IBM Tivoli Storage Productivity Center for Replication for
System z
Version 5.2.9

User's Guide

IBM

IBM Tivoli Storage Productivity Center for Replication for
System z
Version 5.2.9

User's Guide

IBM

Contents

Tables

About this guide

This guide provides task-oriented administration and user information for IBM®
Tivoli® Storage Productivity Center for Replication for System z®.

The Tivoli Storage Productivity Center for Replication for System z family of
products consists of the following licenses:
- IBM Tivoli Storage Productivity Center for Replication for System z
- IBM Tivoli Storage Productivity Center for Replication Basic Edition for
  System z

Who should read this guide

This publication is for users of Tivoli Storage Productivity Center for Replication
for System z.

Users should be familiar with the following topics:
- Tivoli Storage Productivity Center for Replication for System z copy services
  concepts
- General principles of AIX®, Linux, Windows, and the z/OS® operating systems
- Simple Network Management Protocol (SNMP) concepts
- Storage Area Network (SAN) concepts

Publications

A number of publications are provided with IBM Spectrum Control™.

IBM Tivoli Storage Productivity Center for Replication is no longer installed with
IBM Spectrum Control. For current Tivoli Storage Productivity Center for
Replication users who upgrade to IBM Spectrum Control, Tivoli Storage
Productivity Center for Replication will remain available. See IBM Knowledge
Center for any earlier versions of Tivoli Storage Productivity Center for Replication
publications.

The following section describes how to access these publications online.

Accessing publications online

Information about installing, configuring, using, upgrading, and uninstalling IBM
Spectrum Control and related products is available online.

Use Table 1 on page x to view and download these publications. Translated
documents are available for some products.

Table 1. Locations of publications for IBM Spectrum Control and related products

Product	Online location
IBM Spectrum Control	http://www.ibm.com/support/knowledgecenter/SS5R93/ In the navigation pane, click IBM Spectrum Control for the appropriate release. To obtain PDF documents, in the navigation pane, click **IBM Spectrum Control** > **Printable documentation**.
IBM Copy Services Manager (formerly known as IBM Tivoli Storage Productivity Center for Replication)	http://www.ibm.com/support/knowledgecenter/SSESK4
IBM WebSphere® Application Server	http://www.ibm.com/support/knowledgecenter/SSEQTP/mapfiles/product_welcome_was.html
Jazz™ for Service Management	http://www.ibm.com/support/knowledgecenter/SSEKCU/welcome
IBM System Storage® DS3000, IBM System Storage DS4000® or IBM System Storage DS5000™	http://www.ibm.com/support/knowledgecenter/ Go to the preceding website and go to **System Storage** > **Disk systems**.
IBM System Storage DS6000™	http://www.redbooks.ibm.com/abstracts/sg246781.html The DS6000 series is no longer supported. For information, see the preceding publication.
IBM System Storage DS8000®	http://www.ibm.com/support/knowledgecenter/ Go to the preceding website and go to **System Storage** > **Disk systems** > **Enterprise Storage Servers**. Select the DS8000 series system that you want.
IBM System Storage DS® Open Application Programming Interface publications	http://www.ibm.com/support/docview.wss?uid=ssg1S7001160 Use these publications for information about how to install, configure, and use the CIM agent.
IBM System Storage SAN Volume Controller	http://www.ibm.com/support/knowledgecenter/STPVGU
IBM Storwize® V3500	http://www.ibm.com/support/knowledgecenter/STLM6B/landing/v3500_welcome.html
IBM Storwize V3700	http://www.ibm.com/support/knowledgecenter/STLM5A/landing/V3700_welcome.html
IBM Storwize V5000	http://www.ibm.com/support/knowledgecenter/STHGUJ/landing/V5000_welcome.html
IBM Storwize V7000	http://www.ibm.com/support/knowledgecenter/ST3FR7/landing/V7000_welcome.html
IBM Storwize V7000 Unified	http://www.ibm.com/support/knowledgecenter/ST5q4U/landing/v7000_unified_welcome.html
IBM Scale Out Network Attached Storage (IBM SONAS)	http://www.ibm.com/support/knowledgecenter/STAV45/landing/sonas_151_kc_welcome.html
IBM XIV® Storage System	http://www.ibm.com/support/knowledgecenter/STJTAG

Table 1. Locations of publications for IBM Spectrum Control and related products (continued)

Product	Online location
IBM Spectrum Accelerate™	http://www.ibm.com/support/knowledgecenter/ STZSWD/welcome
IBM Spectrum Scale™	http://www.ibm.com/support/knowledgecenter/ SSFKCN/gpfs_welcome.html
IBM DB2® Database for Linux UNIX and Windows	http://www.ibm.com/support/knowledgecenter/ SSEPGG/welcome
IBM System Storage N series	http://www.ibm.com/systems/storage/network/ index.html For more information about IBM System Storage N series, see http://www.ibm.com/systems/storage/ network/software/.
IBM Systems Director	http://www.ibm.com/systems/software/director/ index.html
VMware vSphere	http://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-pubs.html
IBM Cognos® Business Intelligence	http://www.ibm.com/support/knowledgecenter/ SSEP7J
IBM FlashSystem™ V840 Enterprise Performance Solution	http://www.ibm.com/systems/storage/flash/v840/ overview.html
IBM FlashSystem V9000	http://www.ibm.com/systems/storage/flash/v9000/ overview.html

IBM Redbooks

The IBM Redbooks® are publications about specialized topics.

You can order publications through your IBM representative or the IBM branch office serving your locality. You can also search for and order books of interest to you by visiting the IBM Redbooks home page at http://www.redbooks.ibm.com.

Translation

Translated publications are available from IBM Knowledge Center, which is available in certain translated languages. It is displayed in the language that is appropriate for the browser locale setting.

When a locale does not have a translated version, the information is displayed in English, which is the default language. Translations of the PDFs are available when the information is translated.

Providing feedback about publications

Your feedback is important to help IBM provide the highest quality information.

To provide comments or suggestions about the product documentation, go to the following location for IBM Spectrum Control: ➦ http://www.ibm.com/support/ knowledgecenter/SS5R93/.Click **Add Comment** on the welcome page or on individual topic pages in a specific release. You must have an IBM ID and password to add a comment.

Service Management Connect community

Connect, learn, and share with Service Management professionals: product support technical experts who provide their perspectives and expertise.

Access Service Management Connect at https://www.ibm.com/developerworks/servicemanagement/. Use Service Management Connect in the following ways:

- Become involved with transparent development, an ongoing, open engagement between other users and developers of IBM products. You can access early designs, sprint demonstrations, product roadmaps, and prerelease code.
- Connect one-on-one with the experts to collaborate and network about IBM and the Storage Management community.
- Read blogs to benefit from the expertise and experience of others.
- Use wikis and forums to collaborate with the broader user community.

Conventions used in this guide

Information is given about the conventions that are used in this publication.

This publication uses several conventions for special terms and actions, and for operating system-dependent commands and paths.

The following typeface conventions are used in this publication:

Bold

- Flags that display with text
- Graphical user interface (GUI) elements (except for titles of windows and dialogs)
- Names of keys

Italic

- Variables
- Values that you must provide
- New terms
- Words and phrases that are emphasized
- Titles of documents

monospace

- Commands and command options
- Flags that display on a separate line
- Code examples and output
- Message text
- Names of files and directories
- Text strings that you must type, when they display within text
- Names of Oracle Java™ methods and classes
- HTML and XML tags that display like this, in monospace type

For syntax notations, remember the following details.

- In AIX, the prompt for the root user is #.
- In AIX and Linux, the commands are case-sensitive, so you must type commands exactly as they are shown.

Contacting IBM Software Support

You can contact IBM Software Support by phone, and you can register for support notifications at the technical support website.

- Go to the IBM Spectrum Control technical support website at http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Productivity_Center.

  To receive future support notifications, sign in under **Subscribe to support notifications**. You are required to enter your IBM ID and password. After you are authenticated, you can configure your subscription for IBM Spectrum Control technical support website updates.
- Customers in the United States can call 1-800-IBM-SERV (1-800-426-7378).
- For international customers, go to the IBM Spectrum Control technical support website to find support by country. Expand **Contact support** and click **Directory of worldwide contacts**.

You can also review the *IBM Software Support Handbook*, which is available at http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html.

The support website offers extensive information, including a guide to support services; frequently asked questions (FAQs); and documentation for all IBM Software products, including Redbooks and white papers. Translated documents are also available for some products.

When you contact IBM Software Support, be prepared to provide identification information for your company so that support personnel can readily assist you. Company identification information might also be needed to access various online services available on the website. See "Reporting a problem."

Reporting a problem

Provide the IBM Support Center with information about the problems that you report.

Have the following information ready when you report a problem:
- The IBM Tivoli Storage Productivity Center for Replication for System z version, release, modification, and service level number.
- The communication protocol (for example, TCP/IP), version, and release number that you are using.
- The activity that you were doing when the problem occurred, listing the steps that you followed before the problem occurred.
- The exact text of any error messages.

Chapter 1. Product overview

IBM Tivoli Storage Productivity Center for Replication for System z manages copy services in storage environments. Copy services are features that are used by storage systems such as IBM System Storage DS8000 to configure, manage, and monitor data-copy functions. Copy services include IBM FlashCopy®, Metro Mirror, Global Mirror, and Metro Global Mirror.

Tivoli Storage Productivity Center for Replication for System z runs on System z and uses Fibre Channel connection (FICON®) or TCP/IP communications to connect to and manage count key data (CKD) volumes. Tivoli Storage Productivity Center for Replication for System z uses TCP/IP to manage fixed-block architecture (FBA) volumes.

You can use Tivoli Storage Productivity Center for Replication for System z to complete the following data replication tasks:
- Plan for replication when provisioning storage
- Keep data on multiple related volumes consistent across storage systems in a planned or unplanned outage
- Monitor and track replication operations
- Automate the mapping of source volumes to target volumes
- Reduce downtime of critical applications

Tivoli Storage Productivity Center for Replication for System z product family

The Tivoli Storage Productivity Center for Replication for System z family of products consists of the following products: Tivoli Storage Productivity Center for Replication Basic Edition for System z and Tivoli Storage Productivity Center for Replication for System z.

These products provide the following features:
- Tivoli Storage Productivity Center for Replication Basic Edition for System z provides high availability for a single site if a disk storage system failure occurs.
- Tivoli Storage Productivity Center for Replication for System z provides high availability and disaster recovery for multiple sites.

Tivoli Storage Productivity Center for Replication Basic Edition for System z

Tivoli Storage Productivity Center for Replication Basic Edition for System z provides copy services management for the storage systems and session types that are listed in the following table.

Storage system	Session type
System Storage DS8000	• Basic HyperSwap • FlashCopy
System Storage DS6000	• Basic HyperSwap • FlashCopy

Storage system	Session type
TotalStorage Enterprise Storage Server® Model 800	• Basic HyperSwap • FlashCopy
SAN Volume Controller	• FlashCopy
Storwize V3500, Storwize V3700, and Storwize V7000	• FlashCopy
Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)	• FlashCopy
The XIV	• Snapshot

Tivoli Storage Productivity Center for Replication for System z

Tivoli Storage Productivity Center for Replication for System z provides copy services management for the storage systems and session types that are listed in the following table.

Storage system	Session type
System Storage DS8000	• Basic HyperSwap • FlashCopy • Metro Mirror Single Direction • Metro Mirror Failover/Failback • Metro Mirror Failover/Failback with Practice • Global Mirror Single Direction • Global Mirror Either Direction with Two-Site Practice • Global Mirror Failover/Failback • Global Mirror Failover/Failback with Practice • Metro Global Mirror • Metro Global Mirror with Practice • Metro Mirror - Metro Mirror (This session type is available only for storage systems with a microcode level that supports single source to multi-target relationships. To determine whether you can use this session type, refer to the System Storage DS8000 documentation for the microcode level that you are using.)
System Storage DS6000	• Basic HyperSwap • FlashCopy • Metro Mirror Single Direction • Metro Mirror Failover/Failback • Metro Mirror Failover/Failback with Practice • Global Mirror Single Direction • Global Mirror Either Direction with Two-Site Practice • Global Mirror Failover/Failback • Global Mirror Failover/Failback with Practice

Storage system	Session type
TotalStorage Enterprise Storage Server Model 800	• Basic HyperSwap • FlashCopy • Metro Mirror Single Direction • Metro Mirror Failover/Failback • Metro Mirror Failover/Failback with Practice • Global Mirror Single Direction • Global Mirror Either Direction with Two-Site Practice • Global Mirror Failover/Failback • Global Mirror Failover/Failback with Practice • Metro Global Mirror (only H1 site) • Metro Global Mirror with Practice (only H1 site)
SAN Volume Controller	• FlashCopy • Metro Mirror Single Direction • Metro Mirror Failover/Failback • Metro Mirror Failover/Failback with Practice • Global Mirror Single Direction • Global Mirror Failover/Failback • Global Mirror Failover/Failback with Practice • Global Mirror Failover/Failback with Change Volumes
Storwize V3500	• FlashCopy
Storwize V3700 and Storwize V7000	• FlashCopy • Metro Mirror Single Direction • Metro Mirror Failover/Failback • Metro Mirror Failover/Failback with Practice • Global Mirror Single Direction • Global Mirror Failover/Failback • Global Mirror Failover/Failback with Practice • Global Mirror Failover/Failback with Change Volumes
Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)	• FlashCopy • Metro Mirror Single Direction • Metro Mirror Failover/Failback • Metro Mirror Failover/Failback with Practice • Global Mirror Single Direction • Global Mirror Failover/Failback • Global Mirror Failover/Failback with Practice • Global Mirror Failover/Failback with Change Volumes
The XIV	• Snapshot • Metro Mirror Failover/Failback • Global Mirror Failover/Failback

Key components of Tivoli Storage Productivity Center for Replication for System z

Tivoli Storage Productivity Center for Replication for System z consists of several key components.

Tivoli Storage Productivity Center for Replication for System z server
> The Tivoli Storage Productivity Center for Replication for System z server is the control point for the sessions, storage systems, hosts, and other components that are managed by Tivoli Storage Productivity Center for Replication for System z.

Database
> A single database instance serves as the repository for all Tivoli Storage Productivity Center for Replication for System z data.
>
> An embedded database is configured automatically with Tivoli Storage Productivity Center for Replication for System z.

GUI	The graphical user interface (GUI) is used to manage all Tivoli Storage Productivity Center for Replication for System z functions.

CLI	The command-line interface (CLI) is used to issue commands for key Tivoli Storage Productivity Center for Replication for System z functions.

Interfaces for Tivoli Storage Productivity Center for Replication for System z

Tivoli Storage Productivity Center for Replication for System z 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 for System z contains the following GUI components:

Menu bar
> Use the menu bar to complete tasks such as navigating the GUI, identifying the ID for the user who is logged on to the GUI, logging out of the GUI, and finding help and other product information.

> **Navigation**
> > The menu bar contains items that open individual pages in the GUI. For example, if you click **Sessions**, the Session page opens. If you click **Storage** > **Storage Systems**, the Storage Systems page opens. To view the items that are in a menu, hover over the menu.
> >
> > You can open pages in the GUI by clicking the following items in the menu bar:
> >
> > **Overview**
> > > Click this item to open the Overview page. This page shows a status summary for all sessions, storage systems, host systems, and management servers that Tivoli Storage Productivity Center for Replication for System z is managing.
> > >
> > > The Overview page is displayed when you log on to the GUI.

Sessions
>
> Click this item to open the Sessions page. Use this page to create and manage sessions.

Storage or Storage > Storage Systems
>
> Click either of these items to open the Storage Systems page. Use this page to create and manage storage system connections and to configure the use of the System Storage DS8000 Easy Tier® heat map transfer function.

Storage > Host Connections
>
> Click this item to open the Host Connections page. Use this page to create and manage connections to host systems.

Storage > Volumes
>
> Click this item to open the Volumes page. Use this page to view the volumes that are associated with a storage system.

Paths or Paths > ESS/DS Paths
>
> Click either of these items to open the ESS/DS Paths page. Use this page to view and manage paths between source logical subsystems and target logical subsystems for TotalStorage Enterprise Storage Server Model 800, System Storage DS6000, and System Storage DS8000 storage systems.

Console
>
> Click this item to open the **Console** window. Use this window to view detailed information about actions that are taken by users, errors that occur during normal operation, and hardware error indications.

Settings or Settings > Advanced Tools
>
> Click either of these items to open the Advanced Tools page. Use this page to create a log file package for diagnostic purposes, change the automatic refresh rate of the GUI, and enable or disable the Metro Mirror heartbeat.

Settings > Management Servers
>
> Click this item to open the Management Servers page. Use this page to manage the active and standby management servers in a high-availability relationship.

Settings > Administration
>
> Click this item to open the Administration page. Use this page to view and manage the users and user groups for Tivoli Storage Productivity Center for Replication for System z and their access privileges.

If one or more of the components that are shown on the correlating GUI page have an Unknown, Severe, or Warning status, one of the following icons is displayed next to the item in the menu bar:

- Unknown status
- Severe status
- Warning status

If there are multiple components on a page and the components have different statuses, the highest level status icon is displayed next to the item in the menu bar. Unknown is the highest level status and Warning is lowest.

User ID
The user ID shows who is logged on to the GUI. An icon next to the user ID indicates the user role. To view the role, hover over the icon.

To log out of the GUI, hover over the user ID and click **Log Out**.

Help icon
The **Help** icon provides the links to the following items:
- Online help for the active page, window, or notebook in the GUI.
- Production documentation.
- Information about the product.

Breadcrumb
The breadcrumb shows your location in the GUI. You can click an element in the breadcrumb to jump to the corresponding location in the GUI.

For example, when you view the details page for a session, the following breadcrumb is displayed: **Sessions** > *Session Name*. To return to the Sessions page, click **Sessions** in the breadcrumb.

Content area
The content area displays a page that correlates to the component that you selected in the menu bar.

Dynamic images
The GUI provides dynamic images that show the status of the sites and role pairs that are in the session. The images present a pictorial view of information such as the status, progress, and direction of the replication between sites and whether volumes are consistent between sites. Consistent volumes are blue while inconsistent volumes are gray.

Data copying symbols indicate the type of copy that occurs between the sites and the status of the copy. The direction that the symbol is displayed in the GUI depends on the direction of the copy. For example, the following image represents a FlashCopy session. The FlashCopy operation occurs from the H1 volumes to the T1 volumes as indicated by the lightning bolt symbol. The symbol is gray because the FlashCopy operation is inactive and the T1 volumes are gray because they are not consistent.

Command-line interface

You can use the Tivoli Storage Productivity Center for Replication for System z CLI by using the **csmcli** utilities. You can use the CLI directly or as a script to automate functions.

For security, the CLI runs only on the Tivoli Storage Productivity Center for Replication for System z management server, which is the central point of control for managing data replication. You can run the CLI remotely by using a remote-access utility, for example, Secure Shell (SSH) or Telnet.

For on Windows systems, you can specify remote access to Linux or AIX terminals if you enabled Telnet on your Windows server.

Related reference:

The 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.

Key concepts

To help you understand and use Tivoli Storage Productivity Center for Replication effectively, an overview of key concepts is provided.

Management server

The *management server* is a server that has Tivoli Storage Productivity Center for Replication installed. The management server provides a central point of control for managing data replication.

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.

For more information about the storage systems that you can use with Tivoli Storage Productivity Center for Replication, see "Storage systems" on page 39.

Host systems

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

A connection to a host system is required in the following scenarios:
- A connection to a z/OS host system is required if you want to enable z/OS features such as HyperSwap® and hardened freeze in Tivoli Storage Productivity Center for Replication sessions.
- A connection to an AIX host system is required if you want to use the Open HyperSwap feature. This feature enables the automatic swap of input/output (I/O) to the volumes on a secondary site when a failure occurs when I/O is written to the primary site. For the swap to occur, the storage systems that contain the volumes must be attached to an AIX host and the AIX host must be connected to Tivoli Storage Productivity Center for Replication. The Open HyperSwap feature is optional for the Metro Mirror with Failover/Failback session type.

You can add both AIX and z/OS host connections in Tivoli Storage Productivity Center for Replication regardless of the operating system on which the application is installed. For example, you can add a connection to a z/OS host in a Tivoli Storage Productivity Center for Replication instance that is running on Windows, Linux, or AIX.

For more information about the host systems and how to manage them in Tivoli Storage Productivity Center for Replication, see Chapter 5, "Managing host systems," on page 55.

Ports

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

Tivoli Storage Productivity Center for Replication

Port numbers are defined during the installation of Tivoli Storage Productivity Center for Replication. To determine the port number for the ports that are used for Tivoli Storage Productivity Center for Replication, refer to the following table.

Table 2. Where to find Tivoli Storage Productivity Center for Replication port numbers

Tivoli Storage Productivity Center for Replication component	Location of port information
Web browser	To view the port number for the web browser, open the `bootstrap.properties` file in the *path_prefix*`/opt/Tivoli/RM/wlp/usr/servers/replicationServer/` directory.

Table 2. Where to find Tivoli Storage Productivity Center for Replication port numbers (continued)

Tivoli Storage Productivity Center for Replication component	Location of port information
Standby management server	To view the port number for each management server, click **Management Servers** in the navigation panel of the Tivoli Storage Productivity Center for Replication GUI. You can also run the `lshaservers` command from the command-line interface to view the ports for the management servers. **Important:** The port number for the standby management server must be the same on both the management server and the standby management server in a high-availability relationship. If you change the standby management server port number on either the management server or the standby management server, you must also change the port number on the other server.
Client and CLI port	To view the port number for the client, click **About** in the navigation panel of the Tivoli Storage Productivity Center for Replication GUI. You can also run the `whoami` command from the command-line interface to view the port for the client. **Important:** The client port number must be the same on both the management server and the standby management server in a high-availability relationship. If you change the client port number on either the management server or the standby management server, you must also change the port number on the other server.

Storage system ports

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

Table 3. Storage system default ports

Storage System	Connection Type	Port
• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000	Direct Connection	2433
• System Storage DS8000	Hardware Management Console Connection	1751 and 1750
• SAN Volume Controller • Storwize V3500 • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified	Direct Connection	443 and 22

Table 3. Storage system default ports (continued)

Storage System	Connection Type	Port
• The XIV system	Direct Connection	7778

Ensure that your network configuration is set up so that Tivoli Storage Productivity Center for Replication can send outgoing TCP/IP packets to the storage controllers.

Because there are typically multiple applications that are running on the management server, it is possible that port conflicts might arise if other applications attempt to use the same ports that Tivoli Storage Productivity Center for Replication uses. 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. You can accept the default port or enter another port number.

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.

Sessions

A *session* manages a specific type of data replication against a specific set of volumes. The type of data replication that is associated with the session determines the replication actions that are available for the session. For example, the options for FlashCopy sessions are different from the options for Metro Mirror sessions.

For more information about sessions, see "Sessions" on page 63.

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*.

For more information about copy sets, see "Copy sets" on page 63.

Volume roles

A *volume role* is 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 at the primary site.

For more information about volume roles, see "Volume roles" on page 67.

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 indicated by an arrow. For example, H1>H2 denotes that H1 is the source and H2 is the target.

For more information about role pairs, see "Role pairs" on page 68.

Practice volumes

You can use a *practice volume* to practice for disaster recovery without interrupting current data replication. Practice volumes are available in Metro Mirror, Global Mirror, and Metro Global Mirror sessions.

You can practice disaster recovery actions without using practice volumes. However, without practice volumes, you cannot continue to copy data changes between volumes while you are practicing.

For more information about practice volumes, see "Practice volumes" on page 69.

Consistency groups

For Global Mirror and Metro Global Mirror sessions, 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 were updated to preserve write order and are recoverable.

For more information about consistency groups, see "Consistency groups" on page 69.

Failover and failback operations

If a session has *failover* and *failback* capabilities, you can move the operation of Tivoli Storage Productivity Center for Replication from one site to another, and change the direction of the copy without completing a full copy.

Failover is the process of temporarily switching production to a backup facility by following a scheduled maintenance period or a disaster at a production site. A failover operation can be followed by a failback operation, which is the process of returning production to its original location and resuming the copy between the two sites.

For more information about failover and failback operations, see "Failover and failback operations" on page 69.

Session commands

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

For descriptions of the session commands for each session type, see "Session commands" on page 106.

Session types

Tivoli Storage Productivity Center for Replication provides several methods to replicate data. The type of data replication that is associated with a session is known as the *session type*.

The session types that are available depend on the storage system type and the edition of IBM Spectrum Control or Tivoli Storage Productivity Center for Replication that you are using as shown in this topic.

Basic HyperSwap

Basic HyperSwap is a copy services solution for IBM z/OS version 1.9 and later. Basic HyperSwap provides high availability of data if a storage system failure occurs.

The Basic HyperSwap session type enables Tivoli Storage Productivity Center for Replication to manage Basic HyperSwap by using Metro Mirror or Metro Global Mirror replication.

Basic HyperSwap sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 4. Basic HyperSwap

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication Basic Edition for System z • Tivoli Storage Productivity Center for Replication for System z	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000

FlashCopy

FlashCopy replication creates a point-in-time copy of the data on a source volume to a target volume. Data that existed on the target volume is replaced by the copied data.

FlashCopy sessions are available for the following IBM Spectrum Control and Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 5. FlashCopy

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication Basic Edition for System z • Tivoli Storage Productivity Center for Replication for System z • All editions of IBM Spectrum Control	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3500 • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)

Snapshot

Snapshot replication 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.

Snapshot sessions are available for the following IBM Spectrum Control and Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 6. Snapshot

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication Basic Edition for System z • Tivoli Storage Productivity Center for Replication for System z • All editions of IBM Spectrum Control	The XIV system

Metro Mirror Single Direction

Metro Mirror is synchronous replication that operates between two sites that are up to 300 km apart. The source volumes are on one storage system and the target volumes are on another storage system.

Metro Mirror replication maintains identical data in both the source and target volumes. When a write operation is issued to a source volume, the changes are propagated to the target volume before the write operation finishes processing.

With the Metro Mirror Single Direction session type, Metro Mirror replication is available only from the primary site.

Metro Mirror Single Direction sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 7. Metro Mirror Single Direction

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of IBM Spectrum Control	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)

Metro Mirror Failover/Failback

Metro Mirror Failover/Failback replication enables you to switch the direction of the data flow so that you can use your secondary site as your production site. You can then copy changes that are made at the secondary site back to the primary site.

Metro Mirror Failover/Failback sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 8. Metro Mirror Failover/Failback

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of IBM Spectrum Control	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage) • The XIV system

Metro Mirror Failover/Failback with Practice

Metro Mirror Failover/Failback with Practice replication combines Metro Mirror Failover/Failback and FlashCopy capabilities to provide a point-in-time copy of the data on the secondary site. This session type provides volumes that you can use to practice for disaster recovery without losing your disaster recovery capability.

Metro Mirror Failover/Failback with Practice sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 9. Metro Mirror Failover/Failback with Practice

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of IBM Spectrum Control	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)

Global Mirror Single Direction

Global Mirror is asynchronous replication that operates between two sites that are over 300 km apart. The source volumes are on one storage system and the target volumes are on another storage system.

Global Mirror replication maintains identical data in both the source and target volumes. When a write operation is issued to a source volume, the changes are typically propagated to the target volume a few seconds after the data is written to the source volume.

With the Global Mirror Single Direction session type, Global Mirror replication is available only from the primary site.

Global Mirror Single Direction sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 10. Global Mirror Single Direction

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of IBM Spectrum Control	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)

Global Mirror Either Direction with Two-Site Practice

Global Mirror Either Direction with Two-Site Practice replication enables you to run Global Mirror replication from either the primary or secondary site. This session type provides volumes on the primary and secondary site that you can use to practice for disaster recovery without losing your disaster recovery capability.

Global Mirror Either Direction with Two-Site Practice sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 11. Global Mirror Either Direction with Two-Site Practice

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of IBM Spectrum Control	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000

Global Mirror Failover/Failback

Global Mirror Failover/Failback replication enables you to switch the direction of the data flow so that you can use your secondary site as your production site. You can then copy changes that are made at the secondary site back to the primary site.

Global Mirror Failover/Failback sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 12. Global Mirror Failover/Failback

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of IBM Spectrum Control	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage) • The XIV system

Global Mirror Failover/Failback with Change Volumes

Global Mirror Failover/Failback with Change Volumes replication provides the same capabilities as the Global Mirror Failover/Failback session type. However, Global Mirror Failover/Failback with Change Volumes sessions also provide the option of enabling or disabling the use of change volumes. Change volumes are available in SAN Volume Controller and Storwize storage systems.

Global Mirror Failover/Failback with Change Volumes sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 13. Global Mirror Failover/Failback with Change Volumes

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of IBM Spectrum Control	• SAN Volume Controller Version 6.3 and later • Storwize V3700 Version 6.4.1 and later • Storwize V7000 Version 6.3 and later • Storwize V7000 Unified Version 1.3 and later (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)

Global Mirror Failover/Failback with Practice

Global Mirror Failover/Failback with Practice replication combines Global Mirror Failover/Failback and FlashCopy capabilities to provide a point-in-time copy of the data on the secondary site. This session type provides volumes that you can use to practice for disaster recovery without losing your disaster recovery capability.

Global Mirror Failover/Failback with Practice sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 14. Global Mirror Failover/Failback with Practice

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of IBM Spectrum Control	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)

Metro Global Mirror

Metro Global Mirror replication provides continuous, remote data replication that operates between three sites that are 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. Metro Global Mirror replication enables you to switch the direction of the data flow so that you can use your secondary or tertiary site as your production site.

Metro Global Mirror sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 15. Metro Global Mirror

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of IBM Spectrum Control	• TotalStorage Enterprise Storage Server Model 800 (only H1 site) • System Storage DS8000

Metro Global Mirror with Practice

Metro Global Mirror with Practice replication combines Metro Mirror, Global Mirror, and FlashCopy capabilities to provide a point-in-time copy of the data on the tertiary site. This session type provides volumes that you can use to practice for disaster recovery without losing your disaster recovery capability.

Metro Global Mirror with Practice sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 16. Metro Global Mirror with Practice

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of IBM Spectrum Control	• TotalStorage Enterprise Storage Server Model 800 (only H1 site) • System Storage DS8000

Related concepts:

"Basic HyperSwap" on page 78
Basic HyperSwap is a copy services solution for z/OS version 1.9 and later. It provides high availability of data if a disk storage system failure occurs. Basic HyperSwap does not replace a disaster recovery solution. If a session is suspended but the suspend operation was not caused by a HyperSwap trigger, the session is not frozen to ensure that it is consistent.

"FlashCopy" on page 79
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.

"Snapshot" on page 80
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.

"Metro Mirror" on page 81
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 in one storage system and the target is in another storage system.

"Global Mirror" on page 88
Global Mirror is asynchronous replication that operates between two sites that are over 300 km apart. The source volumes are on one storage system and the target volumes are on another storage system. Global Mirror replication maintains identical data in both the source and target volumes.

"Metro Global Mirror" on page 93
Metro Global Mirror is a method of continuous, remote data replication that operates between three sites that are 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.

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. When determining whether to use the Metro Mirror heartbeat, analyze your business needs. If you disable the Metro Mirror heartbeat, data might become inconsistent. If you enable the Metro Mirror heartbeat and a freeze occurs on the storage system, your applications cannot perform write operations during the freeze.

For more information about the Metro Mirror heartbeat, see "Metro Mirror heartbeat" on page 151.

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. Site awareness prevents a session relationship from being established in the wrong direction.

For more information about site awareness, see "Site awareness" on page 45.

Users and groups

For authentication and authorization, Tivoli Storage Productivity Center for Replication uses users and groups.

For information about users and groups for Tivoli Storage Productivity Center for Replication for System z, see the IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide.

For information about users and groups for Tivoli Storage Productivity Center for Replication, see "Users and groups" on page 203.

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.

IBM Tivoli Storage Productivity Center for Replication provides the following predefined user roles: monitor, session operator, and administrator.

For more information users roles, see "User roles" on page 204.

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 Tivoli Storage Productivity Center for Replication

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

Starting the Tivoli Storage Productivity Center for Replication server

The commands that are required to start the Tivoli Storage Productivity Center for Replication server depend on the operating system.

The Tivoli Storage Productivity Center for Replication server is started by default after the application is installed.

Starting Tivoli Storage Productivity Center for Replication on Windows

To start Tivoli Storage Productivity Center for Replication, issue the following command from the command prompt:

```
TPC_install_directory\scripts\startTPCReplication.bat
```

Where the variable *TPC_install_directory* is the installation directory. The default directory is `C:\Program Files\IBM\TPC`.

Starting Tivoli Storage Productivity Center for Replication on AIX or Linux

To start Tivoli Storage Productivity Center for Replication, issue the following command from the command line:

```
/TPC_install_directory/scripts/startTPCReplication.sh
```

Where *TPC_install_directory* is the installation directory. The default directory is `/opt/IBM/TPC`.

Starting Tivoli Storage Productivity Center for Replication on z/OS:
To start Tivoli Storage Productivity Center for Replication for System z, complete the following steps:

1. Ensure that the Tivoli Storage Productivity Center for Replication for System z HFS is mounted on the UNIX System Services. The root production directory is *path_prefix*`/opt/Tivoli/RM`.
2. Issue the command from the IBM z/OS System Display and Search Facility (SDSF) panel to start Tivoli Storage Productivity Center for Replication for System z, as shown in the following example. The address spaces that are shown might be different in your environment.
   ```
   /START IWNSRV
   ```

The initiator is IWNSRV and the job name is IWNSRV. For information about the IWNSRV job, see the IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide.

