

**IBM Storage-based high
availability and disaster
recoveryStorage-based high
availability and disaster
recoveryStorage-based high
availability and disaster
recovery**

Contents

Highlighting.....	6
Case-sensitivity in AIX®	6
ISO 9000.....	6
Related information	6

Storage-based high availability and disaster recovery for PowerHA® SystemMirror® Enterprise Edition 7

What's new in Storage-based high availability and disaster recovery for PowerHA® SystemMirror® Enterprise Edition for AIX®	7
How to see what's new or changed	7
July 2025.....	7
December 2024.....	7
December 2023.....	7
December 2022.....	7
July 2021.....	7
December 2019.....	7
June 2018	8
DS8000® PPRC replication resources	8
Metro Mirror.....	8
Global Mirror	67
SVC replication resources.....	74
Overview of SVC management	74
Planning for SVC management	75
Installing PowerHA® SystemMirror® Enterprise Edition Metro Mirror for SVC.....	81
Configuring PowerHA® SystemMirror® Enterprise Edition Metro Mirror for SVC	81
Changing an PowerHA® SystemMirror® Enterprise Edition for Metro Mirror SVC configuration.....	87
Troubleshooting PowerHA® SystemMirror® Enterprise Edition for Metro Mirror for SVC	90
Configuring PowerHA® SystemMirror® Enterprise Edition for SVC policy based replication	92
Adding an SVC policy configuration.....	93
SVC Policy Based Replication resources.....	93
Overview of SVC policy-based replication management.....	93
Configuring PowerHA® SystemMirror® Enterprise Edition policy based replication for SVC flash storage ..	94
Changing an PowerHA® SystemMirror Enterprise Edition for SVC PBR configuration	98
Troubleshooting PowerHA® SystemMirror® Enterprise Edition for PBR for SVC	100
Disaster recovery management with EMC SRDF replication	101
EMC SRDF concepts	101
Planning for SRDF replication	104
Installing SRDF filesets	108
Configuring SRDF replicated resources	109
Monitoring EMC SRDF/Metro SmartDR configuration	112
Changing SRDF replicated resources	113
Removing a SRDF replicated resource.....	114
Hitachi Truecopy/HUR replication resources.....	114
Overview of Truecopy/HUR Management.....	114
Planning for Truecopy/HUR Management	115
Configuring volume groups and filesystems on Truecopy/HUR protected disks	117
Installing Truecopy management filesets	119
Configuring Truecopy/HUR resources.....	120
Changing Truecopy/HUR-Managed replicated resource	122
Removing Truecopy/HUR-Managed replicated resource.....	122
XIV® Remote Mirror replication resources	123
XIV® Remote Mirror concepts	123
Planning for XIV® Remote Mirror	124
Installing XIV® Remote Mirror	125
Configuring XIV® Remote Mirror	125
HyperSwap® for PowerHA® SystemMirror®	131
HyperSwap® for PowerHA® SystemMirror® concepts	131
Planning for HyperSwap® for PowerHA® SystemMirror®	142
Installing HyperSwap® for PowerHA® SystemMirror®	145
Configuring HyperSwap® for PowerHA® SystemMirror®	145
Managing mirror groups for HyperSwap®	157
Troubleshooting HyperSwap® configurations	161
Failover rehearsal.....	163
Example.....	163
Configuring Fibre Channel disks	163
Troubleshooting manual recovery for storage subsystems.....	164

Notices..... 166
 Privacy policy considerations167
 Trademarks.....167

Index 168

Note

Before using this information and the product it supports, read the information in [“Notices” on page 166](#).

This edition applies to IBM® PowerHA® SystemMirror® 7.2 Enterprise Edition for AIX® and to all subsequent releases and modifications until otherwise indicated in new editions.

About this document

This document describes the steps that are required to plan, install, and configure different types of replicated recourses to work in your environment with PowerHA® SystemMirror® Enterprise Edition for AIX®.

Highlighting

The following highlighting conventions are used in this document:

Bold	Identifies commands, subroutines, keywords, files, structures, directories, and other items whose names are predefined by the system. Also identifies graphical objects such as buttons, labels, and icons that the user selects.
<i>Italics</i>	Identifies parameters whose actual names or values are to be supplied by the user.
Monospace	Identifies examples of specific data values, examples of text similar to what you might see displayed, examples of portions of program code similar to what you might write as a programmer, messages from the system, or information you should actually type.

