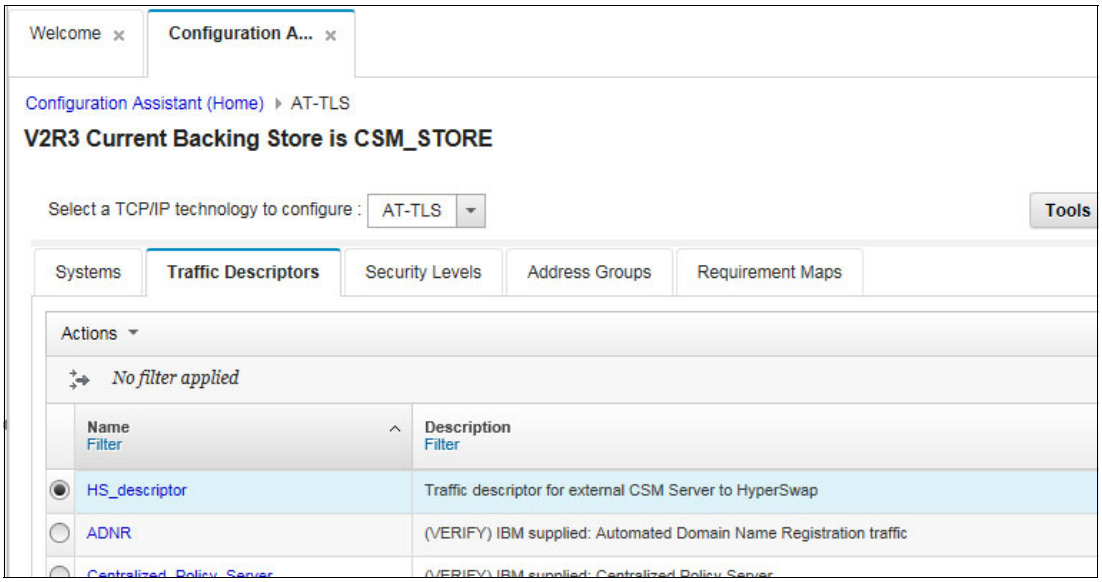

Figure 7-20 Traffic Descriptor HS_descriptor listed in the CSM_STORE

9. Next we start to configure the system(s) as shown in Figure 7-21.

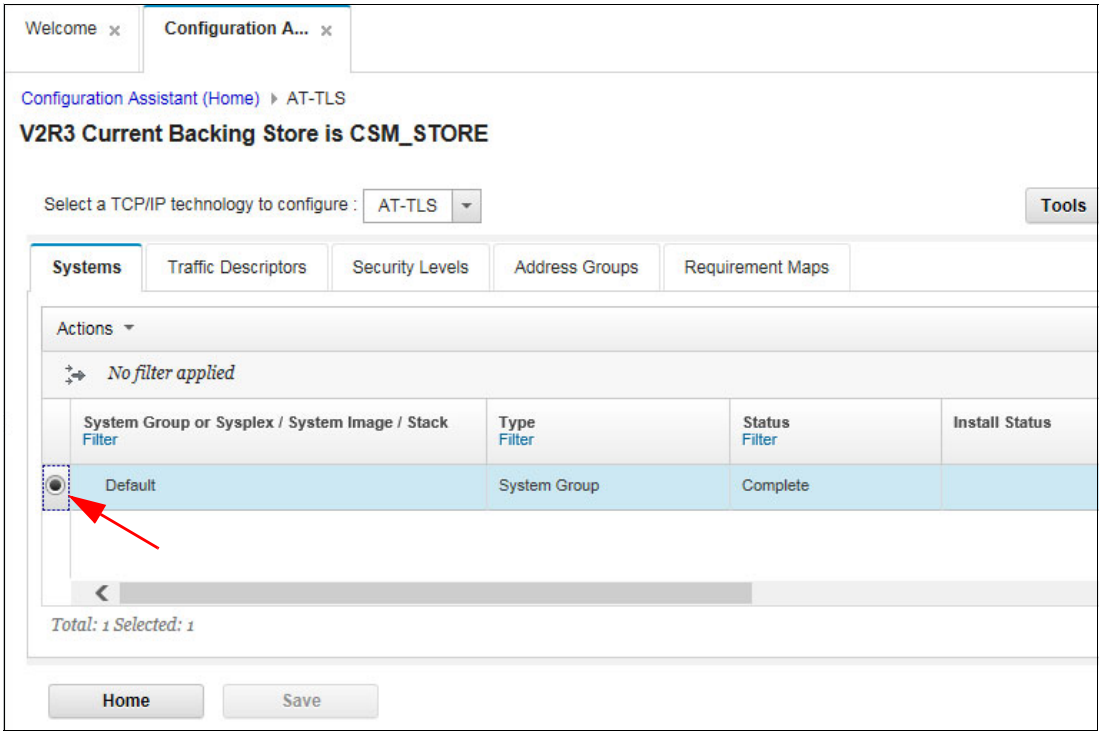

Figure 7-21 Starting Systems configuration

10. We define a *z/OS image* as shown in Figure 7-22. If you need to define a Parallel Sysplex (or a group of systems) you need to add a *z/OS Group*.