Stopping Tivoli Storage Productivity Center for Replication server

The commands that are required to stop the Tivoli Storage Productivity Center for Replication server depend on the operating system.

Stopping Tivoli Storage Productivity Center for Replication on Windows

To stop the Tivoli Storage Productivity Center for Replication server, issue the following command from the command prompt:

`TPC_install_directory\scripts\stopTPCReplication.bat`

Where the variable *TPC_install_directory* is the installation directory. The default directory is `C:\Program Files\IBM\TPC`.

Stopping Tivoli Storage Productivity Center for Replication on AIX or Linux

To stop Tivoli Storage Productivity Center for Replication, issue the following command from the command line:

`/TPC_install_directory/scripts/stopTPCReplication.sh`

Where *TPC_install_directory* is the installation directory. The default directory is `/opt/IBM/TPC`.

Stopping Tivoli Storage Productivity Center for Replication on z/OS

To stop Tivoli Storage Productivity Center for Replication for System z, complete the following steps:

1. Issue a command from the IBM z/OS System Display and Search Facility (SDSF) panel to stop Tivoli Storage Productivity Center for Replication for System z, as shown in the following example command. The address space that is shown might be different in your environment.

   `/STOP IWNSRV`

   Where **`/STOP IWNSRV`** stops the server address space.
2. Issue the **`/D A,L`** command to show the processes that are running and verify that the stop command completed. If the IWNSRV address space is displayed, issue the **`/CANCEL`** command for the address space.

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

You can verify that the Tivoli Storage Productivity Center for Replication server is running in several ways.

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

- Start the Tivoli Storage Productivity Center for Replication GUI or command-line interface shell. If either of these methods are successful, then the Tivoli Storage Productivity Center for Replication server is running.

- View the `csmTrace.log` file in the `path_prefix/opt/Tivoli/RM/wlp/usr/servers/replicationServer/logs/CSM/` directory. If the `csmTrace.log` file is updated regularly and is increasing in size, then Tivoli Storage Productivity Center for Replication is running.

Starting the Tivoli Storage Productivity Center for Replication for System z GUI

You can start the Tivoli Storage Productivity Center for Replication for System z graphical user interface (GUI) by opening a web browser and entering a web address for the Tivoli Storage Productivity Center for Replication for System z logon page.

Before you start Tivoli Storage Productivity Center for Replication for System z, ensure that you are using a supported web browser. For a list of web browsers that you can use with IBM Spectrum Control, see IBM Spectrum Control interoperability matrix. In the **Agents, Servers and Browser** column, click the version of IBM Spectrum Control that is installed on your system. On the next page, click **Web Browsers** to find the web browsers that you can use.

To start the Tivoli Storage Productivity Center for Replication for System z GUI, use these steps:

1. Start a web browser and enter the following address in the address field. The address field is case-sensitive.

   `https://hostname:port/CSM`

   In the preceding address field, specify these values:

   hostname

   The Tivoli Storage Productivity Center for Replication for System z server. You can specify the host name as an IP address or a Domain Name System (DNS) name.

   port

   The port number for Tivoli Storage Productivity Center for Replication for System z. You can find the port number in the `bootstrap.properties` file in the `path_prefix/opt/Tivoli/RM/wlp/usr/servers/replicationServer/` directory on the Tivoli Storage Productivity Center for Replication for System z server.

2. From the Tivoli Storage Productivity Center for Replication for System z logon page, enter your user ID and password, and click the **Log in** arrow. The Tivoli Storage Productivity Center for Replication for System z GUI opens in the browser.

   If you are logging on for the first time after installation: Ensure that you use the user ID that was specified in the IWNRACF job and the password for that user.

Identifying the version of Tivoli Storage Productivity Center for Replication

The Tivoli Storage Productivity Center for Replication version determines the features that are available in the product.

To view the version of Tivoli Storage Productivity Center for Replication from the GUI, hover over the **?** icon in the menu bar, and click **About**.

To view the version of Tivoli Storage Productivity Center for Replication from the command-line interface, use the **-ver** command.

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 back up 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 you export copy sets, Tivoli Storage Productivity Center for Replication creates a backup of the session 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. 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. 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_HHmmssSSS*`.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 following location:

Tivoli Storage Productivity Center for Replication
> *TPC_install_directory*`\wlp\usr\servers\replicationServer\database\`
> `backup\`

Tivoli Storage Productivity Center for Replication for System z
> *path_prefix*`/opt/Tivoli/RM/wlp/usr/servers/replicationServer/`
> `database/backup/`

You can change the default location by editing the **db.backup.location** property in the `rmserver.properties` file, which is in the following location:

Tivoli Storage Productivity Center for Replication
> *TPC_install_directory*`\wlp\usr\servers\replicationServer\properties\`

Tivoli Storage Productivity Center for Replication for System z
> *path_prefix*`/opt/Tivoli/RM/wlp/usr/servers/replicationServer/`
> `properties/`

The backup file contains the Tivoli Storage Productivity Center for Replication database data at the time the backup was completed. 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.

Create a new backup file in the following situations:

- After you change 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 or Set Production to Site command is 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 are 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 are restored and a subsequent Start command to the session creates new relationships. Therefore, you must remove the deprecated copy sets before you issue the Start command.

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

Backing up the Tivoli Storage Productivity Center for Replication database

The database contains information about the storage systems, sessions, copy sets, user administration, and management server configuration in Tivoli Storage Productivity Center for Replication.

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

```
csmcli> mkbackup
```

You must have Administrator privileges to run this command.

By default, the backup file is stored in the following location:

Tivoli Storage Productivity Center for Replication
> *TPC_install_directory*\wlp\usr\servers\replicationServer\database\backup\

Tivoli Storage Productivity Center for Replication for System z
> *path_prefix*/opt/Tivoli/RM/wlp/usr/servers/replicationServer/database/backup/

You can change the default location by editing the **db.backup.location** property in the rmserver.properties file, which is in the following location:

Tivoli Storage Productivity Center for Replication
> *TPC_install_directory*\wlp\usr\servers\replicationServer\properties\

Tivoli Storage Productivity Center for Replication for System z
> *path_prefix*/opt/Tivoli/RM/wlp/usr/servers/replicationServer/properties/

Restoring the Tivoli Storage Productivity Center for Replication database

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

Restoring the database does not require Administrator privileges. However, you must be able to access the files on the Tivoli Storage Productivity Center for Replication server that are listed in the procedure.

Complete these steps to restore the Tivoli Storage Productivity Center for Replication database from a backup database:

1. Stop Tivoli Storage Productivity Center for Replication on the active management server as described in "Stopping Tivoli Storage Productivity Center for Replication server" on page 22.

2. Copy the *path_prefix*/opt/Tivoli/RM/wlp/usr/servers/replicationServer/ database/ directory and all of its contents to another location. You can delete the copy of the database directory after you restore the database.

3. Delete the *path_prefix*/opt/Tivoli/RM/wlp/usr/servers/replicationServer/ database/csmdb/ directory and all of its contents.

4. Uncompress the backup database in the *path_prefix*/opt/Tivoli/RM/wlp/usr/ servers/replicationServer/database/ directory. The backup database is in the *path_prefix*/opt/Tivoli/RM/wlp/usr/servers/replicationServer/database/ backup/ directory.

5. Issue the following commands to change the permissions of the csmdb directory:
   - chgrp -R *tpcr_id* csmdb
   - chmod -R u+rwx csmdb
   - chmod -R g+rwx csmdb
   - chmod -R o+r csmdb

   Where *tpcr_id* is the Tivoli Storage Productivity Center for Replication user that is defined in the IWNRACF job or a user that has authority to run z/OS UNIX System Services. For more information about the IWNRACF job, see the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

6. Restart Tivoli Storage Productivity Center for Replication on the active management server as described in "Starting the Tivoli Storage Productivity Center for Replication server" on page 21.

7. Resolve any changes that occurred since the backup was created.

8. Start the Tivoli Storage Productivity Center for Replication sessions by using the appropriate start commands. The start commands re-establish the relationship between the volume pairs and synchronize data on those volumes.

9. If you have a standby management server, re-establish 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.

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

1. In the menu bar, click **Sessions**.

2. On the Sessions page, select the session for which you want to export copy sets.

3. From the **Session Actions** list, select **Export** > **Export Copy Sets**. The Export Copy Sets window displays the status of the export and a link to the exported file if the export completed.

4. Right-click the link and save the file to a local system.

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 menu bar, click **Sessions**.

2. On the Sessions page, select the session for which you want to import copy sets.
3. From the **Session Actions** list, select **View/Modify** > **Add Copy Sets**. The Add Copy Sets wizard is displayed.
4. Select **Use a CSV file to import copy sets**.
5. Click **Browse**, select the CSV file, and then click **Next**.
6. Verify that the matching results were successful, and then click **Next**.
7. Select the copy sets that you want to add, and then click **Next**.
8. Confirm the number of copy sets that you want to create, and then click **Next**.
9. View the information on the Results page of the wizard, and then click **Finish**.

Verifying the Java heap size

The Java heap size determines the number of volume pairs that Tivoli Storage Productivity Center for Replication can manage.

Verifying the Java heap size for Tivoli Storage Productivity Center for Replication

The default Java heap size is 1024 MB, which supports fewer than 75,000 pairs. If you increase the Java heap size by 256 MB, support increases by 25,000 pairs. You can continue to increase the Java heap size until you reach the maximum heap size.

For more information about the maximum heap size, see the information about sizing the Java heap size in the IBM SDK, Java Technology Knowledge Center at http://www.ibm.com/developerworks/java/jdk/docs.html.

For more information about how to set up the Java heap size, see the *WebSphere Application Server V8.5 Liberty profile* technote at http://www.ibm.com/support/docview.wss?uid=swg21596474.

Verifying the Java heap size for Tivoli Storage Productivity Center for Replication for System z

The default Java heap size is 768 MB, which supports fewer than 50,000 pairs. If you increase the Java heap size by 256 MB, support increases by 25,000 pairs. You can continue to increase the Java heap size until you reach the maximum heap size.

For more information about the maximum heap size, see the information about sizing the Java heap size in the IBM SDK, Java Technology Knowledge Center at http://www.ibm.com/developerworks/java/jdk/docs.html.

For more information about how to set up the Java heap size, see the *WebSphere Application Server V8.5 Liberty profile* technote at http://www.ibm.com/support/docview.wss?uid=swg21596474.

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.

Tivoli Storage Productivity Center for Replication Basic Edition for System z does not support active and standby management servers. If you are using this product, the menu item **Management Servers** is not displayed on the **Settings** menu of the GUI.

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.

Problems might occur during the synchronization process. For example, the management server relationship might become disconnected for a number of reasons, including a connectivity problem or a problem with the alternate server.

If a problem occurs during the synchronization process, the alternate server database is restored to its original state before the synchronization process began. Issue the Reconnect command to restore synchronization.

Completing a takeover on the standby management server

If you must complete 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 complete 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 severs 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.

Setting exceptions for ephemeral port numbers

For management servers that use an Internet Protocol version 4 (IPv4) connection, the port that is shown on the Management Servers page of the Tivoli Storage Productivity Center for Replication GUI is the *listener* port. This port was set when Tivoli Storage Productivity Center for Replication was installed.

However, when the management server initiates a connection to standby server, the management server uses an *ephemeral* port.

An ephemeral port is a temporary port number that is assigned for the duration the connection. When the connection terminates, the ephemeral port is available for reuse. When the management server reconnects to the standby server, a new ephemeral port number is assigned.

If you are using a firewall, you must create exceptions that enable the management server to connect to the standby server by using outbound ephemeral ports.

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 able to communicate with the storage systems by using a z/OS connection. Therefore, another connection must be made by 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 overwrite the same data in the Tivoli Storage Productivity Center for Replication database.

Tivoli Storage Productivity Center for Replication Basic Edition for System z does not support active and standby management servers. If you are using this product, the menu item **Management Servers** is not displayed on the **Settings** menu of the GUI.

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
* Target Available
* Suspending
* (Metro Global Mirror only) SuspendedH2H3
* (Metro Global Mirror only) SuspendedH1H3

An SNMP trap is also sent when a recovery point objective (RPO) threshold is exceeded for a role pair that is in the session.

Table 17. 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.
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.

Table 17. Session state change traps (continued)

Object ID (OID)	Description
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.
1.3.6.1.4.1.2.6.208.0.26	The recovery point objective for the role pair of *X* in session *Y* has passed the warning threshold of *Z* seconds.
1.3.6.1.4.1.2.6.208.0.27	The recovery point objective for the role pair of *X* in session *Y* has passed the severe threshold of *Z* seconds.

Related reference:

"Session states" on page 167
You can view the health and status of a session in the Tivoli Storage Productivity
Center for Replication GUI.

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 18. 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

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

Suspending-event notification SNMP traps indicate that a session has transitioned
to a Severe status due to an unexpected error.

Table 19. 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 20. 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 21. 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).
1.3.6.1.4.1.2.6.208.0.15	The IBM Tivoli Storage Productivity Center for ReplicationServer 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).

Table 21. Management Servers state-change traps (continued)

Object ID (OID)	Description
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

You can define the server on which you are currently logged in as the standby management server.

Attention: When you set a management server as the standby server, all the information on that management server is cleared. The operation cannot be undone.

Important: The port number for the standby management server must be the same on both the management server and the standby management server in a high-availability relationship. If you change the standby management server port number on either the management server or the standby management server, you must also change the port number on the other server.

To set the local management server as the standby server, complete the following steps:
1. In the menu bar, click **Settings** > **Management Servers**.
2. On the Management Servers page, from the **Select Action** list, select **Set this Server as Standby**.
3. Enter the domain name or IP address of the server that you want to use as the active management server.
4. Click **OK** to connect to the active server. The server on which you are logged in is now the standby server.

Setting a remote management server as the standby server

You can define a remote server as the standby management server.

Attention: When you set a management server as the standby server, all the information on that management server is cleared. The operation cannot be undone.

Important: The port number for the standby management server must be the same on both the management server and the standby management server in a

high-availability relationship. If you change the standby management server port number on either the management server or the standby management server, you must also change the port number on the other server.

To set a remote management server as the standby server, complete the following steps:

1. In the menu bar, click **Settings** > **Management Servers**.
2. On the Management Servers page, from the **Select Action** list, select **Define Standby**.
3. Enter the domain name or IP address of the server that you want to use as the standby management server. Log in to the standby management server by entering the user name and password.
4. Click **OK** to connect to the standby management server.

Applying maintenance to an active management server

If an active management server develops a problem during an active session, you can reinstall Tivoli Storage Productivity Center for Replication on the server without affecting the session.

Follow these steps to reinstall Tivoli Storage Productivity Center for Replication on the active management server.

1. If the heartbeat is enabled, disable it:
   a. In the menu bar, click **Settings** > **Advanced Tools**.
   b. On the Advanced Tools page, 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 Tivoli Storage Productivity Center for Replication on Server 1.
5. Reinstall 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 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 Tivoli Storage Productivity Center for Replication on Server 2, skip steps 9 and 10.
9. Uninstall Tivoli Storage Productivity Center for Replication from Server 2.
10. Reinstall 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.

To reconnect the active and standby management servers, complete the following steps:

1. In the menu bar, click **Settings** > **Management Servers**.
2. On the Management Servers page, from the **Select Action** list, select **Reconnect**.

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 on to the Tivoli Storage Productivity Center for Replication GUI that is running on the standby management server.
3. In the menu bar, click **Settings** > **Management Servers**.
4. On the Management Servers page, from the **Select Action** list, select **Takeover**.
5. To reestablish high-availability, complete 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 "Setting up a standby management server" on page 34.
   - 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 `rmserver.properties` file. This file is in the *path_prefix*`/opt/Tivoli/RM/wlp/usr/servers/replicationServer/properties/` directory.

Adding SNMP managers

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

Tivoli Storage Productivity Center for Replication uses management information base (MIB) files to provide a textual description of each SNMP alert that is sent by 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 files are on the installation DVD in the `replication\CSM-Client\etc` directory. Follow the directions that are provided by your SNMP manager application to configure it to use the MIB files.

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 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 port number for the standby management server must be the same on both the management server and the standby management server in a high-availability relationship. If you change the standby management server port number on either the management server or the standby management server, you must also change the port number on the other server.

1. Open the `rmserver.properties` file in the *path_prefix*`/opt/Tivoli/RM/wlp/usr/servers/replicationServer/properties/` directory.
2. Modify the port number for the following property:

   `communications.haPort=`*port_number*
3. Restart Tivoli Storage Productivity Center for Replication. You must restart Tivoli Storage Productivity Center for Replication to activate property changes. Properties are not synchronized between the Tivoli Storage Productivity Center for Replication management servers and must be done on each 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 the management server and the standby management server in a high-availability relationship. If you change the client port number on either the management server or the standby management server, you must also change the port number on the other server.

1. Open the `rmserver.properties` file in the *path_prefix*`/opt/Tivoli/RM/wlp/usr/servers/replicationServer/properties/` directory.

2. Modify the port number for the following property:

   `communications.port=`*port_number*

3. Open the `repcli.properties` file in the *path_prefix*`/opt/Tivoli/RM/cli/` directory.

4. Modify the port number for the following property:

   `port=`*port_number*

5. Restart Tivoli Storage Productivity Center for Replication. You must restart Tivoli Storage Productivity Center for Replication to activate property changes. Properties are not synchronized between the Tivoli Storage Productivity Center for Replication management servers and must be maintained on each Tivoli Storage Productivity Center for Replication management server.

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.

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. You can then omit storage systems for which Tivoli Storage Productivity Center for Replication does not 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 IBM z/OS connections.

You can use the following storage systems:
- IBM TotalStorage Enterprise Storage Server Model 800
- IBM System Storage DS6000
- IBM System Storage DS8000
- IBM System Storage SAN Volume Controller
- IBM Storwize family
- IBM Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)
- IBM XIV Storage System

SAN Volume Controller and Storwize systems can virtualize various storage systems. Although Tivoli Storage Productivity Center for Replication does not support all storage systems, you can manage other storage systems through a single SAN Volume Controller or Storwize cluster interface. Tivoli Storage Productivity Center for Replication connects directly to the SAN Volume Controller or Storwize clusters.

You can define a location for each storage system and for each site in a session. When you are 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.

Easy Tier heat map transfer

The IBM System Storage DS8000 Easy Tier heat map transfer function transfers heat map information from a source storage system to one or more target storage

systems. Each target storage system then generates volume migration plans based on the heat map data and the physical configuration of the storage system.

This process ensures that the performance characteristics of the target storage systems are consistently updated to reflect the performance characteristics of the source storage system.

The Easy Tier heat map transfer function is available for System Storage DS8000 Release 7.1 and later.

You can enable heat map transfers for System Storage DS8000 storage systems from the IBM Tivoli Storage Productivity Center for Replication graphical user interface (GUI). The storage systems must meet the following requirements:

- The source and target storage systems must be connected to Tivoli Storage Productivity Center for Replication by using a Hardware Management Console (HMC) connection.
- The Easy Tier heat map transfer function must be enabled on the source and target storage systems. If the function is disabled on either the source or target storage system, an error is displayed when the next transfer is attempted.

To enable, disable, and configure heat map transfers, see "Configuring the use of the Easy Tier heat map transfer function" on page 52.

Important:

When you enable or disable the use of the Easy Tier heat map transfer function in Tivoli Storage Productivity Center for Replication, the function is not enabled or disabled on the storage systems that are connected to Tivoli Storage Productivity Center for Replication. The configuration options that you set for Easy Tier heat map transfer in Tivoli Storage Productivity Center for Replication are used only by Tivoli Storage Productivity Center for Replication.

Do not manage heat map transfers for the same System Storage DS8000 storage systems from multiple Tivoli Storage Productivity Center for Replication servers. Doing so causes transfers to occur more frequently than the default transfer schedule, which is 12 hours after the previous transfer.

Easy Tier heat map transfer log files

Easy Tier heat map transfer log files are in the *path_prefix*/opt/Tivoli/RM/wlp/ usr/servers/replicationServer/logs/ethmt/log/ directory. If a problem related to the Easy Tier heat map transfer function occurs, you can send these log files to IBM Software Support for troubleshooting purposes.

Storage connections

You must create a connection from the IBM Tivoli Storage Productivity Center for Replication management server to each storage system.

The following types of connection are available. The connection type that you can use depends on the storage system type.

Direct Connection
    This connection type is available for all storage systems other than IBM System Storage DS8700 or later storage systems or System Storage DS8000 storage systems that are on an Internet Protocol version 6 (IPv6) network.

Hardware Management Console (HMC) Connection
This connection type is available for System Storage DS8700 or later storage systems and System Storage DS8000 storage systems that are on an IPv4 or IPv6 network.

IBM z/OS Connection
This connection type is available for System Storage DS8000, System Storage DS6000, and IBM TotalStorage Enterprise Storage Server Model 800 storage systems. If the storage system is connected to a z/OS system on which Tivoli Storage Productivity Center for Replication is installed, a native connection is automatically added. If the storage system is connected to another z/OS system, you can use an IP connection to connect to the system.

A single storage system can be connected by using multiple connections for redundancy. For example, you can connect to a System Storage DS8000 storage system by using an HMC connection and a z/OS connection.

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 and the standby server configuration.

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 TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, or System Storage DS6000 storage system by 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.

Direct connection

By using a TCP/IP connection, the Tivoli Storage Productivity Center for Replication management server can connect directly to certain storage systems. The TCP/IP connection is required to discover a system's configuration (such as LSSs, volumes, volume size, and format), issue queries, and receive asynchronous events.

The Tivoli Storage Productivity Center for Replication management server can connect directly with the following systems:
- TotalStorage Enterprise Storage Server
- DS6000
- DS8000
- SAN Volume Controller
- IBM Storwize family
- Storwize V7000 Unified
- The XIV system

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 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 clusters 0 and 1:
- IP addresses or domain names
- Ports
- User names
- Passwords

SAN Volume Controller and Storwize systems can virtualize various storage systems. Although Tivoli Storage Productivity Center for Replication does not support all storage systems, you can manage other storage systems through a single SAN Volume Controller or Storwize cluster interface. Tivoli Storage Productivity Center for Replication connects directly to the SAN Volume Controller or Storwize clusters.

When you add a direct connection to a cluster to the Tivoli Storage Productivity Center for Replication configuration, specify the cluster IP address of the cluster. 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 storage system user name must have privileges to maintain SSH keys. For information about troubleshooting Secure Shell connections to the storage system, 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. 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 Tivoli Storage Productivity Center for Replication. The default value is two minutes for CKD 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. In addition, the user ID must have Copy Services access to any volumes that are in resource groups on the DS8000 storage system.
- Although you can use either an IP address or a host name to connect to an HMC, the use of a host name ensures that connection to the HMC is not lost if the IP address changes.
- You cannot modify the primary IP address or host name for the HMC in Tivoli Storage Productivity Center for Replication after you add the connection. If you want to change the IP address or host name, you must add a new connection to the HMC and then remove the old connection.
- 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 connection

This connection type is available for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server Model 800 storage systems. If the storage system is connected to an IBM z/OS system on which Tivoli Storage Productivity Center for Replication is installed, a native connection is automatically added. If the storage system is connected to another z/OS system, you must you use an IP connection to connect to the z/OS system.

To add a z/OS connection, the system must be listed on the Host Connections page and in the Connected state. You can add a connection to a z/OS system from Tivoli Storage Productivity Center for Replication that is running on any operating system. For example, you can add a connection to a z/OS system in Tivoli Storage Productivity Center for Replication that is running on Windows, Linux, or AIX. Or, you can add a connection in Tivoli Storage Productivity Center for Replication that is running on z/OS to another z/OS system.

The z/OS connection is used to issue replication commands and queries for attached count key data (CKD) volumes over an existing Fibre Channel network and to receive asynchronous events. When a storage system is added to Tivoli Storage Productivity Center for Replication through a z/OS connection, all CKD volumes that are attached to the Tivoli Storage Productivity Center for Replication management system are added to the Tivoli Storage Productivity Center for Replication configuration. CKD volumes that are not attached to the Tivoli Storage Productivity Center for Replication z/OS management server are not added to the Tivoli Storage Productivity Center for Replication configuration through the z/OS connection.

Ensure that all volumes in the logical storage subsystem (LSS) that you want to manage through a z/OS connection are attached to a z/OS system. 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 Tivoli Storage Productivity Center for Replication to properly manage queries to the hardware. For example, if there are two CKD volumes in an LSS, and one volume is attached to Tivoli Storage Productivity Center for Replication by using a z/OS connection and the other is attached through a direct

connection, Tivoli Storage Productivity Center for Replication issues commands to both volumes over the Fibre Channel network. However, commands issued to the direct connection volume fail, and Tivoli Storage Productivity Center for Replication shows that the copy set that contains that volume has an error.

Tip: It is recommended that you create both a Hardware Management Console (HMC) and direct connection for CKD volumes to allow for greater storage accessibility.

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 CKD volumes and fixed block volumes are added to the Tivoli Storage Productivity Center for Replication configuration.

The Metro Mirror heartbeat

The Metro Mirror heartbeat is not supported through a z/OS connection. To use the Metro Mirror heartbeat, the storage systems must be added by using an HMC or direct connection. If the Metro Mirror heartbeat is enabled and a storage system is added through a direct connection and a z/OS connection, a suspend results if the direct connection becomes disconnected because there is no heartbeat through the z/OS connection.

Removing connections to a z/OS system

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

- If you remove direct and HMC connections first, the fixed-block and non-attached CKD volumes that are attached through these connection types are removed from the Tivoli Storage Productivity Center for Replication configuration.
- The remaining CKD volumes that are attached through the z/OS host connection remain in the Tivoli Storage Productivity Center for Replication configuration until the z/OS host connection is removed.
- If you remove the z/OS host connection first and there is an HMC or direct connection to volumes, those volumes are not removed from the Tivoli Storage Productivity Center for Replication configuration.

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 must add connections to storage systems before you can use Tivoli Storage Productivity Center for Replication to manage data replication.

To add a storage system connection, you must belong to the Administrator role.

Adding a connection to a System Storage DS8000 storage system by using an HMC connection

You must use a Hardware Management Console (HMC) connection to connect to a System Storage DS8000 storage system if the system is a DS8700 or later or is on an Internet Protocol version 6 (IPv6) network.

You can also optionally use an HMC connection if the storage system is on an IPv4 network.

Complete the following steps to add a storage system connection:

1. In the menu bar, click **Storage** > **Storage Systems**.
2. On the Storage Systems page, click **Add Storage Connection**.
3. Click the **DS8000** icon.
4. On the Connection page of the Add Storage System wizard, click **HMC** and complete the following fields.

   Primary HMC
   > Specify the following information about the primary HMC that is used to connect to the storage system.

   > **IP Address/Domain Name**
   > > Enter the IP address or host name for the primary HMC. Although either an IP address or a host name are acceptable, the use of a host name ensures that connection to the HMC is not lost if the IP address changes.

   > > You cannot modify the primary IP address or host name for the HMC in Tivoli Storage Productivity Center for Replication after you add the connection. If you want to change the IP address or host name, you must add a new connection to the HMC and then remove the old connection.

   > **Username**
   > > Enter the user name for the HMC.

   > **Password**
   > > Enter the password for the HMC.

   Secondary HMC (Optional)
   > Specify information about the secondary HMC, if applicable. A secondary HMC is used for redundancy.

   > Both HMCs must be configured identically and must have the same user name and password.

   > **IP Address/Domain Name**
   > > Enter the IP address or host name for the secondary HMC.

5. Click **Next** and follow the instructions in the wizard.

Adding a connection to a System Storage DS8000 storage system by using a direct connection

You can connect to a System Storage DS8000 storage system by using a direct connection if the system is earlier than a DS8700 and is on an Internet Protocol version 4 (IPv4) network.

Direct connection requires that each of the storage system clusters contains an Ethernet card.

Complete the following steps to add a storage system connection:

1. In the menu bar, click **Storage** > **Storage Systems**.
2. On the Storage Systems page, click **Add Storage Connection**.
3. Click the **DS8000** icon.
4. On the Connection page of the Add Storage System wizard, click **Direct Connect** and complete the following fields.

    Cluster 0
    > Specify the following information about connection to cluster 0:

    > **IP Address/Domain Name**
    >> Enter the IP address or host name for the cluster.

    > **Port**   Enter the port number for the cluster. The default value is 2433.

    > **Username**
    >> Enter the user name for the cluster.

    > **Password**
    >> Enter the password for the cluster.

    Cluster 1
    > Specify the following information about connection to cluster 1. The port, user name, and password for cluster 0 are automatically entered in the corresponding fields for cluster 1. You can edit this information, if required for your environment.

    > **IP Address/Domain Name**
    >> Enter the IP address or host name for the cluster.

    > **Port**   Enter the port number for the cluster. The default value is 2433.

    > **Username**
    >> Enter the user name for the cluster.

    > **Password**
    >> Enter the password for the cluster.

5. Click **Next** and follow the instructions in the wizard.

Adding a connection to a System Storage DS6000 or TotalStorage Enterprise Storage Server storage system

You can connect to a System Storage DS6000 or TotalStorage Enterprise Storage Server storage system by using a direct connection to the storage system.

Complete the following steps to add a storage system connection:

1. In the menu bar, click **Storage** > **Storage Systems**.
2. On the Storage Systems page, click **Add Storage Connection**.
3. Click the **DS6000** or **ES800** icon.
4. On the Connection page of the Add Storage System wizard, enter the following information for the storage system clusters:

    Cluster 0
    > Specify the following information about connection to cluster 0:

    > **IP Address/Domain Name**
    >> Enter the IP address or host name for the cluster.

    > **Port**   Enter the port number for the cluster. The default value is 2433.

Username

Enter the user name for the cluster.

Password

Enter the password for the cluster.

Cluster 1

Specify the following information about connection to cluster 1. The port, user name, and password for cluster 0 are automatically entered in the corresponding fields for cluster 1. You can edit this information, if required for your environment.

IP Address/Domain Name

Enter the IP address or host name for the cluster.

Port Enter the port number for the cluster. The default value is 2433.

Username

Enter the user name for the cluster.

Password

Enter the password for the cluster.

5. Click **Next** and follow the instructions in the wizard.

Adding a connection to a SAN Volume Controller or Storwize storage system

You can connect to a SAN Volume Controller or Storwize storage system by using a direct connection to the storage system.

Complete the following steps to add a storage system connection:

1. In the menu bar, click **Storage** > **Storage Systems**.
2. On the Storage Systems page, click **Add Storage Connection**.
3. Click the **SAN Volume Controller**; **Storwize Family** for storage systems in the IBM Storwize family; or **Storwize V7000 Unified** icon.
4. On the Connection page of the Add Storage System wizard, enter the following information for the storage system cluster:

Cluster IP / Domain Name

Enter the IP address or host name of the storage system node.

Username

Enter the user name for the node.

Password

Enter the password for the node.

5. Click **Next** and follow the instructions in the wizard.

Adding a connection to an XIV

You can connect to a IBM XIV Storage System by using a direct connection to the storage system.

Complete the following steps to add a storage system connection:

1. In the menu bar, click **Storage** > **Storage Systems**.
2. On the Storage Systems page, click **Add Storage Connection**.
3. Click the **XIV** icon.

4. On the Connection page of the Add Storage System wizard, enter the following information for a node on the storage system. When you specify the information for one node, the remaining nodes are discovered automatically.

IP Address/Domain Name
　　　Enter the IP address or host name for a node in the storage system.

Username
　　　Enter the user name for the node.

Password
　　　Enter the password for the node.

5. Click **Next** and follow the instructions in the wizard.

Adding a connection to a z/OS system

You can connect to storage systems that are connected to an IBM z/OS system regardless of the operating system on which Tivoli Storage Productivity Center for Replication is installed.

If a z/OS system is defined on the Host Connections page, you can add a connection to a z/OS from any operating system. For example, you can add a connection to a z/OS system in Tivoli Storage Productivity Center for Replication that is running on Windows, Linux, or AIX.

Complete the following steps to add a z/OS system connection:
1. In the menu bar, click **Storage** > **Storage Systems**.
2. On the Storage Systems page, click **Add Storage Connection**.
3. Click the **z/OS Direct Connection** icon.
4. On the Connection page of the Add Storage System wizard, select the storage systems that you want to add.
5. Click **Next** and follow the instructions in the wizard.

Removing a storage system

You can remove a storage system from the 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 menu bar, click **Storage** > **Storage Systems**.
2. On the Storage Systems page, 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. From the **Select Action** list, select **Remove Storage System**.
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 menu bar, click **Storage** > **Storage Systems**.
2. On the Storage Systems page, 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 menu bar, click **Storage** > **Storage Systems**.
2. On the Storage Systems page, click the **Connections** tab.
3. Click the storage connection ID.

4. Modify the appropriate settings to match the settings for the storage system.
5. Click **Apply**.

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 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 menu bar, click **Storage** > **Storage Systems**.
2. On the Storage Systems page, select the storage system for which you want to refresh the configuration.
3. From the **Select Action** list, select **Refresh Configuration**.

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 menu bar, click **Storage** > **Storage Systems**.
2. On the Storage Systems page, click **Volume Protection**. The Volume Protection wizard is opened.
3. The Volume Protection wizard, select a storage system and complete the following steps:
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 (J*x*) is the source volume.
   - The host volume (H*x* or I*x* if using a session with Practice capabilities) is the target volume.
   - *x* is the site the **Recover** command was issued to.

