

IBM Tivoli Storage Productivity Center for Replication for
System z
Version 4.2.2.1

User's Guide



IBM Tivoli Storage Productivity Center for Replication for
System z
Version 4.2.2.1

User's Guide



Note

Before using this information and the product it supports, read the information in "Notices" on page 223.

This edition applies to version 4, release 2, modification 2, fix pack 1 of IBM Tivoli Storage Productivity Center for Replication for System z (product number 5698-B30 and 5698-B31) and to all subsequent releases and modifications until otherwise indicated in new editions.

This edition replaces SC27-2322-05.

© **Copyright IBM Corporation 2005, 2011.**

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

Contents

Figures	vii
--------------------------	------------

Tables	ix
-------------------------	-----------

About this guide **xi**

Intended audience	xi
Accessibility features for Tivoli Storage Productivity Center for Replication	xi
Accessing the IBM Tivoli Storage Productivity Center for Replication Information Center	xiv
Publications and related information for IBM Tivoli Storage Productivity Center for Replication for System z publications.	xiv
Web resources	xvi
Providing feedback about publications.	xvii

New for Tivoli Storage Productivity Center for Replication for System z 4.2.2.1 **xix**

New for Tivoli Storage Productivity Center for Replication for System z 4.2.2 **xxi**

Chapter 1. Product overview **1**

Introducing IBM Tivoli Storage Productivity Center for Replication for System z	1
Tivoli Storage Productivity Center for Replication for System z	1
IBM Tivoli Storage Productivity Center for Replication Basic Edition for System z (z/OS)	1
Tivoli Storage Productivity Center for Replication Two Site Business Continuity	2
Tivoli Storage Productivity Center for Replication Three Site Business Continuity	3
Architecture	5
Tivoli Storage Productivity Center for Replication user interfaces	5
Key concepts	6
Management servers.	6
Storage systems	8
Ports	9
Storage connections.	11
Sessions	14
Session types	20
Session commands	51
Metro Mirror heartbeat	65
Site awareness	68
Users and groups	68
User roles	69
Planning for Open HyperSwap replication	71

Chapter 2. Administering **75**

Starting and stopping IBM Tivoli Storage Productivity Center for Replication	75
Starting IBM Tivoli Storage Productivity Center for Replication	75
Stopping IBM Tivoli Storage Productivity Center for Replication	76
Starting and stopping DB2	77
Verifying that components are running	77
Verifying that IBM WebSphere Application Server is running	77
Verifying that the IBM Tivoli Storage Productivity Center for Replication server is running.	77
Verifying that DB2 is running	78
Starting the IBM Tivoli Storage Productivity Center for Replication GUI.	78
Identifying the version of IBM Tivoli Storage Productivity Center for Replication	79
Backing up and restoring IBM Tivoli Storage Productivity Center for Replication configuration data	79
Back up and recovery	79
Backing up the Tivoli Storage Productivity Center for Replication database	81
Restoring the IBM Tivoli Storage Productivity Center for Replication database.	81
Exporting copy set data	82
Importing copy set data	83

Chapter 3. Managing management servers **85**

Management servers	85
Ports	86
SNMP alerts	88
Session state change SNMP trap descriptions	88
Configuration change SNMP trap descriptions.	89
Suspending-event notification SNMP trap descriptions	89
Communication-failure SNMP trap descriptions	90
Management Servers state-change SNMP trap descriptions	90
Setting up a standby management server	91
Setting the local management server as the standby server	91
Setting a remote management server as the standby server	91
Reinstalling the primary server during an active session	92
Reconnecting the active and standby management servers	93
Performing a takeover on the standby management server	93
Configuring SNMP.	94
Adding SNMP managers	94

Changing the standby management server port number	94
Changing the client port number	94
Changing the time zone in z/OS	95

Chapter 4. Managing storage systems 97

Ports	97
Storage systems	99
Storage connections.	99
Direct connection	101
Hardware Management Console connection	101
z/OS connections	102
Protected volumes.	103
Site awareness	104
Adding a storage connection	104
Removing a storage connection	106
Removing a storage system.	106
Modifying the location of storage systems	107
Modifying storage connection properties	107
Refreshing the storage system configuration	108
Setting volume protection	108
Restoring data from a journal volume	109

Chapter 5. Managing host systems 111

Adding a host system connection.	111
Modifying a host system connection.	112
Removing a host system connection	112
Removing a session from a host system connection	112

Chapter 6. Managing logical paths . . 115

Viewing logical paths.	115
Adding logical paths	115
Adding logical paths using a CSV file	116
Removing logical paths	117

Chapter 7. Setting up data replication 119

Sessions	119
Copy sets.	119
Volume roles	123
Role pairs	124
Practice volumes	124
Consistency groups	124
Session types	125
Basic HyperSwap (ESS, DS6000, and DS8000)	130
FlashCopy	131
Snapshot	132
Metro Mirror	133
Global Mirror	139
Metro Global Mirror (ESS 800 and DS8000)	144
Managing a session with HyperSwap and Open HyperSwap replication	147
Session commands	156
Basic HyperSwap commands	156
FlashCopy commands	157
Snapshot commands	157
Metro Mirror commands	158
Metro Mirror with Practice commands	160
Global Mirror commands	161
Global Mirror with Practice commands	162
Metro Global Mirror commands	163

Metro Global Mirror with Practice commands	166
Site awareness	169
Preserve Mirror option	170
Creating sessions and adding copy sets.	172
Creating a FlashCopy session and adding copy sets.	172
Creating a Snapshot session and adding copy sets.	173
Creating a Metro Mirror session and adding copy sets	174
Creating a Global Mirror session and adding copy sets	176
Creating a Metro Global Mirror session and adding copy sets	177
Using the Metro Mirror heartbeat	178
Metro Mirror heartbeat	178
Enabling and disabling the Metro Mirror heartbeat	181
Exporting copy set data	181
Importing copy set data	181
Modifying the location of session sites	182
Removing sessions	182
Removing copy sets	182
Migrating an existing configuration to Tivoli Storage Productivity Center for Replication	183
Metro Mirror	184
Global Mirror	184
Assimilating Metro Mirror pairs into a Three Site Metro Global Mirror session	185
Assimilating Three Site pairs into a Three Site Metro Global Mirror session	185
Global Mirror and Metro Mirror assimilation for SAN Volume Controller, Storwize V7000, Storwize V7000 Unified or the XIV system	186

Chapter 8. Practicing disaster recovery 187

Practice volumes	187
Practicing disaster recovery for a Metro Mirror Failover/Failback with Practice session.	187
Practicing disaster recovery for a Global Mirror either Direction with two-site Practice session	188
Practicing disaster recovery for a Global Mirror Failover/Failback with Practice session.	188
Practicing disaster recovery for a Metro Global Mirror Failover/Failback with Practice session	189

Chapter 9. Monitoring health and status. 191

Viewing the health summary	191
Viewing SNMP alerts.	191
Viewing sessions	191
Session status icons	191
Session images	192
Session states	194
Session properties	197
Role pair status and progress	208
Viewing session properties	209
Viewing session details	209
Viewing storage system details	210

Viewing storage connection details	211
Viewing volume details	211
Viewing logical paths.	211
Viewing console messages	212

Chapter 10. Security 213

Users and groups	213
User roles	214
Adding the IBM Tivoli Storage Productivity Center for Replication Administrator role to the IBM Tivoli Storage Productivity Center Superuser group.	216
Granting access privileges for a user.	216
Viewing access privileges for a user	217
Modifying access privileges for a user	217
Removing access privileges for a user	218

**Appendix. Using the system logger in
a Tivoli Storage Productivity Center
for Replication for System z
environment 219**

Configuring the system logger for use in the Tivoli Storage Productivity Center for Replication for System z environment	219
Reintroducing frozen system logger CDSs into your sysplex.	221

Notices 223

Trademarks	225
----------------------	-----

Index 227

Figures

Tables

1. Storage system default ports	10	20. Storage system default ports	98
2. Support number of role pairs and volumes per copy set for each session type	15	21. Support number of role pairs and volumes per copy set for each session type.	120
3. Session type summary	21	22. Session type summary	125
4. Basic HyperSwap commands.	52	23. Basic HyperSwap commands	156
5. FlashCopy commands	52	24. FlashCopy commands.	157
6. Snapshot session commands	53	25. Snapshot session commands	157
7. Snapshot group commands	53	26. Snapshot group commands	158
8. Metro Mirror commands	54	27. Metro Mirror commands.	158
9. Metro Mirror with Practice commands	56	28. Metro Mirror with Practice commands	160
10. Global Mirror commands	57	29. Global Mirror commands	161
11. Global Mirror with Practice commands	58	30. Global Mirror with Practice commands	162
12. Metro Global Mirror commands.	59	31. Metro Global Mirror commands	164
13. Metro Global Mirror with Practice commands	62	32. Metro Global Mirror with Practice commands	166
14. Storage system default ports	87	33. Session status icons	191
15. Session state change traps.	88	34. Volume role symbols	192
16. Configuration change traps	89	35. Data copying symbols	193
17. Suspending-event notification traps	90	36. Session states	194
18. Communication-failure traps.	90	37. Detailed status messages for Participating and Non-Participating role pairs.	208
19. Management Servers state-change traps	90		

About this guide

This guide provides information about using the IBM® Tivoli® Storage Productivity Center for Replication family. All Tivoli Storage Productivity Center for Replication solutions provide continuous availability and disaster recovery solutions by using the following replication methods:

- Point-in-time replication, which includes FlashCopy®
- Continuous replication, which includes Metro Mirror, Global Mirror, and Metro Global Mirror

This product is available in the following versions:

- IBM Tivoli Storage Productivity Center for Replication Two Site Business Continuity
- IBM Tivoli Storage Productivity Center for Replication Three Site Business Continuity
- IBM Tivoli Storage Productivity Center for Replication for System z®
- IBM Tivoli Storage Productivity Center for Replication Basic Edition for System z

Intended audience

This publication is intended for users of IBM Tivoli Storage Productivity Center for Replication.

Accessibility features for Tivoli Storage Productivity Center for Replication

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

The following list includes the major accessibility features in Tivoli Storage Productivity Center for Replication:

- Keyboard-only operation
- Interfaces that are commonly used by screen readers
- Keys that are discernible by touch but do not activate just by touching them
- Industry-standard devices for ports and connectors
- The attachment of alternative input and output devices

See the IBM Human Ability and Accessibility Center website at www.ibm.com/able for more information about the commitment that IBM has for accessibility.

Accessibility and keyboard shortcuts in the information center

Accessibility features help users with physical disabilities, such as restricted mobility or limited vision, to use software products successfully. Using the major accessibility features in this product, users can perform these tasks:

- Use assistive technologies, such as screen-reader software and digital speech synthesizer, to hear what is displayed on the screen. Consult the product documentation of the assistive technology for details on using those technologies with this product.

- Operate specific or equivalent features by using only the keyboard.
- Magnify what is displayed on the screen.

In addition, the documentation was modified to include the following features to aid accessibility:

- All documentation is available in HTML formats to give the maximum opportunity for users to apply screen-reader software technology.
- All images in the documentation are provided with alternative text so that users with vision impairments can understand the contents of the images.

Use the following key combinations to navigate the interface by keyboard:

- To go directly to the Topic pane (the right side), press Alt+K, and then press Tab.
- In the Topic pane, to go to the next link, press Tab.
- To go directly to the Search Results view in the left side, press Alt+R, and then press Enter or Up Arrow to enter the view.
- To go directly to the Navigation (Table of Contents) view in the left side, press Alt+C, and then press Enter or Up Arrow to enter the view.
- To expand and collapse a node in the navigation tree, press the Right and Left Arrow.
- To move to the next topic node, press the Down Arrow or Tab.
- To move to the previous topic node, press the Up Arrow or Shift+Tab.
- To go to the next link, button, or topic node from inside on of the views, press Tab.
- To scroll all the way up or down in a pane, press Home or End.
- To go back, press Alt+Left Arrow; to go forward, press Alt+Right Arrow.
- To go to the next pane, press F6.
- To move to the previous pane, press Shift+F6.
- To print the active pane, press Ctrl+P.

Related accessibility information for sight-impaired users

The following list contains hints and tips that can help you more fully use the graphical user interface:

Drop-down lists are positioned directly above or before the radio button that activates it.

If you use a screen reader, you should be aware that there are radio buttons to activate drop-down lists for several GUI pages. The way to activate the drop-down list is by selecting the associated radio button. The drop-down list is positioned directly above or before the radio button that activates it. When you use a screen reader that processes the fields and controls of a page sequentially, you might select the radio button, but not know that the associated drop-down list has been activated. The screen reader processes inactive drop-down lists first, and then processes the next radio button. The drop-down list is activated if you select the radio button.

On the following pages, keep in mind that radio buttons activate a drop-down list:

- Administration
- ESS/DS Paths
- Sessions
- Session Details

- Storage Systems

Tables are best understood by reviewing the surrounding text and the table row and column number of the table.

On some graphical user pages, tables use the header or row ID attributes when reading a single cell. The screen reader reads the table row and column number, along with cell data. Therefore, you can infer the column header and row ID.

Experiment with and fine-tune the way your screen reader pronounces some of the product abbreviations.

Your screen reader might pronounce abbreviations as if they were words. For example, the common abbreviation for Enterprise Storage Server[®] is ESS. Your screen reader might read ESS as the word "ess". With some screen readers you can hear alternate pronunciations. If you frequently use the software you might prefer to fine-tune such associations in your settings. When an association is created, the screen reader can recognize the abbreviation as a word. If you can add dictionary words with your screen reader, replace the capitalized character sequence with the sequence E space S space S.

Typically, this abbreviation is used in the combination form of ESS/DS. This term refers to the Enterprise Storage Server 800, the DS6000[™], or the DS8000[®].

Some decorative artifacts might persist if the cascading style sheet is disabled.

Enable cascading style sheets when possible; otherwise, some decorative elements might persist in the Firefox and Internet Explorer GUIs. These artifacts do not affect performance. If they become too distracting, consider using the command-line interface instead.

For efficiency, confirmation dialogs place initial focus on the Yes button.

When a confirmation dialog box is displayed, focus is given to the **Yes** button. Therefore, the screen reader reads "Yes" but does not read the confirmation text. The software processes the information in this way when you do the following types of tasks:

- Perform an action on a session
- Remove a connection to a storage system
- Click the **About link**
- Create a high-availability connection

To read the confirmation text before clicking the **Yes**, **No**, or **OK** button, view the previous heading before the button.

Dojo components are not read by all screen readers.

The Job Access for Windows and Speech (JAWS) screen reader does not read some Dojo components on Internet Explorer 7. Use the command-line interface instead of the GUI with JAWS on Internet Explorer 7.

Firefox is the preferred browser for use with a screen reader.

Use Firefox as the screen reader because other browsers might not fully expose assistive technology content to the screen reader.

Accessing the IBM Tivoli Storage Productivity Center for Replication Information Center

This topic explains how to access the IBM Tivoli Storage Productivity Center for Replication Information Center.

You can access the information center in the following ways:

- On the publications CD, a readme.txt file describes how to start the information center depending on platform and mode.
- The IBM Tivoli Storage Productivity Center for Replication graphical user interface includes a link to the information center.
- Go to the Web at <http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp>:

Publications and related information for IBM Tivoli Storage Productivity Center for Replication for System z publications

This topic lists the publications in the IBM Tivoli Storage Productivity Center for Replication library and other related publications.

Information Centers

You can browse product documentation in the IBM Tivoli Storage Productivity Center for Replication for System z Information Center at:

<http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp>

Publications

The IBM Publications Center website offers customized search functions to help you find the publications that you need. Some publications are available for you to view or download free of charge. You can also order publications. The publications center displays prices in your local currency. You can access the IBM Publications Center on the web at www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss

The IBM Publications Center website also offers you a notification system for IBM publications. Register and you can create your own profile of publications that interest you. The publications notification system sends you a daily email that contains information about new or revised publications that are based on your profile. Access the publications notification system from the IBM Publications Center on the web at www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss to subscribe.

The following publications make up the IBM Tivoli Storage Productivity Center for Replication for System z library:

IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide

This guide contains instructions for installing and configuring the product on z/OS®.

Program Directory for IBM Tivoli Storage Productivity Center for Replication Basic Edition for System z

This Program Directory includes installation instructions associated with IBM Tivoli Storage Productivity Center for Replication Basic Edition for System z.

Program Directory for IBM Tivoli Storage Productivity Center for Replication for System z

This Program Directory presents information concerning the material and procedures associated with the installation of IBM Tivoli Storage Productivity Center for Replication for System z.

Program Directory for IBM WebSphere® Application Server for z/OS

This Program Directory presents information related to installing IBM WebSphere Application Server for z/OS V6.1.0.

Program Directory for IBM WebSphere Application Server OEM Edition

This Program Directory presents information related to installing IBM WebSphere Application Server OEM Edition for z/OS V6.1.0.

IBM WebSphere Application Server OEM Edition for z/OS Configuration Guide

This guide contains configuration instructions associated with IBM WebSphere Application Server OEM Edition for z/OS.

IBM Tivoli Storage Productivity Center for Replication for System z User's Guide

This guide contains task-oriented instructions for using the product graphical user interface (GUI) to manage copy services.

IBM Tivoli Storage Productivity Center for Replication for System z Command-Line Interface User's Guide

This guide provides information about how to use the product's command-line interface (CLI).

IBM Tivoli Storage Productivity Center for Replication for System z Problem Determination Guide (GC27-2320)

This guide assists administrators or users who are troubleshooting problems with the product.

WebSphere Application Server for z/OS product website

This website provides information about WebSphere Application Server for z/OS, including links to sources of related information such as redbooks, white papers, and ebooks. To view the website, go to http://www-01.ibm.com/software/webservers/appserv/zos_os390/.

Redbooks and white papers

Performance Monitoring and Best Practices for WebSphere on z/OS (SG24-7269)

This IBM Redbooks® publication provides a structure that you can use to set up an environment that is tuned to meet best performance and catch eventual performance bottlenecks.

DB2® for z/OS and WebSphere: The Perfect Couple (SG24-6319)

This IBM Redbooks publication provides a broad understanding of the installation, configuration, and use of the IBM DB2 Universal Driver for SQLJ and JDBC in a DB2 for z/OS and OS/390® Version 7, and DB2 for z/OS Version 8 environment, with IBM WebSphere Application Server for z/OS for z/OS Version 5.02. It describes both type 2 and type 4 connectivity (including the XA transaction support) from a WebSphere Application Server on z/OS to a DB2 for z/OS and OS/390 database server.

Web resources

Listed here are the websites and information center topics that relate to IBM Tivoli Storage Productivity Center for Replication.

Websites

- IBM Tivoli Storage Productivity Center
www.ibm.com/systems/storage/software/center/standard/index.html
This website describes the feature, benefits, and specifications of Tivoli Storage Productivity Center. It also provides a link to product support, data sheets, resource library, and white papers.
- Tivoli Storage Productivity Center for Replication
www.ibm.com/systems/storage/software/center/replication/index.html
This website describes the feature, benefits, and specifications of Tivoli Storage Productivity Center for Replication. It also provides a link to the Software Online Catalog to purchase the product and licenses.
- Tivoli Storage Productivity Center Technical Support
www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Productivity_Center_Standard_Edition
This website provides links to downloads and documentation for all currently supported versions of Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication.
- Supported Storage Products List
<http://www-01.ibm.com/support/docview.wss?uid=swg21386446>
This website provides links to the supported storage products for each version of Tivoli Storage Productivity Center for Replication.
- IBM WebSphere Application Server
www.ibm.com/software/webservers/appserv/was/
This website describes the IBM WebSphere Application Server offerings and provides links for downloading a trial version, purchasing IBM WebSphere Application Server, and viewing online publications and demos.
- IBM DB2 Software
www.ibm.com/software/data/db2/
This website describes the DB2 offerings and provides links for downloading a trial version, purchasing DB2, and viewing analyst reports, online publications, and demos.
- IBM System Storage® Disk Systems
www.ibm.com/servers/storage/disk/
This website provides links to learn more about the IBM System Storage disk systems products and offerings, including DS6000 and DS8000. It also provides links for viewing support and services, software and solutions, and other resources.
- IBM System Storage SAN Volume Controller
www.ibm.com/servers/storage/software/virtualization/svc/index.html
This website describes the IBM System Storage SAN Volume Controller offering and provides links for requesting a quote for and purchasing System Storage SAN Volume Controller and viewing online publications, white papers, and case studies.
- IBM Storwize V7000

www.ibm.com/systems/storage/disk/storwize_v7000/index.html

This website describes the Storwize® V7000 offerings and provides links for requesting a quote and viewing online publications and white papers.

- IBM XIV Storage System

www.ibm.com/systems/storage/disk/xiv

This website describes the XIV® system offering and provides links for requesting a quote for an XIV system and viewing online publications, white papers, and demos.

- System z (and z/OS)

www.ibm.com/systems/z/

This website provides links to learn more about IBM System z offerings and software. It also includes information about upcoming webcasts, blogs, and demos.

Forums

- Tivoli Forums

www.ibm.com/developerworks/forums/tivoli_forums.jspa

This website provides a forum that you can use to discuss issues pertaining to Tivoli Storage Productivity Center, Tivoli Storage Productivity Center for Replication, and other Tivoli products. This website includes a link for obtaining the forum using a Rich Site Summary (RSS) feed.

- Technical Exchange Webcasts

www-01.ibm.com/software/sysmgmt/products/support/supp_tech_exch.html

This website provides webcasts in which technical experts share their knowledge and answer your questions. Visit this site often to see upcoming topics and presenters or to listen to previous webcasts.

Providing feedback about publications

Your feedback is important to help IBM provide the highest quality information. You can provide comments or suggestions about the documentation from the IBM Tivoli Storage Productivity Center for Replication Information Center.

Go to the information center at <http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp> and click the **Feedback** link at the bottom of the information center Welcome page or topic pages.

New for Tivoli Storage Productivity Center for Replication for System z 4.2.2.1

Use this information to learn about new features and enhancements in Tivoli Storage Productivity Center for Replication for System z version 4.2.2.1. This section highlights the changes since Tivoli Storage Productivity Center for Replication for System z 4.2.2.

IBM Storwize V7000 Unified

Storwize V7000 Unified is a virtualized storage system that includes Storwize V7000 and Storwize V7000 File Module. Storwize V7000 Unified is designed to consolidate block and file workloads into a single storage system for simplicity of management, reduced cost, highly scalable capacity, performance, and high availability. Storwize V7000 Unified also offers improved efficiency and flexibility through built-in solid state drive (SSD) optimization, thin provisioning and non-disruptive migration of data from existing storage. The system can virtualize and reuse existing disk systems offering a greater potential return on investment.

Data for Storwize V7000 Unified storage systems is collected, monitored, displayed, and reported in Tivoli Storage Productivity Center and data replication is supported by Tivoli Storage Productivity Center for Replication.

New for Tivoli Storage Productivity Center for Replication for System z 4.2.2

Use this information to learn about new features and enhancements in Tivoli Storage Productivity Center for Replication for System z version 4.2.2. This section highlights the changes since Tivoli Storage Productivity Center for Replication for System z 4.2.1.

Tivoli Storage Productivity Center for Replication for System z 4.2.2 supports IBM XIV Storage System. You can use the following session types for an XIV system:

Snapshot

Snapshot is a new session type that creates a point-in-time copy (snapshot) of a volume or set of volumes without having to define a specific target volume. The target volumes of a Snapshot session are automatically created when the snapshot is created.

Metro Mirror Failover/Failback

Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 kilometers apart. You can use failover and failback to switch the direction of the data flow.

Global Mirror Failover/Failback

Global Mirror is a method of asynchronous, remote data replication that operates between two sites that are over 300 kilometers apart. You can use failover and failback to switch the direction of the data flow.

Tivoli Storage Productivity Center for Replication for System z support for XIV system includes the following new features.

Support for volume nickname

For XIV system sessions, you can provide the volume ID or the volume nickname as a parameter value when you add or remove copy sets by using the command line interface (CLI) commands **mkcpset** and **rmcpset**.

In addition, you can include the XIV system volume ID or the volume nickname in a comma-separated value (CSV) file that is used to import copy set information. You can import the CSV file by using the **importcsv** command or the Tivoli Storage Productivity Center for Replication for System z graphical user interface (GUI). CSV files that are exported from Tivoli Storage Productivity Center for Replication for System z for XIV system sessions include the volume nickname rather than the volume ID. CSV files are exported by using the **exportcsv** command.

New CLI commands

The following CLI commands are new. For more information about new and updated CLI commands for Tivoli Storage Productivity Center for Replication for System z 4.2.2, see the IBM Tivoli Storage Productivity Center for Replication for System z Command-line Interface User's Guide.

cmdsnapgrp

Use the **cmdsnapgrp** command to run a specific action against a snapshot group that is in an XIV system Snapshot session. A snapshot group is a grouping of snapshots of individual volumes in a consistency group at a specific point in time.

lssnapgrp

Use the **lssnapgrp** command to view snapshot groups that are in an XIV system Snapshot session.

lssnapgrpactions

Use the **lssnapgrpactions** command to specify the session and snapshot group name that you want to view available actions for.

lssnapshots

Use the **lssnapshots** command to view snapshots that are in a snapshot group in an XIV system session.

Chapter 1. Product overview

This section provides an overview of the Tivoli Storage Productivity Center for Replication, describes the key concepts necessary to use the product and its components, and contains several scenarios that illustrate how to perform specific types of replication.

Introducing IBM Tivoli Storage Productivity Center for Replication for System z

This section provides an overview of IBM Tivoli Storage Productivity Center for Replication for System z, and IBM Tivoli Storage Productivity Center for Replication Basic Edition for System z.

Tivoli Storage Productivity Center for Replication for System z

IBM Tivoli Storage Productivity Center for Replication for System z provides replication management for IBM System Storage DS8000, IBM System Storage DS6000, IBM TotalStorage Enterprise Storage Server Model 800, IBM System Storage SAN Volume Controller, IBM Storwize V7000 Unified, IBM Storwize V7000, and XIV system storage systems.

You can use Tivoli Storage Productivity Center for Replication for System z for replication management regardless of whether the type of data on the system is extended count key data or fixed-block architecture.

You can use the following functionality for replication management:

- Volume protection to exclude any volumes from being used for disaster-protection copy operations.
- Command prompting to confirm storage administrator actions before running the copy services commands.
- User roles for administrative levels of access.
- Site awareness to indicate site locations of the storage volumes and to help assure copies are done correctly.
- Metro Global Mirror support for System Storage DS8000, providing failover and failback support, fast re-establishment of three-site mirroring, quick resynchronization of mirrored sites using incremental changes only, and data currency at the remote site.

IBM Tivoli Storage Productivity Center for Replication Basic Edition for System z (z/OS)

IBM Tivoli Storage Productivity Center for Replication Basic Edition for System z provides a disaster-recovery solution that helps to protect you from storage system failures.

You can use Tivoli Storage Productivity Center for Replication Basic Edition for System z with Basic HyperSwap® to perform the following tasks:

- Monitoring for events that indicate a storage device has failed
- Determining if the failing storage device is part of a Metro Mirror (synchronous peer-to-peer remote copy [PPRC]) pair

- Determining from policy, the action to be taken
- Ensuring that data consistency is not violated
- Swapping the I/O between the primary logical devices in the consistency group with the secondary logical devices in the consistency group (performing a HyperSwap for IBM System Storage DS8000, System Storage DS6000, IBM TotalStorage Enterprise Storage Server 800).
- Performing FlashCopy point-in-time replication for IBM System Storage DS8000, System Storage DS6000, IBM TotalStorage Enterprise Storage Server 800.
- Allowing only z/OS attached count key data (CKD) volumes to be added to the HyperSwap session.

Tivoli Storage Productivity Center for Replication Basic Edition for System z provides only HyperSwap and FlashCopy sessions and not the full functionality of the other IBM Tivoli Storage Productivity Center for Replication products.

Tivoli Storage Productivity Center for Replication Basic Edition for System z is available at no cost. If you want to use Tivoli Storage Productivity Center for Replication Two Site Business Continuity or Tivoli Storage Productivity Center for Replication Three Site Business Continuity a license is required for each product.

The z/OS HyperSwap license is required for Basic HyperSwap.

Tivoli Storage Productivity Center for Replication Two Site Business Continuity

You can use Tivoli Storage Productivity Center for Replication Two Site Business Continuity to obtain continuous availability and disaster recovery solutions by using point-in-time replication, which includes FlashCopy and Snapshot, and continuous replication, which includes Metro Mirror and Global Mirror.

You can set up Metro Mirror and Global Mirror sessions to replicate data in both the forward and reverse directions.

Use Tivoli Storage Productivity Center for Replication Two Site Business Continuity to create and manage the following session types:

- FlashCopy:
 - FlashCopy replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server Model 800
 - FlashCopy replication for SAN Volume Controller
 - FlashCopy replication for Storwize V7000
 - FlashCopy replication for Storwize V7000 Unified
- Snapshot:
 - Snapshot replication for the XIV system
- Global Mirror:
 - Global Mirror Single Direction replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server 800
 - Global Mirror Single Direction replication for SAN Volume Controller
 - Global Mirror Single Direction replication for Storwize V7000
 - Global Mirror Single Direction replication for Storwize V7000 Unified
 - Global Mirror Failover/Failback replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server 800
 - Global Mirror Failover/Failback replication for SAN Volume Controller

- Global Mirror Failover/Failback replication for Storwize V7000
- Global Mirror Failover/Failback replication for Storwize V7000 Unified
- Global Mirror Failover/Failback replication for XIV system
- Global Mirror Failover/Failback with Practice replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server 800
- Global Mirror Failover/Failback with Practice replication for SAN Volume Controller
- Global Mirror Failover/Failback with Practice replication for Storwize V7000
- Global Mirror Failover/Failback with Practice replication for Storwize V7000 Unified
- Global Mirror Either Direction with Two-Site Practice replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server 800
- Metro Mirror:
 - Metro Mirror or Global Copy replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server 800
 - Metro Mirror for Single Direction replication for SAN Volume Controller
 - Metro Mirror for Single Direction replication for Storwize V7000
 - Metro Mirror for Single Direction replication for Storwize V7000 Unified
 - Metro Mirror for Failover/Failback replication for SAN Volume Controller
 - Metro Mirror for Failover/Failback replication for Storwize V7000
 - Metro Mirror for Failover/Failback replication for Storwize V7000 Unified
 - Metro Mirror Failover/Failback replication for XIV system
 - Metro Mirror for Failover/Failback with Practice replication for SAN Volume Controller
 - Metro Mirror for Failover/Failback with Practice replication for Storwize V7000
 - Metro Mirror for Failover/Failback with Practice replication for Storwize V7000 Unified
 - Metro Mirror Failover/Failback or Global Copy replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server 800
 - Metro Mirror Failover/Failback replication for HyperSwap for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server 800
 - Metro Mirror Failover/Failback replication for Open HyperSwap for the System Storage DS8000

Tivoli Storage Productivity Center for Replication Three Site Business Continuity

You can use Tivoli Storage Productivity Center for Replication Three Site Business Continuity to obtain continuous availability and disaster recovery solutions by using point-in-time replication, which includes FlashCopy and Snapshot, and continuous replication, which includes Metro Mirror, Global Mirror, and Metro Global Mirror to secondary and tertiary sites.

You can set up Metro Mirror and Global Mirror sessions to replicate data in both the forward and reverse directions. With Tivoli Storage Productivity Center for Replication Three Site Business Continuity, you can also use Metro Global Mirror

(with failover and failback to switch production sites between each of the three sites and return to the original configuration.

Important: Tivoli Storage Productivity Center for Replication Three Site Business Continuity requires that Tivoli Storage Productivity Center for Replication Two Site Business Continuity is installed and a separate license is required for both products.

Use Tivoli Storage Productivity Center for Replication Three Site Business Continuity to create and manage the following session types:

- FlashCopy:
 - FlashCopy replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server Model 800
 - FlashCopy replication for SAN Volume Controller
 - FlashCopy replication for Storwize V7000
 - FlashCopy replication for Storwize V7000 Unified
- Snapshot:
 - Snapshot replication for the XIV system
- Global Mirror:
 - Global Mirror Single Direction replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server 800
 - Global Mirror Single Direction replication for SAN Volume Controller
 - Global Mirror Single Direction replication for Storwize V7000
 - Global Mirror Single Direction replication for Storwize V7000 Unified
 - Global Mirror Failover/Failback replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server 800
 - Global Mirror Failover/Failback replication for SAN Volume Controller
 - Global Mirror Failover/Failback replication for Storwize V7000
 - Global Mirror Failover/Failback replication for Storwize V7000 Unified
 - Global Mirror Failover/Failback replication for XIV system
 - Global Mirror Failover/Failback with Practice replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server 800
 - Global Mirror Failover/Failback with Practice replication for SAN Volume Controller
 - Global Mirror Failover/Failback with Practice replication for Storwize V7000
 - Global Mirror Failover/Failback with Practice replication for Storwize V7000 Unified
 - Global Mirror Either Direction with Two-Site Practice replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server 800
- Metro Mirror:
 - Metro Mirror or Global Copy replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server 800
 - Metro Mirror for Single Direction replication for SAN Volume Controller
 - Metro Mirror for Single Direction replication for Storwize V7000
 - Metro Mirror for Single Direction replication for Storwize V7000 Unified
 - Metro Mirror for Failover/Failback replication for SAN Volume Controller
 - Metro Mirror for Failover/Failback replication for Storwize V7000

- Metro Mirror for Failover/Failback replication for Storwize V7000 Unified
- Metro Mirror Failover/Failback replication for XIV system
- Metro Mirror for Failover/Failback with Practice replication for SAN Volume Controller
- Metro Mirror for Failover/Failback with Practice replication for Storwize V7000
- Metro Mirror for Failover/Failback with Practice replication for Storwize V7000 Unified
- Metro Mirror Failover/Failback or Global Copy replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server 800
- Metro Mirror Failover/Failback with HyperSwap replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server 800
- Metro Mirror Failover/Failback with Open HyperSwap replication for System Storage DS8000
- Metro Global Mirror:
 - Metro Global Mirror replication for System Storage DS8000 (with failover and failback)
 - Metro Global Mirror replication for System Storage DS8000 and TotalStorage Enterprise Storage Server 800 (with failover and failback)
 - Metro Global Mirror with Practice replication for System Storage DS8000 and TotalStorage Enterprise Storage Server 800 (with failover and failback)
 - Metro Global Mirror with HyperSwap replication for System Storage DS8000 and TotalStorage Enterprise Storage Server 800 (TotalStorage Enterprise Storage Server 800 can be only in the H1 role.)

Architecture

IBM Tivoli Storage Productivity Center for Replication for System z consists of several key components. This topic identifies these components and shows how they are related.

IBM Tivoli Storage Productivity Center for Replication for System z server

Database

A single database instance serves as the repository for all IBM Tivoli Storage Productivity Center for Replication for System z data.

GUI You can manage and monitor replication from the graphical user interface (GUI).

CLI You can issue commands for major IBM Tivoli Storage Productivity Center for Replication for System z functions from the command-line interface (CLI).

Tivoli Storage Productivity Center for Replication user interfaces

IBM Tivoli Storage Productivity Center for Replication provides a graphical user interface (GUI) and a command line interface (CLI) for managing data replication and disaster recovery.

Graphical user interface

Tivoli Storage Productivity Center for Replication uses a GUI with the following features:

Navigation tree

The left panel provides categories of tasks that you can perform in Tivoli Storage Productivity Center for Replication. Clicking a task opens a main page in the content panel.

Health Overview

This area is below the navigation tree and shows a status summary for all sessions, storage systems, host systems, and management servers that Tivoli Storage Productivity Center for Replication is managing.

Content area

The right panel displays content based on the item you selected in the navigation tree.

You can view help for the currently displayed panel by clicking the ? icon in the upper-right corner.

You can view the information center on the Web by clicking the ⓘ icon in the upper-right corner. You must have Internet access to view the information center.

When you log on to the GUI, by default you see the Health Overview panel in the content area.

Command line interface

You can use the Tivoli Storage Productivity Center for Replication CLI interactively using the `csmlcli` utilities. This CLI can be used either as an efficient way to accomplish simple tasks directly or as a script for automating functions.

Note: For security, the CLI runs only on the management server. You can run the CLI remotely using a remote-access utility, such as secure shell (SSH) or Telnet.

For Tivoli Storage Productivity Center for Replication on Windows, you can specify remote access to Linux or AIX® terminals if you have enabled Telnet on your Windows server.

Key concepts

This section describes key concepts to help you understand and effectively use IBM Tivoli Storage Productivity Center for Replication.

Management servers

The *management server* is a system that has IBM Tivoli Storage Productivity Center for Replication 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 Tivoli Storage Productivity Center for Replication that runs on a different physical system, but is continuously synchronized with the primary (or active) Tivoli Storage Productivity Center for Replication 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. If the active management server fails, you can issue the Takeover command to make the standby management server take over.

Connecting the active management server to the standby management server

Ensure that the active management server is connected to the standby management server. This connection creates the *management server relationship* that begins the synchronization process. Each management server can be in only one management server relationship.

A management server relationship might become disconnected for a number of reasons, including a connectivity problem or a problem with the alternate server. Issue the Reconnect command to restore synchronization.

Performing a takeover on the standby management server

If you must perform a takeover and use the standby server, ensure that you shut down the active management server first. You must ensure that you do not have two active management servers. If there are two active management servers and a condition occurs on the storage systems, both management servers respond to the same conditions, which might lead to unexpected behavior.

If you perform an action on the active management server when the servers are disconnected, the servers will be out of synch.

Viewing the status of the management servers

You can view the status of the active and standby management servers from the Management Servers panel in the Tivoli Storage Productivity Center for Replication graphical user interface (GUI). If you are logged on to the active management server, the icons on this panel show the status of the standby management server. If you are logged on to the standby management server, the icons on this panel show the status of the active management server.

When the status is Synchronized, the standby management server contains the same data that the active management server contains. Any update to the active management server database is replicated to the standby server database.

Managing volumes on storage systems

When you add direct connections, Hardware Management Console (HMC) connections, or z/OS connections on the active management server, Tivoli Storage Productivity Center for Replication automatically enables the management of attached extended count key data (ECKD™) volumes, non-attached count key data (CKD) volumes, and all fixed-block volumes on the storage system. To disable management of volumes on the storage system, use the volume protection function.

Information specific to management servers in z/OS environments

If the standby management server is not in the active server z/OS sysplex, the standby server is not be able to communicate with the storage systems using a z/OS connection; therefore, another connection must be made using a TCP/IP connection.

If DB2 is configured for data sharing mode across the z/OS sysplex, one of the Tivoli Storage Productivity Center for Replication servers must be configured to use the zero-administration embedded repository. If the embedded repository is not used, the two servers will overwrite the same data in the Tivoli Storage Productivity Center for Replication database.

Storage systems

A *storage system* is a hardware device that contains data storage. Tivoli Storage Productivity Center for Replication can control data replication within and between various storage systems.

To replicate data among storage systems using Tivoli Storage Productivity Center for Replication, you must manually add a connection to each storage system in the Tivoli Storage Productivity Center for Replication configuration. This allows you to omit storage systems for which Tivoli Storage Productivity Center for Replication is not to manage replication and omit storage systems that are being managed by another Tivoli Storage Productivity Center for Replication management server.

For redundancy, you can connect a single storage system using a combination of direct, Hardware Management Console (HMC), and z/OS connections.

You can use the following storage systems:

- IBM TotalStorage Enterprise Storage Server (ESS) Model 800
- IBM System Storage DS6000
- IBM System Storage DS8000
- IBM System Storage SAN Volume Controller
- IBM Storwize V7000
- IBM Storwize V7000 Unified
- IBM XIV Storage System

A SAN Volume Controller can virtualize a variety of storage systems. Although Tivoli Storage Productivity Center for Replication does not support all storage systems, you can manage these storage systems through a single SAN Volume Controller cluster interface. Tivoli Storage Productivity Center for Replication connects directly to the SAN Volume Controller clusters.

You can define a location for each storage system and for each site in a session. When adding copy sets 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.

Notes:

- Tivoli Storage Productivity Center for Replication does not automatically discover the physical locations of storage systems. You can manually assign a location to a storage system from the GUI and CLI.

- Throughout this document, *ESS/DS* refers to the following models:
 - IBM TotalStorage Enterprise Storage Server Model 800
 - IBM System Storage DS8000
 - IBM System Storage DS6000

Ports

Tivoli Storage Productivity Center for Replication use ports for communication with the management servers in a high-availability relationship, graphical user interface (GUI), command-line interface (CLI), and storage systems.

Web browser ports

To launch the Tivoli Storage Productivity Center for Replication GUI, use one of these default ports:

	HTTP port	HTTPS port
WebSphere Application Server	9080	9443
IBM System Services Runtime Environment for z/OS or WebSphere Application Server OEM Edition for z/OS	32208	32209

You can verify the ports that are correct for your installation in the *install_root/AppServer/profiles/profile_name/properties/portdef.props* file. The ports is defined by the *WC_defaulthost* (HTTP port) and *WC_defaulthost_secure* (HTTPS port) properties within the file.

Standby management server port

Tivoli Storage Productivity Center for Replication uses the default port 5120 for communication between the active and standby management server. This port number is initially set at installation time.

Important: The standby management server port number must be the same on both management servers in a high-availability relationship. If you change the port number on one management server, you must also change it on the other.

You can view the current port for each management server by clicking **Management Servers** in the navigation tree or from the **Health Overview** panel in the GUI or using the *lshaservers* command from the command line interface.

Client port

IBM Tivoli Storage Productivity Center for Replication client uses the default port 5110 to communicate with the graphical user interface and command line interface from a remote system. This port number is initially set at installation time.

Important: The client port number must be the same on both management servers in a high-availability relationship. If you change the port number on one management server, you must also change it on the other.

You can view the client port number on the local management server by clicking **About** in the navigation tree in the GUI or using the **whoami** command from the command line interface.

Storage system ports

The following table lists the default ports for each storage type.

Table 1. Storage system default ports

Storage System	Connection Type	Port
<ul style="list-style-type: none"> • TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 	Direct Connection	2433
<ul style="list-style-type: none"> • System Storage DS8000 	Hardware Management Console Connection	1750
<ul style="list-style-type: none"> • System Storage SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified 	Direct Connection	443 and 22
<ul style="list-style-type: none"> • The XIV system 	Direct Connection	7778

Ensure that your network configuration is set up so that Tivoli Storage Productivity Center can send outgoing TCP/IP packets to the storage controllers. It is possible when adding the storage controllers to Tivoli Storage Productivity Center to set a specific port number for your storage controller.

Because there are typically multiple applications running on the management server, it is possible that port conflicts might arise if other applications attempt to use the same ports that IBM Tivoli Storage Productivity Center for Replication is using. Use the **netstat** command to verify which ports the various applications on the management server are using.

When you add a storage system to the Tivoli Storage Productivity Center for Replication configuration, the port field is automatically populated with the appropriate value. If you want to use different ports, you can change them by clicking **Storage Systems** located in the navigation tree, clicking the storage system that you want to change, and then changing the port value in the View/Modify Details panel.

Note: The storage system must not be in a Connected state if you want to change port values.

If firewalls are being used in your configuration, ensure that none of these ports are being blocked. Also ensure that not only is the Tivoli Storage Productivity Center for Replication server granted access to reach the other components, but that the other components are granted access to reach the Tivoli Storage Productivity Center for Replication server.

Storage connections

You must create a connection from the IBM Tivoli Storage Productivity Center for Replication management server to each storage system. You can connect either directly or through a Hardware Management Console (HMC) or IBM z/OS connection.

A single storage system can be connected using multiple connections for redundancy. For example, you can connect a IBM System Storage DS8000 storage system using an HMC connection and a z/OS connection. Tivoli Storage Productivity Center for Replication monitors how a storage system has been added to the configuration.

When you add a storage connection to the Tivoli Storage Productivity Center for Replication configuration, the storage system and the connection are added to the active management server configuration. For direct and HMC connections, the storage system and connection are also added to the standby management server configuration. For z/OS connections, only the storage system is added to the standby management server configuration. The connection is not added because the standby management server might not be running on z/OS and might not have access to the volumes on the storage system through a z/OS connection.

The storage systems are not required to be connected to the standby management server. However, if a storage system does not have a connection on the standby management server, you cannot manage copy services on the storage system from the standby server.

Important: If the Metro Mirror heartbeat is enabled, do not connect to a IBM TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, or System Storage DS6000 storage system using both an HMC connection and a direct connection. If you have both types of connections and the direct connection is lost, the session changes to the suspended state even though the HMC connection is still valid. If both connections are lost and the session is in the suspended state, restart the session when connectivity is regained to synchronize the session with the hardware.

When Tivoli Storage Productivity Center for Replication is running on z/OS and a storage system is added to the Tivoli Storage Productivity Center for Replication configuration through a TCP/IP (direct or HMC) connection, all ECKD volumes that are attached to the management server are managed through the TCP/IP connection. To use the Fibre Channel connection, you must explicitly add the storage system to the Tivoli Storage Productivity Center for Replication configuration through a z/OS connection.

If a storage system was previously added to the Tivoli Storage Productivity Center for Replication configuration through a z/OS connection and later the storage system is added through a TCP/IP connection, all non-attached ECKD volumes and fixed block volumes are added to the Tivoli Storage Productivity Center for Replication configuration.

When you remove a storage system, Tivoli Storage Productivity Center for Replication automatically removes all connections that the storage system is using with exception of the z/OS connection. You can also individually remove each connection through which the storage system is connected.

If Tivoli Storage Productivity Center for Replication has multiple connections to a specific storage system, the order in which you remove the connections produces different results:

- If you remove all direct and HMC connections first, the fixed block and non-attached ECKD volumes are removed from the Tivoli Storage Productivity Center for Replication configuration. The remaining ECKD volumes that are attached through the z/OS connection remain in the Tivoli Storage Productivity Center for Replication configuration until the z/OS connection is removed. Removing the TCP/IP connection also disables the Metro Mirror heartbeat.
- If you remove the z/OS connection first and if there is an HMC or direct connection to volumes, those volumes are not removed from the Tivoli Storage Productivity Center for Replication configuration.
- HyperSwap can run provided that volumes are attached and available to z/OS storage, even if you are using a TCP/IP connection to storage.

Direct connection

The Tivoli Storage Productivity Center for Replication management server can connect directly to TotalStorage Enterprise Storage Server Model 800, DS6000, DS8000, SAN Volume Controller, Storwize V7000, Storwize V7000 Unified or the XIV system storage systems through a TCP/IP connection. The TCP/IP connection is required to discover the storage systems configuration (such as LSSs, volumes, volume size, and format), issue queries, and receive asynchronous events.

DS8000 storage systems on an IPV4 network can be connected directly to the management server. A direct connection requires an Ethernet card in the cluster. DS8000 storage systems on an IPV6 network cannot use a direct connection. They can be connected only through an HMC or z/OS connection.

When you add a direct connection to a DS or ESS cluster, specify the following information for cluster 0 and 1:

- IP addresses or domain names
- Ports
- User names
- Passwords

SAN Volume Controller or Storwize V7000 can virtualize various storage systems. Although Tivoli Storage Productivity Center for Replication does not support all storage systems, you can manage these storage systems through a single SAN Volume Controller or Storwize V7000 cluster interface. Tivoli Storage Productivity Center for Replication connects directly to the SAN Volume Controller or Storwize V7000 clusters. When you add a direct connection to a SAN Volume Controller or Storwize V7000 cluster to the Tivoli Storage Productivity Center for Replication configuration, specify the cluster IP address of the SAN Volume Controller or Storwize V7000 cluster, which in turn points to multiple SAN Volume Controller or Storwize V7000 storage systems. Ensure that the user name and password are correct for the cluster. If incorrect values are used, significant communication problems can occur, such as never advancing to the Prepared state.

Important: The SAN Volume Controller or Storwize V7000 user name must have privileges to maintain SSH keys. For information about troubleshooting Secure Shell connections to the SAN Volume Controller or Storwize V7000, see the Ethernet Connection Restrictions on SAN Volume Controller website at www-01.ibm.com/support/docview.wss?uid=ssg1S1002896.

Hardware Management Console connection

The IBM Tivoli Storage Productivity Center for Replication management server can connect to DS8000 storage systems through a Hardware Management Console (HMC). An HMC can have multiple DS8000 storage systems connected to it. When you add an HMC to the IBM Tivoli Storage Productivity Center for Replication configuration, all DS8000 storage systems that are behind the HMC are also added. You cannot add or remove individual storage systems that are behind an HMC.

You can also add a dual-HMC configuration, in which you have two HMCs for redundancy. This is recommended when the Metro Mirror heartbeat is required. You must configure both HMCs identically, including the user ID and password.

If planned maintenance is necessary on the HMC, it is recommended that you disable the Metro Mirror heartbeat on the management server while the maintenance is performed.

If the HMC needs to go down frequently or reboots frequently, it is recommended that you disable the Metro Mirror heartbeat. If the Metro Mirror heartbeat is required, the direct connection is recommended instead of an HMC connection.

Important: If a DS8000 storage system uses an HMC connection, the Metro Mirror heartbeat could trigger a freeze on the storage system and impact applications for the duration of the long busy timeout timer if the HMC is shut down for any reason, including upgrading microcode. The long busy timeout timer is the time after which the storage system will allow I/O to begin again after a freeze occurs if no run command has been issued by IBM Tivoli Storage Productivity Center for Replication. The default value is two minutes for ECKD volumes or one minute for fixed block volumes.

Notes:

- The user ID that you use to connect to the HMC must have admin, op_storage, or op_copy_services privileges on the DS8000 storage system.
- For minimum microcode requirements to connect to a DS8000 through a management console, see the Supported Storage Products List website at www-01.ibm.com/support/docview.wss?uid=swg21386446.

z/OS connections

An IBM Tivoli Storage Productivity Center for Replication management server that runs on z/OS can connect to IBM TotalStorage Enterprise Storage Server (ESS) Model 800, DS8000, and DS6000 storage systems through a z/OS connection. The z/OS connection is used to issue replication commands and queries for attached ECKD volumes over an existing Fibre Channel network and to receive asynchronous events. When a storage system is added to IBM Tivoli Storage Productivity Center for Replication through the z/OS connection, all ECKD volumes that are attached to the IBM Tivoli Storage Productivity Center for Replication management system are added to the IBM Tivoli Storage Productivity Center for Replication configuration. ECKD volumes that are not attached to the IBM Tivoli Storage Productivity Center for Replication z/OS management server are not added to the IBM Tivoli Storage Productivity Center for Replication configuration through the z/OS connection.

Notes:

- Ensure that all volumes in the logical storage subsystem (LSS) that you want to manage through a z/OS connection are attached to z/OS. Either the entire LSS

must be attached to z/OS or none of the volumes in the LSS should be attached to z/OS for IBM Tivoli Storage Productivity Center for Replication to properly manage queries to the hardware.

- The z/OS connection is limited to storage systems that are connected to an IBM Tivoli Storage Productivity Center for Replication management server running z/OS.
- The Metro Mirror heartbeat is not supported through the z/OS connection. To use the Metro Mirror heartbeat, the storage systems must be added using a direct connection or Hardware Management Console (HMC) connection. If the Metro Mirror heartbeat is enabled, a storage system is added through a direct connection and z/OS connection, and the direct connection becomes disconnected, then a suspend results as there is no heartbeat through the z/OS connection.

If at least one volume in a Logical Storage Subsystem (LSS) is attached through a z/OS connection, then all volumes in that LSS must be similarly attached. For example, if there are two ECKD volumes in an LSS, and one volume is attached to the IBM Tivoli Storage Productivity Center for Replication system using a z/OS connection and the other is attached through a direct connection, IBM Tivoli Storage Productivity Center for Replication would have knowledge of direct-connected volume. IBM Tivoli Storage Productivity Center for Replication issues commands to both volumes over the Fibre Channel network; however, commands issued to the direct-connection volume will fail, and IBM Tivoli Storage Productivity Center for Replication will show that the copy set that contains that volume has an error.

Use the following guidelines to add storage systems through a z/OS connection:

- Use the z/OS connection to manage ECKD volumes that are attached to an IBM Tivoli Storage Productivity Center for Replication management server running z/OS.
- To manage z/OS attached volumes through a z/OS connection (for example, for HyperSwap), you must explicitly add the z/OS connection for that storage system in addition to a TCP/IP connection (either the direct connection or the HMC connection).
- Create a z/OS connection before all TCP/IP connections if you want to continue to have IBM Tivoli Storage Productivity Center for Replication manage only the attached ECKD volumes.

Tip: It is recommended that you create both TCP/IP and z/OS connections for ECKD volumes to allow for greater storage accessibility.

Sessions

A *session* is used to perform a specific type of data replication against a specific set of volumes. 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.

If a session has failover and failback capabilities, you can perform a site switch in which you move the application from one site to another and change the direction of the copy without having to perform a full copy. Without failover and failback capabilities, each time you move the application and writing to a different site in the session, you must initiate a full copy to synchronize the new source with the new target to regain disaster recovery capability. An IBM Tivoli Storage Productivity Center for Replication session with failover and failback capabilities

uses the hardware's ability to track changes after a suspension (where applicable), so only the changed data must be resynchronized.

The type of data replication (also known as the session type) that is associated with the session determines the actions that can be performed against all copy sets in the session, the number of volumes in each copy set, and the role that each volume plays.

Important: Use only the Tivoli Storage Productivity Center for Replication GUI or CLI to manage session relationships, such as volume pairs and copy sets. Do not modify session relationships through individual hardware interfaces, such as the DSCCLI. Modifying relationships through the individual hardware interfaces can result in a loss of consistency across the relationships managed by the session, and might cause the session to be unaware of the state or consistency of the relationships.

Copy sets

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 machine type (for example, 3380 volumes must be used with other 3380 volumes and 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 (or copy type) that is associated with the session to which the copy set belongs.

Important: Use the IBM WebSphere Application Server Administrator Console to check the Java heap size (**Application servers > Server1 > Process Definition > Servant > Java Virtual Machine**) for the IBM z/OS servant region. The size of this region affects the performance of IBM Tivoli Storage Productivity Center for Replication. The default Java heap size is 512 MB, which supports fewer than 25,000 role pairs. Increasing the Java heap size to 768 MB increases support to a maximum of 50,000 role pairs. For more information on how to set up the Java heap size, see the WebSphere Application Server information center at <http://publib.boulder.ibm.com/infocenter/wasinfo/v6r1/index.jsp>.

The following table lists the estimated number of role pairs and volumes per copy set that are supported for each session type.

Table 2. Support number of role pairs and volumes per copy set for each session type

Session	Role Pairs	Volumes
Basic HyperSwap	1	2
FlashCopy	1	2
Snapshot ¹	0	1
Metro Mirror	1	2
Metro Mirror with Practice	3	3
Global Mirror (ESS/DS)	3	3
Global Mirror (SAN Volume Controller)	1	2
Global Mirror with Practice (ESS/DS)	5	4
Global Mirror with Practice (SAN Volume Controller)	3	3

Table 2. Support number of role pairs and volumes per copy set for each session type (continued)

Session	Role Pairs	Volumes
Global Mirror Two-Site Practice	8	6
Metro Global Mirror	6	4
Metro Global Mirror with Practice	8	5
1. An XIV Snapshot session requires that the user define only the H1 volumes. All target volumes are created on the same storage pool as the source volumes.		

Use the Add Copy Sets wizard to add copy sets to an existing session. You can select a storage system; a logical subsystem (LSS), I/O group, or pool; or single volume for each role and then create one or more copy sets for the session.