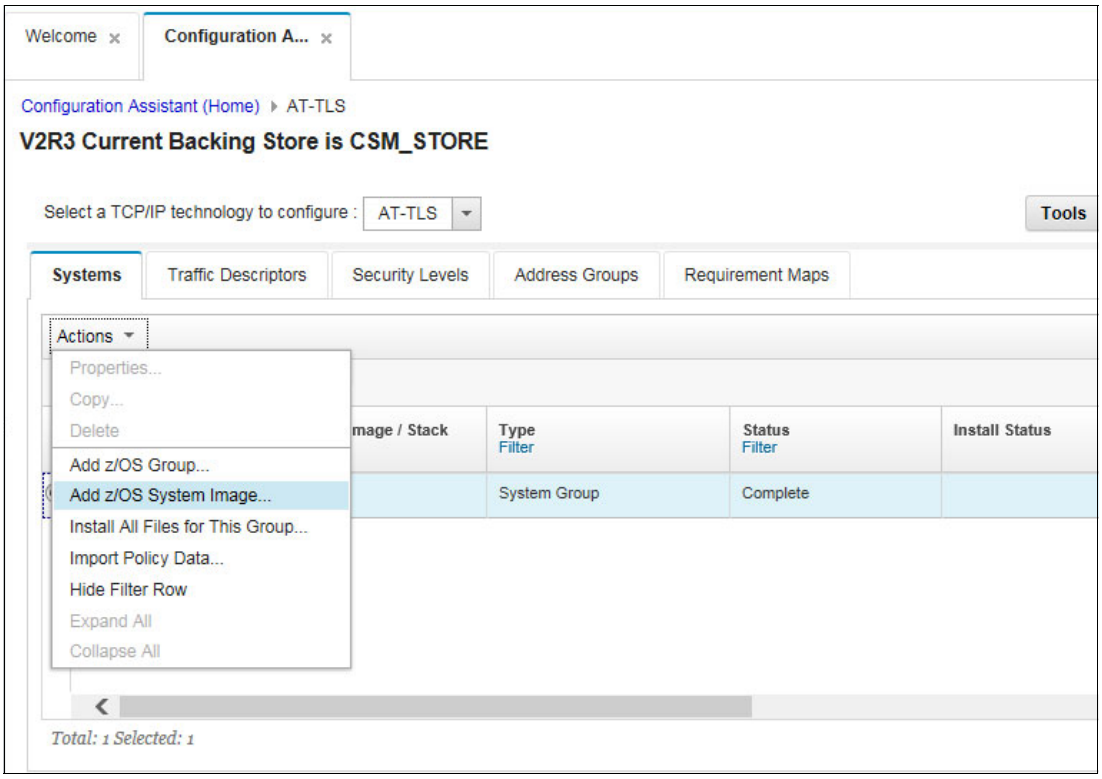

Figure 7-22 Adding a z/OS System image (single system)

11. We define the z/OS System image (running z/OS V2R3) and the key ring (to be used by AT-TLS at System image level) as shown in Figure 7-23 on page 177. When he Configuration Assistant prompts for configuring a TCP/IP Stack we proceed to the next step shown in Figure 7-24 on page 177.