Case-sensitivity in AIX®

Everything in the AIX® operating system is case-sensitive, which means that it distinguishes between uppercase and lowercase letters. For example, you can use the **ls** command to list files. If you type **LS**, the system responds that the command is not found. Likewise, **FILEA**, **FiLea**, and **filea** are three distinct file names, even if they reside in the same directory. To avoid causing undesirable actions to be performed, always ensure that you use the correct case.

ISO 9000

ISO 9000 registered quality systems were used in the development and manufacturing of this product.

Related information

- The PowerHA® SystemMirror® Version 7.2 for AIX® PDF documents are available in the [PowerHA SystemMirror 7.2 PDFs](#) topic.
- The PowerHA® SystemMirror® Version 7.2 for AIX® release notes are available in the [PowerHA SystemMirror 7.2 release notes](#) topic.

Storage-based high availability and disaster recovery for PowerHA® SystemMirror® Enterprise Edition



This information describes the necessary steps to plan, install, configure, and maintain Metro Mirror, Global Mirror, SRDF management replicated resources, HyperSwap®, and Hitachi Truecopy/HUR management.

What's new in Storage-based high availability and disaster recovery for PowerHA® SystemMirror® Enterprise Edition for AIX®

Read about new or significantly changed information for the storage-based high availability disaster recovery topic collection.

How to see what's new or changed

To help you see where technical changes have been made, the information center uses:

- The  image to mark where new or changed information begins.
- The  image to mark where new or changed information ends.

July 2025

The following information is a summary of the updates that were made to PowerHA SystemMirror Version 7.2.9 Service Pack 1:

Added new topics about configuring, changing and troubleshooting policy-based replication (PBR) replication services in the [“SVC Policy Based Replication resources” on page 93](#).

December 2024

Added information about configuring the policy-based replication capability in the [“Configuring PowerHA SystemMirror Enterprise Edition for SVC policy based replication” on page 92](#) topic.

December 2023

Updated the steps to install SVC filesets for PowerHA SystemMirror Enterprise Edition. For more information see. [“Installing PowerHA SystemMirror Enterprise Edition Metro Mirror for SVC” on page 81](#)

December 2022

PowerHA® SystemMirror® Version 7.2.7 added support for DS8800 Metro Global Mirror support in integration with Copy Services Manager (CSM). To know more, see [Support for Metro Global Mirror configuration](#) topic.

July 2021

PowerHA® SystemMirror® Version 7.2.5 SP1 added [EMC SRDF/Metro SmartDR configuration](#), which is a two-region High Availability and Disaster Recovery (HADR) framework, that integrates SRDF/Metro and SRDF/Async replicated resources.

Note: EMC SRDF/Metro SmartDR configuration is available only in PowerHA® SystemMirror® 7.2.5 Enterprise Edition for AIX®.

December 2019

Updated information about non-root user support for SVC mirror management in the following topics:

- [“Changing an SVC cluster configuration” on page 87](#)
- [“SVC and PowerHA SystemMirror configuration requirements” on page 82](#)
- [“Adding an SVC cluster configuration” on page 82](#)

June 2018

Upgraded the version information and added a note in the [Installing SRDF filesets](#) topic.

DS8000® PPRC replication resources

DS8000® series of products belong to IBM® System Storage® (DS) Family Storage Solutions. System Storage® family products provide for an a robust, highly scalable storage solution. These products provide for rich set of storage management functions and additionally support both data mirroring for both shorter distances (100 to 200 KMs) and longer distances (1000's of KMs).

Shorter distance mirroring is typically performed synchronously and is called Metro mirroring. Longer distance mirroring is performed asynchronously and is called Global Mirror. For the DS8000® storage family you can use mirroring with PowerHA® SystemMirror® Enterprise Edition to provide High Availability and Disaster Recovery Solution across sites.

Metro Mirror

Review the following information provides to plan, install, configure, and maintain the PowerHA® SystemMirror® Enterprise Edition for Metro Mirror (previously known as Synchronous PPRC).