You can use one of the following volume pairing options to automatically create multiple copy sets in the same session.

Storage system matching (System Storage DS8000, System Storage DS6000, or TotalStorage Enterprise Storage Server Model 800 Metro Mirror sessions only)

Creates copy sets by matching volumes (based on the volume IDs) across all logical subsystems (LSSs) for the selected storage systems. For example, volume 01 on the source LSS is matched with volume 01 on the target LSS.

You cannot select the storage system and select **All Logical Subsystems** in the list of LSSs. You can also do auto-matching at the LSS level for Metro Mirror sessions.

LSS, I/O group, or pool matching

Creates copy sets by matching all volumes based on the selected LSS, I/O group, or pool for each role in copy set.

Select the storage system and LSS, I/O group, or pool, and then select **All Volumes** in the **Volume** list.

If you do not want to use the auto-generated volume pairing for a copy set, clear that copy set so that it is not added during the wizard. Then, add the remaining copy sets and reopen the Add Copy Set wizard and manually enter the volume pairings that you want.

Invalid copy sets are not added to the session. Copy sets can be invalid if their volumes are not the same type or size.

You can remove copy sets that you do not want to add to the session, even if they are valid. This enables you to filter and eliminate unwanted copy sets before they are added to the session.

You can export the copy sets to take a snapshot of your session at a particular point in time for backup purposes.

Note: You can copy an entire storage system only for Metro Mirror sessions.

Considerations for adding copy sets

When you create a copy set, a warning is displayed if one or more selected volumes already exist in another session. Whether it is safe to add the created copy set to the session depends on the environment. For example, if you created one

session for normal replication and another session for a disaster recovery practice scenario, you must use the same target volumes from the original session as the source volumes in the practice session. If the volume you selected is already in another session, confirm that this is the configuration you want.

Considerations for removing copy sets

You remove a copy set or range of copy sets by selecting the source volume; LSS, I/O group, or pool; or storage system. When the list of copy sets that meet your criteria is displayed, you can select the copy sets that you want to remove.

The consequence of removing copy sets varies depending on the state of the session:

Defined

There is no relationship on the hardware. This removes the copy set from Tivoli Storage Productivity Center for Replication data store.

Preparing or Prepared

The copy set is currently copying data, so Tivoli Storage Productivity Center for Replication terminates the hardware relationship for the copy set. The rest of the copy sets continue to run uninterrupted.

Suspended or Target Available

Any existing relationships on the hardware are removed for the copy set.

Before removing all copy sets from that session, terminate the session. Removing the copy sets when the session is active can considerably increase the amount of time it takes for the copy set removal to complete. Copy sets are removed one at a time, and when the session is active, this requires commands being issued to the hardware. However, if you terminate the session first, then commands are not issued to the hardware and the removal process completes faster.

Tip: When you a remove copy set from Tivoli Storage Productivity Center for Replication, you might want keep hardware relationships on the storage systems. These relationships are useful when you want to migrate from one session type to another or when resolving problems. For more information about keeping the hardware relationships when removing copy sets, see [Removing Copy Sets](#).

The behavior that occurs when a copy set is removed varies depending on the storage system:

ESS 800, DS6000, and DS8000:

- The complete copy set is removed from Tivoli Storage Productivity Center for Replication.
- Any peer-to-peer remote copy (PPRC) pair that is part of a Global Mirror consistency group is removed from the consistency group on the storage system.
- If the PPRC pair is part of a Global Mirror consistency group and is the last remaining source volume in a subordinate session, the subordinate session is removed from the storage system.
- If the PPRC pair is the last remaining participant in a Global Mirror session, the Global Mirror session is removed from the storage system.
- Any PPRC relationship remains on the storage system.

- A Metro Mirror (synchronous PPRC) pair that is in a HyperSwap configuration is removed from that configuration but the pair remains on the hardware.
- FlashCopy relationship remains on the storage system if the hardware has not already completed any background copy.

SAN Volume Controller, Storwize V7000, Storwize V7000 Unified, or the XIV system

- The complete copy set is removed from Tivoli Storage Productivity Center for Replication.
- FlashCopy, Metro Mirror, and Global Mirror relationships are pulled out of their consistency group. If they are the last remaining relationship in a consistency group, that consistency group is removed from the hardware.

When you specify the force removal option, all knowledge of the specified copy set is removed from Tivoli Storage Productivity Center for Replication, even if the relationship itself still exists. If this occurs, you are not able to remove the relationship using Tivoli Storage Productivity Center for Replication, because no information about the relationship exists. If you force a removal of a copy set and the removal fails, you must manually remove the relationship from the hardware. If you do not, you cannot to create new relationships.

One benefit of forcing a removal of the copy sets is that Tivoli Storage Productivity Center for Replication does not manage the consistency of copy sets that it has no knowledge of. This means that additional commands to the session do not affect the removed copy sets, even though they are still in a relationship on the hardware.

If you do not specify the force removal option and an error occurs that prevents the hardware relationships from being removed, the copy set will not be removed from Tivoli Storage Productivity Center for Replication. The copy set remains as part of the session, and you can still perform actions on it.

To re-add the copy set to the session, you must perform a full copy of the data.

Volume roles

Volume roles are given to every volume in the copy set. The role defines how the volume is used in the copy set and, for multi-site sessions, the site location of the volume. For example, the H1 role is made up of host-attached volumes that are located at the primary site.

The site determines the location of the volumes. The number of sites in a copy set is determined by the type of data replication (also known as the session type) that is associated with the session. IBM Tivoli Storage Productivity Center for Replication supports up to three sites:

- Site 1** The location of the primary storage that contain 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 that receives the copy updates from the primary storage.
- Site 3** (Metro Global Mirror only) The location of the tertiary storage that receives data updates from the secondary storage.

The volume roles that are needed in a copy set are determined by the type of replication that is associated with the session. IBM Tivoli Storage Productivity Center for Replication supports these volume roles:

Host volume

A volume that is connected to a server that reads and writes I/O. A host volume can be the source of updated tracks when the server connected to the host volume is actively issuing read and write input/output (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 volume that stores data that has 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 is created between the host or intermediate volume and the corresponding journal volume after a recover request is initiated to create another consistent version of the data.

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

Intermediate volume

A volume that 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 being used, data on intermediate volumes might not be consistent.

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

Target volume

(FlashCopy only) A volume that receives data from a source, either 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).

Role pairs

A *role pair* is the association of two volume 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 host volumes at the primary site and host volumes at the secondary site (H1-H2).

The flow of data in the role pair is shown using an arrow. For example, H1>H2 denotes that H1 is the source and H2 is the target.

Participating role pairs are role pairs that are currently participating in the session's copy.

Non-participating role pairs are role pairs that are not actively participating in the session's copy.

Snapshot sessions do not use role pairs.

Practice volumes

You can use a *practice volume* to practice what you would do in the event of a disaster, without interrupting current data replication. Practice volumes are available in Metro Mirror, Global Mirror, and Metro Global Mirror sessions.

To use the practice volumes, the session must be in the prepared state. Issuing the Flash command against the session while in the Prepared state creates a usable practice copy of the data on the target site.

Note: You can test disaster-recovery actions without using practice volumes; however, without practice volumes, you cannot continue to copy data changes between volumes while testing disaster-recovery actions.

Consistency groups

For Global Mirror and Metro Global Mirror sessions, IBM Tivoli Storage Productivity Center for Replication manages the consistency of dependant writes by creating a consistent point-in-time copy across multiple volumes or storage systems. A *consistency group* is a set of target volumes in a session that have been updated to preserve write order and are therefore recoverable.

Data exposure is the period when data is written to the storage at the primary site until data is replicated to storage at the secondary site. Data exposure is influenced by factors such as:

- Requested consistency-group interval time
- Type of storage systems
- Physical distance between the storage systems
- Available bandwidth of the data link
- Input/output (I/O) load on the storage systems

To manage data exposure, you can change the consistency group interval time. The *consistency group time interval* specifies how often a Global Mirror and Metro Global Mirror session attempts to form a consistency group. When you reduce this value, it might be possible to reduce the data exposure of the session. A lower value causes the session to attempt to create consistency groups more frequently, which might also increase the processing load and message-traffic load on the storage systems.

Session types

Tivoli Storage Productivity Center for Replication supports several methods of data replication. The type of data replication that is associated with a given session is known as the *session type* (also known as a copy type).

The following table describes the session types that are supported by Tivoli Storage Productivity Center for Replication. Depending on the edition of Tivoli Storage Productivity Center for Replication that you are using, some of these session types might not be available.

Table 3. Session type summary

Copy type	Supported Software	Supported storage systems	Description
Basic HyperSwap	Tivoli Storage Productivity Center for Replication Basic Edition for System z and Tivoli Storage Productivity Center for Replication for System z	<ul style="list-style-type: none"> • TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 	Basic HyperSwap replication is a special Metro Mirror replication method designed to provide high availability in the case of a disk storage system failure. Using Basic HyperSwap with Metro Mirror, you can configure and manage your synchronous Peer-to-Peer Remote Copy (PPRC) pairs.
FlashCopy	Tivoli Storage Productivity Center all editions	<ul style="list-style-type: none"> • TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified 	<p>FlashCopy replication creates a point-in-time copy in which the target volume contains a copy of the data that was on the source volume when the FlashCopy was established. Using FlashCopy, your data exists on the second set of volumes in the same storage system, and can be restored to the first set of volumes.</p> <p>SAN Volume Controller or Storwize V7000 FlashCopy sessions are managed by using FlashCopy consistency groups. Sessions for IBM TotalStorage Enterprise Storage Server (ESS) and IBM DS6000 and DS8000 are not managed by using FlashCopy consistency groups.</p>
Snapshot	Tivoli Storage Productivity Center all editions	The XIV system	Snapshot is a session type that creates a point-in-time copy of a volume or set of volumes without having to define a specific target volume. The target volumes of a Snapshot session are automatically created when the snapshot is created.

Table 3. Session type summary (continued)

Copy type	Supported Software	Supported storage systems	Description
Metro Mirror Single Direction	Tivoli Storage Productivity Center for Replication Two Site Business Continuity	<ul style="list-style-type: none"> • TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified 	<p>Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 KM apart. The source is located in one storage system and the target is located in another storage system. Using Metro Mirror, your data exists on the second site that is less than 300 KM away, and can be restored to the first site.</p>
Metro Mirror Failover/Failback	Tivoli Storage Productivity Center for Replication Two Site Business Continuity	<ul style="list-style-type: none"> • TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified • The XIV system 	<p>Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 KM apart. Using Metro Mirror Failover / Failback, your data exists on the second site that is less than 300 KM away. You can use failover and failback to switch the direction of the data flow. This ability enables you to run your business from the secondary site.</p> <p>Using Metro Mirror with HyperSwap, your data exists on the second site that is less than 300 KM away. The data can be restored to the first site. You can also use failover for a backup copy of the data if your primary volumes encounter a permanent I/O error.</p>

Table 3. Session type summary (continued)

Copy type	Supported Software	Supported storage systems	Description
Metro Mirror Failover/Failback with Practice	Tivoli Storage Productivity Center for Replication Two Site Business Continuity	<ul style="list-style-type: none"> • TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified 	Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 KM apart. The source is located in one storage system and the target is located in another storage system. Metro Mirror Failover / Failback with Practice combines Metro Mirror and FlashCopy to provide a point-in-time copy of the data on the remote site.
Global Mirror Single Direction	Tivoli Storage Productivity Center for Replication Two Site Business Continuity	<ul style="list-style-type: none"> • TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified 	Global Mirror is a method of asynchronous, remote data replication between two sites that are over 300 KM apart. It maintains identical data in both the source and target, where the source is located in one storage system and the target is located in another storage system. Using Global Mirror, your data exists on the second site that is more than 300 KM away, and can be restored to the first site.
Global Mirror Either Direction with Two-Site Practice	Tivoli Storage Productivity Center for Replication Two Site Business Continuity	<ul style="list-style-type: none"> • TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 	Global Mirror is a method of asynchronous, remote data replication between two sites that are over 300 KM apart. It maintains identical data in both the source and target, where the source is located in one storage system and the target is located in another storage system. Global Mirror Either Direction with Two-Site Practice combines Global Mirror and FlashCopy to provide a point-in-time copy of the data on either the primary or secondary sites that are over 300 KM apart.

Table 3. Session type summary (continued)

Copy type	Supported Software	Supported storage systems	Description
Global Mirror Failover/Failback	Tivoli Storage Productivity Center for Replication Two Site Business Continuity	<ul style="list-style-type: none"> • TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified • The XIV system 	Global Mirror is a method of asynchronous, remote data replication between two sites that are over 300 KM apart. It maintains identical data in both the source and target, where the source is located in one storage system and the target is located in another storage system. Using Global Mirror Failover / Failback, your data exists on the second site that is more than 300 KM away, and you can use failover and failback to switch the direction of the data flow. This ability enables you to run your business from the secondary site.
Global Mirror Failover/Failback with Practice	Tivoli Storage Productivity Center for Replication Two Site Business Continuity	<ul style="list-style-type: none"> • TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified 	Global Mirror is a method of asynchronous, remote data replication between two sites that are over 300 KM apart. It maintains identical data in both the source and target, where the source is located in one storage system and the target is located in another storage system. Global Mirror Failover / Failback with Practice combines Global Mirror and FlashCopy to provide a point-in-time copy of the data on a remote site at a distance over 300 KM away from your first site.

Table 3. Session type summary (continued)

Copy type	Supported Software	Supported storage systems	Description
Metro Global Mirror	Tivoli Storage Productivity Center for Replication Three Site Business Continuity	<ul style="list-style-type: none"> • TotalStorage Enterprise Storage Server Model 800 (only H1 site) • System Storage DS8000 	<p>Metro Global Mirror is a method of continuous, remote data replication that operates between three sites that varying distances apart. Metro Global Mirror combines Metro Mirror synchronous copy and Global Mirror asynchronous copy into a single session, where the Metro Mirror target is the Global Mirror source. Using Metro Global Mirror and Metro Global Mirror with HyperSwap, your data exists on a second site that is less than 300 KM away, and a third site that is more than 300 KM away. Metro Global Mirror uses both Metro Mirror and Global Mirror Failover / Failback to switch the direction of the data flow. This ability enables you to run your business from the secondary or tertiary sites.</p> <p>Using Basic HyperSwap with Metro Global Mirror, you can configure and manage three-site continuous replication needed in a disaster recovery event.</p>

Table 3. Session type summary (continued)

Copy type	Supported Software	Supported storage systems	Description
Metro Global Mirror with Practice	Tivoli Storage Productivity Center for Replication Three Site Business Continuity	<ul style="list-style-type: none"> TotalStorage Enterprise Storage Server Model 800 (only H1 site) System Storage DS8000 	Using Metro Global Mirror with Practice, you can practice your disaster recovery actions while maintaining disaster recovery capabilities. Your data exists on a second site that is less than 300 KM away, and a third site that is more than 300 KM away. Metro Global Mirror uses both Metro Mirror and Global Mirror Failover / Failback to switch the direction of the data flow; as a result, you can run your business from the secondary or tertiary sites, and simulate a disaster.

Basic HyperSwap (ESS, DS6000, and DS8000)

Basic HyperSwap is an entitled copy services solution for z/OS version 1.9 and later. It provides high availability of data in the case of a disk storage system failure. Basic HyperSwap is not a disaster recovery solution. If a session is suspended but the suspend was not caused by a HyperSwap trigger, no freeze is done to ensure consistency of the session.

When HyperSwap is combined with Metro Mirror and Metro Global Mirror replication, you can prepare for disaster recovery and ensure high availability. If a session is suspended but the suspend was not caused by a HyperSwap trigger, a freeze is done to ensure consistency of the session.

Note: This replication method is available on only ESS, DS6000, and DS8000 storage systems, and on management servers running IBM Tivoli Storage Productivity Center for Replication Basic Edition for System z or IBM Tivoli Storage Productivity Center for Replication for System z.

Basic HyperSwap replication performs the following actions:

- Manages CKD volumes in Metro Mirror (synchronous peer-to-peer remote copy [PPRC]) relationships.
- Permits only CKD volumes to be added to the HyperSwap session. The graphical user interface (GUI) shows only CKD volumes when you add a copy set. The command-line interface (CLI) fails to add a copy set if a fixed block volume is specified.
- Monitoring for events that indicate a storage device has failed.
- Determining whether the failing storage device is part of a Metro Mirror (synchronous PPRC) pair.
- Determining the action to be taken from policy.
- Ensuring that data consistency is not violated.

- Swapping the I/O between the primary logical devices in the consistency group with the secondary logical devices in the consistency group. A swap can be from the preferred logical devices to the alternate logical devices or from the alternate logical devices to the preferred logical devices.

Metro Mirror Failover/Failback with HyperSwap

Metro Mirror Failover/Failback uses HyperSwap to configure and manage synchronous Peer-to-Peer Remote Copy (PPRC) pairs.

Metro Global Mirror with HyperSwap

Metro Global Mirror with HyperSwap is a z/OS replication feature that provides the three-site continuous replication needed in a disaster recovery event.

Important: If HyperSwap occurs by event when running a Metro Global Mirror with a HyperSwap session, a full copy of the data occurs to return to a full three-site configuration. If you issue a HyperSwap command when running a Metro Global Mirror with a HyperSwap session, a full copy does not occur. A full copy is required only for an unplanned HyperSwap or a HyperSwap initiated using the z/OS SETHS SWAP command.

Example

Jane is using multiple DS8000 storage systems. Her host applications run on z/OS and her z/OS environment has connectivity to the DS8000 storage systems. She has a site in Manhattan and a bunker in Hoboken. While she does not need a disaster recovery solution, she does need a high-availability solution to keep her applications running around the clock. Jane is worried that if a volume fails on her DS8000 in Manhattan, her database application will not be able to operate. Even a small downtime can cost Jane thousands of dollars. Jane uses a Basic HyperSwap session to mirror her data in Manhattan to her secondary DS8000 in Hoboken. If a volume at the Manhattan site fails, Basic HyperSwap automatically directs application I/O to the mirrored volumes in Hoboken.

FlashCopy


FlashCopy replication creates a point-in-time copy in which the target volume contains a copy of the data that was on the source volume when the FlashCopy was established.

The ESS, DS6000, and DS8000 platforms provide multiple logical subsystems (LSSs) within a single physical subsystem, while the following platforms provide multiple I/O groups:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

All platforms support local replication in which the source volume is located in one LSS or I/O group and the target volume is located in the same or another LSS or I/O group. Using FlashCopy, you can reference and update the source volume and target volume independently.

The following figure illustrates how a FlashCopy session works.



Example

Jane uses FlashCopy to make a point-in-time copy of the customer data in existing international accounts. Every night, the bank's servers perform batch processing. Jane uses FlashCopy to create checkpoint restarts for the batch processing in case the batch processing fails. In her batch processing, the first step is to balance all international accounts, with a FlashCopy taken of the resulting data. The second step in her batch processing is to process the international disbursements. If the second step in the batch process fails, Jane can use the FlashCopy made of the first step to repeat the second step, instead of beginning the entire process again. Jane also writes a CLI script that performs a FlashCopy every night at 11:59 PM, and another script that quiesces the database. She backs this data up on tape on her target storage system, and then sends the tape to the bank's data facility in Oregon for storage.

Snapshot

Snapshot is a session type that creates a point-in-time copy of a volume or set of volumes. You do not have to define a specific target volume. The target volumes of a Snapshot session are automatically created when the snapshot is created.

The XIV system uses advanced snapshot architecture to create a large number of volume copies without affecting performance. By using the snapshot function to create a point-in-time copy, and to manage the copy, you can save storage. With the XIV system snapshots, no storage capacity is used by the snapshot until the source volume (or the snapshot) is changed.

The following figure illustrates how a Snapshot session works.



Example

Jane's host applications are using an XIV system for their back-end storage. With the XIV system, Jane can create a large number of point-in-time copies of her data. The snapshot function ensures that if data becomes corrupted, she can restore the data to any number of points in time.

Jane sets up a Snapshot session by using Tivoli Storage Productivity Center for Replication and specifies the volumes on the XIV system that are used by her host applications. Jane does not have to provision target volumes for all the snapshots she intends to make. She can quickly get a single Snapshot session configured and ready.

When the session is configured, Jane writes a CLI script that performs a Create Snapshot command to the session every two hours. If a problem occurs, such as data becoming corrupted, Jane can find a snapshot of the data before the problem occurred. She can restore the data to that point.

By creating a set of snapshots of the data, Jane can also schedule batch processing against that data every day. She can use the batch processing to analyze certain trends in the market without causing any effect to the host applications.

Metro Mirror

Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 KM apart. The source is located in one storage system and the target is located in another storage system.

Attention: For Tivoli Storage Productivity Center for Replication for System z sessions containing Metro Mirror relationships, ensure that the session does not contain system volumes (such as paging volumes) unless the session is enabled for HyperSwap. If HyperSwap is not enabled, a freeze that is issued through Tivoli Storage Productivity Center for Replication might cause Tivoli Storage Productivity Center for Replication processing to freeze. This situation might prevent the session from ensuring that the data is consistent.


Metro Mirror replication maintains identical data in both the source and target. When a write is issued to the source copy, the changes made to the source data are propagated to the target before the write finishes posting. If the storage system goes down, Metro Mirror provides zero loss if data must be used from the recovery site.

A Metro Mirror session in Global Copy mode creates an asynchronous relationship to accommodate the high volume of data migration. As a result, the data on the target system might no longer be consistent with the source system. The Metro Mirror session switches back to a synchronous relationship when a Metro Mirror Start command is reissued. In addition, you can start a Metro Mirror session in Global Copy mode and toggle between Metro Mirror and Global Copy modes to accommodate time periods in which you value host input/output (I/O) response time over data consistency.

Tip: To determine if there is any out of synch data that must to be copied over before the session is consistent, check the percent that is complete in the session details panel.

Metro Mirror Single Direction


The following figure illustrates how a Metro Mirror Single Direction session works.



Metro Mirror Failover/Failback

Using Metro Mirror with failover/failback, your data exists on the second site that is less than 300 KM away, and you can use failover/failback to switch the direction of the data flow. This session type enables you to run your business from the secondary site, and to copy changes made at the second site back to the primary site when you want to resume production at the primary site.

The following figure illustrates how a Metro Mirror with Failover/Failback session works.




Metro Mirror Failover/Failback with Practice

A Metro Mirror Failover/Failback with Practice session combines Metro Mirror and FlashCopy to provide a point-in-time copy of the data on the remote site. You can use this session type to practice what you might do if a disaster occurred, without losing your disaster recovery capability.

This solution consists of two host volumes and an intermediate volume.

The following figure illustrates how a Metro Mirror Failover/Failback with Practice session works.



Metro Mirror Failover/Failback with HyperSwap

A Metro Mirror Failover/Failback session can be enabled to have HyperSwap capabilities. To enable HyperSwap the following circumstances must apply:

- The session is running on a Tivoli Storage Productivity Center for Replication server running on IBM z/OS.
- The volumes are only for TotalStorage Enterprise Storage Server, System Storage DS8000, and DS6000 systems.
- The volumes are count key data (CKD) volumes that are attached to the z/OS system.


Metro Mirror Failover/Failback with HyperSwap combines the high availability of Basic HyperSwap with the redundancy of a two-site Metro Mirror

Failover/Failback solution when managing count key data (CKD) volumes on z/OS. If the primary volumes encounter a permanent I/O error, the I/O is automatically swapped to the secondary site with little to no impact on the application.

A swap can be planned or unplanned. A planned swap occurs when you issue a HyperSwap command from the **Select Action** list in the graphical user interface (GUI) or when you issue a **cmdsess -action hyperswap** command.

The following figure illustrates how a Metro Mirror Failover/Failback session enabled for HyperSwap works.

Metro Mirror Failover/Failback



For more information about enabling HyperSwap, see “Managing a session with HyperSwap and Open HyperSwap replication” on page 42.

Metro Mirror Failover/Failback with Open HyperSwap

A Metro Mirror Failover/Failback session can be enabled to have Open HyperSwap capabilities. To enable Open HyperSwap the following circumstances must apply:


- The volumes in the session are System Storage DS8000 5.1 or later volumes.
- The volumes in the session are fixed block and are mounted to IBM AIX 5.3 or AIX 6.1 hosts with the following modules installed:
 - Subsystem Device Driver Path Control Module (SDDPCM) version 3.0.0.0 or later
 - Multi-Path Input/Output (MPIO) module (the version that is provided with AIX version 5.3 or 6.1)
- The connections between the AIX host systems and the Tivoli Storage Productivity Center for Replication server have been established.

Metro Mirror Failover/Failback with Open HyperSwap combines the high availability of Basic HyperSwap on z/OS for fixed block AIX volumes with the redundancy of a two-site Metro Mirror Failover/Failback solution. If the primary volumes encounter a permanent I/O error, the I/O is automatically swapped to the secondary site with little to no impact on the application.

A swap can be planned or unplanned. A planned swap occurs when you issue a HyperSwap command from the **Select Action** list in the GUI or when you issue a **cmdsess -action hyperswap** command.

The following figure illustrates how a Metro Mirror Failover/Failback session enabled for Open HyperSwap works.

Metro Mirror Failover/Failback



For more information about enabling Open HyperSwap, see “Managing a session with HyperSwap and Open HyperSwap replication” on page 42.

Examples

Metro Mirror Single Direction

At the beginning of a work week, Jane is notified that between 10:00 AM and 11:00 AM on the next Friday, power in her building is going to be shut off. Jane does not want to lose any transactions during the power outage, so she decides to transfer operations to the backup site during the outage. She wants a synchronous copy method with no data loss for her critical business functions, so she chooses Metro Mirror, which can be used between locations that are less than 300 KM apart.

In a synchronous copy method, when a write is issued to change the source, the change is propagated to the target before the write is completely posted. This method of replication maintains identical data in both the source and target. The advantage of this is when a disaster occurs, there is no data loss at the recovery site because both writes must complete before signaling completion of a write to the source application. Because the data must be copied to both System Storage DS8000 devices before the write is completed, Jane can be sure that her data is safe.

The night before the planned outage, Jane quiesces her database and servers in San Francisco and starts the database and servers in Oakland. To accomplish this task, Jane issues the Suspend and Recover commands, and then issues the Start command on the secondary site. She powers off her equipment in San Francisco to avoid any power spikes during reboot after the power is turned back on.

Metro Mirror in Global Copy mode

At the beginning of a work week, Jane is notified that between 10:00 AM and 11:00 AM on the next Friday, power in her building is going to be shut off. Jane does not want to lose any transactions during the power outage, so she decides to transfer operations to their backup site during the outage. She wants a synchronous copy method with no data loss for her critical business functions, so she chooses Metro Mirror, which can be used between locations that are less than 300 KM apart.

Jane wants to limit her application impact while completing the initial Metro Mirror synchronization, so she begins her session in Global Copy mode. After she sees that about 70% of the data has been copied, Jane decides to switch the session into Metro Mirror mode, assuring data consistency.

Metro Mirror with Practice

Jane wants to run a Metro Mirror with Practice from San Francisco to Oakland. She wants to verify her recovery procedure for the Oakland site,

but she cannot afford to stop running her Metro Mirror session while she takes time to practice a recovery. By using a Metro Mirror with Practice session, Jane is able to practice her disaster recovery scenario in Oakland while her Metro Mirror session runs uninterrupted. By practicing running her applications at the Oakland site, Jane is better prepared to make a recovery if a disaster ever strikes the San Francisco site.

While her session is running in a Prepared state, Jane practices a recovery at her Oakland site by issuing the Flash command. This command momentarily pauses the session and starts a FlashCopy to the H2 volumes. As soon as the FlashCopy is started, her session will be restarted. These FlashCopy files create a consistent version of the data on the H2 volume that she can use for recovery testing, while her session continues to replicate data from San Francisco to Oakland. As a result, she can carry out her recovery testing without stopping her replication for any extended duration of time.

If at some point in time, the Metro Mirror session suspends due to a failure, Jane can use the practice session to restart her data replication while maintaining a consistent copy of the data at the Oakland site, in case of a failure during the resynchronization process. When the session is suspended, she can issue a Recover command to create a consistent version of the data on the H2 volumes. After the Recover command completes, she can issue the Start H1->H2 command to resynchronize the data from the San Francisco site to the Oakland site. If a failure occurs before her restarted session is in the Prepared state, she has a consistent version of the data on the H2 volumes. She must simply issue the Recover command to put the session into Target Available state and make the H2 volumes accessible from her servers. If the session was not in the Prepared state when it suspended, the subsequent Recover command does not issue the FlashCopy files to put the data on the H2 volumes. This means that the consistent data on the H2 volumes are not overwritten if the data to be copied to them is not consistent.

Metro Mirror Failover/Failback enabled for Open HyperSwap

Jane wants to run a Metro Mirror with Practice from San Francisco to Oakland. She wants to verify her recovery procedure for the Oakland site, but she cannot afford to stop running her Metro Mirror session while she takes time to practice a recovery. By using a Metro Mirror with Practice session, Jane is able to practice her disaster recovery scenario in Oakland while her Metro Mirror session runs uninterrupted. By practicing running her applications at the Oakland site, Jane is better prepared to make a recovery if a disaster ever strikes the San Francisco site.

While her session is running in a Prepared state, Jane practices a recovery at her Oakland site by issuing the Flash command. This command momentarily pauses the session and starts a FlashCopy to the H2 volumes. As soon as the FlashCopy is started, the session is restarted. These FlashCopy files create a consistent version of the data on the H2 volume that she can use for recovery testing, while her session continues to replicate data from San Francisco to Oakland. As a result, she can carry out her recovery testing without stopping her replication for any extended duration of time.

If the Metro Mirror session suspends due to a failure, Jane can use the practice session to restart her data replication while maintaining a consistent copy of the data at the Oakland site, in case of a failure during the resynchronization process. When the session is suspended, she can

issue a Recover command to create a consistent version of the data on the H2 volumes. After the Recover command completes, she can issue the Start H1->H2 command to resynchronize the data from the San Francisco site to the Oakland site. If a failure occurs before her restarted session is in the Prepared state, she has a consistent version of the data on the H2 volumes. She must simply issue the Recover command to put her session into Target Available state and make the H2 volumes accessible from her servers. If the session was not in the Prepared state when it suspended, the subsequent Recover command does not issue the FlashCopy files to put the data on the H2 volumes. This means that the consistent data on the H2 volumes are not overwritten if the data to be copied to the volumes is not consistent.

Selecting a HyperSwap session

A global insurance company has elected to use Tivoli Storage Productivity Center for Replication to manage its disaster recovery environment. Jane wants minimal data exposure, both for planned outages such as routine maintenance, and for unplanned disasters. They have CKD volumes on System Storage DS8000 devices, and use z/OS mainframes. They have two data centers in New York.

Jane reviews the Tivoli Storage Productivity Center for Replication documentation, and chooses a Metro Mirror recovery solution, since her priority is being protected against regional disasters. Jane decides to use Metro Mirror solution, because her company has two data centers located near each other. Jane realizes that because she uses z/OS, CKD, and System Storage DS8000 hardware, she is also able to use a HyperSwap solution. Using Metro Mirror Failover/Failback with HyperSwap, Jane can minimize application impact, while maintaining seamless failover to her secondary site. Jane decides Metro Mirror Failover/Failback with HyperSwap is best for the needs of her company.

After installing and configuring Tivoli Storage Productivity Center for Replication on z/OS, Jane starts the Tivoli Storage Productivity Center for Replication GUI. She adds the Tivoli Storage Productivity Center for Replication storage devices she intends to use on all three sites. From the Session Overview panel, Jane launches the Create Session wizard, and selects the **Metro Mirror Failover/Failback** session type. Then, she continues the wizard, she selects the **Manage H1H2 with HyperSwap** option. After finishing the wizard, Jane clicks **Launch Add Copy Sets Wizard**. She completes this wizard, and issues a **Start H1->H2** command. After the initial copy is completed, Jane is safely replicating her data between both sites. She can also issue a HyperSwap between sites 1 and 2, enabling her to switch sites with minimal application impact during either a disaster or maintenance period.

Performing a planned HyperSwap

Jane's company has successfully been using Metro Mirror Failover/Failback with HyperSwap sessions for the past three months. However, Jane needs to perform maintenance on an H1 box. During this time, Jane does not want her applications or replication to be interrupted. To prevent this from happening, shortly before the maintenance is scheduled to begin, Jane decides to use the Tivoli Storage Productivity Center for Replication GUI to perform a HyperSwap to the H2 volumes. This transitions the applications so that they write to H2. To perform a planned HyperSwap, Jane issues a HyperSwap command.

Understanding what happens when an unplanned HyperSwap occurs

Several weeks after the planned maintenance at Jane's company is completed, an incident occurs at the H1 site. A disk controller fails, causing one of the H1 volumes to encounter a permanent I/O error. Fortunately, Jane's data is safe because she used Metro Mirror Failover/Failback with HyperSwap, and her H2 volume is an exact duplicate of the H1 volume. When the permanent I/O error is detected, a HyperSwap is triggered. The application seamlessly transitions to writing to the H2 volumes. Her data is safe, and her applications are not interrupted.

Jane configured an Simple Network Management Protocol (SNMP) listener to alert her to any events, so she receives the SNMP event indicating that a HyperSwap occurred. Jane investigates the cause of the HyperSwap and uses the z/OS console to identify the volume that triggered the HyperSwap. Jane replaces the faulty disk controller. Then, to recover from the unplanned HyperSwap, Jane issues the **Start H2->H1** command.

Global Mirror


Global mirror is a method of asynchronous, remote data replication between two sites that are over 300 kilometers (km) apart. It maintains identical data in both the source and target, where the source is located in one storage system and the target is located in another storage system.

The data on the target is typically written a few seconds after the data is written to the source volumes. When a write is issued to the source copy, the change is propagated to the target copy, but subsequent changes are allowed to the source before the target verifies that it has received the change. Because consistent copies of data are formed on the secondary site at set intervals, data loss is determined by the amount of time since the last consistency group was formed. If your system stops, Global Mirror might lose some data that was being transmitted when the disaster occurred. Global Mirror still provides data consistency and data recoverability in the event of a disaster.

Global Mirror Single Direction

A Global Mirror single direction session allows you to run your Global Mirror replication from only the primary site.

For ESS, DS6000, and DS8000 storage systems, each copy set in the Global Mirror single direction session consists of two host volumes and a journal volume. The following figure illustrates how a Global Mirror single direction session works on an ESS, DS6000, and DS8000 storage systems:



For the following storage systems, each copy set in the Global Mirror Single Direction session consists of two host volumes:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

The following figure illustrates how a Global Mirror Single Direction session works on these storage systems:




Global Mirror Either Direction with Two-Site Practice (ESS, DS6000, and DS8000)

A Global Mirror Either Direction with Two Site Practice session allows you to run Global Mirror replication from either the primary or secondary site. It combines Global Mirror and FlashCopy to provide a point-in-time copy of the data on a remote site at a distance over 300 km away from your first site. This practice session allows you to create practice volumes on both the primary and secondary site to practice what you might do if a disaster occurred, without losing your disaster recovery capability.

Note: This replication method is available on only ESS, DS6000, and DS8000 storage systems.


The session consists of two host volumes, two intermediate volumes, and two journal volumes. The following figure illustrates how a Global Mirror either direction with two-site practice session works:



Global Mirror Failover/Failback

Using Global Mirror failover/failback, your data exists on the second site that is more than 300 km away, and you can use failover/failback to switch the direction of the data flow. This enables you to run your business from the secondary site.

For ESS, DS6000, and DS8000 storage systems, each copy set in the Global Mirror failover/failback session consists of two host volumes and a journal volume. The following figure illustrates how a Global Mirror failover/failback session works on an ESS, DS6000, or DS8000 storage systems:



For the following storage systems, each copy set in the Global Mirror failover/failback session consists of two host volumes:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified
- The XIV system


The following figure illustrates how a Global Mirror failover/failback session works on these storage systems.



Global Mirror Failover/Failback with Practice

A Global Mirror failover/failback with practice session combines Global Mirror and FlashCopy to provide a point-in-time copy of the data on a remote site at a distance over 300 km away from your primary site. You can use this to practice what you might do if a disaster occurred, without losing your disaster recovery capability.


For ESS, DS6000, and DS8000 storage systems, each copy set in the Global Mirror failover/failback with practice session consists of two host volumes, an intermediate volume, and a journal volume. The following figure illustrates how a Global Mirror failover/failback with practice session works on an ESS, DS6000, or DS8000 storage system:



For the following storage systems, each copy set in the Global Mirror failover/failback with Practice session consists of two host volumes and an intermediate volume:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

The following figure illustrates how a Global Mirror Failover/Failback with Practice session works on these storage systems:



Examples

Global Mirror Single Direction

Although Jane's FlashCopy and Metro Mirror copies were both planned, Jane realizes that sometimes unforeseen things happen, and she wants to make sure her data is safe. Because Jane works in San Francisco, she wants her other site to be far away in case of a localized disaster. Her other site is based in Houston. Jane's foresight pays off when a minor earthquake occurs in San Francisco and power and communications both go down. Fortunately, Jane has arranged for the data on customer accounts that have recently opened or closed to be asynchronously copied in Houston, using

Global Mirror. Jane risks losing the bytes of data that were being processed when the tremor disrupted the San Francisco process, but she views that as a minor inconvenience when weighed next to the value of backing up her data in a non-earthquake zone.

Global Mirror with Practice

Jane wants to run a Global Mirror with practice from San Francisco to Houston. She wants to verify her recovery procedure for the Houston site, but she cannot afford to stop running her Global Mirror session while she takes time to practice a recovery. By using a Global Mirror with Practice session, Jane is able to practice her disaster recovery scenario in Houston while her Global Mirror session runs uninterrupted. By practicing running her applications at the Houston site, Jane will be better prepared to make a recovery if a disaster ever strikes the San Francisco site.

Global Mirror Either Direction with Two-Site Practice

Jane wants to run a Global Mirror with practice from San Francisco to Houston. She wants to verify her recovery procedure for the Houston site, but she cannot afford to stop running her Global Mirror session while she takes time to practice a recovery. By using a Global Mirror either direction two-site with practice session, Jane is able to practice her disaster recovery scenario in Houston while her Global Mirror session runs uninterrupted. By practicing running her applications at the Houston site, Jane will be better prepared to make a recovery if a disaster ever strikes the San Francisco site.

Jane can use the Global Mirror either direction with two-site practice session to run asynchronous consistent data replication from either the San Francisco site or the Houston site. (She can practice her disaster recovery at the target site, no matter what her current production site is.) Jane's business is able to run a consistent Global Mirror session from its Houston site back to San Francisco while running a production at the Houston site.

Setting up Global Mirror for Resource Groups on System Storage DS8000

If resource groups are defined on a System Storage DS8000, Global Mirror session IDs might be defined for some users. Tivoli Storage Productivity Center for Replication does not automatically determine which session IDs are valid. To determine which session IDs are valid, you must modify the `rmserver.properties` file and add the following property:

```
gm.master.sessionid.gm_role,session_name = xx
```

where *gm_role* is the role that has the master volume (for example, H1 in a Global Mirror failover/failback session), *session_name* is the name of the session that uses the session ID, and *xx* is the decimal number for the session ID.

Important: System Storage DS8000 represents session IDs as a two-digit hexadecimal number. Use the decimal version of that number. For example, if you want a Global Mirror failover/failback session to use a session ID of 0F, the decimal number is 15 as shown in the following example:

```
gm.master.sessionid.H2.11194_wprac=15
```

Metro Global Mirror (ESS 800 and DS8000)

Metro Global Mirror is a method of continuous, remote data replication that operates between three sites that varying distances apart. Metro Global Mirror

combines Metro Mirror synchronous copy and Global Mirror asynchronous copy into a single session, where the Metro Mirror target is the Global Mirror source.

Notes:

- This replication method is available on only ESS 800 and DS8000 storage systems.
- You can select ESS 800 storage systems in only the H1 volume role. All other volume roles must use DS8000 volumes.
- You can mix ESS 800 and DS8000 volumes in the H1 volume role. If ESS 800 and DS8000 storage systems are both used in the H1 role, the DS8000 storage system performs Incremental Resync (IR), and the ESS 800 storage system performs a full copy. Because ESS 800 does not support the IR function, a full copy is required when changing from H1->H2->H3 to H1>H3 and from H2->H1->H3 to H2->H3.

Metro Global Mirror maintains a consistent copy of data at the remote site, with minimal impact to applications at the local site. This remote mirroring function works in combination with FlashCopy to meet the requirements of a disaster-recovery solution by providing the following features:

- Fast failover and failback
- Rapid reestablishment of three-site mirroring, without production outages
- Data currency at the remote site with minimal lag behind at the local site, an average of only 3 - 5 seconds for many environments
- Quick resynchronization of mirrored sites using only incremental changes

If IBM Tivoli Storage Productivity Center for Replication is running on z/OS, you can configure a Metro Global Mirror session to control the Metro Mirror relationship between the primary and secondary site using HyperSwap. With HyperSwap enabled, a failure on the primary storage system causes an automatic HyperSwap, transparently redirecting application I/O to the auxiliary storage system. The Global Mirror relationship continues to run uninterrupted throughout this process. With this configuration, you can achieve near zero data loss at larger distances.

Using synchronous mirroring, you can switch from local site H1 to remote site H2 during a planned or unplanned outage. It also provides continuous disaster recovery protection of site H2 through site H3, without the necessity of additional reconfiguration, if a switch from site H1 occurs. With this configuration, you can reestablish H2->H1->H3 recoverability while production continues to run at site H2. Additionally, this cascaded setup can reduce the load on-site H1 as compared to some multi-target (non-cascaded) three-site mirroring environments.


Important:

- If HyperSwap occurs by event when running a Metro Global Mirror with a HyperSwap session, a full copy of the data occurs to return to a full three-site configuration. If you issue a HyperSwap command when running a Metro Global Mirror with a HyperSwap session, a full copy does not occur. A full copy is required only for an unplanned HyperSwap or a HyperSwap initiated using the z/OS SETHS SWAP command.
- In Metro Global Mirror and Metro Global Mirror with Practice sessions, when the H1 is on an ESS 800, you might risk filling up the space efficient journal volumes. Because incremental resynchronization is not supported on the ESS 800, full copies are performed in many of the transitions.

Metro Global Mirror

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

The following figure illustrates how a Metro Global Mirror session works.




Metro Global Mirror with Practice

A Metro Global Mirror session with Practice combines Metro Mirror, Global Mirror, and FlashCopy across three sites to provide a point-in-time copy of the data on the third site. You can use this session to practice what you might do if a disaster occurred without losing your disaster recovery capability.

The session consists of three host volumes, an intermediate volume, and a journal volume. The following figure illustrates how a Metro Global Mirror with Practice session works.

Note: A Metro Global Mirror with practice session can be created when the three-site license has been applied to the server.



Note: In Metro Global Mirror and Metro Global Mirror with Practice sessions, when the H1 is on an ESS 800, you might risk filling up the space efficient journal volumes. Because incremental resynchronization is not supported on the ESS 800, full copies are performed in many of the transitions.

Examples

Metro Global Mirror

Although Jane works in San Francisco, she wants to give herself the ability to run her business from either Oakland (her secondary site) or Houston (her tertiary site). Jane can use Metro Global Mirror with failover/failback to switch the direction of the data flow, so that she can run her business from either Oakland or Houston. Metro Global Mirror means that Jane has zero data loss backup at her secondary site, and minimal data loss at her tertiary site.

Metro Global Mirror with Practice

Jane wants to run a Metro Global Mirror with Practice from San Francisco to Houston. She wants to verify her recovery procedure for the Houston site. However, she cannot afford to stop running her Metro Global Mirror session while she takes time to practice a recovery. By using a Metro Global Mirror with Practice session, Jane is able to practice her disaster recovery scenario in Houston while her Metro Global Mirror session runs uninterrupted. By practicing running her applications at the Houston site, and being prepared to run her applications at the Oakland site if necessary, Jane will be better prepared to make a recovery if a disaster ever strikes the San Francisco site.

Jane can use Metro Global Mirror with Practice to switch the direction of the data flow, so that she can run her business from either Oakland or Houston. Using Metro Global Mirror, Jane has zero data loss backup at her secondary site; and minimal data loss at her tertiary site.

Managing a session with HyperSwap and Open HyperSwap replication

HyperSwap and Open HyperSwap provide high availability of data in the case of a primary disk storage system failure. When a failure occurs in writing input/output (I/O) to the primary storage system, the failure is detected by IOS, and IOS automatically swaps the I/O to the secondary site with no user interaction and little or no application impact.

Sessions that can be enabled for HyperSwap or Open HyperSwap:

You can create sessions that have HyperSwap or Open HyperSwap capabilities. Enabling swapping provides a session with a highly available business continuity solution.

Sessions that can enable HyperSwap

The following session types can enable HyperSwap:

- Basic HyperSwap
- Metro Mirror with Failover/Failback
- Metro Global Mirror

To enable HyperSwap, the following circumstances must apply:

- The session is running on an Tivoli Storage Productivity Center for Replication server that is running on IBM z/OS.

- The volumes are only for TotalStorage Enterprise Storage Server, System Storage DS8000, and DS6000 systems.
- The volumes are count key data (CKD) volumes that are attached to the z/OS system.

Sessions that can enable Open HyperSwap

Only Metro Mirror with Failover/Failback session type can enable Open HyperSwap.

To enable Open HyperSwap, the following circumstances must apply:

- The volumes in the session are only for System Storage DS8000 5.1 or later volumes.
- The volumes in the session are fixed block and mounted to IBM AIX 5.3 or AIX 6.1 hosts with the following modules installed:
 - Subsystem Device Driver Path Control Module (SDDPCM) version 3.0.0.0 or later
 - Multi-Path Input/Output (MPIO) module (the version that is provided with AIX version 5.3 or 6.1)
- The connections between the AIX host systems and the Tivoli Storage Productivity Center for Replication server have been established.

Setting up the environment for HyperSwap:

You must set up an environment that supports HyperSwap before attempting to enable HyperSwap for a IBM Tivoli Storage Productivity Center for Replication session.

The following steps must be completed before HyperSwap can be enabled. For more information about these steps, see the IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide

1. Install IBM Tivoli Storage Productivity Center for Replication for System z.
2. Perform the post installation tasks of setting up the data store and other necessary system settings.
3. Ensure that all RESERVEs are converted to global enqueues (ENQs).
4. Ensure that all volumes in the session that you are enabling for HyperSwap are attached to the IBM z/OS system that is running Tivoli Storage Productivity Center for Replication.

Setting up the environment for Open HyperSwap:

You must set up an environment that supports Open HyperSwap before attempting to enable Open HyperSwap for a IBM Tivoli Storage Productivity Center for Replication session.

The following steps must be completed before Open HyperSwap can be enabled:

1. Ensure that the IBM AIX hosts and IBM System Storage DS8000 to meet the following hardware and software requirements:

AIX requirements

Open HyperSwap support requires AIX version 5.3 or 6.1. (You can find the supported AIX version for each Tivoli Storage Productivity Center for Replication release in the support matrix at <http://www-01.ibm.com/support/docview.wss?rs=40&context=SSBSEX&context=SSMN28&context=SSMMUP&context=SS8JB5>)

&context=SS8JFM&uid=swg21386446&loc=en_US&cs=utf-8&lang=en.
Click the link for the applicable release under Agents, Servers and GUI.)

You must have the following AIX modules installed:

- Subsystem Device Driver Path Control Module (SDDPCM) version 3.0.0.0 or later
- Multi-Path Input/Output (MPIO) module (the version that is provided with AIX version 5.3 or 6.1)

System Storage DS8000 hardware requirements

Only System Storage DS8000 storage systems are supported. Open HyperSwap requires System Storage DS8000 5.1 or later.

Open HyperSwap does not support High Availability Cluster Multi-Processing (HACMP[™]).

2. Create connections from Tivoli Storage Productivity Center for Replication to the AIX hosts (see “Adding a host system connection” on page 111).
3. Assign copy set volumes from the storage device to the host using the System Storage DS8000 command-line interface (CLI) or graphical user interface (GUI).
4. Run the AIX `cfgmgr` command to discover the volumes assigned to the host.

Considerations for Open HyperSwap and the AIX host:

- A single session that has Open HyperSwap enabled can manage multiple hosts; however, each host can be associated with only one session. Multiple hosts can share the same session.
- For AIX 5.3, a single host can manage a maximum of 1024 devices that have been enabled for Open HyperSwap on the host, with 8 logical paths configured for each copy set in the session. For AIX 6.1, a single host can manage a maximum of 1024 devices that have been enabled for Open HyperSwap on the host, with 16 logical paths configured for each copy set in the session.
- If an application on the host has opened a device, a Tivoli Storage Productivity Center for Replication session for that device cannot be terminated. The **Terminate** command fails. To terminate the session, you must either close the application or remove the copy sets from the session. If you remove copy sets from the session, you must ensure that the application writes to the correct volume when the copy set relationship is restored.
- It is possible for Open HyperSwap to fail on a subset of hosts for the session and work on the remaining hosts for the same session. In this situation, you must determine the best action to take if the application is writing to volumes on the source system as well as volumes on the target system. Contact the IBM Support Center if you need assistance determining the best solution for this issue.
- To enable support for Open HyperSwap on the host, refer to the *IBM System Storage Multipath Subsystem Device Driver User's Guide*.

Configuring timers to support Open HyperSwap:

There are configurable timeout values for the storage system, IBM Tivoli Storage Productivity Center for Replication, and IBM AIX hosts systems that can affect the operation of Open HyperSwap.

The following list describes the various timeout values that can affect Open HyperSwap:

Storage system quiesce timeout value

This quiesce timeout timer begins when the storage system starts a quiesce operation. When the timer value expires, input/output (I/O) is resumed on the primary device. The default timeout value is two minutes, but the value can be set from 30 to 600 seconds. To set the quiesce timeout value, see the information about the **chdev** command in the *IBM System Storage Multipath Subsystem Device Driver User's Guide*.

Storage system long busy timeout value

This timeout value is the time in seconds that the logical subsystem (LSS) consistency group volume stays in the long busy state after a remote mirror and copy error is reported.

Timeout values for the applications that are on the host

The various applications that are running on the host have timeout values. The timeout values vary depending on the application.

Considerations for setting timers

Consider the following information for setting timers:

- If the host quiesce timer is set to a shorter value than the Tivoli Storage Productivity Center for Replication response timer, it is possible that an I/O swap failure can occur. If a storage system triggers an unplanned failover and if the storage system quiesce timer expires before Tivoli Storage Productivity Center for Replication responds, the host attempts to write I/O to the primary volume where the loss of access occurred. If the hardware condition that caused the loss of access continues, the attempt to write I/O fails again and an unplanned Open HyperSwap is not performed.
- If the host quiesce timer is set to a longer value than the Tivoli Storage Productivity Center for Replication response timer, an application timeout might occur if Open HyperSwap takes too long to complete.

Enabling a Session for HyperSwap or Open HyperSwap:

Enabling HyperSwap or Open HyperSwap for a session provides a combined business recovery and business continuity solution.

To ensure that your environment supports HyperSwap or Open HyperSwap, see “Setting up the environment for HyperSwap” on page 43 or “Setting up the environment for Open HyperSwap” on page 43.

Perform these steps to enable HyperSwap or Open HyperSwap for a session.

1. In the IBM Tivoli Storage Productivity Center for Replication navigation tree, select **Sessions**. Click the radio button next to the session that you want to enable.
2. From the **Select Action** menu, select **View/Modify Properties** and click **Go**. If you have not created the session, click **Create Session**. You can enable HyperSwap or Open HyperSwap on the Properties page
3. Under **ESS / DS Metro Mirror Options**, select from the following HyperSwap or Open HyperSwap options:
 - **Manage H1-H2 with HyperSwap**. This option enables a session to manage the H1-H2 sequence using HyperSwap. If you select this option, select from the following additional options.
 - **Disable HyperSwap**. Select this option to prevent a HyperSwap from occurring by command or event.

- **On Configuration Error.** Choose one of the following options:
 - **Partition the system(s) out of the sysplex.** Select this option to partition out of the sysplex when a new system is added to the sysplex and encounters an error in loading the configuration. A restart of your system is required if you select this option.
 - **Disable HyperSwap.** Select this option to prevent a HyperSwap from occurring by command or event.
 - **On Planned HyperSwap Error.** Choose one of the following options:
 - **Partition out the failing system(s) and continue swap processing on the remaining system(s).** Select this option to partition out the failing system and continues the swap processing on any remaining systems.
 - **Disable HyperSwap after attempting backout.** Select this option to stop the HyperSwap action, and disable the HyperSwap commands or events.
 - **On Unplanned HyperSwap.** Choose one of the following options:
 - **Partition out the failing system(s) and continue swap processing on the remaining system(s).** Select this option to partition out the failing systems and continues the HyperSwap processing on the remaining systems when a new system is added to the sysplex and HyperSwap does not complete. A restart of your system is required if you select this option.
 - **Disable HyperSwap after attempting backout.** Select this option to stop the HyperSwap action, and disable the HyperSwap commands or events.
 - **Manage H1-H2 with Open HyperSwap.** If volumes are attached to an IBM AIX host, Tivoli Storage Productivity Center for Replication can manage the H1-H2 sequence of a Metro Mirror session using Open HyperSwap. If this option is selected, a failure on the host accessible volumes triggers a swap, which redirects application I/O to the secondary volumes. Only volumes that are currently attached to the host systems that are defined on the Tivoli Storage Productivity Center for Replication Host Systems panel are eligible for Open HyperSwap.
 - **Disable Open HyperSwap.** Select this option to prevent a swap from occurring by a command or event while keeping the configuration on the host system and all primary and secondary volumes coupled.
4. Click **OK** to apply the selected options.

Restarting an AIX Host System that is enabled for Open HyperSwap:

When an IBM AIX host system is restarted, the host automatically attempts to open any volumes for input/output (I/O) that were open prior to the restart. If Open HyperSwap was enabled for a set of volumes on the host system, the host must determine which storage system is the primary system before the host can allow the volumes to be opened.

If the Metro Mirror relationship for the set of volumes is in a Prepared or Suspended state and the host has connectivity to both the primary and secondary storage systems, the host can determine through the hardware which storage system is the primary system. In this situation, the host automatically opens the volumes.

If the Metro Mirror relationship for the set of volumes is in a Prepared state and the host has connectivity to only the secondary storage system, all I/O to the

volumes might be blocked on the host system until the host is able to verify the primary volume in the relationship. The AIX command varyonvg will fail to open the volumes for I/O to prevent the application from writing to the incorrect site. If the host can determine which volume is the primary volume in the relationship and connectivity to the primary storage system is still lost, a Hyperswap event is triggered. This event causes all I/O to be automatically opened and directed to the secondary storage system.

If the Metro Mirror relationship for the set of volumes is in a Target Available state after a Hyperswap or a Recover command has been issued for the session, or the host system does not have the connectivity necessary to determine which site is the primary site and all I/O to the volumes are blocked on the host system. The varyonvg command will fail to open the volumes for I/O to prevent the application from writing to the incorrect site.

Unblocking I/O on the host system after a host system restart

When any of the previous scenarios cause I/O to be blocked, manual actions might be necessary to remove the block.

If the relationships are in a Target Available state on the hardware, issue a Start command to the session in the desired direction of the relationship. This action defines the primary storage system for the host. The host system can allow the volumes to be opened to I/O.

If the relationships cannot be restarted, or the host cannot determine the primary storage system, it might be necessary to manually decouple the volumes on the host system.

To decouple the volumes the following options are available:

- Option 1: Terminate the session or remove the copy set. This option would require a full copy when the relationships are restarted.
- Option 2: Remove Object Data Manager (ODM) entries using the following command:

```
odmdelete -o Volume_Equivalency
```

CAUTION:

This command should be used only for this scenario because the command deletes copy set information.