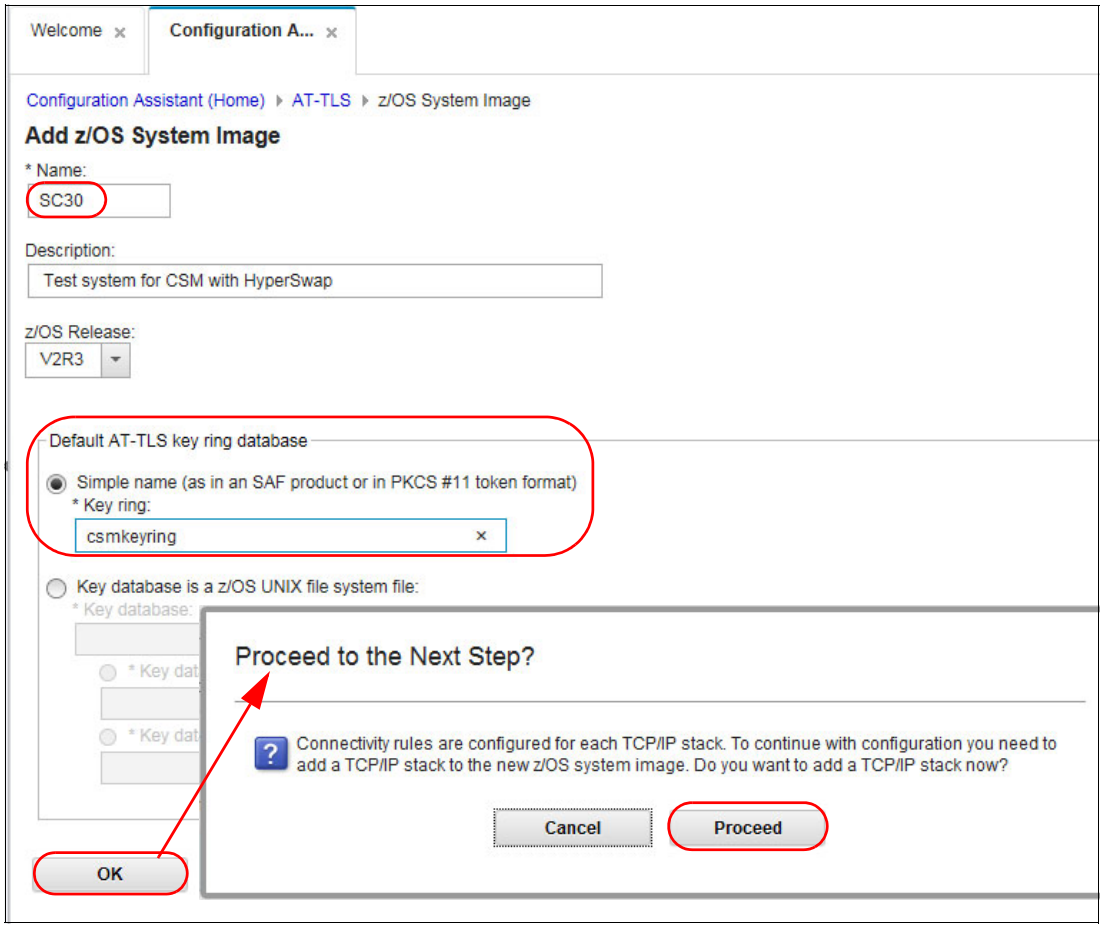

Figure 7-23 Defining system details: name, z/OS release and key ring

12. We define the TCP/IP Stack configured on our system, as shown in Figure 7-24.

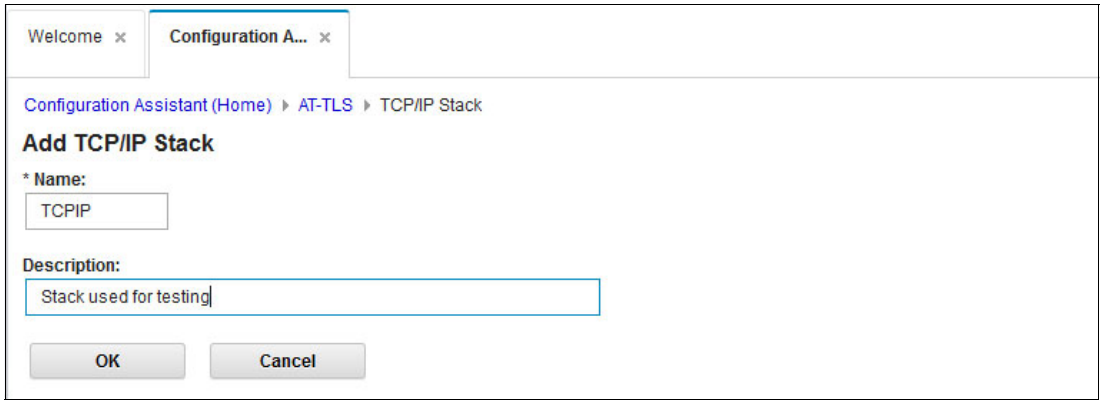

Figure 7-24 Configuring a TCP/IP Stack

13. For the TCP/IP stack previously defined we need to add connectivity rules as prompted in Figure 7-25.

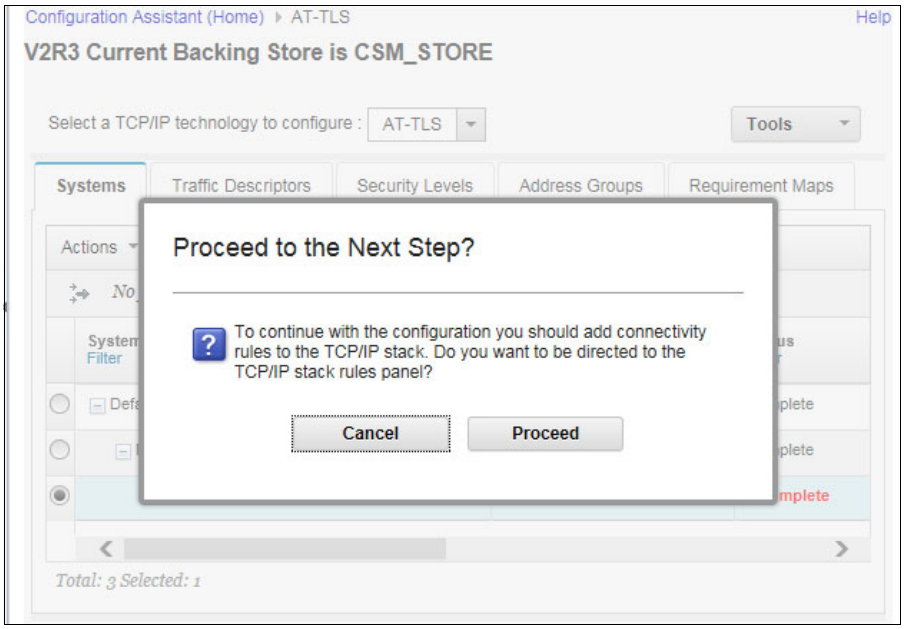

Figure 7-25 Configuration Assistant prompt for connectivity rules

14. The Configuration Assistant next prompts for starting the connectivity rules wizard, as shown in Figure 7-26.

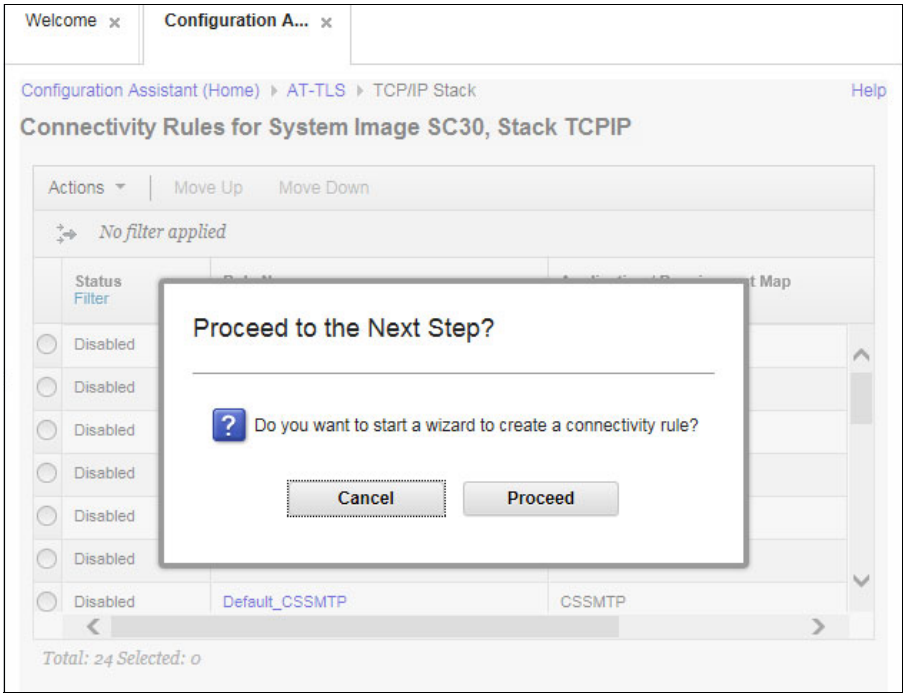

Figure 7-26 Connectivity rules wizard prompt

15. The Connectivity Rule wizard is started as shown in Figure 7-27 on page 179. We define a name for the Connectivity Rule and traffic endpoints information.