Overview of PowerHA® SystemMirror® Enterprise Edition for Metro Mirror

PowerHA® SystemMirror® Enterprise Edition for synchronous Peer-to-Peer Remote Copy (PPRC), now known as and referred to as Metro Mirror increases data availability for IBM® TotalStorage® volumes that use to copy data to a remote site for disaster recovery purposes.

PowerHA® SystemMirror® Enterprise Edition for Metro Mirror takes advantage of the following components to reduce downtime and recovery time during disaster recovery:

- PowerHA® SystemMirror® cluster management
- PPRC failover and fallback functions
- Optional components:
 - DSCLI PPRC support
 - SAN Volume Controller (SVC) PPRC support

Prerequisites

Before using PowerHA® SystemMirror® Enterprise Edition for Metro Mirror documentation, you should be familiar with:

- PowerHA® SystemMirror® installation and administration
- General background and Metro Mirror support for the type of implementation you intend to configure:
 - Direct management PPRC (ESS systems)
 - DSCLI PPRC (ESS or DS systems)
 - San Volume Controller (SVC) PPRC

PowerHA® SystemMirror® Enterprise Edition for Metro Mirror features

PowerHA® SystemMirror® Enterprise Edition for Metro Mirror extends PowerHA® SystemMirror® cluster management for highly available applications and servers to support the disaster recovery mechanism supplied by PPRC.

PPRC is a hardware mirroring technique that the IBM® TotalStorage® Enterprise Storage Server® uses to replicate data. It allows mirroring to be suspended and restarted without affecting data integrity.

For all types of Metro Mirror support, PowerHA® SystemMirror® helps to manage PPRC instances. Direct management PPRC is the longest-running type of support through PowerHA® SystemMirror® Enterprise Edition, designed to provide basic PPRC management for ESS systems. In response to demand for simpler management

interfaces, PowerHA® SystemMirror® also provides more automated PPRC management (notably in the area of path and instance (pair) creation) via the following configurations:

- DSCLI management for PowerHA® SystemMirror® Enterprise Edition Metro Mirror allows additional flexibility by supporting both ESS and DS storage subsystems with the same interface.
- SAN Volume Controller (SVC) Management for PowerHA® SystemMirror® Enterprise Edition Metro Mirror provides storage virtualization and an additional layer of disaster recovery via the SVC cluster and hardware configuration.

In general, the styles of Metro Mirror support mentioned previously (DSCLI and SVC) are intended to make Metro Mirror management easier to install and manage. Direct management PPRC has the advantage of *not* having "middle man" software and hardware to maintain, but, in exchange, there is a relatively heavy amount of configuration and maintenance tasks for the administrator.

High Availability and Disaster Recovery Support Features

PowerHA® SystemMirror® Enterprise Edition for Metro Mirror provides high availability and disaster recovery with:

- Automatic failover of PPRC-protected volume pairs between nodes within a site
- Automatic failover of PPRC-protected volume pairs between sites
- Automatic recovery/reintegration of PPRC-protected volume pairs between sites
- Support for user-defined policy-based resource groups
- Support for the following Inter-Site Management Policies for resource groups:
 - Prefer Primary Site
 - Online on Either Site
- Support for the Subsystem Device Driver Path Control Module (SDDPCM).
- Support for cluster verification and synchronization
- Support for C-SPOC for cluster administration.

PPRC mirroring

PPRC is a mirroring technique used to maintain consistent copies of data across two ESS systems that are at different locations. It mirrors at the disk subsystem level, making it transparent to hosts. PPRC is also called Remote Copy and Mirror.

The three types of remote copy and mirror are:

- **Metro Mirror**
Synchronous mirroring
- **Global Mirror**
Asynchronous mirroring
- **Global Copy**
Extended Distance mirroring

In this topic collection, the term PPRC is synonymous with synchronous mirroring.

Disk subsystems

IBM® TotalStorage® Disk subsystems (Enterprise Storage System - ESS, or DS) are large disk subsystems that are typically configured to provide RAID-5 and RAID-10 volumes. They contain two System p® servers connected to the system backplane to manage and monitor the system.