Configuring the use of the Easy Tier heat map transfer function

You can enable Easy Tier heat map transfers for System Storage DS8000 storage systems that are connected to IBM Tivoli Storage Productivity Center for Replication from the Tivoli Storage Productivity Center for Replication graphical user interface (GUI).

The heat map transfer configuration options that are available in Tivoli Storage Productivity Center for Replication include adding storage systems for which you want to transfer heat map data and enabling or disabling the transfer of Easy Tier heat map data for those storage systems.

1. In the menu bar, click **Storage** > **Storage Systems**.
2. Click the **Easy Tier Heat Map Transfer** tab.
3. The following configuration options are available:

   Status Indicates whether the transfer of Easy Tier heat map data is enabled or disabled.

   Enable Transfer
   Click to enable the transfer of Easy Tier heat map data, and then click **Yes** in the confirmation window. Easy Tier heat map data is copied from the storage systems that are listed on this page to the target storage systems that are associated with those storage systems.

   Disable Transfer
   Click to disable the transfer of Easy Tier heat map data, and then click

Yes in the confirmation window. The storage systems that are listed on this page remain displayed with a status of **Inactive**.

Add Storage Systems
Click to add System Storage DS8000 storage systems that you want to include in the Easy Tier heat map transfer configuration. If there are available systems, a list of storage systems is displayed. Select the storage systems from this list. Only those storage systems that meet the following requirements are displayed:

- The storage system must be version 7.1 or later.
- The storage system must be connected to IBM Tivoli Storage Productivity Center for Replication by using a Hardware Management Console (HMC) connection.
- The storage system is not already displayed on this page.

Select Action
Select a storage system that is listed on this page and then select one of the following actions:

Remove Storage System
Click to remove a storage system from this page. The Easy Tier heat map transfer process from the storage system to the target storage systems is stopped.

View Transfer Details
Click to view detailed information about the transfer from the source storage system to its target storage systems. This information includes whether the transfer succeeded or failed and the date and time that the transfer started and stopped. This action is not available if the source storage system does not have target storage systems.

Storage system details
The following information is shown for each storage system:

Storage System
Lists each source storage system and its target storage systems, if available.

Connection Status
One of the following states is displayed for each storage system connection:

Connecting
Indicates that connection to the storage system is in progress.

Connected
Indicates that connection to the storage system is established.

Disconnected
Indicates that connection to the storage system was lost. The connection can be lost for several reasons, including:

- The storage system is not available.
- The user name and password that is used to connect to the storage system was changed.

- There are firewall issues that prevent connection to the storage system.

Inactive
Indicates that the transfer of Easy Tier heat map data is disabled.

Related concepts:

"Easy Tier heat map transfer" on page 39
The IBM System Storage DS8000 Easy Tier heat map transfer function transfers heat map information from a source storage system to one or more target storage systems. Each target storage system then generates volume migration plans based on the heat map data and the physical configuration of the storage system.

Transfer Results for Storage System window

Use the Transfer Results for Storage System window to view details about the transfer of Easy Tier heat map data to target storage systems.

The following information is displayed in this window. This information includes the number of successful and failed transfers and the last start and stop time that a transfer occurred.

Target Storage System
Select the storage system for which you want to view the transfer details.

Attempts
Shows the number of attempts that were made to transfer the Easy Tier heat map data to the target storage system.

Successful Transfers
Shows the number transfers that were successful.

Failed Transfers
Shows the number of transfers that failed.

Last Transfer Start Time
Shows the date and time that the last transfer started. If a transfer is in progress, **In Progress** is shown.

Last Transfer Stop Time
Shows the date and time that the last transfer stopped. If a transfer is in progress, **Unknown** is shown.

Last Transfer Result
Provides a message that explains the results of the last transfer. If a transfer is in progress, a message specifying that the transfer is in progress is shown.

Next Transfer Time
Shows the date and time that the transfer is scheduled to occur again. If a transfer is in progress, **Unknown** is shown. Transfers are automatically scheduled for 12 hours after the previous transfer.

Chapter 5. Managing host systems

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

A connection to a z/OS host system is required if you want to enable z/OS features such as HyperSwap and hardened freeze in Tivoli Storage Productivity Center for Replication sessions.

If Tivoli Storage Productivity Center for Replication is installed on the host system, the host system connection is automatically added to Tivoli Storage Productivity Center for Replication. This connection is referred to as the native z/OS connection.

If Tivoli Storage Productivity Center for Replication is not installed on the host system, you must add the connection to the host system by using an IP connection. To use an IP connection, you must set certain Resource Access Control Facility (RACF®) settings on the z/OS host as described in the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

A connection to an AIX host system is required if you want to use the Open HyperSwap feature. This feature enables the automatic swap of input/output (I/O) to the volumes on a secondary site when a failure occurs when I/O is written to the primary site. For the swap to occur, the storage systems that contain the volumes must be attached to an AIX host and the AIX host must be connected to Tivoli Storage Productivity Center for Replication. The Open HyperSwap feature is optional for the Metro Mirror with Failover/Failback session type.

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

You can add both z/OS and AIX host connections in Tivoli Storage Productivity Center for Replication regardless of the operating system on which the application is installed. For example, you can add a connection to a z/OS host system in Tivoli Storage Productivity Center for Replication that is running on Windows, Linux, or AIX.

Adding a host system connection

You can add a connection to one or more host systems to the Tivoli Storage Productivity Center for Replication configuration.

Prerequisites: You must have Administrator privileges to add a host system connection.

Follow these steps to add a host system connection:
1. In the menu bar, click **Storage** > **Host Connections**.
2. On the Host Connections page, click **Add Host Connection**.
3. In the Add Host Connection window, enter the following information:

   Connection Type

   Select the host system type: z/OS or AIX.

Host name / IP address
Enter the host name or IP address for the host system.

Port   Enter the port for the host system.

For z/OS, the port number must be the same as the port number that is specified for the HyperSwap management address space IOSHMCTL SOCKPORT parameter on the z/OS system. For more information about the SOCKPORT parameter, see the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

For AIX, use the default port number 9930 unless the port was modified in Subsystem Device Driver Path Control Module (SDDPCM).

User name
Enter the user name. The **User name** field is displayed only for z/OS host systems. The user name and password must be the same name and password that are specified in the Resource Access Control Facility (RACF) settings on the host system. For more information about the RACF settings, see the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

Password
Enter the password. The **Password** field is displayed only for z/OS host systems.

4. Click **Add Host**.

The host system is displayed on the Host Connections page. If connection to the host system is successful, the status **Connected** is displayed. If the connection to the host system failed, the status **Disconnected** is displayed.

Modifying a host system connection

You can modify host system connections that are in the 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 the **Disconnected** state.

Follow these steps to modify a host system connection:
1. In the menu bar, click **Storage** > **Host Connections**.
2. On the Host Connections page, click the link for the host system in the **Host System** column. Complete the information in the following fields. The fields that are displayed depend on the host system and the type of connection to the host system.

Note: All fields are presented for informational purposes; however, some fields are information-only fields.

Host name or IP address
For AIX host systems or IBM z/OS host systems that are connected by using an IP connection, enter the host name or IP address for the host system.

If Tivoli Storage Productivity Center for Replication is installed on the z/OS host system, **ZOS_NATIVE_CONNECTION** is displayed.

Port   Enter the port for the host system.

For z/OS, the port number must be the same as the port number that is specified for the HyperSwap management address space IOSHMCTL SOCKPORT parameter on the z/OS system. For more information about the SOCKPORT parameter, see the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

For AIX, use the default port number 9930 unless the port was modified in Subsystem Device Driver Path Control Module (SDDPCM).

Type Shows the type of host system connection.

System name
Shows the name of the z/OS host system. If there is no connection to the host system, **Unknown** is displayed in this field.

Sysplex name
Shows the name of the z/OS sysplex that the host system is in or specifies that the system is not in a sysplex. If there is no connection to the host system, **Unknown** is displayed.

User name
Enter the user name for the z/OS host system. The user name and password must be the same name and password that are specified in the Resource Access Control Facility (RACF) settings on the host system. For more information about the RACF settings, see the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

Password
Enter the user password for the z/OS host system. If you change the user name, you must change the password.

Sessions
Shows the session that is associated with the AIX host system. The session is displayed only if the Open HyperSwap feature is enabled for the session.

The session name that is displayed might not be a session that is defined on the current Tivoli Storage Productivity Center for Replication server. For example, another Tivoli Storage Productivity Center for Replication server was used to connect to the host and a session name was defined on that server, but was not defined on the current server.

3. Click **Update Host**.

The host system is displayed on the Host Connections page. If connection to the host system is successful, the status **Connected** is displayed. If the connection to the host system failed, the status **Disconnected** is displayed.

Removing a host system connection

You can remove host system connections from the 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 features that are supported by the host such as Open HyperSwap or HyperSwap. Sessions that use the host system to provide these features can no longer communicate with the host and the features are disabled for the entire session.

Follow these steps to remove a host system connection:

1. In the menu bar, click **Storage** > **Host Connections**.
2. On the Host Connections page, select the host system connection that you want to remove.
3. From the **Select Action** list, select **Remove Host Connection**.
4. Click **OK** to remove the host system connection.

The host system is removed from the Host Connections page.

Removing a session from an AIX host system connection

You can remove a session that is associated with an IBM AIX host system from the Tivoli Storage Productivity Center for Replication configuration. When you remove a session, the host no longer recognizes that the session is managing the volumes that are attached to that host. This function is to be used primarily for cleanup purposes.

Prerequisites: You must have Administrator privileges to remove a session association.

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 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 **View/Modify Host Connection** window for the host system. To open this window, click the link for the host system in the **Host System** column on the Host Connections page.

Follow these steps to remove a session association from a host system connection:

1. In the menu bar, click **Storage** > **Host Connections**.
2. On the Host Connections page, select the host system connection that contains the session that you want to remove.
3. From the **Select Action** list, select **Remove Session Association**.
4. Click **OK** to remove the session association.

The session name is removed from the **Sessions** field in the **View/Modify Host Connection** window.

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 for Replication for System z.

To configure partnerships for the following storage systems, use the graphical user interface (GUI) or command-line interface (CLI) for the storage system:

- SAN Volume Controller
- Storwize
- The XIV

Viewing logical paths

You can view all logical paths that are defined on an TotalStorage Enterprise Storage Server, System Storage DS8000, or System Storage DS6000 storage system.

Complete one of these procedures to view logical paths:

- From the ESS/DS Paths page of Tivoli Storage Productivity Center for Replication:
  1. In menu bar, click **Paths**.
  2. Click the storage system ID to display logical paths for that storage system.
- From the Storage Systems page:
  1. In menu bar, click **Storage** > **Storage Systems**.
  2. Select the storage system for which you want to view logical paths.
  3. From the **Select Action** list, select **View Paths**. The paths page is displayed with a list of defined logical paths.

Adding logical paths

You can add logical paths for IBM TotalStorage Enterprise Storage Server, IBM System Storage DS8000, and IBM System Storage DS6000 storage systems.

Ensure that you have defined the appropriate storage systems on the Storage Systems page.

Perform these steps to add logical paths:

1. In menu bar, click **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 subsystem, target storage system, and target logical subsystem. 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 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 by 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. Tivoli Storage Productivity Center for Replication uses the port pairings that are defined in the CSV file to establish logical paths.

You can define the following types of port pairing for the source and target storage systems:

- Port pairings for all logical subsystems that are in a session.
- Port pairings for a logical subsystem to logical subsystem mapping.
- Port pairings for all logical subsystems. You can define a single set of port pairings for the logical subsystems or you can define different pairings for all even logical subsystems and all odd logical subsystems.

Complete these steps to add IBM TotalStorage Enterprise Storage Server, IBM System Storage DS8000, and IBM System Storage DS6000 logical paths by using a CSV file:

1. Create a CSV file named `portpairings.csv` in the *path_prefix*/opt/Tivoli/RM/ wlp/usr/servers/replicationServer/properties/ directory. You can create the CSV file in a spreadsheet such as Microsoft Excel or in a text editor. The `properties` directory contains an example CSV file that is named `portpairings.csv.sample`, which contains the following example port pairs:

```
# 1. Each line in the file represents a storage system pairing
#    and its associated port pairings.
#    There are multiple formats that define
#    different port pairings types:
#
# Specific Session:  This format can be used if you
# want to use a different set of port pairings for each session.
# This may be necessary in some environments in order to prevent
# the paths of of session to affect another.
#
# Format:
# storageSystemA:storageSystemB:sessionName,
# port1:port2,...,portN:portM
#
# where storageSystemA and storageSystemB are the storage
# system IDs used in TPC-R to identify the storage systems
# and sessionName is the name of the session you wish this
# port pairing to be used for.  port1, port2, portN, portM
# are hexadecimal numbers indicating the storage system
# communication ports.
#
# A valid example would be:
# 2107.04131:2107.01532:myPayRollSession,0x0331:0x0024,
# 0x0330:0x0100,0x0331:0x000C
#
#
# Specific LSS to LSS pairing:  This format can be used
# if you want to specify that a specific LSS pairing use a
# specific set of port pairings.
#
# Format:
# storageSystemA:systemALSS:storageSystemB:systemBLSS,
# port1:port2,...,portN:portM
#
```

```
# where storageSystemA and storageSystemB are the storage system
# IDs used in TPC-R to identify the storage systems, and
# systemALSS/systemBLSS are LSSs on storage systems as in
# hexidecimal format.  port1, port2, portN, portM
# are hexadecimal numbers indicating the storage system
# communication ports.
#
# A valid example would be:
# 2107.04131:0x52:2107.01532:0x5f,0x0331:0x0024,
# 0x0330:0x0100,0x0331:0x000C
#
#
# Box to Box pairing:  This format can be used if you want
# to specify a set of port pairings that will be used for
# all LSS pairings across the same source and target box.
# This pairing has multiple options. A single set of port pairings
# can be defined for all LSS pairings between the boxes,
# OR the port pairings can define that the even source LSSs
# should use one set and the odd LSSs should use a different set.
# There are some scenarios particularly with Global Mirror
# support where performance can be improved by mapping
# even and odd LSSs.
#
# Format:
# Option 1:
# storageSystemA:storageSystemB,port1:port2,...,portN:portM
#
# where storageSystemA and storageSystemB are the storage system
# IDs used in TPC-R to identify the storage systems.
# port1:port2,... portN:portM are hexadecimal numbers indicating
# the storage system communication ports.
#
# A valid example would be:
# 2107.04131:2107.01532,0x0331:0x0024,0x0330:0x0100,0x0331:0x000C
#
# Option 2:
# storageSystemA:storageSystemB,even, port1:port2,...,portN:portM,
# odd,port1:port2,...portX:portY
#
# where storageSystemA and storageSystemB are the storage system IDs
# used in TPC-R to identify the storage systems.
# port1:port2,...portN:portM is the set of port pairings to use when
# the source is on an even LSS
# port1:port2,...portX:portY is the set of port pairings to use when
# the source is on an odd LSS
# All ports are hexadecimal numbers indicating the storage system
# communication ports.
#
# A valid example would be:
# 2107.04131:2107.01532,even,0x0331:0x0024,0x0330:0x0100,odd,0x0331:0x000C,
# 0x0334:0x0330
#
```

2. To enable the changes in the file, complete 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 bidirectional. 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 on the ESS/DS Paths page:
  - 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 that are 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

You can remove logical paths that are defined on an TotalStorage Enterprise Storage Server, System Storage DS8000, or System Storage DS6000 storage system.

1. In menu bar, click **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 **Select Action** list, select **Remove**.

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

The type of data replication that is associated with the session determines the actions that you can perform against all copy sets in the session, the number of volumes that a copy set can contain, and the role that each volume plays.

Attention: Use only the Tivoli Storage Productivity Center for Replication graphical user interface (GUI) or command-line interface CLI to manage session relationships, such as volume pairs and copy sets. Do not modify session relationships through other interfaces such as the System Storage DS CLI. If you modify relationships through other interfaces, a loss of consistency can occur across the relationships that are managed by the session. The exceptions to this requirement are failover operations that are managed by external applications for certain session and storage system types as described in "Failover and failback operations" on page 69.

Single-target and multi-target sessions

Sessions are referred to as single-target or multi-target sessions. With single-target sessions, the source volume site can have only one target site. Data replication occurs from the source to the target.

With multi-target sessions, the source volume site can have multiple target sites. Data replication can occur from the source to an individual target or to all targets simultaneously.

Related tasks:

"Completing session administration tasks" on page 150
Use session commands to complete session tasks such as starting, stopping, suspending, or terminating sessions.

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

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

Table 22. 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 Single Direction	1	2
Metro Mirror Failover/Failback	1	2
Metro Mirror Failover/Failback with Practice	3	3
Metro Mirror - Metro Mirror	3	3
Global Mirror Single Direction (TotalStorage Enterprise Storage Server, System Storage DS8000, System Storage DS6000)	3	3
Global Mirror Single Direction (SAN Volume Controller, Storwize)	1	2
Global Mirror Failover/Failback (TotalStorage Enterprise Storage Server, System Storage DS8000, System Storage DS6000)	3	3
Global Mirror Failover/Failback (SAN Volume Controller, Storwize)	1	1
Global Mirror Failover/Failback with Practice (TotalStorage Enterprise Storage Server, System Storage DS8000, System Storage DS6000)	5	4
Global Mirror Failover/Failback with Practice (SAN Volume Controller, Storwize)	3	3
Global Mirror Either Direction with Two-Site Practice	8	6
Global Mirror Failover/Failback with Change Volumes	1	4
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 a session. You can select a storage system; a logical subsystem (LSS), I/O group, or pool; or a 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 LSSs for the selected storage systems. For example, volume 01 on the source LSS is matched with volume 01 on the target LSS.

You can 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 are valid. This process 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.

Adding copy sets

When you create a copy set for a session, a warning is displayed if one or more of the selected volumes exist in another session. If the volume you selected is in another session, confirm whether you want to add the volume to this session. Ensure that you intend to have the volume in multiple sessions because having the volume in multiple sessions can create conflicts. A scenario in which you would want the same volumes in multiple sessions is one in which you create a session for normal replication and then create another session for a disaster recovery practice. You must use the same target volumes in both sessions.

You can use extent space-efficient volumes as copy set volumes depending on the capabilities of the storage system.

Removing copy sets

You remove a copy set or range of copy sets by selecting the following items:
- Source volume
- LSS, I/O group, or pool
- 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. The copy set is removed from the 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, commands are 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 to 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.

Depending on the storage system, the following actions occur when a copy set is removed:

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 completed any background copy.

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

- 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. In this situation, you are not able to remove the relationship by 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 is not 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.

Related tasks:

"Creating sessions and adding copy sets" on page 137
A *session* completes a specific type of data replication for a specific set of volumes. During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a copy set. A session can contain one or more copy sets.

"Removing copy sets" on page 156
This topic describes how to remove copy sets.

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 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 session type. 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   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. Tivoli Storage Productivity Center for Replication supports these volume roles:

Host volume

A host volume is 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 journal volume 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 session can be 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

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

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

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

Target volume

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

Change volume (only SAN Volume Controller or Storwize storage systems)

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

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

Related concepts:

"Session types" on page 12
Tivoli Storage Productivity Center for Replication provides several methods to replicate data. The type of data replication that is associated with a session is known as the *session type*.

"Copy sets" on page 63
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*.

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 smaller interval 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.

Failover and failback operations

If a session has *failover* and *failback* capabilities, you can perform a site switch in which you move the operation of Tivoli Storage Productivity Center for Replication from one site to another, and change the direction of the copy without completing a full copy.

Failover is the process of temporarily switching production to a backup facility following a scheduled maintenance period or a disaster at a production site. A failover operation can be followed by a failback operation, which is the process of returning production to its original location and resuming the copy between the two sites.

The failover and failback operations use remote mirror and copy functions for these tasks:

- To reduce the time that is required to synchronize volumes after switching sites during planned or unplanned outages.
- To eliminate the need to perform a full volume copy from the recovery site to the production site, which can reduce the time that is required to resume operations at the production site.

Failover operations that are managed by other applications

Applications such as the IBM Series i Toolkit, VMware Site Recovery Manager, and Veritas Cluster Server can also manage failover operations for the session types that are listed in the following table.

Session type	Storage systems
Metro Mirror Single Direction	- System Storage DS8000 - System Storage DS6000 - SAN Volume Controller - Storwize V3700 - Storwize V7000 - Storwize V7000 Unified
Metro Mirror Failover/Failback	- System Storage DS8000 - System Storage DS6000 - SAN Volume Controller - Storwize V3700 - Storwize V7000 - Storwize V7000 Unified - The XIV
Metro Mirror Failover/Failback with Practice	- System Storage DS8000 - System Storage DS6000 - SAN Volume Controller - Storwize V3700 - Storwize V7000 - Storwize V7000 Unified
Metro Mirror - Metro Mirror	System Storage DS8000 (This session type is available only for System Storage DS8000 storage systems with a microcode level that supports single source to multi-target relationships. To determine whether you can use this session type, refer to the System Storage DS8000 documentation for the microcode level that you are using.)
Global Mirror Single Direction	- System Storage DS8000 - System Storage DS6000 - SAN Volume Controller - Storwize V3700 - Storwize V7000 - Storwize V7000 Unified
Global Mirror Either Direction with Two-Site Practice	- System Storage DS8000 - System Storage DS6000

Session type	Storage systems
Global Mirror Failover/Failback	• System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified • The XIV
Global Mirror Failover/Failback with Practice	• System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified
Global Mirror Failover/Failback with Change Volumes	• SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified

If an application completes a failover operation for a session, the Severe status is displayed for the session on the Session Details page in Tivoli Storage Productivity Center for Replication and an error message is generated for the role pairs for which the failover occurred. To view the error message:

1. On the Session Details page, click the link for each role pair for which the failover occurred.
2. On the Role Pair Details page, click the **Show** link for each role pair. This link is in the **Details** column.

To change the session status to Normal, use the application to confirm that the role pairs that are in the session are consistent and then restart the session.

Session types

Tivoli Storage Productivity Center for Replication provides several methods to replicate data. The type of data replication that is associated with a session is known as the *session type*.

The session types that are available depend on the storage system type and the edition of IBM Spectrum Control or Tivoli Storage Productivity Center for Replication that you are using as shown in this topic.

Basic HyperSwap

Basic HyperSwap is a copy services solution for IBM z/OS version 1.9 and later. Basic HyperSwap provides high availability of data if a storage system failure occurs.

The Basic HyperSwap session type enables Tivoli Storage Productivity Center for Replication to manage Basic HyperSwap by using Metro Mirror or Metro Global Mirror replication.

Basic HyperSwap sessions are available for the following Tivoli Storage
Productivity Center for Replication editions and storage systems.

Table 23. Basic HyperSwap

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication Basic Edition for System z • Tivoli Storage Productivity Center for Replication for System z	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000

FlashCopy

FlashCopy replication creates a point-in-time copy of the data on a source volume
to a target volume. Data that existed on the target volume is replaced by the
copied data.

FlashCopy sessions are available for the following IBM Spectrum Control and
Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 24. FlashCopy

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication Basic Edition for System z • Tivoli Storage Productivity Center for Replication for System z • All editions of IBM Spectrum Control	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3500 • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)

Snapshot

Snapshot replication 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.

Snapshot sessions are available for the following IBM Spectrum Control and Tivoli
Storage Productivity Center for Replication editions and storage systems.

Table 25. Snapshot

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication Basic Edition for System z • Tivoli Storage Productivity Center for Replication for System z • All editions of IBM Spectrum Control	The XIV system

Metro Mirror Single Direction

Metro Mirror is synchronous replication that operates between two sites that are up to 300 km apart. The source volumes are on one storage system and the target volumes are on another storage system.

Metro Mirror replication maintains identical data in both the source and target volumes. When a write operation is issued to a source volume, the changes are propagated to the target volume before the write operation finishes processing.

With the Metro Mirror Single Direction session type, Metro Mirror replication is available only from the primary site.

Metro Mirror Single Direction sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 26. Metro Mirror Single Direction

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of IBM Spectrum Control	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)

Metro Mirror Failover/Failback

Metro Mirror Failover/Failback replication enables you to switch the direction of the data flow so that you can use your secondary site as your production site. You can then copy changes that are made at the secondary site back to the primary site.

Metro Mirror Failover/Failback sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 27. Metro Mirror Failover/Failback

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of IBM Spectrum Control	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage) • The XIV system

Metro Mirror Failover/Failback with Practice

Metro Mirror Failover/Failback with Practice replication combines Metro Mirror Failover/Failback and FlashCopy capabilities to provide a point-in-time copy of the data on the secondary site. This session type provides volumes that you can use to practice for disaster recovery without losing your disaster recovery capability.

Metro Mirror Failover/Failback with Practice sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 28. Metro Mirror Failover/Failback with Practice

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of IBM Spectrum Control	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)

Global Mirror Single Direction

Global Mirror is asynchronous replication that operates between two sites that are over 300 km apart. The source volumes are on one storage system and the target volumes are on another storage system.

Global Mirror replication maintains identical data in both the source and target volumes. When a write operation is issued to a source volume, the changes are typically propagated to the target volume a few seconds after the data is written to the source volume.

With the Global Mirror Single Direction session type, Global Mirror replication is available only from the primary site.

Global Mirror Single Direction sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 29. Global Mirror Single Direction

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of IBM Spectrum Control	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)

Global Mirror Either Direction with Two-Site Practice

Global Mirror Either Direction with Two-Site Practice replication enables you to run Global Mirror replication from either the primary or secondary site. This session type provides volumes on the primary and secondary site that you can use to practice for disaster recovery without losing your disaster recovery capability.

Global Mirror Either Direction with Two-Site Practice sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 30. Global Mirror Either Direction with Two-Site Practice

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of IBM Spectrum Control	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000

Global Mirror Failover/Failback

Global Mirror Failover/Failback replication enables you to switch the direction of the data flow so that you can use your secondary site as your production site. You can then copy changes that are made at the secondary site back to the primary site.

Global Mirror Failover/Failback sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 31. Global Mirror Failover/Failback

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of IBM Spectrum Control	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage) • The XIV system

Global Mirror Failover/Failback with Change Volumes

Global Mirror Failover/Failback with Change Volumes replication provides the same capabilities as the Global Mirror Failover/Failback session type. However, Global Mirror Failover/Failback with Change Volumes sessions also provide the option of enabling or disabling the use of change volumes. Change volumes are available in SAN Volume Controller and Storwize storage systems.

Global Mirror Failover/Failback with Change Volumes sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 32. Global Mirror Failover/Failback with Change Volumes

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of IBM Spectrum Control	• SAN Volume Controller Version 6.3 and later • Storwize V3700 Version 6.4.1 and later • Storwize V7000 Version 6.3 and later • Storwize V7000 Unified Version 1.3 and later (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)

Global Mirror Failover/Failback with Practice

Global Mirror Failover/Failback with Practice replication combines Global Mirror Failover/Failback and FlashCopy capabilities to provide a point-in-time copy of the data on the secondary site. This session type provides volumes that you can use to practice for disaster recovery without losing your disaster recovery capability.

Global Mirror Failover/Failback with Practice sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 33. Global Mirror Failover/Failback with Practice

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of IBM Spectrum Control	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)

Metro Global Mirror

Metro Global Mirror replication provides continuous, remote data replication that operates between three sites that are 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. Metro Global Mirror replication enables you to switch the direction of the data flow so that you can use your secondary or tertiary site as your production site.

Metro Global Mirror sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 34. Metro Global Mirror

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of IBM Spectrum Control	• TotalStorage Enterprise Storage Server Model 800 (only H1 site) • System Storage DS8000

Metro Global Mirror with Practice

Metro Global Mirror with Practice replication combines Metro Mirror, Global Mirror, and FlashCopy capabilities to provide a point-in-time copy of the data on the tertiary site. This session type provides volumes that you can use to practice for disaster recovery without losing your disaster recovery capability.

Metro Global Mirror with Practice sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 35. Metro Global Mirror with Practice

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of IBM Spectrum Control	• TotalStorage Enterprise Storage Server Model 800 (only H1 site) • System Storage DS8000

Related concepts:

"Basic HyperSwap"
Basic HyperSwap is a copy services solution for z/OS version 1.9 and later. It provides high availability of data if a disk storage system failure occurs. Basic HyperSwap does not replace a disaster recovery solution. If a session is suspended but the suspend operation was not caused by a HyperSwap trigger, the session is not frozen to ensure that it is consistent.

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.

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.

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 in one storage system and the target is in another storage system.

Global Mirror is asynchronous replication that operates between two sites that are over 300 km apart. The source volumes are on one storage system and the target volumes are on another storage system. Global Mirror replication maintains identical data in both the source and target volumes.

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

Basic HyperSwap

Basic HyperSwap is a copy services solution for z/OS version 1.9 and later. It provides high availability of data if a disk storage system failure occurs. Basic HyperSwap does not replace a disaster recovery solution. If a session is suspended but the suspend operation was not caused by a HyperSwap trigger, the session is not frozen to ensure that it is consistent.

When HyperSwap is combined with Metro Mirror and Metro Global Mirror replication, you can prepare your system for disaster recovery and ensure high availability of data. If a session is suspended but the suspend operation was not caused by a HyperSwap trigger, the session is frozen to ensure that it is consistent.

Restriction: This replication method is available only for ESS, DS6000, and DS8000 storage systems.

Basic HyperSwap replication performs the following actions:
- Manages CKD volumes in Metro Mirror to manage 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) does not add a copy set if a fixed block volume is specified.
- Monitors events that indicate a storage device failed.

- Determines whether the failing storage device is part of a Metro Mirror synchronous PPRC pair.
- Determines the action to be taken from policy.
- Ensures that data remains consistent.
- Swaps the I/O between the primary logical devices in the consistency group with the secondary logical devices in the consistency group. A swap can occur from the preferred logical devices to the alternate logical devices or from the alternate logical devices to the preferred logical devices.

Example

Jane is using multiple DS8000 storage systems. The host applications run on a z/OS operating system and the z/OS environment is connected to the DS8000 storage systems. She has a site in Manhattan and a secondary DS8000 in Hoboken. While it is not required that she has a disaster recovery solution in effect, she does need to have a high-availability solution to keep her applications running all the time. Jane is worried that if a volume fails on the DS8000 in Manhattan, her database application might fail. Even a small downtime can be costly to Jane.

Jane uses a Basic HyperSwap session to mirror the data on the DS8000 in Manhattan to the 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 provide multiple logical subsystems (LSSs) within a single physical subsystem, while the following platforms provide multiple I/O groups:
- SAN Volume Controller
- Storwize V3500
- Storwize V3700
- Storwize V7000
- Storwize V7000 Unified

All platforms can use 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 the FlashCopy feature, you can reference and update the source volume and target volume independently.

The following figure illustrates a FlashCopy session.

Example

Jane works for a bank. Jane uses a FlashCopy session to make a point-in-time copy of the customer data in existing international accounts. Every night, the bank's servers start batch processing. Jane uses the FlashCopy session to create checkpoint restarts for the batch processing if the batch processing fails. In the batch processing, the first step is to balance all international accounts and to make a FlashCopy point-in-time copy of the resulting data. The second step in the batch processing is to complete the international disbursements.

If the second step in the batch process fails, Jane can use the data from the FlashCopy session that was taken of the first step to repeat the second step. In this way, Jane does not have to begin the entire process again. Jane also uses a CLI script that performs a FlashCopy operation every night at 11:59 PM, and another script that quiesces the database. She backs up the data to tape on the target storage system, and then sends the tape to the bank's data facility 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 a snapshot session to create a large number of volume copies without affecting system 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 a Snapshot session.

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 the data. The snapshot function ensures that if data becomes corrupted, she can restore the data to any point 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 the host applications. Jane does not have to provision target volumes for all the snapshots she intends to make. She can quickly configure a single Snapshot session.

When the session is configured, Jane uses a CLI script that runs 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 from a time 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 affecting 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 in one storage system and the target is in another storage system.

Attention: If you have sessions that contain Metro Mirror relationships, ensure that the session does not contain system volumes (such as paging volumes) unless you select the **Manage H1-H2 with HyperSwap** or the **Enable Hardened Freeze** option for the session. By using these options, z/OS Input/Output Supervisor (IOS) manages freeze operations for the volumes in the session, which prevents Tivoli Storage Productivity Center for Replication from freezing the volumes and possibly freezing itself.

Metro Mirror replication maintains identical data on both the source and target. When a write operation is issued to the source copy, the changes that are made to the source data are propagated to the target before the write operation finishes processing. If the storage system ends, no data is lost when you use Metro Mirror 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 that is migrated. As a result, the data on the target system might no longer be consistent with the source system. The Metro Mirror session switches back a synchronous relationship when Metro Mirror reissues a **Start** command. In addition, you can start a Metro Mirror session in Global Copy mode and toggle between Metro Mirror and Global Copy modes to accommodate periods in which you require host I/O response time over data consistency.

Tip: To determine whether there is any out-of-sync data that must to be copied before the session is consistent, check the percent that is complete in the session details page.