Planned and unplanned swaps:

Once a session has been enabled for HyperSwap or Open HyperSwap and reaches the Prepared state, IBM Tivoli Storage Productivity Center for Replication loads the configuration of volumes that are capable of being swapped onto IBM z/OS or AIX.

When the load is complete, the session is capable of a planned or unplanned swap. The H1-H2 role pair on the session show a type of HS. An H is displayed over the connection in the dynamic image for that role pair as shown in the following

Metro Mirror Failover/Failback




figure.

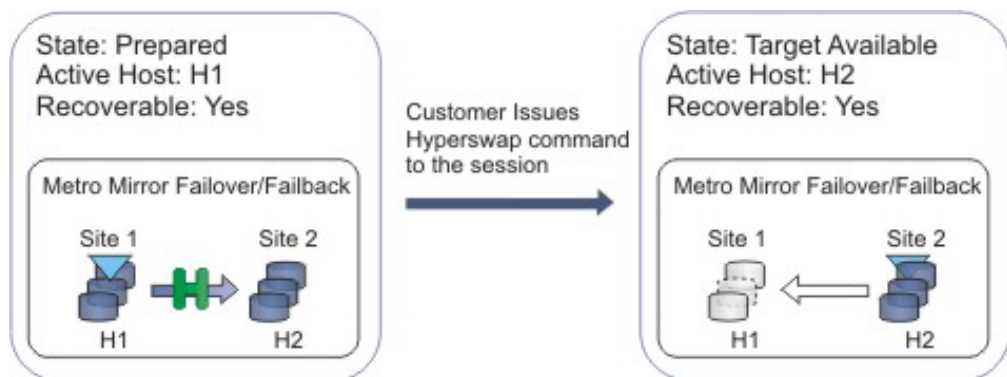
Performing a Planned Swap

Once the session configuration is loaded on z/OS for HyperSwap or AIX for Open HyperSwap, the session is considered swap capable. There may be cases such as a planned maintenance or a migration from the primary storage, in which a planned swap might be required. Once the session is in a swap capable state, a planned swap can be executed by issuing the HyperSwap command against the session.

Once a planned swap is run for z/OS HyperSwap and Open HyperSwap, the session is transitioned to a Target Available state and all the H1-H2 pairs are in a Target Available state. If the H1-H2 role pair was consistent at the time of the swap, the session will have a status of Normal and will indicate that H1-H2 is consistent. If the H1-H2 role pair was not consistent at the time of the swap, the session might display a status of SEVERE because the session is inconsistent. The active host on the session is then displayed as H2.

All input/output (I/O) should have been redirected to the H2 volumes. After a successful swap to site 2, it is not possible to re-enable copy to site 2. Therefore, it is not possible to issue a Start H1->H2 command. The only way to restart the copy is a Start H2->H1 command. To have the volumes protected with high availability and disaster recovery again, the error that caused the swap must be fixed and then the session must be manually restarted to begin copying to the other site.

The following figure illustrates a planned swap.



What happens when an unplanned swap occurs

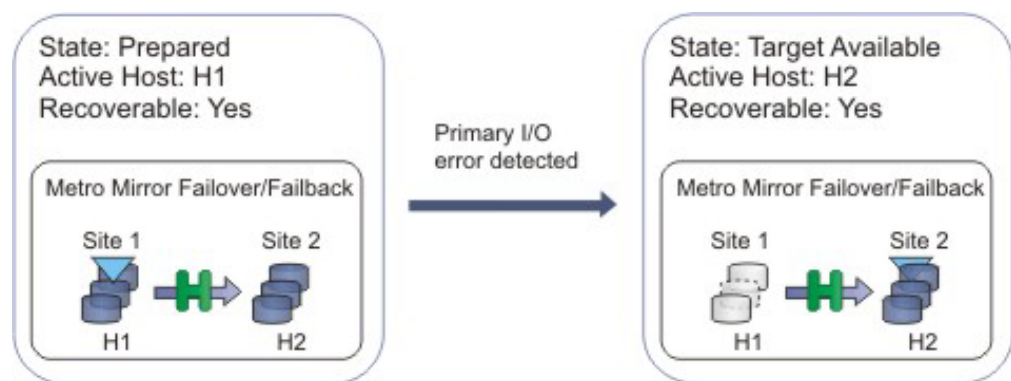
Once the session configuration is loaded on z/OS for HyperSwap or AIX for Open HyperSwap, the session is considered swap capable. In the event of a primary I/O error, a swap occurs automatically. For HyperSwap, z/OS performs the entire swap

and then alerts Tivoli Storage Productivity Center for Replication that a swap has occurred. For Open HyperSwap, Tivoli Storage Productivity Center for Replication and the AIX host work together to perform the swap.

Once an unplanned swap occurs for HyperSwap and Open HyperSwap, the session is transitioned to a Target Available state and all the H1-H2 pairs are in a Target Available state. If the H1-H2 role pair was consistent at the time of the swap, the session will have a Status of Normal and will indicate that H1-H2 is consistent. If the H1-H2 role pair was not consistent at the time of the swap, the session might display a status of SEVERE because the session is inconsistent. The active host on the session is then displayed as H2.

All I/O should have been redirected to the H2 volumes. After a successful swap to site 2, it is not possible to re-enable copy to site 2. Therefore, it is not possible to issue a Start H1->H2 command. The only way to restart the copy is a Start H2->H1 command. To have the volumes protected with high availability and disaster recovery again, the error that caused the swap must be fixed and then the session must be manually restarted to begin copying to the other site.

The following figure illustrates an unplanned swap.



Scenarios requiring a full copy in Metro Global Mirror with HyperSwap sessions

In the following cases, a full copy is required to return to the three-site configuration after a swap:

- If you are running a Metro Global Mirror session with HyperSwap and you issue the HyperSwap command using the z/OS HyperSwap API rather than the Tivoli Storage Productivity Center for Replication graphical user interface (GUI).
- If you are running a Metro Global Mirror session with HyperSwap and an unplanned swap occurs.

Verifying that a session is capable of a planned or unplanned swap:

You can verify whether a sessions is capable of a planned or unplanned swap from the IBM z/OS console (HyperSwap) or the IBM AIX host (Open HyperSwap).

Perform these steps to verify the status of HyperSwap from the z/OS console:

1. Issue the `ds hs,status` command for the overall status of the HyperSwap session. For example:

```

15.03.06 SYSTEM1 d hs,status
15.03.06 SYSTEM1 STC00063 IOSHM0303I HyperSwap Status 531
Replication Session: SR_HS
HyperSwap enabled
New member configuration load failed: Disable
Planned swap recovery: Disable
Unplanned swap recovery: Disable
FreezeAll: No
Stop: No

```

2. Issue the `ds hs,config(detail,all)` command to verify all the volumes in the configuration. For example:

```

15.03.51 SYSTEM1 d hs,config(detail,all)
15.03.51 SYSTEM1 STC00063 IOSHM0304I HyperSwap Configuration 534
Replication Session: SR_HS
Prim. SSID UA DEV# VOLSER Sec. SSID UA DEV# Status
06 02 00F42 8K3602 06 04 00FA2
06 01 00F41 8K3601 06 03 00FA1
06 00 00F40 8K3600 06 02 00FA0

```

Perform these steps to verify the status of Open HyperSwap from the AIX host:

1. Issue the `pcmpath query device` command to see the session association and which path the input/output (I/O) is currently being routed to, which is indicated by an asterisk. For example:

```

host1> pcmpath query device 14

DEV#: 14 DEVICE NAME: hdisk14 TYPE: 2107900 ALGORITHM: Load Balance
SESSION NAME: session1
OS Direction: H1<-H2
=====
PRIMARY SERIAL: 25252520000
-----
Path# Adapter/Path Name State Mode Select Errors
0 fscsi0/path0 CLOSE NORMAL 6091 0
1 fscsi0/path2 CLOSE NORMAL 6300 0
2 fscsi1/path4 CLOSE NORMAL 6294 0
3 fscsi1/path5 CLOSE NORMAL 6187 0

SECONDARY SERIAL: 34343430000 *
-----
Path# Adapter/Path Name State Mode Select Errors
4 fscsi0/path1 CLOSE NORMAL 59463 0
5 fscsi0/path3 CLOSE NORMAL 59250 0
6 fscsi1/path6 CLOSE NORMAL 59258 0
7 fscsi1/path7 CLOSE NORMAL 59364 0

```

Temporarily disabling HyperSwap or Open HyperSwap:

In some situations, it might be necessary to temporarily disable the HyperSwap or Open HyperSwap capabilities for a session.

You might want to disable HyperSwap or Open HyperSwap under the following circumstances:

- Performing maintenance
- Inability for one sysplex member to communicate with one or more volumes

Perform these steps to disable Open HyperSwap for a specific session:

1. In the navigation IBM Tivoli Storage Productivity Center for Replication tree, select **Sessions**. The Sessions panel is displayed.
2. Select the sessions for which you want to disable HyperSwap or Open HyperSwap.

3. Select **View/Modify Properties** from the **Select Actions** list, and click **Go**.
4. Select **Disable HyperSwap** or **Disable Open HyperSwap** and click **OK**.

Tip: On management servers that run IBM z/OS, you can also disable HyperSwap from an MVS command prompt by entering **SETHS DISABLE**.

Using active and standby Tivoli Storage Productivity Center for Replication servers with HyperSwap or Open HyperSwap:

To ensure that there is an IBM Tivoli Storage Productivity Center for Replication server available in the event of a disaster, an active and standby management server configuration can be set up. You can enable HyperSwap and Open HyperSwap for sessions while maintaining an active and standby server configuration.

Active and standby servers with HyperSwap

When the storage system is set up to connect through the z/OS interface, the connection information is automatically sent to the standby server and a connection is attempted. It is possible that the connection will fail if the standby server is not running on z/OS or does not have access to the same volumes. If the connection fails, any takeover done on the standby server will not be able to manage the HyperSwap. On z/OS, if the session configuration had been successfully loaded prior to the HyperSwap, the z/OS system is still capable of performing the HyperSwap. If the z/OS system swaps the volumes but there is no communication possible to the Tivoli Storage Productivity Center for Replication server, the session recognizes that the pairs became suspended and the session will go into a Suspended/Severe state. From this state the customer can clear the **Manage H1-H2 with Hyperswap** option and issue the Recover command for the session to get the session into a Target Available state.

Active and standby servers with Open HyperSwap

When there is an active and standby management server configuration and a host system connection is added to the active server, the host system connection is automatically sent to the standby server and a connection is attempted. Once the session configuration is loaded on IBM AIX, Open HyperSwap is possible only if there is continual communication between AIX and the Tivoli Storage Productivity Center for Replication server. If a takeover is performed on a standby server that is unable to connect to the host system that is managing the swap, the session is no longer Open HyperSwap capable. Communication to the host system must be activated before the session can become Open HyperSwap capable again.

Related tasks:

Chapter 3, “Managing management servers,” on page 85

This section provides information about how to set up active and standby management servers, restore a lost connection between the management servers, or perform a takeover on the standby management server.

Session commands

The commands that are available for a session depend on the session type.

Commands are issued synchronously to IBM Tivoli Storage Productivity Center for Replication sessions. Any subsequent command issued to an individual session is not processed until the first command has completed. Some commands can take an extended amount of time to complete, such as the **Start** command as it sets up the

hardware. The GUI continues to allow you to issue commands to other sessions and not hold up functionality. When a command has completed, the console displays the results of the command.

Basic HyperSwap commands

Use this information to learn about commands available for Basic HyperSwap sessions.

Note: Individual suspend and recover commands are not available in HyperSwap.

Table 4. Basic HyperSwap commands

Command	Action
HyperSwap	Triggers a HyperSwap where I/O is redirected from the source volume to the target volume, without affecting the application using those volumes. You can use this command if you want to perform maintenance on the original source volumes.
Start H1->H2	Starts copying data synchronously from H1 to H2 in a Metro Mirror session. Note: A session might go into a Severe state with error code 1000000 before the session returns to Normal/Prepared State and HyperSwap Capable. The duration of the Severe state depends on how large of a session is running.
Start H2->H1	Starts copying data synchronously from H2 to H1 in a Metro Mirror session. You can issue this command only after the session has been swapped and the production site is H2. To enable data protection when the H1 volumes are available again, start I/O to the H2 volumes, and issue this command to replicate data from the H2 volumes to H1 volumes.
Stop	Suspends updates to all the targets of pairs in a session. You can issue this command at any time during an active session. Note: After you issue the stop command, targets might not be consistent.
Terminate	Removes all physical copies and relationships from the hardware during an active session.

FlashCopy commands

Use this information to learn about commands available for FlashCopy sessions.

Table 5. FlashCopy commands

Command	Action
Start	Performs any steps necessary to define the relationship before performing a FlashCopy operation. For ESS, DS6000, and DS8000, this command is not an option. Issue this command to put the session in the prepared state for the following storage systems: <ul style="list-style-type: none"> • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified

Table 5. FlashCopy commands (continued)

Command	Action
Flash	<p>Performs the FlashCopy operation using the specified options. Issue the Flash command to create a data consistent point-in-time copy for a FlashCopy Session with volumes on the following storage systems:</p> <ul style="list-style-type: none"> • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified <p>For a FlashCopy session containing ESS, DS6000, and DS8000 volumes, the Flash command by itself is not sufficient to create a consistent copy. To create a consistent copy using the ESS, DS6000, and DS8000 Flash commands, you must quiesce application I/O before issuing the Flash command.</p>
InitiateBackgroundCopy	<p>Copies all tracks from the source to the target immediately, instead of waiting until the source track is written to. This command is valid <i>only</i> when the background copy is not already running.</p>
Terminate	<p>Removes all active physical copies and relationships from the hardware during an active session.</p> <p>If you want the targets to be data consistent before removing their relationship, you must issue the InitiateBackgroundCopy command if NOCOPY was specified, and then wait for the background copy to complete by checking the copying status of the pairs.</p>

Snapshot commands

Use this information to learn about commands that are available for Snapshot sessions and groups. A snapshot group is a grouping of snapshots of individual volumes in a consistency group at a specific point in time.

Table 6. Snapshot session commands

Command	Action
Create Snapshot	Creates a snapshot of the volumes in the session
Restore	Restores the H1 volumes in the session from a set of snapshot volumes. You must have at least one snapshot group to restore from. When you issue this command in the Tivoli Storage Productivity Center for Replication graphical user interface (GUI), you are prompted to select the snapshot group.

Table 7. Snapshot group commands

Command	Action
Delete	Deletes the snapshot group and all the individual snapshots that are in the group from the session and from the XIV system. If the deleted snapshot group is the last snapshot group that is associated with the session, the session returns to the Defined state.

Table 7. Snapshot group commands (continued)

Command	Action
Disband	Disbands the snapshot group. When a snapshot group is disbanded, the snapshot group no longer exists. All snapshots in the snapshot group become individual snapshots that are no longer associated to the consistency group or the session. After a snapshot group is disbanded, it is no longer displayed in or managed by Tivoli Storage Productivity Center for Replication. If the disbanded snapshot group is the last snapshot group that is associated with the session, the session returns to the Defined state.
Duplicate	Duplicates the snapshot group. When a snapshot group is duplicated, a new snapshot group is created with new snapshots for all volumes that are in the duplicated group. The name of the duplicated snapshot group is generated automatically by the XIV system.
Lock	Locks a snapshot group. If the snapshot group is locked, write operations to the snapshots that are in the snapshot group are prevented. By default, a snapshot group is locked when it is created. This action is valid only if the snapshot group is unlocked.
Overwrite	Overwrites the snapshot group to reflect the data that is on the H1 volume.
Rename	Renames the snapshot group to a name that you provide. The name can be a maximum of 64 alphanumeric characters.
Restore	Restores the contents of a snapshot group by using another snapshot group in the session. Both of the snapshot groups must contain the same subset of volumes.
Set Priority	Sets the priority in which a snapshot group is deleted. The value can be the number 1 - 4. A value of 1 specifies that the snapshot group is deleted last. A value of 4 specifies that the snapshot group is deleted first.
Unlock	Unlocks a snapshot group. If the snapshot group is unlocked, write operations to the snapshots that are in the snapshot group are enabled and the snapshot group is displayed as modified. This action is valid only if the snapshot group is locked.

Metro Mirror commands

Use this information to learn about commands available for Metro Mirror sessions.

Table 8. Metro Mirror commands

Command	Action
Enable Copy to Site 1	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the Start H2->H1 command becomes available.
Enable Copy to Site 2	Run this command and confirm that you want to reverse the direction of replication before reversing the direction of copying in a failover and failback session. After you issue this command, the Start H1->H2 command becomes available.

Table 8. Metro Mirror commands (continued)

Command	Action
HyperSwap	Triggers a HyperSwap where I/O is redirected from the source volume to the target volume, without affecting the application using those volumes. You can use this command if you want to perform maintenance on the original source volumes.
Start	Establishes a single-direction session with the hardware and begins the synchronization process between the source and target volumes.
Start H1->H2	Establishes Metro Mirror relationships between the H1 volumes and the H2 volumes, and begins data replication from H1 to H2.
Start H2->H1	Establishes Metro Mirror relationships between the H2 volumes and the H1 volumes and starts data replication from H2 to H1. Indicates direction of a failover and failback between two hosts in a Metro Mirror session. If the session has been recovered such that the production site is now H2, you can issue the Start H2->H1 command to start production on H2 and start data replication.
Stop	Inconsistently suspends updates to all the targets of pairs in a session. This command can be issued at any point during an active session. Note: Targets after the suspend are not considered to be consistent.
StartGC	Establishes Global Copy relationships between the H1 volumes and the H2 volumes, and begins asynchronous data replication from H1 to H2. While in the Preparing state, it will not change to the Prepared state unless you switch to Metro Mirror.
Suspend	Causes all target volumes to remain at a data-consistent point and stops all data that is moving to the target volumes. This command can be issued at any point during a session when the data is actively being copied. Note: It is recommended that you avoid using the same LSS pairs for multiple Metro Mirror sessions. Metro Mirror uses a freeze command on ESS, DS6000, and DS8000 storage systems to create the data-consistent point. If there are other Metro Mirror sessions overlapping the same LSS pairs as in this session, those sessions are also suspended. When a Suspend command is issued to a source volume in an LSS that has source volumes in another active Metro Mirror session, the other source volumes are affected only if they have the same target LSS. The primary volumes are suspended, but volumes in the same source LSS that have target volumes in a different LSS are not affected because they use a different PPRC path connection.
Recover	Issues the Recover command to suspended sessions. This command performs the steps necessary to make the target available as the new primary site. Upon completion of this command, the session becomes Target Available.
Terminate	Removes all copy relationships from the hardware during an active session. If you want the targets to be data consistent before removing their relationship, you must issue the Suspend command, then the Recover command, and then the Terminate command.

Metro Mirror with Practice commands

Use this information to learn about commands available for Metro Mirror with Practice sessions.

Table 9. Metro Mirror with Practice commands

Command	Action
Enable Copy to Site 1	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the Start H2->H1 command becomes available.
Enable Copy to Site 2	Run this command and confirm that you want to reverse the direction of replication before reversing the direction of copying in a failover and failback session. After you issue this command, the Start H1->H2 command becomes available.
Flash	Creates a FlashCopy image from I2 volumes to H2 volumes. The amount of time for this to occur will vary depending on the number of copy sets in the session. Note: For ESS, DS6000, DS8000 storage systems, the Flash command uses the freeze and thaw processing to create a data consistent point for the FlashCopy. If there is another Metro Mirror session overlapping on one or more of the same LSS pairs, that session will be suspended. It is also possible that the suspension of the other session might cause the Metro Mirror session to remain suspended after the flash command is issued instead of returning to Prepared state. Avoid using the same LSS pairs for multiple Metro Mirror sessions if possible.
Start H1->H2	Establishes a Metro Mirror relationship from the H1 volumes to the I2 volumes, and begins data replication.
Start H2->H1	Establishes a Metro Mirror relationship from H2 to H1 and begins data replication.
StartGC_H1H2	Distinguishes when the session is in the Preparing state from H1 to I2 and begins the asynchronous process between the source and target volumes. While in the Preparing state the session will not change to the Prepared state unless you switch to Metro Mirror.
StartGC_H2H1	Distinguishes when the session is in the Preparing state from H2 to H1 and begins the asynchronous process between the source and target volumes. While in the Preparing state the session will not change to the Prepared state unless you switch to Metro Mirror.

Table 9. Metro Mirror with Practice commands (continued)

Command	Action
Suspend	<p>Causes all target volumes to remain at a data-consistent point and stops all data that is moving to the target volumes. This command can be issued at any point during a session when the data is actively being copied.</p> <p>Note: The Metro Mirror command uses a freeze command on the ESS, DS6000, or DS8000 devices to create the data-consistent point. If there are other Metro Mirror sessions overlapping the same LSS pairs as in this session, those sessions will also become suspended. Avoid using the same LSS pairs for multiple Metro Mirror sessions.</p> <p>When a Suspend command is issued to a source volume in an LSS that has source volumes in another active Metro Mirror session, the other source volumes are affected only if they have the same target LSS. The primary volumes are suspended, but volumes in the same source LSS that have target volumes in a different LSS are not affected because they use a different PPRC path connection.</p>
Stop	<p>Inconsistently suspends updates to all the targets of pairs in a session. This command can be issued at any point during an active session.</p> <p>Note: Targets after the suspend are not considered to be consistent.</p>
Terminate	Terminates all copy relationships on the hardware.
Recover	<p>Takes a point-in-time copy of the data on I2 to the H2 volumes, enabling the application to be attached and run from the H2 volumes on site 2.</p> <p>Note: The point-in-time copy is performed when the session is in a recoverable state to avoid that previous consistent data on H2 are overwritten by inconsistent data upon Recover. You can issue the Flash command if you want to force a point-in-time copy from I2 to JH2 volumes afterwards.</p>

Global Mirror commands

Use this information to learn about commands available for Global Mirror sessions.

Table 10. Global Mirror commands

Command	Action
Enable Copy to Site 1	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the Start H2->H1 command becomes available.
Enable Copy to Site 2	Run this command and confirm that you want to reverse the direction of replication before reversing the direction of copying in a failover and failback session. After you issue this command, the Start H1->H2 command becomes available.
Start	Establishes all relationships in a single-direction session and begins the process necessary to start forming consistency groups on the hardware.

Table 10. Global Mirror commands (continued)

Command	Action
Start H1->H2	Starts copying data from H1 to H2 in a Global Mirror failover and failback session. Establishes the necessary relationships in the session and begins the process necessary to start copying data from the H1 site to the H2 site and to start forming consistency groups.
Start H2->H1	Starts copying data from H2 to H1 in a failover and failback session for ESS, DS6000 and DS8000 sessions. If a recover has been performed on a session such that the production site is now H2, you can issue a Start H2->H1 to start moving data back to Site 1. However, this start does not provide consistent protection as it copies only asynchronously back because of the long distance. A Global Copy relationship is used. When you are ready to move production back to Site 1, issue a suspend to the session; this puts the relationships into a synchronized state and suspends them consistently. Sessions are consistent when copying H2->H1 for the following storage systems: <ul style="list-style-type: none"> • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified
Suspend	Stops all consistency group formation when the data is actively being copied and suspends the H1->H2 Global Copy pairs. To issue the pause command to the Global Mirror session on the hardware without suspending the Global Copy pairs, open the <code>rmserver.properties</code> file and add the following property to disable the Global Copy suspension on the Suspend command: <code>csm.server.sus_gc_pairs_on_gm_pause = false</code> . The default property is true and the Suspend command automatically suspends the Global Copy pairs.
Recover	Issue this command to recover the session to the target site. This command performs the steps necessary to make the target host volumes consistent and available for access as the new primary site. Upon completion of this command, the session becomes Target Available. Do not access H2 volumes until the Recover command is completed and the session displays Target Available and Recoverable. A Recover to H2 also establishes a point-in-time copy to J2 to preserve the last consistent data.
Terminate	Removes all physical copies and relationships from the hardware during an active session. If you want the targets to be data consistent before removing their relationship, you must issue the Suspend command, the Recover command, and then the Terminate command.

Global Mirror with Practice commands

Use this information to learn about commands available for Global Mirror with Practice sessions.

Table 11. Global Mirror with Practice commands

Command	Action
Enable Copy to Site 1	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the Start H2->H1 command becomes available.

Table 11. Global Mirror with Practice commands (continued)

Command	Action
Enable Copy to Site 2	Run this command and confirm that you want to reverse the direction of replication before reversing the direction of copying in a failover and failback session. After you issue this command, the Start H1->H2 command becomes available.
Flash	The Flash command on a Global Mirror Practice session for ESS, DS6000, and DS8000 temporarily pauses the formation of consistency groups, ensure that all I2s are consistent, and then flash the data from I2 to the H2 volumes. After the flash is complete, the Global Mirror session is automatically restarted, and the session begins forming consistency groups on I2. You can then use the H2 volumes to practice your disaster recovery procedures.
Start H1->H2	Starts copying data from H1 to H2. After the first pass of the copy is complete for all pairs, the session establishes the I2->J2 FlashCopy pairs, and starts the Global Mirror master so that the hardware will begin forming consistency groups, to ensure consistent data is at site 2.
Start H2->H1	Starts copying data from H2 to H1 in a failover and failback session. If a recover has been performed on a session such that the production site is now H2, you can issue a Start H2->H1 to start moving data back to Site 1. However, this start does not provide consistent protection as it copies only asynchronously back because of the long distance. Note: ESS, DS6000, and DS8000 volumes are not consistent for the Start H2->H1 command. A Global Copy relationship is used. When you are ready to move production back to Site 1, issue a suspend to the session; this puts the relationships into a synchronized state and suspends them consistently.
Terminate	Removes all physical copies and relationships on the hardware.
Suspend	Stops all consistency group formation when the data is actively being copied and suspends the H1-I2 Global Copy pairs. To issue the pause command to the Global Mirror session on the hardware without suspending the Global Copy pairs, open the <code>rmserver.properties</code> file and add the following property to disable the Global Copy suspension on the Suspend command: <code>csm.server.sus_gc_pairs_on_gm_pause = false</code> . The default property is true and the Suspend command automatically suspends the Global Copy pairs..
Recover	Restores consistent data on I2 volumes and takes a point-in-time copy of the data on I2 to the H2 volumes, enabling the application to be attached and run from the H2 volumes on site 2. The I2 Volumes will continue to hold the consistent data and can be flashed again to H2 using the Flash command.

Metro Global Mirror commands

Use this information to learn about commands available for Metro Global Mirror sessions.

Table 12. Metro Global Mirror commands

Command	Action
Enable Copy to Site 1	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the Start H2->H1->H3 command becomes available.

Table 12. Metro Global Mirror commands (continued)

Command	Action
Enable Copy to Site 2	Run this command and confirm that you want to reverse the direction of replication before reversing the direction of copying in a failover and failback session. After you issue this command, the Start H1->H2->H3 command becomes available.
HyperSwap	<p>Causes a site switch, equivalent to a suspend and recover for a Metro Mirror with failover and failback individual suspend and recover commands are not available.</p> <p>The Metro Global Mirror for HyperSwap supports IBM Tivoli Storage Productivity Center for Replication installed on z/OS with the IBM Tivoli Storage Productivity Center for Replication Three Site Business Continuity license.</p>
Start H1->H2->H3	<p>(This is the Metro Global Mirror initial start command.) Establishes Metro Mirror relationships between H1 and H2, and Global Mirror relationships between H2 and H3. For Metro Global Mirror, this includes the J3 volume to complete the Global Mirror configuration. (The J3 volume role is the journal volume at site 3.). Start H1->H2->H3 can be used from some Metro Global Mirror configurations to transition back to the starting H1->H2->H3 configuration.</p> <p>This command is valid only when the session is in a defined, preparing, prepared, or suspended state.</p>
Start H1->H3	<p>From the H1->H2->H3 configuration, this command changes the session configuration to a Global Mirror-only session between H1 and H3, with H1 as the source. Use this command in case of an H2 failure with transition bitmap support provided by incremental resynchronization. It can be used when a session is in preparing, prepared, and suspended states because there is not a source host change involved.</p> <p>This command allows you to bypass the H2 volume in case of an H2 failure and copy only the changed tracks and tracks in flight from H1 to H3. After the incremental resynchronization is performed, the session is running Global Mirror from H1 to H3 and thus loses the near-zero data loss protection achieved with Metro Mirror when running H1->H2->H3. However, data consistency is still maintained at the remote site with the Global Mirror solution.</p> <p>From H2->H1->H3 configuration, this command changes the session configuration to a Global Mirror-only session configuration between H1 and H3, with H1 as the source. Use this command when the source site has a failure and production is moved to the H1 site. This can be done for unplanned HyperSwap. The Global Mirror session is continued. This is a host-volume change so this command is valid only when restarting the H1->H3 configuration or from the TargetAvailable H2->H1->H3 state.</p>

Table 12. Metro Global Mirror commands (continued)

Command	Action
Start H2->H3	<p>From the H1->H2->H3 configuration, this command moves the session configuration to a configuration between H2 and H3, with H2 as the source. Use this command when the source site has a failure and production is moved to the H2 site, for example, for unplanned HyperSwap. The Global Mirror session is continued. This session is a host-volume change so this command is valid only when restarting the H1->H3 configuration or from the TargetAvailable H2->H1>H3 state.</p> <p>From the H2->H1->H3 configuration, this command changes the session configuration to a configuration between H2 and H3 with H2 as the source. Use this command in case of an H1 failure with transition bitmap support provided by incremental resynchronization. Because there is not a source-host change involved, it can be used when the session is in preparing, prepared, and suspended states Start H2->H1->H3 can be used to transition back to the starting H2->H1->H3 configuration.</p>
Start H2->H1->H3	<p>(This is the Metro Global Mirror start command.) This is the configuration that completes the HyperSwap processing. This command creates Metro Mirror relationships between H2 and H1 and Global Mirror relationships between H1 and H3. For Metro Global Mirror, this includes the J3 volume to complete the Global Mirror configuration.</p>
Start H3->H1->H2	<p>After recovering to H3, this command sets up the hardware to allow the application to begin writing to H3, and the data is copied back to H1 and H2. However, issuing this command does not guarantee consistency in the case of a disaster because only Global Copy relationships are established to cover the long distance copy back to site 1.</p> <p>To move the application back to H1, you can issue a suspend while in this state to drive all the relationships to a consistent state and then issue a freeze to make the session consistent. You can then issue a Rcover followed by a Start H1->H2->H3 to go back to the original configuration.</p>
SuspendH2H3	<p>When running H1->H2->H3, this command issues a pause to the Global Mirror master and causes the Global Mirror master to stop forming consistency groups.</p> <p>This command is valid only when the session is in a prepared state. Stops all consistency group formation when the data is actively being copied and suspends the H2>H3 Global Copy pairs. To issue the pause command to the Global Mirror session on the hardware without suspending the Global Copy pairs, open the <code>rmserver.properties</code> file and add the following property to disable the Global Copy suspension on the Suspend command: csm.server.sus_gc_pairs_on_gm_pause = false. The default property is true and the Suspend command automatically suspends the Global Copy pairs..</p>

Table 12. Metro Global Mirror commands (continued)

Command	Action
SuspendH1H3	<p>When running H2->H1->H3, this command issues a pause to the Global Mirror master and causes the Global Mirror master to stop forming consistency groups.</p> <p>This command is valid only when the session is in a prepared state. Stops all consistency group formation when the data is actively being copied and suspends the H2->H3 Global Copy pairs. To issue the pause command to the Global Mirror session on the hardware without suspending the Global Copy pairs, open the <code>rmsserver.properties</code> file and add the following property to disable the Global Copy suspension on the Suspend command: <code>csm.server.sus_gc_pairs_on_gm_pause = false</code>. The default property is true and the Suspend command automatically suspends the Global Copy pairs..</p>
RecoverH1	<p>Specifying H1 makes the H1 volume TargetAvailable. Metro Global Mirror (when running H2->H1->H3) can move production to either the H1 or H3 set of volumes. IBM Tivoli Storage Productivity Center for Replication processing differs, depending on the recovery site. Therefore, the site designation is added to the Recover command so IBM Tivoli Storage Productivity Center for Replication can set up for the failback.</p>
RecoverH2	<p>Specifying H2 makes the H2 volume TargetAvailable. Metro Global Mirror (when running H1->H2->H3) can move production to either the H2 or H3 set of volumes. IBM Tivoli Storage Productivity Center for Replication processing is different depending on the recovery site. Therefore the site designation is added to the Recover command so IBM Tivoli Storage Productivity Center for Replication can prepare for the failback.</p>
RecoverH3	<p>Specifying H3 makes the H3 volume TargetAvailable. Metro Global Mirror (when running H1->H2->H3) can then move production to the H3 set of volumes. Because IBM Tivoli Storage Productivity Center for Replication processing differs depending on the recovery site, the site designation is added to the Recover command so that IBM Tivoli Storage Productivity Center for Replication can prepare for the failback.</p> <p>This command prepares H3 so that you can start the application on H3. H3 becomes the active host, and you then have the option start H3->H1->H2 to perform a Global Copy copy back. The recovery establishes point-in-time copy to J3 volumes to preserve the last consistent data.</p>

Metro Global Mirror with Practice commands

Use this information to learn about commands available for Metro Global Mirror with Practice sessions.

Table 13. Metro Global Mirror with Practice commands

Command	Action
Enable Copy to Site 1	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session.
Enable Copy to Site 2	Run this command and confirm that you want to reverse the direction of replication before reversing the direction of copying in a failover and failback session.

Table 13. Metro Global Mirror with Practice commands (continued)

Command	Action
Flash	<p>This command is available in the following states:</p> <ul style="list-style-type: none"> • Target Available state when the active host is H3 Note: Use this command if the FlashCopy portion of the Recover command from I3 to H3, fails for any reason. The problem can be addressed; and a Flash command issued to complete the flash of the consistent data from I3 to H3. • Prepared state when the active host is H1 and data is copying H1 to H2 to I3, or the active host is H2 and data is copying H2 to H1 to H3. • Prepared state when the active host is H2 and data is copying H2 to I3. • Prepared state when the active host is H1 and data is copying H1 to I3. <p>Use this command if the FlashCopy portion of the Recover command from I3 to H3, fails for any reason. The problem can be addressed, and a Flash command issued to complete the flash of the consistent data from I3 to H3.</p> <p>Issuing a Flash command on a Global Mirror Practice session for ESS, DS6000, and DS8000 will temporarily pause the formation of consistency groups, ensure that all I3s are consistent, and then flash the data from I3 to the H3 volumes. After the flash is complete, the Global Mirror session will be automatically restarted, and the session will begin forming consistency groups on I3. You can then use the H3 volumes to practice your disaster recovery procedures.</p>
RecoverH1	<p>Specifying H1 makes the H1 volume TargetAvailable. Metro Global Mirror (when running H1->H2->H3) can move production to either the H2 or H3 set of volumes. IBM Tivoli Storage Productivity Center for Replication processing differs, depending on the recovery site. Therefore the site designation is added to the Recover command so IBM Tivoli Storage Productivity Center for Replication can prepare for the failback. The FlashCopy creates a consistent copy of the data on the H3 volumes so that an application can recover to those volumes and begin writing I/O. When the FlashCopy is complete, the session will reach a Target Available state, and you can attach your volumes on Site 3.</p>
RecoverH2	<p>Specifying H2 makes the H2 volume TargetAvailable. Metro Global Mirror (when running H1->H2->H3) can move production to either the H2 or H3 set of volumes. IBM Tivoli Storage Productivity Center for Replication processing differs, depending on the recovery site. Therefore the site designation is added to the Recover command so IBM Tivoli Storage Productivity Center for Replication can prepare for the failback.</p>
RecoverH3	<p>Specifying H3 makes the H3 volume the TargetAvailable. When running H1->H2->H3, Metro Global Mirror can move production to either the H2 or H3 set of volumes. IBM Tivoli Storage Productivity Center for Replication processing differs, depending on the recovery site; therefore, the site designation is added to the Recover command so IBM Tivoli Storage Productivity Center for Replication can prepare for the failback.</p>
Re-enable Copy to Site 1	<p>After issuing a RecoverH1 command, you can run this command to restart the copy to the original the direction of replication in a failover and failback session.</p>

Table 13. Metro Global Mirror with Practice commands (continued)

Command	Action
Re-enable Copy to Site 2	After issuing a RecoverH2 command, you can run this command to restart the copy to the original the direction of replication in a failover and failback session.
Re-enable Copy to Site 3	After issuing a RecoverH3 command, you can run this command to restart the copy to the original the direction of replication in a failover and failback session.
Start H1->H2->H3	<p>Metro Global Mirror initial start command. This command creates Metro Mirror relationships between H1 and H2, and Global Mirror relationships between H2 and H3. For Metro Global Mirror, this includes the J3 volume to complete the Global Mirror configuration. (The J3 volume role is the journal volume at site 3.). Start H1->H2->H3 can be used from some Metro Global Mirror configurations to return to the starting H1>H2>H3 configuration.</p> <p>This command is valid only when the session is in a defined, preparing, prepared, target-available, or suspended state.</p>
Start H1->H3	<p>From the H1->H2->H3 configuration, this command changes the session configuration to a Global-Mirror-only session between H1 and H3, with H1 as the source. Use this command in case of an H2 failure with transition bitmap support provided by incremental resynchronization. Because there is not a source host change involved, it can be used when a session is in preparing, prepared, and suspended states.</p> <p>You can use this command to bypass the H2 volume in case of an H2 failure and copy only the changed tracks and tracks in flight from H1 to H3. After the incremental resynchronization is performed, the session is running Global Mirror from H1 to H3 and thus loses the near-zero data loss protection achieved with Metro Mirror when running H1->H2->H3. However, data consistency is still maintained at the remote site with the Global Mirror solution.</p> <p>From H2->H1->H3 configuration, this command changes the session configuration to a Global-Mirror-only session configuration between H1 and H3, with H1 as the source. Use this command when the source site has a failure and production is moved to the H1 site. This can be done for unplanned HyperSwap. The Global Mirror session is continued. This is a host-volume change so this command is valid only when restarting the H1->H3 configuration or from the TargetAvailable H2->H1->H3 state.</p>

Table 13. Metro Global Mirror with Practice commands (continued)

Command	Action
Start H2->H3	<p>From the H1->H2->H3 configuration, this command moves the session configuration to a configuration between H2 and H3, with H2 as the source. Use this command when the source site has a failure and production is moved to the H2 site. This can be done for unplanned HyperSwap. The Global Mirror session is continued. This is a host-volume change so this command is valid only when restarting the H1>H3 configuration or from the TargetAvailable H2->H1->H3 state.</p> <p>From the H2->H1->H3 configuration, this command changes the session configuration to a configuration between H2 and H3 with H2 as the source. Use this command in case of an H1 failure with transition bitmap support provided by incremental resynchronization. Because there is not a source-host change involved it can be used when the session is in preparing, prepared, and suspended states. Start H2->H1->H3 can be used to return to the starting H2->H1->H3 configuration.</p>
SuspendH2H3	<p>When running H1->H2->H3, this command issues a pause to the Global Mirror master and causes the Global Mirror master to stop forming consistency groups.</p> <p>This command is valid only when the session is in a prepared state. It stops all consistency group formation when the data is actively being copied and suspends the H2->H3 Global Copy pairs. To issue the pause command to the Global Mirror session on the hardware without suspending the Global Copy pairs, open the <code>rmsvr.properties</code> file and add the following property to disable the Global Copy suspension on the Suspend command: <code>csm.server.sus_gc_pairs_on_gm_pause = false</code>. The default property is true and the Suspend command automatically suspends the Global Copy pairs..</p>
SuspendH1H3	<p>When running H2->H1->H3, this command issues a pause to the Global Mirror master and causes the Global Mirror master to stop forming consistency groups.</p> <p>This command is valid only when the session is in a prepared state. Stops all consistency group formation when the data is actively being copied and suspends the H2->H3 Global Copy pairs. To issue the pause command to the Global Mirror session on the hardware without suspending the Global Copy pairs, open the <code>rmsvr.properties</code> file and add the following property to disable the Global Copy suspension on the Suspend command: <code>csm.server.sus_gc_pairs_on_gm_pause = false</code>. The default property is true and the Suspend command automatically suspends the Global Copy pairs..</p>
Terminate	This command terminates all copy relationships on the hardware.

Metro Mirror heartbeat

The *heartbeat* is a Metro Mirror function. When the Metro Mirror heartbeat is disabled, data consistency across multiple storage systems is not guaranteed if the IBM Tivoli Storage Productivity Center for Replication management server cannot communicate with one or more storage systems. The problem occurs as a result of the Hardware Freeze Timeout Timer function within the storage system. If the controlling software loses connection to a storage system, the Metro Mirror relationships that it is controlling stay established and there is no way to *freeze*

those pairs to create consistency across the multiple storage systems. When the freeze times out, dependent I/O is written to the target storage systems, which might corrupt data consistency. Freeze refers to a Metro Mirror (peer-to-peer remote copy [PPRC]) freeze function.

When determining whether to use the Metro Mirror heartbeat, analyze your business needs. Disabling the Metro Mirror heartbeat might result in data inconsistency. If you enable the Metro Mirror heartbeat and a freeze occurs, your applications will be unable to write during the freeze.

Metro Mirror heartbeat is disabled by default.

Metro Mirror heartbeat is not available for Metro Mirror with HyperSwap or Metro Global Mirror with HyperSwap.

There are two cases where lost communication between the coordination software (controller) and one or more storage systems can result in data consistency loss:

Freeze event not detected by a disconnected storage system

Consider a situation with four storage system machines in a primary site and four in a secondary site. One of the four storage systems on the primary loses the connection to the target site. This causes the affected storage system to prevent any writes from occurring, for a period determined by the Freeze timeout timer. At the same time, the affected storage controller loses communication with the controlling software and cannot communicate the Freeze event to the software.

Unaware of the problem, the controlling software does not issue the Freeze command to the remaining source storage systems. The freeze will stop dependent writes from being written to connected storage systems. However, once the Freeze times out and the long-busy is terminated, dependent write I/Os continue to be copied from the storage systems that did not receive the Freeze command. The Metro Mirror session is left in a state where one storage system has suspended copying while the other three storage systems are still copying data. This state causes inconsistent data on the target storage systems.

Freeze event detected, but unable to propagate the Freeze command to all storage systems

Consider a situation with four storage system machines in a primary site and four in a secondary site. One of the four storage systems on the primary loses the connection to the target site. This causes the affected storage system to issue long-busy to the applications for a period determined by the Freeze timeout timer. At the same time, one of the remaining three source systems loses communications with the controlling software.

The storage system that had an error writing to its target cannot communicate the Freeze event to the controlling software. The controlling software issues the Freeze command to all but the disconnected storage system (the one that lost communication with the software). The long-busy stops dependent writes from being written to the connected storage systems.

However, once the Freeze times out on the frozen storage system and the long-busy is terminated, dependent write I/Os continue to the target storage system from the source storage system that lost communication and did not receive the Freeze command. The Metro Mirror session is left a

state where three storage systems have suspended copying and one storage system is still copying data. This state causes inconsistent data on the target storage systems.

Before IBM Tivoli Storage Productivity Center for Replication V3.1, if the controlling software within a Metro Mirror environment detected that a managed storage system lost its connection to its target, the controlling software stopped all the other source systems to ensure consistency across all the targets. However, if the controlling software lost communication with any of the source subsystems during the failure, it could not notify those storage systems of the freeze event or ensure data consistency. The Metro Mirror heartbeat helps to overcome this problem. In a high-availability configuration, the Metro Mirror heartbeat is continued by the standby server after the Takeover command is issued on the standby, enabling you to perform actions on the standby server without causing a freeze.

IBM Tivoli Storage Productivity Center for Replication registers with the managed ESS 800, DS6000 or DS8000 storage systems within a Metro Mirror session when the start command is issued to the session. After this registration occurs, a constant heartbeat is sent to the storage system. If the storage system does not receive a heartbeat from the IBM Tivoli Storage Productivity Center for Replication management server within the allotted time (a subset of lowest LSS timeout value across all the source LSSs), the storage system initiates a freeze. If IBM Tivoli Storage Productivity Center for Replication did not successfully communicate with the storage system, it initiates a freeze on the remaining storage system after the allotted time is expired.

Note: It is recommended that you avoid using the same LSS pairs for multiple Metro Mirror sessions. Metro Mirror uses a freeze command on ESS, DS6000, and DS8000 storage systems to create the data-consistent point. If there are other Metro Mirror sessions overlapping the same LSS pairs as in this session, those sessions are also suspended.

When you are using the Metro Mirror heartbeat, be aware that:

- The Metro Mirror heartbeat can cause a single point of failure: if an error occurs on just the management server and not the storage system, a freeze might occur.
- When the Metro Mirror heartbeat timeout occurs, the storage system remains in a long busy state for the duration of the LSS freeze timeout.

Note: If Metro Mirror heartbeat is enabled for storage systems that are connected through a HMC connection, a connection loss might cause lost Metro Mirror heartbeats, resulting in Freeze actions with application I/O impact for configured Extended Long Busy timeout.

The Metro Mirror heartbeat is supported on storage systems connected through a TCP/IP (direct connect or HMC) connection. It is not supported on storage systems connected through a z/OS connection. Enabling the Metro Mirror heartbeat with a z/OS connection does not fail; however, a warning message is displayed specifying that the Metro Mirror heartbeat function does not work unless you have an IP connection.

If Metro Mirror heartbeat is enabled for storage systems that are connected through a TCP/IP (either direct connect or HMC) connection and z/OS connection,

and the TCP/IP connection fails, IBM Tivoli Storage Productivity Center for Replication suspends the Metro Mirror session because there is no heartbeat through the z/OS connection.

If Metro Mirror heartbeat is enabled for storage systems that are connected through a TCP/IP connection and z/OS connection and you remove all TCP/IP connections, IBM Tivoli Storage Productivity Center for Replication suspends the Metro Mirror sessions and the applications using those volume will be in Extended Long Busy timeout until the storage system's internal timeout timer expires. Ensure that you disable the Metro Mirror heartbeat for all Metro Mirror sessions before removing the last TCP/IP connection to avoid the Extended Long Busy timeout.

Site awareness

You can associate a location with each storage system and each site in a session. This *site awareness* ensures that only the volumes whose location matches the location of the site are allowed for selection when you add copy sets to the session. This prevents a session relationship from being established in the wrong direction.

Note: To filter the locations for site awareness, you must first assign a site location to each storage system.

IBM Tivoli Storage Productivity Center for Replication does not perform automatic discovery of locations. Locations are user-defined and specified manually.

You can change the location associated with a storage system that has been added to the IBM Tivoli Storage Productivity Center for Replication configuration. You can choose an existing location or add a new one. Locations are deleted when there is no longer a storage system with an association to that location.

When adding a copy set to a session, a list of candidate storage systems is presented, organized by location. Storage systems that do not have a location are displayed and available for use when you create a copy set.

You can also change the location for any site in a session. Changing the location of a session does not affect the location of the storage systems that are in the session.

Changing the location of a storage system might have consequences. When a session has a volume role with a location that is linked to the location of the storage system, changing the location of the storage system could change the session's volume role location. For example, if there is one storage system with the location of A_Location and a session with the location of A_Location for its H1 role, changing the location of the storage system to a different location, such as B_Location, also changes the session's H1 location to Site 1. However, if there is a second storage system that has the location of A_Location, the session's role location is not changed.

Important: Location matching is enabled only when adding copy sets. If you change the location of a storage system or volume role, IBM Tivoli Storage Productivity Center for Replication does not audit existing copy sets to confirm or deny location mismatches.

Users and groups

For authentication and authorization, IBM Tivoli Storage Productivity Center for Replication uses users and groups that are defined in a configured user registry on

the management server, which is associated with either the local operating system or a Lightweight Directory Access Protocol (LDAP) server.

IBM Tivoli Storage Productivity Center for Replication does not provide the capability to create, update, or delete users or groups in the user registry. To manage users or groups, you must use the appropriate tool associated with the user registry in which the users and groups are stored.

IBM Tivoli Storage Productivity Center for Replication uses roles to authorize users to manage certain sessions and perform certain actions.

For more information about authentication, see information about single sign-on in the IBM Tivoli Storage Productivity Center documentation.

Primary administrative ID

If you switch the authentication method, either from the local operating system to an LDAP server or vice versa, IBM Tivoli Storage Productivity Center for Replication removes all access to existing users and user groups. This occurs because the user IDs might not be on the same local operating system and the LDAP server; however, you must have at least one user ID that can log in to IBM Tivoli Storage Productivity Center for Replication.

When you change the authentication method using Tivoli Integrated Portal, you can specify a primary administrative ID for both local operating system and LDAP authentication. Use this primary administrator to log in to IBM Tivoli Storage Productivity Center for Replication and manually add user IDs requiring access to IBM Tivoli Storage Productivity Center for Replication.

You can log in to both IBM Tivoli Storage Productivity Center for Replication and IBM Tivoli Storage Productivity Center using the primary administrative ID and password.

You cannot use the following characters for the IBM Tivoli Storage Productivity Center for Replication administrative password:

- square brackets ([and])
- semicolon (;)
- backward slash (\)

User roles

A *user role* is a set of privileges that is assigned to a user or user group to allow the user or user group to perform certain tasks and manage certain sessions.

To be assigned to a role, each user or group of users must have a valid user ID or group ID in the user registry on the management server.

Both individual users and a group of users can be assigned to a role. All users in a group are assigned the role of the group. If a user is assigned to one role as an individual and a different role as a member of a group, the user has access to the permissions of the role that has greater access.

Restricting access to sessions prevents unwarranted administrative access. This is especially useful in an open environment, where there can be many storage administrators who are responsible for their servers, applications, databases, file systems, and so on.

IBM Tivoli Storage Productivity Center for Replication provides a set of predefined user roles: monitor, session operator, and administrator.

Monitor

Monitors can view the health and status in the IBM Tivoli Storage Productivity Center for Replication GUI and CLI; however, they cannot modify or perform any commands or actions.

Monitors can view the following information:

- All storage systems and storage system details
- All connections and connection details
- All sessions and session details
- All path information
- Management server status and details

Operator

Session operators can manage sessions to which they have been assigned, including:

- Adding or removing a session. The user ID that created the session is automatically granted access to manage that session.
- Performing actions on an assigned session, such as start, flash, terminate, and suspend.
- Modifying session properties.
- Adding copy sets to a session. The session operator can add volumes to a copy set only when the volume is not protected and not in another session.
- Removing copy sets from a session.
- Adding Peer To Peer Remote Copy (PPRC) paths, and removing paths with no hardware relationships. PPRC paths are a common resource used in IBM Tivoli Storage Productivity Center for Replication sessions and also in an ESS, DS6000, or DS8000 storage-system relationship that is established between two common logical subsystems (LSSs).

Note: The session operator cannot issue a force removal of a path.

Note: A path can also be auto-generated when starting a session.

- Monitoring health and status, including viewing the following information:
 - All storage systems and storage system details
 - All connections and connection details
 - All sessions and session details
 - All path information
 - Management server status and details

Note: Session operators can make changes only to the volumes that they own. They are not able to make changes to volumes being managed by other users.

Administrator

During installation of IBM Tivoli Storage Productivity Center for Replication, the installation wizard requests an ID to use for the initial administrator user ID.

Administrators have unrestricted access. They can manage all sessions and perform all actions associated with IBM Tivoli Storage Productivity Center for Replication, including:

- Granting permissions to users and groups of users.
- Adding or removing a session. The user ID that created the session is automatically granted access manage that session.
- Performing actions on all sessions, such as start, flash, terminate, and suspend.
- Modifying session properties.
- Adding and removing copy sets from a session. The administrator can add volumes to a copy set only when the volume is not protected and not in another session.
- Protecting volumes and removing volume protection.
- Adding or removing storage system connections.
- Modifying connection properties.
- Assigning or changing storage system locations.
- Adding PPRC paths and removing paths with no hardware relationships. PPRC paths are a common resource used in IBM Tivoli Storage Productivity Center for Replication sessions and also in an ESS, DS6000, or DS8000 storage-system relationship that is established between two common logical subsystems (LSSs).

Note: A path can also be auto-generated when starting a session.

- Managing management servers. The standby management server is a common resource that is available to multiple sessions.
- Packaging program error (PE) log files.
- Monitoring health and status, including viewing the following information:
 - All storage systems and storage system details
 - All connections and connection details
 - All sessions and session details
 - All path information
 - Management server status and details

Important: IBM Tivoli Storage Productivity Center supports multiple user roles, including the Superuser role. A superuser can perform all IBM Tivoli Storage Productivity Center functions. For IBM Tivoli Storage Productivity Center superusers to have full access to IBM Tivoli Storage Productivity Center for Replication, the Superuser group must be added to the IBM Tivoli Storage Productivity Center for Replication and assigned to the Administrator role. Then, you can manage the IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication products by groups, instead of by user IDs.

Note: Administrators cannot revoke their own administrative access rights.

Planning for Open HyperSwap replication

Open HyperSwap replication is a special Metro Mirror replication method designed to automatically failover I/O from the primary logical devices to the secondary logical devices in the event of a primary disk storage system failure. This function can be done with minimal disruption to the applications that are using the logical devices.

Overview

Open HyperSwap replication applies to both planned and unplanned replication sessions. When a session has Open HyperSwap enabled, an I/O error on the primary site automatically causes the I/O to switch to the secondary site without any user interaction and with minimal application impact. In addition, while Open HyperSwap is enabled, the Metro Mirror session supports disaster recovery. If a write is successful on the primary site but is unable to get replicated on the secondary site, IBM Tivoli Storage Productivity Center for Replication suspends the entire set of data consistency checking, thus ensuring that a consistent copy of the data exists on the secondary site. If the system fails, this data might not be the latest data, but the data should be consistent and allow the user to manually switch host servers to the secondary site.

You can control Open HyperSwap from any system running IBM Tivoli Storage Productivity Center for Replication (AIX, Windows, Linux, or z/OS). However, the volumes that are involved with Open HyperSwap must be attached to an AIX system. The AIX system is then connected to Tivoli Storage Productivity Center for Replication.

Software and hardware requirements

There are several requirements for Open HyperSwap support:

AIX requirements

Open HyperSwap support requires AIX version 5.3 or 6.1. (You can find the supported AIX version for each Tivoli Storage Productivity Center for Replication release in the support matrix at http://www-01.ibm.com/support/docview.wss?rs=40&context=SSBSEX&context=SSMN28&context=SSMMUP&context=SS8JB5&context=SS8JFM&uid=swg21386446&loc=en_US&cs=utf-8&lang=en. Click the link for the applicable release under Agents, Servers and GUI.)

You must have the following AIX modules installed:

- Subsystem Device Driver Path Control Module (SDDPCM) version 3.0.0.0 or later
- Multi-Path Input/Output (MPIO) module (the version that is provided with AIX version 5.3 or 6.1)

DS8000 hardware requirements

Only DS8000 storage systems are supported. Open HyperSwap requires DS8000 5.1 or later.

Note: Open HyperSwap does not support PowerHA[®] (previously known as High Availability Cluster Multi-processing (HACMP)).