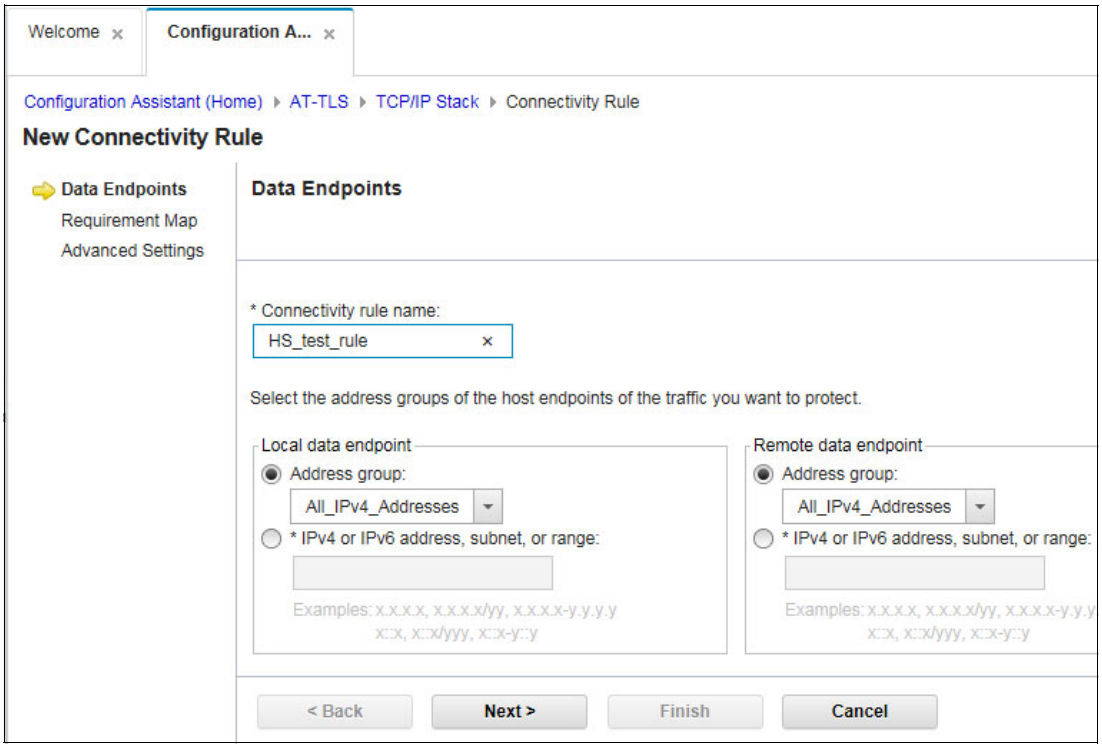

Figure 7-27 Connectivity Rule wizard

16. After defining traffic endpoints we define the traffic descriptor to be used for this rule, as shown in Figure 7-28.

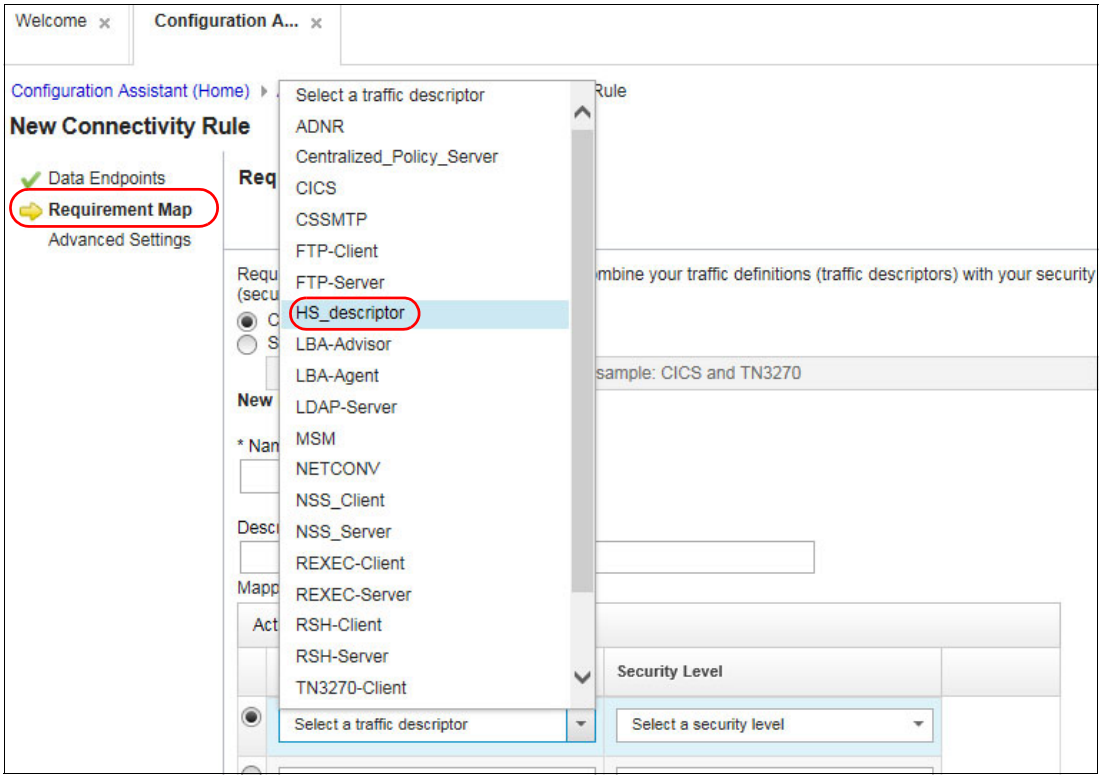

Figure 7-28 Selecting Traffic Descriptor for the Connectivity Rule

17. We enter the name for the Connectivity Rule, as shown in Figure 7-29.

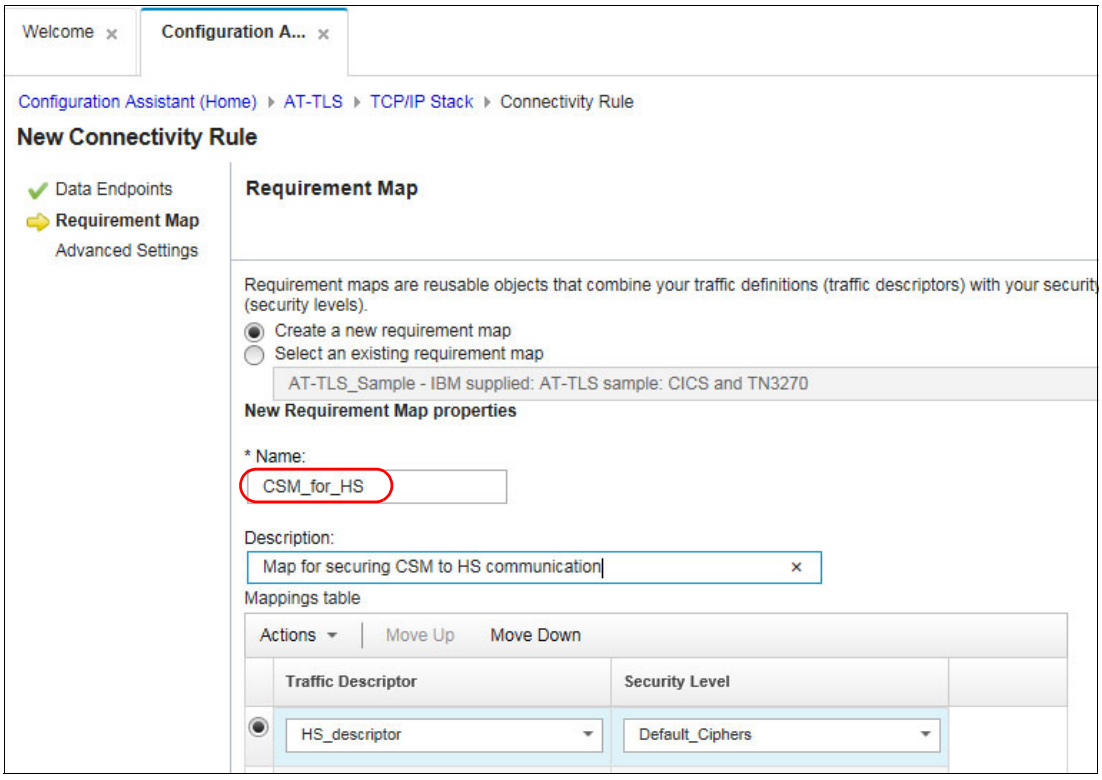

Figure 7-29 Connectivity Rule name

18. We finalize the creation of the Connectivity rule as shown in Figure 7-30.