Metro Mirror Single Direction

The following figure illustrates a Metro Mirror Single Direction session.
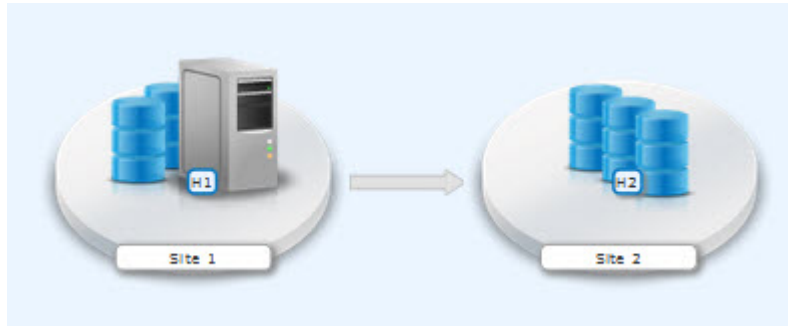
Metro Mirror Failover/Failback

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

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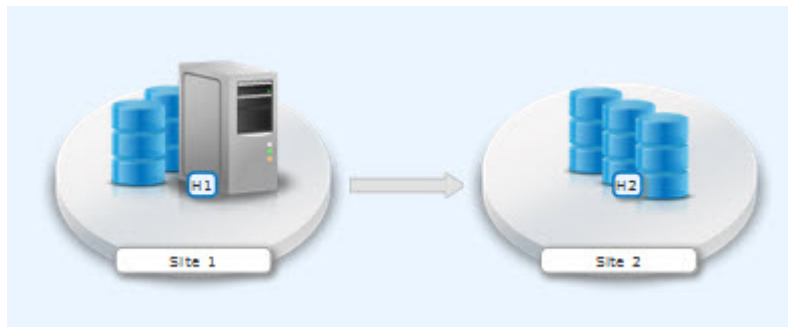
Metro Mirror Failover/Failback with Practice

A Metro Mirror Failover/Failback with Practice session combines Metro Mirror and FlashCopy features 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 occurs, without losing your disaster recovery capability.

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

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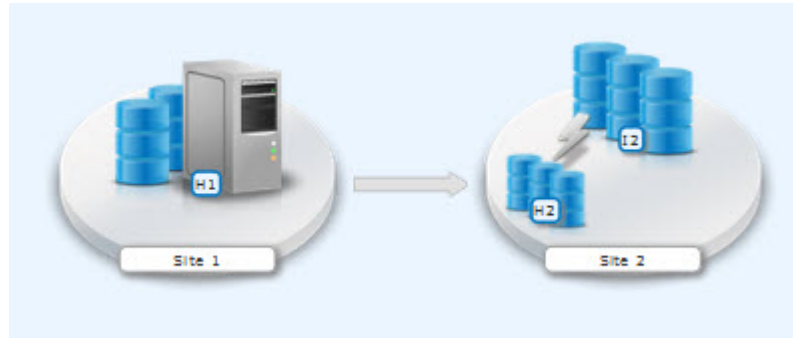
Metro Mirror Failover/Failback with HyperSwap

You can enable a Metro Mirror Failover/Failback session to have HyperSwap capabilities. To enable HyperSwap processing, see "Setting up the environment for HyperSwap" on page 97.

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 for 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 minimal effect on the application.

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

The following figure illustrates a Metro Mirror Failover/Failback session that is enabled for HyperSwap.
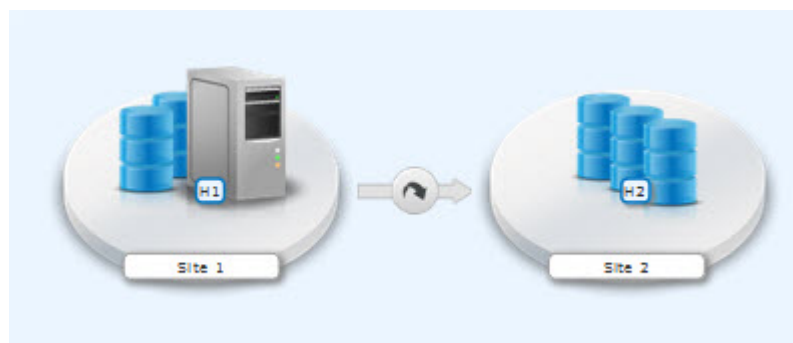
Metro Mirror Failover/Failback with Open HyperSwap

You can enable a Metro Mirror Failover/Failback session to have Open HyperSwap capabilities. To enable Open HyperSwap processing, see "Setting up the environment for Open HyperSwap" on page 98.

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 minimal effect on the application.

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

The following figure illustrates a Metro Mirror Failover/Failback session that is enabled for Open HyperSwap.
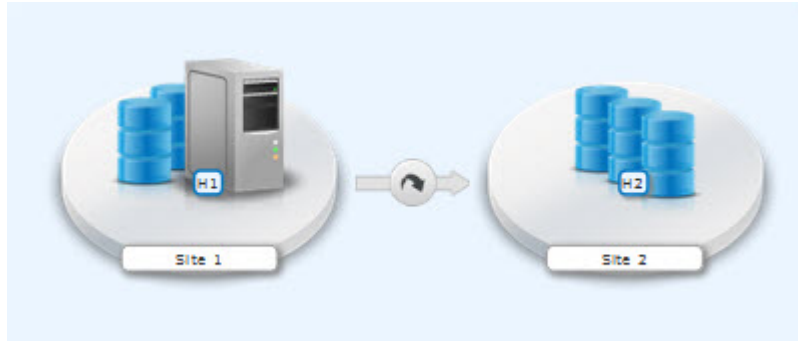
Metro Mirror - Metro Mirror

A Metro Mirror - Metro Mirror session is a multi-target session that consists of three sites. You can define any of the sites as the primary site and then run Metro Mirror replication from the primary site to either of the other sites individually or both sites simultaneously. For example, if Site 1 is the primary site, data replication can occur between the H1 and H2 volumes and the H1 and H3 volumes separately or at the same time.

This session type is available only for System Storage DS8000 storage systems with a microcode level that supports single source to multi-target relationships. To determine whether you can use this session type, refer to the System Storage DS8000 documentation for the microcode level that you are using.

The following figure illustrates a Metro Mirror - Metro Mirror session.
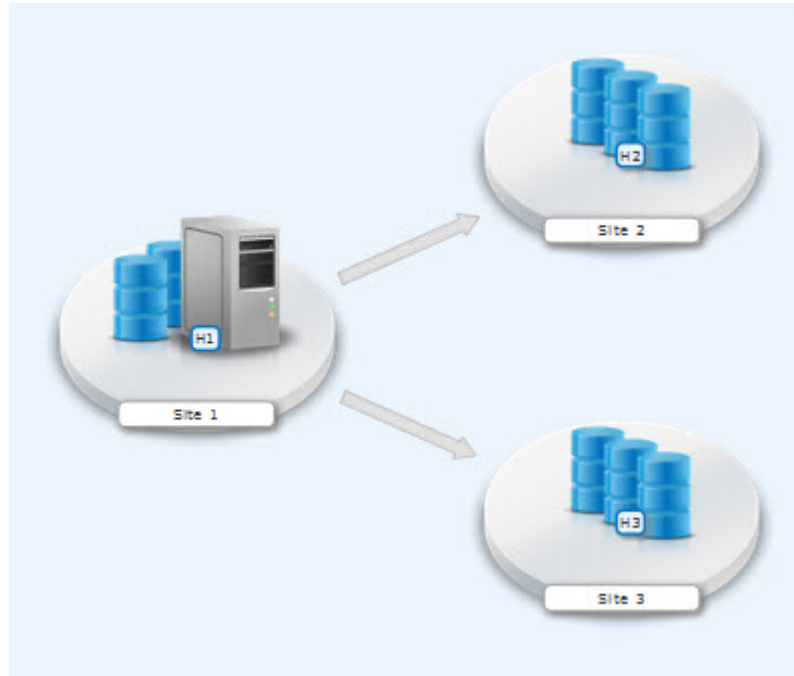
Examples

Read the following scenarios for information on using Metro Mirror for synchronous, remote data replication between two sites.

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 the 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 posted. This method of replication maintains identical data in both the source and target. The advantage of this method 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 the data is safe.

The night before the planned outage, Jane quiesces the 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 shuts down the equipment in San Francisco to avoid any power spikes when she restarts the system after the power is turned 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 the backup site during the outage. She wants a synchronous copy method with no data loss for the 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 the effect on any applications while completing the initial Metro Mirror synchronization, so she begins the session in Global Copy mode. After she sees that approximately 70% of the data is copied, Jane decides to switch the session to 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 the recovery procedure for the Oakland site, but she cannot stop running the Metro Mirror session while she takes time to practice a recovery. By using a Metro Mirror with Practice session, Jane can practice the disaster recovery scenario in Oakland while the Metro Mirror session runs uninterrupted. By practicing running the applications at the Oakland site, Jane is better prepared to recover data if a disaster occurs at the San Francisco site.

While her session is running in Prepared state, Jane practices a recovery at the 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 the session continues to replicate data from San Francisco to Oakland. As a result, she can carry out the recovery testing without stopping the replication for any extended duration of time.

If, at some point, the Metro Mirror session is suspended because of a failure, Jane can use the practice session to restart the data replication process. She maintains 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 the restarted session is in Prepared state, she has a consistent version of the data on the H2 volumes. She only has to issue the `Recover` command to put the session into Target Available state and make the H2 volumes accessible from the servers. If the session was not in Prepared state when it was 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 is 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 the recovery procedure for the Oakland site, but she cannot stop running the Metro Mirror session while she takes time to practice a recovery. By using a Metro Mirror with Practice session, Jane can practice the disaster recovery scenario in Oakland while the Metro Mirror session runs uninterrupted. By practicing running the applications at the Oakland site, Jane is better prepared to recover data if a disaster occurs at the San Francisco site.

While the session is running in a Prepared state, Jane practices a recovery at the 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 the session continues to replicate data from San Francisco to Oakland. As a result, she can carry out the recovery testing without stopping the replication for any extended duration of time.

If the Metro Mirror session is suspended because of a failure, Jane can use the practice session to restart the data replication process while she maintains 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 the restarted session is in Prepared state, she has a consistent version of the data on the H2 volumes. She only has to issue the **Recover** command to put the session into Target Available state and make the H2 volumes accessible from the servers. If the session was not in Prepared state when it was 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 is not overwritten if the data to be copied to the volumes is not consistent.

Selecting a HyperSwap session

A global insurance company decided 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 operating systems. They have two data centers in New York.

Jane chooses a Metro Mirror recovery solution because her priority is to protect the system from regional disasters. Jane decides to use Metro Mirror solution because her company has two data centers that are located near each other. Jane realizes that because she uses a z/OS operating system, CKD, and System Storage DS8000 hardware, she can also use a HyperSwap solution. Using Metro Mirror Failover/Failback with HyperSwap, Jane can minimize the effects on any applications, while she maintains the failover process to the secondary site. Jane decides Metro Mirror Failover/Failback with HyperSwap is the best solution.

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 that she intends to use on all sites. From the Sessions page, Jane opens the Create Session window and selects the **Metro Mirror Failover/Failback** session type. After completing the information in the window, Jane clicks **Launch Add Copy Sets Wizard**. After she completes the wizard, she selects the **Manage H1-H2 with HyperSwap** option in the View/Modify Properties notebook.

Jane then issues a **Start H1->H2** command. After the initial copy is completed, Jane can safely replicate the data between both sites. She can

also issue the **HyperSwap** command between sites 1 and 2 to switch sites with minimal effect on the application during either a disaster or maintenance period.

Performing a planned HyperSwap

Jane's company used Metro Mirror Failover/Failback with HyperSwap sessions for the past three months. However, Jane must perform maintenance on an H1 volume. During this time, Jane does not want the applications or replication to be interrupted. To prevent this interruption, before the maintenance is scheduled to begin, Jane uses the Tivoli Storage Productivity Center for Replication GUI to perform a HyperSwap operation to the H2 volumes. This process changes the applications so that the data is written to H2. To perform a planned HyperSwap operation, 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. Jane's data is safe because she used Metro Mirror Failover/Failback with HyperSwap, and the H2 volume is an exact duplicate of the H1 volume. When the permanent I/O error is detected, a HyperSwap is triggered. The application changes to write data to the H2 volumes. The applications are not interrupted.

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

Global Mirror

Global Mirror is asynchronous replication that operates between two sites that are over 300 km apart. The source volumes are on one storage system and the target volumes are on another storage system. Global Mirror replication maintains identical data in both the source and target volumes.

When a write operation is issued to a source volume, the changes are typically propagated to the target volume a few seconds after the data is written to the source volume. However, changes can occur on the source volume before the target volume verifies that it 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 the system fails, Global Mirror might lose some data that was transmitted when the failure occurred.

Global Mirror Single Direction

The Global Mirror Single Direction session type enables you to run Global Mirror replication from the primary site.

For TotalStorage Enterprise Storage Server Model 800, System Storage DS6000, and System Storage DS8000 storage systems, each copy set in the single direction

session consists of two host volumes and a journal volume. The following figure
illustrates a Global Mirror Single Direction session for these storage systems.

For SAN Volume Controller and Storwize storage systems, each copy set in the
Global Mirror Single Direction session consists of two host volumes. The following
figure illustrates a Global Mirror Single Direction session for these storage systems.
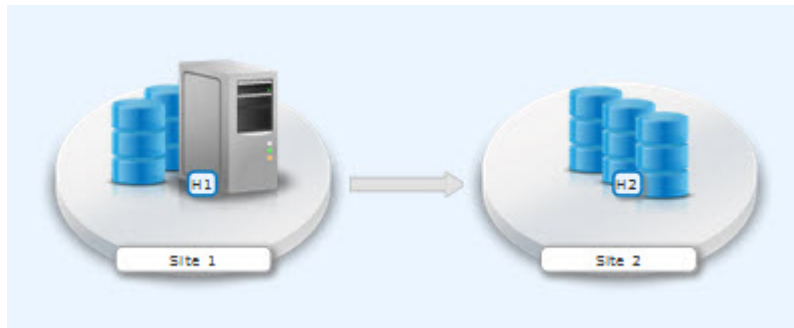
Global Mirror Failover/Failback

The Global Mirror Failover/Failback with Practice session type combines Global
Mirror Failover/Failback and FlashCopy capabilities to provide a point-in-time
copy of the data on the secondary site. This session type provides volumes that
you can use to practice for disaster recovery without losing your disaster recovery
capability.

For TotalStorage Enterprise Storage Server Model 800, System Storage DS6000, and
System Storage 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 a Global Mirror Failover/Failback session for these
storage systems.

For SAN Volume Controller and Storwize storage systems, each copy set in the Global Mirror Failover/Failback session consists of two host volumes. The following figure illustrates a Global Mirror Failover/Failback session for these storage systems.

Global Mirror Failover/Failback with Change Volumes

Global Mirror Failover/Failback with Change Volumes replication provides the same capabilities as the Global Mirror Failover/Failback session type. However, Global Mirror Failover/Failback with Change Volumes sessions also provide the option of enabling or disabling the use of change volumes. Change volumes are available in SAN Volume Controller and Storwize storage systems.

Change volumes are composed of a source change volume and a target change volume that contain a point-in-time image of the data from the source and target volumes, respectively.

A FlashCopy operation occurs between the source volume and the source change volume. The frequency of the FlashCopy operation is determined by the cycle period. For information about the cycle period and how it is set, see "Global Mirror Failover/Failback with Change Volumes session properties" on page 197. The data on the source change volume is then replicated to the target volume, and finally to the target change volume.

Because the data that is replicated between sites contains point-in-time changes rather than all changes, a lower bandwidth link is required between the sites. However, when you use change volumes, your data exposure can increase. Therefore, you might want to include or exclude change volumes depending on your network traffic or business requirements.

This session type is available only for SAN Volume Controller and Storwize storage systems. The following figure illustrates a Global Mirror Failover/Failback with Change Volumes session for these storage systems.

Global Mirror Failover/Failback with Practice

The Global Mirror Failover/Failback with Practice session type combines Global Mirror Failover/Failback and FlashCopy capabilities to provide a point-in-time copy of the data on the secondary site. This session type provides volumes that you can use to practice for disaster recovery without losing your disaster recovery capability.

For TotalStorage Enterprise Storage Server Model 800, System Storage DS6000, and System Storage 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 a Global Mirror Failover/Failback with Practice session for these storage systems.

For SAN Volume Controller and Storwize storage systems, each copy set in the Global Mirror failover/failback with Practice session consists of two host volumes and an intermediate volume. The following figure illustrates a Global Mirror Failover/Failback with Practice session for these storage systems.

Global Mirror Either Direction with Two-Site Practice

Global Mirror Either Direction with Two-Site Practice enables you to run Global Mirror replication from either the primary or secondary site. This session type provides volumes on the primary and secondary site that you can use to practice for disaster recovery without losing your disaster recovery capability.

This session type is available only for TotalStorage Enterprise Storage Server Model 800, System Storage DS6000, and System Storage DS8000 storage systems.

The session consists of two host volumes, two intermediate volumes, and two journal volumes. The following figure illustrates a Global Mirror Either Direction with Two-Site Practice session for these storage systems.
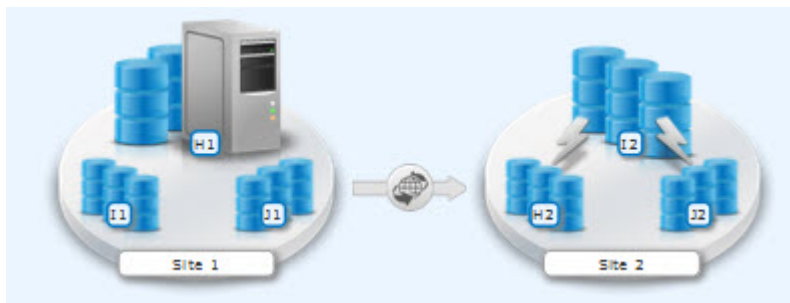
Examples

Read the following scenarios for information about using Global Mirror for asynchronous, remote data replication between two sites.

Global Mirror Single Direction
> Although Jane's FlashCopy and Metro Mirror copies were both planned, Jane realizes that sometimes a failure can occur, and she wants to ensure that the data is safe. Because Jane works in San Francisco, she wants her other site to be far away if a disaster occurs locally. Her other site is based in Houston. A minor earthquake occurs in San Francisco and power and communications both fails. Jane arranged for the data on customer accounts that 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 processing in San Francisco. However, she understands that it is minor inconvenience when weighed next to the value of backing up her data to an area that is not prone to earthquakes.

Global Mirror with Practice

Jane wants to run a Global Mirror with practice from San Francisco to Houston. She plans to verify her recovery procedure for the Houston site, but she cannot stop running the Global Mirror session while she takes time to practice a recovery. By using a Global Mirror with Practice session, Jane can practice her disaster recovery scenario in Houston while the Global Mirror session runs uninterrupted. When she practices running the applications at the Houston site, Jane is better prepared to recover the data 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 plans to verify her recovery procedure for the Houston site, but she cannot stop running the Global Mirror session while she takes time to practice a recovery. By using a Global Mirror Either Direction with Two-Site Practice session, Jane can practice her disaster recovery scenario in Houston while the Global Mirror session runs uninterrupted. When she practices running the applications at the Houston site, Jane is better prepared to recover the data 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 the disaster-recovery tasks at the target site, regardless of the location of the current production site. Jane's business can run a consistent Global Mirror session from its Houston site back to San Francisco while it runs a production site at Houston.

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. 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.

Tip: 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

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

This replication method is available on only TotalStorage Enterprise Storage Server Model 800 and System Storage DS8000 storage systems. You can select ESS storage systems in only the H1 volume role. All other volume roles must use DS8000 volumes.

You can use both ESS and DS8000 volumes in the H1 volume role. If ESS andDS8000 storage systems are both used in the H1 role, the DS8000 storage system performs Incremental Resync processing, and the ESS storage system performs a full copy operation. Because you cannot use the Incremental Resync function with ESS, a full copy is required when you change from H1->H2->H3 to H1>H3 and from H2->H1->H3 to H2->H3.

Attention: If you have Tivoli Storage Productivity Center for Replication for System z sessions that contain Metro Global Mirror relationships, ensure that the session does not contain system volumes (such as paging volumes) unless you select the **Manage H1-H2 with HyperSwap** or the **Enable Hardened Freeze** option for the session. By using these options, z/OS Input/Output Supervisor (IOS) manages freeze operations for the volumes in the session, which prevents Tivoli Storage Productivity Center for Replication from freezing the volumes and possibly freezing itself.

Metro Global Mirror maintains a consistent copy of data at the remote site, with minimal effect on applications at the local site. This remote mirroring function works in combination with FlashCopy to provide a disaster-recovery solution that includes the following features:
- Fast failover and failback
- Rapid reestablishment of three-site mirroring, without production outages
- Data currency at the remote site with minimal time lag at the local site, an average of only 3 - 5 seconds for many environments
- Quick resynchronization of mirrored sites using only incremental changes

If Tivoli Storage Productivity Center for Replication runs on a z/OS operating system, you can configure a Metro Global Mirror session to control the Metro Mirror relationship between the primary and secondary site by using the HyperSwap feature. With HyperSwap enabled, a failure on the primary storage system causes an automatic HyperSwap operation, which transparently redirects 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 almost a 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 the H2 and H3 sites. You do not have to configure H2, 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 setup can reduce the workload on site H1.

Important: In Metro Global Mirror and Metro Global Mirror with Practice sessions, when the H1 is on an ESS storage system, you might risk filling up the space efficient journal volumes. Because incremental resynchronization is not available on the ESS storage system, 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 a Metro Global Mirror session.
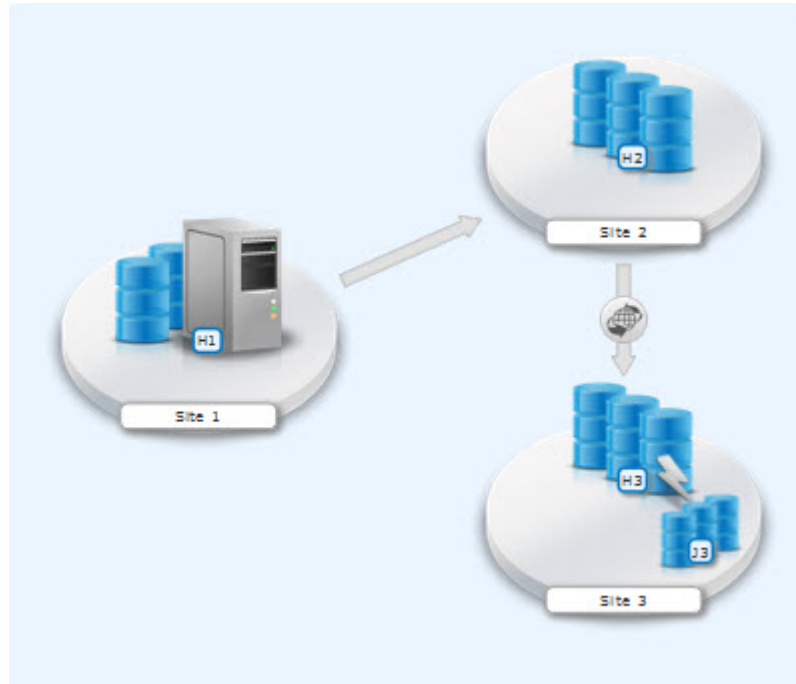
Metro Global Mirror with Practice

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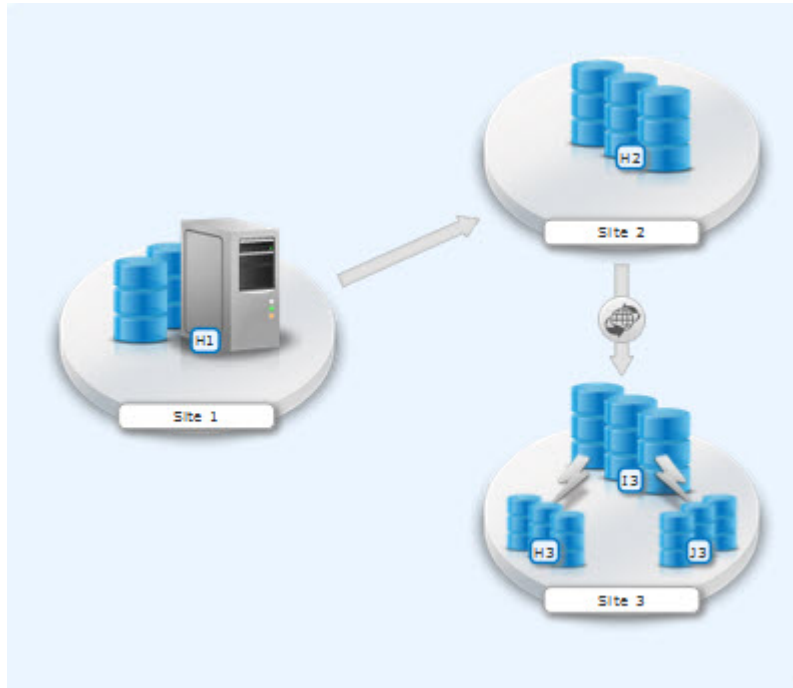
Note: In Metro Global Mirror and Metro Global Mirror with Practice sessions, when the H1 is on an ESS storage system, you might risk filling up the space efficient journal volumes. Because incremental resynchronization is not available on the ESS storage system, full copies are performed in many of the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

Examples

Read the following scenarios for information on using Metro Global Mirror for continuous, remote data replication between three sites.

Metro Global Mirror
> Although Jane works in San Francisco, she wants to run the business from either Oakland, the secondary site, or Houston, the tertiary site. Jane can use Metro Global Mirror with Failover/Failback to switch the direction of the data flow, so that she can run the business from either Oakland or Houston. Using Metro Global Mirror, Jane has zero data loss at the secondary site and minimal data loss at the tertiary site.

Metro Global Mirror with Practice
> Jane wants to run a Metro Global Mirror with Practice from San Francisco to Houston. She plans to verify her recovery procedure for the Houston site. However, she cannot stop running the Metro Global Mirror session while she takes time to practice a recovery. By using a Metro Global Mirror with Practice session, Jane can practice her disaster recovery scenario in Houston while the Metro Global Mirror session runs uninterrupted. When she practices running the applications at the Houston site and is prepared to run the applications at the Oakland site, if necessary, Jane is prepared to recover data if a disaster occurs at 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 the business from either Oakland or Houston. Using Metro Global Mirror, Jane has zero data loss at the secondary site and minimal data loss at the tertiary site.

Managing a session with HyperSwap and Open HyperSwap replication

HyperSwap and Open HyperSwap provide high availability of data if a primary disk storage system failure occurs. When a failure occurs in writing 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 minimal application effect.

Sessions that can be enabled for HyperSwap or Open HyperSwap

You can create sessions that enable swapping, which provides a session with a highly available business continuity solution.

Sessions that can enable HyperSwap

The following session types can enable HyperSwap processing:
* Basic HyperSwap
* Metro Mirror with Failover/Failback
* Metro Global Mirror
* Metro Global Mirror with Practice

To enable HyperSwap processing, the following conditions must apply:
* The volumes that are managed by the session are only TotalStorage Enterprise Storage Server, System Storage DS8000, and DS6000 storage systems volumes.
* The volumes are count key data (CKD) volumes that are attached to an IBM z/OS system.
* The z/OS system is listed on the Tivoli Storage Productivity Center for Replication Host Connections page and is in the Connected state.

Sessions that can enable Open HyperSwap

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

To enable Open HyperSwap processing, the following conditions 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 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 are established.

Setting up the environment for HyperSwap

You must set up an environment in which you can run HyperSwap processing before you can enable HyperSwap for a Tivoli Storage Productivity Center for Replication session.

To enable HyperSwap processing, the following conditions must exist:
* The HSIB and IOSHSAPI address spaces are started on the IBM z/OS system to which the storage systems are connected.

- All reserve volumes are converted to global enqueues (ENQs).
- The storage system volumes are count key data (CKD) volumes that are attached to a z/OS system. The volumes can be attached to the same z/OS system or different systems.
- If the storage system volumes are not connected to a z/OS system on which Tivoli Storage Productivity Center for Replication is installed, the required Resource Access Control Facility (RACF) settings are set on the z/OS system and the `SOCKPORT` parameter is defined for the HyperSwap management address space IOSHMCTL. This set up enables Tivoli Storage Productivity Center for Replication to connect to the z/OS system by using an IP connection. For the required RACF settings, see the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

Complete the following tasks to set up an environment in which you can run HyperSwap processing:

1. Add a connection for each z/OS system in the Tivoli Storage Productivity Center for Replication GUI. To add a connection, open the Host Connections page and click **Add Host Connection**. If the z/OS system is the system on which Tivoli Storage Productivity Center for Replication is installed, a native connection is automatically added in the GUI. If the z/OS system does not have Tivoli Storage Productivity Center for Replication installed, you must you use an IP connection to connect to the z/OS system.
2. Add a connection to the storage systems that are attached to the z/OS systems. To add a connection, open the Storage Systems page and click **Add Storage Connection**. You can connect to a storage system by using a Hardware Management Console (HMC), direct, or z/OS connection.
3. Create Basic HyperSwap, Metro Mirror with Failover/Failback, Metro Global Mirror, or Metro Global Mirror with Practice sessions and enable HyperSwap. To create a session, open the Sessions page and click **Create Session**.

Setting up the environment for Open HyperSwap

You must set up an environment in which you can run Open HyperSwap processing before you can enable Open HyperSwap for a IBM Tivoli Storage Productivity Center for Replication session.

Complete the following steps before you enable Open HyperSwap.

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

   AIX requirements

   > For a list of the supported AIX versions for each Tivoli Storage
   >
   > Productivity Center for Replication release, see ↪ IBM Spectrum Control interoperability matrix at http://www.ibm.com/support/docview.wss?uid=swg21386446 and go to the *Agents, Servers and Browsers* section.
   >
   > 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)

   System Storage DS8000 hardware requirements
   > The Open HyperSwap function requires System Storage DS8000 5.1 or later.

You cannot use Open HyperSwap with High Availability Cluster Multi-Processing (HACMP™).

2. Create connections from Tivoli Storage Productivity Center for Replication to the AIX hosts by following the instructions at "Adding a host system connection" on page 55.

3. Assign copy set volumes from the storage device to the host by using the System Storage DS8000 command-line interface (CLI) or the graphical user interface (GUI).

4. Run the AIX **cfgmgr** command to discover the volumes that are assigned to the host.

Considerations for Open HyperSwap and the AIX host:

- A single session that has the Open HyperSwap function enabled can manage multiple hosts; however, each host can be associated with only one session. Multiple hosts can share the session.

- For AIX 5.3, a single host can manage a maximum of 1024 devices that are enabled for Open HyperSwap processing on the host. A maximum of eight logical paths can be configured for each copy set in the session. For AIX 6.1, a single host can manage a maximum of 1024 devices that are enabled for Open HyperSwap processing on the host. A maximum of 16 logical paths can be configured for each copy set in the session.

- If an application on the host opened a device, a Tivoli Storage Productivity Center for Replication session for that device cannot be ended. The **Terminate** command fails. To end 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 the data to the correct volume when the copy set relationship is restored.

- Open HyperSwap processing can fail on a subset of hosts for the session and continue to run on the remaining hosts for the same session. In this situation, you must determine the best action to take if the application is writing data to volumes on the source system and to volumes on the target system.

- To enable Open HyperSwap processing on the host, see the *IBM System Storage Multipath Subsystem Device Driver User's Guide*.

Configuring timers to support Open HyperSwap:

You can configure timeout values for the storage system, Tivoli Storage Productivity Center for Replication, and AIX hosts systems. These values can affect the operation of Open HyperSwap.

The following various timeout values can affect Open HyperSwap:

Storage system quiesce timeout value
> This value is the time when the storage system starts a quiesce operation. When the timer value expires, I/O is resumed on the primary device. The default timeout value is 2 minutes, but the value can be set from 30 to 600 seconds. To set the quiesce timeout value, use the **chdev** command in the *IBM System Storage Multipath Subsystem Device Driver User's Guide*.

Storage system long busy timeout value
> This 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, 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 write I/O operation 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 HyperSwap or Open HyperSwap for a session

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

Ensure that you can use the HyperSwap or Open HyperSwap functions in your environment. See "Setting up the environment for HyperSwap" on page 97 or "Setting up the environment for Open HyperSwap" on page 98.

To enable HyperSwap or Open HyperSwap for a session, complete the following steps:

1. In the Tivoli Storage Productivity Center for Replication menu bar, click **Sessions**.
2. On the Sessions page, select the session.
3. From the **Session Actions** list, select **View/Modify** > **Properties**.
4. In the **View/Modify Properties** notebook, select the following properties on the applicable role pair options tab:

   For HyperSwap

   > **System or sysplex**
   > > Select the z/OS system or sysplex that you want to associate with the session. Commands for z/OS features are issued to this system or sysplex. Volumes in the session must be attached to the system or sysplex to enable the features.

   > **Manage H1-H2 with HyperSwap**
   > > Select this option to trigger a HyperSwap operation, which redirects application I/O to the target volumes when there is a failure on the host accessible volumes. Tivoli Storage Productivity Center for Replication uses HyperSwap to manage the H1-H2 sequence of a Metro Mirror or Metro Global Mirror session.

   > > **On Configuration Error:**

   > > > **Partition the system(s) out of the sysplex**
   > > > > Select this option to partition a new system out

of the sysplex when an error occurs because the system cannot be added to the HyperSwap configuration.

Disable HyperSwap
Select this option to prevent a HyperSwap operation from occurring.

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 enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

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 the HyperSwap operation does not complete.

Requirement: If you select this option, you must restart the system .

Disable HyperSwap after attempting backout
Select this option to enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

For Open HyperSwap

Manage H1-H2 with Open HyperSwap
Select this option to trigger an Open HyperSwap operation for volumes that are attached to an IBM AIX host. This option redirects application I/O to the target volumes when there is a failure on the host accessible volumes.