The disk drives are configured to provide groups of logical disks that are divided into partitions. A data partition in a disk group is referred to as a volume. Each volume is associated with a logical unit number (LUN) that is a subdivision of the physical disk on a storage system.

The disk subsystems have connections to:

- System p® (or other) host servers over SCSI or Fibre Channel links
- Other disk subsystems over ESCON® or Fibre Channel links. (PPRC uses these physical links to mirror data.)
- ESSNet, the local management Ethernet network for the ESS
Commands sent over this network configure and manage PPRC through ESS Copy Services.
- Nodes that access stored data.

Note: PowerHA® SystemMirror® Enterprise Edition for Metro Mirror requires that inter-storage links be available to carry data in both directions at the same time. If ESCON® links are used, a minimum of two ESCON® links is required because each link can carry data in only one direction at a time. You should have at least four links to improve throughput and to provide redundancy of the ESCON® cables and adapters.

PPRC overview

When PPRC is activated, it establishes a mirror between specified volumes on two storage systems. The PPRC mirrored volumes are referred to as a *PPRC pair*, *instance*, or as a *PPRC-protected volume*.

PPRC can mirror data in one of two ways:

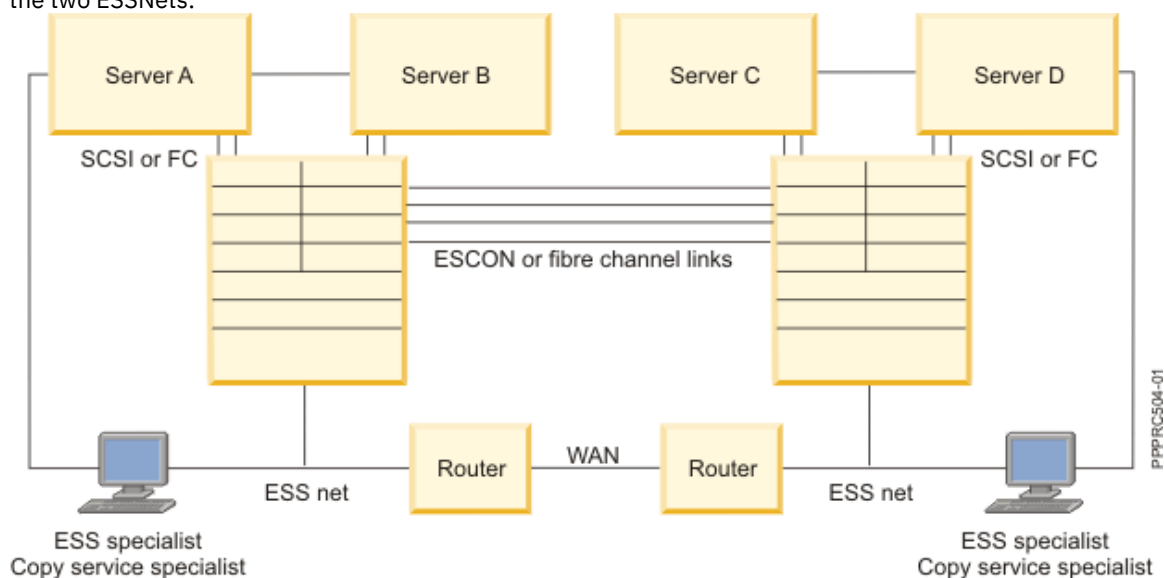
- Synchronous mirroring provides concurrent copy over ESCON® links. This is supported by PowerHA® SystemMirror® Enterprise Edition .
Note that ESCON® links transmit data in one direction and have distance limitations.
- Non-synchronous mirroring, also called Global Mirror mirroring, provides copying over longer distances. Extended distance mirroring ensures that volumes are consistent with each other at a point-in-time. This method is not supported in PowerHA® SystemMirror® Enterprise Edition.

Note: FlashCopy® is an optional feature on the ESS. It is an instant (point-in-time) copy of the data.

Hardware configuration for PPRC mirroring

The following illustrations show sample hardware configurations that support PPRC-mirroring. The illustrations also show the connections from the ESS to:

- Nodes over SCSI or Fibre Channel connections
- Fibre channel or ESCON® links that connect the two ESS systems to mirror data
- The ESS Specialist and ESS Copy Services Specialist over ESSNet and a wide-area network (WAN) that joins the two ESSNets.