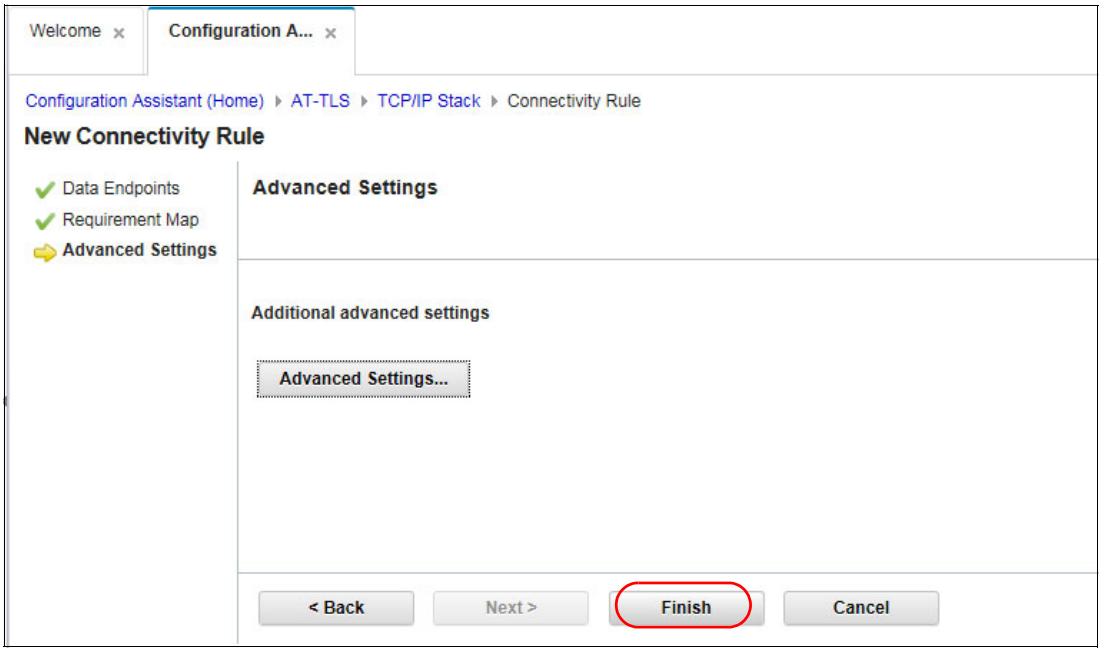

Figure 7-30 Finalizing the creation of the Connectivity Rule

19. Figure 7-31 shows the result for matching the Connectivity Rule previously defined with the TCP/IP stack belonging to SC30 system image.

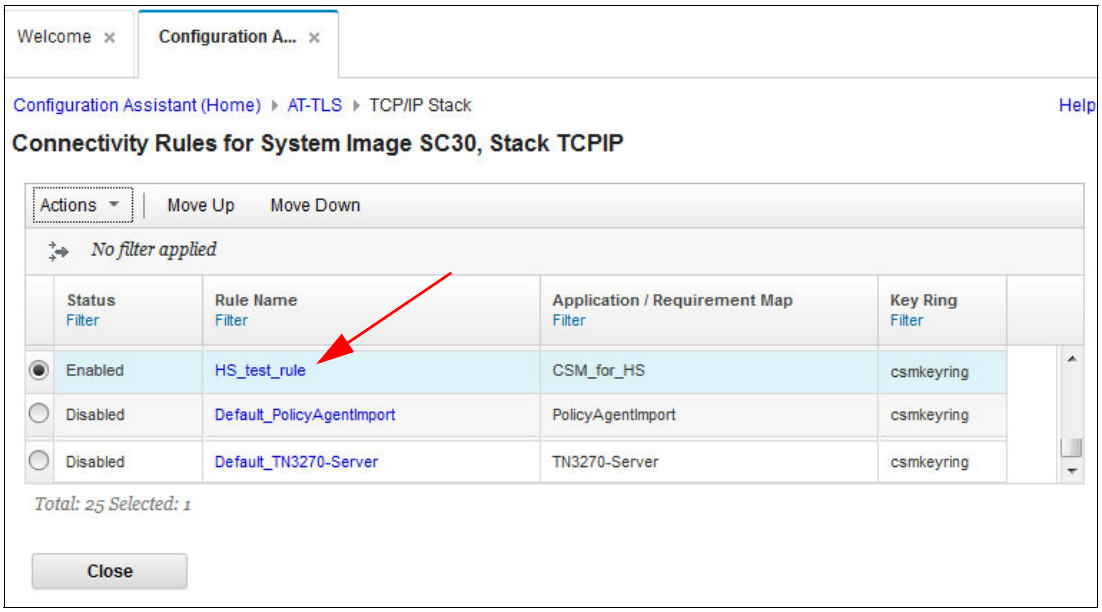

Figure 7-31 Connectivity rule added to the configuration

20. When the configuration is complete, the Policy Agent must be made aware of the changes. Figure 7-32 shows the policy configured but *not* installed yet.