General tasks

Before you can use Open HyperSwap, you must set up your environment for this function. The general steps are:

1. Prepare the AIX system for Open HyperSwap. Use the AIX configuration manager (**cfgmgr**) to identify all volumes that are involved with the Open HyperSwap session.
2. Set up the host connection of Tivoli Storage Productivity Center for Replication to the AIX system. Use the Tivoli Storage Productivity Center for Replication

- user interface to manually set up the connection to the AIX system. Use the Host Systems page to enter the IP address and port number for the AIX system.
3. Set up the Tivoli Storage Productivity Center for Replication Metro Mirror Failover/Failback session, selecting the function **Manage H1-H2 with Open HyperSwap**.
 4. Add the copy sets to the session where all the volumes in the copy sets are volumes that are on the AIX system that is connected to Tivoli Storage Productivity Center for Replication.
 5. You can now start your Open HyperSwap session the same as other sessions.

Chapter 2. Administering

Administer IBM Tivoli Storage Productivity Center for Replication to authorize users, start and use the graphical user interface (GUI), start and stop services, and many other administrative tasks.

Starting and stopping IBM Tivoli Storage Productivity Center for Replication

Use these procedures to start and stop IBM Tivoli Storage Productivity Center for Replication, including the embedded IBM WebSphere Application Server and DB2.

Starting IBM Tivoli Storage Productivity Center for Replication

You start IBM Tivoli Storage Productivity Center for Replication by starting the embedded IBM WebSphere Application Server.

Starting IBM Tivoli Storage Productivity Center for Replication on Windows

To start IBM Tivoli Storage Productivity Center for Replication on Windows, perform one of these procedures:

- From the desktop, perform these steps:
 1. Click **Start > Control Panel > Administrative Tools > Services**.
 2. Right-click IBM WebSphere Application Server, and then click **Start**.
- From the command line, enter the following command:

```
install_root/AppServer/profiles/default/bin/startServer.bat server1
```

Starting IBM Tivoli Storage Productivity Center for Replication on AIX and Linux

To start IBM Tivoli Storage Productivity Center for Replication on AIX and Linux, enter the following command from the command line:

```
install_root/AppServer/profiles/default/bin/startServer.sh server1
```

Starting IBM Tivoli Storage Productivity Center for Replication on z/OS

For IBM Tivoli Storage Productivity Center for Replication Basic Edition for System z or IBM Tivoli Storage Productivity Center for Replication for System z that uses IBM WebSphere Application Server OEM Edition for z/OS, see the *IBM WebSphere Application Server OEM Edition for z/OS Configuration Guide* for information about how to start IBM WebSphere Application Server OEM Edition for z/OS.

For IBM Tivoli Storage Productivity Center for Replication for System z that uses IBM WebSphere Application Server on z/OS, perform these steps to start IBM Tivoli Storage Productivity Center for Replication:

1. Ensure that the IBM WebSphere Application Server hierarchical file system (HFS), DB2 HFS, and IBM Tivoli Storage Productivity Center for Replication HFS are all mounted on the UNIX System Services (USS). The root installation

directory for IBM Tivoli Storage Productivity Center for Replication is
-PathPrefix-/usr/lpp/Tivoli/RM. The root production directory is
-PathPrefix-/var/lpp/Tivoli/RM.

2. Run the **TSO ISH** command to obtain the name of the USS mount point for the instance as the ENV parameter (for example, SYS7.SYS7.BBOS001). This parameter is typically found in the /WAS_HOME/ directory. This parameter contains the path to several IBM WebSphere Application Server scripts.
3. Run the following command from the System Display and Search Facility (SDSF) panel to start the IBM WebSphere Application Server:

```
/START BB06ACR,JOBNAME=BBOS001,ENV=ENV_parameter
```

The initiator is BB06ACR, and the job name is BBOS001. When you start a BB06ACR, the BBODMNB also starts.

Stopping IBM Tivoli Storage Productivity Center for Replication

You stop IBM Tivoli Storage Productivity Center for Replication by stopping the embedded IBM WebSphere Application Server.

Stopping IBM Tivoli Storage Productivity Center for Replication on Windows

To stop IBM Tivoli Storage Productivity Center for Replication on Windows, perform one of these procedures:

- From the desktop, perform these steps:
 1. Click **Start > Control Panel > Administrative Tools > Services**.
 2. Right-click IBM IBM WebSphere Application Server, and then click **Stop**.
- From the command line, enter the following command:

```
install_root/AppServer/profiles/default/bin/stopServer.bat server1
```

Stopping IBM Tivoli Storage Productivity Center for Replication on AIX and Linux

To stop IBM Tivoli Storage Productivity Center for Replication on AIX and Linux, enter the following command from the command line:

```
install_root/AppServer/profiles/default/bin/stopServer.sh -username user_name  
-password password server1
```

Where *user_name* and *password* are your IBM WebSphere Application Server user name and password.

Stopping IBM Tivoli Storage Productivity Center for Replication on z/OS

For IBM Tivoli Storage Productivity Center for Replication Basic Edition for System z or IBM Tivoli Storage Productivity Center for Replication for System z that uses IBM WebSphere Application Server OEM Edition for z/OS, see the *IBM WebSphere Application Server OEM Edition for z/OS Configuration Guide* for information about how to stop IBM WebSphere Application Server OEM Edition for z/OS.

For IBM Tivoli Storage Productivity Center for Replication for System z that uses IBM WebSphere Application Server on z/OS, perform these steps to stop IBM Tivoli Storage Productivity Center for Replication:

1. Run the following command from the System Display and Search Facility (SDSF) panel to stop the IBM WebSphere Application Server daemon process:

```
/STOP BBOS001
```

Attention: For WebSphere Application Server OEM Edition for z/OS, use the following command:

```
/STOP WASOM1
```

2. Run the following command from the System Display and Search Facility (SDSF) panel to stop the daemon process:

```
/STOP BBODMNB
```

Attention: For WebSphere Application Server OEM Edition for z/OS, use the following command:

```
/STOP WASOM1D
```

3. Run the **/D A,L** command to display the running processes and verify that the stop command completed successfully. If you see the BBOS001 or BBODMNB processes, run the **/CANCEL** command.

Starting and stopping DB2

Use the DB2 Command Line Processor to start and stop DB2.

Note: If you are using the zero administration embedded repository (Apache Derby), you do not need to start or stop the repository.

To start DB2, run the **db2start** command from the DB2 Command Line Processor.

To stop DB2, run the **db2stop** command from the DB2 Command Line Processor.

Important: All connections to DB2 must be removed, or the **db2stop** command will fail. It is recommended that you stop the embedded IBM WebSphere Application Server, IBM Tivoli Storage Productivity Center for Replication, and all other applications accessing DB2 to ensure there are no connections to DB2 before attempting to stop DB2.

Verifying that components are running

This information describes how to verify that the IBM Tivoli Storage Productivity Center for Replication components are running.

Verifying that IBM WebSphere Application Server is running

This information describes how to verify that IBM WebSphere Application Server is running in your environment.

Use the System Display and Search Facility (SDSF) panel to run the **/D A,L** command. If the BBOS001 and BBODMNB processes are running, then the IBM WebSphere Application Server is running. If the WASOM1S, WASOM1 and WASOM1D processes are running, then the IBM WebSphere Application Server OEM Edition for z/OS is running

Verifying that the IBM Tivoli Storage Productivity Center for Replication server is running

This information describes how to verify whether the IBM Tivoli Storage Productivity Center for Replication server is running.

Perform one or more of these tasks to determine whether the IBM Tivoli Storage Productivity Center for Replication server is running:

- Start the IBM Tivoli Storage Productivity Center for Replication GUI or start up the IBM Tivoli Storage Productivity Center for Replication command line interface shell. If either of these methods is successful, then the IBM Tivoli Storage Productivity Center for Replication server is running.
- Determine whether IBM WebSphere Application Server is running. If IBM WebSphere Application Server is not running, then the IBM Tivoli Storage Productivity Center for Replication server is also not running.
- If IBM WebSphere Application Server is running, perform these steps to determine if IBM Tivoli Storage Productivity Center for Replication is configured correctly in IBM WebSphere Application Server:
 1. From a command prompt, change to the *install_root*/AppServer/profiles/default/bin directory.
 2. Enter the **wsadmin** command.
 3. Enter the **\$AdminApp list** command.
If you see CSM and CSMGUI in the resulting list, then IBM Tivoli Storage Productivity Center for Replication is configured correctly.
- View the *esmTrace.log* file in the *install_root*/AppServer/profiles/default/logs directory. If the *esmTrace.log* is being updated regularly and increasing in size, then IBM Tivoli Storage Productivity Center for Replication is running.

Verifying that DB2 is running

This information describes how to verify that DB2 is running in your environment.

If IBM Tivoli Storage Productivity Center for Replication is configured to use DB2 and DB2 is not running, IBM Tivoli Storage Productivity Center for Replication cannot start correctly. If this is the case, you will find log entries similar to the following in the *esmTrace.log* files:

```
[2007-02-23 16:30:30.369-05:00] server.startup : 0 RepMgr D
com.ibm.csm.server.CSM databaseReady TRACE: The database is not ready for
connections. Will retry in 10 seconds
```

```
[2007-02-23 16:30:40.403-05:00] server.startup : 0 RepMgr D DBService
isDatabaseReady TRACE: [IBM][CLI Driver] SQL1032N No start database manager
command was issued. SQLSTATE=57019 DSRA0010E: SQL State = 57019,
Error Code = -1,032
```

If DB2 is running and IBM Tivoli Storage Productivity Center for Replication can successfully connect to it, the following log entry can be found in the *esmTrace.log* file:

```
[2007-02-23 16:30:55.965-05:00] server.startup : 0 RepMgr D
com.ibm.csm.server.CSM databaseReady TRACE: The database is ready to go.
```

Starting the IBM Tivoli Storage Productivity Center for Replication GUI

After IBM Tivoli Storage Productivity Center for Replication has been installed, you can access the IBM Tivoli Storage Productivity Center for Replication GUI.

Go to the following URL to start the IBM Tivoli Storage Productivity Center for Replication GUI. Note that the Web address is case sensitive.

`https://host_name:port/CSM`

Note:

- If you upgraded from IBM Tivoli Storage Productivity Center for Replication version 3.3 or earlier, the default HTTP port is 9080 and the default HTTPS port is 9443.
- For IBM WebSphere Application Server for z/OS, the default port is 9443. For IBM WebSphere Application Server OEM Edition for z/OS, the default port is or 32209. Upgrading IBM Tivoli Storage Productivity Center for Replication does not change your original port settings.
- You can verify the ports that are correct for your installation in the `install_root/AppServer/profiles/profile_name/properties/portdef.props` file. The ports is defined by the `WC_defaulthost` (HTTP port) and `WC_defaulthost_secure` (HTTPS port) properties within the file.

To log in to the command line interface, use the CLI user ID and password that you entered during the installation of IBM Tivoli Storage Productivity Center for Replication. This ID and password are the same as the Administrator ID and password that you entered when you installed the product.

Identifying the version of IBM Tivoli Storage Productivity Center for Replication

This topic tells how to view the IBM Tivoli Storage Productivity Center for Replication version.

The version of code on the IBM Tivoli Storage Productivity Center for Replication server determines the available features and enhancements added to the product. By clicking **About**, located in the Main menu, you can view the version and release of the IBM Tivoli Storage Productivity Center for Replication. The `ver` command described in the *IBM Tivoli Storage Productivity Center for Replication for System z Command-line Interface User's Guide*, also displays the current version of the product.

Backing up and restoring IBM Tivoli Storage Productivity Center for Replication configuration data

You can back up the entire IBM Tivoli Storage Productivity Center for Replication database. You can also back up copy sets in a specific session. You can then use the backup files to restore a previous configuration or recover from a disaster.

Back up and recovery

You can backup and recover copy set data for a specific session and the complete Tivoli Storage Productivity Center for Replication database.

Copy sets

You can export data about all copy sets in a specific session to maintain a backup copy that you can use to recover if you lose the session or to upgrade to a different management server.

When exporting copy sets, Tivoli Storage Productivity Center for Replication takes a snapshot of the session at that point in time and saves the data in a comma separated value (CSV) file, which you can view or edit in a spreadsheet program such as Microsoft Excel. The exported CSV file includes the session name, session type, date that the data was exported, and the copy sets for each role pair. There is

one line per copy set, and the volumes in the copy set are separated by a comma (for example: ESS:2105.FCA57:VOL:17C7,ESS:2105.12043:VOL:17C7).

The following example illustrates the content of the CSV file for a FlashCopy session. Note that the first valid row must contain the appropriate role names for the session. The order of the copy sets does not matter, and you can include extra roles. A copy set is created from each row that follows the role names. All rows must have data in each column to be a valid row. Note that the number sign (#) indicates that the line is a comment. Lines that are comments are ignored.

```
#Session1,  
#FlashCopy,  
#Oct 2, 2009 10:03:18 AM  
  
H1,T1  
DS8000:2107.FRLL1:VOL:1004,DS8000:2107.FRLL1:VOL:1104  
DS8000:2107.FRLL1:VOL:1011,DS8000:2107.FRLL1:VOL:1101  
DS8000:2107.FRLL1:VOL:1005,DS8000:2107.FRLL1:VOL:1105
```

Important: You must manually save this file on the local system when you export copy sets from the Tivoli Storage Productivity Center for Replication Web interface.

IBM Tivoli Storage Productivity Center for Replication database

Tivoli Storage Productivity Center for Replication database contains all product data, including data about storage systems, sessions, copy sets, paths, user administration and management servers. You can back up this data and use the backup file to recover from a disaster or restore a previous configuration.

Important: You must have Administrator privileges to back up and recover the database.

The current data is stored in a new file each time you create a backup. The backup file is named *yyyyMMdd_HHmmsSSS.zip*, where *yyyy* is the year, *MM* is the month, *dd* is the day, *HH* is the hour, *mm* is the minute, *ss* is the seconds, *SSS* is the milliseconds when the backup command was run. It is your responsibility to delete backup versions that are no longer needed.

By default, the backup file is stored in the *install_root/AppServer/profiles/default/database/backup* directory. You can change the default location by editing the **db.backup.location** property in *rmserver.properties* file, which is located in the *websphere_home/AppServer/profiles/websphere_profile/properties* directory.

The backup file contains the Tivoli Storage Productivity Center for Replication database data at the time the backup was performed. Any changes that were made after the backup are not reflected when the backup files are used to restore an Tivoli Storage Productivity Center for Replication database. It is recommended that you create a new backup file:

- After changing the Tivoli Storage Productivity Center for Replication database data, such as adding or deleting a storage system, changing properties, and changing user privileges
- After an Tivoli Storage Productivity Center for Replication session changes direction. For example, if a Metro Mirror session was copying data from H1 to H2 when the backup was taken, and later, the session was started in the H2 to H1 direction. The session must be in the Prepared state before you create the backup.

- After a site switch has been declared and the Enable Copy To Site command has been issued. After you create a backup, consider deleting the previous backup to prevent Tivoli Storage Productivity Center for Replication from starting the copy in the wrong direction.

When you create a backup, ensure that all Tivoli Storage Productivity Center for Replication sessions are either in the Defined, Prepared or Target Available state.

Restoring the Tivoli Storage Productivity Center for Replication database from a backup copy puts Tivoli Storage Productivity Center for Replication back to the point in time when the backup was made. Relationships that exist on the storage systems that were created by Tivoli Storage Productivity Center for Replication after the backup was made will no longer be managed by Tivoli Storage Productivity Center for Replication until you add the copy set to the session and Tivoli Storage Productivity Center for Replication assimilates the relationship into the session. Copy sets that were deleted after the backup will be restored and a subsequent Start command to the session will create new relationships; therefore, you must remove the deprecated copy sets before issuing the Start command.

After restoring a Global Mirror session, you must stop the Global Mirror master and subordinates before restarting the Global Mirror session. Refer to your DS6000 and DS8000 storage system documentation for more information.

Backing up the Tivoli Storage Productivity Center for Replication database

This topic describes how to create a backup of the IBM Tivoli Storage Productivity Center for Replication database, including data about storage systems, sessions, copy sets, user administration and management server configuration.

To back up the Tivoli Storage Productivity Center for Replication database, run the **mkbackup** command from the command line, for example:

```
csmdi> mkbackup
```

Prerequisites:

- You must have Administrator privileges to run this command.
- This procedure applies to only the zero-administration embedded repository. This procedure *is not* applicable when DB2 is being used as the persistent datastore for the IBM Tivoli Storage Productivity Center for Replication database. For information about restoring your DB2 environment, refer to your DB2 documentation.

By default, the backup file is stored in the *install_root/AppServer/profiles/default/database/backup* directory. You can change the default location by editing the **db.backup.location** property in *rmserver.properties* file, which is located in the *websphere_home/AppServer/profiles/websphere_profile/properties* directory.

Restoring the IBM Tivoli Storage Productivity Center for Replication database

You can restore IBM Tivoli Storage Productivity Center for Replication database that was previously backed up to the local system.

Important:

- Restoring the database does not require Administrator privileges. However, you must be able to access the files on the IBM Tivoli Storage Productivity Center for Replication server that are listed in the procedure.
- This procedure applies to only the zero-administration embedded repository. This procedure *is not* applicable when DB2 is being used as the persistent datastore for the IBM Tivoli Storage Productivity Center for Replication database. For information about restoring your DB2 environment, refer to your DB2 documentation.

Perform these steps to restore the IBM Tivoli Storage Productivity Center for Replication database from a backed up version:

1. Stop IBM Tivoli Storage Productivity Center for Replication on the active management server by running the **stopServer** command from a command line.
2. Delete the *install_root/AppServer/profiles/default/database/csmdb* directory and all contents in it.
3. Uncompress the backup file into the *install_root/AppServer/profiles/default/database* directory.
4. If IBM Tivoli Storage Productivity Center for Replication is running on z/OS, change the permissions of the csmdb directory by running the following commands:
 - `chgrp -R $WAS_GROUP csmdb`
 - `chmod -R u+rwx csmdb`
 - `chmod -R g+rwx csmdb`
 - `chmod -R o+r csmdb`
5. Restart IBM Tivoli Storage Productivity Center for Replication on the active management server by running the **startServer** command from a command line.
6. Resolve any changes that might have occurred since the backup was created.
7. Start the IBM Tivoli Storage Productivity Center for Replication sessions using the appropriate start commands. The start commands reestablish the relationship between the volume pairs and synchronize data on those volumes.
8. If you have a standby management server, reestablish that standby relationship to update the database on the standby server.

Exporting copy set data

You can export data about all copy sets in a specific session, to maintain a backup copy that can be used to recover if you lose your session or upgrade to a different server.

Perform these steps to export the copy sets in a specific session:

1. In the navigation tree, select **Sessions**. The Sessions panel is displayed
2. Select the session for which you want to export copy sets.
3. Select **Export Copy Sets** from the **Actions** list, and click **Go**. The Export Copy Set wizard displays the status of the export and a link to the exported file.
4. Click that link and save the file to the local system.

Important: You must save the file to your local system. After you close the panel, the data will be lost.

5. Click **Finish**.

Importing copy set data

You can import copy set data that was previously exported to a comma separated value (CSV) file.

Perform the following steps to import copy sets into an existing session:

1. In the navigation tree, select **Sessions**. The Session panel is displayed.
2. Select the session for which you want to import copy sets.
3. Select **Add Copy Sets** from the **Actions** list, and click **Go**. The Add Copy Sets wizard is displayed.
4. Select **Use a CSV file to import copy sets**.
5. Type the location and name of the CSV file to import, or use **Browse** to select the file. Then, click **Next**.
6. Verify that the matching results were successful, and then click **Next**.
7. Select the copy sets you want to add, and then click **Next**.
8. Confirm the number of copy sets that you want to create, and click **Next**. A progress bar displays.
9. Click **Next**.
10. Verify the matches, and click **Finish**.

Chapter 3. Managing management servers

This section provides information about how to set up active and standby management servers, restore a lost connection between the management servers, or perform a takeover on the standby management server.

Management servers

The *management server* is a system that has IBM Tivoli Storage Productivity Center for Replication 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 Tivoli Storage Productivity Center for Replication that runs on a different physical system, but is continuously synchronized with the primary (or active) Tivoli Storage Productivity Center for Replication 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. If the active management server fails, you can issue the Takeover command to make the standby management server take over.

Connecting the active management server to the standby management server

Ensure that the active management server is connected to the standby management server. This connection creates the *management server relationship* that begins the synchronization process. Each management server can be in only one management server relationship.

A management server relationship might become disconnected for a number of reasons, including a connectivity problem or a problem with the alternate server. Issue the Reconnect command to restore synchronization.

Performing a takeover on the standby management server

If you must perform a takeover and use the standby server, ensure that you shut down the active management server first. You must ensure that you do not have two active management servers. If there are two active management servers and a condition occurs on the storage systems, both management servers respond to the same conditions, which might lead to unexpected behavior.

If you perform an action on the active management server when the servers are disconnected, the servers will be out of synch.

Viewing the status of the management servers

You can view the status of the active and standby management servers from the Management Servers panel in the Tivoli Storage Productivity Center for Replication graphical user interface (GUI). If you are logged on to the active management server, the icons on this panel show the status of the standby

management server. If you are logged on to the standby management server, the icons on this panel show the status of the active management server.

When the status is Synchronized, the standby management server contains the same data that the active management server contains. Any update to the active management server database is replicated to the standby server database.

Managing volumes on storage systems

When you add direct connections, Hardware Management Console (HMC) connections, or z/OS connections on the active management server, Tivoli Storage Productivity Center for Replication automatically enables the management of attached extended count key data (ECKD) volumes, non-attached count key data (CKD) volumes, and all fixed-block volumes on the storage system. To disable management of volumes on the storage system, use the volume protection function.

Information specific to management servers in z/OS environments

If the standby management server is not in the active server z/OS sysplex, the standby server is not be able to communicate with the storage systems using a z/OS connection; therefore, another connection must be made using a TCP/IP connection.

If DB2 is configured for data sharing mode across the z/OS sysplex, one of the Tivoli Storage Productivity Center for Replication servers must be configured to use the zero-administration embedded repository. If the embedded repository is not used, the two servers will overwrite the same data in the Tivoli Storage Productivity Center for Replication database.

Ports

Tivoli Storage Productivity Center for Replication use ports for communication with the management servers in a high-availability relationship, graphical user interface (GUI), command-line interface (CLI), and storage systems.

Web browser ports

To launch the Tivoli Storage Productivity Center for Replication GUI, use one of these default ports:

	HTTP port	HTTPS port
WebSphere Application Server	9080	9443
IBM System Services Runtime Environment for z/OS or WebSphere Application Server OEM Edition for z/OS	32208	32209

You can verify the ports that are correct for your installation in the *install_root/AppServer/profiles/profile_name/properties/portdef.props* file. The ports is defined by the `WC_defaulthost` (HTTP port) and `WC_defaulthost_secure` (HTTPS port) properties within the file.

Standby management server port

Tivoli Storage Productivity Center for Replication uses the default port 5120 for communication between the active and standby management server. This port number is initially set at installation time.

Important: The standby management server port number must be the same on both management servers in a high-availability relationship. If you change the port number on one management server, you must also change it on the other.

You can view the current port for each management server by clicking **Management Servers** in the navigation tree or from the **Health Overview** panel in the GUI or using the **lshaservers** command from the command line interface.

Client port

IBM Tivoli Storage Productivity Center for Replication client uses the default port 5110 to communicate with the graphical user interface and command line interface from a remote system. This port number is initially set at installation time.

Important: The client port number must be the same on both management servers in a high-availability relationship. If you change the port number on one management server, you must also change it on the other.

You can view the client port number on the local management server by clicking **About** in the navigation tree in the GUI or using the **whoami** command from the command line interface.

Storage system ports

The following table lists the default ports for each storage type.

Table 14. Storage system default ports

Storage System	Connection Type	Port
<ul style="list-style-type: none">• TotalStorage Enterprise Storage Server Model 800• System Storage DS8000• System Storage DS6000	Direct Connection	2433
<ul style="list-style-type: none">• System Storage DS8000	Hardware Management Console Connection	1750
<ul style="list-style-type: none">• System Storage SAN Volume Controller• Storwize V7000• Storwize V7000 Unified	Direct Connection	443 and 22
<ul style="list-style-type: none">• The XIV system	Direct Connection	7778

Ensure that your network configuration is set up so that Tivoli Storage Productivity Center can send outgoing TCP/IP packets to the storage controllers. It is possible when adding the storage controllers to Tivoli Storage Productivity Center to set a specific port number for your storage controller.

Because there are typically multiple applications running on the management server, it is possible that port conflicts might arise if other applications attempt to

use the same ports that IBM Tivoli Storage Productivity Center for Replication is using. Use the **netstat** command to verify which ports the various applications on the management server are using.

When you add a storage system to the Tivoli Storage Productivity Center for Replication configuration, the port field is automatically populated with the appropriate value. If you want to use different ports, you can change them by clicking **Storage Systems** located in the navigation tree, clicking the storage system that you want to change, and then changing the port value in the View/Modify Details panel.

Note: The storage system must not be in a Connected state if you want to change port values.

If firewalls are being used in your configuration, ensure that none of these ports are being blocked. Also ensure that not only is the Tivoli Storage Productivity Center for Replication server granted access to reach the other components, but that the other components are granted access to reach the Tivoli Storage Productivity Center for Replication server.

SNMP alerts

This topic describes the SNMP alerts that are sent by IBM Tivoli Storage Productivity Center for Replication and the associated object IDs (OIDs).

SNMP alerts are sent during the following general events:

- Session state change
- Configuration change
- Suspending-event notification
- Communication failure
- Management Server state change

Session state change SNMP trap descriptions

This topic lists the SNMP traps that are sent during a session state change. A different trap is sent for each state change. These alerts are sent by only the active management server.

A session state change SNMP trap is sent each time the session changes to one of the following states:

- Defined
- Preparing
- Prepared
- Suspended
- Recovering
- Flashing
- TargetAvailable
- Suspending
- (Metro Global Mirror only) SuspendedH2H3
- (Metro Global Mirror only) SuspendedH1H3

Table 15. Session state change traps

Object ID (OID)	Description
1.3.6.1.4.1.2.6.208.0.1	The state of session X has transitioned to Defined.
1.3.6.1.4.1.2.6.208.0.2	The state of session X has transitioned to Preparing.

Table 15. Session state change traps (continued)

Object ID (OID)	Description
1.3.6.1.4.1.2.6.208.0.3	The state of session X has transitioned to Prepared.
1.3.6.1.4.1.2.6.208.0.4	The state of session X has transitioned to Suspended.
1.3.6.1.4.1.2.6.208.0.5	The state of session X has transitioned to Recovering.
1.3.6.1.4.1.2.6.208.0.6	The state of session X has transitioned to Target Available.
1.3.6.1.4.1.2.6.208.0.19	The state of session X has transitioned to Suspending.
1.3.6.1.4.1.2.6.208.0.20	The state of session X has transitioned to SuspendedH2H3.
1.3.6.1.4.1.2.6.208.0.21	The state of session X has transitioned to SuspendedH1H3.
1.3.6.1.4.1.2.6.208.0.22	The state of session X has transitioned to Flashing.
1.3.6.1.4.1.2.6.208.0.23	The state of session X has transitioned to Terminating.

Configuration change SNMP trap descriptions

This topic lists the SNMP traps that are sent when the configuration changes. These alerts are sent by only the active management server.

Configuration change SNMP traps are sent after the following configurations changes are made:

- One or more copy sets have been added or deleted from a session
An alert is sent for each set of copy sets added to or removed from a session. Note that an alert for copy set changes is sent only once within 15 minutes of a configuration change, so you might not see alerts from successive changes that occur within that 15-minute period. For example, if you make a copy set configuration change that causes an alert to be sent at 10:41:01, and if you were to make additional copy set changes at 10:42:04 and 10:50:09, no alerts would be sent for these two changes because they occurred within the 15-minute minimum interval from the first alert.
- PPRC path definitions have been changed
An alert is sent for each path configuration change made.

Table 16. Configuration change traps

Object ID (OID)	Description
1.3.6.1.4.1.2.6.208.0.7	One or more copy sets have been added or deleted from this session. Note: An event is sent for each session at least every 15 minutes.
1.3.6.1.4.1.2.6.208.0.8	Peer-to-Peer Remote Copy (PPRC) path definitions have been changed. An event is sent for each path configuration change.

Suspending-event notification SNMP trap descriptions

This topic lists the SNMP traps that are sent during a suspending-event notification. These alerts are sent by the active and standby management server.

Configuration change SNMP indicates that a session has transitioned to a Severe status due to an unexpected error. This event is raised only for an error with peer-to-peer remote copy (PPRC) processing.

Table 17. Suspending-event notification traps

Object ID (OID)	Description
1.3.6.1.4.1.2.6.208.0.9	The session is in a Severe state due to an unexpected error.

Communication-failure SNMP trap descriptions

This topic lists the SNMP traps that are sent during a communication-failure. These alerts are sent by both the active and standby management servers.

Communication-failure SNMP traps are sent after the following events occur:

- A server times out attempting to communicate with a storage system.
- A server encounters errors attempting to communicate with a storage system.
- An active server terminates communication with a standby server as a result of communication errors.
- A standby encounters communication errors with an active server.

After an SNMP trap for a given failure is sent, it is not resent unless communication has been reestablished and failed again.

Table 18. Communication-failure traps

Object ID (OID)	Description
1.3.6.1.4.1.2.6.208.0.10	Server X has timed out attempting to communicate with storage system Y.
1.3.6.1.4.1.2.6.208.0.11	Server X has encountered errors attempting to communicate with storage system Y.
1.3.6.1.4.1.2.6.208.0.12	Active server X has terminated communication with standby server Y as a result of communication errors.
1.3.6.1.4.1.2.6.208.0.13	Standby server X has encountered communication errors with active server Y.

Management Servers state-change SNMP trap descriptions

This topic lists the SNMP traps that are sent when the state of the management server changes. These alerts are sent by both the active and standby management servers.

A management server state change SNMP trap is sent each time the management server changes to one of the following states:

- Unknown
- Synchronization Pending
- Synchronized
- Disconnected Consistent
- Disconnected

Table 19. Management Servers state-change traps

Object ID (OID)	Description
1.3.6.1.4.1.2.6.208.0.14	The IBM Tivoli Storage Productivity Center for Replication Server Management Server connection X->Y has changed state to Unknown (previously Offline).

Table 19. Management Servers state-change traps (continued)

Object ID (OID)	Description
1.3.6.1.4.1.2.6.208.0.15	The IBM Tivoli Storage Productivity Center for Replication Server Management Server connection X->Y has changed state to Synchronized.
1.3.6.1.4.1.2.6.208.0.16	The IBM Tivoli Storage Productivity Center for Replication Server Management Server connection X->Y has changed state to Disconnected Consistent (previously Consistent Offline).
1.3.6.1.4.1.2.6.208.0.17	The IBM Tivoli Storage Productivity Center for Replication Server Management Server connection X->Y has changed state to Synchronization Pending.
1.3.6.1.4.1.2.6.208.0.18	The IBM Tivoli Storage Productivity Center for Replication Server Management Server connection X->Y has changed state to Disconnected.

Setting up a standby management server

You can set up a standby management server in two ways: setting up the management server you are logged in to as the standby management server, or designating another server as the standby management server.

Note: When you define a standby management server, the IBM Tivoli Storage Productivity Center for Replication code must be at the same level on both the standby and active management servers.

Setting the local management server as the standby server

This topic describes how to set the management server on which you are currently logged in as the standby management server.

Attention: When you set the local management server, the server you are logged in at will be wiped of all session information, and replaced with the session information belonging to the server you specified.

Important: The standby management server port number must be the same on both management servers in a high-availability relationship. If you change the port number on one management server, you must also change it on the other.

Perform these steps to set the local management server as the standby server:

1. In the navigation tree, select **Management Servers**.
2. From the **Action** menu, select **Set This Server as Standby**, and click **Go**. This opens the Set This Server As Standby panel.
3. Enter the domain name or IP address of the desired active management server.
4. Click **OK** to connect to the active server. You have now designated the server you are logged in at as the standby server.

Setting a remote management server as the standby server

This topic describes how to set up a management server on which you are not logged in as the standby server.

Attention: When you set a remote management server as the standby server, the remote management server is wiped of all session information and replaced with the session information belonging to the management server you specified.

Important: The standby management server port number must be the same on both management servers in a high-availability relationship. If you change the port number on one management server, you must also change it on the other.

Perform these steps to set a remote management server as the standby server:

1. In the navigation tree, select **Management Servers**.
2. Select **Define Standby** from the drop-down menu, and click **Go**. The Define Standby panel appears.
3. Type the domain name or IP address of the server that you are defining as the new standby management server. Log in to the standby management server by entering the username and password.
4. Click **OK** to connect to the standby management server.

Reinstalling the primary server during an active session

This topic provides the steps to reinstall the primary IBM Tivoli Storage Productivity Center for Replication server if it develops a problem during an active session.

If there is an active session running when you need to reinstall the primary server, follow the steps listed here to reinstall the server without affecting the active session.

1. If the heartbeat is enabled, disable it:
 - a. In the navigation tree, select **Advanced Tools**.
 - b. Click **Disable Heartbeat**.
2. On the standby server, Server 2, issue a take-over. This makes Server 2 the active server. It is possible that the original active server, Server 1, is still listed on the Server 2 Management Servers page. If so, select **Remove Standby**.
3. Disable the heartbeat on Server 2, in case there are any problems.
4. Uninstall IBM Tivoli Storage Productivity Center for Replication on Server 1.
5. Reinstall IBM Tivoli Storage Productivity Center for Replication on Server 1.

Note: If no changes have been made to the configuration while Server 1 was being reinstalled, steps 6 and 7 are not necessary.

6. When IBM Tivoli Storage Productivity Center for Replication is running on Server 1, log into Server 2 and set Server 1 as the standby server for Server 2. This step copies the configuration from Server 2 to Server 1. This process takes a few minutes.
7. When the management servers status is synchronized, issue a take-over on Server 1.
This makes Server 1 an active server, able to control sessions.

Note: It is possible that Server 2 is still listed on the Server 1 Management Servers page. If this is the case, select **Remove Standby**.

8. Disable the heartbeat on Server 1 to make sure this active server does not have any problems.

Note: If you do not need to reinstall IBM Tivoli Storage Productivity Center for Replication on Server 2, skip steps 9 and 10.

9. Uninstall IBM Tivoli Storage Productivity Center for Replication from Server 2.

10. Reinstall IBM Tivoli Storage Productivity Center for Replication on Server 2.
11. On Server 2, go to the Management Servers page and select the **Set This Server As Standby** option, entering the information for Server 1. When this step is complete, Server 1 is the active server, and Server 2 is the standby server.
12. When you are confident that the active server is running without any problems, enable the heartbeat again, if desired.

Reconnecting the active and standby management servers

If the active and standby management servers become disconnected, reestablish that connection.

Perform these steps to cause the standby management server to become the active management server:

1. In the navigation tree, select **Management Servers**. The Management Servers panel is displayed.
2. Select **Reconnect** from the **Actions** list, and click **Go**.

Performing a takeover on the standby management server

If the active management server fails, you can force the standby management server to take over monitoring and managing replication responsibilities.

Important: If the current active management server is still active, you must not attempt to control the replication environment simultaneously from both management servers. Instead, either reconfigure the current active management server to be a standby management server, or shut it down.

Perform these steps to cause the standby management server to become the active management server:

1. If the active management server is functioning, take it offline so you do not have two active management servers managing the same sessions.
2. Log in to the IBM Tivoli Storage Productivity Center for Replication Web interface running on the standby management server.
3. In the navigation tree, select **Management Servers**. The Management Servers panel is displayed.
4. Select **Takeover** from the **Actions** list, and click **Go**.
5. To reestablish high-availability, perform one of these steps:
 - Choose another server to be the standby management server. See instructions for setting up a standby management server.
 - Bring the failed management server back online, and then make that server the standby management server. See instructions for setting up a standby management server.
 - Bring the failed management server back online, and then make that server the active management server to return to the original configuration. Repeat the steps in this section and then add the original standby server as the standby server.

Important: Do *not* use the **Reconnect** command if you perform a takeover. You would use the **Reconnect** command when the active server loses its connection with the standby server; it reconnects the two servers. Do *not* use the **Reconnect** command after a takeover to reconnect to the original active server.

Configuring SNMP

The SNMP community name has a default value of `public`.

To change the community name, modify or add the `csm.server.snmp_community_string` property in the `rmservers.properties` file, which is located in the `websphere_home/AppServer/profiles/websphere_profile/properties` directory.

Adding SNMP managers

Use the `mksnmp` command to add an SNMP manager to the list of servers to which IBM Tivoli Storage Productivity Center for Replication sends SNMP alerts.

IBM Tivoli Storage Productivity Center for Replication uses management information base (MIB) files to provide a textual description of each SNMP alert sent by IBM Tivoli Storage Productivity Center for Replication. You must configure the SNMP manager to use both the `SYSAPPL-MIB.mib` and `ibm-TPC-Replication.mib` files. These MIB files are located in the `install_root\Scripts` directory. Follow the directions provided by your SNMP manager application to configure it to use the MIB files.

Tip: You can also find the MIB files on the installation CD in the `TPCRM/CSM-Client/etc` directory.

Note: IBM Tivoli Storage Productivity Center for Replication sends all SNMP alerts to each registered SNMP manager. SNMP alerts are not specific to any particular session, and all alerts for any session are sent. You cannot choose to send a subset of SNMP alerts.

Changing the standby management server port number

The standby management server port is used for communication between the active and standby management server. This port is initially defined during the installation. You can manually change this port after installation.

Important: The standby management server port number must be the same on both management servers in a high-availability relationship. If you change the port number on one management server, you must also change it on the other.

1. Open the `rmservers.properties` file in the `websphere_home/AppServer/profiles/websphere_profile/properties` directory.
2. Modify the port number for the following property:
`communications.haPort=5120`
3. Restart IBM Tivoli Storage Productivity Center for Replication. You must restart IBM Tivoli Storage Productivity Center for Replication to activate property changes. Properties are not synchronized between the IBM Tivoli Storage Productivity Center for Replication management servers and must be done on each IBM Tivoli Storage Productivity Center for Replication management server.

Changing the client port number

The client port is used to log in to the graphical user interface and command line interface from a remote system. This port is initially defined during the installation. You can manually change this port after installation.

Important: The client port number must be the same on both management servers in a high-availability relationship. If you change the port number on one management server, you must also change it on the other.

1. Open the `rmservers.properties` file in the `websphere_home/AppServer/profiles/websphere_profile/properties` directory.
2. Modify the port number for the following property:
`communications.port=5110`
3. Open the `repcli.properties` file in the `install_root/CLI` directory.
4. Modify the port number for the following property:
`port=5110`
5. Restart IBM Tivoli Storage Productivity Center for Replication. You must restart IBM Tivoli Storage Productivity Center for Replication to activate property changes. Properties are not synchronized between the IBM Tivoli Storage Productivity Center for Replication management servers and must be maintained on each IBM Tivoli Storage Productivity Center for Replication management server.

Changing the time zone in z/OS

Timestamps automatically default to Greenwich Mean Time (GMT) on z/OS. This topic describes how to change to the Eastern Daylight Time (EDT).

To change the time zone to EDT instead of GMT, perform the following steps from the IBM WebSphere Application Server or IBM WebSphere Application Server OEM Edition for z/OS Administration Console:

1. Go to **Environment > WebSphere Variables > Cell Scope**.
2. Click **New**.
3. Type the name as `TZ`.
4. Type the value as `EST5EDT`.
5. Apply the changes.
6. Save the changes.
7. Restart IBM WebSphere Application Server or IBM WebSphere Application Server OEM Edition for z/OS.

Note: For instructions on changing to a time zone other than EDT, see the IBM WebSphere Application Server Express[®] information center on the web at publib.boulder.ibm.com/infocenter/wasinfo/v6r0/topic/com.ibm.websphere.express.doc/info/exp/ae/rrun_svr_timezones.html.

Chapter 4. Managing storage systems

To replicate data among storage systems using IBM Tivoli Storage Productivity Center for Replication, you must add connections to the storage systems. After a storage system is added, you can associate a location, modify connection properties, set volume protection, and refresh the storage configuration for that storage system.

Ports

Tivoli Storage Productivity Center for Replication use ports for communication with the management servers in a high-availability relationship, graphical user interface (GUI), command-line interface (CLI), and storage systems.

Web browser ports

To launch the Tivoli Storage Productivity Center for Replication GUI, use one of these default ports:

	HTTP port	HTTPS port
WebSphere Application Server	9080	9443
IBM System Services Runtime Environment for z/OS or WebSphere Application Server OEM Edition for z/OS	32208	32209

You can verify the ports that are correct for your installation in the *install_root/AppServer/profiles/profile_name/properties/portdef.props* file. The ports is defined by the *WC_defaulthost* (HTTP port) and *WC_defaulthost_secure* (HTTPS port) properties within the file.

Standby management server port

Tivoli Storage Productivity Center for Replication uses the default port 5120 for communication between the active and standby management server. This port number is initially set at installation time.

Important: The standby management server port number must be the same on both management servers in a high-availability relationship. If you change the port number on one management server, you must also change it on the other.

You can view the current port for each management server by clicking **Management Servers** in the navigation tree or from the **Health Overview** panel in the GUI or using the **lshaservers** command from the command line interface.

Client port

IBM Tivoli Storage Productivity Center for Replication client uses the default port 5110 to communicate with the graphical user interface and command line interface from a remote system. This port number is initially set at installation time.

Important: The client port number must be the same on both management servers in a high-availability relationship. If you change the port number on one management server, you must also change it on the other.

You can view the client port number on the local management server by clicking **About** in the navigation tree in the GUI or using the **whoami** command from the command line interface.

Storage system ports

The following table lists the default ports for each storage type.

Table 20. Storage system default ports

Storage System	Connection Type	Port
<ul style="list-style-type: none"> • TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 	Direct Connection	2433
<ul style="list-style-type: none"> • System Storage DS8000 	Hardware Management Console Connection	1750
<ul style="list-style-type: none"> • System Storage SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified 	Direct Connection	443 and 22
<ul style="list-style-type: none"> • The XIV system 	Direct Connection	7778

Ensure that your network configuration is set up so that Tivoli Storage Productivity Center can send outgoing TCP/IP packets to the storage controllers. It is possible when adding the storage controllers to Tivoli Storage Productivity Center to set a specific port number for your storage controller.

Because there are typically multiple applications running on the management server, it is possible that port conflicts might arise if other applications attempt to use the same ports that IBM Tivoli Storage Productivity Center for Replication is using. Use the **netstat** command to verify which ports the various applications on the management server are using.

When you add a storage system to the Tivoli Storage Productivity Center for Replication configuration, the port field is automatically populated with the appropriate value. If you want to use different ports, you can change them by clicking **Storage Systems** located in the navigation tree, clicking the storage system that you want to change, and then changing the port value in the View/Modify Details panel.

Note: The storage system must not be in a Connected state if you want to change port values.

If firewalls are being used in your configuration, ensure that none of these ports are being blocked. Also ensure that not only is the Tivoli Storage Productivity Center for Replication server granted access to reach the other components, but that the other components are granted access to reach the Tivoli Storage Productivity Center for Replication server.

Storage systems

A *storage system* is a hardware device that contains data storage. Tivoli Storage Productivity Center for Replication can control data replication within and between various storage systems.

To replicate data among storage systems using Tivoli Storage Productivity Center for Replication, you must manually add a connection to each storage system in the Tivoli Storage Productivity Center for Replication configuration. This allows you to omit storage systems for which Tivoli Storage Productivity Center for Replication is not to manage replication and omit storage systems that are being managed by another Tivoli Storage Productivity Center for Replication management server.

For redundancy, you can connect a single storage system using a combination of direct, Hardware Management Console (HMC), and z/OS connections.

You can use the following storage systems:

- IBM TotalStorage Enterprise Storage Server (ESS) Model 800
- IBM System Storage DS6000
- IBM System Storage DS8000
- IBM System Storage SAN Volume Controller
- IBM Storwize V7000
- IBM Storwize V7000 Unified
- IBM XIV Storage System

A SAN Volume Controller can virtualize a variety of storage systems. Although Tivoli Storage Productivity Center for Replication does not support all storage systems, you can manage these storage systems through a single SAN Volume Controller cluster interface. Tivoli Storage Productivity Center for Replication connects directly to the SAN Volume Controller clusters.

You can define a location for each storage system and for each site in a session. When adding copy sets 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.

Notes:

- Tivoli Storage Productivity Center for Replication does not automatically discover the physical locations of storage systems. You can manually assign a location to a storage system from the GUI and CLI.
- Throughout this document, *ESS/DS* refers to the following models:
 - IBM TotalStorage Enterprise Storage Server Model 800
 - IBM System Storage DS8000
 - IBM System Storage DS6000

Storage connections

You must create a connection from the IBM Tivoli Storage Productivity Center for Replication management server to each storage system. You can connect either directly or through a Hardware Management Console (HMC) or IBM z/OS connection.

A single storage system can be connected using multiple connections for redundancy. For example, you can connect a IBM System Storage DS8000 storage system using an HMC connection and a z/OS connection. Tivoli Storage Productivity Center for Replication monitors how a storage system has been added to the configuration.

When you add a storage connection to the Tivoli Storage Productivity Center for Replication configuration, the storage system and the connection are added to the active management server configuration. For direct and HMC connections, the storage system and connection are also added to the standby management server configuration. For z/OS connections, only the storage system is added to the standby management server configuration. The connection is not added because the standby management server might not be running on z/OS and might not have access to the volumes on the storage system through a z/OS connection.

The storage systems are not required to be connected to the standby management server. However, if a storage system does not have a connection on the standby management server, you cannot manage copy services on the storage system from the standby server.

Important: If the Metro Mirror heartbeat is enabled, do not connect to a IBM TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, or System Storage DS6000 storage system using both an HMC connection and a direct connection. If you have both types of connections and the direct connection is lost, the session changes to the suspended state even though the HMC connection is still valid. If both connections are lost and the session is in the suspended state, restart the session when connectivity is regained to synchronize the session with the hardware.

When Tivoli Storage Productivity Center for Replication is running on z/OS and a storage system is added to the Tivoli Storage Productivity Center for Replication configuration through a TCP/IP (direct or HMC) connection, all ECKD volumes that are attached to the management server are managed through the TCP/IP connection. To use the Fibre Channel connection, you must explicitly add the storage system to the Tivoli Storage Productivity Center for Replication configuration through a z/OS connection.

If a storage system was previously added to the Tivoli Storage Productivity Center for Replication configuration through a z/OS connection and later the storage system is added through a TCP/IP connection, all non-attached ECKD volumes and fixed block volumes are added to the Tivoli Storage Productivity Center for Replication configuration.

When you remove a storage system, Tivoli Storage Productivity Center for Replication automatically removes all connections that the storage system is using with exception of the z/OS connection. You can also individually remove each connection through which the storage system is connected.

If Tivoli Storage Productivity Center for Replication has multiple connections to a specific storage system, the order in which you remove the connections produces different results:

- If you remove all direct and HMC connections first, the fixed block and non-attached ECKD volumes are removed from the Tivoli Storage Productivity Center for Replication configuration. The remaining ECKD volumes that are attached through the z/OS connection remain in the Tivoli Storage Productivity

Center for Replication configuration until the z/OS connection is removed. Removing the TCP/IP connection also disables the Metro Mirror heartbeat.

- If you remove the z/OS connection first and if there is an HMC or direct connection to volumes, those volumes are not removed from the Tivoli Storage Productivity Center for Replication configuration.
- HyperSwap can run provided that volumes are attached and available to z/OS storage, even if you are using a TCP/IP connection to storage.

Direct connection

The Tivoli Storage Productivity Center for Replication management server can connect directly to TotalStorage Enterprise Storage Server Model 800, DS6000, DS8000, SAN Volume Controller, Storwize V7000, Storwize V7000 Unified or the XIV system storage systems through a TCP/IP connection. The TCP/IP connection is required to discover the storage systems configuration (such as LSSs, volumes, volume size, and format), issue queries, and receive asynchronous events.

DS8000 storage systems on an IPV4 network can be connected directly to the management server. A direct connection requires an Ethernet card in the cluster. DS8000 storage systems on an IPV6 network cannot use a direct connection. They can be connected only through an HMC or z/OS connection.

When you add a direct connection to a DS or ESS cluster, specify the following information for cluster 0 and 1:

- IP addresses or domain names
- Ports
- User names
- Passwords

SAN Volume Controller or Storwize V7000 can virtualize various storage systems. Although Tivoli Storage Productivity Center for Replication does not support all storage systems, you can manage these storage systems through a single SAN Volume Controller or Storwize V7000 cluster interface. Tivoli Storage Productivity Center for Replication connects directly to the SAN Volume Controller or Storwize V7000 clusters. When you add a direct connection to a SAN Volume Controller or Storwize V7000 cluster to the Tivoli Storage Productivity Center for Replication configuration, specify the cluster IP address of the SAN Volume Controller or Storwize V7000 cluster, which in turn points to multiple SAN Volume Controller or Storwize V7000 storage systems. Ensure that the user name and password are correct for the cluster. If incorrect values are used, significant communication problems can occur, such as never advancing to the Prepared state.

Important: The SAN Volume Controller or Storwize V7000 user name must have privileges to maintain SSH keys. For information about troubleshooting Secure Shell connections to the SAN Volume Controller or Storwize V7000, see the Ethernet Connection Restrictions on SAN Volume Controller website at www-01.ibm.com/support/docview.wss?uid=ssg1S1002896.

Hardware Management Console connection

The IBM Tivoli Storage Productivity Center for Replication management server can connect to DS8000 storage systems through a Hardware Management Console (HMC). An HMC can have multiple DS8000 storage systems connected to it. When you add an HMC to the IBM Tivoli Storage Productivity Center for Replication configuration, all DS8000 storage systems that are behind the HMC are also added. You cannot add or remove individual storage systems that are behind an HMC.

You can also add a dual-HMC configuration, in which you have two HMCs for redundancy. This is recommended when the Metro Mirror heartbeat is required. You must configure both HMCs identically, including the user ID and password.

If planned maintenance is necessary on the HMC, it is recommended that you disable the Metro Mirror heartbeat on the management server while the maintenance is performed.

If the HMC needs to go down frequently or reboots frequently, it is recommended that you disable the Metro Mirror heartbeat. If the Metro Mirror heartbeat is required, the direct connection is recommended instead of an HMC connection.

Important: If a DS8000 storage system uses an HMC connection, the Metro Mirror heartbeat could trigger a freeze on the storage system and impact applications for the duration of the long busy timeout timer if the HMC is shut down for any reason, including upgrading microcode. The long busy timeout timer is the time after which the storage system will allow I/O to begin again after a freeze occurs if no run command has been issued by IBM Tivoli Storage Productivity Center for Replication. The default value is two minutes for ECKD volumes or one minute for fixed block volumes.

Notes:

- The user ID that you use to connect to the HMC must have admin, op_storage, or op_copy_services privileges on the DS8000 storage system.
- For minimum microcode requirements to connect to a DS8000 through a management console, see the Supported Storage Products List website at www-01.ibm.com/support/docview.wss?uid=swg21386446.

z/OS connections

An IBM Tivoli Storage Productivity Center for Replication management server that runs on z/OS can connect to IBM TotalStorage Enterprise Storage Server (ESS) Model 800, DS8000, and DS6000 storage systems through a z/OS connection. The z/OS connection is used to issue replication commands and queries for attached ECKD volumes over an existing Fibre Channel network and to receive asynchronous events. When a storage system is added to IBM Tivoli Storage Productivity Center for Replication through the z/OS connection, all ECKD volumes that are attached to the IBM Tivoli Storage Productivity Center for Replication management system are added to the IBM Tivoli Storage Productivity Center for Replication configuration. ECKD volumes that are not attached to the IBM Tivoli Storage Productivity Center for Replication z/OS management server are not added to the IBM Tivoli Storage Productivity Center for Replication configuration through the z/OS connection.

Notes:

- Ensure that all volumes in the logical storage subsystem (LSS) that you want to manage through a z/OS connection are attached to z/OS. Either the entire LSS must be attached to z/OS or none of the volumes in the LSS should be attached to z/OS for IBM Tivoli Storage Productivity Center for Replication to properly manage queries to the hardware.
- The z/OS connection is limited to storage systems that are connected to an IBM Tivoli Storage Productivity Center for Replication management server running z/OS.
- The Metro Mirror heartbeat is not supported through the z/OS connection. To use the Metro Mirror heartbeat, the storage systems must be added using a

direct connection or Hardware Management Console (HMC) connection. If the Metro Mirror heartbeat is enabled, a storage system is added through a direct connection and z/OS connection, and the direct connection becomes disconnected, then a suspend results as there is no heartbeat through the z/OS connection.

If at least one volume in a Logical Storage Subsystem (LSS) is attached through a z/OS connection, then all volumes in that LSS must be similarly attached. For example, if there are two ECKD volumes in an LSS, and one volume is attached to the IBM Tivoli Storage Productivity Center for Replication system using a z/OS connection and the other is attached through a direct connection, IBM Tivoli Storage Productivity Center for Replication would have knowledge of direct-connected volume. IBM Tivoli Storage Productivity Center for Replication issues commands to both volumes over the Fibre Channel network; however, commands issued to the direct-connection volume will fail, and IBM Tivoli Storage Productivity Center for Replication will show that the copy set that contains that volume has an error.

Use the following guidelines to add storage systems through a z/OS connection:

- Use the z/OS connection to manage ECKD volumes that are attached to an IBM Tivoli Storage Productivity Center for Replication management server running z/OS.
- To manage z/OS attached volumes through a z/OS connection (for example, for HyperSwap), you must explicitly add the z/OS connection for that storage system in addition to a TCP/IP connection (either the direct connection or the HMC connection).
- Create a z/OS connection before all TCP/IP connections if you want to continue to have IBM Tivoli Storage Productivity Center for Replication manage only the attached ECKD volumes.

Tip: It is recommended that you create both TCP/IP and z/OS connections for ECKD volumes to allow for greater storage accessibility.

Protected volumes

You can mark volumes as protected if you do not want those volumes used for replication.

When a volume is marked as protected, you cannot include that volume in a copy set. This protection applies only to IBM Tivoli Storage Productivity Center for Replication.

You might want to protect a volume in the following instances:

- The volume contains data that you never want to be copied to another volume. For example, the volume is secure, but if the data is copied to an unsecured volume, the data could be read. For this reason, the volume should not be the source for a relationship.
- The volume contains data that you do not want to be overwritten. For this reason, the volume should not be the target of a relationship.

Only administrators can change the volume protection settings.

Site awareness

You can associate a location with each storage system and each site in a session. This *site awareness* ensures that only the volumes whose location matches the location of the site are allowed for selection when you add copy sets to the session. This prevents a session relationship from being established in the wrong direction.

Note: To filter the locations for site awareness, you must first assign a site location to each storage system.

IBM Tivoli Storage Productivity Center for Replication does not perform automatic discovery of locations. Locations are user-defined and specified manually.

You can change the location associated with a storage system that has been added to the IBM Tivoli Storage Productivity Center for Replication configuration. You can choose an existing location or add a new one. Locations are deleted when there is no longer a storage system with an association to that location.

When adding a copy set to a session, a list of candidate storage systems is presented, organized by location. Storage systems that do not have a location are displayed and available for use when you create a copy set.

You can also change the location for any site in a session. Changing the location of a session does not affect the location of the storage systems that are in the session.

Changing the location of a storage system might have consequences. When a session has a volume role with a location that is linked to the location of the storage system, changing the location of the storage system could change the session's volume role location. For example, if there is one storage system with the location of A_Location and a session with the location of A_Location for its H1 role, changing the location of the storage system to a different location, such as B_Location, also changes the session's H1 location to Site 1. However, if there is a second storage system that has the location of A_Location, the session's role location is not changed.

Important: Location matching is enabled only when adding copy sets. If you change the location of a storage system or volume role, IBM Tivoli Storage Productivity Center for Replication does not audit existing copy sets to confirm or deny location mismatches.