Tivoli Storage Productivity Center for Replication uses Open HyperSwap to manage the H1-H2 sequence of a Metro Mirror session.

Only volumes that are attached to host systems that are defined in the Tivoli Storage Productivity Center for Replication Host Systems panel are eligible for Open HyperSwap.

Restarting an AIX Host System that is enabled for Open HyperSwap

When an AIX host system is restarted, the host automatically opens any volumes for I/O that were open before the system was restarted. 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 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 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. The I/O might be blocked until the host can verify the primary volume in the relationship. The AIX `varyonvg` command fails to open the volumes for I/O to prevent the application from writing to an incorrect volume. 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 Target Available state or the host system does not have connectivity and is unable to determine which site is the primary site, all I/O to the volumes are blocked on the host system. The Target Available state occurs after a Hyperswap or a Recover command is issued for the session. The `varyonvg` command fails to open the volumes for I/O to prevent the application from writing to an incorrect volume.

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

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, complete one of the following tasks:
- Stop the session or remove the copy set. This action requires a full copy when the relationships are restarted.
- Remove Object Data Manager (ODM) entries by using the following command:
  ```
  odmdelete -o Volume_Equivalency
  ```

  Restriction: This command deletes copy set information; use it only for this scenario.

Planned and unplanned swaps
After a session is enabled for HyperSwap or Open HyperSwap processing and reaches Prepared state, Tivoli Storage Productivity Center for Replication loads the configuration of volumes that can be swapped to a z/OS or AIX operating system.

When the load is complete, the session can start a planned or unplanned swap. A hyperswap symbol is displayed over the arrow that shows replication between sites, as shown in the following figure.
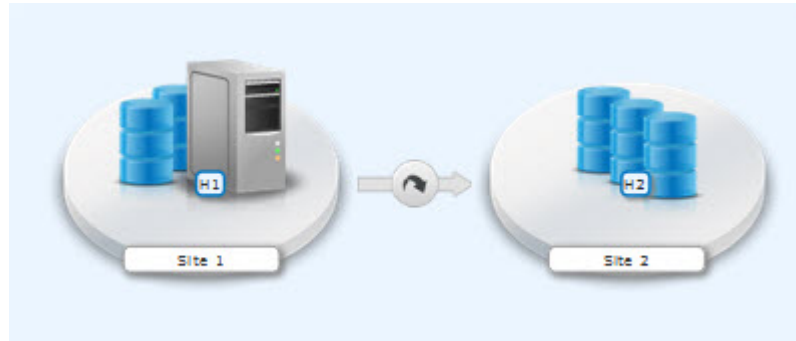
Completing a Planned Swap

When 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. When the session is in swap capable state, a planned swap can be completed by issuing the **HyperSwap** command for the session.

After you run a planned swap for z/OS HyperSwap and Open HyperSwap, the session transitions to Target Available state and all the H1-H2 pairs are in Target Available state. If the H1-H2 role pair was consistent at the time of the swap, the session has a status of Normal, which indicates that H1-H2 is consistent. If the H1-H2 role pair was not consistent at the time of the swap, the session has a status of SEVERE. The active host on the session is shown as H2.

All I/O is redirected to the H2 volumes. After a successful swap to site 2, you cannot enable the copy function to site 2. The only way to restart the copy processing is by issuing the **Start H2->H1** command. To have the volumes protected with high availability and disaster recovery again, correct the error that caused the swap to occur, and manually restart the session to begin copying data to the other site.

The following figure illustrates a planned swap operation.
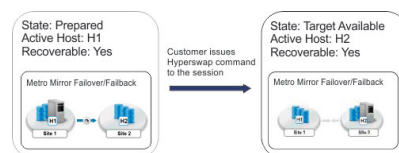
Unplanned swaps

When the session configuration is loaded on z/OS for HyperSwap processing or on AIX for Open HyperSwap processing, the session is considered swap capable. If a primary I/O error occurs, a swap occurs automatically. For HyperSwap, z/OS completes the swap and then alerts Tivoli Storage Productivity Center for Replication that a swap occurred. For Open HyperSwap, Tivoli Storage Productivity Center for Replication and the AIX host work together to complete the swap.

When an unplanned swap occurs for HyperSwap and Open HyperSwap, the session transitions to Target Available state and all the H1-H2 pairs are in Target Available state. If the H1-H2 role pair was consistent at the time of the swap, the session has a status of Normal and indicates that H1-H2 is consistent. If the H1-H2 role pair was not consistent at the time of the swap, the session might show a status of SEVERE. The active host on the session is shown as H2.

All I/O is redirected to the H2 volumes. After a successful swap to site 2, you cannot enable the copy function to site 2. Therefore, it is not possible to issue a **Start H1->H2** command. The only way to restart the copy processing is by issuing a **Start H2->H1** command. To have the volumes protected with high availability and disaster recovery again, correct the error that caused the swap to occur, and manually restart the session to begin copying data to the other site.

The following figure illustrates an unplanned swap.
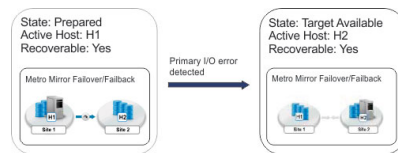
Verifying a session for a planned or unplanned swap:

You can verify whether a session is capable of a planned or unplanned swap from the z/OS console on HyperSwap or the AIX host on Open HyperSwap.

Verify the status of HyperSwap from the z/OS console:

1. View the overall status of the HyperSwap session. Issue the ds hs,status command. 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. Verify all the volumes in the configuration. Issue the `ds hs,config(detail,all)` command. 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
```

Verify the status of Open HyperSwap from the AIX host:

1. View the session association and the path that the I/O is being routed to. The path is indicated by an asterisk. Issue the `pcmpath query device` command. 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 function for a session.

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

- If you are performing maintenance on the system
- If one sysplex member cannot communicate with one or more volumes

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

1. In the Tivoli Storage Productivity Center for Replication menu bar, click **Sessions**.
2. On the Sessions page, select the sessions for which you want to disable the HyperSwap or Open HyperSwap function.
3. From the **Session Actions** list, select **View/Modify** > **Properties**.
4. On the applicable tab of the View/Modify Properties notebook, select **Disable HyperSwap** or **Disable Open HyperSwap**.

   Tip: On management servers that run on a z/OS operating system, you can disable the HyperSwap function from an MVS command prompt window by entering the **SETHS DISABLE** command.

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

You can set up an active and standby management server configuration with HyperSwap or Open HyperSwap.

Active and standby servers with HyperSwap

When the storage system is set up to connect through the z/OS interface, the connection information is sent to the standby server and a connection is attempted. The connection might fail if the standby server does not run on a z/OS operating system or does not have access to the same volumes. If the connection fails and the standby server resumes the processing, the standby server cannot manage the HyperSwap processing.

On the z/OS system, if the session configuration was successfully loaded before the HyperSwap processing, the z/OS system can manage the HyperSwap processing. If the z/OS system swaps the volumes but cannot communicate with the Tivoli Storage Productivity Center for Replication server, the session is set to Suspended or Severe state. To set the session to Target Available state, clear the **Manage H1-H2 with Hyperswap** option, and issue the **Recover** command for the session.

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, a connection is attempted between the host system and the standby server. After AIX loads the session configuration, Open HyperSwap processing is possible only if there is communication between AIX and the Tivoli Storage Productivity Center for Replication server. If a standby server resumes the processing and is unable to connect to the host system that manages the swap, the session cannot run Open HyperSwap processing. You must activate communication to the host system before the session can run Open HyperSwap processing.

Related tasks:

Chapter 3, "Managing management servers," on page 29
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, and for some commands, the storage system type.

Commands are issued to Tivoli Storage Productivity Center for Replication sessions. Any subsequent command that is issued to an individual session is not processed until the first command is completed. Some commands, such as the **Start** command, can take an extended amount of time to complete because the commands set up the hardware relationships. You can continue to issue commands to other sessions during this time. When a command is completed, the console shows 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 36. 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.
Refresh States	Refreshes the states of the role pairs that are in the session (if applicable for the session type) and refreshes the state of the session if it is incorrect. This command queries the states of the copy sets on the hardware. You are not required 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 Tivoli Storage Productivity Center for Replication to update itself. This command triggers multiple queries on the hardware, which can impact hardware performance. Do not run this command more frequently than every few minutes in each session. This command is not available if the session is in the Defined state.
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 37. FlashCopy commands

Command	Action
Flash	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: • SAN Volume Controller • Storwize V3500 • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified 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.
InitiateBackgroundCopy	Copies all tracks from the source to the target immediately, instead of waiting until the source track is written to. This command is valid *only* when the background copy is not already running.
Refresh States	Refreshes the states of the role pairs that are in the session (if applicable for the session type) and refreshes the state of the session if it is incorrect. This command queries the states of the copy sets on the hardware. You are not required 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 Tivoli Storage Productivity Center for Replication to update itself. This command triggers multiple queries on the hardware, which can impact hardware performance. Do not run this command more frequently than every few minutes in each session. This command is not available if the session is in the Defined state.
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: • SAN Volume Controller • Storwize V3500 • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified

Table 37. FlashCopy commands (continued)

Command	Action
Terminate	Removes all active 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 **InitiateBackgroundCopy** command if NOCOPY was specified, and then wait for the background copy to complete by checking the copying status of the pairs.

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 38. Snapshot session commands

Command	Action
Create Snapshot	Creates a snapshot of the volumes in the session
Refresh States	Refreshes the states of the role pairs that are in the session (if applicable for the session type) and refreshes the state of the session if it is incorrect. This command queries the states of the copy sets on the hardware. You are not required 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 Tivoli Storage Productivity Center for Replication to update itself. This command triggers multiple queries on the hardware, which can impact hardware performance. Do not run this command more frequently than every few minutes in each session. This command is not available if the session is in the Defined state.
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 39. 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. If the deleted snapshot group is the last snapshot group that is associated with the session, the session returns to the Defined state.

Table 39. 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.
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 40. 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 you reverse the direction of copying in a failover and failback session. After you issue this command, the **Start H1->H2** command becomes available.

Table 40. 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. This command is available only for Tivoli Storage Productivity Center for Replication for System z.
Refresh States	Refreshes the states of the role pairs that are in the session (if applicable for the session type) and refreshes the state of the session if it is incorrect. This command queries the states of the copy sets on the hardware. You are not required 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 Tivoli Storage Productivity Center for Replication to update itself.

This command triggers multiple queries on the hardware, which can impact hardware performance. Do not run this command more frequently than every few minutes in each session.

This command is not available if the session is in the Defined state. |
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.
Release I/O	Enables the source volume to receive data after a copy relationship is suspended. This command is not displayed if the **Release I/O after Suspend** property is selected for the session.
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.

Table 40. Metro Mirror commands (continued)

Command	Action
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.
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 41. 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 you reverse the direction of copying in a failover and failback session. After you issue this command, the **Start H1->H2** command becomes available.

Table 41. Metro Mirror with Practice commands (continued)

Command	Action
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, and 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.
Recover	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. **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 H2 volumes afterward.
Refresh States	Refreshes the states of the role pairs that are in the session (if applicable for the session type) and refreshes the state of the session if it is incorrect. This command queries the states of the copy sets on the hardware. You are not required 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 Tivoli Storage Productivity Center for Replication to update itself. This command triggers multiple queries on the hardware, which can impact hardware performance. Do not run this command more frequently than every few minutes in each session. This command is not available if the session is in the Defined state.
Release I/O	Enables the source volume to receive data after a copy relationship is suspended. This command is not displayed if the **Release I/O after Suspend** property is selected for the session.
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 41. Metro Mirror with Practice commands (continued)

Command	Action
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:** The Metro Mirror command uses a freeze command on the ESS, DS6000, and 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. 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.
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.
Terminate	Terminates all copy relationships on the hardware.
Recover	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. **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 H2 volumes afterward.

Global Mirror commands

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

Table 42. 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 you reverse the direction of copying in a failover and failback session. After you issue this command, the **Start H1->H2** command becomes available.

Table 42. Global Mirror commands (continued)

Command	Action
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 state 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.
Refresh States	Refreshes the states of the role pairs that are in the session (if applicable for the session type) and refreshes the state of the session if it is incorrect. This command queries the states of the copy sets on the hardware. You are not required 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 Tivoli Storage Productivity Center for Replication to update itself. This command triggers multiple queries on the hardware, which can impact hardware performance. Do not run this command more frequently than every few minutes in each session. This command is not available if the session is in the Defined state.
Start	Establishes all relationships in a single-direction session and begins the process necessary to start forming consistency groups on the hardware.
Start H1->H2	Starts copying data from H1 to H2 in a Global Mirror Failover/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 TotalStorage Enterprise Storage Server Model 800, System Storage DS6000, and System Storage DS8000. If a recover was performed on a session such that the production site is now H2, you can issue a **Start H2->H1** command 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 action puts the relationships into a synchronized state and suspends them consistently. Sessions are consistent when a H2->H1 copy occurs for the following storage systems: • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified

Table 42. Global Mirror commands (continued)

Command	Action
StartGC H1->H2 This command is available only for Global Mirror Failover/Failback and Global Mirror Failover/Failback with Practice sessions.	Establishes Global Copy relationships between site 1 and site 2 and begins asynchronous data replication from H1 to H2. To change the session state from Preparing to Prepared, you must issue the **Start H1->H2** command and the session must begin to form consistency groups. There is no disaster recovery protection for Global Copy relationships. If a disaster such as the loss of a primary storage system or a link failure between the sites occurs, the session might be inconsistent when you issue the **Recover** command. This command is available for the following storage systems: • TotalStorage Enterprise Storage Server Model 800 • System Storage DS6000 • System Storage DS8000

Table 42. Global Mirror commands (continued)

Command	Action
Suspend	Pauses the Global Mirror master session, which causes the session to stop forming consistency groups, and suspends the H1 to H2 Global Copy pairs. **If you are using a System Storage DS8000:** The **Suspend** command invokes a pause command for System Storage DS8000. The pause command that is invoked depends on the System Storage DS8000 microcode level. If the following conditions are true, the **Suspend** command invokes a command that is equivalent to the System Storage DS8000 command `pausegmir -withsecondary`: • All primary volumes that are in the session are on a System Storage DS8000 storage system with a microcode level that includes the pause with secondary consistency option. To determine whether this option is available for your storage system, refer to the System Storage DS8000 documentation for the microcode level that you are using. • The pause with secondary consistency option is enabled on the storage system. If these conditions are not true, the **Suspend** command invokes a command that is equivalent to the System Storage DS8000 command `pausegmir`. Both commands temporarily pause the formation of consistency groups after the current consistency group is formed. However, the command for a pause with secondary consistency creates a consistent data set on the secondary volumes. If you do not want to suspend the Global Copy pairs when a Global Mirror session is suspended, open the `path_prefix`/opt/Tivoli/RM/wlp/usr/servers/ replicationServer/properties/rmserver.properties file and add the following property: `csm.server.sus_gc_pairs_on_gm_pause = false` **Tip:** The pause with secondary consistency option requires that the Global Copy pairs are suspended. If you do not want to use the pause with secondary consistency option for storage systems on which the option is enabled, set the `csm.server.sus_gc_pairs_on_gm_pause` property to `false`.
Terminate	Removes all physical copies and relationships from the hardware during an active session. If you want the targets to be data consistent before you remove 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 43. 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.
Enable Copy to Site 2	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 H1->H2** command becomes available.
Flash	The Flash command ensures that all I2s are consistent, and then flashes 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. For Global Mirror with Practice sessions for TotalStorage Enterprise Storage Server Model 800, System Storage DS6000, and System Storage DS8000, the Flash command temporarily pauses the formation of consistency groups.
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 continue to hold the consistent data and can be flashed again to H2 by using the **Flash** command.
Refresh States	Refreshes the states of the role pairs that are in the session (if applicable for the session type) and refreshes the state of the session if it is incorrect. This command queries the states of the copy sets on the hardware. You are not required 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 Tivoli Storage Productivity Center for Replication to update itself. This command triggers multiple queries on the hardware, which can impact hardware performance. Do not run this command more frequently than every few minutes in each session. This command is not available if the session is in the Defined state.
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. The session starts the Global Mirror master so that the hardware begins to form consistency groups, to ensure that consistent data is at site 2.

Table 43. Global Mirror with Practice commands (continued)

Command	Action
Start H2->H1	Starts copying data from H2 to H1 in a failover and failback session. If a recover was performed on a session such that the production site is now H2, you can issue a **Start H2->H1** command 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:** TotalStorage Enterprise Storage Server Model 800, System Storage DS6000, and System Storage 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 command to the session. The relationships are put into a synchronized state and are suspended consistently.
StartGC H1->H2 This command is available only for Global Mirror Failover/Failback and Global Mirror Failover/Failback with Practice sessions.	Establishes Global Copy relationships between site 1 and site 2 and begins asynchronous data replication from H1 to I2. To change the session state from Preparing to Prepared, you must issue the `Start H1->H2` command and the session must begin to form consistency groups. There is no disaster recovery protection for Global Copy relationships. If a disaster such as the loss of the primary Tivoli Storage Productivity Center for Replication server occurs, the session might be inconsistent when you issue the `Recover` command. This command is available for the following storage systems: • TotalStorage Enterprise Storage Server Model 800 • System Storage DS6000 • System Storage DS8000

Table 43. Global Mirror with Practice commands (continued)

Command	Action
Suspend	Pauses the Global Mirror master session, which causes the session to stop forming consistency groups, and suspends the H1 to H2 Global Copy pairs. **If you are using a System Storage DS8000:** The **Suspend** command invokes a pause command for System Storage DS8000. The pause command that is invoked depends on the System Storage DS8000 microcode level. If the following conditions are true, the **Suspend** command invokes a command that is equivalent to the System Storage DS8000 command **pausegmir -withsecondary**: • All primary volumes that are in the session are on a System Storage DS8000 storage system with a microcode level that includes the pause with secondary consistency option. To determine whether this option is available for your storage system, refer to the System Storage DS8000 documentation for the microcode level that you are using. • The pause with secondary consistency option is enabled on the storage system. If these conditions are not true, the **Suspend** command invokes a command that is equivalent to the System Storage DS8000 command **pausegmir**. Both commands temporarily pause the formation of consistency groups after the current consistency group is formed. However, the command for a pause with secondary consistency creates a consistent data set on the secondary volumes. If you do not want to suspend the Global Copy pairs when a Global Mirror session is suspended, open the *path_prefix*/opt/Tivoli/RM/wlp/usr/servers/replicationServer/ properties/rmserver.properties file and add the following property: csm.server.sus_gc_pairs_on_gm_pause = false **Tip:** The pause with secondary consistency option requires that the Global Copy pairs are suspended. If you do not want to use the pause with secondary consistency option for storage systems on which the option is enabled, set the **csm.server.sus_gc_pairs_on_gm_pause** property to **false**.
Terminate	Removes all physical copies and relationships that are on the hardware.

Metro Global Mirror commands

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

Table 44. 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 you reverse the direction of copying in a failover and failback session. After you issue this command, the **Start H1->H2->H3** command becomes available.
HyperSwap	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. This command is available only for Tivoli Storage Productivity Center for Replication for System z.
RecoverH1	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.
RecoverH2	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.
RecoverH3	Specifying H3 makes the H3 volume TargetAvailable. Metro Global Mirror (when H1->H2->H3 is run) can then move production to the H3 set of volumes. Because Tivoli Storage Productivity Center for Replication processing differs depending on the recovery site, the site designation is added to the **Recover** command so that Tivoli Storage Productivity Center for Replication can prepare for the failback.

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. |
Re-enable Copy to Site 1	After you issue 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.
Re-enable Copy to Site 2	After you issue 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 you issue 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.

Table 44. Metro Global Mirror commands (continued)

Command	Action
Refresh States	Refreshes the states of the role pairs that are in the session (if applicable for the session type) and refreshes the state of the session if it is incorrect. This command queries the states of the copy sets on the hardware. You are not required 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 Tivoli Storage Productivity Center for Replication to update itself. This command triggers multiple queries on the hardware, which can impact hardware performance. Do not run this command more frequently than every few minutes in each session. This command is not available if the session is in the Defined state.
Release I/O	Enables the source volume to receive data after a copy relationship is suspended. This command is not displayed if the **Release I/O after Suspend** property is selected for the session.
Start H1->H2->H3 This command 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 relationship includes the J3 volume to complete the Global Mirror configuration. (The J3 volume role is the journal volume at site 3.) The **Start H1->H2->H3** command can be used from some Metro Global Mirror configurations to transition back to the starting H1->H2->H3 configuration. This command is valid only when the session is in a defined, preparing, prepared, target available, or suspended state.
Start H1->H3	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. 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 loses the near-zero data loss protection that is achieved with Metro Mirror when H1->H2->H3 is run. However, data consistency is still maintained at the remote site with the Global Mirror solution. 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. For example, for an unplanned HyperSwap. The Global Mirror session is continued. This is a host-volume change so this command is valid only when you restart the H1->H3 configuration or from the TargetAvailable H2->H1->H3 state.

Table 44. Metro Global Mirror commands (continued)

Command	Action
Start H2->H3	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 an unplanned HyperSwap. The Global Mirror session is continued. This session is a host-volume change so this command is valid only when you restart the H1->H3 configuration or from the TargetAvailable H2->H1>H3 state.
	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, this command can be used when the session is in the preparing, prepared, and suspended states. The **Start H2->H1->H3** command can be used to transition back to the starting H2->H1->H3 configuration.
Start H2->H1->H3 This command is the Metro Global Mirror start command.	This configuration 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, the J3 volume completes the Global Mirror configuration.
Start H3->H1->H2	After a recovery to H3 is completed, 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.
	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.

Table 44. Metro Global Mirror commands (continued)

Command	Action
SuspendH2H3	Pauses the Global Mirror master session, which causes the session to stop forming consistency groups, and suspends the H2 to H3 Global Copy pairs. **If you are using a System Storage DS8000:** The **Suspend** command invokes a pause command for System Storage DS8000. The pause command that is invoked depends on the System Storage DS8000 microcode level. If the following conditions are true, the **Suspend** command invokes a command that is equivalent to the System Storage DS8000 command `pausegmir -withsecondary`: • All primary volumes that are in the session are on a System Storage DS8000 storage system with a microcode level that includes the pause with secondary consistency option. To determine whether this option is available for your storage system, refer to the System Storage DS8000 documentation for the microcode level that you are using. • The pause with secondary consistency option is enabled on the storage system. If these conditions are not true, the **Suspend** command invokes a command that is equivalent to the System Storage DS8000 command `pausegmir`. Both commands temporarily pause the formation of consistency groups after the current consistency group is formed. However, the command for a pause with secondary consistency creates a consistent data set on the secondary volumes. If you do not want to suspend the Global Copy pairs when a Global Mirror session is suspended, open the *path_prefix*`/opt/Tivoli/RM/wlp/usr/servers/replicationServer/properties/rmserver.properties` file and add the following property: `csm.server.sus_gc_pairs_on_gm_pause = false` **Tip:** The pause with secondary consistency option requires that the Global Copy pairs are suspended. If you do not want to use the pause with secondary consistency option for storage systems on which the option is enabled, set the `csm.server.sus_gc_pairs_on_gm_pause` property to `false`.
SuspendH1H3	Pauses the Global Mirror master session, which causes the session to stop forming consistency groups, and suspends the H1 to H3 Global Copy pairs. The information that is specific to System Storage DS8000 in the previous description for the **SuspendH2H3** command is also applicable to this command.

Metro Global Mirror with Practice commands

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

Table 45. 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 you reverse the direction of copying in a failover and failback session.
Flash	This command is available in the following states: • 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 can be 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. 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 can be issued to complete the flash of the consistent data from I3 to H3. Issuing a **Flash** command on a Global Mirror Practice session for ESS, DS6000, and DS8000 temporarily pauses 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.
HyperSwap	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. This command is available only for Tivoli Storage Productivity Center for Replication for System z.
RecoverH1	Specifying H1 makes the H1 volume TargetAvailable. When H1->H2->H3 is run, 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. 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 reaches a Target Available state, and you can attach your volumes on Site 3.

Table 45. Metro Global Mirror with Practice commands (continued)

Command	Action
RecoverH2	Specifying H2 makes the H2 volume TargetAvailable. When H1->H2->H3 is run, 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.
RecoverH3	Specifying H3 makes the H3 volume the TargetAvailable. When H1->H2->H3 is run, 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.
Re-enable Copy to Site 1	After you issue 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.
Re-enable Copy to Site 2	After you issue 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 you issue 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.
Refresh States	Refreshes the states of the role pairs that are in the session (if applicable for the session type) and refreshes the state of the session if it is incorrect. This command queries the states of the copy sets on the hardware. You are not required 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 Tivoli Storage Productivity Center for Replication to update itself.

This command triggers multiple queries on the hardware, which can impact hardware performance. Do not run this command more frequently than every few minutes in each session.

This command is not available if the session is in the Defined state. |
| **Release I/O** | Enables the source volume to receive data after a copy relationship is suspended. This command is not displayed if the **Release I/O after Suspend** property is selected for the session. |
| **Start H1->H2->H3**

This command is the 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 relationship includes the J3 volume to complete the Global Mirror configuration. (The J3 volume role is the journal volume at site 3.) The **Start H1->H2->H3** command can be used from some Metro Global Mirror configurations to return to the starting H1>H2>H3 configuration.

This command is valid only when the session is in a defined, preparing, prepared, target available, or suspended state. |

Table 45. Metro Global Mirror with Practice commands (continued)

Command	Action
Start H1->H3	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. 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 completed, the session is running Global Mirror from H1 to H3 and loses the near-zero data loss protection that is achieved with Metro Mirror when H1->H2->H3 is run. However, data consistency is still maintained at the remote site with the Global Mirror solution. 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. For example, for an unplanned HyperSwap. The Global Mirror session is continued. This is a host-volume change so this command is valid only when you restart the H1->H3 configuration or from the TargetAvailable H2->H1->H3 state.
Start H2->H3	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 an unplanned HyperSwap. The Global Mirror session is continued. This is a host-volume change so this command is valid only when you restart the H1>H3 configuration or from the TargetAvailable H2->H1->H3 state. 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, this command can be used when the session is in the preparing, prepared, and suspended states. The **Start H2->H1->H3** command can be used to return to the starting H2->H1->H3 configuration.
Start H2->H1->H3 This command is the Metro Global Mirror start command.	This configuration 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, the J3 volume completes the Global Mirror configuration.

Table 45. Metro Global Mirror with Practice commands (continued)

Command	Action
Start H3->H1->H2	After a recovery to H3 is completed, 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. 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.
SuspendH2H3	Pauses the Global Mirror master session, which causes the session to stop forming consistency groups, and suspends the H2 to H3 Global Copy pairs. **If you are using a System Storage DS8000:** The **Suspend** command invokes a pause command for System Storage DS8000. The pause command that is invoked depends on the System Storage DS8000 microcode level. If the following conditions are true, the **Suspend** command invokes a command that is equivalent to the System Storage DS8000 command `pausegmir -withsecondary`: • All primary volumes that are in the session are on a System Storage DS8000 storage system with a microcode level that includes the pause with secondary consistency option. To determine whether this option is available for your storage system, refer to the System Storage DS8000 documentation for the microcode level that you are using. • The pause with secondary consistency option is enabled on the storage system. If these conditions are not true, the **Suspend** command invokes a command that is equivalent to the System Storage DS8000 command `pausegmir`. Both commands temporarily pause the formation of consistency groups after the current consistency group is formed. However, the command for a pause with secondary consistency creates a consistent data set on the secondary volumes. If you do not want to suspend the Global Copy pairs when a Global Mirror session is suspended, open the *path_prefix*`/opt/Tivoli/RM/wlp/usr/servers/replicationServer/properties/rmserver.properties` file and add the following property: `csm.server.sus_gc_pairs_on_gm_pause = false` **Tip:** The pause with secondary consistency option requires that the Global Copy pairs are suspended. If you do not want to use the pause with secondary consistency option for storage systems on which the option is enabled, set the `csm.server.sus_gc_pairs_on_gm_pause` property to `false`.

Table 45. Metro Global Mirror with Practice commands (continued)

Command	Action
SuspendH1H3	Pauses the Global Mirror master session, which causes the session to stop forming consistency groups, and suspends the H1 to H3 Global Copy pairs. The information that is specific to System Storage DS8000 in the previous description for the **SuspendH2H3** command is also applicable to this command.
Terminate	This command terminates all copy relationships on the hardware.

Metro Mirror multi-target commands

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

Table 46. Metro Mirror multi-target command

Command	Action
Confirm Production at Site 1	Confirms that you want to establish site 1 as the production site after you issue suspend and recover commands for the session.
Confirm Production at Site 2	Confirms that you want to establish site 2 as the production site after you issue suspend and recover commands for the session.
Confirm Production at Site 3	Confirms that you want to establish site 3 as the production site after you issue suspend and recover commands for the session.
FailoverH1	Completes a failover operation to the H1 target site. This command enables read and write access to the target volumes, while maintaining a recording of changes to the volumes that prevents a full copy when the relationship is restarted. You can use this command to practice your disaster recovery capabilities without initiating a site switch. To switch sites, the RecoverH1 command is required.
FailoverH2	Completes a failover operation to the H2 target site. This command enables read and write access to the target volumes, while maintaining a recording of changes to the volumes that prevents a full copy when the relationship is restarted. You can use this command to practice your disaster recovery capabilities without initiating a site switch. To switch sites, the RecoverH2 command is required.
FailoverH3	Completes a failover operation to the H3 target site. This command enables read and write access to the target volumes, while maintaining a recording of changes to the volumes that prevents a full copy when the relationship is restarted. You can use this command to practice your disaster recovery capabilities without initiating a site switch. To switch sites, the RecoverH3 command is required.
HyperSwap H1	Triggers a HyperSwap operation where I/O is redirected from the source volume to the target volume H1, without affecting the application using those volumes.
HyperSwap H2	Triggers a HyperSwap operation where I/O is redirected from the source volume to the target volume H2, without affecting the application using those volumes.
HyperSwap H3	Triggers a HyperSwap operation where I/O is redirected from the source volume to the target volume H3, without affecting the application using those volumes.

Table 46. Metro Mirror multi-target command (continued)

Command	Action
RecoverH1	Recovers the session to the H1 target site. This command makes the target host volumes consistent and available for access as the new production site. Upon completion of this command, the session becomes Target Available.
RecoverH2	Recovers the session to the H2 target site. This command makes the target host volumes consistent and available for access as the new production site. Upon completion of this command, the session becomes Target Available.
RecoverH3	Recovers the session to the H3 target site. This command makes the target host volumes consistent and available for access as the new production site. Upon completion of this command, the session becomes Target Available.
Re-enable Original Direction	Sets the production site back to the original site and re-enables the start commands in the original direction. You can use this command after you issue the suspend, recover, and confirm commands.
Refresh States	Refreshes the states of the role pairs that are in the session (if applicable for the session type) and refreshes the state of the session if it is incorrect. This command queries the states of the copy sets on the hardware. You are not required 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 Tivoli Storage Productivity Center for Replication to update itself.

This command triggers multiple queries on the hardware, which can impact hardware performance. Do not run this command more frequently than every few minutes in each session.

This command is not available if the session is in the Defined state. |
Release I/O	Enables the source volume to receive data after a copy relationship is suspended. This command is not displayed if the **Release I/O after Suspend** property is selected for the session.
Set Production to Site 1	Establishes Site 1 as the production site while the session is in the Defined state. After you issue this command, the **Start H1->H2** and **Start H1->H3** commands become available.
Set Production to Site 2	Establishes Site 2 as the production site while the session is in the Defined state. After you issue this command, the **Start H2->H1** and **Start H2->H3** commands become available.
Set Production to Site 3	Establishes Site 3 as the production site while the session is in the Defined state. After you issue this command, the **Start H3->H1** and **Start H3->H2** commands become available.
Start	Establishes Metro Mirror relationships between the volumes on the production site and the volumes at the other sites in the session, and begins data replication to those sites. For example, if Site 1 is the production site, data replication occurs between the H1 and H2 volumes and the H1 and H3 volumes at the same time.
Start H1->H2	Establishes Metro Mirror relationships between the H1 volumes and the H2 volumes, and begins data replication from H1 to H2.

Table 46. Metro Mirror multi-target command (continued)

Command	Action
Start H1->H3	Establishes Metro Mirror relationships between the H1 volumes and the H3 volumes, and begins data replication from H1 to H3.
Start H2->H1	Establishes Metro Mirror relationships between the H2 volumes and the H1 volumes and starts data replication from H2 to H1.
Start H2->H3	Establishes Metro Mirror relationships between the H2 volumes and the H3 volumes and starts data replication from H2 to H3.
Start H3->H1	Establishes Metro Mirror relationships between the H3 volumes and the H1 volumes and starts data replication from H3 to H1.
Start H3->H2	Establishes Metro Mirror relationships between the H3 volumes and the H2 volumes and starts data replication from H3 to H2.
StartGC	Establishes Global Copy relationships between the volumes on the production site and the volumes at the other sites in the session, and begins asynchronous data replication to those sites. For example, if Site 1 is the production site, data replication occurs between the H1 and H2 volumes and the H1 and H3 volumes at the same time. The session remains in the Preparing state until you issue a Start, Suspend, or Terminate command.
StartGC H1->H2	Establishes Global Copy relationships between the H1 volumes and the H2 volumes, and begins asynchronous data replication from H1 to H2. The session remains in the Preparing state until you issue a Start, Suspend, or Terminate command.
StartGC H1->H3	Establishes Global Copy relationships between the H1 volumes and the H3 volumes, and begins asynchronous data replication from H1 to H3. The session remains in the Preparing state until you issue a Start, Suspend, or Terminate command.
StartGC H2->H1	Establishes Global Copy relationships between the H2 volumes and the H1 volumes, and begins asynchronous data replication from H2 to H1. The session remains in the Preparing state until you issue a Start, Suspend, or Terminate command.
StartGC H2->H3	Establishes Global Copy relationships between the H2 volumes and the H3 volumes, and begins asynchronous data replication from H2 to H3. The session remains in the Preparing state until you issue a Start, Suspend, or Terminate command.
StartGC H3->H1	Establishes Global Copy relationships between the H3 volumes and the H1 volumes, and begins asynchronous data replication from H3 to H1. The session remains in the Preparing state until you issue a Start, Suspend, or Terminate command.

Table 46. Metro Mirror multi-target command (continued)

Command	Action
StartGC H3->H2	Establishes Global Copy relationships between the H3 volumes and the H2 volumes, and begins asynchronous data replication from H3 to H2. The session remains in the Preparing state until you issue a Start, Suspend, or Terminate command.
Stop	Stops updates to all the target volumes of the pairs in a session. The following information applies to this and the following **Stop** commands: • Targets are not consistent after the stop. • You can issue the command at any point during an active session.
StopH1H2	Stops updates from H1 to H2 volumes in a session.
StopH1H3	Stops updates from H1 to H3 volumes in a session.
StopH2H1	Stops updates from H2 to H1 volumes in a session.
StopH2H3	Stops updates from H2 to H3 volumes in a session.
StopH3H1	Stops updates from H3 to H1 volumes in a session.
StopH3H2	Stops updates from H3 to H2 volumes in a session.
Suspend	Suspends updates to all the target volumes of the pairs in a session. The following information applies to this and the following **Suspend** commands: • You can issue the command at any point during an active session. • 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.
SuspendH1H2	Suspends updates from H1 to H2 volumes in a session.
SuspendH1H3	Suspends updates from H1 to H3 volumes in a session.
SuspendH2H1	Suspends updates from H2 to H1 volumes in a session.
SuspendH2H3	Suspends updates from H2 to H3 volumes in a session.
SuspendH3H1	Suspends updates from H3 to H1 volumes in a session.
SuspendH3H2	Suspends updates from H3 to H2 volumes in a session.

Table 46. Metro Mirror multi-target command (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 you remove their relationship, issue the following commands before you issue the **Terminate** command: • **Suspend** • **RecoverHx** Where x is the site to which you want to recover.
TerminateH1H2	Removes all H1>H2 or H2>H1 relationships from the hardware during an active session. If you want the target to be data consistent before you remove their relationship, issue the following commands before you issue the **Terminate** command: • **SuspendH1H2** • **Recoverx** Where x is the site to which you want to recover.
TerminateH1H3	Removes all H1>H3 or H3>H1 relationships from the hardware during an active session. If you want the target to be data consistent before you remove their relationship, issue the following commands before you issue the **Terminate** command: • **SuspendH1H3** • **Recoverx** Where x is the site to which you want to recover.
TerminateH2H3	Removes all H2>H3 or H3>H2 relationships from the hardware during an active session. • **SuspendH2H3** • **Recoverx** Where x is the site to which you want to recover.

Example: Setting the production site and starting a Metro Mirror copy to both targets

In this example, the production site is site 1 and H1>H2 and H1>H3 are in the Defined state. To change to the production site to site 3 and start a Metro Mirror copy to both H1 and H2, issue the commands in the order that is presented:

1. **Set Production to Site 3**
2. **Start**

Example: Setting the production site and starting a Global Copy to both targets

In this example, the production site is site 1 and H1>H2 and H1>H3 are in the Defined state. To change to the production site to site 3 and start a Global Copy to both H1 and H2, issue the commands in the order that is presented:

1. **Set Production to Site 3**
2. **Start GC**

Example: Setting the production site and starting a Metro Mirror copy to a single target

In this example, the production site is site 1 and H1>H2 and H1>H3 are in the Defined state. To change to the production site to site 3 and start a Metro Mirror copy only to H2, issue the commands in the order that is presented:

1. **Set Production to Site 3**
2. **Start H3->H2**

Example: Setting the production site and starting Metro Mirror copy to one target and a Global Copy to the second target

In this example, the production site is site 1 and H1>H2 and H1>H3 are in the Defined state. To change to the production site to site 3 and start a Metro Mirror copy to H2 and a Global Copy to H1, issue the commands in the order that is presented:

1. **Set Production to Site 3**
2. **Start H3->H2**
3. **Start GC H3->H1**

Example: Changing the production site after recovering to the site

In this example, the production site is site 1 and H1>H2 and H1>H3 are active. To recover to site 2 and confirm the production site as site 2, issue the

1. **Suspend**
2. **RecoverH2**
3. **Confirm Production at Site 2**

Example: Re-enabling the original production site

In this example, the production site is site 2 and H2>H1 and H2>H3 are active after issuing the commands in the preceding example. To re-enable H1 as the production site, issue the **Re-enable Original Direction** command.

Example: Completing a failover to a single target for practice

In this example, the production site is site 1 and H1>H2 and H1>H3 are active. To failover to site 3, issue the following commands in the order that is presented:

1. **SuspendH1H3**
2. **FailoverH3**

If the production site is site 2 and H2>H1 and H2>H3 are active, issue the following commands to failover to site 3:

1. **SuspendH2H3**
2. **FailoverH3**

Example: Suspending the copy of data to both target sites

In this example, the production site is site 1 and H1>H2 and H1>H3 are active. To suspend the copy of data to both target sites, issue the **Suspend** command.

Example: Suspending the copy of data to a single target site

In this example, the production site is site 1 and H1>H2 and H1>H3 are active. To suspend the copy of data only from H1 to H3, issue the **SuspendH1H3** command.

Example: Stopping the copy of data to both target sites

In this example, the production site is site 1 and H1>H2 and H1>H3 are active. To stop the copy of data to both target sites, issue the **Stop** command.

Example: Stopping the copy of data to a single target site

In this example, the production site is site 1 and H1>H2 and H1>H3 are active. To stop the copy of data only from H1 to H3, issue the **StopH1H3** command.

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

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

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 **No Copy** 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.

To add a FlashCopy session, you create the session, add copy sets to the session, and then select the properties for the session.

1. Follow these steps to create a FlashCopy session:

   a. In the menu bar, click **Sessions**.

   b. On the Sessions page, click **Create Session**.

   c. In the Create Session window, select the following options:

      Hardware type
      Select the type of storage system for the session:
      - DS8000, DS6000, ESS 800
      - SAN Volume Controller
      - Storwize Family (for storage systems in the IBM Storwize family)
      - Storwize V7000 Unified

      Session type
      Select **FlashCopy**.

      Session name
      Enter a name to identify the session. The session name can be up to 250 alphanumeric characters. The default name for a new session is `newsession`.

      Site 1 Location
      For each site that is in the session, select the location for the site or **None**. The site location, if available, is set by storage system on the Storage Systems page.

      You can associate a location with each storage system and each site in a session. This site awareness ensures that you can select only storage systems that have either a matching site location for each role or have no defined location when you add copy sets to a session.

      Session image
      Shows an image that represents the session. The session image is displayed when you select a session type. The image is a visual aid to help you create your session. The lightning bolt in the image represents FlashCopy replication.

      If you selected a location for the site that is in the session, the site is labeled with the site location. If you did not select site locations, the sites are labeled with **Site** x, where x is the site number. You can hover the mouse pointer over the site location to view the text in a larger format.

   d. Click **OK**.

   e. If the session was successfully created, click **Launch Add Copy Sets Wizard** and continue to the next step. If the session was not created, refer to the displayed message text to fix the problem.

2. Follow these steps to add copy sets to the session:

   a. On the Select Host1 page of the Add Copy Sets wizard, complete the following information. The field names that are displayed depend on the storage system type. When you complete the information, click **Next**.

      Storage system
      Select a storage system. If the volume role has a site location that is assigned to it, you can select a storage system that is assigned to the same location as the role or a storage system that is not assigned to a location. Storage systems that have a location are listed under the

location. Storage systems that do not have a location are listed under **None**. If the role does not have a site location, you can select any storage system.

Logical storage system or I/O Group
Select a logical subsystem (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.

The use of extent space-efficient volumes in volume roles is dependent on the capabilities of the storage system.

Session image
Shows an image that represents the session in which the role for which you are selecting volumes is highlighted. This image shows how many roles 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 **Browse** to select the CSV file.

b. On the Choose Target1 page, select the target storage system, LSS or I/O group, and volume. Click **Next**.

c. 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. On this page, the **Storage system** and **Logical storage system or I/O Group** lists are populated with the values from the previously selected copy set. Repeat the previous steps to complete the Choose Host1 and Choose Target1 pages and add additional copy sets.

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.

   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**.

3. Follow these steps to add properties for the session:

   a. On the Sessions page, select the session.

   b. From the **Session Actions** list, select **View/Modify** > **Properties**.

   c. In the View/Modify Properties notebook, select the options that you want. For a description of the options, see "Viewing session properties" on page 172.

   d. Click **OK**.

Related concepts:

"Sessions" on page 63
A *session* completes a specific type of data replication for a specific set of volumes. During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a copy set. A session can contain one or more copy sets.

"Session types" on page 12
Tivoli Storage Productivity Center for Replication provides several methods to replicate data. The type of data replication that is associated with a session is known as the *session type*.

Creating a Snapshot session and adding copy sets

A Snapshot session 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.

To add a Snapshot session, you create the session and then add copy sets to the session.

1. Follow these steps to create a Snapshot session:

   a. In the menu bar, click **Sessions**.

   b. On the Sessions page, click **Create Session**.

   c. In the Create Session window, select the following options:

      Hardware type
         Select **XIV**.

      Session type
         Select **Snapshot**.

      Session name
         Enter a name to identify the session. The session name can be up to 58 alphanumeric characters. The default name for a new session is `newsession`.

      Site 1 Location
         For each site that is in the session, select the location for the site or **None**. The site location, if available, is set by storage system on the Storage Systems page.

You can associate a location with each storage system and each site in a session. This site awareness ensures that only storage systems that have either a matching site location for each role or have no defined location are displayed in the copy set wizard.

Session image

Shows an image that represents the session. The session image is displayed when you select a session type. The image is a visual aid to help you create your session.

If you selected a location for the site that is in the session, the site is labeled with the site location. If you did not select site locations, the sites are labeled with **Site** x, where x is the site number. You can hover the mouse pointer over the site location to view the text in a larger format.

d. Click **OK**.

e. If the session was successfully created, click **Launch Add Copy Sets Wizard** and continue to the next step. If the session was not created, refer to the displayed message text to fix the problem.

2. Follow these steps to add copy sets to the session:

a. From 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.

b. From the **Host1 pool** list, select the pool that contains the volumes.

c. From the **Host1 volume** list, select the volumes. To select multiple volumes, press Ctrl or Shift and click the volumes in the list.

d. If you want to import copy sets from a comma-separated value (CSV) file, click **Use a CSV file to import copy sets**. Click **Browse** to select the CSV file and click **Next**.

e. On the Matching Results page, click **Next** if the match was successful.

f. 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 lists on this page are populated with the values from the previously selected copy set. Repeat the previous steps to complete the Choose Host1 page and add additional copy sets.

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.

g. On the Confirm page, the number of copy sets to be added is displayed. Click **Next**.

h. A progress bar is displayed. When the copy sets are added, review the results and click **Finish**.

If you want to add a description for the session, select the session on the Sessions page. From the **Session Actions** list, select **View/Modify** > **Properties**.

Related concepts:

"Sessions" on page 63
A *session* completes a specific type of data replication for a specific set of volumes. During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a copy set. A session can contain one or more copy sets.

"Session types" on page 12
Tivoli Storage Productivity Center for Replication provides several methods to replicate data. The type of data replication that is associated with a session is known as the *session type*.

Creating a Metro Mirror session and adding copy sets

A Metro Mirror session is a method of synchronous, remote data replication that operates between two sites that are up to 300 KM apart.

To add a Metro Mirror session, you create the session, add copy sets to the session, and then select the properties for the session.

1. Follow these steps to create a Metro Mirror session:
   a. In the menu bar, click **Sessions**.
   b. On the Sessions page, click **Create Session**.
   c. In the Create Session window, select the following options:

   Hardware type
   Select the type of storage system for the session:
   - DS8000, DS6000, ESS 800
   - SAN Volume Controller
   - Storwize Family (for storage systems in the IBM Storwize family)
   - Storwize V7000 Unified
   - XIV

   Session type
   Select a Metro Mirror session type.

Session name

Enter a name to identify the session. For sessions that contain an IBM XIV Storage System, the session name can be up to 58 alphanumeric characters. For sessions that contain other storage system types, the session name can be up to 250 alphanumeric characters. The default name for a new session is `newsession`.

Site x Location

For each site that is in the session, select the location for the site or **None**. The site location, if available, is set by storage system on the Storage Systems page.

You can associate a location with each storage system and each site in a session. This site awareness ensures that you can select only storage systems that have either a matching site location for each role or have no defined location when you add copy sets to a session.

Session image