Optional Metro Mirror configurations

These topics are a light overview of each of the optional PowerHA® SystemMirror® Enterprise Edition for Metro Mirror configurations available.

DSCLI management

DSCLI management is one type of optional configuration of PowerHA® SystemMirror® Enterprise Edition for Metro Mirror. It provides a simplified PPRC interface for both the ESS and DS storage hardware.

The DSCLI interface provides simplified management of PPRC paths and instances in the following manners:

- Provides a simplified interface to IBM® TotalStorage® PPRC services on ESS or DS storage systems in order to allow management and reporting on PPRC instances and paths.
- Monitors the status of PPRC relationships and consistency groups between the volumes being mirrored. It reports any change in status, such as a volume moving to an offline state.

The DSCLI client software interfaces with the ESSNI server on the HMC, SMC or controller connected to the storage to which the DSCLI client is connected.

The following list highlights a few of the specific types of functions that you can perform with the DS command-line interface:

- Check and verify your storage unit configuration.
- Check the current Copy Services configuration that is used by the storage unit.
- Create new logical storage and Copy Services configuration settings.
- Modify or delete logical storage and Copy Services configuration settings.

For more information on DSCLI configuration, refer to the current DSCLI online or otherwise published documentation:

(http://publib.boulder.ibm.com/infocenter/dsichelp/ds8000ic/index.jsp?topic=/com.ibm.storage.ssic.help.doc/f2c_cliesscli_1kx2so.html)

SVC management

SVC management is another optional configuration of PowerHA® SystemMirror® Enterprise Edition for Metro Mirror.

SVC provides both storage virtualization and PPRC management services in the following ways:

- Provides virtualized storage that interfaces with the TotalStorage® PPRC services to provide management services and reporting for PPRC relationships (instances,) and consistency groups.
- Monitors the status of PPRC relationships and consistency groups between the volumes being mirrored. It reports any change in status, such as a volume moving to an offline state.
- Responds to a site failure (or possible site failure), by suspending all mirroring, and, if necessary, activating the mirror copy at the backup site for host access. All data changes are tracked until the mirrors can be synchronized again.
- Synchronizes the volumes.

PowerHA® SystemMirror® Enterprise Edition for Metro Mirror in an PowerHA® SystemMirror® cluster

PowerHA® SystemMirror® Enterprise Edition for Metro Mirror allows you to include PPRC-mirrored volumes in an PowerHA® SystemMirror® cluster.

This requires two PowerHA® SystemMirror® sites—an PowerHA® SystemMirror® component to which you assign nodes. Cluster nodes access the same shared volume groups, but the nodes at each site access them from different physical volumes—the two volumes in a single PPRC pair. This is different from a single-site PowerHA® SystemMirror® environment, in which all cluster nodes sharing volume groups have physical connections to the same set of disks.

PPRC replicated resources

A PPRC replicated resource is an PowerHA® SystemMirror® resource that manages a PPRC pair (has a primary and secondary instance that is copied from one site to another). The managing resource group definition includes the volume groups and filesystems built on top of the volume groups defined to a PPRC replicated resource.

PowerHA® SystemMirror® Enterprise Edition for Metro Mirror supports three types of replicated resources:

- *Direct management PPRC replicated resources.* A PPRC replicated resource is a PPRC pair associated with an PowerHA® SystemMirror® site that is included in an PowerHA® SystemMirror® cluster. The definition for a PPRC replicated resource contains a volume identifier and the name of the associated ESS. PowerHA® SystemMirror® has knowledge of which volumes are mirrors of each other for each PPRC replicated resource. Synchronous mirroring is supported.
- *DSCLI-managed PPRC replicated resources.* A DSCLI-managed PPRC replicated resource is a definition of a set of volume pairs and the paths needed to communicate between them. The resource group definition includes the volume groups built on top of the PPRC replicated volumes. Synchronous mirroring is supported.
- *SVC-managed PPRC replicated resources.* An SVC-managed PPRC replicated resource contains the SVC virtual disk (vDisk) volume pair information combined with the SVC Cluster name bundled in PPRC relationships and consistency group. Synchronous mirroring is supported.