Adding a storage connection

You can add one or more connections to a storage system. To replicate data among storage systems by using Tivoli Storage Productivity Center for Replication, you must add connections to the storage systems.

Prerequisites: You must have Administrator privileges to add a storage connection.

A single storage system can be connected by a combination of direct, management-console, and z/OS connections. Perform these steps to add a storage system connection:

1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is displayed in the Storage Systems view.

2. Click **Add Storage Connection**. The Add Storage System wizard Type page is displayed.
3. Select the connection type that you want to use to add the storage systems and click **Next**. The Connection page is displayed.
4. Perform one of these steps depending on the type of storage connection that you selected:

- **DS8000 / ESS 800 / DS6000 (Direct Connection)**: Enter the IP address or domain name, port number, user name, and password for both cluster 0 and cluster 1 in the appropriate fields.

In a System Storage DS8000 environment, if resource groups are defined on the System Storage DS8000 that you are connecting to, the user name you define here must have the appropriate access level on the System Storage DS8000 to manage copy services for the volumes that are used by Tivoli Storage Productivity Center for Replication.

- **DS8000 (HMC Connection)**: Enter the primary HMC IP address or domain, user name, password, and optionally the secondary HMC IP address or domain name in the appropriate fields.
- **SAN Volume Controller / Storwize V7000 / Storwize V7000 Unified (Direct Connection)**: Enter the cluster IP address or domain name, user name, and password in the appropriate fields.

The following are considerations for connection:

- a. Authentication with the storage system cluster is performed by public key exchange. When adding a storage system cluster to the Tivoli Storage Productivity Center for Replication configuration, specify the user ID and password of an administrator on the cluster that has sufficient privileges to maintain the Secure Shell (SSH) keys.
- b. When a valid user name and password are specified, Tivoli Storage Productivity Center for Replication audits the public keys installed on the storage system cluster to ensure that the Tivoli Storage Productivity Center for Replication key has the correct access levels. If the key is not installed on the cluster, Tivoli Storage Productivity Center for Replication attempts to install the key. If you install the key yourself, install the public key found in the *install_root/AppServer/profiles/default/etc/public* directory on the storage system cluster by using the SSH Key Maintenance panel on the storage system graphical user interface. The public key ID is installed under tpcr.
- c. You must use a storage system cluster superuser user name and password to correctly receive indications from the storage system cluster.
- d. For Storwize V7000, Storwize V7000 Unified, or SAN Volume Controller version 5.0 or later, you must create a user ID through the storage system administration console on the cluster to which you want to attach Tivoli Storage Productivity Center for Replication. This user ID must have local authentication type and have the Administrator role authority.

- **XIV (Direct Connection)**: Enter the IP address or domain name of the XIV system node. Enter a user name and password of a user with appropriate access rights to the XIV system.
- **z/OS (FICON Connection)**: Select one or more storage systems from the list of candidate systems that match the selected storage type, or click **All** to select all storage systems.

5. Click **Next**.
6. Click **Finish**.

Removing a storage connection

You can remove a single connection to a storage system from the IBM Tivoli Storage Productivity Center for Replication configuration.

Prerequisites: You must have Administrator privileges to remove a storage connection.

When removing a connection, the storage system might have other connections and therefore still be connected to IBM Tivoli Storage Productivity Center for Replication.

When removing a connection, all storage systems that rely on only that connection are removed. If a storage system is removed, the volumes for that storage system are removed from management server control. All copy sets with a volume on the removed storage systems are also removed from their respective sessions, making the target volume unrecoverable.

If Tivoli Storage Productivity Center for Replication has multiple connections to a specific storage system, the order in which you remove the connections produces different results:

- If you remove all direct and HMC connections first, the fixed block and non-attached ECKD volumes are removed from the Tivoli Storage Productivity Center for Replication configuration. The remaining ECKD volumes that are attached through the z/OS connection remain in the Tivoli Storage Productivity Center for Replication configuration until the z/OS connection is removed. Removing the TCP/IP connection also disables the Metro Mirror heartbeat.
- If you remove the z/OS connection first and if there is an HMC or direct connection to volumes, those volumes are not removed from the Tivoli Storage Productivity Center for Replication configuration.
- HyperSwap can run provided that volumes are attached and available to z/OS storage, even if you are using a TCP/IP connection to storage.

Perform one of these procedures to remove a storage connection:

1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is displayed in the Storage Systems view.
2. Click the **Connections** tab.
3. Select the storage connection that you want to remove.

Important: If you choose to delete an HMC connection, all storage systems that share the HMC connection will also be removed.

4. Select **Remove Connection** from the **Actions** list, and click **Go**.
5. Click **Yes** to remove the storage system.

Removing a storage system

You can remove a storage system from the IBM Tivoli Storage Productivity Center for Replication configuration.

Prerequisites: You must have Administrator privileges to remove a storage system.

Removing a storage system removes all volumes on that storage system from management server control. All copy sets with a volume on the removed storage system are removed from their respective sessions, making the target volume

unrecoverable. All connections to the removed storage system are removed, and any storage systems sharing these connections are also removed.

Perform these steps to remove a storage system:

1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is displayed in the Storage Systems view.
2. Select the storage system that you want to remove.

Important: All connections to this storage system will be removed, all volumes on the storage system will be removed from management server control, and all copy sets that have a volume on this storage system will be removed from their respective sessions, leaving the target volume unrecoverable. Any storage systems sharing these connections will be removed as well.

3. Select **Remove Storage System** from the **Actions** list, and click **Go**.
4. Click **Yes** to remove the storage system.

Modifying the location of storage systems

You can associate a location with a storage system after a connection has been made to that storage system.

Prerequisites: You must have Administrator privileges to modify the location of a storage system.

Changing the location of a storage system might have consequences. When a session has a volume role with a location that is linked to the location of the storage system, changing the location of the storage system could change the session's volume role location. For example, if there is one storage system with the location of A_Location and a session with the location of A_Location for its H1 role, changing the location of the storage system to a different location, such as B_Location, also changes the session's H1 location to Site 1. However, if there is a second storage system that has the location of A_Location, the session's role location is not changed.

Perform these steps to modify the location of a storage system:

1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is displayed in the Storage Systems view.
2. Change the location of the storage system by selecting a previously defined location from the drop-down list or type a new name in the table cell.

To disable site awareness, set the location to None.

Note: Locations are deleted from the drop-down list when there is no longer a storage system with an association to that location.

Modifying storage connection properties

You can modify the connection properties for a storage system, including IP addresses, user name and password.

Prerequisites:

- You must have Administrator privileges to modify storage connection properties.
- The storage system must be in the Disconnected state to change most storage connection parameters. You can add a secondary HMC to an existing HMC connection without the HMC being disconnected.

A storage system can lose connection to the management server, for example, if a port is blocked by a firewall or the user name or password is changed on the storage system. If the storage system loses connection, you might need to modify parameters (for example, user name or password) manually on the storage system, and then update the parameters in IBM Tivoli Storage Productivity Center for Replication.

Perform these steps to modify storage connection properties:

1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is displayed in the Storage Systems view.
2. Click the **Connections** tab.
3. Perform one of these steps to view details for a specific storage connection:
 - Click the storage connection ID.
 - Select the storage connection, click **View/modify Connection Details** from the actions list, and then click **Go**.
4. Modify the appropriate settings match the settings for the storage system.
5. Click **Apply** to continue making changes, and click **OK** when finished.

Refreshing the storage system configuration

You can refresh the storage system configuration to query the storage system for changes, such as which volumes are contained in an LSS. You might do this when you reconfigure a storage system and you want IBM Tivoli Storage Productivity Center for Replication to be aware of the changes.

Prerequisites: You must have Administrator privileges to modify storage connection settings.

Perform these steps to refresh the storage configuration:

1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is displayed in the Storage Systems view.
2. Select the storage system for which you want to refresh the configuration.
3. Select **Refresh Configuration** from the **Actions** list, and click **Go**.

Setting volume protection

To ensure that data on a volume is not overwritten, you set its status to protected. Protected volumes are excluded from replication.

You must have Administrator privileges to change the protection setting of a volume.

1. In the navigation tree, select **Storage Systems**. The Storage Systems page is opened.
2. Click **Volume Protection**. The Volume Protection wizard is opened.
3. Select a storage system.
4. Optional: Depending on the type of storage system:
 - a. Select **All IO Groups** or a specific I/O group.
 - b. Select **All Logical Storage Subsystems** or a specific logical storage subsystem.
 - c. Select **All Pools** or a specific pool.
5. Optional: In the **Volume** field, select a single volume.

6. Optional: In the **Volume Mask** field, enter a sequence of characters and wildcards that match user-defined or system-defined volume IDs. To protect a specific volume, enter the volume ID such as ESS:2105.FCA57:VOL:1000. To use a pattern to retrieve one or more volume IDs, you can enter a partial volume ID and use the wildcard character (*) to represent zero or more characters. For example, to retrieve all volume IDs that contain the characters FCA57, you enter *FCA57*.
7. Click **Next**.
8. Verify the search results, and click **Next**.
9. Click **Select All** to protect all the volumes. Alternatively, select a check box next to the volumes that you want to protect.
10. Click **Next**.
11. Click **Finish**.

Restoring data from a journal volume

This topic provides information on restoring data from a journal (J) volume. It describes how to restore data from journal volume used as part of a ESS/DS6000/DS8000 Global Mirror session, or as part of a Metro Global Mirror session, if data was corrupted on a host volume after you issued a **Recover** command. Following these steps will enable you to return to a consistent copy of the data on the host volume.

Perform the following steps to move the data from the journal volume back to the host volume:

Note: Follow these instructions only if you have already issued a **Recover** command to the site containing the journal volume. After the **Recover** command is issued, the journal volume will hold a copy of the consistent data at the time the command was issued.

1. Outside of IBM Tivoli Storage Productivity Center for Replication, using the DS8000 GUI /CLI, issue withdraw initiate background copy (issue a **rmflash -cp** command) on pairs containing the journal volume (for example, H2J2). This copies the remaining uncopied tracks from the host to the journal. Then, ensure all the Out of Sync (OOS) tracks reach zero.
2. Create a separate FlashCopy session either with IBM Tivoli Storage Productivity Center for Replication, or with the DS8000 GUI /CLI (issue a **mkflash** command with background copy), with the following conditions:
 - The journal volume (Jx) is the source volume.
 - The host volume (Hx or Lx if using a session with Practice capabilities) is the target volume.
 - x is the site the **Recover** command was issued to.

Chapter 5. Managing host systems

Host system refers to an IBM AIX 5.3 or 6.1 server that is connected to IBM System Storage DS8000 devices. A connection from Tivoli Storage Productivity Center for Replication to the host system is used to enable the automatic swap of input/output (I/O) from the primary storage unit to the secondary storage unit in the case of a primary error.

A connection to the host system is required to use the Tivoli Storage Productivity Center for Replication Open HyperSwap feature. For the software and hardware required to support Open HyperSwap, see “Setting up the environment for Open HyperSwap” on page 43.

Connecting to a host system requires the IP address or host name of the system and the port number for communication. All connections are secured via a Secure Socket Layer (SSL) connection.

Restriction: Open HyperSwap is not supported for AIX host servers that are in a clustered environment such as PowerHA (previously known as HACMP).

Related tasks:

“Managing a session with HyperSwap and Open HyperSwap replication” on page 42

HyperSwap and Open HyperSwap provide high availability of data in the case of a primary disk storage system failure. When a failure occurs in writing input/output (I/O) to the primary storage system, the failure is detected by IOS, and IOS automatically swaps the I/O to the secondary site with no user interaction and little or no application impact.

Adding a host system connection

You can add a connection to one or more host systems to the IBM Tivoli Storage Productivity Center for Replication configuration.

Prerequisites: You must have Administrator privileges to add a host system connection.

For the software and hardware required to support Open HyperSwap, see “Setting up the environment for Open HyperSwap” on page 43.

Perform these steps to add a host system connection:

1. In the navigation tree, select **Host Systems**. The Host Systems panel is displayed.
2. Click **Add Host Connection**. The Add Host Connection dialog box is displayed.
3. Enter the host name or IP address and the port for the host system and click **Add Host**. The host system is displayed on the Host Systems panel. The default port is 9930. Unless the port has been modified in Subsystem Device Driver Path Control Module (SDDPCM), use the default port.

The host system is displayed in the Host System table. If the connection is successful, the status **Connected** is displayed for the connection. If the connection was not successful, the status **Disconnected** is displayed.

Modifying a host system connection

You can modify host system connections in the IBM Tivoli Storage Productivity Center for Replication configuration.

Prerequisites: You must have Administrator privileges to modify a host system connection and the connection must be in a disconnected state.

Perform these steps to modify a host system connection:

1. In the navigation tree, select **Host Systems**. The Host Systems panel is displayed.
2. Select the host system connection that you want to modify.
3. Select **Modify Host Connection** from the **Select Action** list, and click **Go**.
4. Modify the information that is presented for the host system and click **Update Host**. The updated host system information is displayed on the Host Systems panel.

The updated host system information is displayed in the Host System table. If the connection is successful, the status **Connected** is displayed for the connection. If the connection was not successful, the status **Disconnected** is displayed.

Removing a host system connection

You can remove host system connections from the IBM Tivoli Storage Productivity Center for Replication configuration.

Prerequisites: You must have Administrator privileges to remove a host system.

Removing a host system connection disables the ability to use Open HyperSwap. Any session using the host system to provide Open HyperSwap capabilities can no longer communicate with the host and Open HyperSwap is disabled for the entire session.

Perform these steps to remove a host system connection:

1. In the navigation tree, select **Host Systems**. The Host Systems panel is displayed.
2. Select the host system connection that you want to remove.
3. Select **Remove Host Connection** from the **Select Action** list, and click **Go**.
4. Click **OK** to remove the host system connection.

Removing a session from a host system connection

You can remove a session that is associated with a host system from the IBM Tivoli Storage Productivity Center for Replication configuration. Once removed, the host no longer recognizes the session that is managing the volumes attached to that host. This function is to be used primarily for cleanup purposes.

When a session has Open HyperSwap enabled, the session communicates with the host system and the host system stores an association to that session on the IBM Tivoli Storage Productivity Center for Replication server. If the Tivoli Storage

Productivity Center for Replication server that made the association becomes inaccessible, it might be necessary to clean up and remove the session association from a different Tivoli Storage Productivity Center for Replication server.

If a host system has an associated session, the session name is displayed in the **Sessions** column of the Host Systems table. If the session is a session that is currently defined on the Tivoli Storage Productivity Center for Replication server, the session name is displayed as a link. The link opens the Session Details panel.

If the session name is not a session on the server, an icon is displayed. This session must be removed because the host system can support only a single session association. The session association must be removed before a Tivoli Storage Productivity Center for Replication server can re-establish capabilities with the host system.

Prerequisites: You must have Administrator privileges to remove a host system.

Perform these steps to remove session association from a host system connection:

1. In the navigation tree, select **Host Systems**. The Host Systems panel is displayed.
2. Select the host system connection that contains the session that you want to remove.
3. Select **Remove Session Association** from the **Select Action** list, and click **Go**.
4. Click **OK** to remove the host system connection.

The session name is removed from the **Session** column on the Host Systems table.

Chapter 6. Managing logical paths

Logical paths define the relationship between a source logical subsystem (LSS) and a target LSS that is created over a physical path. To configure logical paths for TotalStorage Enterprise Storage Server, System Storage DS8000, and DS6000, use the ESS/DS Paths panel in Tivoli Storage Productivity Center.

To configure partnerships for the following storage systems, use the graphical user interface (GUI) or command-line interface (CLI):

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified
- XIV system

Viewing logical paths

You can view all logical paths that are defined on an IBM TotalStorage Enterprise Storage Server, IBM System Storage DS8000, or IBM System Storage DS6000 storage system.

Perform one of these procedures to view logical paths:

- From the ESS/DS Paths panel of IBM Tivoli Storage Productivity Center for Replication:
 1. In the navigation tree, select **ESS/DS Paths**. The ESS/DS Paths panel is displayed.
 2. Click the storage system ID to display logical paths for that storage system.
- From the Storage Systems panel:
 1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is displayed in the Storage Systems view.
 2. Select an ESS, DS6000, or DS8000 storage system for which you want to view logical paths.
 3. Select **View Paths** from the **Select Action** list, and click **Go**. The ESS/DS Paths panel is displayed with a list of defined logical paths.

Adding logical paths

This topic describes how to add IBM TotalStorage Enterprise Storage Server, IBM System Storage DS8000, and IBM System Storage DS6000 logical paths.

Ensure that you have defined the appropriate storage systems on the Storage Systems panel.

Perform these steps to add logical paths:

1. In the IBM Tivoli Storage Productivity Center for Replication navigation tree, select **ESS/DS Paths**.
2. Click **Manage Paths**. The Path Management wizard is displayed.
3. From the drop-down boxes in the Path Management wizard, select the source storage system, source logical storage system, target storage system, and target logical storage system. Then, click **Next**.

4. From the drop-down boxes in the Path Management wizard, select the source port and target port and click **Add**. You can add multiple paths between the logical storage subsystems, or just one at a time. When you have made your selections, click **Next**.
5. Confirm your selections and click **Next**.
6. Verify the remaining wizard panels and click **Next**.
7. Click **Finish**.

Adding logical paths using a CSV file

You can create a comma separated (CSV) file to define logical paths. The CSV file specifies storage systems pairings and associated port pairings that are used for replication. IBM Tivoli Storage Productivity Center for Replication uses the port pairings defined in the CSV file to establish logical paths

Perform these steps to add IBM TotalStorage Enterprise Storage Server, IBM System Storage DS8000, and IBM System Storage DS6000 logical paths using a CSV file:

1. Create a CSV file named `portpairings.csv` in the `websphere_home/AppServer/profiles/websphere_profile/properties` directory. You can create the CSV file in a spreadsheet such as Microsoft Excel or in a text editor. An example CSV file is as follows:

```
#
# Example CSV file
#
2107.04131:2107.01532,0x0331:0x0024,0x0330:0x0100,0x0331:0x000C
2107.05131:2107.01532,0x0330:0x0029,0x0331:0x0001
```

Each line represents a storage system to storage system pairing. The first value represents the storage systems, which are delimited by a colon. The remaining values are the port pairs, which are delimited by a colon. All values are separated by a comma. Commented lines must start with #,

2. To enable the changes in the file, perform a task that requires new paths to be established. For example, suspend a session to remove the logical paths and then issue the **Start H1->H2** command to enable the paths to use the port pairings in the CSV file.

Considerations when creating and using the CSV file:

- The CSV file does not affect Global Mirror control paths.
- Port mapping is bi-directional. A logical path is established from system A to system B and from system B to system A depending on the direction of the pairs on the hardware.
- If the CSV file contains multiple lines that specify the same storage system to storage system pairing, Tivoli Storage Productivity Center for Replication uses the last line. This rule applies regardless of the order of the storage system pairing. For example, if you have storage systems 2107.04131:2107.01532 defined on the first line of the CSV file and then have 2107.01532:2107.04131 defined on the second line, Tivoli Storage Productivity Center for Replication uses second line.
- If a line in the CSV contains information that is not formatted correctly, the line is ignored. This rule includes lines that specify storage systems but do not include ports or include ports that are not formatted correctly.

- If the CSV file contains valid and invalid port pairs, the valid port pairs might or might not be established. Invalid port pairs can cause the following errors to be displayed in the Tivoli Storage Productivity Center for Replication console and in the ESS/DS Paths panel:

- Return Code F52: This error is displayed if a port is invalid.
- Return Code 0400: This error is displayed if a port is invalid and out of the range for the device.

Other storage system error codes might be displayed also, depending on the path topology, types of paths, and the incorrect port pairings specified in the CSV file.

- If the CSV file contains no valid port pairs, no logical paths are established and subsequent commands to the storage systems that require logical paths might fail. If there are existing logical paths for a storage system, those paths are used until they are removed.

Removing logical paths

This topic describes how to remove IBM TotalStorage Enterprise Storage Server, IBM System Storage DS8000, and IBM System Storage DS6000 logical paths.

1. In the IBM Tivoli Storage Productivity Center for Replication navigation tree, select **ESS/DS Paths**.
2. Click the link for the storage system that contains the paths that you want to remove.
3. Select the paths that you want to remove.
4. From the drop-down box, select **Remove**.
5. Click **Go**.

Chapter 7. Setting up data replication

This topic describes the how to set up data replication in your environment, including creating sessions and adding copy sets to those sessions.

A *session* is a container of multiple copy sets managed by a replication manager. A *copy set* is a set of volumes that contain copies of the same data. All the volumes in a copy set are the same format (count key data [CKD] or fixed block) and size. In a replication session, the number of volumes in a copy set and the role that each volume in the copy set plays are determined by the session type.

Sessions

A *session* is used to perform a specific type of data replication against a specific set of volumes. 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.

If a session has failover and failback capabilities, you can perform a site switch in which you move the application from one site to another and change the direction of the copy without having to perform a full copy. Without failover and failback capabilities, each time you move the application and writing to a different site in the session, you must initiate a full copy to synchronize the new source with the new target to regain disaster recovery capability. An IBM Tivoli Storage Productivity Center for Replication session with failover and failback capabilities uses the hardware's ability to track changes after a suspension (where applicable), so only the changed data must be resynchronized.

The type of data replication (also known as the session type) that is associated with the session determines the actions that can performed against all copy sets in the session, the number of volumes in each copy set, and the role that each volume plays.

Important: Use only the Tivoli Storage Productivity Center for Replication GUI or CLI to manage session relationships, such as volume pairs and copy sets. Do not modify session relationships through individual hardware interfaces, such as the **DSCLI**. Modifying relationships through the individual hardware interfaces can result in a loss of consistency across the relationships managed by the session, and might cause the session to be unaware of the state or consistency of the relationships.

Copy sets

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 machine type (for example, 3380 volumes must be used with other 3380 volumes and 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 (or copy type) that is associated with the session to which the copy set belongs.

Important: Use the IBM WebSphere Application Server Administrator Console to check the Java heap size (**Application servers > Server1 > Process Definition > Servant > Java Virtual Machine**) for the IBM z/OS servant region. The size of this region affects the performance of IBM Tivoli Storage Productivity Center for Replication. The default Java heap size is 512 MB, which supports fewer than 25,000 role pairs. Increasing the Java heap size to 768 MB increases support to a maximum of 50,000 role pairs. For more information on how to set up the Java heap size, see the WebSphere Application Server information center at <http://publib.boulder.ibm.com/infocenter/wasinfo/v6r1/index.jsp>.

The following tables lists the estimated number of role pairs and volumes per copy set that are supported for each session type.

Table 21. Support number of role pairs and volumes per copy set for each session type

Session	Role Pairs	Volumes
Basic HyperSwap	1	2
FlashCopy	1	2
Snapshot ¹	0	1
Metro Mirror	1	2
Metro Mirror with Practice	3	3
Global Mirror (ESS/DS)	3	3
Global Mirror (SAN Volume Controller)	1	2
Global Mirror with Practice (ESS/DS)	5	4
Global Mirror with Practice (SAN Volume Controller)	3	3
Global Mirror Two-Site Practice	8	6
Metro Global Mirror	6	4
Metro Global Mirror with Practice	8	5
1. An XIV Snapshot session requires that the user define only the H1 volumes. All target volumes are created on the same storage pool as the source volumes.		

Use the Add Copy Sets wizard to add copy sets to an existing session. You can select a storage system; a logical subsystem (LSS), I/O group, or pool; or single volume for reach role and then create one or more copy sets for the session.

You can use one of the following volume pairing options to automatically create multiple copy sets in the same session.

Storage system matching (System Storage DS8000, System Storage DS6000, or TotalStorage Enterprise Storage Server Model 800 Metro Mirror sessions only)

Creates copy sets by matching volumes (based on the volume IDs) across all logical subsystems (LSSs) for the selected storage systems. For example, volume 01 on the source LSS is matched with volume 01 on the target LSS.

You cannot select the storage system and select **All Logical Subsystems** in the list of LSSs. You can also do auto-matching at the LSS level for Metro Mirror sessions.

LSS, I/O group, or pool matching

Creates copy sets by matching all volumes based on the selected LSS, I/O group, or pool for each role in copy set.

Select the storage system and LSS, I/O group, or pool, and then select **All Volumes** in the **Volume** list.

If you do not want to use the auto-generated volume pairing for a copy set, clear that copy set so that it is not added during the wizard. Then, add the remaining copy sets and reopen the Add Copy Set wizard and manually enter the volume pairings that you want.

Invalid copy set are not added to the session. Copy sets can be invalid if their volumes are not the same type or size.

You can remove copy sets that you do not want to add to the session, even if they valid. This enables you to filter and eliminate unwanted copy sets before they are added to the session.

You can export the copy sets to take a snapshot of your session at a particular point in time for backup purposes.

Note: You can copy an entire storage system only for Metro Mirror sessions.

Considerations for adding copy sets

When you create a copy set, a warning is displayed if one or more selected volumes already exist in another session. Whether is safe to add the created copy set to the session depends on the environment. For example, if you created one session for normal replication and another session for a disaster recovery practice scenario, you must use the same target volumes from the original session as the source volumes in the practice session. If the volume you selected is already in another session, confirm that this is the configuration you want.

Considerations for removing copy sets

You remove a copy set or range of copy sets by selecting the source volume; LSS, I/O group, or pool; or storage system. When the list of copy sets that meet your criteria is displayed, you can select the copy sets that you want to remove.

The consequence of removing copy sets varies depending on the state of the session:

Defined

There is no relationship on the hardware. This removes the copy set from Tivoli Storage Productivity Center for Replication data store.

Preparing or Prepared

The copy set is currently copying data, so Tivoli Storage Productivity Center for Replication terminates the hardware relationship for the copy set. The rest of the copy sets continue to run uninterrupted.

Suspended or Target Available

Any existing relationships on the hardware are removed for the copy set.

Before removing all copy sets from that session, terminate the session. Removing the copy sets when the session is active can considerably increase the amount of time it takes for the copy set removal to complete. Copy sets are removed one at a time, and when the session is active, this requires commands being issued to the hardware. However, if you terminate the session first, then commands are not issued to the hardware and the removal process completes faster.

Tip: When you a remove copy set from Tivoli Storage Productivity Center for Replication, you might want keep hardware relationships on the storage systems.

These relationships are useful when you want to migrate from one session type to another or when resolving problems. For more information about keeping the hardware relationships when removing copy sets, see *Removing Copy Sets*.

The behavior that occurs when a copy set is removed varies depending on the storage system:

ESS 800, DS6000, and DS8000:

- The complete copy set is removed from Tivoli Storage Productivity Center for Replication.
- Any peer-to-peer remote copy (PPRC) pair that is part of a Global Mirror consistency group is removed from the consistency group on the storage system.
- If the PPRC pair is part of a Global Mirror consistency group and is the last remaining source volume in a subordinate session, the subordinate session is removed from the storage system.
- If the PPRC pair is the last remaining participant in a Global Mirror session, the Global Mirror session is removed from the storage system.
- Any PPRC relationship remains on the storage system.
- A Metro Mirror (synchronous PPRC) pair that is in a HyperSwap configuration is removed from that configuration but the pair remains on the hardware.
- FlashCopy relationship remains on the storage system if the hardware has not already completed any background copy.

SAN Volume Controller, Storwize V7000, Storwize V7000 Unified, or the XIV system

- The complete copy set is removed from Tivoli Storage Productivity Center for Replication.
- FlashCopy, Metro Mirror, and Global Mirror relationships are pulled out of their consistency group. If they are the last remaining relationship in a consistency group, that consistency group is removed from the hardware.

When you specify the force removal option, all knowledge of the specified copy set is removed from Tivoli Storage Productivity Center for Replication, even if the relationship itself still exists. If this occurs, you are not able to remove the relationship using Tivoli Storage Productivity Center for Replication, because no information about the relationship exists. If you force a removal of a copy set and the removal fails, you must manually remove the relationship from the hardware. If you do not, you cannot to create new relationships.

One benefit of forcing a removal of the copy sets is that Tivoli Storage Productivity Center for Replication does not manage the consistency of copy sets that it has no knowledge of. This means that additional commands to the session do not affect the removed copy sets, even though they are still in a relationship on the hardware.

If you do not specify the force removal option and an error occurs that prevents the hardware relationships from being removed, the copy set will not be removed from Tivoli Storage Productivity Center for Replication. The copy set remains as part of the session, and you can still perform actions on it.

To re-add the copy set to the session, you must perform a full copy of the data.

Volume roles

Volume roles are given to every volume in the copy set. The role defines how the volume is used in the copy set and, for multi-site sessions, the site location of the volume. For example, the H1 role is made up of host-attached volumes that are located at the primary site.

The site determines the location of the volumes. The number of sites in a copy set is determined by the type of data replication (also known as the session type) that is associated with the session. IBM Tivoli Storage Productivity Center for Replication supports up to three sites:

- Site 1** The location of the primary storage that contain 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 that receives the copy updates from the primary storage.
- Site 3** (Metro Global Mirror only) The location of the tertiary storage that receives data updates from the secondary storage.

The volume roles that are needed in a copy set are determined by the type of replication that is associated with the session. IBM Tivoli Storage Productivity Center for Replication supports these volume roles:

Host volume

A volume that is connected to a server that reads and writes I/O. A host volume can be the source of updated tracks when the server connected to the host volume is actively issuing read and write input/output (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 H x , where x identifies the site.

Journal volume

A volume that stores data that has 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 is created between the host or intermediate volume and the corresponding journal volume after a recover request is initiated to create another consistent version of the data.

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

Intermediate volume

A volume that 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 being used, data on intermediate volumes might not be consistent.

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

Target volume

(FlashCopy only) A volume that receives data from a source, either 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).

Role pairs

A *role pair* is the association of two volume 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 host volumes at the primary site and host volumes at the secondary site (H1-H2).

The flow of data in the role pair is shown using an arrow. For example, H1>H2 denotes that H1 is the source and H2 is the target.

Participating role pairs are role pairs that are currently participating in the session's copy.

Non-participating role pairs are role pairs that are not actively participating in the session's copy.

Snapshot sessions do not use role pairs.

Practice volumes

You can use a *practice volume* to practice what you would do in the event of a disaster, without interrupting current data replication. Practice volumes are available in Metro Mirror, Global Mirror, and Metro Global Mirror sessions.

To use the practice volumes, the session must be in the prepared state. Issuing the Flash command against the session while in the Prepared state creates a usable practice copy of the data on the target site.

Note: You can test disaster-recovery actions without using practice volumes; however, without practice volumes, you cannot continue to copy data changes between volumes while testing disaster-recovery actions.

Consistency groups

For Global Mirror and Metro Global Mirror sessions, IBM Tivoli Storage Productivity Center for Replication manages the consistency of dependant writes by creating a consistent point-in-time copy across multiple volumes or storage systems. A *consistency group* is a set of target volumes in a session that have been updated to preserve write order and are therefore recoverable.

Data exposure is the period when data is written to the storage at the primary site until data is replicated to storage at the secondary site. Data exposure is influenced by factors such as:

- Requested consistency-group interval time
- Type of storage systems
- Physical distance between the storage systems
- Available bandwidth of the data link
- Input/output (I/O) load on the storage systems

To manage data exposure, you can change the consistency group interval time. The *consistency group time interval* specifies how often a Global Mirror and Metro Global Mirror session attempts to form a consistency group. When you reduce this value, it might be possible to reduce the data exposure of the session. A lower value

causes the session to attempt to create consistency groups more frequently, which might also increase the processing load and message-traffic load on the storage systems.

Session types

Tivoli Storage Productivity Center for Replication supports several methods of data replication. The type of data replication that is associated with a given session is known as the *session type* (also known as a copy type).

The following table describes the session types that are supported by Tivoli Storage Productivity Center for Replication. Depending on the edition of Tivoli Storage Productivity Center for Replication that you are using, some of these session types might not be available.

Table 22. Session type summary

Copy type	Supported Software	Supported storage systems	Description
Basic HyperSwap	Tivoli Storage Productivity Center for Replication Basic Edition for System z and Tivoli Storage Productivity Center for Replication for System z	<ul style="list-style-type: none"> TotalStorage Enterprise Storage Server Model 800 System Storage DS8000 System Storage DS6000 	Basic HyperSwap replication is a special Metro Mirror replication method designed to provide high availability in the case of a disk storage system failure. Using Basic HyperSwap with Metro Mirror, you can configure and manage your synchronous Peer-to-Peer Remote Copy (PPRC) pairs.
FlashCopy	Tivoli Storage Productivity Center all editions	<ul style="list-style-type: none"> TotalStorage Enterprise Storage Server Model 800 System Storage DS8000 System Storage DS6000 SAN Volume Controller Storwize V7000 Storwize V7000 Unified 	<p>FlashCopy replication creates a point-in-time copy in which the target volume contains a copy of the data that was on the source volume when the FlashCopy was established. Using FlashCopy, your data exists on the second set of volumes in the same storage system, and can be restored to the first set of volumes.</p> <p>SAN Volume Controller or Storwize V7000 FlashCopy sessions are managed by using FlashCopy consistency groups. Sessions for IBM TotalStorage Enterprise Storage Server (ESS) and IBM DS6000 and DS8000 are not managed by using FlashCopy consistency groups.</p>

Table 22. Session type summary (continued)

Copy type	Supported Software	Supported storage systems	Description
Snapshot	Tivoli Storage Productivity Center all editions	The XIV system	Snapshot is a session type that creates a point-in-time copy of a volume or set of volumes without having to define a specific target volume. The target volumes of a Snapshot session are automatically created when the snapshot is created.
Metro Mirror Single Direction	Tivoli Storage Productivity Center for Replication Two Site Business Continuity	<ul style="list-style-type: none"> • TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified 	Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 KM apart. The source is located in one storage system and the target is located in another storage system. Using Metro Mirror, your data exists on the second site that is less than 300 KM away, and can be restored to the first site.
Metro Mirror Failover/Failback	Tivoli Storage Productivity Center for Replication Two Site Business Continuity	<ul style="list-style-type: none"> • TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified • The XIV system 	<p>Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 KM apart. Using Metro Mirror Failover / Failback, your data exists on the second site that is less than 300 KM away. You can use failover and failback to switch the direction of the data flow. This ability enables you to run your business from the secondary site.</p> <p>Using Metro Mirror with HyperSwap, your data exists on the second site that is less than 300 KM away. The data can be restored to the first site. You can also use failover for a backup copy of the data if your primary volumes encounter a permanent I/O error.</p>

Table 22. Session type summary (continued)

Copy type	Supported Software	Supported storage systems	Description
Metro Mirror Failover/Failback with Practice	Tivoli Storage Productivity Center for Replication Two Site Business Continuity	<ul style="list-style-type: none"> • TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified 	Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 KM apart. The source is located in one storage system and the target is located in another storage system. Metro Mirror Failover / Failback with Practice combines Metro Mirror and FlashCopy to provide a point-in-time copy of the data on the remote site.
Global Mirror Single Direction	Tivoli Storage Productivity Center for Replication Two Site Business Continuity	<ul style="list-style-type: none"> • TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified 	Global Mirror is a method of asynchronous, remote data replication between two sites that are over 300 KM apart. It maintains identical data in both the source and target, where the source is located in one storage system and the target is located in another storage system. Using Global Mirror, your data exists on the second site that is more than 300 KM away, and can be restored to the first site.
Global Mirror Either Direction with Two-Site Practice	Tivoli Storage Productivity Center for Replication Two Site Business Continuity	<ul style="list-style-type: none"> • TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 	Global Mirror is a method of asynchronous, remote data replication between two sites that are over 300 KM apart. It maintains identical data in both the source and target, where the source is located in one storage system and the target is located in another storage system. Global Mirror Either Direction with Two-Site Practice combines Global Mirror and FlashCopy to provide a point-in-time copy of the data on either the primary or secondary sites that are over 300 KM apart.

Table 22. Session type summary (continued)

Copy type	Supported Software	Supported storage systems	Description
Global Mirror Failover/Failback	Tivoli Storage Productivity Center for Replication Two Site Business Continuity	<ul style="list-style-type: none"> • TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified • The XIV system 	Global Mirror is a method of asynchronous, remote data replication between two sites that are over 300 KM apart. It maintains identical data in both the source and target, where the source is located in one storage system and the target is located in another storage system. Using Global Mirror Failover / Failback, your data exists on the second site that is more than 300 KM away, and you can use failover and failback to switch the direction of the data flow. This ability enables you to run your business from the secondary site.
Global Mirror Failover/Failback with Practice	Tivoli Storage Productivity Center for Replication Two Site Business Continuity	<ul style="list-style-type: none"> • TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified 	Global Mirror is a method of asynchronous, remote data replication between two sites that are over 300 KM apart. It maintains identical data in both the source and target, where the source is located in one storage system and the target is located in another storage system. Global Mirror Failover / Failback with Practice combines Global Mirror and FlashCopy to provide a point-in-time copy of the data on a remote site at a distance over 300 KM away from your first site.

Table 22. Session type summary (continued)

Copy type	Supported Software	Supported storage systems	Description
Metro Global Mirror	Tivoli Storage Productivity Center for Replication Three Site Business Continuity	<ul style="list-style-type: none"> • TotalStorage Enterprise Storage Server Model 800 (only H1 site) • System Storage DS8000 	<p>Metro Global Mirror is a method of continuous, remote data replication that operates between three sites that varying distances apart. Metro Global Mirror combines Metro Mirror synchronous copy and Global Mirror asynchronous copy into a single session, where the Metro Mirror target is the Global Mirror source. Using Metro Global Mirror and Metro Global Mirror with HyperSwap, your data exists on a second site that is less than 300 KM away, and a third site that is more than 300 KM away. Metro Global Mirror uses both Metro Mirror and Global Mirror Failover / Failback to switch the direction of the data flow. This ability enables you to run your business from the secondary or tertiary sites.</p> <p>Using Basic HyperSwap with Metro Global Mirror, you can configure and manage three-site continuous replication needed in a disaster recovery event.</p>

Table 22. Session type summary (continued)

Copy type	Supported Software	Supported storage systems	Description
Metro Global Mirror with Practice	Tivoli Storage Productivity Center for Replication Three Site Business Continuity	<ul style="list-style-type: none"> TotalStorage Enterprise Storage Server Model 800 (only H1 site) System Storage DS8000 	Using Metro Global Mirror with Practice, you can practice your disaster recovery actions while maintaining disaster recovery capabilities. Your data exists on a second site that is less than 300 KM away, and a third site that is more than 300 KM away. Metro Global Mirror uses both Metro Mirror and Global Mirror Failover / Failback to switch the direction of the data flow; as a result, you can run your business from the secondary or tertiary sites, and simulate a disaster.

Basic HyperSwap (ESS, DS6000, and DS8000)

Basic HyperSwap is an entitled copy services solution for z/OS version 1.9 and later. It provides high availability of data in the case of a disk storage system failure. Basic HyperSwap is not a disaster recovery solution. If a session is suspended but the suspend was not caused by a HyperSwap trigger, no freeze is done to ensure consistency of the session.

When HyperSwap is combined with Metro Mirror and Metro Global Mirror replication, you can prepare for disaster recovery and ensure high availability. If a session is suspended but the suspend was not caused by a HyperSwap trigger, a freeze is done to ensure consistency of the session.

Note: This replication method is available on only ESS, DS6000, and DS8000 storage systems, and on management servers running IBM Tivoli Storage Productivity Center for Replication Basic Edition for System z or IBM Tivoli Storage Productivity Center for Replication for System z.

Basic HyperSwap replication performs the following actions:

- Manages CKD volumes in Metro Mirror (synchronous peer-to-peer remote copy [PPRC]) relationships.
- Permits only CKD volumes to be added to the HyperSwap session. The graphical user interface (GUI) shows only CKD volumes when you add a copy set. The command-line interface (CLI) fails to add a copy set if a fixed block volume is specified.
- Monitoring for events that indicate a storage device has failed.
- Determining whether the failing storage device is part of a Metro Mirror (synchronous PPRC) pair.
- Determining the action to be taken from policy.
- Ensuring that data consistency is not violated.

- Swapping the I/O between the primary logical devices in the consistency group with the secondary logical devices in the consistency group. A swap can be from the preferred logical devices to the alternate logical devices or from the alternate logical devices to the preferred logical devices.

Metro Mirror Failover/Failback with HyperSwap

Metro Mirror Failover/Failback uses HyperSwap to configure and manage synchronous Peer-to-Peer Remote Copy (PPRC) pairs.

Metro Global Mirror with HyperSwap

Metro Global Mirror with HyperSwap is a z/OS replication feature that provides the three-site continuous replication needed in a disaster recovery event.

Important: If HyperSwap occurs by event when running a Metro Global Mirror with a HyperSwap session, a full copy of the data occurs to return to a full three-site configuration. If you issue a HyperSwap command when running a Metro Global Mirror with a HyperSwap session, a full copy does not occur. A full copy is required only for an unplanned HyperSwap or a HyperSwap initiated using the z/OS SETHS SWAP command.

Example

Jane is using multiple DS8000 storage systems. Her host applications run on z/OS and her z/OS environment has connectivity to the DS8000 storage systems. She has a site in Manhattan and a bunker in Hoboken. While she does not need a disaster recovery solution, she does need a high-availability solution to keep her applications running around the clock. Jane is worried that if a volume fails on her DS8000 in Manhattan, her database application will not be able to operate. Even a small downtime can cost Jane thousands of dollars. Jane uses a Basic HyperSwap session to mirror her data in Manhattan to her secondary DS8000 in Hoboken. If a volume at the Manhattan site fails, Basic HyperSwap automatically directs application I/O to the mirrored volumes in Hoboken.

FlashCopy


FlashCopy replication creates a point-in-time copy in which the target volume contains a copy of the data that was on the source volume when the FlashCopy was established.

The ESS, DS6000, and DS8000 platforms provide multiple logical subsystems (LSSs) within a single physical subsystem, while the following platforms provide multiple I/O groups:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

All platforms support local replication in which the source volume is located in one LSS or I/O group and the target volume is located in the same or another LSS or I/O group. Using FlashCopy, you can reference and update the source volume and target volume independently.

The following figure illustrates how a FlashCopy session works.



Example

Jane uses FlashCopy to make a point-in-time copy of the customer data in existing international accounts. Every night, the bank's servers perform batch processing. Jane uses FlashCopy to create checkpoint restarts for the batch processing in case the batch processing fails. In her batch processing, the first step is to balance all international accounts, with a FlashCopy taken of the resulting data. The second step in her batch processing is to process the international disbursements. If the second step in the batch process fails, Jane can use the FlashCopy made of the first step to repeat the second step, instead of beginning the entire process again. Jane also writes a CLI script that performs a FlashCopy every night at 11:59 PM, and another script that quiesces the database. She backs this data up on tape on her target storage system, and then sends the tape to the bank's data facility in Oregon for storage.

Snapshot

Snapshot is a session type that creates a point-in-time copy of a volume or set of volumes. You do not have to define a specific target volume. The target volumes of a Snapshot session are automatically created when the snapshot is created.

The XIV system uses advanced snapshot architecture to create a large number of volume copies without affecting performance. By using the snapshot function to create a point-in-time copy, and to manage the copy, you can save storage. With the XIV system snapshots, no storage capacity is used by the snapshot until the source volume (or the snapshot) is changed.

The following figure illustrates how a Snapshot session works.



Example

Jane's host applications are using an XIV system for their back-end storage. With the XIV system, Jane can create a large number of point-in-time copies of her data. The snapshot function ensures that if data becomes corrupted, she can restore the data to any number of points in time.

Jane sets up a Snapshot session by using Tivoli Storage Productivity Center for Replication and specifies the volumes on the XIV system that are used by her host applications. Jane does not have to provision target volumes for all the snapshots she intends to make. She can quickly get a single Snapshot session configured and ready.

When the session is configured, Jane writes a CLI script that performs a Create Snapshot command to the session every two hours. If a problem occurs, such as data becoming corrupted, Jane can find a snapshot of the data before the problem occurred. She can restore the data to that point.

By creating a set of snapshots of the data, Jane can also schedule batch processing against that data every day. She can use the batch processing to analyze certain trends in the market without causing any effect to the host applications.

Metro Mirror

Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 KM apart. The source is located in one storage system and the target is located in another storage system.

Attention: For Tivoli Storage Productivity Center for Replication for System z sessions containing Metro Mirror relationships, ensure that the session does not contain system volumes (such as paging volumes) unless the session is enabled for HyperSwap. If HyperSwap is not enabled, a freeze that is issued through Tivoli Storage Productivity Center for Replication might cause Tivoli Storage Productivity Center for Replication processing to freeze. This situation might prevent the session from ensuring that the data is consistent.


Metro Mirror replication maintains identical data in both the source and target. When a write is issued to the source copy, the changes made to the source data are propagated to the target before the write finishes posting. If the storage system goes down, Metro Mirror provides zero loss if data must be used from the recovery site.

A Metro Mirror session in Global Copy mode creates an asynchronous relationship to accommodate the high volume of data migration. As a result, the data on the target system might no longer be consistent with the source system. The Metro Mirror session switches back to a synchronous relationship when a Metro Mirror Start command is reissued. In addition, you can start a Metro Mirror session in Global Copy mode and toggle between Metro Mirror and Global Copy modes to accommodate time periods in which you value host input/output (I/O) response time over data consistency.

Tip: To determine if there is any out of synch data that must be copied over before the session is consistent, check the percent that is complete in the session details panel.

Metro Mirror Single Direction


The following figure illustrates how a Metro Mirror Single Direction session works.



Metro Mirror Failover/Failback

Using Metro Mirror with failover/failback, your data exists on the second site that is less than 300 KM away, and you can use failover/failback to switch the direction of the data flow. This session type enables you to run your business from the secondary site, and to copy changes made at the second site back to the primary site when you want to resume production at the primary site.

The following figure illustrates how a Metro Mirror with Failover/Failback session works.




Metro Mirror Failover/Failback with Practice

A Metro Mirror Failover/Failback with Practice session combines Metro Mirror and FlashCopy to provide a point-in-time copy of the data on the remote site. You can use this session type to practice what you might do if a disaster occurred, without losing your disaster recovery capability.

This solution consists of two host volumes and an intermediate volume.

The following figure illustrates how a Metro Mirror Failover/Failback with Practice session works.



Metro Mirror Failover/Failback with HyperSwap

A Metro Mirror Failover/Failback session can be enabled to have HyperSwap capabilities. To enable HyperSwap the following circumstances must apply:


- The session is running on an Tivoli Storage Productivity Center for Replication server running on IBM z/OS.
- The volumes are only for TotalStorage Enterprise Storage Server, System Storage DS8000, and DS6000 systems.
- The volumes are count key data (CKD) volumes that are attached to the z/OS system.

Metro Mirror Failover/Failback with HyperSwap combines the high availability of Basic HyperSwap with the redundancy of a two-site Metro Mirror Failover/Failback solution when managing count key data (CKD) volumes on z/OS. If the primary volumes encounter a permanent I/O error, the I/O is automatically swapped to the secondary site with little to no impact on the application.

A swap can be planned or unplanned. A planned swap occurs when you issue a HyperSwap command from the **Select Action** list in the graphical user interface (GUI) or when you issue a `cmdsess -action hyperswap` command.

The following figure illustrates how a Metro Mirror Failover/Failback session enabled for HyperSwap works.

Metro Mirror Failover/Failback



For more information about enabling HyperSwap, see “Managing a session with HyperSwap and Open HyperSwap replication” on page 42.

Metro Mirror Failover/Failback with Open HyperSwap

A Metro Mirror Failover/Failback session can be enabled to have Open HyperSwap capabilities. To enable Open HyperSwap the following circumstances must apply:

- The volumes in the session are System Storage DS8000 5.1 or later volumes.
- The volumes in the session are fixed block and are mounted to IBM AIX 5.3 or AIX 6.1 hosts with the following modules installed:
 - Subsystem Device Driver Path Control Module (SDDPCM) version 3.0.0.0 or later
 - Multi-Path Input/Output (MPIO) module (the version that is provided with AIX version 5.3 or 6.1)
- The connections between the AIX host systems and the Tivoli Storage Productivity Center for Replication server have been established.


Metro Mirror Failover/Failback with Open HyperSwap combines the high availability of Basic HyperSwap on z/OS for fixed block AIX volumes with the redundancy of a two-site Metro Mirror Failover/Failback solution. If the primary

volumes encounter a permanent I/O error, the I/O is automatically swapped to the secondary site with little to no impact on the application.

A swap can be planned or unplanned. A planned swap occurs when you issue a HyperSwap command from the **Select Action** list in the GUI or when you issue a `cmdsess -action hyperswap` command.

The following figure illustrates how a Metro Mirror Failover/Failback session enabled for Open HyperSwap works.

Metro Mirror Failover/Failback



For more information about enabling Open HyperSwap, see “Managing a session with HyperSwap and Open HyperSwap replication” on page 42.

Examples

Metro Mirror Single Direction

At the beginning of a work week, Jane is notified that between 10:00 AM and 11:00 AM on the next Friday, power in her building is going to be shut off. Jane does not want to lose any transactions during the power outage, so she decides to transfer operations to the backup site during the outage. She wants a synchronous copy method with no data loss for her critical business functions, so she chooses Metro Mirror, which can be used between locations that are less than 300 KM apart.

In a synchronous copy method, when a write is issued to change the source, the change is propagated to the target before the write is completely posted. This method of replication maintains identical data in both the source and target. The advantage of this is when a disaster occurs, there is no data loss at the recovery site because both writes must complete before signaling completion of a write to the source application. Because the data must be copied to both System Storage DS8000 devices before the write is completed, Jane can be sure that her data is safe.

The night before the planned outage, Jane quiesces her database and servers in San Francisco and starts the database and servers in Oakland. To accomplish this task, Jane issues the Suspend and Recover commands, and then issues the Start command on the secondary site. She powers off her equipment in San Francisco to avoid any power spikes during reboot after the power is turned back on.

Metro Mirror in Global Copy mode

At the beginning of a work week, Jane is notified that between 10:00 AM and 11:00 AM on the next Friday, power in her building is going to be shut off. Jane does not want to lose any transactions during the power outage, so she decides to transfer operations to their backup site during the outage. She wants a synchronous copy method with no data loss for her critical

business functions, so she chooses Metro Mirror, which can be used between locations that are less than 300 KM apart.

Jane wants to limit her application impact while completing the initial Metro Mirror synchronization, so she begins her session in Global Copy mode. After she sees that about 70% of the data has been copied, Jane decides to switch the session into Metro Mirror mode, assuring data consistency.

Metro Mirror with Practice

Jane wants to run a Metro Mirror with Practice from San Francisco to Oakland. She wants to verify her recovery procedure for the Oakland site, but she cannot afford to stop running her Metro Mirror session while she takes time to practice a recovery. By using a Metro Mirror with Practice session, Jane is able to practice her disaster recovery scenario in Oakland while her Metro Mirror session runs uninterrupted. By practicing running her applications at the Oakland site, Jane is better prepared to make a recovery if a disaster ever strikes the San Francisco site.

While her session is running in a Prepared state, Jane practices a recovery at her Oakland site by issuing the Flash command. This command momentarily pauses the session and starts a FlashCopy to the H2 volumes. As soon as the FlashCopy is started, her session will be restarted. These FlashCopy files create a consistent version of the data on the H2 volume that she can use for recovery testing, while her session continues to replicate data from San Francisco to Oakland. As a result, she can carry out her recovery testing without stopping her replication for any extended duration of time.

If at some point in time, the Metro Mirror session suspends due to a failure, Jane can use the practice session to restart her data replication while maintaining a consistent copy of the data at the Oakland site, in case of a failure during the resynchronization process. When the session is suspended, she can issue a Recover command to create a consistent version of the data on the H2 volumes. After the Recover command completes, she can issue the Start H1->H2 command to resynchronize the data from the San Francisco site to the Oakland site. If a failure occurs before her restarted session is in the Prepared state, she has a consistent version of the data on the H2 volumes. She must simply issue the Recover command to put the session into Target Available state and make the H2 volumes accessible from her servers. If the session was not in the Prepared state when it suspended, the subsequent Recover command does not issue the FlashCopy files to put the data on the H2 volumes. This means that the consistent data on the H2 volumes are not overwritten if the data to be copied to them is not consistent.

Metro Mirror Failover/Failback enabled for Open HyperSwap

Jane wants to run a Metro Mirror with Practice from San Francisco to Oakland. She wants to verify her recovery procedure for the Oakland site, but she cannot afford to stop running her Metro Mirror session while she takes time to practice a recovery. By using a Metro Mirror with Practice session, Jane is able to practice her disaster recovery scenario in Oakland while her Metro Mirror session runs uninterrupted. By practicing running her applications at the Oakland site, Jane is better prepared to make a recovery if a disaster ever strikes the San Francisco site.

While her session is running in a Prepared state, Jane practices a recovery at her Oakland site by issuing the Flash command. This command

momentarily pauses the session and starts a FlashCopy to the H2 volumes. As soon as the FlashCopy is started, the session is restarted. These FlashCopy files create a consistent version of the data on the H2 volume that she can use for recovery testing, while her session continues to replicate data from San Francisco to Oakland. As a result, she can carry out her recovery testing without stopping her replication for any extended duration of time.

If the Metro Mirror session suspends due to a failure, Jane can use the practice session to restart her data replication while maintaining a consistent copy of the data at the Oakland site, in case of a failure during the resynchronization process. When the session is suspended, she can issue a Recover command to create a consistent version of the data on the H2 volumes. After the Recover command completes, she can issue the Start H1->H2 command to resynchronize the data from the San Francisco site to the Oakland site. If a failure occurs before her restarted session is in the Prepared state, she has a consistent version of the data on the H2 volumes. She must simply issue the Recover command to put her session into Target Available state and make the H2 volumes accessible from her servers. If the session was not in the Prepared state when it suspended, the subsequent Recover command does not issue the FlashCopy files to put the data on the H2 volumes. This means that the consistent data on the H2 volumes are not overwritten if the data to be copied to the volumes is not consistent.

Selecting a HyperSwap session

A global insurance company has elected to use Tivoli Storage Productivity Center for Replication to manage its disaster recovery environment. Jane wants minimal data exposure, both for planned outages such as routine maintenance, and for unplanned disasters. They have CKD volumes on System Storage DS8000 devices, and use z/OS mainframes. They have two data centers in New York.

Jane reviews the Tivoli Storage Productivity Center for Replication documentation, and chooses a Metro Mirror recovery solution, since her priority is being protected against regional disasters. Jane decides to use Metro Mirror solution, because her company has two data centers located near each other. Jane realizes that because she uses z/OS, CKD, and System Storage DS8000 hardware, she is also able to use a HyperSwap solution. Using Metro Mirror Failover/Failback with HyperSwap, Jane can minimize application impact, while maintaining seamless failover to her secondary site. Jane decides Metro Mirror Failover/Failback with HyperSwap is best for the needs of her company.

After installing and configuring Tivoli Storage Productivity Center for Replication on z/OS, Jane starts the Tivoli Storage Productivity Center for Replication GUI. She adds the Tivoli Storage Productivity Center for Replication storage devices she intends to use on all three sites. From the Session Overview panel, Jane launches the Create Session wizard, and selects the **Metro Mirror Failover/Failback** session type. Then, she continues the wizard, she selects the **Manage H1H2 with HyperSwap** option. After finishing the wizard, Jane clicks **Launch Add Copy Sets Wizard**. She completes this wizard, and issues a **Start H1->H2** command. After the initial copy is completed, Jane is safely replicating her data between both sites. She can also issue a HyperSwap between sites 1 and 2, enabling her to switch sites with minimal application impact during either a disaster or maintenance period.