Shows an image that represents the session. The session image is displayed when you select a session type. The image is a visual aid to help you create your session.

This image shows the number of roles in the session and how the roles are distributed between the sites. The arrows indicate the type of copy between roles:

- Straight arrows = synchronous replication
- Straight arrows with an asynchronous replication symbol = asynchronous replication
- Lightning bolts = FlashCopy replication

If you selected a location for the sites that are in the session, the sites are labeled with the site location. If you did not select site locations, the sites are labeled with **Site x**, where x is the site number. You can hover the mouse pointer over the site location to view the text in a larger format.

d. Click **OK**.

e. If the session was successfully created, click **Launch Add Copy Sets Wizard** and continue to the next step. If the session was not created, refer to the displayed message text to fix the problem.

2. Follow these steps to add copy sets to the session:

a. In the Add Copy Sets wizard, enter the following information for the choose volume pages. The field names that are displayed depend on the storage system type. When you complete the information on each page, click **Next**.

Storage system

Select a storage system. If the volume role has a site location that is assigned to it, you can select a storage system that is assigned to the same location as the role or a storage system that is not assigned to a location. Storage systems that have a location are listed under the location. Storage systems that do not have a location are listed under **None**. If the role does not have a site location, you can select any storage system.

Logical storage system, I/O Group, or Pool

Select a logical subsystem (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.

The use of extent space-efficient volumes in volume roles is dependent on the capabilities of the storage system.

Session image

Shows an image that represents the session in which the role for which you are selecting volumes is highlighted. This image shows how many roles 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.

b. 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. On this page, the **Storage system** and **Logical storage system, I/O Group, or Pool** lists are populated with the values from the previously selected copy set. Repeat the previous steps to complete the choose volume pages and add additional copy sets.

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.

c. On the Confirm page, the number of copy sets to be added is displayed. Click **Next**.

d. A progress bar is displayed. When the copy sets are added, review the results and click **Finish**.

3. Follow these steps to add properties for the session:

   a. On the Sessions page, select the session.

   b. From the **Session Actions** list, select **View/Modify** > **Properties**.

   c. In the View/Modify Properties notebook, select the options that you want. For a description of the options, see "Viewing session properties" on page 172.

   d. Click **OK**.

Related concepts:

"Sessions" on page 63
A *session* completes a specific type of data replication for a specific set of volumes. During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a copy set. A session can contain one or more copy sets.

"Session types" on page 12
Tivoli Storage Productivity Center for Replication provides several methods to replicate data. The type of data replication that is associated with a session is known as the *session type*.

Creating a Global Mirror session and adding copy sets

A Global Mirror session is a method of asynchronous, remote data replication between two sites that are over 300 KM apart.

To add a Global Mirror session, you create the session, add copy sets to the session, and then select the properties for the session.

1. Follow these steps to create a Global Mirror session:

   a. In the menu bar, click **Sessions**.

   b. On the Sessions page, click **Create Session**.

   c. In the Create Session window, select the following options:

   Hardware type
   Select the type of storage system for the session:
   - DS8000, DS6000, ESS 800
   - SAN Volume Controller
   - Storwize Family (for storage systems in the IBM Storwize family)
   - Storwize V7000 Unified
   - XIV

   Session type
   Select a Global Mirror session type.

   Session name
   Enter a name to identify the session. For sessions that contain an IBM XIV Storage System, the session name can be up to 58 alphanumeric characters. For sessions that contain other storage system types, the session name can be up to 250 alphanumeric characters. The default name for a new session is `newsession`.

   Site *x* Location
   For each site that is in the session, select the location for the site or **None**. The site location, if available, is set by storage system on the Storage Systems page.

   You can associate a location with each storage system and each site in a session. This site awareness ensures that you can select only

storage systems that have either a matching site location for each role or have no defined location when you add copy sets to a session.

Session image
Shows an image that represents the session. The session image is displayed when you select a session type. The image is a visual aid to help you create your session.

This image shows the number of roles in the session and how the roles are distributed between the sites. The arrows indicate the type of copy between roles:
- Straight arrows = synchronous replication
- Straight arrows with an asynchronous replication symbol = asynchronous replication
- Lightning bolts = FlashCopy replication

If you selected a location for the sites that are in the session, the sites are labeled with the site location. If you did not select site locations, the sites are labeled with **Site** *x*, where *x* is the site number. You can hover the mouse pointer over the site location to view the text in a larger format.

  d. Click **OK**.

  e. If the session was successfully created, click **Launch Add Copy Sets Wizard** and continue to the next step. If the session was not created, refer to the displayed message text to fix the problem.

2. Follow these steps to add copy sets to the session:

  a. In the Add Copy Sets wizard, enter the following information for the choose volume pages. The field names that are displayed depend on the storage system type. When you complete the information on each page, click **Next**.

Storage system
Select a storage system. If the volume role has a site location that is assigned to it, you can select a storage system that is assigned to the same location as the role or a storage system that is not assigned to a location. Storage systems that have a location are listed under the location. Storage systems that do not have a location are listed under **None**. If the role does not have a site location, you can select any storage system.

Logical storage system, I/O Group, or Pool
Select a logical subsystem (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.

The use of extent space-efficient volumes in volume roles is dependent on the capabilities of the storage system.

Session image
Shows an image that represents the session in which the role for which you are selecting volumes is highlighted. This image shows how many roles 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.

  b. 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. On this page, the **Storage system** and **Logical storage system, I/O Group, or Pool** lists are populated with the values from the previously selected copy set. Repeat the previous steps to complete the choose volume pages and add additional copy sets.
>
> **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.

  c. On the Confirm page, the number of copy sets to be added is displayed. Click **Next**.

  d. A progress bar is displayed. When the copy sets are added, review the results and click **Finish**.

3. Follow these steps to add properties for the session:

  a. On the Sessions page, select the session.

  b. From the **Session Actions** list, select **View/Modify** > **Properties**.

  c. In the View/Modify Properties notebook, select the options that you want. For a description of the options, see "Viewing session properties" on page 172.

  d. Click **OK**.

Related concepts:

"Sessions" on page 63
A *session* completes a specific type of data replication for a specific set of volumes. During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a copy set. A session can contain one or more copy sets.

"Session types" on page 12
Tivoli Storage Productivity Center for Replication provides several methods to replicate data. The type of data replication that is associated with a session is known as the *session type*.

Creating a Metro Global Mirror session and adding copy sets

A Metro Global Mirror session is a method of continuous, remote data replication that operates between three sites of 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.

To add a Metro Global Mirror session, you create the session, add copy sets to the session, and then select the properties for the session.

1. Follow these steps to create a Metro Global Mirror session:

   a. In the menu bar, click **Sessions**.

   b. On the Sessions page, click **Create Session**.

   c. In the Create Session window, select the following options:

   Hardware type
   > Select the hardware type **DS8000, DS6000, ESS 800**.

   Session type
   > Select a Metro Global Mirror session type.

   Session name
   > Enter a name to identify the session. The session name can be up to 250 alphanumeric characters. The default name for a new session is `newsession`.

   Site *x* Location
   > For each site that is in the session, select the location for the site or **None**. The site location, if available, is set by storage system on the Storage Systems page.
   >
   > You can associate a location with each storage system and each site in a session. This site awareness ensures that you can select only storage systems that have either a matching site location for each role or have no defined location when you add copy sets to a session.

   Session image
   > Shows an image that represents the session. The session image is displayed when you select a session type. The image is a visual aid to help you create your session.
   >
   > This image shows the number of roles in the session and how the roles are distributed between the sites. The arrows indicate the type of copy between roles:
   > - Straight arrows = synchronous replication
   > - Straight arrows with an asynchronous replication symbol = asynchronous replication

- Lightning bolts = FlashCopy replication

    If you selected a location for the sites that are in the session, the
    sites are labeled with the site location. If you did not select site
    locations, the sites are labeled with **Site** *x*, where *x* is the site
    number. You can hover the mouse pointer over the site location to
    view the text in a larger format.

   d. Click **OK**.

   e. If the session was successfully created, click **Launch Add Copy Sets Wizard**
      and continue to the next step. If the session was not created, refer to the
      displayed message text to fix the problem.

2. Follow these steps to add copy sets to the session:

   a. In the Add Copy Sets wizard, enter the following information for the
      choose volume pages. The field names that are displayed depend on the
      storage system type. When you complete the information on each page,
      click **Next**.

      Storage system
      Select a storage system. If the volume role has a site location that is
      assigned to it, you can select a storage system that is assigned to the
      same location as the role or a storage system that is not assigned to
      a location. Storage systems that have a location are listed under the
      location. Storage systems that do not have a location are listed
      under **None**. If the role does not have a site location, you can select
      any storage system.

      Logical storage system
      Select a logical subsystem (LSS).

      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.

      The use of extent space-efficient volumes in volume roles is
      dependent on the capabilities of the storage system.

      Session image
      Shows an image that represents the session in which the role for
      which you are selecting volumes is highlighted. This image shows
      how many roles 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 **Browse** to select the CSV file.

   b. 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. On this page, the **Storage system** and **Logical storage system, I/O Group, or Pool** lists are populated with the values from the previously selected copy set. Repeat the previous steps to complete the choose volume pages and add additional copy sets.

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.

c. On the Confirm page, the number of copy sets to be added is displayed. Click **Next**.

d. A progress bar is displayed. When the copy sets are added, review the results and click **Finish**.

3. Follow these steps to add properties for the session:

a. On the Sessions page, select the session.

b. From the **Session Actions** list, select **View/Modify** > **Properties**.

c. In the View/Modify Properties notebook, select the options that you want. For a description of the options, see "Viewing session properties" on page 172.

d. Click **OK**.

Related concepts:

"Sessions" on page 63
A *session* completes a specific type of data replication for a specific set of volumes. During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a copy set. A session can contain one or more copy sets.

"Session types" on page 12
Tivoli Storage Productivity Center for Replication provides several methods to replicate data. The type of data replication that is associated with a session is known as the *session type*.

Completing session administration tasks

Use session commands to complete session tasks such as starting, stopping, suspending, or terminating sessions.

You can issue session commands by using the GUI, as described in this topic, or you can use the CLI to issue commands. For a description of the CLI session

commands, see the IBM Tivoli Storage Productivity Center for Replication for System z Command-line Interface User's Guide.

To issue a command for a session, complete the following steps:

1. In the menu bar, click **Sessions**.
2. On the Sessions page, select the session for which you want to issue a command.
3. Click **Session Actions** > **Commands**.

   If the session is a single-target session type, all available commands are displayed under **Commands**. Click the command that you want to issue.

   If the session is a multi-target session type, some of the commands are grouped by overall function. For example, **Start**, **Stop**, or **Suspend**. When you click a command, a window is displayed that provides specific commands for that function. For example, for Metro Mirror - Metro Mirror sessions, if you select **Start**, a window is displayed that contains the specific start commands that you can select. Select the command that you want, review the message text for that command, and then click **Yes** if you want the command to run.

Related reference:

"Session commands" on page 106
The commands that are available for a session depend on the session type, and for some commands, the storage system type.

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 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 remains 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 remains 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 remains in 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 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.

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 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 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 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 though a TCP/IP (direct connect or HMC) connection. It is not supported on storage systems connected though 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, 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, 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 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 menu bar, click **Settings** > **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.

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

1. In the menu bar, click **Sessions**.
2. On the Sessions page, select the session for which you want to export copy sets.
3. From the **Session Actions** list, select **Export** > **Export Copy Sets**. The Export Copy Sets window displays the status of the export and a link to the exported file if the export completed.
4. Right-click the link and save the file to a local system.

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 menu bar, click **Sessions**.
2. On the Sessions page, select the session for which you want to import copy sets.
3. From the **Session Actions** list, select **View/Modify** > **Add Copy Sets**. The Add Copy Sets wizard is displayed.
4. Select **Use a CSV file to import copy sets**.
5. Click **Browse**, select the CSV file, and then click **Next**.
6. Verify that the matching results were successful, and then click **Next**.
7. Select the copy sets that you want to add, and then click **Next**.
8. Confirm the number of copy sets that you want to create, and then click **Next**.
9. View the information on the Results page of the wizard, and then click **Finish**.

Exporting historical data for Global Mirror role pairs

You can export data for a Global Mirror role pair that is in a TotalStorage Enterprise Storage Server, System Storage DS6000, or System Storage DS8000 session to a comma-separated value (CSV) file. You can then use the data in the CSV file to analyze trends in your storage environment that affect your recovery point objective (RPO).

Complete these steps to export data to a CSV file:

1. In the menu bar, click **Sessions**.
2. On the Sessions page, click the Global Mirror session for which you want to export data.
3. From the **Session Actions** list, select **Export** > **Export Global Mirror Data**.
4. In the Export Historical Data for Global Mirror window, complete the following information, and then click **OK**:

Select the role pair
> Select the role pair for the data that you want to show in the CSV file.

Select the type of data
> Select the type of data that you want to export to a CSV file. Depending on the data type that you select, the CSV file contains data about the RPO or data about logical subsystem (LSS) out-of-sync tracks.
>
> To better analyze trends, you can create a CSV file with one data type and then create a file with the other data type. For example, the file that contains data for the RPO might show that the RPO threshold is often exceeded on a particular day and time. You can then view the file that contains data for LSS out-of-sync tracks to see whether a particular LSS or set of LSSs have high out-of-sync track values for that day and time.
>
> **RPO** Select this option to export data for the RPO. The data that is shown in the export file includes the average RPO for the dates that you select and information related to the formation of consistency groups.
>
> **LSS Out-of-Sync Tracks**
> > Select this option to export data for the out-of-sync tracks that are in the LSSs.

Start date
> Select the start date for the data that you want to show in the export file.
>
> By default, the date range maximum for RPO files is 31 days of data and the maximum for LSS out-of-sync track files is 7 days of data.

End date
> Select the end date for the data that you want to show in the export file.

After you click **OK** to export the file, a results window displays the status of the export and a link to the exported file if the export completed. Download and save the file to a local system.

Modifying the location of session sites

You can change the location that is associated with each site in session.

Prerequisites: You must have Administrator privileges or Operator privileges for the session to modify the location of a site in a session.

Changing the location of a site in a session does not affect the location of the storage systems that are associated with that site.

Complete the following steps to modify the location of a site:
1. In the menu bar, click **Sessions**.
2. On the Sessions page, select the session whose site locations you want to change.
3. From the **Session Actions** list, select **View/Modify** > **Site Location(s)**.
4. In the Modify Site Locations window, select a site location or **None** for each site that you want to change. An image that represents the sites in the session is displayed. If you selected locations for the sites that are in the session, the

sites are labeled with the site location. If you did not select site locations, the sites are labeled with **Site *n***, where *n* is the site number. You can hover the mouse pointer over the site name to view the name in a larger format.

5. Click **OK**.

Related concepts:

"Site awareness" on page 45
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.

Removing sessions

This topic describes how to remove sessions.

Important: You can remove only sessions that are in the Defined state.

Complete these steps to remove a session:

1. In the menu bar, click **Sessions**.
2. On the Sessions page, select the session that you want to remove.
3. From the **Session Actions** list, select **Remove Session**.
4. Click **Yes** to confirm that you want to remove the session.

Removing copy sets

This topic describes how to remove copy sets.

Complete these steps to remove a copy set:

1. In the menu bar, click **Sessions**.
2. On the Sessions page, select the session hat you want to remove copy sets from.
3. From the **Session Actions** list, select **View/Modify** > **Remove Copy Sets**. This starts the Remove Copy Sets wizard.
4. 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**.
5. Select the copy sets that you want to remove and click **Next**.
6. 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) remain 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.

7. After the copy sets are removed, click **Finish**.

Related concepts:

"Copy sets" on page 63
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*.

Migrating existing hardware relationships to Tivoli Storage Productivity Center for Replication

You can convert existing hardware relationships for copy services to a Tivoli Storage Productivity Center for Replication session.

For session types other than multi-target sessions, you can either complete this action manually as described in this topic or use the data migration utility for Tivoli Storage Productivity Center for Replication. The data migration utility is not available for multi-target sessions.

To download the data migration utility, go to ➡ http://www.ibm.com/support/docview.wss?uid=swg24000625. The data migration utility produces the Tivoli Storage Productivity Center for Replication CLI command script files and DSCLI script files that are necessary to migrate the configuration to Tivoli Storage Productivity Center for Replication.

To manually migrate role pairs in a hardware relationship, complete the following steps:

1. Identify the hardware relationships that you want to migrate.

2. Create a Tivoli Storage Productivity Center for Replication session and add copy sets to the session.

3. If you are migrating from one of the following storage systems, complete the applicable step:

    a. For DS8000, DS6000, and TotalStorage Enterprise Storage Server storage systems with Global Mirror relationships, terminate the Global Mirror master.

    b. For SAN Volume Controller, Storwize, or the XIV storage systems, remove all volumes from the consistency group.

4. Issue a start command to start the Tivoli Storage Productivity Center for Replication session.

Related reference:

"Session commands" on page 106
The commands that are available for a session depend on the session type, and for some commands, the storage system type.

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.

Important: 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:
- System Storage DS8000
- System Storage DS6000
- TotalStorage Enterprise Storage Server Model 800
- SAN Volume Controller,
- Storwize V3700
- Storwize V7000
- Storwize V7000 Unified

Perform these steps to practice disaster recover 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, and 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:

- System Storage DS8000
- System Storage DS6000
- TotalStorage Enterprise Storage Server Model 800
- SAN Volume Controller,
- Storwize V3700
- Storwize V7000
- Storwize V7000 Unified

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 Overview page to view overall health and status of sessions, storage systems, host systems, and management servers.

The Overview page is the first page that you see after you log on. You can display this page by selecting **Overview** in the Tivoli Storage Productivity Center for Replication menu bar. 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 exit.

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
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.

Viewing SNMP alerts

You can view Tivoli Storage Productivity Center for Replication for System z SNMP trap descriptions from the Alerts page.

To view the alerts for the devices that Tivoli Storage Productivity Center for Replication for System z is monitoring , go to **Home** > **Alerts**.

Viewing sessions

The Tivoli Storage Productivity Center for Replication GUI includes icons, images, messages, and other visual and informational aids to help you to determine the status and state of your sessions.

Customizing the session list

The Session page provides a list of created sessions. This list is organized in columns and rows. There is a row for each session and columns for information that is related to that session, such as the session name, status, state, and type. You can customize the session list by filtering information, resizing columns, and sorting rows.

Filtering the session list

Use filtering when you want to limit the rows that are shown in the session list based on the session name or type. To apply a filter, enter the filter text in the **Filter** field on the Session page. You can enter complete text or partial text. Only those rows that contain matching text in the **Name** or **Type** column are displayed in the list.

To remove the filtering for the list, click the **X** icon in the **Filter** field.

Resizing the columns in the session list

You can resize columns to more easily view the information in the list. To resize a column, drag the column heading border until the column is at the width that you want.

Sorting the session list

By default, the sessions in the list are sorted by session name in ascending order. You can change the value that is used for sorting by clicking the appropriate column heading. For example, if you want to sort by the session type, click the **Type** heading. The list sorts by the value in that column. You can choose an ascending or descending sort order by clicking the column heading.

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 47. Session status icons

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

Session images

The 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 48. 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).
	Active host volumes with change volumes	This symbol represents volumes that contain the source of updated tracks to which the application is actively issuing read and write I/O and change volumes.
	Recoverable volumes	This symbol represents volumes that contain a consistent copy of the data.
	Recoverable volumes with change volumes	This symbol represents volumes and change volumes that contain a consistent copy of the data.
	Inconsistent volumes	This symbol represents the volumes that do not contain a consistent copy of the data.
	Inconsistent volumes with change volumes	This symbol represents the volumes and change volumes that do not contain a consistent copy of the data.

Data copying symbols

The data copying symbols indicate the type of copy that occurs between the volume roles. The direction that the symbol is displayed in the Tivoli Storage Productivity Center for Replication GUI depends on the direction of the copy.

Table 49. Data copying symbols

Symbol	Type of Copy	Meaning
	FlashCopy	This symbol represents a FlashCopy operation.
	FlashCopy with errors	This symbol represents a FlashCopy operation with errors on one or more pair.
	FlashCopy inactive	This symbol represents an inactive FlashCopy operation.
	FlashCopy inactive with errors	This symbol represents an inactive FlashCopy operation with errors on one or more pair.
	Synchronous	This symbol represents a synchronous copy.
	Synchronous with errors	This symbol represents a synchronous copy with errors on one or more pairs.
	Synchronous inactive	This symbol represents an inactive synchronous copy.
	Synchronous inactive with errors	This symbol represents an inactive synchronous copy with errors on one or more pair.
	Asynchronous	This symbol represents an asynchronous copy.
	Asynchronous with errors	This symbol represents an asynchronous copy with errors on one or more pair.
	Asynchronous inactive	This symbol represents an inactive asynchronous copy.
	Asynchronous inactive with errors	This symbol represents an inactive asynchronous copy with errors on one or more pair.

Table 49. Data copying symbols (continued)

Symbol	Type of Copy	Meaning
	HyperSwap or Open HyperSwap	This symbol indicates that the HyperSwap or Open HyperSwap feature is enabled for the relationship. If a failure occurs when I/O is being written to the primary storage system, these features automatically swap the I/O to the secondary site with no user interaction and little or no application impact.
	Suspended	This symbol represents a suspended copy relationship.
	Failed over	This symbol represents a failover copy relationship.

Session states

You can view the health and status of a session in the Tivoli Storage Productivity Center for Replication GUI.

Attention: Use only the Tivoli Storage Productivity Center for Replication graphical user interface (GUI) or command-line interface CLI to manage session relationships, such as volume pairs and copy sets. Do not modify session relationships through other interfaces such as the System Storage DS CLI. If you modify relationships through other interfaces, a loss of consistency can occur across the relationships that are managed by the session. The exceptions to this requirement are failover operations that are managed by external applications for certain session and storage system types as described in "Failover and failback operations" on page 69.

The `Refresh States` command is used to refresh the states of the role pairs that are in a session (if applicable for the session type) and refresh the state of the session if it is incorrect. Issue this command to query the states of the copy sets on the hardware. You are not required 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 Tivoli Storage Productivity Center for Replication to update itself. This command triggers multiple queries on the hardware, which can impact hardware performance. Do not run this command more frequently than every few minutes in each session.

The following table describes each session state.

Table 50. Session states

State	Session type	Description
Defined	All	The session exists but is inactive.

Table 50. Session states (continued)

State	Session type	Description
Flashing	All	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. In a Metro Global Mirror session, data copying is temporarily suspended while a consistent practice copy of data is being prepared on site 3.
Prepared	All	The source to target data transfer is active. In a Metro Mirror, Global Mirror, or Metro Global Mirror session, the data that is written to the source is transferred to the target, and all volumes are consistent and recoverable. In a FlashCopy session, the volumes are not yet consistent, but the flash is ready to begin. **Important:** 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: • SAN Volume Controller • Storwize V3500 (use only FlashCopy sessions for this storage system) • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified For example, if a Metro Mirror session with one copy set is in the Prepared state, and you stop the role pair, the session is still displayed in the Prepared state.
Preparing	All	The volumes are initializing, synchronizing, or resynchronizing. 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. In a FlashCopy session, the volumes are initializing. The preparing state for FlashCopy sessions applies only to SAN Volume Controller or Storwize storage systems.

Table 50. Session states (continued)

State	Session type	Description
Recovering	All	The session is in the process of recovering.
Suspended	All	Data copying has temporarily stopped. **Important:** The suspended state applies only to Global Mirror, Metro Mirror, and Metro Global Mirror sessions.
Suspended (Partial)	Multi-target sessions	Data copying between the source site and one of the target sites has suspended.
SuspendedH1H3 MGM	Metro Global Mirror	Data copying between site 1 and site 3 is suspended.
SuspendedH1H2 MGM	Metro Global Mirror	Data copying between site 1 and site 2 is suspended.
Suspending	All	The session is changing into a Suspended state. **Important:** The Suspending state applies only to Metro Mirror Sessions in Global Copy mode, Global Mirror, and Metro Global Mirror sessions and does not apply to SAN Volume Controller or Storwize storage systems.
Target available	All	Target volumes are available for application updates.
Terminating	FlashCopy	The session is being terminated because you issued a Terminate action under the following conditions: • You permitted the target to be Metro Mirror or Global Copy. • You set the Require or Attempt to Preserve Mirror option. The session displays as Terminating until the FlashCopy background copy is complete and no longer exists on the hardware.

Role pair status and progress

In a session, a *role pair* is the association of two volume roles that take part in a copy relationship. Tivoli Storage Productivity Center for Replication provides detailed role pair status and progress messages for sessions.

The role pair status and progress messages are updated to provide a message that indicates what the session is doing at the time. By hovering over a progress bar, you can see specific information about the 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 are displayed in the Session Details and Role Pair Details page.

Table 51. Detailed status messages for Participating and Non-Participating role pairs

Supported session type	Status message
FlashCopy Metro Mirror Global Copy Global Mirror	Starting *role_pair_name* relationships on the hardware
Metro Mirror Global Mirror	Waiting for all pairs in the role pair *role_pair_name* to reach state of *state*
FlashCopy Metro Mirror Global Copy Global Mirror	Terminating all pairs in role pair *role_pair_name*
FlashCopy Metro Mirror Global Mirror	Recovering all pairs in role pair *role_pair_name*
Metro Mirror Global Copy Global Mirror	Suspending all pairs in role pair *role_pair_name*
FlashCopy	Background copy is running for role pair *role_pair_name*
Global Copy	Waiting for all pairs in role pair *role_pair_name* to become consistent
Global Copy	Waiting for all pairs in role pairs *role_pair_name* to complete the initial copy
FlashCopy	Waiting for all pairs in role pairs *role_pair_name* to complete FRR
Global Mirror	Waiting for all pairs in role pairs *role_pair_name* to join the Global Mirror session

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 menu bar, click **Sessions**.
2. On the Sessions page, select the session that you want to view.
3. From the **Session Actions** list, select **View/Modify** > **View Details**.

Viewing additional details for Global Mirror sessions

Additional detail information is available for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server Global Mirror sessions that do not have an Inactive status. This information includes details about the Global Mirror master, consistency groups that have been formed, and data exposure.

Complete these steps to view additional details for Global Mirror sessions:

1. In the menu bar, click **Sessions**.
2. On the Sessions page, select the Global Mirror session that you want to view.
3. From the **Session Actions** list, select **View/Modify** > **View Details**.
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 tracked in the **Data Exposure** chart.

Show Data Exposure over

Data exposure that is over the value that is entered in this field is shown in the **Data Exposure** chart.

Show Data Exposure under

Data exposure that is under the value that is entered in this field is shown in the **Data Exposure** chart.

Viewing session properties

You can view and modify the description and options for a session. The options that are presented depend on the storage system type and the session type that you select.

To view or modify the properties for a session, complete the following steps:

1. In the menu bar, click **Sessions**.
2. On the Sessions page, select the session with the properties that you want to view.
3. From the **Session Actions** list, click **View/Modify** > **Properties**. The properties that are displayed depend on the session type as described in the following sections.

TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 sessions

The properties for IBM TotalStorage Enterprise Storage Server Model 800, IBM System Storage DS8000, and IBM System Storage DS6000 sessions depend on the session type.

The following sections describe the properties for each session type.

FlashCopy session properties

Use the View/Modify Properties notebook to view or modify properties for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 FlashCopy sessions.

The properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session. Use the **H1-T1 Options** tab to set properties that apply to the role pair.

Session Options

Enter the properties for the session.

Description

Type the description for this session.

H1-T1 Options

Enter the properties for the role pair.

Incremental

Select this option to apply incremental changes to the target

volume. After the initial FlashCopy operation, only data that changed on the source volume since the last FlashCopy operation was performed is copied to the target volume.

If you select this option, a persistent FlashCopy relationship is created regardless of whether you select the **Persistent** check box.

Persistent
Select this option to keep the FlashCopy relationship established on the hardware after all source tracks are copied to the target volume. If you do not select this option, the local replication relationship ends after the target volume contains a complete point-in-time image of the source volume.

No Copy
Select this option if you do not want the hardware to write the background copy until the source track is written to. Data is not copied to the target volume until the blocks or tracks of the source volume are modified.

This option is required for space-efficient volumes.

Allow FlashCopy target to be Metro Mirror source
Select this option to enable the FlashCopy operation if the target volume of the FlashCopy relationship is also the source volume of a Metro Mirror relationship. If this option is not selected, the FlashCopy operation fails.

Requirement: This option requires that the IBM Remote Pair FlashCopy option is available for your IBM System Storage DS8000 storage system. To determine whether you can use the IBM Remote Pair FlashCopy option with your System Storage DS8000 storage system, refer to the System Storage DS8000 documentation for microcode level that you are using.

Select one of the following options to specify whether you want to maintain consistency, if possible:

Do not attempt to preserve Metro Mirror consistency
Click this option if you want the FlashCopy operation to complete without preserving consistency of the Metro Mirror relationship on the remote site. The FlashCopy operation does not occur on the remote site.

Attempt to preserve Metro Mirror consistency, but allow FlashCopy even if Metro Mirror target consistency cannot be preserved
Click this option 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 the 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 FlashCopy operations are performed on both sites if possible.

Attempt to preserve Metro Mirror consistency, but fail FlashCopy if Metro Mirror target consistency cannot be preserved
Click this option to prevent a full copy from being performed over the Metro Mirror link. Instead, if possible,

parallel FlashCopy operations are performed 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.

Metro Mirror Single Direction session properties

Use the View/Modify Properties notebook to view or modify properties for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 Metro Mirror Single Direction sessions.

The properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session. Use the **H1-H2 Options** tab to set properties that apply to the role pair.

Session Options
> Enter the properties for the session.

> **Description**
>> Type the description for this session.

> **Reset Secondary Reserves**
>> Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets when a Start command is issued for the session.

>> **Attention:** This option causes the session to overwrite all data that is on the target volume.

> **Fail MM/GC if target is online (CKD only)**
>> Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online state. For more information about this state, refer to the documentation for the storage system.

>> This option applies only to count key data (CKD) volumes.

H1-H2 Options
> Enter the properties for the role pair.

> **Metro Mirror Suspend Policy**
>> Select one of the following options to specify the policy for holding or releasing input/output (I/O) after a Metro Mirror relationship is suspended. When a relationship is in a Suspended state, write operations from the source volume are no longer mirrored to the target volume.

>> **Hold I/O after Suspend**
>>> Click this option to prevent the source volume from receiving new data after the relationship is suspended. Use this option if you want to determine the scope of the suspension before new data is written to the source volume. This option helps to ensure that the data on the source and target volumes remains the same.

>>> To enable write operations to the source volume to continue, issue the Tivoli Storage Productivity Center for Replication `Release I/O` command for the session. If you do not issue this command, write operations are enabled when the hardware timeout value on the storage system expires.

Release I/O after Suspend

Click this option to enable the source volume to automatically receive data after the relationship is suspended. Use this option if you want to limit the effect on the applications that are writing to the source volume. This option is enabled by default.

Attention: This option can cause the data on the source volume to be different from the data on target volume.

z/OS Management

Select the following options to manage IBM z/OS features for the role pair.

The **z/OS Management** options are displayed only if Tivoli Storage Productivity Center for Replication is installed on or connected to a z/OS host system. You can add a z/OS host connections in Tivoli Storage Productivity Center for Replication regardless of the operating system on which the application is installed.

When you are connected to a z/OS host system, you can manage z/OS features for volumes that are attached to the host system. In addition, you can manage these features for volumes that are attached to other z/OS systems that are connected to the host system through a sysplex.

System or sysplex

Select the z/OS system or sysplex that you want to associate with the session. Commands for z/OS features are issued to this system or sysplex. Volumes in the session must be attached to the system or sysplex to enable the features.

Enable Hardened Freeze

Select this option to enable the z/OS Input/Output Supervisor (IOS) to manage freeze operations.

Requirement: This option requires the z/OS address spaces Basic HyperSwap Management and Basic HyperSwap API. For instructions about how to start these address spaces, see the information about preparing to use Basic HyperSwap from z/OS in the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

If you select this option, IOS can freeze volumes regardless of whether the Tivoli Storage Productivity Center for Replication server is started or stopped.

In addition, this option enables you to include z/OS system volumes such as paging, database, and IBM WebSphere Application Server hierarchical file system (HFS) as Metro Mirror volumes in the session. When you select this option, IOS manages the freeze operations for all Metro Mirror volumes in the session, which prevents Tivoli Storage Productivity Center for Replication from freezing the volumes and possibly freezing itself. This option does not enable IOS to manage freeze operations for Global Mirror volumes.

Metro Mirror Failover/Failback session properties

Use the View/Modify Properties notebook to view or modify properties for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 Metro Mirror Failover/Failback sessions.

The properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session. Use the **H1-H2 Options** tab to set properties that apply to the role pair.

Session Options

Enter the properties for the session.

Description

Type the description for this session.

Reset Secondary Reserves

Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets when a Start command is issued for the session.

Attention: This option causes the session to overwrite all data that is on the target volume.

Fail MM/GC if target is online (CKD only)

Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online state. For more information about this state, refer to the documentation for the storage system.

This option applies only to count key data (CKD) volumes.

H1-H2 Options

Enter the properties for the role pair.

Metro Mirror Suspend Policy

Select one of the following options to specify the policy for holding or releasing input/output (I/O) after a Metro Mirror relationship is suspended. When a relationship is in a Suspended state, write operations from the source volume are no longer mirrored to the target volume.

Hold I/O after Suspend

Click this option to prevent the source volume from receiving new data after the relationship is suspended. Use this option if you want to determine the scope of the suspension before new data is written to the source volume. This option helps to ensure that the data on the source and target volumes remains the same.

To enable write operations to the source volume to continue, issue the Tivoli Storage Productivity Center for Replication `Release I/O` command for the session. If you do not issue this command, write operations are enabled when the hardware timeout value on the storage system expires.

Release I/O after Suspend

Click this option to enable the source volume to automatically receive data after the relationship is suspended. Use this option if you want to limit the effect

on the applications that are writing to the source volume. This option is enabled by default.

Attention: This option can cause the data on the source volume to be different from the data on target volume.

z/OS Management

Select the following options to manage IBM z/OS features for the role pair.

The **z/OS Management** options are displayed only if Tivoli Storage Productivity Center for Replication is installed on or connected to a z/OS host system. You can add a z/OS host connections in Tivoli Storage Productivity Center for Replication regardless of the operating system on which the application is installed.

When you are connected to a z/OS host system, you can manage z/OS features for volumes that are attached to the host system. In addition, you can manage these features for volumes that are attached to other z/OS systems that are connected to the host system through a sysplex.

System or sysplex

Select the z/OS system or sysplex that you want to associate with the session. Commands for z/OS features are issued to this system or sysplex. Volumes in the session must be attached to the system or sysplex to enable the features.

Enable Hardened Freeze

Select this option to enable the z/OS Input/Output Supervisor (IOS) to manage freeze operations.

Requirement: This option requires the z/OS address spaces Basic HyperSwap Management and Basic HyperSwap API. For instructions about how to start these address spaces, see the information about preparing to use Basic HyperSwap from z/OS in the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

If you select this option, IOS can freeze volumes regardless of whether the Tivoli Storage Productivity Center for Replication server is started or stopped.

In addition, this option enables you to include z/OS system volumes such as paging, database, and IBM WebSphere Application Server hierarchical file system (HFS) as Metro Mirror volumes in the session. When you select this option, IOS manages the freeze operations for all Metro Mirror volumes in the session, which prevents Tivoli Storage Productivity Center for Replication from freezing the volumes and possibly freezing itself. This option does not enable IOS to manage freeze operations for Global Mirror volumes.

If you select the **Manage H1-H2 with HyperSwap** option for the session, this option is ignored. HyperSwap includes IOS for managing freeze operations. The **Enable Hardened**

Freeze option ensures data integrity if Tivoli Storage Productivity Center for Replication freezes and HyperSwap are not enabled for a session.

Manage H1-H2 with HyperSwap

Select this option to trigger a HyperSwap operation, which redirects application I/O to the target volumes when there is a failure on the host accessible volumes. Tivoli Storage Productivity Center for Replication uses HyperSwap to manage the H1-H2 sequence of a Metro Mirror or Metro Global Mirror session.

Disable HyperSwap

Select this option to prevent a HyperSwap operation from occurring.

On Configuration Error:

Partition the system(s) out of the sysplex

Select this option to partition a new system out of the sysplex when an error occurs because the system cannot be added to the HyperSwap configuration.

Disable HyperSwap

Select this option to prevent a HyperSwap operation from occurring.

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 enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

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 the HyperSwap operation does not complete.

Requirement: If you select this option, you must restart the system.

Disable HyperSwap after attempting backout

Select this option to enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

Metro Mirror Failover/Failback with Practice session properties

Use the View/Modify Properties notebook to view or modify properties for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 Metro Mirror Failover/Failback with Practice sessions.

The properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session or multiple role pairs in the session. Use the **H2-I2 Options** tab to set properties that apply to that role pair.

Session Options
> Enter the properties for the session.

> **Description**
>> Type the description for this session.

> **Reset Secondary Reserves**
>> Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets when a Start command is issued for the session.

>> **Attention:** This option causes the session to overwrite all data that is on the target volume.

> **Fail MM/GC if target is online (CKD only)**
>> Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online state. For more information about this state, refer to the documentation for the storage system.

>> This option applies only to count key data (CKD) volumes.

> **Metro Mirror Suspend Policy**
>> Select one of the following options to specify the policy for holding or releasing input/output (I/O) after a Metro Mirror relationship is suspended. When a relationship is in a Suspended state, write operations from the source volume are no longer mirrored to the target volume.

>> **Hold I/O after Suspend**
>>> Click this option to prevent the source volume from receiving new data after the relationship is suspended. Use this option if you want to determine the scope of the suspension before new data is written to the source volume. This option helps to ensure that the data on the source and target volumes remains the same.

>>> To enable write operations to the source volume to continue, issue the Tivoli Storage Productivity Center for Replication `Release I/O` command for the session. If you do not issue this command, write operations are enabled when the hardware timeout value on the storage system expires.

>> **Release I/O after Suspend**
>>> Click this option to enable the source volume to automatically receive data after the relationship is suspended. Use this option if you want to limit the effect on the applications that are writing to the source volume. This option is enabled by default.

Attention: This option can cause the data on the source volume to be different from the data on target volume.

z/OS Management

Select the following options to manage IBM z/OS features for the H1-I2 and H2-H1 role pairs.

The **z/OS Management** options are displayed only if Tivoli Storage Productivity Center for Replication is installed on or connected to a z/OS host system. You can add a z/OS host connection in Tivoli Storage Productivity Center for Replication regardless of the operating system on which the application is installed.

When you are connected to a z/OS host system, you can manage z/OS features for volumes that are attached to the host system. In addition, you can manage these features for volumes that are attached to other z/OS systems that are connected to the host system through a sysplex.

System or sysplex

Select the z/OS system or sysplex that you want to associate with the session. Commands for z/OS features are issued to this system or sysplex. Volumes in the session must be attached to the system or sysplex to enable the features.

Enable Hardened Freeze

Select this option to enable the z/OS Input/Output Supervisor (IOS) to manage freeze operations.

Requirement: This option requires the z/OS address spaces Basic HyperSwap Management and Basic HyperSwap API. For instructions about how to start these address spaces, see the information about preparing to use Basic HyperSwap from z/OS in the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

If you select this option, IOS can freeze volumes regardless of whether the Tivoli Storage Productivity Center for Replication server is started or stopped.

In addition, this option enables you to include z/OS system volumes such as paging, database, and IBM WebSphere Application Server hierarchical file system (HFS) as Metro Mirror volumes in the session. When you select this option, IOS manages the freeze operations for all Metro Mirror volumes in the session, which prevents Tivoli Storage Productivity Center for Replication from freezing the volumes and possibly freezing itself. This option does not enable IOS to manage freeze operations for Global Mirror volumes.

H2-I2 Options

Enter the properties for the role pair.

Persistent

Select this option to keep FlashCopy pairs persistent on the hardware.

This option is available only for IBM System Storage DS8000 version 4.2 or later.

Global Mirror Single Direction and Failover/Failback session properties

Use the View/Modify Properties notebook to view or modify properties for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 Global Mirror Single Direction and Global Mirror Failover/Failback sessions.

The properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session. Use the **H1-J2 Options** tab to set properties that apply to the role pair.

Session Options

Enter the properties for the session.

Description

Type the description for this session.

Reset Secondary Reserves

Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets when a Start command is issued for the session.

Attention: This option causes the session to overwrite all data that is on the target volume.

Fail MM/GC if target is online (CKD only)

Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online state. For more information about this state, refer to the documentation for the storage system.

This option applies only to count key data (CKD) volumes.

H1-J2 Options

Enter the properties for the role pair.

Consistency group interval time (seconds)

Type how often, in seconds, the Global Mirror session attempts to form a consistency group. A lower value possibly reduces the data exposure of the session. However, a lower value also causes the session to attempt to create consistency groups more frequently, which can increase network traffic.

Recovery Point Objective Alerts

Specify the length of time that you want to set for the recovery point objective (RPO) thresholds. The values determine whether a Warning or Severe alert is generated when the RPO threshold is exceeded for a role pair. The RPO represents the length of time in seconds of data exposure that is acceptable if a disaster occurs.

Use the following options to set the RPO threshold values. For both options, you can specify an RPO threshold in the range of 0 - 65535 seconds. The default is 0 seconds, which specifies that no alerts are generated.