Resource groups that include PPRC replicated resources

An PowerHA® SystemMirror® resource group is a collection of resources that comprise the operating environment for an application. Applications, as resources in a resource group, are made highly available.

Resource group management policies direct which node hosts the resource group during normal operation and when the host node fails or goes offline. With PowerHA® SystemMirror® Enterprise Edition for Metro Mirror, resource group configuration is the same as for other resource groups. In addition, the resource group includes

- A shared volume group and the PPRC replicated resources associated with the individual volumes in the volume group
- An intersite management policy to handle a resource group during site recovery.

Limitations and restrictions on resource groups that include PPRC replicated resources

Refer to the section for the PowerHA® SystemMirror® Enterprise Edition Metro Mirror configuration you are implementing for additional Limitations lists. In general, the following restrictions apply to PowerHA® SystemMirror® resource groups that will manage PPRC replicated resources because of the way that PPRC instances are managed (i.e. source site nodes will have I/O access, target site nodes will not.) The outcome of this is that any PowerHA® SystemMirror® policy that allows a resource group to come online on more than one site at a time is not supported:

- Inter-Site Management Policy **Online Both Sides** is not supported.
- Startup Policies of **Online Using Distribution Policy** and **Online on All Available Nodes** are not supported.
- Fallover Policy **Fallover Using Dynamic Node Priority** is not supported.

PowerHA® SystemMirror® sites

PowerHA® SystemMirror® Enterprise Edition for Metro Mirror support requires the use of sites. PowerHA® SystemMirror® supports two sites. The primary site is the active site, and the secondary site is the standby site.

The Inter-Site Management Policy for a resource group directs how a resource group and its resources fall over in response to an outage, and how they fall back if configured to do so. For each resource group, one site is an active production site and the other a backup site. If the nodes at the active production site become unavailable, the backup site becomes the active production site.

For non direct-management types of PowerHA® SystemMirror® Enterprise Edition for Metro Mirror support (DSCLI, SVC PPRC), each site contains at least one storage system and the nodes attached to it. For PowerHA® SystemMirror® Enterprise Edition Direct management environments, each site contains only one ESS.

Resource groups have two types of management policies:

- Resource group management policies determine fallover behavior if a node becomes unavailable.

- Site management policies determine fallover behavior if all of the nodes at a site are *not* available.

Fallover and fallback

PowerHA® SystemMirror® Enterprise Edition for Metro Mirror handles the automation of fallover from one site to another in response to an outage at a production site, minimizing recovery time. When a site fails, the resource group configuration determines whether source volumes are accessible from the secondary site.

PowerHA® SystemMirror® automates application recovery by managing replicated resources defined to PowerHA® SystemMirror® resource groups in the following manner:

- The fallover of nodes within a site based on node priority (as identified in the nodelist for a resource group)
- The fallover between sites (as specified by the site management policy for a resource group)
- The fallback of a resource group or site as configured.

When an application is running on an active production site:

- Updates to the application data are made to the disks associated with the active production site.
- Data is mirrored using PPRC to the backup disks.

If the node or the disks at the production site become unavailable:

- The application moves to a server at the backup site.
- The application continues operation using the mirrored copy of the data.

When the initial production site becomes active again, resource group and site management policies determine whether or not the application moves back to the previous site:

- The direction of mirroring may be reversed.
- The application may be stopped and restarted on another node.
- Manual intervention may be required to bring the application back to a nominal functioning state.

PowerHA® SystemMirror® Enterprise Edition for Metro Mirror general planning

These topics describe planning tasks required for all the types of PowerHA® SystemMirror® Enterprise Edition Metro Mirror support.

PowerHA® SystemMirror® Enterprise Edition for Metro Mirror management types

PowerHA® SystemMirror® Enterprise Edition for Metro Mirror supports both synchronous and extended-distance PPRC mirrored volumes. The method used to manage PPRC depends on the hardware configuration available.

Using the following table, you can decide what type of PPRC configuration you will need to manage your cluster.