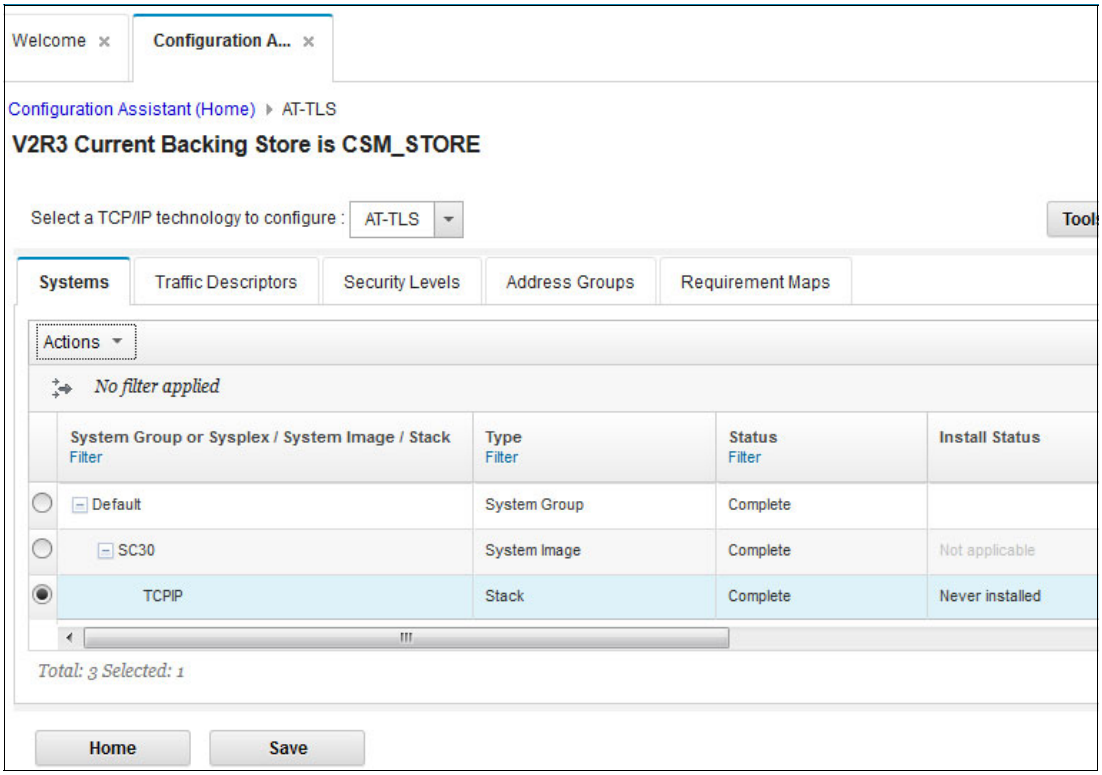

Figure 7-32 Backing store CSM_STORE before installing the policy

21. We initiate the policy files installation process as shown in Figure 7-33.

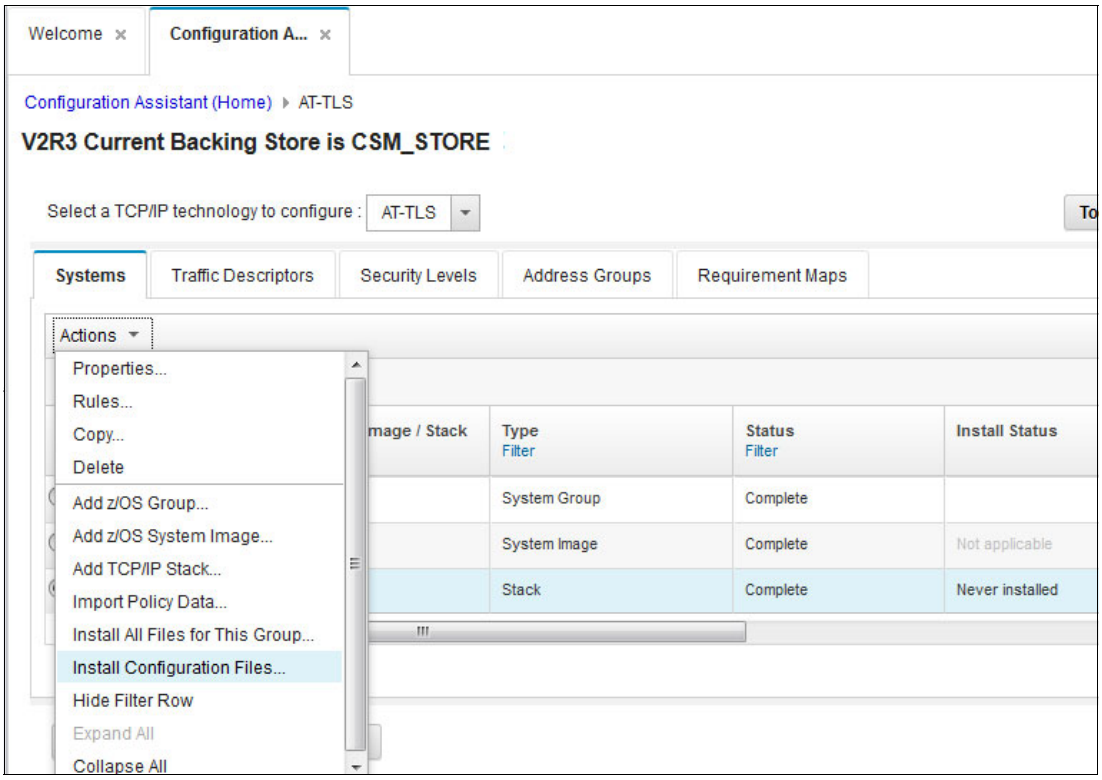

Figure 7-33 Initiating the installation of configuration files

22. Figure 7-34 shows the files that need to be installed.

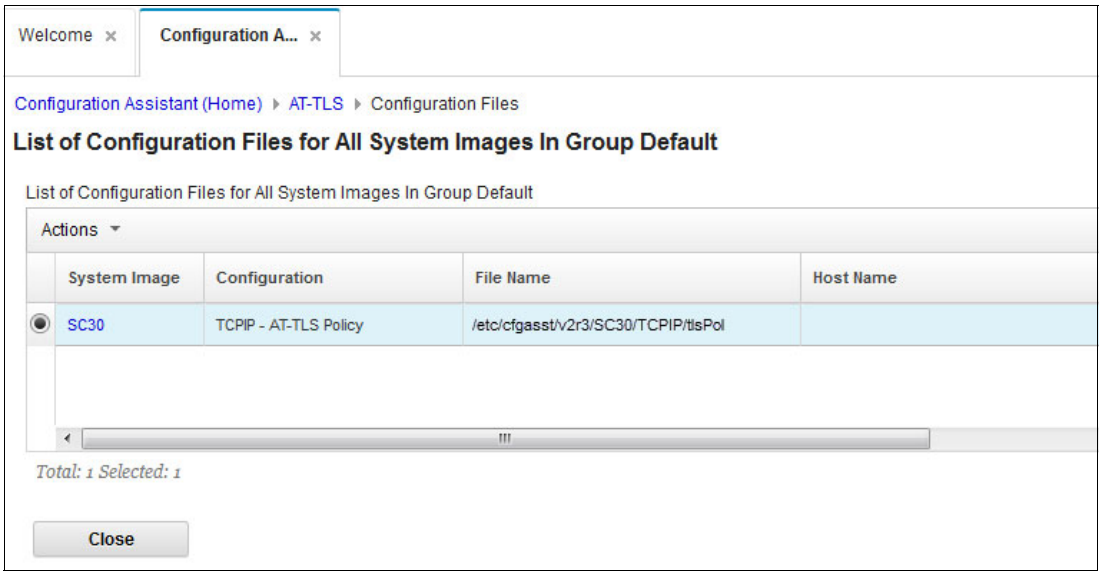

Figure 7-34 Policy files that need to be installed

23. Click **Actions** → **Install** to start the installation process, as shown in Figure 7-35.