Performing a planned HyperSwap

Jane's company has successfully been using Metro Mirror Failover/Failback with HyperSwap sessions for the past three months. However, Jane needs to perform maintenance on an H1 box. During this time, Jane does not want her applications or replication to be interrupted. To prevent this from happening, shortly before the maintenance is scheduled to begin, Jane decides to use the Tivoli Storage Productivity Center for Replication GUI to perform a HyperSwap to the H2 volumes. This transitions the applications so that they write to H2. To perform a planned HyperSwap, Jane issues a HyperSwap command.

Understanding what happens when an unplanned HyperSwap occurs

Several weeks after the planned maintenance at Jane's company is completed, an incident occurs at the H1 site. A disk controller fails, causing one of the H1 volumes to encounter a permanent I/O error. Fortunately, Jane's data is safe because she used Metro Mirror Failover/Failback with HyperSwap, and her H2 volume is an exact duplicate of the H1 volume. When the permanent I/O error is detected, a HyperSwap is triggered. The application seamlessly transitions to writing to the H2 volumes. Her data is safe, and her applications are not interrupted.

Jane configured an Simple Network Management Protocol (SNMP) listener to alert her to any events, so she receives the SNMP event indicating that a HyperSwap occurred. Jane investigates the cause of the HyperSwap and uses the z/OS console to identify the volume that triggered the HyperSwap. Jane replaces the faulty disk controller. Then, to recover from the unplanned HyperSwap, Jane issues the **Start H2->H1** command.

Global Mirror


Global mirror is a method of asynchronous, remote data replication between two sites that are over 300 kilometers (km) apart. It maintains identical data in both the source and target, where the source is located in one storage system and the target is located in another storage system.

The data on the target is typically written a few seconds after the data is written to the source volumes. When a write is issued to the source copy, the change is propagated to the target copy, but subsequent changes are allowed to the source before the target verifies that it has received the change. Because consistent copies of data are formed on the secondary site at set intervals, data loss is determined by the amount of time since the last consistency group was formed. If your system stops, Global Mirror might lose some data that was being transmitted when the disaster occurred. Global Mirror still provides data consistency and data recoverability in the event of a disaster.

Global Mirror Single Direction

A Global Mirror single direction session allows you to run your Global Mirror replication from only the primary site.

For ESS, DS6000, and DS8000 storage systems, each copy set in the Global Mirror single direction session consists of two host volumes and a journal volume. The following figure illustrates how a Global Mirror single direction session works on an ESS, DS6000, and DS8000 storage systems:



For the following storage systems, each copy set in the Global Mirror Single Direction session consists of two host volumes:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

The following figure illustrates how a Global Mirror Single Direction session works on these storage systems:




Global Mirror Either Direction with Two-Site Practice (ESS, DS6000, and DS8000)

A Global Mirror Either Direction with Two Site Practice session allows you to run Global Mirror replication from either the primary or secondary site. It combines Global Mirror and FlashCopy to provide a point-in-time copy of the data on a remote site at a distance over 300 km away from your first site. This practice session allows you to create practice volumes on both the primary and secondary site to practice what you might do if a disaster occurred, without losing your disaster recovery capability.

Note: This replication method is available on only ESS, DS6000, and DS8000 storage systems.


The session consists of two host volumes, two intermediate volumes, and two journal volumes. The following figure illustrates how a Global Mirror either direction with two-site practice session works:



Global Mirror Failover/Failback

Using Global Mirror failover/failback, your data exists on the second site that is more than 300 km away, and you can use failover/failback to switch the direction of the data flow. This enables you to run your business from the secondary site.

For ESS, DS6000, and DS8000 storage systems, each copy set in the Global Mirror failover/failback session consists of two host volumes and a journal volume. The following figure illustrates how a Global Mirror failover/failback session works on an ESS, DS6000, or DS8000 storage systems:



For the following storage systems, each copy set in the Global Mirror failover/failback session consists of two host volumes:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified
- The XIV system


The following figure illustrates how a Global Mirror failover/failback session works on these storage systems.



Global Mirror Failover/Failback with Practice

A Global Mirror failover/failback with practice session combines Global Mirror and FlashCopy to provide a point-in-time copy of the data on a remote site at a distance over 300 km away from your primary site. You can use this to practice what you might do if a disaster occurred, without losing your disaster recovery capability.


For ESS, DS6000, and DS8000 storage systems, each copy set in the Global Mirror failover/failback with practice session consists of two host volumes, an intermediate volume, and a journal volume. The following figure illustrates how a Global Mirror failover/failback with practice session works on an ESS, DS6000, or DS8000 storage system:



For the following storage systems, each copy set in the Global Mirror failover/failback with Practice session consists of two host volumes and an intermediate volume:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

The following figure illustrates how a Global Mirror Failover/Failback with Practice session works on these storage systems:



Examples

Global Mirror Single Direction

Although Jane's FlashCopy and Metro Mirror copies were both planned, Jane realizes that sometimes unforeseen things happen, and she wants to make sure her data is safe. Because Jane works in San Francisco, she wants her other site to be far away in case of a localized disaster. Her other site is based in Houston. Jane's foresight pays off when a minor earthquake occurs in San Francisco and power and communications both go down. Fortunately, Jane has arranged for the data on customer accounts that have recently opened or closed to be asynchronously copied in Houston, using Global Mirror. Jane risks losing the bytes of data that were being processed when the tremor disrupted the San Francisco process, but she views that as a minor inconvenience when weighed next to the value of backing up her data in a non-earthquake zone.

Global Mirror with Practice

Jane wants to run a Global Mirror with practice from San Francisco to Houston. She wants to verify her recovery procedure for the Houston site, but she cannot afford to stop running her Global Mirror session while she takes time to practice a recovery. By using a Global Mirror with Practice session, Jane is able to practice her disaster recovery scenario in Houston while her Global Mirror session runs uninterrupted. By practicing running her applications at the Houston site, Jane will be better prepared to make a recovery if a disaster ever strikes the San Francisco site.

Global Mirror Either Direction with Two-Site Practice

Jane wants to run a Global Mirror with practice from San Francisco to Houston. She wants to verify her recovery procedure for the Houston site, but she cannot afford to stop running her Global Mirror session while she takes time to practice a recovery. By using a Global Mirror either direction two-site with practice session, Jane is able to practice her disaster recovery scenario in Houston while her Global Mirror session runs uninterrupted. By practicing running her applications at the Houston site, Jane will be better prepared to make a recovery if a disaster ever strikes the San Francisco site.

Jane can use the Global Mirror either direction with two-site practice session to run asynchronous consistent data replication from either the San Francisco site or the Houston site. (She can practice her disaster recovery at the target site, no matter what her current production site is.) Jane's business is able to run a consistent Global Mirror session from its Houston site back to San Francisco while running a production at the Houston site.

Setting up Global Mirror for Resource Groups on System Storage DS8000

If resource groups are defined on a System Storage DS8000, Global Mirror session IDs might be defined for some users. Tivoli Storage Productivity Center for Replication does not automatically determine which session IDs are valid. To determine which session IDs are valid, you must modify the `rmserver.properties` file and add the following property:

```
gm.master.sessionid.gm_role,session_name = xx
```

where `gm_role` is the role that has the master volume (for example, H1 in a Global Mirror failover/failback session), `session_name` is the name of the session that uses the session ID, and `xx` is the decimal number for the session ID.

Important: System Storage DS8000 represents session IDs as a two-digit hexadecimal number. Use the decimal version of that number. For example, if you want a Global Mirror failover/failback session to use a session ID of 0F, the decimal number is 15 as shown in the following example:

```
gm.master.sessionid.H2.11194_wprac=15
```

Metro Global Mirror (ESS 800 and DS8000)

Metro Global Mirror is a method of continuous, remote data replication that operates between three sites that varying distances apart. Metro Global Mirror combines Metro Mirror synchronous copy and Global Mirror asynchronous copy into a single session, where the Metro Mirror target is the Global Mirror source.

Notes:

- This replication method is available on only ESS 800 and DS8000 storage systems.
- You can select ESS 800 storage systems in only the H1 volume role. All other volume roles must use DS8000 volumes.
- You can mix ESS 800 and DS8000 volumes in the H1 volume role. If ESS 800 and DS8000 storage systems are both used in the H1 role, the DS8000 storage system performs Incremental Resync (IR), and the ESS 800 storage system performs a full copy. Because ESS 800 does not support the IR function, a full copy is required when changing from H1->H2->H3 to H1>H3 and from H2->H1->H3 to H2->H3.

Metro Global Mirror maintains a consistent copy of data at the remote site, with minimal impact to applications at the local site. This remote mirroring function works in combination with FlashCopy to meet the requirements of a disaster-recovery solution by providing the following features:

- Fast failover and failback
- Rapid reestablishment of three-site mirroring, without production outages
- Data currency at the remote site with minimal lag behind at the local site, an average of only 3 - 5 seconds for many environments
- Quick resynchronization of mirrored sites using only incremental changes

If IBM Tivoli Storage Productivity Center for Replication is running on z/OS, you can configure a Metro Global Mirror session to control the Metro Mirror relationship between the primary and secondary site using HyperSwap. With HyperSwap enabled, a failure on the primary storage system causes an automatic HyperSwap, transparently redirecting application I/O to the auxiliary storage

system. The Global Mirror relationship continues to run uninterrupted throughout this process. With this configuration, you can achieve near zero data loss at larger distances.

Using synchronous mirroring, you can switch from local site H1 to remote site H2 during a planned or unplanned outage. It also provides continuous disaster recovery protection of site H2 through site H3, without the necessity of additional reconfiguration, if a switch from site H1 occurs. With this configuration, you can reestablish H2->H1->H3 recoverability while production continues to run at site H2. Additionally, this cascaded setup can reduce the load on-site H1 as compared to some multi-target (non-cascaded) three-site mirroring environments.


Important:

- If HyperSwap occurs by event when running a Metro Global Mirror with a HyperSwap session, a full copy of the data occurs to return to a full three-site configuration. If you issue a HyperSwap command when running a Metro Global Mirror with a HyperSwap session, a full copy does not occur. A full copy is required only for an unplanned HyperSwap or a HyperSwap initiated using the z/OS SETHS SWAP command.
- In Metro Global Mirror and Metro Global Mirror with Practice sessions, when the H1 is on an ESS 800, you might risk filling up the space efficient journal volumes. Because incremental resynchronization is not supported on the ESS 800, full copies are performed in many of the transitions.

Metro Global Mirror

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

The following figure illustrates how a Metro Global Mirror session works.




Metro Global Mirror with Practice

A Metro Global Mirror session with Practice combines Metro Mirror, Global Mirror, and FlashCopy across three sites to provide a point-in-time copy of the data on the third site. You can use this session to practice what you might do if a disaster occurred without losing your disaster recovery capability.

The session consists of three host volumes, an intermediate volume, and a journal volume. The following figure illustrates how a Metro Global Mirror with Practice session works.

Note: A Metro Global Mirror with practice session can be created when the three-site license has been applied to the server.



Note: In Metro Global Mirror and Metro Global Mirror with Practice sessions, when the H1 is on an ESS 800, you might risk filling up the space efficient journal volumes. Because incremental resynchronization is not supported on the ESS 800, full copies are performed in many of the transitions.

Examples

Metro Global Mirror

Although Jane works in San Francisco, she wants to give herself the ability to run her business from either Oakland (her secondary site) or Houston (her tertiary site). Jane can use Metro Global Mirror with failover/failback to switch the direction of the data flow, so that she can run her business from either Oakland or Houston. Metro Global Mirror means that Jane has zero data loss backup at her secondary site, and minimal data loss at her tertiary site.

Metro Global Mirror with Practice

Jane wants to run a Metro Global Mirror with Practice from San Francisco to Houston. She wants to verify her recovery procedure for the Houston site. However, she cannot afford to stop running her Metro Global Mirror session while she takes time to practice a recovery. By using a Metro Global Mirror with Practice session, Jane is able to practice her disaster recovery scenario in Houston while her Metro Global Mirror session runs uninterrupted. By practicing running her applications at the Houston site, and being prepared to run her applications at the Oakland site if necessary, Jane will be better prepared to make a recovery if a disaster ever strikes the San Francisco site.

Jane can use Metro Global Mirror with Practice to switch the direction of the data flow, so that she can run her business from either Oakland or Houston. Using Metro Global Mirror, Jane has zero data loss backup at her secondary site; and minimal data loss at her tertiary site.

Managing a session with HyperSwap and Open HyperSwap replication

HyperSwap and Open HyperSwap provide high availability of data in the case of a primary disk storage system failure. When a failure occurs in writing input/output (I/O) to the primary storage system, the failure is detected by IOS, and IOS automatically swaps the I/O to the secondary site with no user interaction and little or no application impact.

Sessions that can be enabled for HyperSwap or Open HyperSwap

You can create sessions that have HyperSwap or Open HyperSwap capabilities. Enabling swapping provides a session with a highly available business continuity solution.

Sessions that can enable HyperSwap

The following session types can enable HyperSwap:

- Basic HyperSwap
- Metro Mirror with Failover/Failback
- Metro Global Mirror

To enable HyperSwap, the following circumstances must apply:

- The session is running on an Tivoli Storage Productivity Center for Replication server that is running on IBM z/OS.
- The volumes are only for TotalStorage Enterprise Storage Server, System Storage DS8000, and DS6000 systems.
- The volumes are count key data (CKD) volumes that are attached to the z/OS system.

Sessions that can enable Open HyperSwap

Only Metro Mirror with Failover/Failback session type can enable Open HyperSwap.

To enable Open HyperSwap, the following circumstances must apply:

- The volumes in the session are only for System Storage DS8000 5.1 or later volumes.
- The volumes in the session are fixed block and mounted to IBM AIX 5.3 or AIX 6.1 hosts with the following modules installed:
 - Subsystem Device Driver Path Control Module (SDDPCM) version 3.0.0.0 or later
 - Multi-Path Input/Output (MPIO) module (the version that is provided with AIX version 5.3 or 6.1)
- The connections between the AIX host systems and the Tivoli Storage Productivity Center for Replication server have been established.

Setting up the environment for HyperSwap

You must set up an environment that supports HyperSwap before attempting to enable HyperSwap for a IBM Tivoli Storage Productivity Center for Replication session.

The following steps must be completed before HyperSwap can be enabled. For more information about these steps, see the IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide

1. Install IBM Tivoli Storage Productivity Center for Replication for System z.
2. Perform the post installation tasks of setting up the data store and other necessary system settings.
3. Ensure that all RESERVEs are converted to global enqueues (ENQs).
4. Ensure that all volumes in the session that you are enabling for HyperSwap are attached to the IBM z/OS system that is running Tivoli Storage Productivity Center for Replication.

Setting up the environment for Open HyperSwap

You must set up an environment that supports Open HyperSwap before attempting to enable Open HyperSwap for a IBM Tivoli Storage Productivity Center for Replication session.

The following steps must be completed before Open HyperSwap can be enabled:

1. Ensure that the IBM AIX hosts and IBM System Storage DS8000 to meet the following hardware and software requirements:

AIX requirements

Open HyperSwap support requires AIX version 5.3 or 6.1. (You can find the supported AIX version for each Tivoli Storage Productivity Center for Replication release in the support matrix at http://www-01.ibm.com/support/docview.wss?rs=40&context=SSBSEX&context=SSMN28&context=SSMMUP&context=SS8JB5&context=SS8JFM&uid=swg21386446&loc=en_US&cs=utf-8&lang=en. Click the link for the applicable release under Agents, Servers and GUI.)

You must have the following AIX modules installed:

- Subsystem Device Driver Path Control Module (SDDPCM) version 3.0.0.0 or later
- Multi-Path Input/Output (MPIO) module (the version that is provided with AIX version 5.3 or 6.1)

System Storage DS8000 hardware requirements

Only System Storage DS8000 storage systems are supported. Open HyperSwap requires System Storage DS8000 5.1 or later.

Open HyperSwap does not support High Availability Cluster Multi-Processing (HACMP).

2. Create connections from Tivoli Storage Productivity Center for Replication to the AIX hosts (see “Adding a host system connection” on page 111).
3. Assign copy set volumes from the storage device to the host using the System Storage DS8000 command-line interface (CLI) or graphical user interface (GUI).
4. Run the AIX `cfgmgr` command to discover the volumes assigned to the host.

Considerations for Open HyperSwap and the AIX host:

- A single session that has Open HyperSwap enabled can manage multiple hosts; however, each host can be associated with only one session. Multiple hosts can share the same session.
- For AIX 5.3, a single host can manage a maximum of 1024 devices that have been enabled for Open HyperSwap on the host, with 8 logical paths configured for each copy set in the session. For AIX 6.1, a single host can manage a maximum of 1024 devices that have been enabled for Open HyperSwap on the host, with 16 logical paths configured for each copy set in the session.

- If an application on the host has opened a device, a Tivoli Storage Productivity Center for Replication session for that device cannot be terminated. The **Terminate** command fails. To terminate the session, you must either close the application or remove the copy sets from the session. If you remove copy sets from the session, you must ensure that the application writes to the correct volume when the copy set relationship is restored.
- It is possible for Open HyperSwap to fail on a subset of hosts for the session and work on the remaining hosts for the same session. In this situation, you must determine the best action to take if the application is writing to volumes on the source system as well as volumes on the target system. Contact the IBM Support Center if you need assistance determining the best solution for this issue.
- To enable support for Open HyperSwap on the host, refer to the *IBM System Storage Multipath Subsystem Device Driver User's Guide*.

Configuring timers to support Open HyperSwap:

There are configurable timeout values for the storage system, IBM Tivoli Storage Productivity Center for Replication, and IBM AIX hosts systems that can affect the operation of Open HyperSwap.

The following list describes the various timeout values that can affect Open HyperSwap:

Storage system quiesce timeout value

This quiesce timeout timer begins when the storage system starts a quiesce operation. When the timer value expires, input/output (I/O) is resumed on the primary device. The default timeout value is two minutes, but the value can be set from 30 to 600 seconds. To set the quiesce timeout value, see the information about the **chdev** command in the *IBM System Storage Multipath Subsystem Device Driver User's Guide*.

Storage system long busy timeout value

This timeout value is the time in seconds that the logical subsystem (LSS) consistency group volume stays in the long busy state after a remote mirror and copy error is reported.

Timeout values for the applications that are on the host

The various applications that are running on the host have timeout values. The timeout values vary depending on the application.

Considerations for setting timers

Consider the following information for setting timers:

- If the host quiesce timer is set to a shorter value than the Tivoli Storage Productivity Center for Replication response timer, it is possible that an I/O swap failure can occur. If a storage system triggers an unplanned failover and if the storage system quiesce timer expires before Tivoli Storage Productivity Center for Replication responds, the host attempts to write I/O to the primary volume where the loss of access occurred. If the hardware condition that caused the loss of access continues, the attempt to write I/O fails again and an unplanned Open HyperSwap is not performed.
- If the host quiesce timer is set to a longer value than the Tivoli Storage Productivity Center for Replication response timer, an application timeout might occur if Open HyperSwap takes too long to complete.

Enabling a Session for HyperSwap or Open HyperSwap

Enabling HyperSwap or Open HyperSwap for a session provides a combined business recovery and business continuity solution.

To ensure that your environment supports HyperSwap or Open HyperSwap, see “Setting up the environment for HyperSwap” on page 43 or “Setting up the environment for Open HyperSwap” on page 43.

Perform these steps to enable HyperSwap or Open HyperSwap for a session.

1. In the IBM Tivoli Storage Productivity Center for Replication navigation tree, select **Sessions**. Click the radio button next to the session that you want to enable.
2. From the **Select Action** menu, select **View/Modify Properties** and click **Go**. If you have not created the session, click **Create Session**. You can enable HyperSwap or Open HyperSwap on the Properties page
3. Under **ESS / DS Metro Mirror Options**, select from the following HyperSwap or Open HyperSwap options:
 - **Manage H1-H2 with HyperSwap**. This option enables a session to manage the H1-H2 sequence using HyperSwap. If you select this option, select from the following additional options.
 - **Disable HyperSwap**. Select this option to prevent a HyperSwap from occurring by command or event.
 - **On Configuration Error**. Choose one of the following options:
 - **Partition the system(s) out of the sysplex**. Select this option to partition out of the sysplex when a new system is added to the sysplex and encounters an error in loading the configuration. A restart of your system is required if you select this option.
 - **Disable HyperSwap**. Select this option to prevent a HyperSwap from occurring by command or event.
 - **On Planned HyperSwap Error**. Choose one of the following options:
 - **Partition out the failing system(s) and continue swap processing on the remaining system(s)**. Select this option to partition out the failing system and continues the swap processing on any remaining systems.
 - **Disable HyperSwap after attempting backout**. Select this option to stop the HyperSwap action, and disable the HyperSwap commands or events.
 - **On Unplanned HyperSwap**. Choose one of the following options:
 - **Partition out the failing system(s) and continue swap processing on the remaining system(s)**. Select this option to partition out the failing systems and continues the HyperSwap processing on the remaining systems when a new system is added to the sysplex and HyperSwap does not complete. A restart of your system is required if you select this option.
 - **Disable HyperSwap after attempting backout**. Select this option to stop the HyperSwap action, and disable the HyperSwap commands or events.
 - **Manage H1-H2 with Open HyperSwap**. If volumes are attached to an IBM AIX host, Tivoli Storage Productivity Center for Replication can manage the H1-H2 sequence of a Metro Mirror session using Open HyperSwap. If this option is selected, a failure on the host accessible volumes triggers a swap, which redirects application I/O to the secondary volumes. Only volumes

that are currently attached to the host systems that are defined on the Tivoli Storage Productivity Center for Replication Host Systems panel are eligible for Open HyperSwap.

- **Disable Open HyperSwap.** Select this option to prevent a swap from occurring by a command or event while keeping the configuration on the host system and all primary and secondary volumes coupled.

4. Click **OK** to apply the selected options.

Restarting an AIX Host System that is enabled for Open HyperSwap

When an IBM AIX host system is restarted, the host automatically attempts to open any volumes for input/output (I/O) that were open prior to the restart. If Open HyperSwap was enabled for a set of volumes on the host system, the host must determine which storage system is the primary system before the host can allow the volumes to be opened.

If the Metro Mirror relationship for the set of volumes is in a Prepared or Suspended state and the host has connectivity to both the primary and secondary storage systems, the host can determine through the hardware which storage system is the primary system. In this situation, the host automatically opens the volumes.

If the Metro Mirror relationship for the set of volumes is in a Prepared state and the host has connectivity to only the secondary storage system, all I/O to the volumes might be blocked on the host system until the host is able to verify the primary volume in the relationship. The AIX command `varyonvg` will fail to open the volumes for I/O to prevent the application from writing to the incorrect site. If the host can determine which volume is the primary volume in the relationship and connectivity to the primary storage system is still lost, a Hyperswap event is triggered. This event causes all I/O to be automatically opened and directed to the secondary storage system.

If the Metro Mirror relationship for the set of volumes is in a Target Available state after a Hyperswap or a Recover command has been issued for the session, or the host system does not have the connectivity necessary to determine which site is the primary site and all I/O to the volumes are blocked on the host system. The `varyonvg` command will fail to open the volumes for I/O to prevent the application from writing to the incorrect site.

Unblocking I/O on the host system after a host system restart

When any of the previous scenarios cause I/O to be blocked, manual actions might be necessary to remove the block.

If the relationships are in a Target Available state on the hardware, issue a Start command to the session in the desired direction of the relationship. This action defines the primary storage system for the host. The host system can allow the volumes to be opened to I/O.

If the relationships cannot be restarted, or the host cannot determine the primary storage system, it might be necessary to manually decouple the volumes on the host system.

To decouple the volumes the following options are available:

- Option 1: Terminate the session or remove the copy set. This option would require a full copy when the relationships are restarted.
- Option 2: Remove Object Data Manager (ODM) entries using the following command:

```
odmdelete -o Volume_Equivalency
```

CAUTION:

This command should be used only for this scenario because the command deletes copy set information.

Planned and unplanned swaps

Once a session has been enabled for HyperSwap or Open HyperSwap and reaches the Prepared state, IBM Tivoli Storage Productivity Center for Replication loads the configuration of volumes that are capable of being swapped onto IBM z/OS or AIX.

When the load is complete, the session is capable of a planned or unplanned swap. The H1-H2 role pair on the session show a type of HS. An H is displayed over the connection in the dynamic image for that role pair as shown in the following




figure.

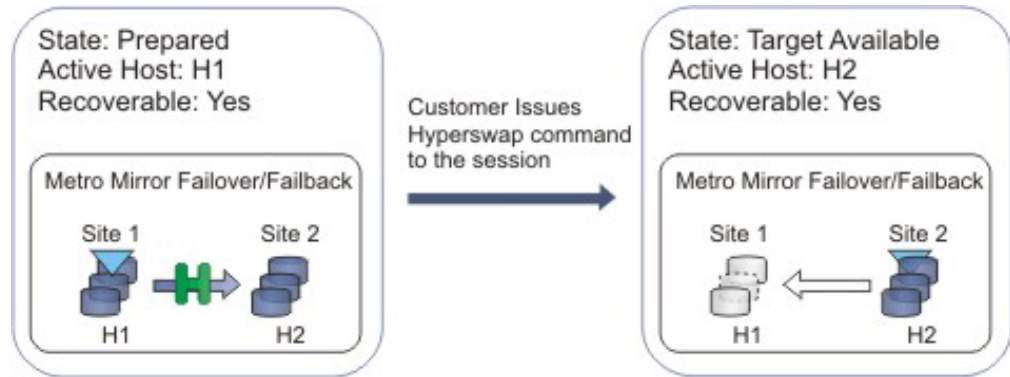
Performing a Planned Swap

Once the session configuration is loaded on z/OS for HyperSwap or AIX for Open HyperSwap, the session is considered swap capable. There may be cases such as a planned maintenance or a migration from the primary storage, in which a planned swap might be required. Once the session is in a swap capable state, a planned swap can be executed by issuing the HyperSwap command against the session.

Once a planned swap is run for z/OS HyperSwap and Open HyperSwap, the session is transitioned to a Target Available state and all the H1-H2 pairs are in a Target Available state. If the H1-H2 role pair was consistent at the time of the swap, the session will have a status of Normal and will indicate that H1-H2 is consistent. If the H1-H2 role pair was not consistent at the time of the swap, the session might display a status of SEVERE because the session is inconsistent. The active host on the session is then displayed as H2.

All input/output (I/O) should have been redirected to the H2 volumes. After a successful swap to site 2, it is not possible to re-enable copy to site 2. Therefore, it is not possible to issue a Start H1->H2 command. The only way to restart the copy is a Start H2->H1 command. To have the volumes protected with high availability and disaster recovery again, the error that caused the swap must be fixed and then the session must be manually restarted to begin copying to the other site.

The following figure illustrates a planned swap.



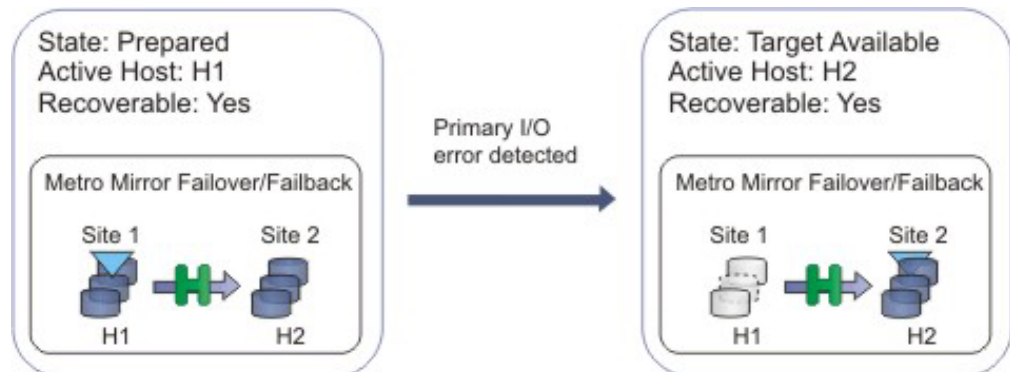
What happens when an unplanned swap occurs

Once the session configuration is loaded on z/OS for HyperSwap or AIX for Open HyperSwap, the session is considered swap capable. In the event of a primary I/O error, a swap occurs automatically. For HyperSwap, z/OS performs the entire swap and then alerts Tivoli Storage Productivity Center for Replication that a swap has occurred. For Open HyperSwap, Tivoli Storage Productivity Center for Replication and the AIX host work together to perform the swap.

Once an unplanned swap occurs for HyperSwap and Open HyperSwap, the session is transitioned to a Target Available state and all the H1-H2 pairs are in a Target Available state. If the H1-H2 role pair was consistent at the time of the swap, the session will have a Status of Normal and will indicate that H1-H2 is consistent. If the H1-H2 role pair was not consistent at the time of the swap, the session might display a status of SEVERE because the session is inconsistent. The active host on the session is then displayed as H2.

All I/O should have been redirected to the H2 volumes. After a successful swap to site 2, it is not possible to re-enable copy to site 2. Therefore, it is not possible to issue a Start H1->H2 command. The only way to restart the copy is a Start H2->H1 command. To have the volumes protected with high availability and disaster recovery again, the error that caused the swap must be fixed and then the session must be manually restarted to begin copying to the other site.

The following figure illustrates an unplanned swap.



Scenarios requiring a full copy in Metro Global Mirror with HyperSwap sessions

In the following cases, a full copy is required to return to the three-site configuration after a swap:

- If you are running a Metro Global Mirror session with HyperSwap and you issue the HyperSwap command using the z/OS HyperSwap API rather than the Tivoli Storage Productivity Center for Replication graphical user interface (GUI).
- If you are running a Metro Global Mirror session with HyperSwap and an unplanned swap occurs.

Verifying that a session is capable of a planned or unplanned swap:

You can verify whether a sessions is capable of a planned or unplanned swap from the IBM z/OS console (HyperSwap) or the IBM AIX host (Open HyperSwap).

Perform these steps to verify the status of HyperSwap from the z/OS console:

1. Issue the `ds hs,status` command for the overall status of the HyperSwap session. For example:

```
15.03.06 SYSTEM1 d hs,status
15.03.06 SYSTEM1 STC00063 IOSHM0303I HyperSwap Status 531
Replication Session: SR_HS
HyperSwap enabled
New member configuration load failed: Disable
Planned swap recovery: Disable
Unplanned swap recovery: Disable
FreezeAll: No
Stop: No
```
2. Issue the `ds hs,config(detail,all)` command to verify all the volumes in the configuration. For example:

```
15.03.51 SYSTEM1 d hs,config(detail,all)
15.03.51 SYSTEM1 STC00063 IOSHM0304I HyperSwap Configuration 534
Replication Session: SR_HS
Prim. SSID UA DEV# VOLSER Sec. SSID UA DEV# Status
06 02 00F42 8K3602 06 04 00FA2
06 01 00F41 8K3601 06 03 00FA1
06 00 00F40 8K3600 06 02 00FA0
```

Perform these steps to verify the status of Open HyperSwap from the AIX host:

1. Issue the `pcmpath query device` command to see the session association and which path the input/output (I/O) is currently being routed to, which is indicated by an asterisk. For example:

```
host1> pcmpath query device 14

DEV#: 14 DEVICE NAME: hdisk14 TYPE: 2107900 ALGORITHM: Load Balance
SESSION NAME: session1
OS Direction: H1<-H2
=====
PRIMARY SERIAL: 25252520000
-----
Path# Adapter/Path Name State Mode Select Errors
0 fscsi0/path0 CLOSE NORMAL 6091 0
1 fscsi0/path2 CLOSE NORMAL 6300 0
2 fscsi1/path4 CLOSE NORMAL 6294 0
3 fscsi1/path5 CLOSE NORMAL 6187 0

SECONDARY SERIAL: 34343430000 *
-----
Path# Adapter/Path Name State Mode Select Errors
```

4	fscsi0/path1	CLOSE	NORMAL	59463	0
5	fscsi0/path3	CLOSE	NORMAL	59250	0
6	fscsi1/path6	CLOSE	NORMAL	59258	0
7	fscsi1/path7	CLOSE	NORMAL	59364	0

Temporarily disabling HyperSwap or Open HyperSwap

In some situations, it might be necessary to temporarily disable the HyperSwap or Open HyperSwap capabilities for a session.

You might want to disable HyperSwap or Open HyperSwap under the following circumstances:

- Performing maintenance
- Inability for one sysplex member to communicate with one or more volumes

Perform these steps to disable Open HyperSwap for a specific session:

1. In the navigation IBM Tivoli Storage Productivity Center for Replication tree, select **Sessions**. The Sessions panel is displayed.
2. Select the sessions for which you want to disable HyperSwap or Open HyperSwap.
3. Select **View/Modify Properties** from the **Select Actions** list, and click **Go**.
4. Select **Disable HyperSwap** or **Disable Open HyperSwap** and click **OK**.

Tip: On management servers that run IBM z/OS, you can also disable HyperSwap from an MVS command prompt by entering **SETHS DISABLE**.

Using active and standby Tivoli Storage Productivity Center for Replication servers with HyperSwap or Open HyperSwap

To ensure that there is an IBM Tivoli Storage Productivity Center for Replication server available in the event of a disaster, an active and standby management server configuration can be set up. You can enable HyperSwap and Open HyperSwap for sessions while maintaining an active and standby server configuration.

Active and standby servers with HyperSwap

When the storage system is set up to connect through the z/OS interface, the connection information is automatically sent to the standby server and a connection is attempted. It is possible that the connection will fail if the standby server is not running on z/OS or does not have access to the same volumes. If the connection fails, any takeover done on the standby server will not be able to manage the HyperSwap. On z/OS, if the session configuration had been successfully loaded prior to the HyperSwap, the z/OS system is still capable of performing the HyperSwap. If the z/OS system swaps the volumes but there is no communication possible to the Tivoli Storage Productivity Center for Replication server, the session recognizes that the pairs became suspended and the session will go into a Suspended/Severe state. From this state the customer can clear the **Manage H1-H2 with Hyperswap** option and issue the Recover command for the session to get the session into a Target Available state.

Active and standby servers with Open HyperSwap

When there is an active and standby management server configuration and a host system connection is added to the active server, the host system connection is automatically sent to the standby server and a connection is attempted. Once the session configuration is loaded on IBM AIX, Open HyperSwap is possible only if

there is continual communication between AIX and the Tivoli Storage Productivity Center for Replication server. If a takeover is performed on a standby server that is unable to connect to the host system that is managing the swap, the session is no longer Open HyperSwap capable. Communication to the host system must be activated before the session can become Open HyperSwap capable again.

Related tasks:

Chapter 3, “Managing management servers,” on page 85

This section provides information about how to set up active and standby management servers, restore a lost connection between the management servers, or perform a takeover on the standby management server.

Session commands

The commands that are available for a session depend on the session type.

Commands are issued synchronously to IBM Tivoli Storage Productivity Center for Replication sessions . Any subsequent command issued to an individual session is not processed until the first command has completed. Some commands can take an extended amount of time to complete, such as the **Start** command as it sets up the hardware. The GUI continues to allow you to issue commands to other sessions and not hold up functionality. When a command has completed, the console displays the results of the command.

Basic HyperSwap commands

Use this information to learn about commands available for Basic HyperSwap sessions.

Note: Individual suspend and recover commands are not available in HyperSwap.

Table 23. Basic HyperSwap commands

Command	Action
HyperSwap	Triggers a HyperSwap where I/O is redirected from the source volume to the target volume, without affecting the application using those volumes. You can use this command if you want to perform maintenance on the original source volumes.
Start H1->H2	Starts copying data synchronously from H1 to H2 in a Metro Mirror session. Note: A session might go into a Severe state with error code 1000000 before the session returns to Normal/Prepared State and HyperSwap Capable. The duration of the Severe state depends on how large of a session is running.
Start H2->H1	Starts copying data synchronously from H2 to H1 in a Metro Mirror session. You can issue this command only after the session has been swapped and the production site is H2. To enable data protection when the H1 volumes are available again, start I/O to the H2 volumes, and issue this command to replicate data from the H2 volumes to H1 volumes.
Stop	Suspends updates to all the targets of pairs in a session. You can issue this command at any time during an active session. Note: After you issue the stop command, targets might not be consistent.
Terminate	Removes all physical copies and relationships from the hardware during an active session.

FlashCopy commands

Use this information to learn about commands available for FlashCopy sessions.

Table 24. FlashCopy commands

Command	Action
Start	<p>Performs any steps necessary to define the relationship before performing a FlashCopy operation. For ESS, DS6000, and DS8000, this command is not an option. Issue this command to put the session in the prepared state for the following storage systems:</p> <ul style="list-style-type: none"> • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified
Flash	<p>Performs the FlashCopy operation using the specified options. Issue the Flash command to create a data consistent point-in-time copy for a FlashCopy Session with volumes on the following storage systems:</p> <ul style="list-style-type: none"> • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified <p>For a FlashCopy session containing ESS, DS6000, and DS8000 volumes, the Flash command by itself is not sufficient to create a consistent copy. To create a consistent copy using the ESS, DS6000, and DS8000 Flash commands, you must quiesce application I/O before issuing the Flash command.</p>
InitiateBackgroundCopy	<p>Copies all tracks from the source to the target immediately, instead of waiting until the source track is written to. This command is valid <i>only</i> when the background copy is not already running.</p>
Terminate	<p>Removes all active physical copies and relationships from the hardware during an active session.</p> <p>If you want the targets to be data consistent before removing their relationship, you must issue the InitiateBackgroundCopy command if NOCOPY was specified, and then wait for the background copy to complete by checking the copying status of the pairs.</p>

Snapshot commands

Use this information to learn about commands that are available for Snapshot sessions and groups. A snapshot group is a grouping of snapshots of individual volumes in a consistency group at a specific point in time.

Table 25. Snapshot session commands

Command	Action
Create Snapshot	Creates a snapshot of the volumes in the session
Restore	Restores the H1 volumes in the session from a set of snapshot volumes. You must have at least one snapshot group to restore from. When you issue this command in the Tivoli Storage Productivity Center for Replication graphical user interface (GUI), you are prompted to select the snapshot group.

Table 26. Snapshot group commands

Command	Action
Delete	Deletes the snapshot group and all the individual snapshots that are in the group from the session and from the XIV system. If the deleted snapshot group is the last snapshot group that is associated with the session, the session returns to the Defined state.
Disband	Disbands the snapshot group. When a snapshot group is disbanded, the snapshot group no longer exists. All snapshots in the snapshot group become individual snapshots that are no longer associated to the consistency group or the session. After a snapshot group is disbanded, it is no longer displayed in or managed by Tivoli Storage Productivity Center for Replication. If the disbanded snapshot group is the last snapshot group that is associated with the session, the session returns to the Defined state.
Duplicate	Duplicates the snapshot group. When a snapshot group is duplicated, a new snapshot group is created with new snapshots for all volumes that are in the duplicated group. The name of the duplicated snapshot group is generated automatically by the XIV system.
Lock	Locks a snapshot group. If the snapshot group is locked, write operations to the snapshots that are in the snapshot group are prevented. By default, a snapshot group is locked when it is created. This action is valid only if the snapshot group is unlocked.
Overwrite	Overwrites the snapshot group to reflect the data that is on the H1 volume.
Rename	Renames the snapshot group to a name that you provide. The name can be a maximum of 64 alphanumeric characters.
Restore	Restores the contents of a snapshot group by using another snapshot group in the session. Both of the snapshot groups must contain the same subset of volumes.
Set Priority	Sets the priority in which a snapshot group is deleted. The value can be the number 1 - 4. A value of 1 specifies that the snapshot group is deleted last. A value of 4 specifies that the snapshot group is deleted first.
Unlock	Unlocks a snapshot group. If the snapshot group is unlocked, write operations to the snapshots that are in the snapshot group are enabled and the snapshot group is displayed as modified. This action is valid only if the snapshot group is locked.

Metro Mirror commands

Use this information to learn about commands available for Metro Mirror sessions.

Table 27. Metro Mirror commands

Command	Action
Enable Copy to Site 1	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the Start H2->H1 command becomes available.

Table 27. Metro Mirror commands (continued)

Command	Action
Enable Copy to Site 2	Run this command and confirm that you want to reverse the direction of replication before reversing the direction of copying in a failover and failback session. After you issue this command, the Start H1->H2 command becomes available.
HyperSwap	Triggers a HyperSwap where I/O is redirected from the source volume to the target volume, without affecting the application using those volumes. You can use this command if you want to perform maintenance on the original source volumes.
Start	Establishes a single-direction session with the hardware and begins the synchronization process between the source and target volumes.
Start H1->H2	Establishes Metro Mirror relationships between the H1 volumes and the H2 volumes, and begins data replication from H1 to H2.
Start H2->H1	Establishes Metro Mirror relationships between the H2 volumes and the H1 volumes and starts data replication from H2 to H1. Indicates direction of a failover and failback between two hosts in a Metro Mirror session. If the session has been recovered such that the production site is now H2, you can issue the Start H2->H1 command to start production on H2 and start data replication.
Stop	Inconsistently suspends updates to all the targets of pairs in a session. This command can be issued at any point during an active session. Note: Targets after the suspend are not considered to be consistent.
StartGC	Establishes Global Copy relationships between the H1 volumes and the H2 volumes, and begins asynchronous data replication from H1 to H2. While in the Preparing state, it will not change to the Prepared state unless you switch to Metro Mirror.
Suspend	Causes all target volumes to remain at a data-consistent point and stops all data that is moving to the target volumes. This command can be issued at any point during a session when the data is actively being copied. Note: It is recommended that you avoid using the same LSS pairs for multiple Metro Mirror sessions. Metro Mirror uses a freeze command on ESS, DS6000, and DS8000 storage systems to create the data-consistent point. If there are other Metro Mirror sessions overlapping the same LSS pairs as in this session, those sessions are also suspended. When a Suspend command is issued to a source volume in an LSS that has source volumes in another active Metro Mirror session, the other source volumes are affected only if they have the same target LSS. The primary volumes are suspended, but volumes in the same source LSS that have target volumes in a different LSS are not affected because they use a different PPRC path connection.
Recover	Issues the Recover command to suspended sessions. This command performs the steps necessary to make the target available as the new primary site. Upon completion of this command, the session becomes Target Available.

Table 27. Metro Mirror commands (continued)

Command	Action
Terminate	Removes all copy relationships from the hardware during an active session. If you want the targets to be data consistent before removing their relationship, you must issue the Suspend command, then the Recover command, and then the Terminate command.

Metro Mirror with Practice commands

Use this information to learn about commands available for Metro Mirror with Practice sessions.

Table 28. Metro Mirror with Practice commands

Command	Action
Enable Copy to Site 1	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the Start H2->H1 command becomes available.
Enable Copy to Site 2	Run this command and confirm that you want to reverse the direction of replication before reversing the direction of copying in a failover and failback session. After you issue this command, the Start H1->H2 command becomes available.
Flash	Creates a FlashCopy image from I2 volumes to H2 volumes. The amount of time for this to occur will vary depending on the number of copy sets in the session. Note: For ESS, DS6000, DS8000 storage systems, the Flash command uses the freeze and thaw processing to create a data consistent point for the FlashCopy. If there is another Metro Mirror session overlapping on one or more of the same LSS pairs, that session will be suspended. It is also possible that the suspension of the other session might cause the Metro Mirror session to remain suspended after the flash command is issued instead of returning to Prepared state. Avoid using the same LSS pairs for multiple Metro Mirror sessions if possible.
Start H1->H2	Establishes a Metro Mirror relationship from the H1 volumes to the I2 volumes, and begins data replication.
Start H2->H1	Establishes a Metro Mirror relationship from H2 to H1 and begins data replication.
StartGC_H1H2	Distinguishes when the session is in the Preparing state from H1 to I2 and begins the asynchronous process between the source and target volumes. While in the Preparing state the session will not change to the Prepared state unless you switch to Metro Mirror.
StartGC_H2H1	Distinguishes when the session is in the Preparing state from H2 to H1 and begins the asynchronous process between the source and target volumes. While in the Preparing state the session will not change to the Prepared state unless you switch to Metro Mirror.

Table 28. Metro Mirror with Practice commands (continued)

Command	Action
Suspend	<p>Causes all target volumes to remain at a data-consistent point and stops all data that is moving to the target volumes. This command can be issued at any point during a session when the data is actively being copied.</p> <p>Note: The Metro Mirror command uses a freeze command on the ESS, DS6000, or DS8000 devices to create the data-consistent point. If there are other Metro Mirror sessions overlapping the same LSS pairs as in this session, those sessions will also become suspended. Avoid using the same LSS pairs for multiple Metro Mirror sessions.</p> <p>When a Suspend command is issued to a source volume in an LSS that has source volumes in another active Metro Mirror session, the other source volumes are affected only if they have the same target LSS. The primary volumes are suspended, but volumes in the same source LSS that have target volumes in a different LSS are not affected because they use a different PPRC path connection.</p>
Stop	<p>Inconsistently suspends updates to all the targets of pairs in a session. This command can be issued at any point during an active session.</p> <p>Note: Targets after the suspend are not considered to be consistent.</p>
Terminate	Terminates all copy relationships on the hardware.
Recover	<p>Takes a point-in-time copy of the data on I2 to the H2 volumes, enabling the application to be attached and run from the H2 volumes on site 2.</p> <p>Note: The point-in-time copy is performed when the session is in a recoverable state to avoid that previous consistent data on H2 are overwritten by inconsistent data upon Recover. You can issue the Flash command if you want to force a point-in-time copy from I2 to JH2 volumes afterwards.</p>

Global Mirror commands

Use this information to learn about commands available for Global Mirror sessions.

Table 29. Global Mirror commands

Command	Action
Enable Copy to Site 1	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the Start H2->H1 command becomes available.
Enable Copy to Site 2	Run this command and confirm that you want to reverse the direction of replication before reversing the direction of copying in a failover and failback session. After you issue this command, the Start H1->H2 command becomes available.
Start	Establishes all relationships in a single-direction session and begins the process necessary to start forming consistency groups on the hardware.

Table 29. Global Mirror commands (continued)

Command	Action
Start H1->H2	Starts copying data from H1 to H2 in a Global Mirror failover and failback session. Establishes the necessary relationships in the session and begins the process necessary to start copying data from the H1 site to the H2 site and to start forming consistency groups.
Start H2->H1	Starts copying data from H2 to H1 in a failover and failback session for ESS, DS6000 and DS8000 sessions. If a recover has been performed on a session such that the production site is now H2, you can issue a Start H2->H1 to start moving data back to Site 1. However, this start does not provide consistent protection as it copies only asynchronously back because of the long distance. A Global Copy relationship is used. When you are ready to move production back to Site 1, issue a suspend to the session; this puts the relationships into a synchronized state and suspends them consistently. Sessions are consistent when copying H2->H1 for the following storage systems: <ul style="list-style-type: none"> • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified
Suspend	Stops all consistency group formation when the data is actively being copied and suspends the H1->H2 Global Copy pairs. To issue the pause command to the Global Mirror session on the hardware without suspending the Global Copy pairs, open the <code>rmserver.properties</code> file and add the following property to disable the Global Copy suspension on the Suspend command: <code>csm.server.sus_gc_pairs_on_gm_pause = false</code> . The default property is true and the Suspend command automatically suspends the Global Copy pairs.
Recover	Issue this command to recover the session to the target site. This command performs the steps necessary to make the target host volumes consistent and available for access as the new primary site. Upon completion of this command, the session becomes Target Available. Do not access H2 volumes until the Recover command is completed and the session displays Target Available and Recoverable. A Recover to H2 also establishes a point-in-time copy to J2 to preserve the last consistent data.
Terminate	Removes all physical copies and relationships from the hardware during an active session. If you want the targets to be data consistent before removing their relationship, you must issue the Suspend command, the Recover command, and then the Terminate command.

Global Mirror with Practice commands

Use this information to learn about commands available for Global Mirror with Practice sessions.

Table 30. Global Mirror with Practice commands

Command	Action
Enable Copy to Site 1	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the Start H2->H1 command becomes available.

Table 30. Global Mirror with Practice commands (continued)

Command	Action
Enable Copy to Site 2	Run this command and confirm that you want to reverse the direction of replication before reversing the direction of copying in a failover and failback session. After you issue this command, the Start H1->H2 command becomes available.
Flash	The Flash command on a Global Mirror Practice session for ESS, DS6000, and DS8000 temporarily pauses the formation of consistency groups, ensure that all I2s are consistent, and then flash the data from I2 to the H2 volumes. After the flash is complete, the Global Mirror session is automatically restarted, and the session begins forming consistency groups on I2. You can then use the H2 volumes to practice your disaster recovery procedures.
Start H1->H2	Starts copying data from H1 to H2. After the first pass of the copy is complete for all pairs, the session establishes the I2->J2 FlashCopy pairs, and starts the Global Mirror master so that the hardware will begin forming consistency groups, to ensure consistent data is at site 2.
Start H2->H1	Starts copying data from H2 to H1 in a failover and failback session. If a recover has been performed on a session such that the production site is now H2, you can issue a Start H2->H1 to start moving data back to Site 1. However, this start does not provide consistent protection as it copies only asynchronously back because of the long distance. Note: ESS, DS6000, and DS8000 volumes are not consistent for the Start H2->H1 command. A Global Copy relationship is used. When you are ready to move production back to Site 1, issue a suspend to the session; this puts the relationships into a synchronized state and suspends them consistently.
Terminate	Removes all physical copies and relationships on the hardware.
Suspend	Stops all consistency group formation when the data is actively being copied and suspends the H1-I2 Global Copy pairs. To issue the pause command to the Global Mirror session on the hardware without suspending the Global Copy pairs, open the <code>rmserver.properties</code> file and add the following property to disable the Global Copy suspension on the Suspend command: <code>csm.server.sus_gc_pairs_on_gm_pause = false</code> . The default property is true and the Suspend command automatically suspends the Global Copy pairs..
Recover	Restores consistent data on I2 volumes and takes a point-in-time copy of the data on I2 to the H2 volumes, enabling the application to be attached and run from the H2 volumes on site 2. The I2 Volumes will continue to hold the consistent data and can be flashed again to H2 using the Flash command.

Metro Global Mirror commands

Use this information to learn about commands available for Metro Global Mirror sessions.

Table 31. Metro Global Mirror commands

Command	Action
Enable Copy to Site 1	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the Start H2->H1->H3 command becomes available.
Enable Copy to Site 2	Run this command and confirm that you want to reverse the direction of replication before reversing the direction of copying in a failover and failback session. After you issue this command, the Start H1->H2->H3 command becomes available.
HyperSwap	<p>Causes a site switch, equivalent to a suspend and recover for a Metro Mirror with failover and failback individual suspend and recover commands are not available.</p> <p>The Metro Global Mirror for HyperSwap supports IBM Tivoli Storage Productivity Center for Replication installed on z/OS with the IBM Tivoli Storage Productivity Center for Replication Three Site Business Continuity license.</p>
Start H1->H2->H3	<p>(This is the Metro Global Mirror initial start command.) Establishes Metro Mirror relationships between H1 and H2, and Global Mirror relationships between H2 and H3. For Metro Global Mirror, this includes the J3 volume to complete the Global Mirror configuration. (The J3 volume role is the journal volume at site 3.). Start H1->H2->H3 can be used from some Metro Global Mirror configurations to transition back to the starting H1->H2->H3 configuration.</p> <p>This command is valid only when the session is in a defined, preparing, prepared, or suspended state.</p>
Start H1->H3	<p>From the H1->H2->H3 configuration, this command changes the session configuration to a Global Mirror-only session between H1 and H3, with H1 as the source. Use this command in case of an H2 failure with transition bitmap support provided by incremental resynchronization. It can be used when a session is in preparing, prepared, and suspended states because there is not a source host change involved.</p> <p>This command allows you to bypass the H2 volume in case of an H2 failure and copy only the changed tracks and tracks in flight from H1 to H3. After the incremental resynchronization is performed, the session is running Global Mirror from H1 to H3 and thus loses the near-zero data loss protection achieved with Metro Mirror when running H1->H2->H3. However, data consistency is still maintained at the remote site with the Global Mirror solution.</p> <p>From H2->H1->H3 configuration, this command changes the session configuration to a Global Mirror-only session configuration between H1 and H3, with H1 as the source. Use this command when the source site has a failure and production is moved to the H1 site. This can be done for unplanned HyperSwap. The Global Mirror session is continued. This is a host-volume change so this command is valid only when restarting the H1->H3 configuration or from the TargetAvailable H2->H1->H3 state.</p>

Table 31. Metro Global Mirror commands (continued)

Command	Action
Start H2->H3	<p>From the H1->H2->H3 configuration, this command moves the session configuration to a configuration between H2 and H3, with H2 as the source. Use this command when the source site has a failure and production is moved to the H2 site, for example, for unplanned HyperSwap. The Global Mirror session is continued. This session is a host-volume change so this command is valid only when restarting the H1->H3 configuration or from the TargetAvailable H2->H1>H3 state.</p> <p>From the H2->H1->H3 configuration, this command changes the session configuration to a configuration between H2 and H3 with H2 as the source. Use this command in case of an H1 failure with transition bitmap support provided by incremental resynchronization. Because there is not a source-host change involved, it can be used when the session is in preparing, prepared, and suspended states Start H2->H1->H3 can be used to transition back to the starting H2->H1->H3 configuration.</p>
Start H2->H1->H3	<p>(This is the Metro Global Mirror start command.) This is the configuration that completes the HyperSwap processing. This command creates Metro Mirror relationships between H2 and H1 and Global Mirror relationships between H1 and H3. For Metro Global Mirror, this includes the J3 volume to complete the Global Mirror configuration.</p>
Start H3->H1->H2	<p>After recovering to H3, this command sets up the hardware to allow the application to begin writing to H3, and the data is copied back to H1 and H2. However, issuing this command does not guarantee consistency in the case of a disaster because only Global Copy relationships are established to cover the long distance copy back to site 1.</p> <p>To move the application back to H1, you can issue a suspend while in this state to drive all the relationships to a consistent state and then issue a freeze to make the session consistent. You can then issue a Rcover followed by a Start H1->H2->H3 to go back to the original configuration.</p>
SuspendH2H3	<p>When running H1->H2->H3, this command issues a pause to the Global Mirror master and causes the Global Mirror master to stop forming consistency groups.</p> <p>This command is valid only when the session is in a prepared state. Stops all consistency group formation when the data is actively being copied and suspends the H2>H3 Global Copy pairs. To issue the pause command to the Global Mirror session on the hardware without suspending the Global Copy pairs, open the <code>rmserver.properties</code> file and add the following property to disable the Global Copy suspension on the Suspend command: csm.server.sus_gc_pairs_on_gm_pause = false. The default property is true and the Suspend command automatically suspends the Global Copy pairs..</p>

Table 31. Metro Global Mirror commands (continued)

Command	Action
SuspendH1H3	<p>When running H2->H1->H3, this command issues a pause to the Global Mirror master and causes the Global Mirror master to stop forming consistency groups.</p> <p>This command is valid only when the session is in a prepared state. Stops all consistency group formation when the data is actively being copied and suspends the H2->H3 Global Copy pairs. To issue the pause command to the Global Mirror session on the hardware without suspending the Global Copy pairs, open the <code>rmsserver.properties</code> file and add the following property to disable the Global Copy suspension on the Suspend command: <code>csm.server.sus_gc_pairs_on_gm_pause = false</code>. The default property is true and the Suspend command automatically suspends the Global Copy pairs..</p>
RecoverH1	<p>Specifying H1 makes the H1 volume TargetAvailable. Metro Global Mirror (when running H2->H1->H3) can move production to either the H1 or H3 set of volumes. IBM Tivoli Storage Productivity Center for Replication processing differs, depending on the recovery site. Therefore, the site designation is added to the Recover command so IBM Tivoli Storage Productivity Center for Replication can set up for the failback.</p>
RecoverH2	<p>Specifying H2 makes the H2 volume TargetAvailable. Metro Global Mirror (when running H1->H2->H3) can move production to either the H2 or H3 set of volumes. IBM Tivoli Storage Productivity Center for Replication processing is different depending on the recovery site. Therefore the site designation is added to the Recover command so IBM Tivoli Storage Productivity Center for Replication can prepare for the failback.</p>
RecoverH3	<p>Specifying H3 makes the H3 volume TargetAvailable. Metro Global Mirror (when running H1->H2->H3) can then move production to the H3 set of volumes. Because IBM Tivoli Storage Productivity Center for Replication processing differs depending on the recovery site, the site designation is added to the Recover command so that IBM Tivoli Storage Productivity Center for Replication can prepare for the failback.</p> <p>This command prepares H3 so that you can start the application on H3. H3 becomes the active host, and you then have the option <code>start H3->H1->H2</code> to perform a Global Copy copy back. The recovery establishes point-in-time copy to J3 volumes to preserve the last consistent data.</p>

Metro Global Mirror with Practice commands

Use this information to learn about commands available for Metro Global Mirror with Practice sessions.