Mirror type	Hardware Type	PPRC is managed by	How PowerHA® SystemMirror® manages PPRC pairs
Synchronous	ESS 800, etc.	Copy Services Server (CSS, on storage controller)	Directly manages the fallover and resynchronization of the PPRC pairs by issuing commands directly to the ESS systems. In this book, this is referred to as Direct management. In general, it refers to PPRC between ESS storage systems.
Synchronous	ESS (800) or DS (8000, 6000) or intermix of any of these	DSCLI management, via ESSNI Server on either storage controller or Hardware Management Console (HMC)	Relies on the ESSNI server to manage PPRC pairs. Directly manages the fallover and resynchronization of the PPRC pairs by issuing commands directly to the storage systems. In this book, this is referred to as DSCLI management.

Mirror type	Hardware Type	PPRC is managed by	How PowerHA® SystemMirror® manages PPRC pairs
Synchronous	vDisks as provided by SAN Volume Controller. (Hardware as supported by SVC services)	SVC management of PPRC services on SVC-specific hardware	Relies on the Copy Services Server in order to manage PPRC function. Directly manages the fallover and resynchronization of the PPRC pairs by issuing commands directly to the Copy Services Server. In this book, this is referred to as PPRC SVC management.

Coexistence of two solutions

PowerHA® SystemMirror® Enterprise Edition for Metro Mirror solutions can coexist on the same PowerHA® SystemMirror® cluster only if the PPRC pairs are managed by one of the PPRC solutions at a time. Please refer to the latest support information for which PPRC solutions can successfully coexist on a single PowerHA® SystemMirror® cluster.

Planning overview

Once you have decided what type of PowerHA® SystemMirror® Enterprise Edition for Metro Mirror implementation to set up, you can begin planning the configuration.

At this stage, you should be familiar with the planning tasks for base PowerHA® SystemMirror®.

The following planning is required for all PowerHA® SystemMirror® Enterprise Edition for Metro Mirror implementations:

- Identify sites for PowerHA® SystemMirror® Enterprise Edition for Metro Mirror.
- Identify the resource groups needed to manage the PPRC replicated resources (*optional* at this point, can be done at a later stage).
- Identify storage systems to be used in the configuration.
- Plan for connections to the storage units.
- Plan the PowerHA® SystemMirror® Enterprise Edition for Metro Mirror configuration.

Identify the type of PowerHA® SystemMirror® Enterprise Edition for Metro Mirror support to be used.

Planning PowerHA® SystemMirror® sites

Within a resource group, the nodes at one site may handle the PPRC replicated resource differently than the nodes at the other site, especially in cases where the states of the volumes (suspended or full duplex) are different at the two sites.

PowerHA® SystemMirror® Enterprise Edition for Metro Mirror requires two PowerHA® SystemMirror® sites for use within a resource group to control which volume in a PPRC pair a node can access. Although nodes at both sites can access a volume group, access is permitted to only one volume in a PPRC pair at a time - the source volume. This prevents nodes at different sites from accessing the same volume group at the same time. Typically, a number of volumes are mirrored through PPRC from one site to the other.

Planning PowerHA® SystemMirror® Enterprise Edition for Metro Mirror resource groups

In addition to basic resource group planning, you need to plan for resource group attribute such as sites and PPRC replicated resources.

Note: An PowerHA® SystemMirror® resource group can contain multiple volume groups, file systems, PPRC replicated resources, applications, etc.

When configuring a resource group to use a PPRC replicated resource, the resource group includes nodes in both cluster sites. For the resource group to manage nodes and resources across sites, you assign one of the following inter-site management policies:

- **Prefer Primary Site** . In a two-site configuration, replicated resources at startup are on the site with the highest priority, fall over to the other site, and then fall back to the site with the highest priority.
- **Online on Either Site** . Replicated resources are on either site at startup, fall over to the other site and remain on that site after failover. This selection simplifies resource group takeover rules, which is helpful if you have a number of resource groups.

Note: Each PPRC pair will be included in a resource group whose home node is located at the site that is primary for that PPRC pair.

If you want to set up a mutual recovery configuration for ESS volumes, you configure one resource group in which one site is the active site and another resource group in which the other site is the active site. For an example of a mutual recovery configuration, see [Sample configuration for Direct management](#).

Resource group limitations

In PowerHA® SystemMirror® Enterprise Edition for Metro Mirror configurations, all volume groups in a resource group must contain only PPRC-protected disks.