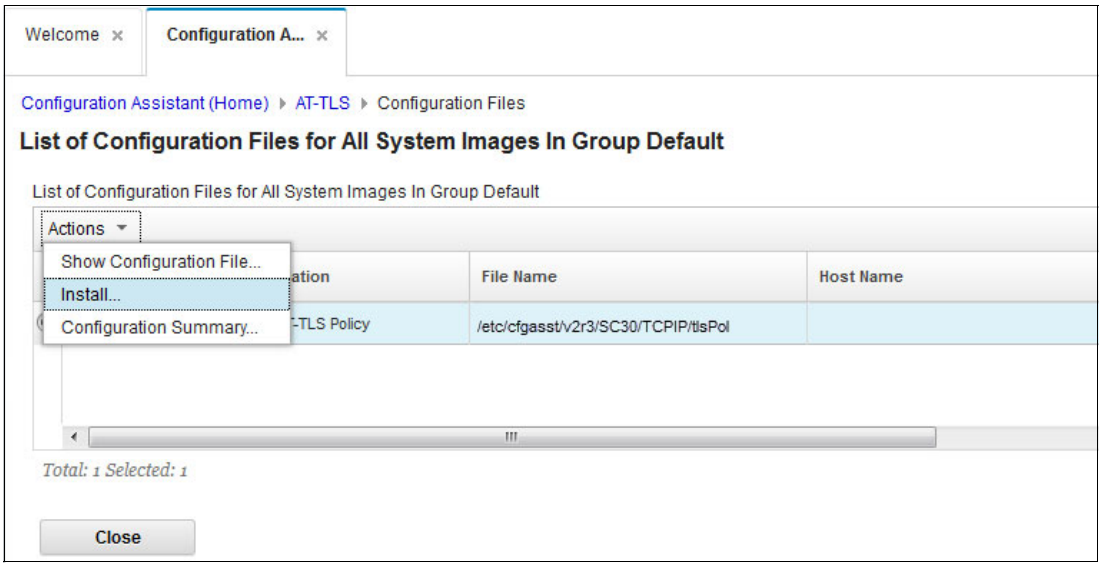

Figure 7-35 Initiating installation of files

24. We chose the location of the file on the TARGET system (SC30 in our case) and chose FTP transfer as shown in Figure 7-36.

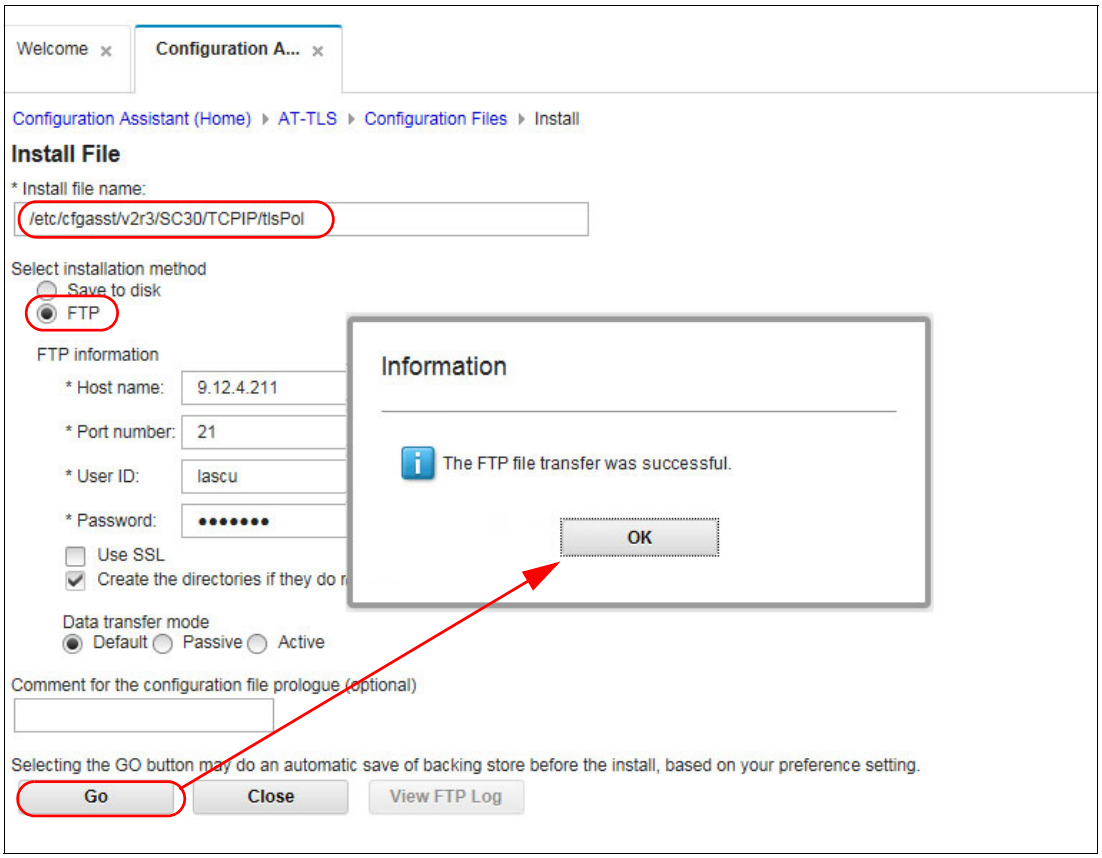

Figure 7-36 Installing file on target system via FTP

Figure 7-37 shows the installed policy.

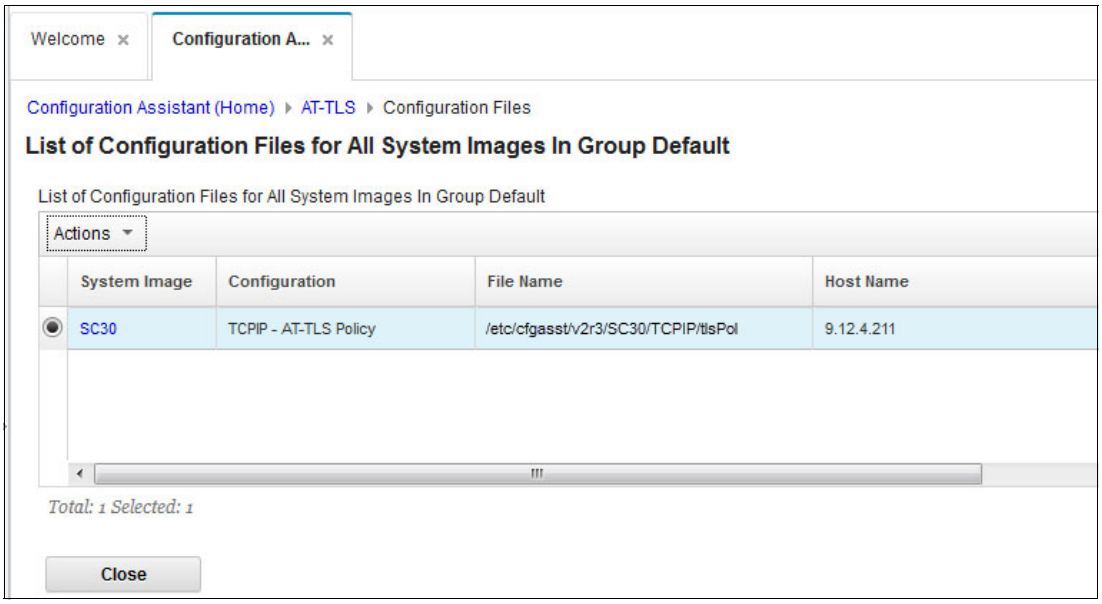

Figure 7-37 List of AT-TLS configuration files

Example 7-15 shows the content of the policy configuration file on our system (SC30), which is installed in the following location:

`/etc/cfgasst/v2r3/SC30/TCPIP/tlsPol`

Example 7-15 Installed policy file

```
##
## AT-TLS Policy Agent Configuration file for:
##      Image: SC30
##      Stack: TCPIP
##
## Created by the IBM Configuration Assistant for z/OS Communications Server
## Version 2 Release 3
## Backing Store = CSM_STORE
## Install History:
## 2017-01-24 20:40:01 : lascu to 9.12.4.211
##
## End of Configuration Assistant information
TTLSRule                                        HS_test~1
{
   LocalAddrSetRef                              addr1
   RemoteAddrSetRef                             addr1
   LocalPortRangeRef                            portR1
   Direction                                    Both
   Priority                                     255
   TTLSGroupActionRef                           gAct1~Basic_HS
   TTLSEnvironmentActionRef                     eAct1~Basic_HS
   TTLSConnectionActionRef                      cAct1~Basic_HS
}
TTLSGroupAction                                 gAct1~Basic_HS
{
```

```
   TTLSEnabled                                On
}
TTLSEnvironmentAction                         eAct1~Basic_HS
{
   HandshakeRole                              Server
   EnvironmentUserInstance                    0
   TTLSKeyringParmsRef                        keyR~SC30
}
TTLSConnectionAction                          cAct1~Basic_HS
{
   HandshakeRole                              Server
   TTLSCipherParmsRef                         cipher1~AT-TLS__Silver
   TTLSConnectionAdvancedParmsRef             cAdv1~Basic_HS
   CtraceClearText                            Off
   Trace                                      2
}
TTLSConnectionAdvancedParms                   cAdv1~Basic_HS
{
   SSLv3                                      Off
   HandshakeTimeout                           60
   SecondaryMap                               Off
   TLSv1.2                                    On
}
TTLSKeyringParms                              keyR~SC30
{
   Keyring                                    csmkeyring
}
TTLSCipherParms                               cipher1~AT-TLS__Silver
{
   V3CipherSuites                             TLS_RSA_WITH_AES_128_CBC_SHA256
   V3CipherSuites                             TLS_RSA_WITH_AES_256_CBC_SHA256
   V3CipherSuites                             TLS_RSA_WITH_AES_128_GCM_SHA256
}
IpAddrSet addr1
{
   Prefix                                     0.0.0.0/0
}
   PortRange                                  portR1
{
   Port                                       5858
}
```

The Configuration Assistant tab now shows the hierarchy for the configuration for our z/OS HyperSwap (see Figure 7-38).

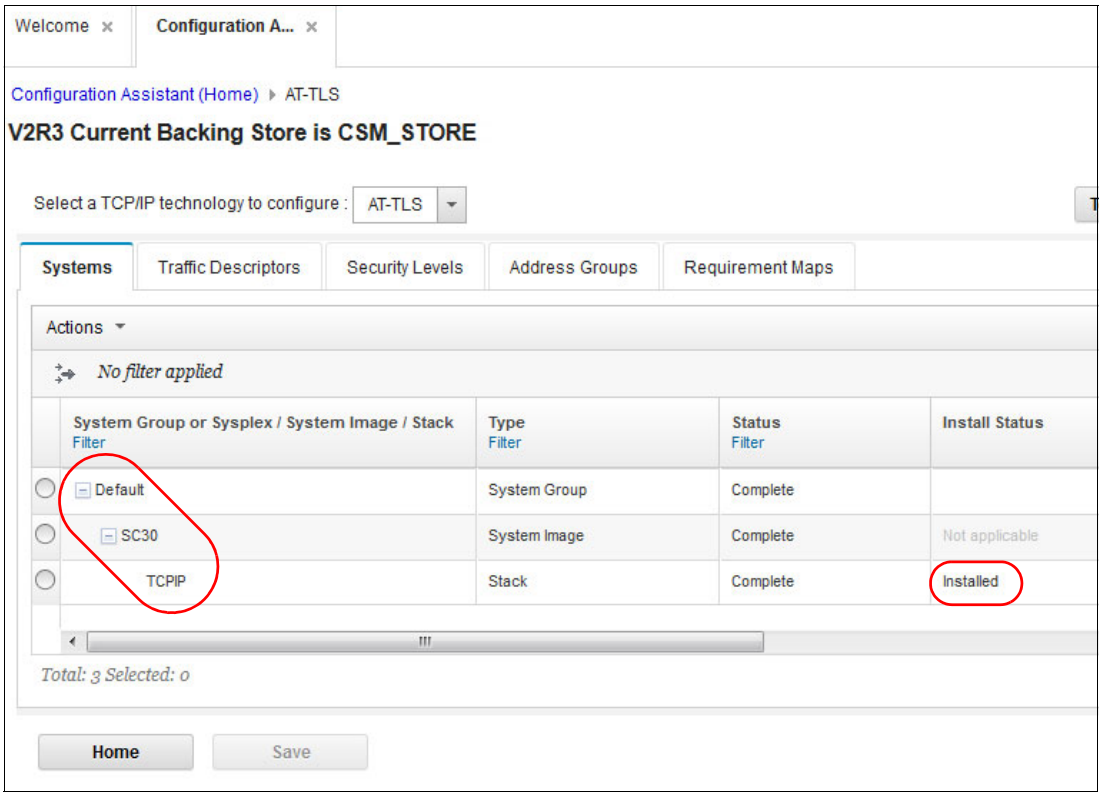

Figure 7-38 Store hierarchy

The procedure is now complete. The SC30 host can be added to the CSM Server configuration through it's IP address, HyperSwap USER ID and password.

Redbooks

IBM Copy Services Manager Implementation Guide

(0.2”spine)
0.17”<->0.473”
90<->249 pages

SG24-8375-00

ISBN DocISBN

Printed in U.S.A.