Table 32. Metro Global Mirror with Practice commands

Command	Action
Enable Copy to Site 1	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session.
Enable Copy to Site 2	Run this command and confirm that you want to reverse the direction of replication before reversing the direction of copying in a failover and failback session.

Table 32. Metro Global Mirror with Practice commands (continued)

Command	Action
Flash	<p>This command is available in the following states:</p> <ul style="list-style-type: none"> • Target Available state when the active host is H3 Note: Use this command if the FlashCopy portion of the Recover command from I3 to H3, fails for any reason. The problem can be addressed; and a Flash command issued to complete the flash of the consistent data from I3 to H3. • Prepared state when the active host is H1 and data is copying H1 to H2 to I3, or the active host is H2 and data is copying H2 to H1 to H3. • Prepared state when the active host is H2 and data is copying H2 to I3. • Prepared state when the active host is H1 and data is copying H1 to I3. <p>Use this command if the FlashCopy portion of the Recover command from I3 to H3, fails for any reason. The problem can be addressed, and a Flash command issued to complete the flash of the consistent data from I3 to H3.</p> <p>Issuing a Flash command on a Global Mirror Practice session for ESS, DS6000, and DS8000 will temporarily pause the formation of consistency groups, ensure that all I3s are consistent, and then flash the data from I3 to the H3 volumes. After the flash is complete, the Global Mirror session will be automatically restarted, and the session will begin forming consistency groups on I3. You can then use the H3 volumes to practice your disaster recovery procedures.</p>
RecoverH1	<p>Specifying H1 makes the H1 volume TargetAvailable. Metro Global Mirror (when running H1->H2->H3) can move production to either the H2 or H3 set of volumes. IBM Tivoli Storage Productivity Center for Replication processing differs, depending on the recovery site. Therefore the site designation is added to the Recover command so IBM Tivoli Storage Productivity Center for Replication can prepare for the failback. The FlashCopy creates a consistent copy of the data on the H3 volumes so that an application can recover to those volumes and begin writing I/O. When the FlashCopy is complete, the session will reach a Target Available state, and you can attach your volumes on Site 3.</p>
RecoverH2	<p>Specifying H2 makes the H2 volume TargetAvailable. Metro Global Mirror (when running H1->H2->H3) can move production to either the H2 or H3 set of volumes. IBM Tivoli Storage Productivity Center for Replication processing differs, depending on the recovery site. Therefore the site designation is added to the Recover command so IBM Tivoli Storage Productivity Center for Replication can prepare for the failback.</p>
RecoverH3	<p>Specifying H3 makes the H3 volume the TargetAvailable. When running H1->H2->H3, Metro Global Mirror can move production to either the H2 or H3 set of volumes. IBM Tivoli Storage Productivity Center for Replication processing differs, depending on the recovery site; therefore, the site designation is added to the Recover command so IBM Tivoli Storage Productivity Center for Replication can prepare for the failback.</p>
Re-enable Copy to Site 1	<p>After issuing a RecoverH1 command, you can run this command to restart the copy to the original the direction of replication in a failover and failback session.</p>

Table 32. Metro Global Mirror with Practice commands (continued)

Command	Action
Re-enable Copy to Site 2	After issuing a RecoverH2 command, you can run this command to restart the copy to the original the direction of replication in a failover and failback session.
Re-enable Copy to Site 3	After issuing a RecoverH3 command, you can run this command to restart the copy to the original the direction of replication in a failover and failback session.
Start H1->H2->H3	<p>Metro Global Mirror initial start command. This command creates Metro Mirror relationships between H1 and H2, and Global Mirror relationships between H2 and H3. For Metro Global Mirror, this includes the J3 volume to complete the Global Mirror configuration. (The J3 volume role is the journal volume at site 3.). Start H1->H2->H3 can be used from some Metro Global Mirror configurations to return to the starting H1>H2>H3 configuration.</p> <p>This command is valid only when the session is in a defined, preparing, prepared, target-available, or suspended state.</p>
Start H1->H3	<p>From the H1->H2->H3 configuration, this command changes the session configuration to a Global-Mirror-only session between H1 and H3, with H1 as the source. Use this command in case of an H2 failure with transition bitmap support provided by incremental resynchronization. Because there is not a source host change involved, it can be used when a session is in preparing, prepared, and suspended states.</p> <p>You can use this command to bypass the H2 volume in case of an H2 failure and copy only the changed tracks and tracks in flight from H1 to H3. After the incremental resynchronization is performed, the session is running Global Mirror from H1 to H3 and thus loses the near-zero data loss protection achieved with Metro Mirror when running H1->H2->H3. However, data consistency is still maintained at the remote site with the Global Mirror solution.</p> <p>From H2->H1->H3 configuration, this command changes the session configuration to a Global-Mirror-only session configuration between H1 and H3, with H1 as the source. Use this command when the source site has a failure and production is moved to the H1 site. This can be done for unplanned HyperSwap. The Global Mirror session is continued. This is a host-volume change so this command is valid only when restarting the H1->H3 configuration or from the TargetAvailable H2->H1->H3 state.</p>

Table 32. Metro Global Mirror with Practice commands (continued)

Command	Action
Start H2->H3	<p>From the H1->H2->H3 configuration, this command moves the session configuration to a configuration between H2 and H3, with H2 as the source. Use this command when the source site has a failure and production is moved to the H2 site. This can be done for unplanned HyperSwap. The Global Mirror session is continued. This is a host-volume change so this command is valid only when restarting the H1>H3 configuration or from the TargetAvailable H2->H1->H3 state.</p> <p>From the H2->H1->H3 configuration, this command changes the session configuration to a configuration between H2 and H3 with H2 as the source. Use this command in case of an H1 failure with transition bitmap support provided by incremental resynchronization. Because there is not a source-host change involved it can be used when the session is in preparing, prepared, and suspended states. Start H2->H1->H3 can be used to return to the starting H2->H1->H3 configuration.</p>
SuspendH2H3	<p>When running H1->H2->H3, this command issues a pause to the Global Mirror master and causes the Global Mirror master to stop forming consistency groups.</p> <p>This command is valid only when the session is in a prepared state. It stops all consistency group formation when the data is actively being copied and suspends the H2->H3 Global Copy pairs. To issue the pause command to the Global Mirror session on the hardware without suspending the Global Copy pairs, open the <code>rmserver.properties</code> file and add the following property to disable the Global Copy suspension on the Suspend command: <code>csm.server.sus_gc_pairs_on_gm_pause = false</code>. The default property is true and the Suspend command automatically suspends the Global Copy pairs..</p>
SuspendH1H3	<p>When running H2->H1->H3, this command issues a pause to the Global Mirror master and causes the Global Mirror master to stop forming consistency groups.</p> <p>This command is valid only when the session is in a prepared state. Stops all consistency group formation when the data is actively being copied and suspends the H2->H3 Global Copy pairs. To issue the pause command to the Global Mirror session on the hardware without suspending the Global Copy pairs, open the <code>rmserver.properties</code> file and add the following property to disable the Global Copy suspension on the Suspend command: <code>csm.server.sus_gc_pairs_on_gm_pause = false</code>. The default property is true and the Suspend command automatically suspends the Global Copy pairs..</p>
Terminate	This command terminates all copy relationships on the hardware.

Site awareness

You can associate a location with each storage system and each site in a session. This *site awareness* ensures that only the volumes whose location matches the location of the site are allowed for selection when you add copy sets to the session. This prevents a session relationship from being established in the wrong direction.

Note: To filter the locations for site awareness, you must first assign a site location to each storage system.

IBM Tivoli Storage Productivity Center for Replication does not perform automatic discovery of locations. Locations are user-defined and specified manually.

You can change the location associated with a storage system that has been added to the IBM Tivoli Storage Productivity Center for Replication configuration. You can choose an existing location or add a new one. Locations are deleted when there is no longer a storage system with an association to that location.

When adding a copy set to a session, a list of candidate storage systems is presented, organized by location. Storage systems that do not have a location are displayed and available for use when you create a copy set.

You can also change the location for any site in a session. Changing the location of a session does not affect the location of the storage systems that are in the session.

Changing the location of a storage system might have consequences. When a session has a volume role with a location that is linked to the location of the storage system, changing the location of the storage system could change the session's volume role location. For example, if there is one storage system with the location of A_Location and a session with the location of A_Location for its H1 role, changing the location of the storage system to a different location, such as B_Location, also changes the session's H1 location to Site 1. However, if there is a second storage system that has the location of A_Location, the session's role location is not changed.

Important: Location matching is enabled only when adding copy sets. If you change the location of a storage system or volume role, IBM Tivoli Storage Productivity Center for Replication does not audit existing copy sets to confirm or deny location mismatches.

Preserve Mirror option

This topic presents recommendations for using the Preserve Mirror option in FlashCopy and Metro Mirror sessions.

When the source of the FlashCopy relationship is a source of a Metro Mirror relationship, and the target of the FlashCopy relationship is the source of a Metro Mirror relationship, the Preserve Mirror option attempts to preserve the consistency of the Metro Mirror relationship at the target of the FlashCopy relationship, preventing a full copy from being performed over the Metro Mirror link. Instead, parallel flashes will be performed (if possible) on both sites. If the consistency cannot be preserved, the Flash for the FlashCopy relationships will fail, and the data of the Metro Mirror relationship at the target of the FlashCopy relationship will not be changed.

Note: This option is available only on DS8000 storage devices with the required code levels installed.

However, in some instances, the Preserve Mirror option can cause a Metro Mirror session to go into a Preparing state, or even a Suspended state. This topic describes the recommended usage of the Preserve Mirror feature. Using this feature in other ways might lead to a Metro Mirror session going into a Preparing or Suspended state.

FlashCopy session

You can use the Preserve Mirror option in FlashCopy sessions in two different methods:

Perform an incremental resynchronization

To perform an incremental resynchronization, select the **Incremental** and **Persistent** options in the FlashCopy session: do *not* select the **No Copy** option.

Perform a single full copy

To perform a single full copy, ensure that the **Incremental**, **Persistent** and **No Copy** options are *not* selected before you issue a **Flash** command. If you use the **No Copy** option, issue either an **Initiate Background Copy** command or **Terminate** command before you issue the Flash command.

Refer to your DS8000 documentation for more information about the Preserve Mirror function.

Metro Mirror session

You can set up your Metro Mirror pairs in two different ways, depending on the level of consistency you need, and your preferences.

Note: For the examples in this section, the source pair is H1a->H2a and the target pair is H1b->H2b. The source pair will contain volumes that will be the source of the FlashCopy relationship and the target pair will always contain volumes that will be the target of the FlashCopy relationship.

Create one Metro Mirror session, and add the Metro Mirror pairs as copy sets to that session

The benefit to this approach is that you do not need to worry about whether the host considers the H1a->H2a and H1b->H2b volumes to be consistent with one another. IBM Tivoli Storage Productivity Center for Replication will ensure that all of the volumes remain consistent.

A drawback to this approach is that when using the **Attempt to preserve Metro Mirror consistency, but fail FlashCopy if Metro Mirror target consistency cannot be preserved** option (Preserve Mirror Required), there is a chance that the target pair (H1b->H2b) might suspend unexpectedly: this causes all other pairs in the Metro Mirror session to suspend (including H1a->H2a). This can occur when a FlashCopy establish or withdraw fails unexpectedly on the remote (H1b->H2b) site. If the host requires the H1a->H2a and H1b->H2b volumes to be consistent, then you should suspend all other volumes.

Create one Metro Mirror session for the H1a->H2a volumes, and another Metro Mirror session for the H1b->H2b volumes

Use this option when the hosts and applications do not require the H1a->H2a and H1b->H2b volumes to be consistent with one another. In this case, you should create one Metro Mirror session for all of the H1a->H2a volumes, and another Metro Mirror session for the H1b->H2b volumes. The H1a->H2a pair is added to the first session, while the H1b->H2b pair is added to the second Metro Mirror session. As long as the host does not require consistency between the H1a and H1b volumes, this option benefits you when you use the **Attempt to preserve Metro Mirror consistency, but fail FlashCopy if Metro Mirror target consistency cannot be preserved** option (Preserve Mirror Required). The benefit is that if one

pair is suspended (such as H1a->H2a), the pairs in the other session will not be affected, since it is in a different Metro Mirror session. Using this method, you can avoid the situation in which a critical application is writing to the source pair (H1a->H2a), while a batch job is writing to the target pair (H1b->H2b), and both pairs are in the same IBM Tivoli Storage Productivity Center for Replication session. These factors cause both applications to receive extended long busy signals, instead of just the batch job.

Creating sessions and adding copy sets

This section describes how to create a session for a specific replication method and then add copy sets to that session.

Creating a FlashCopy session and adding copy sets

FlashCopy replication creates a point-in-time copy in which the target volume contains a copy of the data that was on the source volume when the FlashCopy session was established.

When you create a FlashCopy session for Global Mirror or Metro Global Mirror space-efficient target volumes, you must select the **No Copy** option for the FlashCopy session. With space-efficient volumes, you can use your FlashCopy repository more efficiently. Instead of requiring an equal amount of space to write data to, you can set aside a smaller amount of space in which to write data, where only the tracks that are changed are recorded. When your pool of storage is full, you can no longer perform a FlashCopy operation, and your session goes into a Severe state.

1. Follow these steps to create a FlashCopy session:
 - a. In the navigation tree, select **Sessions**.
 - b. Click **Create Session**.
 - c. In the Create Session wizard, for the **Choose Hardware Type** list, select the item that shows the type of storage system for the session.
 - d. In the **Choose Session Type** list, select **FlashCopy** and click **Next**.
 - e. On the Properties page, enter a session name, and a description. Complete the options that are described in "Session properties" on page 197. The options that are displayed depend on the storage system type. After you enter the information, click **Next**.
 - f. On the Site Locations page, select a location for Site 1 and click **Next**.
 - g. On the Results page, verify that the session was added successfully.
2. Follow these steps to add copy sets to the session:
 - a. On the Results page of the Create Session wizard, click **Launch Add Copy Sets Wizard**.
 - b. In the Add Copy Sets wizard, complete the following information. The field names that are displayed depend on the storage system type. When you have completed the information, click **Next**.

Storage system

Select a storage system. If the role has a location assigned to it, acceptable values for the storage system list are storage systems assigned to the same location as the role, and storage systems assigned to no location. In this case, the storage systems are grouped under different headings. If the role does not have a location, any storage system is acceptable.

Logical storage system or I/O Group

Select a logical storage system (LSS) or I/O group.

Volume

Select one volume or all volumes. The volumes are limited to the volumes within the LSS or I/O group that you selected.

Session image

Shows an image representing the session in which the role for which you are selecting volumes is highlighted. This image shows how many roles there are in the session and how the roles are distributed between the sites.

Volume Details

Shows information about the selected volume, including the volume name, full name, type, capacity, and whether the volume is protected and space efficient.

Use a CSV file to import copy sets

Select this option to import copy sets from a comma-separated value (CSV) file, click **Use a CSV file to import copy sets**. Type the full path name of the CSV file or click **Browse** to select the CSV file.

- c. On the Choose Target page, select the target storage system, LSS or I/O group, and volume. Click **Next**.
- d. On the Select Copy Sets page, select the copy sets that you want to add. You can click **Select All** to select all copy sets, **Deselect All** to select none of the copy sets, or **Add More** to add more copy sets to this session. Click **Next**.
- e. On the Confirm page, the number of copy sets to be added is displayed. Click **Next**.
- f. A progress bar is displayed. When the copy sets are added, review the results and click **Finish**.

Creating a Snapshot session and adding copy sets

A *snapshot* is a session type that creates a point-in-time copy of a volume or set of volumes without having to define a specific target volume. Snapshot sessions are available only for the XIV system.

1. Follow these steps to create a Snapshot session:
 - a. In the navigation tree, expand **Sessions**.
 - b. Click **Create Session**. The Create Session wizard starts.
 - c. In the **Choose Hardware Type** list, select **XIV**.
 - d. In the **Choose Session Type** list, select **Snapshot** and click **Next**.
 - e. On the Properties page, type a session name and description and click **Next**.
 - f. On the Site Locations page, select a location for Site 1 and click **Next**.
 - g. On the Results page, verify that the session was added successfully.
2. Follow these steps to add copy sets to the session:
 - a. On the Results page of the Create Session wizard, click **Launch Add Copy Sets Wizard**.
 - b. In the **Host1 storage system** list, select the storage system that contains the volumes that you want to add.

If the H1 role has an assigned location, only those storage systems that have the same location as the H1 role or storage systems that do not have a set location are displayed for selection.

If the H1 role does not have an assigned location, all storage systems are displayed for selection.

Storage systems that are assigned to a location are listed under the location name. Storage systems that are not assigned to a location are listed under the **None.** column.

- c. From the **Host1 pool** list, select the pool that contains the volumes.
- d. From the **Host1 volume** list, select the volumes. To select multiple volumes, press Ctrl or Shift and click the volumes in the list.
- e. If you want to import copy sets from a comma-separated value (CSV) file, click **Use a CSV file to import copy sets**. Type the full path name of the CSV file or click **Browse** to select the CSV file. Click **Next**.
- f. On the Matching Results page, click **Next** if the match was successful.
- g. On the Select Copy Sets page, select from the following options and click **Next**.

Select All

Click this button to select all of the copy sets in the table.

Deselect All

Click this button to clear all of the copy sets in the table.

Add More

Click this button to add another copy set to the list of copy sets to be created.

When you click **Add More**, you are returned to the Choose Host1 page of the wizard. The **Host1 storage system** and **Host1 pool** lists are populated with the values from the previously selected copy set. When you select the volumes for the various roles, the volumes are matched together and added to the list of copy sets on this page.

Selection check boxes

Select one or more copy sets that you want to create.

Host 1 Lists the volume IDs that are associated with the Host1 role. You can click the link to display information about the volume, including the full name, type, capacity, and whether the volume is protected and space efficient.

Copy Set

Displays the copy set information for the specified copy sets and any warning or error messages that are associated with the copy set.

A warning or icon next to the **Show** button indicates that you cannot create a copy set for the H1 volume. Click **Show** to view the message.

- h. On the Confirm page, the number of copy sets to be added is displayed. Click **Next**.
- i. A progress bar is displayed. When the copy sets are added, review the results and click **Finish**.

Creating a Metro Mirror session and adding copy sets

A Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 KM apart.

Follow the steps in this topic create a Metro Mirror session and add copy sets. If you want to enable HyperSwap or Open HyperSwap for the session, see “Managing a session with HyperSwap and Open HyperSwap replication” on page 42.

1. Follow these steps to create a Metro Mirror session:
 - a. In the navigation tree, select **Sessions**.
 - b. Click **Create Session**.
 - c. In the Create Session wizard, for the **Choose Hardware Type** list, select the item that shows the type of storage system for the session.
 - d. In the **Choose Session Type** list, select a Metro Mirror session type and click **Next**.
 - e. On the Properties page, enter a session name and description. Depending on the session type and storage system, additional options are displayed. Complete the options that are described in “Session properties” on page 197. After you enter the information on the Properties page, click **Next**.
 - f. On the site locations pages, choose a location for Sites 1 and 2 and click **Next**.
 - g. On the Results page, verify that the session was added successfully.
2. Follow these steps to add copy sets to the session:
 - a. On the Results page of the Create Session wizard, click **Launch Add Copy Sets Wizard**.
 - b. In the Add Copy Sets wizard, complete the following information for the Host pages. The field names that are displayed depend on the storage system type. When you have completed the information on each page, click **Next**.

Storage system

Select a storage system. If the role has a location assigned to it, acceptable values for the storage system list are storage systems assigned to the same location as the role, and storage systems assigned to no location. In this case, the storage systems are grouped under different headings. If the role does not have a location, any storage system is acceptable.

Logical storage system, I/O Group, or Pool

Select a logical storage system (LSS), I/O group, or pool.

Volume

Select one volume or all volumes. The volumes are limited to the volumes within the LSS, I/O group, or pool that you selected.

Session image

Shows an image representing the session in which the role for which you are selecting volumes is highlighted. This image shows how many roles there are in the session and how the roles are distributed between the sites.

Volume Details

Shows information about the selected volume, including the volume name, full name, type, capacity, and whether the volume is protected and space efficient.

Use a CSV file to import copy sets

Select this option to import copy sets from a comma-separated value (CSV) file, click **Use a CSV file to import copy sets**. Type the full path name of the CSV file or click **Browse** to select the CSV file.

- c. On the Select Copy Sets page, select the copy sets that you want to add. You can click **Select All** to select all copy sets, **Deselect All** to select none of the copy sets, or **Add More** to add more copy sets to this session. Click **Next**.
- d. On the Confirm page, the number of copy sets to be added is displayed. Click **Next**.
- e. A progress bar is displayed. When the copy sets are added, review the results and click **Finish**.

Creating a Global Mirror session and adding copy sets

A Global Mirror is a method of asynchronous, remote data replication between two sites that are over 300 KM apart.

Follow the steps in this topic create a Global Mirror session and add copy sets.

1. Follow these steps to create a Global Mirror session:
 - a. In the navigation tree, select **Sessions**.
 - b. Click **Create Session**.
 - c. In the Create Session wizard, for the **Choose Hardware Type** list, select the item that shows the type of storage system for the session.
 - d. In the **Choose Session Type** list, select a Global Mirror session type and click **Next**.
 - e. On the Properties page, enter a session name and description. Depending on the session type and storage system, additional options are displayed. Complete the options that are described in "Session properties" on page 197. After you enter the information on the Properties page, click **Next**.
 - f. On the site locations pages, choose a location for Sites 1 and 2 and click **Next**.
 - g. On the Results page, verify that the session was added successfully.
2. Follow these steps to add copy sets to the session:
 - a. On the Results page of the Create Session wizard, click **Launch Add Copy Sets Wizard**.
 - b. In the Add Copy Sets wizard, complete the following information for the Host and Journal pages. The field names that are displayed depend on the storage system type. When you have completed the information on each page, click **Next**.

Storage system

Select a storage system. If the role has a location assigned to it, acceptable values for the storage system list are storage systems assigned to the same location as the role, and storage systems assigned to no location. In this case, the storage systems are grouped under different headings. If the role does not have a location, any storage system is acceptable.

Logical storage system, I/O Group, or Pool

Select a logical storage system (LSS), I/O group, or pool from this list.

Volume

Select one volume or all volumes. The volumes are limited to the volumes within the LSS, I/O group, or pool that you selected.

Session image

Shows an image representing the session in which the role for

which you are selecting volumes is highlighted. This image shows how many roles there are in the session and how the roles are distributed between the sites.

Volume Details

Shows information about the selected volume, including the volume name, full name, type, capacity, and whether the volume is protected and space efficient.

Use a CSV file to import copy sets

Select this option to import copy sets from a comma-separated value (CSV) file, click **Use a CSV file to import copy sets**. Type the full path name of the CSV file or click **Browse** to select the CSV file.

- c. On the Select Copy Sets page, select the copy sets that you want to add. You can click **Select All** to select all copy sets, **Deselect All** to select none of the copy sets, or **Add More** to add more copy sets to this session. Click **Next**.
- d. On the Confirm page, the number of copy sets to be added is displayed. Click **Next**.
- e. A progress bar is displayed. When the copy sets are added, review the results and click **Finish**.

Creating a Metro Global Mirror session and adding copy sets

A Metro Global Mirror is a method of continuous, remote data replication that operates between three sites that varying distances apart. Metro Global Mirror combines Metro Mirror synchronous copy and Global Mirror asynchronous copy into a single session, where the Metro Mirror target is the Global Mirror source.

Follow the steps in this topic create a Metro Global Mirror session and add copy sets.

1. Follow these steps to create a Metro Global Mirror session:
 - a. In the navigation tree, select **Sessions**.
 - b. Click **Create Session**.
 - c. In the Create Session wizard, for the **Choose Hardware Type** list, select the item that shows the type of storage system for the session.
 - d. In the **Choose Session Type** list, select a Metro Global Mirror session type and click **Next**.
 - e. On the Properties page, enter a session name and a description. Complete the options that are described in “Session properties” on page 197. After you enter the information, click **Next**.
 - f. On the site locations pages, choose a location for Sites 1, 2, and 3 and click **Next**.
 - g. On the Results page, verify that the session was added successfully.
2. Follow these steps to add copy sets to the session:
 - a. On the Results page of the Create Session wizard, click **Launch Add Copy Sets Wizard**.
 - b. In the Add Copy Sets wizard, complete the following information for the Host and Journal pages. The field names that are displayed depend on the storage system type. When you have completed the information on each page, click **Next**.

Storage system

Select a storage system. If the role has a location assigned to it,

acceptable values for the storage system list are storage systems assigned to the same location as the role, and storage systems assigned to no location. In this case, the storage systems are grouped under different headings. If the role does not have a location, any storage system is acceptable.

Logical storage system

Select a logical storage system.

Volume

Select one volume or all volumes. The volumes are limited to the volumes within the LSS that you selected.

Session image

Shows an image representing the session in which the role for which you are selecting volumes is highlighted. This image shows how many roles there are in the session and how the roles are distributed between the sites.

Volume Details

Shows information about the selected volume, including the volume name, full name, type, capacity, and whether the volume is protected and space efficient.

Use a CSV file to import copy sets

Select this option to import copy sets from a comma-separated value (CSV) file, click **Use a CSV file to import copy sets**. Type the full path name of the CSV file or click **Browse** to select the CSV file.

- c. On the Select Copy Sets page, select the copy sets that you want to add. You can click **Select All** to select all copy sets, **Deselect All** to select none of the copy sets, or **Add More** to add more copy sets to this session. Click **Next**.
- d. On the Confirm page, the number of copy sets to be added is displayed. Click **Next**.
- e. A progress bar is displayed. When the copy sets are added, review the results and click **Finish**.

Using the Metro Mirror heartbeat

This topic provides information about Metro Mirror heartbeat, including how to enable and disable the heartbeat.

Metro Mirror heartbeat

The *heartbeat* is a Metro Mirror function. When the Metro Mirror heartbeat is disabled, data consistency across multiple storage systems is not guaranteed if the IBM Tivoli Storage Productivity Center for Replication management server cannot communicate with one or more storage systems. The problem occurs as a result of the Hardware Freeze Timeout Timer function within the storage system. If the controlling software loses connection to a storage system, the Metro Mirror relationships that it is controlling stay established and there is no way to *freeze* those pairs to create consistency across the multiple storage systems. When the freeze times out, dependent I/O is written to the target storage systems, which might corrupt data consistency. Freeze refers to a Metro Mirror (peer-to-peer remote copy [PPRC]) freeze function.

When determining whether to use the Metro Mirror heartbeat, analyze your business needs. Disabling the Metro Mirror heartbeat might result in data

inconsistency. If you enable the Metro Mirror heartbeat and a freeze occurs, your applications will be unable to write during the freeze.

Metro Mirror heartbeat is disabled by default.

Metro Mirror heartbeat is not available for Metro Mirror with HyperSwap or Metro Global Mirror with HyperSwap.

There are two cases where lost communication between the coordination software (controller) and one or more storage systems can result in data consistency loss:

Freeze event not detected by a disconnected storage system

Consider a situation with four storage system machines in a primary site and four in a secondary site. One of the four storage systems on the primary loses the connection to the target site. This causes the affected storage system to prevent any writes from occurring, for a period determined by the Freeze timeout timer. At the same time, the affected storage controller loses communication with the controlling software and cannot communicate the Freeze event to the software.

Unaware of the problem, the controlling software does not issue the Freeze command to the remaining source storage systems. The freeze will stop dependent writes from being written to connected storage systems. However, once the Freeze times out and the long-busy is terminated, dependent write I/Os continue to be copied from the storage systems that did not receive the Freeze command. The Metro Mirror session is left in a state where one storage system has suspended copying while the other three storage systems are still copying data. This state causes inconsistent data on the target storage systems.

Freeze event detected, but unable to propagate the Freeze command to all storage systems

Consider a situation with four storage system machines in a primary site and four in a secondary site. One of the four storage systems on the primary loses the connection to the target site. This causes the affected storage system to issue long-busy to the applications for a period determined by the Freeze timeout timer. At the same time, one of the remaining three source systems loses communications with the controlling software.

The storage system that had an error writing to its target cannot communicate the Freeze event to the controlling software. The controlling software issues the Freeze command to all but the disconnected storage system (the one that lost communication with the software). The long-busy stops dependent writes from being written to the connected storage systems.

However, once the Freeze times out on the frozen storage system and the long-busy is terminated, dependent write I/Os continue to the target storage system from the source storage system that lost communication and did not receive the Freeze command. The Metro Mirror session is left a state where three storage systems have suspended copying and one storage system is still copying data. This state causes inconsistent data on the target storage systems.

Before IBM Tivoli Storage Productivity Center for Replication V3.1, if the controlling software within a Metro Mirror environment detected that a managed storage system lost its connection to its target, the controlling software stopped all

the other source systems to ensure consistency across all the targets. However, if the controlling software lost communication with any of the source subsystems during the failure, it could not notify those storage systems of the freeze event or ensure data consistency. The Metro Mirror heartbeat helps to overcome this problem. In a high-availability configuration, the Metro Mirror heartbeat is continued by the standby server after the Takeover command is issued on the standby, enabling you to perform actions on the standby server without causing a freeze.

IBM Tivoli Storage Productivity Center for Replication registers with the managed ESS 800, DS6000 or DS8000 storage systems within a Metro Mirror session when the start command is issued to the session. After this registration occurs, a constant heartbeat is sent to the storage system. If the storage system does not receive a heartbeat from the IBM Tivoli Storage Productivity Center for Replication management server within the allotted time (a subset of lowest LSS timeout value across all the source LSSs), the storage system initiates a freeze. If IBM Tivoli Storage Productivity Center for Replication did not successfully communicate with the storage system, it initiates a freeze on the remaining storage system after the allotted time is expired.

Note: It is recommended that you avoid using the same LSS pairs for multiple Metro Mirror sessions. Metro Mirror uses a freeze command on ESS, DS6000, and DS8000 storage systems to create the data-consistent point. If there are other Metro Mirror sessions overlapping the same LSS pairs as in this session, those sessions are also suspended.

When you are using the Metro Mirror heartbeat, be aware that:

- The Metro Mirror heartbeat can cause a single point of failure: if an error occurs on just the management server and not the storage system, a freeze might occur.
- When the Metro Mirror heartbeat timeout occurs, the storage system remains in a long busy state for the duration of the LSS freeze timeout.

Note: If Metro Mirror heartbeat is enabled for storage systems that are connected through a HMC connection, a connection loss might cause lost Metro Mirror heartbeats, resulting in Freeze actions with application I/O impact for configured Extended Long Busy timeout.

The Metro Mirror heartbeat is supported on storage systems connected through a TCP/IP (direct connect or HMC) connection. It is not supported on storage systems connected through a z/OS connection. Enabling the Metro Mirror heartbeat with a z/OS connection does not fail; however, a warning message is displayed specifying that the Metro Mirror heartbeat function does not work unless you have an IP connection.

If Metro Mirror heartbeat is enabled for storage systems that are connected through a TCP/IP (either direct connect or HMC) connection and z/OS connection, and the TCP/IP connection fails, IBM Tivoli Storage Productivity Center for Replication suspends the Metro Mirror session because there is no heartbeat through the z/OS connection.

If Metro Mirror heartbeat is enabled for storage systems that are connected through a TCP/IP connection and z/OS connection and you remove all TCP/IP connections, IBM Tivoli Storage Productivity Center for Replication suspends the Metro Mirror sessions and the applications using those volume will be in Extended Long Busy timeout until the storage system's internal timeout timer expires. Ensure

that you disable the Metro Mirror heartbeat for all Metro Mirror sessions before removing the last TCP/IP connection to avoid the Extended Long Busy timeout.

Enabling and disabling the Metro Mirror heartbeat

The Metro Mirror heartbeat guarantees data consistency across multiple storage systems when the IBM Tivoli Storage Productivity Center for Replication management server cannot communicate with one or more storage systems. The Metro Mirror heartbeat is disabled by default.

To enable the Metro Mirror heartbeat, perform the following steps:

1. In the navigation tree, select **Advanced Tools**.
2. To enable the Metro Mirror heartbeat, click **Enable Heartbeat**.
3. To disable the Metro Mirror heartbeat, click **Disable Heartbeat**.

Exporting copy set data

You can export data about all copy sets in a specific session, to maintain a backup copy that can be used to recover if you lose your session or upgrade to a different server.

Perform these steps to export the copy sets in a specific session:

1. In the navigation tree, select **Sessions**. The Sessions panel is displayed.
2. Select the session for which you want to export copy sets.
3. Select **Export Copy Sets** from the **Actions** list, and click **Go**. The Export Copy Set wizard displays the status of the export and a link to the exported file.
4. Click that link and save the file to the local system.

Important: You must save the file to your local system. After you close the panel, the data will be lost.

5. Click **Finish**.

Importing copy set data

You can import copy set data that was previously exported to a comma separated value (CSV) file.

Perform the following steps to import copy sets into an existing session:

1. In the navigation tree, select **Sessions**. The Session panel is displayed.
2. Select the session for which you want to import copy sets.
3. Select **Add Copy Sets** from the **Actions** list, and click **Go**. The Add Copy Sets wizard is displayed.
4. Select **Use a CSV file to import copy sets**.
5. Type the location and name of the CSV file to import, or use **Browse** to select the file. Then, click **Next**.
6. Verify that the matching results were successful, and then click **Next**.
7. Select the copy sets you want to add, and then click **Next**.
8. Confirm the number of copy sets that you want to create, and click **Next**. A progress bar displays.
9. Click **Next**.
10. Verify the matches, and click **Finish**.

Modifying the location of session sites

You can change the location associated with each site of in session.

Prerequisites: You must have Administrator privileges to modify the location of a site.

Changing the location of a site in a session does not affect the location of the storage systems associated with that site.

Perform these steps to modify the location of a site:

1. In the navigation tree, select **Sessions**. The Sessions panel is displayed.
2. Select the session whose site locations you want to change.
3. Click **Modify Site Location(s)** from the actions list, and then click **Go**. The Modify Site Locations wizard is displayed.
4. Select the location for the site from the drop-down list, and then click **Next**. Repeat for each available site.

To disable site awareness, set the location to None.

Note: You can select only locations that have already been associated with one or more storage systems.

5. Click **Next**.
6. Click **Finish**.

Removing sessions

This topic describes how to remove sessions.

Important: You can remove only sessions that are in the Defined state.

Perform these steps to remove a session:

1. In the navigation tree, select **Sessions**.
2. Click the radio button next to the session you want to remove.
3. Select **Remove Session** from the **Actions** drop-down menu, and click **Go**.

Removing copy sets

This topic describes how to remove copy sets.

Perform these steps to remove a copy set:

1. In the IBM Tivoli Storage Productivity Center for Replication navigation tree, select **Sessions**. Click the radio button next to the session that you want to remove copy sets from.
2. From the **Select Action** menu, select **Remove Copy Sets** and click **Go**. This starts the Remove Copy Sets wizard.
3. From the drop-down menus in the Remove Copy Sets wizard, select the Host 1 storage system, logical storage subsystem, and volume or select the all option. If you select all for a filter, the lower level filter or filters are disabled. Click **Next**.
4. Select the copy sets that you want to remove and click **Next**.
5. The number of copy sets to be removed is displayed. Select the following options for removing the copy sets and click **Next**:

- **Do you want to keep the base relationships on the hardware, but remove the copy sets from the session?**
 - **Yes.** This option specifies that the base relationships remain on the hardware, but the copy sets are removed from the Tivoli Storage Productivity Center for Replication session. This option supports scenarios in which it might be best to leave the relationship on the hardware to avoid performing a full copy. For example, when you are migrating from one session type to another.
Only the base relationships (Metro Mirror, Global Copy, Snapshot, and FlashCopy) are left on the hardware. The relationships are removed from any consistency groups that are defined on the storage system.
 - **No.** This option specifies that all relationships for the copy sets are removed from the hardware as well as the Tivoli Storage Productivity Center for Replication session. This option is the default.
- **If there are errors removing relationships on the hardware, do you want to force the copy sets to be removed from the session?**
 - **Yes.** This option forces the removal of copy sets despite any errors that occur when removing the relationships from the storage system. Once a forced removal is complete, any relationships that remain on the storage system for that copy set must be removed manually using the storage system interface.
 - **No.** This option does not force the removal of copy sets. This option enables you to correct the errors and try to remove the copy sets again. This option is the default.

6. After the copy sets are removed, click **Finish**.

Important: If an application on the host has opened a device, the copy sets in a Tivoli Storage Productivity Center for Replication session for that device are removed, but the copy sets remain coupled on the host. To decouple the copy sets, see the Troubleshooting and support section of the Tivoli Storage Productivity Center for Replication Information Center at <http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp>.

Migrating an existing configuration to Tivoli Storage Productivity Center for Replication

This topic describes how to convert an existing DS6000 or DS8000 hardware relationship for a Metro Mirror, Global Mirror, or Metro Global Mirror session. This topic also describes how to assimilate pairs in a relationship for a Metro Mirror or Global Mirror session with SAN Volume Controller, Storwize V7000 Unified or Storwize V7000 volumes.

You can either complete this action manually or use the data migration utility for Tivoli Storage Productivity Center for Replication. You can download the data migration utility from the Data Migration Utility for IBM(r) Tivoli(r) Storage Productivity Center for Replication. The data migration utility can handle all session types, and produces the Tivoli Storage Productivity Center for Replication CLI command script files and DSCLI script files necessary to migrate the relationships to Tivoli Storage Productivity Center for Replication.

Tivoli Storage Productivity Center for Replication cannot assimilate pairs that are suspended on the hardware. All pairs must be active for Tivoli Storage Productivity Center for Replication to assimilate the pairs.

Metro Mirror

To manually assimilate pairs in a relationship for a Metro Mirror session, complete the following steps:

1. Determine which existing Metro Mirror pairs you want Tivoli Storage Productivity Center for Replication to manage, and record the names of your source and target volumes.
2. Create a Metro Mirror session in Tivoli Storage Productivity Center for Replication.
3. Using the Add Copy Set Wizard, choose the same source and target volumes that you identified in step 1.
4. Complete the Add Copy Set Wizard, and ensure that all copy sets are created successfully.
5. Issue a **Start** command from Tivoli Storage Productivity Center for Replication. At this time, the session automatically assimilates the relationships on the hardware (for example, if they are already in a Prepared state, Tivoli Storage Productivity Center for Replication will show them as Prepared).

Note: The **Start** command does not fully resynchronize data; instead, it detects the existing relationships and adopts them.

Global Mirror

To manually assimilate pairs in a relationship for a Global Mirror session, complete the following steps:

1. Terminate the Global Mirror session on the hardware using the IBM System Storage DS[®] command-line interface (DSCLI) or another application.
2. Remove the volumes from the Global Mirror session using the DSCLI or another application.
3. Close the session on each LSS upon which it is open using the DSCLI or another application.
4. Determine which existing Global Mirror pairs you want Tivoli Storage Productivity Center for Replication to manage, and record the names of your source and target volumes.
5. Create a Global Mirror session in Tivoli Storage Productivity Center for Replication.
6. Using the Add Copy Set Wizard, choose the same source and target volumes that you identified in step 4.
7. Complete the Add Copy Set Wizard, and ensure that all copy sets are created successfully.
8. Issue a **Start** command from Tivoli Storage Productivity Center for Replication. At this time, the session automatically assimilates the relationships on the hardware.

Note: The **Start** command does not fully resynchronize data; instead, it detects the existing relationships and adopts them.

Assimilating Metro Mirror pairs into a Three Site Metro Global Mirror session

To manually assimilate a set of Metro Mirror pairs into a Three Site Metro Global Mirror session, complete the following steps:

1. Determine which existing Metro Mirror pairs you want Tivoli Storage Productivity Center for Replication to manage, and record the names of your source and target volumes.
2. Create a Metro Global Mirror session in Tivoli Storage Productivity Center for Replication.
3. Using the Add Copy Set Wizard, choose the same source and target volumes for the H1 and H2 roles of the Metro Global Mirror copy set, that you identified in step 1.
4. Complete the Add Copy Set Wizard, and ensure that all copy sets are created successfully.
5. Issue a **Start** command from Tivoli Storage Productivity Center for Replication. At this time, the session automatically assimilates the Metro Mirror relationships on the hardware (for example, if they are already in a Prepared state, Tivoli Storage Productivity Center for Replication will show them as Prepared). In addition, the **Start** command will establish the Global Copy relationships.

Note: The **Start** command does not fully resynchronize data on the Metro Mirror pairs assimilated; instead, it detects the existing relationships and adopts them. However, a full synchronization will be required for the Global Copy pairs.

Assimilating Three Site pairs into a Three Site Metro Global Mirror session

To assimilate a full Three Site set of pairs into a Three Site Metro Global Mirror session, complete the following steps:

1. Terminate the Global Mirror session on the hardware using the IBM System Storage DS command-line interface (DSCLI) or another application.
2. Remove the volumes from the Global Mirror session using the DSCLI or another application
3. Close the session on each LSS upon which it is open using the DSCLI or another application.
4. Determine which existing Metro Mirror and Global Copy pairs you want Tivoli Storage Productivity Center for Replication to manage, and record the names of your source and target volumes.
5. Create a Metro Global Mirror session in Tivoli Storage Productivity Center for Replication.
6. Using the Add Copy Set Wizard, choose the same source and target volumes for the H1, H2 and H3 roles of the Metro Global Mirror copy set, that you identified in step 4.
7. Complete the Add Copy Set Wizard, and ensure that all copy sets are created successfully.
8. Issue a Start command from IBM Tivoli Storage Productivity Center for Replication. At this time, the session automatically assimilates the Metro Mirror and Global Copy relationship on the hardware.

Note: The **Start** command does not fully resynchronize data on the Metro Mirror and Global Copy pairs assimilated; instead, it detects the existing relationships and adopts them.

Global Mirror and Metro Mirror assimilation for SAN Volume Controller, Storwize V7000, Storwize V7000 Unified or the XIV system

To manually assimilate pairs in a relationship for a Metro Mirror or Global Mirror session with SAN Volume Controller, Storwize V7000, Storwize V7000 Unified or the XIV system volumes, complete the following steps:

1. Remove all pairs from the consistency group and delete the consistency group on the storage system using CLI for the storage system.
2. Determine which existing Metro Mirror and Global Mirror pairs you want Tivoli Storage Productivity Center for Replication to manage, and record the names of your source and target volumes.
3. Create a Metro Mirror or Global Mirror session in Tivoli Storage Productivity Center for Replication.
4. Using the Add Copy Set Wizard, choose the same source and target volumes that you identified in step 2.
5. Complete the Add Copy Set Wizard, and ensure that all copy sets are created successfully.
6. Issue a **Start** command from Tivoli Storage Productivity Center for Replication. The session automatically assimilates the relationships on the hardware and adds them to a new consistency group on the storage system.

Note: The **Start** command does not fully resynchronize data; instead, it detects the existing relationships and adopts them placing them in a consistency group managed by the session.

Chapter 8. Practicing disaster recovery

You can use practice volumes to test your disaster recovery actions while maintaining disaster recovery capability. Practice volumes are available in Metro Mirror Failover and Failback sessions, Global Mirror Failover and Failback, Global Mirror Either Direction sessions, and Metro Global Mirror with Practice sessions.

Note: You can test your disaster recovery actions without using practice volumes; however, without practice volumes, you cannot maintain disaster recovery capability while continuing to copy the data.

Practice volumes

You can use a *practice volume* to practice what you would do in the event of a disaster, without interrupting current data replication. Practice volumes are available in Metro Mirror, Global Mirror, and Metro Global Mirror sessions.

To use the practice volumes, the session must be in the prepared state. Issuing the Flash command against the session while in the Prepared state creates a usable practice copy of the data on the target site.

Note: You can test disaster-recovery actions without using practice volumes; however, without practice volumes, you cannot continue to copy data changes between volumes while testing disaster-recovery actions.

Practicing disaster recovery for a Metro Mirror Failover/Failback with Practice session

A Metro Mirror Failover and Failback session with Practice combines Metro Mirror and FlashCopy to provide a point-in-time copy of the data on the remote site. You can use this to practice what you might do if a disaster occurred, without losing your disaster recovery capability.

This function is available on the following storage systems:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified
- TotalStorage Enterprise Storage Server Model 800
- System Storage DS8000
- System Storage DS6000

Perform these steps to practice disaster recovery actions for a Metro Mirror Failover/Failback with Practice session:

1. Start a Metro Mirror with Practice session.
2. When the Metro Mirror session reaches the Prepared state, issue a Flash command to make a point-in-time copy of the data on H2. This creates a consistent point-in-time copy of your data on the H2 volume and then restarts the session so the copying from H1 to I2 continues. This temporarily stops copying of the data from site 1 to site 2, and creates a consistent point-in-time copy. The data replication from H1 to I2 is then restarted.

Note: For ESS, DS6000, DS8000 storage systems, the **Flash** command uses the freeze and thaw processing to create a data consistent point for the FlashCopy. If there is another Metro Mirror session overlapping on one or more of the same LSS pairs, that session will be suspended. It is also possible that the suspension of the other session might cause the Metro Mirror session to remain suspended after the flash command is issued instead of returning to Prepared state. Avoid using the same LSS pairs for multiple Metro Mirror sessions if possible.

3. Practice the same actions you would take in an actual disaster, using H2 as your new host volume, while running the real application on H1.

Practicing disaster recovery for a Global Mirror either Direction with two-site Practice session

A Global Mirror (either direction) with two-site Practice combines Global Mirror and FlashCopy to provide a point-in-time copy of the data on a remote site at a distance over 300 km away from your first site. You can use this to practice what you might do if a disaster occurred, without losing your disaster recovery capability.

Note: This function is available only on ESS, DS6000, and DS8000 storage systems.

Perform these steps to practice disaster recover actions for a Global Mirror either Direction with two-site Practice session:

1. Start a Global Mirror with Practice session.
2. When the session reaches the Prepared state, issue a Flash command to restore consistent data on I2 and make a point-in-time copy of the data on H2. This creates a consistent point-in-time copy of your data on the H2 volume and then restarts the session so the copying from H1 to I2 continues. This temporarily stops copying of the data from site 1 to site 2, and creates a consistent point-in-time copy. The data replication from H1 to I2 is then restarted.

Note: FlashCopy must be always a full copy due to limitations of the hardware.

3. Practice the same actions you would take in an actual disaster, using H2 as your new host volume, while running the real application on H1.

Note: With two directions, you can reverse the direction of your data flow.

Practicing disaster recovery for a Global Mirror Failover/Failback with Practice session

A Global Mirror Failover and Failback with Practice combines Global Mirror and FlashCopy to provide a point-in-time copy of the data on a remote site at a distance over 300 km away from your first site. You can use this to practice what you might do if a disaster occurred.

You can do this practice without losing your disaster recovery capability. The number of volumes used for the device varies, but the steps to conduct a Global Mirror Failover and Failback with Practice are the same for both devices.

This function is available on the following storage systems:

- SAN Volume Controller

- Storwize V7000
- Storwize V7000 Unified
- TotalStorage Enterprise Storage Server Model 800
- System Storage DS8000
- System Storage DS6000

Perform these steps to practice disaster recover actions for a Global Mirror Failover/Failback with Practice session:

1. Start a Global Mirror with Practice session.
2. When the session reaches the Prepared state, issue a Flash command to restore consistent data on I2 and make a point-in-time copy of the data on H2. This creates a consistent point-in-time copy of your data on the H2 volume and then restarts the session so the copying from H1 to I2 continues. This temporarily stops copying of the data from site 1 to site 2, and creates a consistent point-in-time copy. The data replication from H1 to I2 is then restarted.

Note: FlashCopy must be always a full copy due to limitations of the hardware.

3. Practice the same actions you would take in an actual disaster, using H2 as your new host volume, while running the real application on H1.

Practicing disaster recovery for a Metro Global Mirror Failover/Failback with Practice session

A Metro Global Mirror Failover/Failback with Practice session combines Metro Mirror, Global Mirror and FlashCopy across three sites to provide a point-in-time copy of the data on the third site. You can use this to practice what you might do if a disaster occurred without losing your disaster recovery capability.

Note: This function is available on ESS, DS6000, and DS8000 storage systems.

The intermediate volume is on the third site (I3). This maintains disaster recovery capability while a copy is kept on the H3 volume for practice purposes.

Perform these steps to practice disaster recover actions for a Metro Global Mirror Failover/Failback with Practice session:

1. Start a Metro Global Mirror with Practice session.
2. When the session reaches the Prepared state, issue a Flash command to take a point-in-time copy of the data that is on I3, on H3. This creates a consistent point-in-time copy of your data on the H3 volume, and then automatically restarts the session so that copying from H1 to H2 to I3 continues. The Flash command temporarily stops copying the data from site 2 to site 3, in order to create a consistent point-in-time copy on I3, while maintaining disaster recovery capabilities on site 2 using the Metro Mirror portion of the session. Then, data replication from H2 to I3 is restarted.

Note: FlashCopy must be always a full copy due to hardware limitations.

3. Practice the same actions you would take in an actual disaster, using H3 as your practice host volume, while you run the real application on H1. This enables you to use the same scripts and commands to run on H3 that you would use in an actual disaster.

Chapter 9. Monitoring health and status

Viewing the health summary

Use the Health Overview panel to view overall health and status of sessions, storage systems, host systems, and management servers.

The Health Overview panel is the first panel that you see after you log on. You can display this panel by selecting **Health Overview** in the IBM Tivoli Storage Productivity Center for Replication navigation tree. This panel provides the following information:

- Overall session status: indicates session status, which can be normal, warning, or severe. The status can also be inactive, if all sessions are defined or if no sessions exist.
- Overall storage system status: indicates the connection status of storage systems.
- Overall host system status: indicates the connection status of host systems.
- Management server status (applicable only if you are using a Business Continuity license): indicates the status of the standby server if you are logged on to the local server. If you are logged on to the standby server, this status indicates the status of the local server.

Health information is always shown as a mini-panel under the navigation tree.

Viewing SNMP alerts

IBM Tivoli Storage Productivity Center for Replication SNMP trap descriptions can be viewed from the IBM Tivoli Storage Productivity Center Alert panel.

From the IBM Tivoli Storage Productivity Center GUI, in the navigation tree, expand **Alerting > Alert Log > All** and select **Replication**.

Viewing sessions

This section describes how to view sessions details and session properties.

Session status icons

The IBM Tivoli Storage Productivity Center for Replication GUI uses icons to represent the status of each session.

The following table describes each session status icon.

Table 33. 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.

Table 33. Session status icons (continued)

Icon	Meaning	Description
	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, and means that the session is either preparing or is ready for a flash command but targets do not yet have a consistent copy.
	Severe	One or more errors must be dealt with immediately. Possible causes include the following: <ul style="list-style-type: none"> • One or more volumes are suspended • A session is suspended • A volume is not copying correctly







Session images

The IBM Tivoli Storage Productivity Center for Replication GUI 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. It also shows the copy method and direction.

Volume role symbols

The volume role symbols represent the replication status on the volumes.

Table 34. 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).
	Recoverable volumes	This symbol represents 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.
	Selected volumes	This symbol represents the volumes that are selected for an operation (for example, changing location or displaying role pair information).
		
		

Data copying symbols

The data copying symbols are arrows indicate the type of copy that occurs between the volume roles. The direction of the arrow indicates the direction of the copy.

Table 35. Data copying symbols


















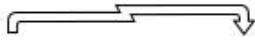



Symbol	Description	Meaning
	FlashCopy copying (the lightning bolt indicates the direction of the FlashCopy)	This symbol represents a FlashCopy relationship in which data is being copied from the host to the target.
		
	FlashCopy copying with errors (the lightning bolt indicates the direction of the FlashCopy)	This symbol represents a FlashCopy relationship in which data is being copied from the host to the target, but there are errors on one or more pairs.
		
	FlashCopy inactive (the lightning bolt indicates the direction of the FlashCopy)	This symbol represents an inactive FlashCopy relationship.
		
	FlashCopy inactive with errors (the lightning bolt indicates the direction of the FlashCopy)	This symbol represents an inactive FlashCopy relationship in which there are errors on one or more pairs.
		
	Metro Mirror copying	This symbol represents a copying Metro Mirror relationship.
	Metro Mirror copying with errors	This symbol represents a Metro Mirror relationship that is copying, but with errors on one or more pairs.
	Metro Mirror inactive	This symbol represents an inactive Metro Mirror relationship.
	Metro Mirror inactive with errors	This symbol represents an inactive Metro Mirror relationship with errors on one or more pairs.
	Global Copy copying	This symbol represents a copying Global Copy relationship.
		
	Global Copy copying with errors	This symbol represents a copying Global Copy relationship with errors on one or more pairs.
		

Table 35. Data copying symbols (continued)

Symbol	Description	Meaning
	Global Copy inactive	This symbol represents an inactive Global Copy relationship.
		
	Global Copy inactive with errors	This symbol represents an inactive Global Copy relationship with errors on one or more pairs.
		
	HyperSwap or Open HyperSwap	This symbol HyperSwap or Open HyperSwap for a session. If a failure occurs when input/output (I/O) is being written to the primary storage system, HyperSwap or Open HyperSwap automatically swap the I/O to the secondary site with no user interaction and little or no application impact.

Session states

You can view the health and status of a session in the Tivoli Storage Productivity Center for Replication GUI.

Important: Use only the Tivoli Storage Productivity Center for Replication GUI or CLI to manage session relationships, such as volume pairs and copy sets. Do not modify session relationships through individual hardware interfaces, such as the DSCLI. Modifying relationships through the individual hardware interfaces can result in a loss of consistency across the relationships managed by the session, and might cause the session to be unaware of the state or consistency of the relationships.

You can use the Refresh States command is used to refresh the states of the session. Issue this command to query the states of the copy sets on the hardware. You do not need to run this command under typical circumstances; Tivoli Storage Productivity Center for Replication refreshes the states of its sessions through multiple means. However, if you discover an inconsistency between Tivoli Storage Productivity Center for Replication and the hardware, you can use this command to enable IBM Tivoli Storage Productivity Center for Replication to update itself. Because this command triggers multiple queries on the hardware (having an adverse impact on hardware performance if you run it too often), you can only execute this command every few minutes in each session.

The following table describes each session state.

Table 36. Session states

State	Session type	Description
Defined	All	The session exists but is inactive.