> ### Warning level threshold (seconds)
> Type the number of seconds that you want to set for the warning level RPO threshold. If the RPO is greater than this value, an alert is generated.
>
> If the value in this field is other than 0, it must be greater than the value in the **Consistency group interval time (seconds)** field and less than the value in the **Severe level threshold (seconds)** field.
>
> ### Severe level threshold (seconds)
> Type the number of seconds that you want to set for the severe level RPO threshold. If the RPO is greater than this value, an alert is generated and the session status changes to Severe.
>
> If the value in this field is other than 0, it must be greater than the value in the **Warning level threshold (seconds)** field.

Global Mirror Failover/Failback with Practice session properties

Use the View/Modify Properties notebook to view or modify properties for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 Global Mirror Failover/Failback with Practice sessions.

The properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session. Use the individual role pair option tabs to set properties that apply to specific role pairs. For example, options for the H1-J2 role pair are contained in the **H1-J2 Options** tab.

Session Options
> Enter the properties for the session.
>
> ### Description
> Type the description for this session.
>
> ### Fail MM/GC if target is online (CKD only)
> Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online state. For more information about this state, refer to the documentation for the storage system.
>
> This option applies only to count key data (CKD) volumes.
>
> ### Reset Secondary Reserves
> Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets when a Start command is issued for the session.
>
> **Attention:** This option causes the session to overwrite all data that is on the target volume.

H1-J2 Options
> Enter the properties for the role pair.
>
> ### Consistency group interval time (seconds)
> Type how often, in seconds, the Global Mirror session attempts to form a consistency group. A lower value possibly reduces the data exposure of the session. However, a lower value also causes the session to attempt to create consistency groups more frequently, which can increase network traffic.

Recovery Point Objective Alerts

Specify the length of time that you want to set for the recovery point objective (RPO) thresholds. The values determine whether a Warning or Severe alert is generated when the RPO threshold is exceeded for a role pair. The RPO represents the length of time in seconds of data exposure that is acceptable if a disaster occurs.

Use the following options to set the RPO threshold values. For both options, you can specify an RPO threshold in the range of 0 - 65535 seconds. The default is 0 seconds, which specifies that no alerts are generated.

Warning level threshold (seconds)

Type the number of seconds that you want to set for the warning level RPO threshold. If the RPO is greater than this value, an alert is generated.

If the value in this field is other than 0, it must be greater than the value in the **Consistency group interval time (seconds)** field and less than the value in the **Severe level threshold (seconds)** field.

Severe level threshold (seconds)

Type the number of seconds that you want to set for the severe level RPO threshold. If the RPO is greater than this value, an alert is generated and the session status changes to Severe.

If the value in this field is other than 0, it must be greater than the value in the **Warning level threshold (seconds)** field.

H2-I2 Options

Enter the properties for the role pair. The following options are available only for IBM System Storage DS8000 version 4.2 or later.

Persistent

Select this option to keep FlashCopy pairs persistent on the hardware.

No Copy

Select this option if you do not want the hardware to write the background copy until the source track is written to. Data is not copied to the H2 volume until the blocks or tracks of the I2 volume are modified.

I2-J2 Options

Enter the properties for the role pair.

Reflash After Recover

Select this option if you want to create a FlashCopy replication between the I2 and J2 volumes after the recovery of a Global Mirror Failover/Failback with Practice session. If you do not select this option, a FlashCopy replication is created only between the I2 and H2 volumes.

This option is available only for IBM System Storage DS8000 version 4.2 or later.

Global Mirror Either Direction with Two-Site Practice session properties

Use the View/Modify Properties notebook to view or modify properties for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 Global Mirror Either Direction with Two-Site Practice sessions.

The properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session. Use the individual role pair option tabs to set properties that apply to specific role pairs. For example, properties for the H1-J2 role pair are contained in the **H1-J2 Options** tab.

Session Options

> Enter the properties for the session.

> **Description**
>> Type the description for this session.

> **Reset Secondary Reserves**
>> Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets when a Start command is issued for the session.

>> **Attention:** This option causes the session to overwrite all data that is on the target volume.

> **Fail MM/GC if target is online (CKD only)**
>> Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online state. For more information about this state, refer to the documentation for the storage system.

>> This option applies only to count key data (CKD) volumes.

H1-J2 Options and H2-J1 Options

> Enter the properties for the role pair.

> **Consistency group interval time (seconds)**
>> Type how often, in seconds, the Global Mirror session attempts to form a consistency group. A lower value possibly reduces the data exposure of the session. However, a lower value also causes the session to attempt to create consistency groups more frequently, which can increase network traffic.

> **Recovery Point Objective Alerts**
>> Specify the length of time that you want to set for the recovery point objective (RPO) thresholds. The values determine whether a Warning or Severe alert is generated when the RPO threshold is exceeded for a role pair. The RPO represents the length of time in seconds of data exposure that is acceptable if a disaster occurs.

>> Use the following options to set the RPO threshold values. For both options, you can specify an RPO threshold in the range of 0 - 65535 seconds. The default is 0 seconds, which specifies that no alerts are generated.

>> **Warning level threshold (seconds)**
>>> Type the number of seconds that you want to set for the warning level RPO threshold. If the RPO is greater than this value, an alert is generated.

If the value in this field is other than 0, it must be greater than the value in the **Consistency group interval time (seconds)** field and less than the value in the **Severe level threshold (seconds)** field.

Severe level threshold (seconds)
Type the number of seconds that you want to set for the severe level RPO threshold. If the RPO is greater than this value, an alert is generated and the session status changes to Severe.

If the value in this field is other than 0, it must be greater than the value in the **Warning level threshold (seconds)** field.

Metro Global Mirror session properties

Use the View/Modify Properties notebook to view or modify properties for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 Metro Global Mirror sessions.

The properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session. Use the individual role pair option tabs to set properties that apply to specific role pairs. For example, properties for the H1-H2 role pair are contained in the **H1-H2 Options** tab.

Session Options
Enter the properties for the session.

Description
Type the description for this session.

Reset Secondary Reserves
Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets when a Start command is issued for the session.

Attention: This option causes the session to overwrite all data that is on the target volume.

Fail MM/GC if target is online (CKD only)
Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online state. For more information about this state, refer to the documentation for the storage system.

This option applies only to count key data (CKD) volumes.

H1-H2 Options
Enter the properties for the role pair.

Metro Mirror Suspend Policy
Select one of the following options to specify the policy for holding or releasing input/output (I/O) after a Metro Mirror relationship is suspended. When a relationship is in a Suspended state, write operations from the source volume are no longer mirrored to the target volume.

Hold I/O after Suspend
Click this option to prevent the source volume from receiving new data after the relationship is suspended. Use this option if you want to determine the scope of the

suspension before new data is written to the source volume. This option helps to ensure that the data on the source and target volumes remains the same.

To enable write operations to the source volume to continue, issue the Tivoli Storage Productivity Center for Replication `Release I/O` command for the session. If you do not issue this command, write operations are enabled when the hardware timeout value on the storage system expires.

Release I/O after Suspend
Click this option to enable the source volume to automatically receive data after the relationship is suspended. Use this option if you want to limit the effect on the applications that are writing to the source volume. This option is enabled by default.

Attention: This option can cause the data on the source volume to be different from the data on target volume.

z/OS Management
Select the following options to manage IBM z/OS features for the role pair.

The **z/OS Management** options are displayed only if Tivoli Storage Productivity Center for Replication is installed on or connected to a z/OS host system. You can add a z/OS host connections in Tivoli Storage Productivity Center for Replication regardless of the operating system on which the application is installed.

When you are connected to a z/OS host system, you can manage z/OS features for volumes that are attached to the host system. In addition, you can manage these features for volumes that are attached to other z/OS systems that are connected to the host system through a sysplex.

System or sysplex
Select the z/OS system or sysplex that you want to associate with the session. Volumes in the session must be attached to the system or sysplex that you choose to enable the z/OS features. Commands for the features are issued to this system or sysplex.

Enable Hardened Freeze
Select this option to enable the z/OS Input/Output Supervisor (IOS) to manage freeze operations.

Requirement: This option requires the z/OS address spaces Basic HyperSwap Management and Basic HyperSwap API. For instructions about how to start these address spaces, see the information about preparing to use Basic HyperSwap from z/OS in the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

If you select this option, IOS can freeze volumes regardless of whether the Tivoli Storage Productivity Center for Replication server is started or stopped.

In addition, this option enables you to include z/OS system volumes such as paging, database, and IBM WebSphere Application Server hierarchical file system (HFS) as Metro Mirror volumes in the session. When you select this option, IOS manages the freeze operations for all Metro Mirror volumes in the session, which prevents Tivoli Storage Productivity Center for Replication from freezing the volumes and possibly freezing itself. This option does not enable IOS to manage freeze operations for Global Mirror volumes.

If you select the **Manage H1-H2 with HyperSwap** option for the session, this option is ignored. HyperSwap includes IOS for managing freeze operations. The **Enable Hardened Freeze** option ensures data integrity if Tivoli Storage Productivity Center for Replication freezes and HyperSwap are not enabled for a session.

Manage H1-H2 with HyperSwap

Select this option to trigger a HyperSwap operation, which redirects application I/O to the target volumes when there is a failure on the host accessible volumes. Tivoli Storage Productivity Center for Replication uses HyperSwap to manage the H1-H2 sequence of a Metro Mirror or Metro Global Mirror session.

Disable HyperSwap

Select this option to prevent a HyperSwap operation from occurring.

On Configuration Error:

Partition the system(s) out of the sysplex

Select this option to partition a new system out of the sysplex when an error occurs because the system cannot be added to the HyperSwap configuration.

Disable HyperSwap

Select this option to prevent a HyperSwap operation from occurring.

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 enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

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 the HyperSwap operation does not complete.

Requirement: If you select this option, you must restart the system.

Disable HyperSwap after attempting backout
Select this option to enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

H1-J3 Options and H2-J3 Options
Enter the properties for the role pair.

Consistency group interval time (seconds)
Type how often, in seconds, the Global Mirror session attempts to form a consistency group. A lower value possibly reduces the data exposure of the session. However, a lower value also causes the session to attempt to create consistency groups more frequently, which can increase network traffic.

Recovery Point Objective Alerts
Specify the length of time that you want to set for the recovery point objective (RPO) thresholds. The values determine whether a Warning or Severe alert is generated when the RPO threshold is exceeded for a role pair. The RPO represents the length of time in seconds of data exposure that is acceptable if a disaster occurs.

Use the following options to set the RPO threshold values. For both options, you can specify an RPO threshold in the range of 0 - 65535 seconds. The default is 0 seconds, which specifies that no alerts are generated.

Warning level threshold (seconds)
Type the number of seconds that you want to set for the warning level RPO threshold. If the RPO is greater than this value, an alert is generated.

If the value in this field is other than 0, it must be greater than the value in the **Consistency group interval time (seconds)** field and less than the value in the **Severe level threshold (seconds)** field.

Severe level threshold (seconds)
Type the number of seconds that you want to set for the severe level RPO threshold. If the RPO is greater than this value, an alert is generated and the session status changes to Severe.

If the value in this field is other than 0, it must be greater than the value in the **Warning level threshold (seconds)** field.

Metro Global Mirror with Practice session properties

Use the View/Modify Properties notebook to view or modify properties for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 Metro Global Mirror with Practice sessions.

The properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session. Use the individual role pair option tabs to set properties that apply to specific role pairs. For example, properties for the H1-H2 role pair are contained in the **H1-H2 Options** tab.

Session Options

Enter the properties for the session.

Description

Type the description for this session.

Reset Secondary Reserves

Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets when a Start command is issued for the session.

Attention: This option causes the session to overwrite all data that is on the target volume.

Fail MM/GC if target is online (CKD only)

Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online state. For more information about this state, refer to the documentation for the storage system.

This option applies only to count key data (CKD) volumes.

H1-H2 Options

Enter the properties for the role pair.

Metro Mirror Suspend Policy

Select one of the following options to specify the policy for holding or releasing input/output (I/O) after a Metro Mirror relationship is suspended. When a relationship is in a Suspended state, write operations from the source volume are no longer mirrored to the target volume.

Hold I/O after Suspend

Click this option to prevent the source volume from receiving new data after the relationship is suspended. Use this option if you want to determine the scope of the suspension before new data is written to the source volume. This option helps to ensure that the data on the source and target volumes remains the same.

To enable write operations to the source volume to continue, issue the Tivoli Storage Productivity Center for Replication `Release I/O` command for the session. If you do not issue this command, write operations are enabled when the hardware timeout value on the storage system expires.

Release I/O after Suspend

Click this option to enable the source volume to automatically receive data after the relationship is suspended. Use this option if you want to limit the effect

on the applications that are writing to the source volume. This option is enabled by default.

Attention: This option can cause the data on the source volume to be different from the data on target volume.

z/OS Management

Select the following options to manage IBM z/OS features for the role pair.

The **z/OS Management** options are displayed only if Tivoli Storage Productivity Center for Replication is installed on or connected to a z/OS host system. You can add a z/OS host connections in Tivoli Storage Productivity Center for Replication regardless of the operating system on which the application is installed.

When you are connected to a z/OS host system, you can manage z/OS features for volumes that are attached to the host system. In addition, you can manage these features for volumes that are attached to other z/OS systems that are connected to the host system through a sysplex.

System or sysplex

Select the z/OS system or sysplex that you want to associate with the session. Commands for z/OS features are issued to this system or sysplex. Volumes in the session must be attached to the system or sysplex to enable the features.

Enable Hardened Freeze

Select this option to enable the z/OS Input/Output Supervisor (IOS) to manage freeze operations.

Requirement: This option requires the z/OS address spaces Basic HyperSwap Management and Basic HyperSwap API. For instructions about how to start these address spaces, see the information about preparing to use Basic HyperSwap from z/OS in the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

If you select this option, IOS can freeze volumes regardless of whether the Tivoli Storage Productivity Center for Replication server is started or stopped.

In addition, this option enables you to include z/OS system volumes such as paging, database, and IBM WebSphere Application Server hierarchical file system (HFS) as Metro Mirror volumes in the session. When you select this option, IOS manages the freeze operations for all Metro Mirror volumes in the session, which prevents Tivoli Storage Productivity Center for Replication from freezing the volumes and possibly freezing itself. This option does not enable IOS to manage freeze operations for Global Mirror volumes.

If you select the **Manage H1-H2 with HyperSwap** option for the session, this option is ignored. HyperSwap includes IOS for managing freeze operations. The **Enable Hardened**

Freeze option ensures data integrity if Tivoli Storage
Productivity Center for Replication freezes and HyperSwap
are not enabled for a session.

Manage H1-H2 with HyperSwap
Select this option to trigger a HyperSwap operation, which
redirects application I/O to the target volumes when there
is a failure on the host accessible volumes. Tivoli Storage
Productivity Center for Replication uses HyperSwap to
manage the H1-H2 sequence of a Metro Mirror or Metro
Global Mirror session.

Disable HyperSwap
Select this option to prevent a HyperSwap
operation from occurring.

On Configuration Error:

Partition the system(s) out of the sysplex
Select this option to partition a new system
out of the sysplex when an error occurs
because the system cannot be added to the
HyperSwap configuration.

Disable HyperSwap
Select this option to prevent a HyperSwap
operation from occurring.

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 enable IOS to back out
the HyperSwap operation, if possible, if an
error occurs during HyperSwap processing.
HyperSwap is disabled.

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
the HyperSwap operation does not
complete.

Requirement: If you select this option, you
must restart the system.

Disable HyperSwap after attempting backout
Select this option to enable IOS to back out
the HyperSwap operation, if possible, if an
error occurs during HyperSwap processing.
HyperSwap is disabled.

H1-J3 Options and H2-J3 Options
Enter the properties for the role pair.

Consistency group interval time (seconds)
Type how often, in seconds, the Global Mirror session attempts to form a consistency group. A lower value possibly reduces the data exposure of the session. However, a lower value also causes the session to attempt to create consistency groups more frequently, which can increase network traffic.

Recovery Point Objective Alerts
Specify the length of time that you want to set for the recovery point objective (RPO) thresholds. The values determine whether a Warning or Severe alert is generated when the RPO threshold is exceeded for a role pair. The RPO represents the length of time in seconds of data exposure that is acceptable if a disaster occurs.

Use the following options to set the RPO threshold values. For both options, you can specify an RPO threshold in the range of 0 - 65535 seconds. The default is 0 seconds, which specifies that no alerts are generated.

Warning level threshold (seconds)
Type the number of seconds that you want to set for the warning level RPO threshold. If the RPO is greater than this value, an alert is generated.

If the value in this field is other than 0, it must be greater than the value in the **Consistency group interval time (seconds)** field and less than the value in the **Severe level threshold (seconds)** field.

Severe level threshold (seconds)
Type the number of seconds that you want to set for the severe level RPO threshold. If the RPO is greater than this value, an alert is generated and the session status changes to Severe.

If the value in this field is other than 0, it must be greater than the value in the **Warning level threshold (seconds)** field.

H3-I3 Options
Enter the properties for the role pair.

No Copy
Select this option if you do not want the hardware to write the background copy until the source track is written to. Data is not copied to the H3 volume until the blocks or tracks of the I3 volume are modified.

This option is available only for IBM System Storage DS8000 version 4.2 or later.

Metro Mirror - Metro Mirror session properties

Use the View/Modify Properties notebook to view or modify properties for System Storage DS8000 Metro Mirror - Metro Mirror sessions.

If Tivoli Storage Productivity Center for Replication is installed on or connected to a z/OS host system, the properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session. Use the individual role pair option tabs to set z/OS management properties that apply to specific role pairs. For example, properties for the H1-H2 role pair are contained in the **H1-H2 Options** tab. Use the **HyperSwap Options** tab to set properties for managing HyperSwap operations for the session.

If Tivoli Storage Productivity Center for Replication is not installed on or connected to a z/OS host system, only the **Session Options** tab is displayed.

Connecting to a z/OS host: You can add a z/OS host connection in Tivoli Storage Productivity Center for Replication regardless of the operating system on which the application is installed. When you are connected to a z/OS host system, you can manage z/OS features for volumes that are attached to the host system. In addition, you can manage these features for volumes that are attached to other z/OS systems that are connected to the host system through a sysplex.

Session Options
Enter the properties for the session.

Description
Type the description for this session.

Fail MM/GC if target is online (CKD only)
Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online state. For more information about this state, refer to the documentation for the storage system.

This option applies only to count key data (CKD) volumes.

Reset Secondary Reserves
Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets when a Start command is issued for the session.

Attention: This option causes the session to overwrite all data that is on the target volume.

Metro Mirror Suspend Policy
Select one of the following options to specify the policy for holding or releasing input/output (I/O) after a Metro Mirror relationship is suspended. When a relationship is in a Suspended state, write operations from the source volume are no longer mirrored to the target volume.

Hold I/O after Suspend
Click this option to prevent the source volume from receiving new data after the relationship is suspended. Use this option if you want to determine the scope of the suspension before new data is written to the source volume. This option helps to ensure that the data on the source and target volumes remains the same.

To enable write operations to the source volume to continue, issue the Tivoli Storage Productivity Center for Replication `Release I/O` command for the session. If you

do not issue this command, write operations are enabled when the hardware timeout value on the storage system expires.

Release I/O after Suspend
Click this option to enable the source volume to automatically receive data after the relationship is suspended. Use this option if you want to limit the effect on the applications that are writing to the source volume. This option is enabled by default.

Attention: This option can cause the data on the source volume to be different from the data on target volume.

System or sysplex
Select the z/OS system or sysplex that you want to associate with the session. Commands for z/OS features are issued to this system or sysplex. Volumes in the session must be attached to the system or sysplex to enable the features.

If you select **No Association**, the options that are on the remaining tabs in the notebook are unavailable.

This option is displayed only if Tivoli Storage Productivity Center for Replication is installed on or connected to a z/OS host system.

H1-H2 Options, H1-H3 Options, and H2-H3 Options
Enter the properties for the role pair.

Manage H*x*-H*x* with HyperSwap
Select this option to trigger a HyperSwap operation, which redirects application I/O to the target volumes when there is a failure on the host accessible volumes. Tivoli Storage Productivity Center for Replication uses HyperSwap to manage the H1-H2, H1-H3, or H3-H2 sequence of the session.

HyperSwap Options
Enter the HyperSwap properties for the role pairs that are in the session. If you select **Manage Hx-Hx with HyperSwap** on the **H1-H2 Options**, **H1-H3 Options**, or **H2-H3 Options** tab, these properties define how HyperSwap Manager determines the target site for the HyperSwap operation, the priority for target sites, and the process that occurs in the event of an error during the operation.

Disable HyperSwap
Select this option to prevent a HyperSwap operation from occurring for all of the role pairs in the session. This option is typically selected to disable HyperSwap temporarily. For example, to disable HyperSwap for planned maintenance.

On Configuration Error:

Partition the system(s) out of the sysplex
Select this option to partition a new system out of the sysplex when an error occurs because the system cannot be added to the HyperSwap configuration.

Disable HyperSwap
Select this option to prevent a HyperSwap operation from occurring.

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 enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

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 the HyperSwap operation does not complete.
>
> If you select this option, you must restart the system.

Disable HyperSwap after attempting backout
> Select this option to enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

HyperSwap Site Selection
> Select one of the following options to specify how HyperSwap Manager selects the target site for a HyperSwap operation.

Allow HyperSwap Manager to determine the HyperSwap site
> Select this option to enable HyperSwap to select the target site for a HyperSwap operation.
>
> HyperSwap Manager selects the target site that keeps the most sysplex members active. If multiple target sites are equal candidates for the HyperSwap operation, HyperSwap Manager uses the **Site HyperSwap Priorities** value to select the target site. For example, if both site 1 and site 3 are equal, but the **Site HyperSwap Priorities** priority is site 3, site 2, and site 1, HyperSwap Manager selects site 3.

Determine the HyperSwap site by specified priorities
> Select this option to direct HyperSwap Manager to select the target site based on the priority list that is specified by the **Site HyperSwap Priorities** option.

HyperSwap Site Priorities
> Specify the priority in which HyperSwap Manager selects the target site for a HyperSwap operation.
>
> The default value priority is site 1, site 2, and site 3. This order indicates that the first priority is site 1, the second priority is site 2, and the third priority is site 3. If site 1 is the active site, then site 1 is ignored and HyperSwap Manager selects site 2 or 3 depending on the load status and availability of the sites, with site 2 being the preferred site.
>
> If you select **Allow HyperSwap Manager to determine the HyperSwap site**, this option is used only when multiple target sites are equal candidates for a HyperSwap operation.

If you select **Determine the HyperSwap site by specified priorities**, this option determines the order in which HyperSwap Manager selects the target site.

SAN Volume Controller and Storwize sessions

The properties for IBM System Storage SAN Volume Controller and IBM Storwize sessions depend on the session type.

The following sections describe the properties for each session type with the exception of session types that include only the **Description** property.

FlashCopy session properties

Use the View/Modify Properties notebook to view or modify properties for SAN Volume Controller and Storwize FlashCopy sessions.

The properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session. Use the **H1-T1 Options** tab to set properties that apply to the role pair.

Session Options
> Enter the properties for the session.
>
> > **Description**
> > > Type the description for this session.

H1-T1 Options
> Enter the properties for the role pair.
>
> > **Incremental**
> > > Select this option to apply incremental changes to the target volume. After the initial FlashCopy operation, only data that changed on the source volume since the last FlashCopy operation was performed is copied to the target volume.
> >
> > **Background copy rate (percentage)**
> > > Type the copy rate that the storage system uses to perform the background copy of the FlashCopy role pair. You can specify a percentage in the range of 0 - 100. The default is 50.
> > >
> > > Specify 0 if you do not want the hardware to write the background copy until the source track is written to. Data is not copied to the target volume until the blocks or tracks of the source volume are modified.
> > >
> > > You can modify this value at any time during the session. If the session is performing a background copy when you change the option, IBM 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.

Metro Mirror and Global Mirror Failover/Failback with Practice session properties

Use the View/Modify Properties notebook to view or modify properties for SAN Volume Controller and Storwize Metro Mirror Failover/Failback with Practice and Global Mirror Failover/Failback with Practice sessions.

The properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session. Use the **H2-I2 Options** tab to set properties that apply to the role pair.

Session Options
>Enter the properties for the session.

>>**Description**
>>>Type the description for this session.

H2-I2 Options
>Enter the properties for the role pair.

>>**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 changed since the previous FlashCopy operation.

>>**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 in the range of 0 - 100. The default is 50.

>>>Specify 0 if you do not want the hardware to write the background copy until the source track is written to. Data is not copied to the target volume until the blocks or tracks of the source volume are modified.

>>>You can modify this value at any time during the session. If the session is performing a background copy when you change the option, IBM 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.

Global Mirror Failover/Failback with Change Volumes session properties

Use the View/Modify Properties notebook to view or modify properties for SAN Volume Controller and Storwize Global Mirror Failover/Failback with Change Volumes sessions.

The properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session. Use the **H1-H2 Options** tab to set properties that apply to the role pair.

Session Options
>Enter the properties for the session.

>>**Description**
>>>Type the description for this session.

H1-H2 Options
>Enter the properties for the role pair.

>>**Enable change volumes**
>>>Select this option to enable the use of change volumes in the copy sets for the session. If this option is not selected, the session functions the same as a Global Mirror Failover/Failback session.

Change volumes are denoted as C*x*, where *x* identifies the site. These volumes contain point-in-time images that are copied from the H1 and H2 volumes. The C1 volume stores changes from the H1 volume. These changes are sent from the C1 volume to the H2 volume, and then to the C2 volume.

Because the data that is replicated between sites contains point-in-time changes rather than all changes, a lower bandwidth link is required between the sites when change volumes are used. However, the use of change volumes can result in an increase to data exposure. Therefore, you might want to enable or disable this option depending on your network traffic or business requirements.

When you select this option, it is the equivalent of setting the `-cyclingmode` parameter to `multi` on a SAN Volume Controller or Storwize storage system. When you clear this option, it is the equivalent of setting the `-cyclingmode` parameter to `none` for these storage systems.

Cycle period (seconds)
Specify the amount of time, in seconds, in which the change volumes will be refreshed with a consistent copy of the data. If a copy does not complete in the cycle period, the next cycle period will not start until the copy is complete. The range of possible values is 60 - 86400. The default is 300.

Recovery Point Objective Alerts
Specify the length of time that you want to set for the recovery point objective (RPO) thresholds. The values determine whether a Warning or Severe alert is generated when the RPO threshold is exceeded for a role pair. The RPO represents the length of time, in seconds, of data exposure that is acceptable if a disaster occurs.

Use the following options to set the RPO threshold values. For both options, you can specify an RPO threshold in the range of 0 - 172800 seconds. The default is 0 seconds, which specifies that no alerts are generated.

Warning level threshold (seconds)
Specify the number of seconds that you want to set for the warning level RPO threshold. If the RPO is greater than this value, an alert is generated.

If the value in this field is other than 0, it must be greater than the value in the **Cycle period (seconds)** field and less than the value in the **Severe level threshold (seconds)** field.

Severe level threshold (seconds)
Specify the number of seconds that you want to set for the severe level RPO threshold. If the RPO is greater than this value, an alert is generated and the session status changes to Severe.

If the value in this field is other than 0, it must be greater than the value in the **Warning level threshold (seconds)** field.

XIV sessions

The properties for the XIV sessions depend on the session type.

The following sections describe the properties for each session type with the exception of session types that include only the **Description** property.

XIV Global Mirror Failover/Failback session properties

Use the View/Modify Properties notebook to view or modify properties for the XIV Global Mirror Failover/Failback sessions.

The properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session. Use the **H1-H2 Options** to set properties that apply to the role pair.

Session Options
> Enter the properties for the session.
>
> > **Description**
> > > Type the description for this session.

H1-H2 Options
> Enter the properties for the role pair.
>
> > **Recovery point objective threshold (seconds)**
> > > Type the number of seconds that you want to set for the recovery point objective (RPO) threshold for the role pair.
> > >
> > > If the XIV determines that the RPO is greater than this threshold value, an alert is generated and the session status changes to Severe. You can specify an RPO in the range of 30 - 86400 seconds. The default is 30 seconds.
> >
> > **Synchronization schedule (HH:MM:SS)**
> > > Select an interval to create a synchronization schedule. The schedule is used to form consistency groups. The XIV 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 does not create consistency groups. When the XIV determines that the RPO threshold is exceeded, the session state becomes Severe.

Basic HyperSwap session properties

Use the View/Modify Properties notebook to view or modify properties for Basic HyperSwap sessions.

The properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session. Use the **H1-H2 Options** tab to set properties that apply to the role pair.

Session Options
> Enter the properties for the session.
>
> > **Description**
> > > Type the description for this session.
> >
> > **Fail MM/GC if target is online (CKD only)**
> > > Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online

state. For more information about this state, refer to the documentation for the storage system.

This option applies only to count key data (CKD) volumes.

H1-H2 Options

Enter the properties for the role pair.

z/OS Management

The **z/OS Management** options are displayed only if Tivoli Storage Productivity Center for Replication is installed on or connected to a z/OS host system. You can add a z/OS host connections in Tivoli Storage Productivity Center for Replication regardless of the operating system on which the application is installed.

When you are connected to a z/OS host system, you can manage z/OS features for volumes that are attached to the host system. In addition, you can manage these features for volumes that are attached to other z/OS systems that are connected to the host system through a sysplex.

System or sysplex

Select the z/OS system or sysplex that you want to associate with the session. Commands for z/OS features are issued to this system or sysplex. Volumes in the session must be attached to the system or sysplex to enable the features.

Disable HyperSwap

Select this option to prevent a HyperSwap operation from occurring.

On Configuration Error:

Partition the system(s) out of the sysplex

Select this option to partition a new system out of the sysplex when an error occurs because the system cannot be added to the HyperSwap configuration.

Disable HyperSwap

Select this option to prevent a HyperSwap operation from occurring.

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 enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

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 the HyperSwap operation does not complete.

Requirement: If you select this option, you must restart the system.

Disable HyperSwap after attempting backout
Select this option to enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

Viewing copy sets in a session

You can view the copy sets that are in a session. You can also export the copy sets to a comma-separated value (CSV) file for backup purposes.

Follow these steps to view or export the copy sets for a session:

1. In the menu bar, click **Sessions**.
2. On the Sessions page, select the session that contains the copy sets.
3. From the **Session Actions** list, select **View/Modify** > **View Copy Sets**.

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 menu bar, click **Storage** > **Storage Systems**.
2. On the Storage Systems page, click the ID for the storage system in **Storage System** column.

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 menu bar, click **Storage** > **Storage Systems**.
2. On the Storage Systems page, click the **Connections** tab.
3. Click the ID for the storage system connection in the **Storage Connection** column.

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 menu bar, click **Storage** > **Volumes**.
2. On the Volumes page, select a storage system.
3. Depending on the type of storage system, complete one of the following actions:
   - Select **All IO Groups** or a specific I/O group.

- Select **All Logical Storage Subsystems** or a specific logical storage subsystem.
- 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 TotalStorage Enterprise Storage Server, System Storage DS8000, or System Storage DS6000 storage system.

Complete one of these procedures to view logical paths:
- From the ESS/DS Paths page of Tivoli Storage Productivity Center for Replication:
    1. In menu bar, click **Paths**.
    2. Click the storage system ID to display logical paths for that storage system.
- From the Storage Systems page:
    1. In menu bar, click **Storage** > **Storage Systems**.
    2. Select the storage system for which you want to view logical paths.
    3. From the **Select Action** list, select **View Paths**. The paths page is displayed with a list of defined logical paths.

Viewing host connection details

You can view and modify the connection information for a host system. You can modify the connection information only if the host system in the Disconnected state.

Follow these steps to view or modify connection information for a host system:
1. In the menu bar, click **Storage** > **Host Connections**.
2. On the Host Connections page, click the link for the host system in the **Host System** column.

Viewing console messages

Tivoli Storage Productivity Center for Replication provides a detailed log of user and system activity.

To view this log, click **Console** in the menu bar. The Console window contains log entries that specify the activity that has occurred and a link to the correlating message for that activity.

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

To log on to Tivoli Storage Productivity Center for Replication, user ID and password is required.

The information in this topic is specific to Tivoli Storage Productivity Center for Replication on open systems. For information about defining users and groups for Tivoli Storage Productivity Center for Replication for System z, see the IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide.

Tivoli Storage Productivity Center for Replication does not maintain a directory of user names and passwords. Instead, the application uses one of the following repositories for user authentication:
- An operating system repository
- A Lightweight Directory Access Protocol (LDAP) repository

IBM Spectrum Control uses the federated repositories feature in WebSphere Application Server and an operating system repository is created by default during installation of IBM Spectrum Control and Tivoli Storage Productivity Center for Replication.

If you choose to use LDAP authentication, you must add the LDAP repository to the federated framework after the installation of IBM Spectrum Control and Tivoli Storage Productivity Center for Replication is complete.

For more information about the operating system and LDAP repositories, see
http://www.ibm.com/support/knowledgecenter/SS5R93_5.2.9/com.ibm.spectrum.sc.doc/fqz0_c_change_user_auth_method.html.

Adding IBM Spectrum Control users and groups to Tivoli Storage Productivity Center for Replication

By default, the common user that was defined during the installation of IBM Spectrum Control and the user tpcFileRegistryUser are granted access to Tivoli Storage Productivity Center for Replication and assigned to the Administrator role.

The user tpcFileRegistryUser is used only for recovery purposes. For example, if you accidentally delete the repository that you are using for authentication, you can access IBM Spectrum Control and Tivoli Storage Productivity Center for Replication by using the user tpcFileRegistryUser.

The password for the user tpcFileRegistryUser is the same as the password that was entered for the common user during the installation of Tivoli Storage Productivity Center.

To ensure smooth integration of IBM Spectrum Control with Tivoli Storage Productivity Center for Replication, complete the following actions:

*   Add all IBM Spectrum Control users and groups (other than those that are assigned by default) to Tivoli Storage Productivity Center for Replication. For example, if you added a TPCSuperuser group to IBM Spectrum Control for LDAP authentication, add that group to Tivoli Storage Productivity Center for Replication as well.
*   Use the same user or group to log on to both applications.

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.

Tivoli Storage Productivity Center for Replication provides a set of predefined user roles: Monitor, Operator, and Administrator.

Monitor

Monitors can view the health and status in the 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

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 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: 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

Administrators have unrestricted access. They can manage all sessions and perform all actions associated with 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 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

Note: Administrators cannot revoke their own administrative access rights.

Granting access privileges for a user

You can assign user roles to a 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. Log on to Tivoli Storage Productivity Center for Replication as a user with administrator privileges.
2. In the menu bar, click **Settings** > **Administration**.
3. On the Administration page, click **Add Access**. The Add Access wizard is displayed.
4. Type the name of the user to whom you want to give access, and click **Next**. The Select Users and Groups page is displayed.

   Tip: You can enter a partial name and use the * wildcard character to represent zero or more characters.
5. Select one or more names from the list of found users.
6. Select the role to associate with this user.
7. If you selected the Operator role, select one or more session that this user can manage, and click **Next**.
8. Click **Next** to confirm this action.
9. Click **Finish**.

Viewing access privileges for a user

You can view a list of all Tivoli Storage Productivity Center for Replication user and their assigned roles. You can also view the assigned sessions for each user.

Complete the following steps to view access privileges for a user:

1. Log in to Tivoli Storage Productivity Center for Replication as a user with administrator privileges.
2. In the menu bar, click **Settings** > **Administration**. The Administration page is displayed with a list of 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. From the **Select Action** list, select **View/Modify Access**. The View/Modify Access page is displayed. This page shows the role assigned to the user and lists the sessions that the user can manage.
5. Click **Cancel** to close the page.

Modifying access privileges for a user

You can change the user role and assigned sessions for a Tivoli Storage Productivity Center for Replication user.

Prerequisite: You must have Administrator privileges to perform this action.

Complete the following steps to modify the access privileges for a user:

1. Log in to Tivoli Storage Productivity Center for Replication as a user with administrator privileges.
2. In the menu bar, click **Settings** > **Administration**. The Administration page is displayed with a list of Tivoli Storage Productivity Center for Replication users and user groups and their associated role.
3. Select the user whose access privileges you want to modify.
4. From the **Select Action** list, select **View/Modify Access**. The View/Modify Access page is displayed. This page shows the role assigned to the user and lists the sessions that the user can manage.
5. Select the role to associate with the user.
6. If you selected the Operator role, select one or more sessions that the user can manage and click **Next**.
7. Click **OK**.

Removing access privileges for a user

You can remove access privileges for a Tivoli Storage Productivity Center for Replication user. When you remove access, the user ID cannot access the 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 Tivoli Storage Productivity Center for Replication as a user with administrator privileges.
2. In the menu bar, click **Settings** > **Administration**. The Administration page is displayed with a list of Tivoli Storage Productivity Center for Replication users and user groups and their associated role.
3. Select the user for which you want to remove access.
4. From the **Select Action** list, select **Remove Access**.

Appendix A. 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.
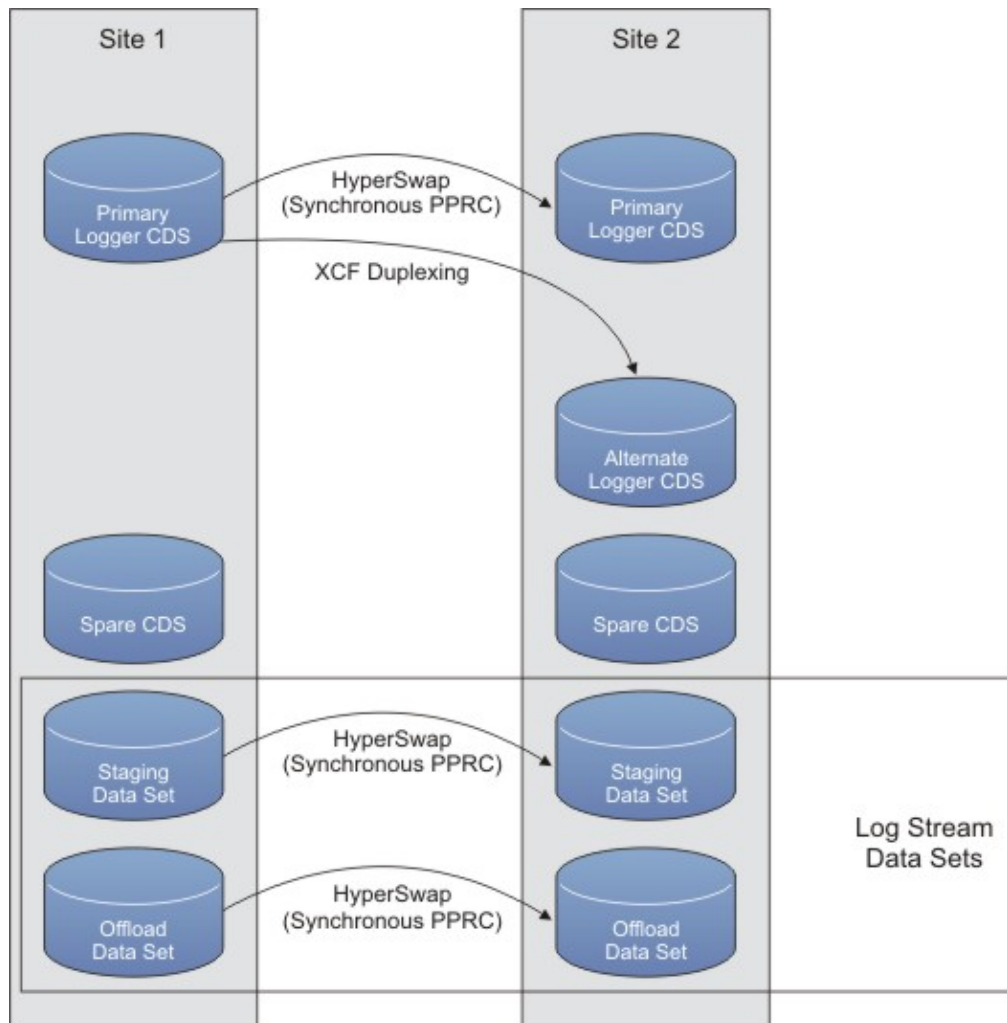
1. In the system logger policy, all CF log streams must be forced to duplex to staging data sets. The following data sets must be in the same Metro Mirror session:
   - Log stream staging data sets that are direct access storage device (DASD)-only
   - CF log stream data sets
   - All of the offload 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.

Appendix B. 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.

Accessibility features

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
- A Knowledge Center that includes the following accessibility features:
  - The Knowledge Center is provided in XHTML 1.0 format, which is viewable in most web browsers. With XHTML, you can view documentation according to the display preferences that are set in your browser. XHTML supports screen readers and other assistive technologies.
  - All documentation for Tivoli Storage Productivity Center for Replication is available in Adobe Portable Document Format (PDF) by using the Adobe Acrobat Reader. To access the PDFs, go to ⤴ ftp://public.dhe.ibm.com/ software/tivoli/tpc/ and select the v527 directory.
  - All images in the Knowledge Center are provided with alternative text so that visually impaired users can understand the contents of the images.

Keyboard navigation

Use the following key combinations to navigate the interface by keyboard:
- To go directly to the topic pane, 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, press Alt+R, and then press the Enter or Up-Arrow key to enter the view.
- To go directly to the Navigation (Table of Contents) view, press Alt+C, and then press the Enter or Up-Arrow key to enter the view.
- To expand and collapse a node in the navigation tree, press the Right and Left-Arrow keys.
- To move to the next topic node, press the Down-Arrow or Tab key.
- To move to the previous topic node, press the Up-Arrow key 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

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

Drop-down lists are positioned directly over 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 over 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 TotalStorage Enterprise Storage Server Model 800, System Storage DS6000, or System Storage 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 web browser GUI. 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 Windows Internet Explorer. Use the command-line interface instead of the GUI with JAWS on Windows Internet Explorer.

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.

IBM and accessibility

For more information about IBM's commitment to accessibility, see the IBM Human Ability and Accessibility Center website at http://www.ibm.com/able.

Notices

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

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

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not 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 character set (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.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, 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
2Z4A/101
11400 Burnet Road
Austin, TX 78758
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 measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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

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

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

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

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at http://www.ibm.com/privacy and IBM's Online Privacy Statement at http://www.ibm.com/privacy/details in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at http://www.ibm.com/software/info/product-privacy.

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.

Adobe is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

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.

Glossary

A glossary is available with terms and definitions for the IBM Spectrum Control family of products.

You can view the glossary in the IBM Spectrum Control product documentation at ➥ http://www.ibm.com/support/knowledgecenter/SS5R93/.

To view glossaries for other IBM products, see ➥ http://www.ibm.com/software/globalization/terminology/.

Index

A

about this document   xi
access privileges
   granting   206
   modifying   207
   removing   207
   viewing   206
accessibility features
   for Tivoli Storage Productivity Center
    for Replication   213
active management servers
   about   29
   reconnecting   36
   reinstalling during active session   35
   taking over   36
adding
   copy sets to a FlashCopy session   137
   copy sets to a Global Mirror
    session   145
   copy sets to a Metro Global Mirror
    session   148
   copy sets to a Metro Mirror
    session   142
   copy sets to a Snapshot session   140
   ESS/DS series logical paths   59
    using a CSV file   60
   host system connections   55
   SNMP managers   37
   storage connections
    DS6000   47, 48
    DS8000   46
    SAN Volume Controller   48
    Storwize   48
    TotalStorage Enterprise Server   47,
     48
    z/OS   49
administering
   users   203
architecture, product   4
awareness, site   45, 135

B

backing up
   considerations for   24
   database   26
Basic HyperSwap
   about   78
   session commands   107

C

client ports   8
comments, sending   xi
communication-failure trap
  descriptions   33
configuration change, SNMP trap
  descriptions   32
configuring
   logical paths   59

configuring *(continued)*
   storage systems   51
connections
   about   40
   direct   41
   Hardware Management Console
    (HMC)   42
   modifying
    host systems   56
   modifying properties   50
   z/OS   43
considerations, back up and recovery   24
consistency group   69
   time interval   69
copy sets
   about   64
   adding   137
   exporting   27, 154
   importing   27, 154
   removing   156
   supported number of role pairs and
    volumes   64
   viewing   201
creating
   FlashCopy session   137
   Global Mirror session   145
   Metro Global Mirror session   148
   Metro Mirror session   142
   Snapshot session   140
customizing session lists   163

D

data copying symbols   165
data exposure   69
data migration utility   157
database
   backing up   26
   restoring   26
disabling
   Metro Mirror heartbeat   153
   Open HyperSwap   105
   OpenSwap   105
disaster recovery, practicing   159

E

Easy Tier heat map transfer
   about   40
   configuring for   52
   log file location   40
   results, viewing   54
enabling, Metro Mirror heartbeat   153
ESS/DS storage systems
   adding logical paths   59
    using a CSV file   60
   removing logical paths   62
exporting
   copy sets   27, 154

exporting *(continued)*
   historical data for Global Mirror role
    pairs   154

F

failback   63, 69
failover   63, 69
FlashCopy
   about   79
   creating a session   137
   session commands   108

G

Global Mirror
   about   88
   creating a session   145
   session commands   114
   session details, viewing   170
Global Mirror with Practice, session
  commands   118
granting
   access privileges   206

H

health, viewing summary of   163
heartbeat, enabling and disabling   153
host systems
   adding a connection to   55
   managing   55
   modifying a connection to   56
   removing   57, 58
   viewing connection information
    for   202

I

IBM Tivoli Storage Productivity Center
  for Replication
   starting   21
IBM Tivoli Storage Productivity Center
  for Replication for System z
  overview   1
icons
   session status   164
   storage systems status   39
identifying the version of Tivoli Storage
  Productivity Center for Replication   24
importing copy sets   27, 154

J

journal volume, restoring data from   52

IBM. ®

Printed in USA