Table 36. Session states (continued)

State	Session type	Description
Flashing	All	<p>In a Metro Mirror or Global Mirror session, data copying is temporarily suspended while a consistent practice copy of data is being prepared on site 2.</p> <p>In a Metro Global Mirror session, data copying is temporarily suspended while a consistent practice copy of data is being prepared on site 3.</p>
Prepared	All	<p>The source to target data transfer is active.</p> <p>In a Metro Mirror, Global Mirror, or Metro Global Mirror session, the data written to the source is transferred to the target, and all volumes are consistent and recoverable.</p> <p>In a FlashCopy session, the volumes are not yet consistent, but the flash is ready to begin.</p> <p>Note: For sessions on the following storage systems, do not alter the relationships on the hardware that you established with Tivoli Storage Productivity Center for Replication:</p> <ul style="list-style-type: none"> • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified <p>For example, if a Metro Mirror session with one copy set is in the Prepared state, and you stop the role pair, the session still displays in the Prepared state.</p>

Table 36. Session states (continued)

State	Session type	Description
Preparing	All	<p>The volumes are initializing, synchronizing, or resynchronizing.</p> <p>In a Metro Mirror, Global Mirror, or Metro Global Mirror session, synchronization occurs after the first Start command is issue on a session. Resynchronization occurs when a volume was prepared and then suspended. The hardware records the changed tracks so that on the next startup, only the changed tracks are copied.</p> <p>In a FlashCopy session, the volumes are initializing. The preparing state for FlashCopy sessions applies only to the following storage systems:</p> <ul style="list-style-type: none"> • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified
Recovering	All	The session is in the process of recovering.
Suspended	All	<p>Data copying has temporarily stopped.</p> <p>Note: The suspended state applies only to Global Mirror, Metro Mirror, and Metro Global Mirror sessions.</p>
SuspendedH1H2 MGM	Metro Global Mirror	Data copying between site 1 and site 2 is suspended.
SuspendedH1H3 MGM	Metro Global Mirror	Data copying between site 1 and site 3 is suspended.
Suspending	All	<p>The session is transitioning into a Suspended state.</p> <p>Note: The Suspending state applies only to Global Mirror and Metro Global Mirror sessions and does not apply to the following storage systems:</p> <ul style="list-style-type: none"> • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified
Target available	All	Target volumes are available for application updates.

Table 36. Session states (continued)

State	Session type	Description
Terminating	FlashCopy	<p>The session is being terminated because you issued a Terminate action under the following conditions:</p> <ul style="list-style-type: none"> • You permitted the target to be Metro Mirror or Global Copy. • You set the Require or Attempt to Preserve Mirror option. <p>The session displays as Terminating until the FlashCopy background copy is complete and no longer exists on the hardware.</p>

Session properties

This topic lists the session properties.

Session properties include the following attributes:

Description

Type the description for this session.

ESS/DS FlashCopy Options

Select the FlashCopy options that you want to associate with this session:

Incremental

Select this option to set up the relationship for recording changes to the H1 volume. Any subsequent FlashCopy operation for that session copies only the tracks that have changed since the last flash. Incremental always assumes persistence.

Note: This option is used along with the Persistent and No-Copy options.

Persistent

Select this option to keep the relationship established on the hardware after all source tracks are copied to the target. If this option is not selected, the local replication relationship ends after the T1 volume contains a complete point-in-time image of the H1 volume.

Note: If you select the Incremental option, the Persistent option is automatically selected as well.

No Copy

Select this option if the hardware is not to write the background copy until the source track is written to. Replication is done by using a copy-on-write technique. This technique does not copy data to the T1 volume until the blocks or tracks of the H1 volume are modified. The point-in-time volume image is composed of the unmodified data on the H1 volume, and the data that was copied

to the T1 volume. If you want a complete point-in-time copy of the H1 volume to be created on the T1 volume, do not use the No Copy option. This option causes the data to be asynchronously copied from the H1 volume to the T1 volume.

Note: Although you can select any space-efficient volume as the target, you cannot change the Permit Space Efficient Target flag. This flag is always set. When selecting space-efficient volumes as targets, you might receive an x0FBD error message if you attempt a full background copy. To avoid this message, select the No Copy option.

Allow FlashCopy target to be Metro Mirror source

Select this option if the target of the FlashCopy relationship is to overwrite all data on the source of a Metro Mirror relationship. If this option is cleared, a flash to a Metro Mirror source volume fails.

If you select this option, you must also select one of the following policies:

Do not attempt to preserve Metro Mirror consistency

Select this option if the Metro Mirror or Global Copy pair at the target of the FlashCopy relationship is to perform a full copy of the data to the secondary of the Metro Mirror or Global Copy pair.

Attempt to preserve Metro Mirror consistency, but allow FlashCopy even if Metro Mirror target consistency cannot be preserved

Select this option to attempt to preserve the consistency of the Metro Mirror relationship at the target of the FlashCopy relationship when both the source and target of the FlashCopy relationship are a source of a Metro Mirror relationship. If the consistency cannot be preserved, a full copy of the Metro Mirror relationship at the target of the FlashCopy relationship is performed. To preserve consistency, parallel flashes are performed (if possible) on both sites.

Note: This option is available only on IBM System Storage DS8000 storage systems with the required code levels installed.

Attempt to preserve Metro Mirror consistency, but fail FlashCopy if Metro Mirror target consistency cannot be preserved

This option prevents a full copy from being performed over the Metro Mirror link. Instead, parallel flashes are performed (if possible) on both sites. If the consistency cannot be preserved, the Flash for the FlashCopy relationships fails, and the data of the Metro Mirror relationship at the target of the FlashCopy relationship is not changed.

Note: This option is available only on System Storage DS8000 storage systems with the required code levels installed.

SAN Volume Controller / Storwize V7000 / Storwize V7000 Unified FlashCopy Options

Select the FlashCopy options that you want to associate with this session.

Incremental

Select this option to set up the relationship for recording changes to the H1 volume. Any subsequent FlashCopy operation for that session copies only the tracks that have changed since the last flash. Incremental always assumes persistence.

Background Copy Rate

Type the copy rate that the Storage System uses to perform the background copy of the FlashCopy role pair. You can specify a percentage between 0 and 100. The default is 50. Specifying 0 is equivalent to specifying the No-Copy option on a TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, or System Storage DS6000 FlashCopy session.

You can modify this value at any time during the session. If the session is performing a background copy when you change the option, Tivoli Storage Productivity Center for Replication immediately modifies the background copy rate of the consistency group on the storage system. The storage system consistency group immediately starts by using this new rate to complete the background copy.

Basic HyperSwap Options

Select the Basic HyperSwap options that you want to associate with this session.

If Tivoli Storage Productivity Center for Replication is running on z/OS, and volumes are attached by a fibre-channel connection, Tivoli Storage Productivity Center for Replication can manage a HyperSwap session. With this option selected, a failure on the host accessible volumes triggers a HyperSwap, redirecting application input/output (I/O) to the secondary volumes.

These options are available only for Basic HyperSwap sessions:

Disable HyperSwap

Select this option to prevent a HyperSwap from occurring by command or event.

On Configuration Error:

Partition the system(s) out of the sysplex

Select this option to partition out of the sysplex when a new system is added to the sysplex and encounters an error in loading the configuration. Restart the system if you select this option.

Disable HyperSwap

Select this option to prevent a HyperSwap from occurring by command or event.

On Planned HyperSwap Error:

Partition out the failing system(s) and continue swap processing on the remaining system(s)

Select this option to partition out the failing system and continue the swap processing on any remaining systems.

Disable HyperSwap after attempting backout

Select this option to stop the HyperSwap action, and disable the HyperSwap commands or events.

On Unplanned HyperSwap Error:

Partition out the failing system(s) and continue swap processing on the remaining system(s)

Select this option to partition out the failing systems and continue HyperSwap processing on the remaining systems when a new system is added to the sysplex and HyperSwap does not complete. You must restart the system if you select this option.

Disable HyperSwap after attempting backout

Select this option to stop the HyperSwap action, and disable the HyperSwap commands or events.

Fail MM/GC if target is online (CKD only)

Select this option to fail any session commands for a copy set if the session is online and if the target volume in the copy set is visible to a host.

ESS/DS Metro Mirror Options

Select the Metro Mirror options that you want to associate with this session. These options are available only for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, or System Storage DS6000 storage systems in a Metro Mirror session:

Reset Secondary Reserve

Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets being started when a Start command is issued for the session.

Note:

Before enabling the Reset Secondary Reserves option, be aware that this action causes the session to overwrite all data on the target volume.

Fail MM/GC if target is online (CKD only)

Select this option to fail any session commands for a copy set if the session is online and if the target volume in the copy set is visible to a host.

Metro Mirror Suspend Policy

You must select either **Hold I/O after Suspend** or **Release I/O after Suspend** when a replication link suspends. A replication link suspends for various reasons. In some cases, it happens because the replication link is down or the auxiliary storage controller has failed. In other cases, it might be the first failure in an eventual full failure at the primary site.

If your primary goal is zero data loss, specify **Hold I/O after Suspend**. Selecting this option means that no additional updates are made to storage until the entire scope of the failure can be determined. If you select this option, the systems by using that storage become unavailable until a decision is made about how to proceed. It might also require restarting the systems.

If your primary goal is availability, specify **Release I/O after Suspend**. This option enables workloads to continue running after all the replication links have been suspended, maintaining a consistent copy of data at the remote site. If the failure spreads through the primary site after the workloads resume running and results in a site failure, data loss occurs. The I/O requests that are issued after the resume and before the failure at the primary site are lost. The storage at the secondary site remains consistent, however, and you can recover it there.

Hold I/O after Suspend

Select this option to block the application from writing while a consistent copy of the data is maintained on the remote site; however, it does not automatically release the application. This option keeps the source equal to the target. You must use the **Release I/O** command on the session or wait for the Hardware Freeze Timeout timer to expire before the application can continue to write to the source.

Tip: If the **Manage H1-H2 with Hyperswap** option is set on the session, you must issue the **SETHS RESUMEIO** command from z/OS. This command allows write operations to continue to the primary site after a disaster, such as a link failure between the primary and auxiliary storage systems. When the suspension occurs, IOS holds the I/O until the **SETHS RESUMEIO** command is issued and either the ELB timer has expired or the Release I/O command is issued from Tivoli Storage Productivity Center for Replication. The swap does not automatically occur in these cases, so you must decide if there is a disaster or a problem keeping the secondary site consistent. If this is a disaster, issue the **Recover** command and complete a re-IPL, by using the storage from the secondary site. If it is not a disaster, you can issue the **SETHS RESUMEIO** command and **Release I/O** command in order for I/O to again continue to the primary site.

There are situations, however, where the system must access some of the HyperSwap managed volumes, such as the paging volumes, and you might not be able to issue the **SETHS RESUMEIO** command or it might not finish. In these situations, restart the system to enable write operations to the primary site to resume. Restart the system after you issue the **Release I/O** command from Tivoli Storage Productivity Center for Replication or after the ELB timer has expired.

Release I/O after Suspend

Select this option to block writes to the application while a consistent copy of the data is formed on the remote site. This operation is followed immediately by releasing the block so that the application can continue writing to the source. This option allows for little application impact, but causes the source to potentially be different from the target.

This option is the default setting for all new sessions.

DS FlashCopy Options for Role pair H2-I2

The following option is available only for System Storage DS8000 version 4.2, or later.

Persistent

Select this option to keep FlashCopy pairs persistent on the hardware.

Manage H1-H2 with HyperSwap

Select this option to trigger a HyperSwap, redirecting application I/O to the secondary volumes, when there is a failure on the host accessible volumes. For Metro Global Mirror sessions, the Global Mirror portion of the session continues to run uninterrupted.

Tivoli Storage Productivity Center for Replication can manage the H1-H2 sequence of a Metro Mirror or Metro Global Mirror session by using HyperSwap if Tivoli Storage Productivity Center for Replication is running on z/OS, and volumes are attached by a fibre-channel connection.

Notes:

- When this option is selected, the **Suspend H1-H2** command is available only if the Disable HyperSwap option is also selected.
- Setting this option automatically sets the Release IO after suspend Metro Mirror policy.

Disable HyperSwap

Select this option to prevent a HyperSwap from occurring by command or event.

On Configuration Error:

Partition the system(s) out of the sysplex

Select this option to partition out of the sysplex when a new system is added to the sysplex and encounters an error in loading the configuration. Restart the system if you select this option.

Disable HyperSwap

Select this option to prevent a HyperSwap from occurring by command or event.

On Planned HyperSwap Error:

Partition out the failing system(s) and continue swap processing on the remaining system(s)

Select this option to partition out the failing system and continue the swap processing on any remaining systems.

Disable HyperSwap after attempting backout

Select this option to stop the HyperSwap action, and disable the HyperSwap commands or events.

On Unplanned HyperSwap Error:

Partition out the failing system(s) and continue swap processing on the remaining system(s)

Select this option to partition out the failing systems and continue HyperSwap processing on the remaining systems when a new system is

added to the sysplex and HyperSwap does not complete. You must restart the system if you select this option.

Disable HyperSwap after attempting backout

Select this option to stop the HyperSwap action, and disable the HyperSwap commands or events.

Manage H1-H2 with Open HyperSwap

If volumes are attached to an IBM AIX host, Tivoli Storage Productivity Center for Replication can manage the H1-H2 sequence of a Metro Mirror session by using Open HyperSwap. If this option is selected, a failure on the host accessible volumes triggers a swap, which redirects application I/O to the secondary volumes. Only volumes that are currently attached to the host systems that are defined on the Tivoli Storage Productivity Center for Replication Host Systems panel are eligible for Open HyperSwap.

Disable Open HyperSwap

Select this option to prevent a swap from occurring by a command or event while keeping the configuration on the host system and all primary and secondary volumes coupled.

SAN Volume Controller / Storwize V7000 / Storwize V7000 Unified Metro Mirror Options

Select the Metro Mirror options that you want to associate with this session. The following options are available only for storage systems that are in a Metro Mirror Failover/Failback with Practice session:

Incremental

Select this option to set up the relationship for recording changes to the practice volume (H2). All subsequent FlashCopy operations between the intermediate volume and the host volume copy only the data that has changed since the previous FlashCopy operation. Incremental always assumes persistence.

Background Copy Rate for H2-I2

Type the copy rate that the Storage System uses to perform the background copy of the FlashCopy role pair. You can specify a percentage between 0 and 100. The default is 50. Specifying 0 is equivalent to specifying the No-Copy option on a TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, or System Storage DS6000 FlashCopy session.

You can modify this value at any time during the session. If the session is performing a background copy when you change the option, Tivoli Storage Productivity Center for Replication immediately modifies the background copy rate of the consistency group on the storage system. The storage system consistency group immediately starts by using this new rate to complete the background copy.

ESS/DS Global Mirror Options

Select the Global Mirror options that you want to associate with this session. These options are available only for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 storage systems in a Global Mirror session:

Consistency group interval time (sec)

Type how often, in seconds, the Global Mirror or Metro Global Mirror session attempts to form a consistency group. When you reduce this value, it might be possible to reduce the data exposure of the session. A lower value causes the session to attempt to create consistency groups more frequently, which might affect the storage systems (for example, by increasing the processing load and message traffic load).

Reset Secondary Reserve

Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets being started when a Start command is issued for the session.

Note:

Before enabling the Reset Secondary Reserves option, be aware that this action causes the session to overwrite all data on the target volume.

Fail MM/GC if target is online (CKD only)

Select this option to fail any session commands for a copy set if the session is online and if the target volume in the copy set is visible to a host.

DS FlashCopy Options for Role pair H2-I2

The following option is available only for System Storage DS8000 version 4.2, or later.

Persistent

Select this option to keep FlashCopy pairs persistent on the hardware.

SAN Volume Controller / Storwize V7000 / Storwize V7000 Unified Global Mirror Options

Select the Global Mirror options that you want to associate with this session. These options are available only for storage systems that are in a Global Mirror Failover/Failback with Practice session:

Incremental

Select this option to set up the relationship for recording changes to the practice volume (H2). All subsequent FlashCopy operations between the intermediate volume and the host volume copy only the data that has changed since the previous FlashCopy operation. Incremental always assumes persistence.

Background Copy Rate for H2-I2

Type the copy rate that the Storage System uses to perform the background copy of the FlashCopy role pair. You can specify a percentage between 0 and 100. The default is 50. Specifying 0 is equivalent to specifying the No-Copy option on a TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, or System Storage DS6000 FlashCopy session.

You can modify this value at any time during the session. If the session is performing a background copy when you change the option, Tivoli Storage Productivity Center for Replication immediately modifies the background copy rate of the consistency group on the storage system. The storage system consistency group immediately starts by using this new rate to complete the background copy.

XIV Global Mirror Options

Select the Global Mirror options that you want to associate with this session. These options are available only for the XIV systems in a Global Mirror Failover/Failback session:

Recovery point objective threshold (seconds)

Type the number of seconds that you want to set for the recovery point objective (RPO) threshold. RPO represents a measure of the maximum data loss that is acceptable in the event of a failure or unavailability of the master.

If the XIV system determines that the RPO is greater than this value, the session state becomes Severe. You can specify an RPO between 30 and 86400 seconds. The default is 30 seconds.

Synchronization schedule (HH:MM:SS)

Select an interval for the creation of the XIV system synchronization schedule. The XIV system attempts to form consistent points of data by taking automatic snapshots of the volumes in the session at this interval. The default is **Minimum Interval**, which is 20 seconds.

If you select **Never**, synchronization is not scheduled and the XIV system does not create consistency groups. When the XIV system determines that the RPO threshold has been passed, the session state becomes Severe.

ESS/DS Global Mirror Options

Select the Metro Global Mirror options that you want to associate with this session. These options are available only for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, or System Storage DS6000 storage systems in a Metro Global Mirror session:

Consistency group interval time (sec)

Type how often, in seconds, the Global Mirror or Metro Global Mirror session attempts to form a consistency group. When you reduce this value, it might be possible to reduce the data exposure of the session. A lower value causes the session to attempt to create consistency groups more frequently, which might affect the storage systems (for example, by increasing the processing load and message traffic load).

Reset Secondary Reserve

Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets being started when a Start command is issued for the session.

Note:

Before enabling the Reset Secondary Reserves option, be aware that this action causes the session to overwrite all data on the target volume.

Fail MM/GC if target is online (CKD only)

Select this option to fail any session commands for a copy set if the session is online and if the target volume in the copy set is visible to a host.

Metro Mirror Suspend Policy

You must select either **Hold I/O after Suspend** or **Release I/O after Suspend** when a replication link suspends. A replication link

suspends for various reasons. In some cases, it happens because the replication link is down or the auxiliary storage controller has failed. In other cases, it might be the first failure in an eventual full failure at the primary site.

If your primary goal is zero data loss, specify **Hold I/O after Suspend**. Selecting this option means that no additional updates are made to storage until the entire scope of the failure can be determined. If you select this option, the systems by using that storage become unavailable until a decision is made about how to proceed. It might also require restarting the systems.

If your primary goal is availability, specify **Release I/O after Suspend**. This option enables workloads to continue running after all the replication links have been suspended, maintaining a consistent copy of data at the remote site. If the failure spreads through the primary site after the workloads resume running and results in a site failure, data loss occurs. The I/O requests that are issued after the resume and before the failure at the primary site are lost. The storage at the secondary site remains consistent, however, and you can recover it there.

Hold I/O after Suspend

Select this option to block the application from writing while a consistent copy of the data is maintained on the remote site; however, it does not automatically release the application. This option keeps the source equal to the target. You must use the **Release I/O** command on the session or wait for the Hardware Freeze Timeout timer to expire before the application can continue to write to the source.

Tip: If the **Manage H1-H2 with Hyperswap** option is set on the session, you must issue the **SETHS RESUMEIO** command from z/OS. This command allows write operations to continue to the primary site after a disaster, such as a link failure between the primary and auxiliary storage systems. When the suspension occurs, IOS holds the I/O until the **SETHS RESUMEIO** command is issued and either the ELB timer has expired or the Release I/O command is issued from Tivoli Storage Productivity Center for Replication. The swap does not automatically occur in these cases, so you must decide if there is a disaster or a problem keeping the secondary site consistent. If this is a disaster, issue the **Recover** command and complete a re-IPL, by using the storage from the secondary site. If it is not a disaster, you can issue the **SETHS RESUMEIO** command and **Release I/O** command in order for I/O to again continue to the primary site.

There are situations, however, where the system must access some of the HyperSwap managed volumes, such as the paging volumes, and you might not be able to issue the **SETHS RESUMEIO** command or it might not finish. In these situations, restart the system to enable write operations to the primary site to resume. Restart the system after you

issue the **Release I/O** command from Tivoli Storage Productivity Center for Replication or after the ELB timer has expired

Release I/O after Suspend

Select this option to block writes to the application while a consistent copy of the data is formed on the remote site. This operation is followed immediately by releasing the block so that the application can continue writing to the source. This option allows for little application impact, but causes the source to potentially be different from the target.

This option is the default setting for all new sessions.

Manage H1-H2 with HyperSwap

Select this option to trigger a HyperSwap, redirecting application I/O to the secondary volumes, when there is a failure on the host accessible volumes. For Metro Global Mirror sessions, the Global Mirror portion of the session continues to run uninterrupted.

Tivoli Storage Productivity Center for Replication can manage the H1-H2 sequence of a Metro Mirror or Metro Global Mirror session by using HyperSwap if Tivoli Storage Productivity Center for Replication is running on z/OS, and volumes are attached by a Fibre Channel connection.

Notes:

- When this option is selected, the Suspend H1-H2 command is available only if the Disable HyperSwap option is also selected.
- Setting this option automatically sets the Release IO after suspend Metro Mirror policy.

Disable HyperSwap

Select this option to prevent a HyperSwap from occurring by command or event.

On Configuration Error:

Partition the system(s) out of the sysplex

Select this option to partition out of the sysplex when a new system is added to the sysplex and encounters an error in loading the configuration. Restart the system if you select this option.

Disable HyperSwap

Select this option to prevent a HyperSwap from occurring by command or event.

On Planned HyperSwap Error:

Partition out the failing system(s) and continue swap processing on the remaining system(s)

Select this option to partition out the failing system and continue the swap processing on any remaining systems.

Disable HyperSwap after attempting backout

Select this option to stop the HyperSwap action, and disable the HyperSwap commands or events.

On Unplanned HyperSwap Error:

Partition out the failing system(s) and continue swap processing on the remaining system(s)

Select this option to partition out the failing systems and continue HyperSwap processing on the remaining systems when a new system is added to the sysplex and HyperSwap does not complete. You must restart the system if you select this option.

Disable HyperSwap after attempting backout

Select this option to stop the HyperSwap action, and disable the HyperSwap commands or events.

Role pair status and progress

This topic describes the status messages for role pair status and progress.

Tivoli Storage Productivity Center for Replication provides detailed role pair status and progress messages for sessions. The status and progress messages are updated to provide a message that indicates what the session is currently doing at the time. By hovering over a progress bar, you can see specific information about the current action running on the session. Some status messages might include an estimated time to completion for the action in hours and minutes.

Role pair status is not provided for the XIV system Snapshot sessions because role pairs are not used for these sessions.

The status messages appear in the Session Details and Role Pair Details panel.

Table 37. Detailed status messages for Participating and Non-Participating role pairs

Supported session type	Status message
FlashCopy Metro Mirror Global Copy Global Mirror	Starting <i>role_pair_name</i> relationships on the hardware
Metro Mirror Global Mirror	Waiting for all pairs in the role pair <i>role_pair_name</i> to reach state of <i>state</i>
FlashCopy Metro Mirror Global Copy Global Mirror	Terminating all pairs in role pair <i>role_pair_name</i>
FlashCopy Metro Mirror Global Mirror	Recovering all pairs in role pair <i>role_pair_name</i>
Metro Mirror Global Copy Global Mirror	Suspending all pairs in role pair <i>role_pair_name</i>
FlashCopy	Background copy is running for role pair <i>role_pair_name</i>
Global Copy	Waiting for all pairs in role pair <i>role_pair_name</i> to become consistent
Global Copy	Waiting for all pairs in role pairs <i>role_pair_name</i> to complete the initial copy
FlashCopy	Waiting for all pairs in role pairs <i>role_pair_name</i> to complete FRR

Table 37. Detailed status messages for Participating and Non-Participating role pairs (continued)

Supported session type	Status message
Global Mirror	Waiting for all pairs in role pairs <i>role_pair_name</i> to join the Global Mirror session

Viewing session properties

This topic describes how to view session properties.

To view a session's properties, perform the following steps:

1. In the navigation tree, select **Sessions**. The Sessions panel is displayed.
2. Click the name of the session with the properties you want to view.
3. Select **View/Modify Properties** from the **Actions** menu and click **Go**.

Viewing session details

You can view detailed information about a session, including role pairs, error count, whether the session is recoverable, copying progress, session type, and the timestamp.

Perform these steps to view session details:

1. In the navigation tree, select **Sessions**.
2. Select the session that you want to view.
3. Select **View Details** from the **Actions** drop-down menu, and click **Go**.

Viewing additional details for Global Mirror sessions

Additional detail information is available for Global Mirror sessions, including information about the Global Mirror master, consistency groups that have been formed, and data exposure.

Perform these steps to view additional details for Global Mirror sessions:

1. In the IBM Tivoli Storage Productivity Center for Replication navigation tree, select **Sessions**.
2. Select the Global Mirror session that you want to view.
3. Select **View Details** from the **Actions** drop-down menu, and click **Go**.
4. Click the **Global Mirror Info** tab. The following information is displayed on the tab:

Global Mirror Master

Shows the name of the storage system acting as the Global Mirror master.

Last Master Consistency Group Time

Shows the time that the last consistency group was formed

Master Time During Last Query

Shows the time on the master storage device when the query was performed,

Data Exposure

Shows the average exposure to potential data loss in seconds over the query interval.

Session ID

Shows the Global Mirror session ID.

Master State

Shows the state of the master session on the hardware.

Unsuccessful CGs since last successful CG

Shows the number of consistency groups that have failed to form since the last successful consistency group was formed.

CG Interval Time

Shows the interval time between attempts to form a consistency group.

Max Coordination Interval

Shows the extended distance consistency maximum coordination interval.

Max CG Drain Time

Shows the maximum time the consistent set of data is allowed to drain at the remote site before failing consistency group formation.

Unsuccessful CGS/Previous Query

Shows the number of consistency groups and percentage of consistency groups that were unsuccessful since the previous query.

Unsuccessful CGS/Total

Shows the total number of unsuccessful consistency groups and percentage of consistency groups that have failed.

Successful CGS/Previous Query

Shows the number of consistency groups and percentage of consistency groups that were successful since the previous query.

Successful CGS/Total

Shows the total number of successful consistency groups and percentage of consistency groups that have been successful.

Consistency Group Failure Messages

Shows the failure messages that have occurred on the Global Mirror session that prevented the formation of a consistency group.

Data Exposure chart

Shows the data exposure values in seconds for the last 15 minutes or 24 hours.

Highlight Data Exposure

Use the following fields to define a value in seconds for which you want data exposure to be tracked in the **Data Exposure** chart.

Show Data Exposure over

Data exposure that is above the value that is entered in this field is shown in the **Data Exposure** chart.

Show Data Exposure under

Data exposure that is below the value that is entered in this field is shown in the **Data Exposure** chart.

Viewing storage system details

You can view detailed information about storage system, including the name, location, type, vendor, and the status of all connections to the storage system.

Perform these steps to view storage system details:

1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is displayed in the Storage Systems view.
2. Perform one of these steps to view details for a specific storage system:
 - Click the storage system ID.
 - Select the storage system, click **View storage system details** from the **Actions** list, and then click **Go**.

Viewing storage connection details

You can view storage connection details and a list of all storage systems that are located behind the connection.

Perform these steps to view storage connection details:

1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is displayed in the Storage Systems view.
2. Click the **Connections** tab.
3. Perform one of these steps to view details for a specific storage connection:
 - Click the storage connection ID.
 - Select the storage connection, click **View/modify Connection Details** from the actions list, and then click **Go**.

Viewing volume details

You can view information about volumes such as the name of the volume, the capacity of the volume, and the type of volume.

1. In the navigation tree, select **Volumes**.
2. Select a storage system.
3. Depending on the type of storage, do one of the following:
 - a. Select **All IO Groups** or a specific I/O group.
 - b. Select **All Logical Storage Subsystems** or a specific logical storage subsystem.
 - c. Select **All Pools** or a specific pool.
4. Click **Perform Query**. Information about the volumes is displayed in a table.

Viewing logical paths

You can view all logical paths that are defined on an IBM TotalStorage Enterprise Storage Server, IBM System Storage DS8000, or IBM System Storage DS6000 storage system.

Perform one of these procedures to view logical paths:

- From the ESS/DS Paths panel of IBM Tivoli Storage Productivity Center for Replication:
 1. In the navigation tree, select **ESS/DS Paths**. The ESS/DS Paths panel is displayed.
 2. Click the storage system ID to display logical paths for that storage system.
- From the Storage Systems panel:
 1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is displayed in the Storage Systems view.

2. Select an ESS, DS6000, or DS8000 storage system for which you want to view logical paths.
3. Select **View Paths** from the **Select Action** list, and click **Go**. The ESS/DS Paths panel is displayed with a list of defined logical paths.

Viewing console messages

This topic describes how to view the console and messages.

IBM Tivoli Storage Productivity Center for Replication provides detailed information about actions taken by users, errors that occur during normal operation, and hardware error indications.

From the graphical user interface, you can view console messages by selecting **Console** in the navigation tree. You can then click the link for the specific message code to get more information on the message.

You can also get detailed information and help for specific messages in the IBM Tivoli Storage Productivity Center for Replication for System z information center at <http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp>.

Chapter 10. Security

The IBM Tivoli Storage Productivity Center for Replication authentication process uses a configured user registry from either the operating system or Lightweight Directory Access Protocol (LDAP) server. To perform a specific action and manage specific sessions in the IBM Tivoli Storage Productivity Center for Replication GUI or CLI, the user must also have the appropriate authorization. Authorization is granted by assigning a specific role to the user account or user group.

Users and groups

For authentication and authorization, IBM Tivoli Storage Productivity Center for Replication uses users and groups that are defined in a configured user registry on the management server, which is associated with either the local operating system or a Lightweight Directory Access Protocol (LDAP) server.

IBM Tivoli Storage Productivity Center for Replication does not provide the capability to create, update, or delete users or groups in the user registry. To manage users or groups, you must use the appropriate tool associated with the user registry in which the users and groups are stored.

IBM Tivoli Storage Productivity Center for Replication uses roles to authorize users to manage certain sessions and perform certain actions.

For more information about authentication, see information about single sign-on in the IBM Tivoli Storage Productivity Center documentation.

Primary administrative ID

If you switch the authentication method, either from the local operating system to an LDAP server or vice versa, IBM Tivoli Storage Productivity Center for Replication removes all access to existing users and user groups. This occurs because the user IDs might not be on the same local operating system and the LDAP server; however, you must have at least one user ID that can log in to IBM Tivoli Storage Productivity Center for Replication.

When you change the authentication method using Tivoli Integrated Portal, you can specify a primary administrative ID for both local operating system and LDAP authentication. Use this primary administrator to log in to IBM Tivoli Storage Productivity Center for Replication and manually add user IDs requiring access to IBM Tivoli Storage Productivity Center for Replication.

You can log in to both IBM Tivoli Storage Productivity Center for Replication and IBM Tivoli Storage Productivity Center using the primary administrative ID and password.

You cannot use the following characters for the IBM Tivoli Storage Productivity Center for Replication administrative password:

- square brackets ([and])
- semicolon (;)
- backward slash (\)

User roles

A *user role* is a set of privileges that is assigned to a user or user group to allow the user or user group to perform certain tasks and manage certain sessions.

To be assigned to a role, each user or group of users must have a valid user ID or group ID in the user registry on the management server.

Both individual users and a group of users can be assigned to a role. All users in a group are assigned the role of the group. If a user is assigned to one role as an individual and a different role as a member of a group, the user has access to the permissions of the role that has greater access.

Restricting access to sessions prevents unwarranted administrative access. This is especially useful in an open environment, where there can be many storage administrators who are responsible for their servers, applications, databases, file systems, and so on.

IBM Tivoli Storage Productivity Center for Replication provides a set of predefined user roles: monitor, session operator, and administrator.

Monitor

Monitors can view the health and status in the IBM Tivoli Storage Productivity Center for Replication GUI and CLI; however, they cannot modify or perform any commands or actions.

Monitors can view the following information:

- All storage systems and storage system details
- All connections and connection details
- All sessions and session details
- All path information
- Management server status and details

Operator

Session operators can manage sessions to which they have been assigned, including:

- Adding or removing a session. The user ID that created the session is automatically granted access to manage that session.
- Performing actions on an assigned session, such as start, flash, terminate, and suspend.
- Modifying session properties.
- Adding copy sets to a session. The session operator can add volumes to a copy set only when the volume is not protected and not in another session.
- Removing copy sets from a session.
- Adding Peer To Peer Remote Copy (PPRC) paths, and removing paths with no hardware relationships. PPRC paths are a common resource used in IBM Tivoli Storage Productivity Center for Replication sessions and also in an ESS, DS6000, or DS8000 storage-system relationship that is established between two common logical subsystems (LSSs).

Note: The session operator cannot issue a force removal of a path.

Note: A path can also be auto-generated when starting a session.

- Monitoring health and status, including viewing the following information:
 - All storage systems and storage system details
 - All connections and connection details
 - All sessions and session details
 - All path information
 - Management server status and details

Note: Session operators can make changes only to the volumes that they own. They are not able to make changes to volumes being managed by other users.

Administrator

During installation of IBM Tivoli Storage Productivity Center for Replication, the installation wizard requests an ID to use for the initial administrator user ID.

Administrators have unrestricted access. They can manage all sessions and perform all actions associated with IBM Tivoli Storage Productivity Center for Replication, including:

- Granting permissions to users and groups of users.
- Adding or removing a session. The user ID that created the session is automatically granted access manage that session.
- Performing actions on all sessions, such as start, flash, terminate, and suspend.
- Modifying session properties.
- Adding and removing copy sets from a session. The administrator can add volumes to a copy set only when the volume is not protected and not in another session.
- Protecting volumes and removing volume protection.
- Adding or removing storage system connections.
- Modifying connection properties.
- Assigning or changing storage system locations.
- Adding PPRC paths and removing paths with no hardware relationships. PPRC paths are a common resource used in IBM Tivoli Storage Productivity Center for Replication sessions and also in an ESS, DS6000, or DS8000 storage-system relationship that is established between two common logical subsystems (LSSs).

Note: A path can also be auto-generated when starting a session.

- Managing management servers. The standby management server is a common resource that is available to multiple sessions.
- Packaging program error (PE) log files.
- Monitoring health and status, including viewing the following information:
 - All storage systems and storage system details
 - All connections and connection details
 - All sessions and session details
 - All path information
 - Management server status and details

Important: IBM Tivoli Storage Productivity Center supports multiple user roles, including the Superuser role. A superuser can perform all IBM Tivoli Storage Productivity Center functions. For IBM Tivoli Storage Productivity Center

superusers to have full access to IBM Tivoli Storage Productivity Center for Replication, the Superuser group must be added to the IBM Tivoli Storage Productivity Center for Replication and assigned to the Administrator role. Then, you can manage the IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication products by groups, instead of by user IDs.

Note: Administrators cannot revoke their own administrative access rights.

Adding the IBM Tivoli Storage Productivity Center for Replication Administrator role to the IBM Tivoli Storage Productivity Center Superuser group

If you use local operating system or Lightweight Directory Access Protocol (LDAP) authentication, you must add the IBM Tivoli Storage Productivity Center superuser group to IBM Tivoli Storage Productivity Center for Replication with administrator privileges.

Prerequisite: You must have Administrator privileges to perform this action.

Perform these steps to add the superuser group to IBM Tivoli Storage Productivity Center for Replication:

1. Log in to IBM Tivoli Storage Productivity Center as the superuser.
2. In the navigation tree, expand **Administrative Services > Configuration** and select **Role-to-Group Mappings**.
3. Locate the name of the superuser user group.
4. Log in to IBM Tivoli Storage Productivity Center for Replication using a user ID with administrator privileges.
5. In the navigation tree, select **Administration**. The Administration panel is displayed.
6. Click **Add Access**. The Add Access wizard is displayed.
7. Type the name of the superuser group in the **User or group names** field, and click **Next**.
8. Select the name of the superuser group, and click **Next**.
9. Select **Administrator privileges**, and click **Next**.
10. Click **Next** to confirm this action.
11. Click **Finish**.

Granting access privileges for a user

You can assign user roles to an IBM Tivoli Storage Productivity Center for Replication user to grant access privileges to individual sessions and tasks.

Perform the following steps to authorize a user:

1. Create the user ID or group ID if it does not already exist in the user registry, either the operating system of the active management server or Lightweight Directory Access Protocol (LDAP) server.
2. Log in to IBM Tivoli Storage Productivity Center for Replication as a user with administrator privileges.
3. In the navigation tree, select **Administration**. The Administration panel is displayed.

4. Click **Add Access**. The Add Access wizard is displayed.
5. Type the name of the user to whom you want to give access, and click **Next**. The Select Users and Groups panel is displayed.

Tip: You can enter a partial name and use the * wildcard character to represent zero or more characters.

6. Select one or more names from the list of found users.
7. Select the role to associate with this user.
8. If you selected the Operator role, select one or more session that this user can manage, and click **Next**.
9. Click **Next** to confirm this action.
10. Click **Finish**.

Viewing access privileges for a user

You can view a list of all IBM Tivoli Storage Productivity Center for Replication user and their assigned roles. You can also view the assigned sessions for each user.

Perform the following steps to view access privileges for a user:

1. Log in to IBM Tivoli Storage Productivity Center for Replication as a user with administrator privileges.
2. In the navigation tree, select **Administration**. The Administration page is displayed with a list of IBM Tivoli Storage Productivity Center for Replication users and user groups and their associated role.
3. Select the user whose access privileges you want to view.
4. Select **View/Modify Access** from the **Actions** drop-down list, and click **Go**. The View/Modify Access panel is displayed. This panel shows the role assigned to the user and lists the sessions that the user can manage.
5. Click **Cancel**.

Modifying access privileges for a user

You can change the user role and assigned sessions for an IBM Tivoli Storage Productivity Center for Replication user.

Prerequisite: You must have Administrator privileges to perform this action.

Perform the following steps to modify the access privileges for a user:

1. Log in to IBM Tivoli Storage Productivity Center for Replication as a user with administrator privileges.
2. In the navigation tree, select **Administration**. The Administration panel is displayed with a list of users and user groups and their associated role.
3. Select the user whose access privileges you want to view.
4. Select **View/Modify Access** from the **Select Action** drop-down list, and click **Go**. The View/Modify Access panel is displayed. This panel shows the role assigned to the user and lists the sessions that the user can manage.
5. Select the role to associate with this user.
6. If you selected the Operator role, select one or more session that this user can manage and click **Next**.
7. Click **OK**.

Removing access privileges for a user

You can remove access privileges for an IBM Tivoli Storage Productivity Center for Replication user. When you remove access, the user ID cannot access the IBM Tivoli Storage Productivity Center for Replication GUI or run commands from the command line.

Prerequisite: You must have Administrator privileges to perform this action.

Perform the following steps to remove user access:

1. Log in to IBM Tivoli Storage Productivity Center for Replication as a user with administrator privileges.
2. In the navigation tree, select **Administration**. The Administration panel is displayed with a list of users and user groups and their associated role.
3. Select the user from which you want to remove access.
4. Select **Remove Access** from the **Actions** list, and click **Go**.

Appendix. Using the system logger in a Tivoli Storage Productivity Center for Replication for System z environment

The system logger is an IBM z/OS component that provides a logging facility for applications running in a single-system or multisystem sysplex. There are many factors to consider when you are using the system logger in a IBM Tivoli Storage Productivity Center for Replication for System z environment and are using Metro Mirror sessions.

Configuring the system logger for use in the Tivoli Storage Productivity Center for Replication for System z environment

When the system logger is used in an IBM Tivoli Storage Productivity Center for Replication for System z environment, steps must be taken to avoid data consistency issues.

The following situations can lead to data consistency issues when using Tivoli Storage Productivity Center for Replication for System z with the system logger:

- The **Release I/O after Suspend** option has been selected for a Metro Mirror session.
- The system logger couple data sets (CDSs) are not part of the Metro Mirror session. In this situation, the data sets are not frozen even though the related application secondary volumes have been frozen.
- The system logger log streams use coupling facility (CF) structures.
- After a suspend event, the primary site fails and you must recover at the alternate site.

If the secondary disks in the Metro Mirror session are frozen and the workload continues to run using the primary disks, the data on the secondary disks is out of sync with the CF structures or the CDSs. If you attempt to restart the applications using the frozen secondary disks, the restart fails because of this inconsistency. For example, Customer Information Control System (CICS) require a cold start instead of an emergency restart, and transaction backout and handling of in-doubt transactions are not possible.

If **Release I/O after Suspend** has been selected for Metro Mirror sessions, the actions that are shown in the following figure are required.



1. In the system logger policy, all CF log streams must be forced to duplex to staging data sets. The following data sets must be in the same Metro Mirror session:
 - Log stream staging data sets that are direct access storage device (DASD)-only
 - CF log stream data sets
 - All of the offload data sets for both types of log streams
2. Four system logger CDSs must be set up as follows:
 - The primary system logger CDS in Site 1 must be in the same Metro Mirror session.
 - The spare system logger CDS in Site 1 must not be in a Metro Mirror session.
 - The alternate system logger CDS in Site 2 must not be in a Metro Mirror session.
 - The spare system logger CDS in Site 2 must not be in a Metro Mirror session.

Set up all CDS types other than the system logger CDSs as required for Tivoli Storage Productivity Center for Replication for System z. That is, the primary system logger CDS should be in Site 1 and the alternate system logger CDS in Site 2. There should be spare CDSs in both sites. The alternate and spare CDSs should be not be in a Metro Mirror session.

By following the preceding steps, the primary system logger CDS, CF log stream staging, and offload data sets are on volumes in the Metro Mirror session. If a freeze occurs, system logger data will be consistent on the secondary devices. If the reason for the freeze requires that you restart from the secondary devices, you can recover and use this frozen copy of the system logger environment.

Important: Ensure that no CF log streams remain allocated in any coupling facilities that the production systems have access to following a disaster. In this situation, recovery occurs from the mirrored copies of the data. If any log streams are allocated, you must force the connections and ensure the structure is deleted before restarting your production systems.

Reintroducing frozen system logger CDSs into your sysplex

In the event that CDSs become frozen, you can correct the issue that resulted in the freeze and re-introduce the CDS into your sysplex.

Reintroducing CDSs after an unplanned swap

After a suspend event, the secondary disks are frozen and you cannot access the disks. To recover at the secondary site, you must make the disks accessible by using IBM Tivoli Storage Productivity Center for Replication for System z to initiate a recover. The Recover command performs the steps necessary to make the target available as the new primary site. Upon completion of this command, the session is in the Target Available state.

If the active Tivoli Storage Productivity Center for Replication for System z server was located at Site 1, and the system the server was running on failed, you must use your standby server to recover. Issue the Takeover command, before initiating the Recover command.

When the session is in the Target Available state, the systems at Site 2 can be restarted using the Site 2 volumes.

Switching Disks Back to Site 1 After an Unplanned Failover to Site 2

To switch disks back to Site 1, see the information about switching from Site 2 to Site 1 in the following sections.

Reintroducing CDSs after a planned swap

Typically, you perform a planned switch from Site 1 to Site 2 for one of the following reasons:

- The Site 1 disk is temporarily unavailable because of a disruptive disk maintenance action.
- Site 1 is temporarily unavailable in its entirety because of a site maintenance activity.

In these situations, switch the disks to Site 2. When the Site 1 disk is available again, switch back to the Site 1 disk when you have the Site 2-to-Site 1 mirroring in full duplex.

Considerations for a Planned Metro Mirror Swap

When the system logger CDS is part of the Metro Mirror session and you plan to switch your primary disks from Site 1 to Site 2, you must complete the following tasks to release the allocation against the system logger CDS:

1. Switch to the system logger CDS that is not in the Metro Mirror session (that is, make the Site 2 alternate system logger CDS the new primary system logger CDS) by issuing the following command:

```
SETXCF COUPLE,TYPE=LOGR,PSWITCH
```

2. Make the Site 2 spare CDS the new alternate data set by issuing the following command:

```
SETXCF COUPLE,TYPE=LOGR,ACOUPLE=(spare cds in site 2)
```

When you switch back from Site 2 to Site 1, switch the Metro Mirror direction and then perform a CDS switch to return to the normal CDS configuration. After you switch the Metro Mirror session direction, perform the following actions to switch the CDS:

1. Make the primary at Site 1 the alternate by issuing the following command:

```
SETXCF COUPLE,TYPE=LOGR,ACOUPLE=(original primary cds in site 1)
```

2. Make the original primary the primary again using the following command:

```
SETXCF COUPLE,TYPE=LOGR,PSWITCH
```

3. Make the original alternate CDS at Site 2 the alternate again using the following command:

```
SETXCF COUPLE,TYPE=LOGR,ACOUPLE=(original alternate cds in site 2)
```

Considerations for Planned HyperSwap

If you are using Tivoli Storage Productivity Center for Replication for System z planned HyperSwap capability and you have your system logger CDSs mirrored, when swapping disks from Site 1 to Site 2, switch your CDS configuration to use only Site 2 CDSs before running the SWAP command to perform the disk swap. When swapping back to the Site 1 disks, use the normal CDS configuration after the HyperSwap has completed successfully.

Notices

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

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

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

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

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan

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

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

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

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

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

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
MW9A/050
5600 Cottle Road
San Jose, CA 95193
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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

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

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

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

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., 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 www.ibm.com/legal/copytrade.shtml.

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

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

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

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.

Index

A

- about this document xvii
- access privileges
 - granting 216
 - modifying 217
 - removing 218
 - viewing 217
- accessibility features for users with disabilities xi
- accessing the server 78
- active management servers
 - about 6, 85
 - reconnecting 93
 - taking over 93
- adding
 - copy sets to a FlashCopy session 172
 - copy sets to a Global Mirror session 176
 - copy sets to a Metro Global Mirror session 177
 - copy sets to a Metro Mirror session 175
 - copy sets to a Snapshot session 173
 - ESS/DS series logical paths 115
 - using a CSV file 116
 - host system connections 111, 112
 - SNMP managers 94
 - storage connections 104
 - superuser groups 216
- administering
 - users 213
- architecture, product 5
- awareness, site 68, 104, 170

B

- backing up
 - considerations for 79
 - database 81
- Basic HyperSwap
 - about 26, 130
 - session commands 52, 156

C

- changes in this release xix, xxi
- changing the time zone in z/OS 95
- client ports 9, 86, 97
- comments, sending xvii
- communication-failure trap
 - descriptions 90
- configuration change, SNMP trap
 - descriptions 89
- configuring
 - logical paths 115
 - storage systems 108
- connections
 - about 11, 100
 - adding
 - host systems 112

- connections (*continued*)
 - direct 12, 101
 - Hardware Management Console (HMC) 13, 102
 - modifying properties 107
 - removing 106
 - z/OS 13, 102
- considerations, back up and recovery 79
- consistency group 20, 124
 - time interval 20, 124
- copy sets
 - about 15, 119
 - adding 172, 173, 175, 176, 177
 - exporting 82, 181
 - importing 83, 181
 - removing 182
 - supported number of role pairs and volumes 15, 119
- creating
 - FlashCopy session 172
 - Global Mirror session 176
 - Metro Global Mirror session 177
 - Metro Mirror session 175
 - Snapshot session 173

D

- data copying symbols 192
- data exposure 20, 124
- data migration utility 183
- database
 - backing up 81
 - restoring 81
- DB2 78
 - starting 77
 - stopping 77
- disabilities, accessibility features for users with xi
- disabling
 - Metro Mirror heartbeat 181
 - Open HyperSwap 50, 155
 - OpenSwap 50, 155
- disaster recovery, practicing 187
- drop-down lists, limitations of to the sight-impaired xi

E

- enabling, Metro Mirror heartbeat 181
- enhancements to this release xix, xxi
- ESS/DS storage systems
 - adding logical paths 115
 - using a CSV file 116
 - removing logical paths 117
- exporting copy sets 82, 181

F

- failback 14, 119
- failover 14, 119

FlashCopy

- about 27, 131
- creating a session 172
- session commands 52, 157

G

- Global Mirror
 - about 35, 139
 - creating a session 176
 - session commands 57, 161
 - session details, viewing 209
- Global Mirror with Practice, session commands 58, 162
- granting
 - access privileges 216

H

- health, viewing summary of 191
- heartbeat, enabling and disabling 181
- host systems
 - adding a connection to 111, 112
 - managing 111
 - removing 112

I

- IBM Tivoli Storage Productivity Center for Replication
 - server, accessing 78
 - starting 75
 - stopping 76
- IBM Tivoli Storage Productivity Center for Replication Basic Edition for System z overview 1
- IBM Tivoli Storage Productivity Center for Replication for System z overview 1
- IBM Tivoli Storage Productivity Center for Replication Three Site Business Continuity overview 3
- IBM Tivoli Storage Productivity Center for Replication Two Site Business Continuity overview 2
- icons
 - session status 191
 - storage systems status 8, 99
- identifying the version of IBM Tivoli Storage Productivity Center for Replication 79
- importing copy sets 83, 181

J

- journal volume, restoring data from 109

L

- locations
 - modifying for a site 182
 - modifying for a storage system 107
- logical paths
 - adding 115
 - using a CSV file 116
 - configuring 115
 - removing 117
 - viewing 115, 211
- logical storage subsystems
 - viewing volumes 211

M

- management servers
 - about 6, 85
 - reconnecting 93
 - setting as the standby server 91
 - setting up a remote server as the standby server 91
 - state change SNMP trap descriptions 90
 - taking over 93
 - viewing health summary 191
- managing
 - host systems 111
 - storage systems 97
- messages, viewing 212
- Metro Global Mirror
 - about 40, 144
 - creating a session 177
 - session commands 59, 164
- Metro Global Mirror with Practice, session commands 62, 166
- Metro Mirror
 - about 29, 133
 - creating a session 175
 - session commands 54, 158
- Metro Mirror heartbeat, enabling and disabling 181
- Metro Mirror with Practice, session commands 56, 160
- modifying
 - access privileges 217
 - site location 182
 - storage connection properties 107
 - storage system location 107

N

- notices, legal 223

O

- Open HyperSwap
 - disabling 50, 155
- OpenSwap
 - disabling 50, 155
- overview
 - IBM Tivoli Storage Productivity Center for Replication Basic Edition for System z 1

overview (*continued*)

- IBM Tivoli Storage Productivity Center for Replication for System z 1
- IBM Tivoli Storage Productivity Center for Replication Three Site Business Continuity 3
- IBM Tivoli Storage Productivity Center for Replication Two Site Business Continuity 2

P

- paths
 - adding 115
 - using a CSV file 116
 - removing 117
- performing
 - takeover 93
- ports 9, 86, 97
- practice volumes 20, 124, 187
- practicing disaster recovery 187
- Preserve Mirror option, recommendations for 170
- primary server, reinstalling during active session 92
- product architecture 5
- properties
 - session 197
 - viewing session 209
- protected volumes, about 103
- protecting volumes 108
- publications, related xiv

R

- reader feedback, sending xvii
- recommendations for using the Preserve Mirror option 170
- reconnecting the active and standby management servers 93
- recovering, considerations for 79
- refreshing, storage system configuration 108
- related websites xvi
- removing
 - copy sets 182
 - host systems 112
 - logical paths 117
 - sessions 182
 - storage connection 106
 - storage systems 106
 - user access 218
- restoring data from a journal volume 109
- restoring, database 81
- role pairs 19, 124
- role pairs, status messages 208

S

- security 213
- sending comments xvii
- session commands
 - Basic HyperSwap 52, 156
 - FlashCopy 52, 157

session commands (*continued*)

- Global Mirror 57, 161
- Global Mirror with Practice 58, 162
- Metro Global Mirror 59, 164
 - Metro Global Mirror with Practice 62, 166
 - Metro Mirror 54, 158
 - Metro Mirror with Practice 56, 160
 - Snapshot 53, 157
- session state change trap descriptions 88
- session types
 - Basic HyperSwap 26, 130
 - FlashCopy 27, 131
 - Global Mirror 35, 139
 - Metro Global Mirror 40, 144
 - Metro Mirror 29, 133
 - Snapshot 28, 132
- sessions
 - about 14, 119
 - FlashCopy 172
 - Global Mirror 176
 - images 192
 - Metro Global Mirror 177
 - Metro Mirror 175
 - properties 197
 - removing 182
 - Snapshot 173
 - states 194
 - status icons 191
 - viewing 191
 - viewing details 209
 - Global Mirror 209
 - viewing health summary 191
 - viewing properties 209
- setting
 - remote standby management server 91
 - SNMP 94
 - standby management server 91
 - volume protection 108
- sight-impaired xi
- site awareness 68, 104, 170
- site, modifying location of 182
- Snapshot
 - about 28, 132
 - creating a session 173
 - session commands 53, 157
- SNMP
 - adding managers 94
 - setting up 94
 - viewing alerts 191
- SNMP trap descriptions
 - communication-failure 90
 - configuration change 89
 - management servers state change 90
 - session state change 88
 - suspending-event notification 89
- standby management servers
 - about 6, 85
 - ports 9, 86, 97
 - setting 91
 - setting a remote server 91
- starting
 - DB2 77
 - IBM Tivoli Storage Productivity Center for Replication 75
- states, session 194

- status messages, role pairs 208
- status, session icons 191
- stopping
 - DB2 77
 - IBM Tivoli Storage Productivity Center for Replication 76
- storage connections
 - removing 106
 - viewing details 211
- storage systems
 - about 8, 99
 - adding a connection to 104
 - icons 8, 99
 - managing 97
 - modifying connection properties 107
 - modifying location of 107
 - ports 9, 86, 97
 - refreshing the configuration 108
 - removing 106
 - removing a connection to 106
 - supported 8, 99
 - viewing details 210
 - viewing health summary 191
 - viewing volumes 211
- summary of changes in this release xix, xxi
- superuser group
 - adding administrator role to 216
- support websites xvi
- suspending-event notification trap descriptions 89
- symbols
 - data copying 192
 - session status 191
 - volume roles 192

T

- taking over management from the active server 93
- time zone in z/OS, changing the 95
- trademarks 225

U

- user administration 213
 - adding superuser groups 216
 - granting access privileges 216
 - modifying access privileges 217
 - removing user access 218
 - viewing access privileges 217
- users
 - about 69, 213
 - groups 69, 213
 - roles 69, 214
- using practice volumes 187

V

- ver command 79
- version of IBM Tivoli Storage Productivity Center for Replication 79
- viewing
 - access privileges 217
 - error messages 212
 - health summary 191

- viewing (*continued*)
 - logical paths 115, 211
 - session details 209
 - Global Mirror 209
 - session properties 209
 - sessions 191
 - SNMP alerts 191
 - storage connection details 211
 - storage system details 210
 - volume details 211
- volume protection
 - about 103
 - setting 108
- volume roles
 - about 18, 123
 - symbols 192
- volumes
 - practice 20, 124, 187
 - role pairs 19, 124
 - roles 18, 123
 - viewing details 211

W

- websites, related xvi
- what's new in this release xix, xxi

Z

- z/OS System Logger, using with Tivoli Storage Productivity Center for Replication for System z 219
- z/OS, changing the time zone in 95



Product Number: 5698-B30,5698-B31

Printed in USA

SC27-2322-06