If PowerHA® SystemMirror® Enterprise Edition for Metro Mirror is not being used, all volume groups in a resource group must contain only non-PPRC-protected disks.

Each PPRC pair will be included in a resource group whose home node is located at the site that is primary for that PPRC pair.

Installing PowerHA® SystemMirror® Enterprise Edition for Metro Mirror

These topics describe how to install base PowerHA® SystemMirror® Enterprise Edition filesets. Details for installing filesets and prerequisites for specific types of PPRC support are contained in subsequent sections.

Note: At runtime, make sure you have at least 9 MB of free space in the `/tmp` directory. PPRC commands use this directory.

PowerHA® SystemMirror® Enterprise Edition for Metro Mirror is installed in addition to the PowerHA® SystemMirror® base filesets.

In order to install PowerHA® SystemMirror® Enterprise Edition for Metro Mirror, you must install the **cluster.xd.license** fileset in addition to the base PowerHA® SystemMirror® **cluster.license** fileset.

Installation prerequisites for PowerHA® SystemMirror® Enterprise Edition Metro Mirror

Before installing PowerHA® SystemMirror® Enterprise Edition Metro Mirror, be sure to have the necessary base PowerHA® SystemMirror® filesets installed.

PowerHA® SystemMirror® Enterprise Edition requires base PowerHA® SystemMirror®. Base PowerHA® SystemMirror® can be installed at the same time as PowerHA® SystemMirror® Enterprise Edition. Installation requisites will ensure the proper version(s) of base PowerHA® SystemMirror® filesets are installed before installing PowerHA® SystemMirror® Enterprise Edition.

At least the following base PowerHA® SystemMirror® filesets should be installed on all cluster nodes:

Note: Fileset versions must reflect the version of PowerHA® SystemMirror® you are installing.

In addition, refer to the following table for a list of filesets to install for each type of PPRC management support:

PPRC management type	Filesets to install
Direct management	<pre>cluster.es.pprc.cmds cluster.es.pprc.rte cluster.msg.en_US.pprc</pre> <p>(and other appropriate language message sets)</p>
DSCLI	<pre>cluster.es.pprc.cmds cluster.es.pprc.rte cluster.es.spprc.cmds cluster.es.spprc.rte cluster.msg.en_US.pprc</pre> <p>(and other appropriate language message sets)</p>
SAN Volume Controller (SVC)	<pre>cluster.es.svcpprc.cmds cluster.es.svcpprc.rte cluster.msg.en_US.svcpprc</pre> <p>(and other appropriate language message sets)</p>

Since each type of PPRC management has different prerequisites, information on installing the particular filesets for specific support types (for example, cluster.es.pprc filesets) is deferred to the section specific to that PPRC management type.

Note: For the latest information about PowerHA® SystemMirror® Enterprise Edition software, see the Release Notes® in the /usr/es/sbin/cluster/release_notes_xd file.

Contents of the installation media

The PowerHA® SystemMirror® Enterprise Edition for Metro Mirror installation media provides the images for installation on each node in the cluster that can take over a PPRC mirrored volume group.

These images include:

Filesets for base install

- cluster.es.pprc
- cluster.es.spprc
- cluster.es.svcpprc
- cluster.es.genxd

Direct management and DSCLI message sets

- cluster.msg.ja_JP.pprc
- cluster.msg.en_US.pprc
- cluster.msg.Ja_JP.pprc
- cluster.msg.En_US.pprc

SVC PPRC message sets

- cluster.msg.ja_JP.svcpprc
- cluster.msg.en_US.svcpprc
- cluster.msg.Ja_JP.svcpprc
- cluster.msg.En_US.svcpprc

User Documentation for all Management Types

PowerHA® SystemMirror® Enterprise Edition for Metro Mirror installation choices

Install the PowerHA® SystemMirror® Enterprise Edition for Metro Mirror software on each cluster node (server).

Installing from an installation server

To install the PowerHA® SystemMirror® Enterprise Edition for Metro Mirror software in a cluster environment, you can create an PowerHA® SystemMirror® Enterprise Edition for Metro Mirror installation server (containing the PowerHA® SystemMirror® Enterprise Edition for Metro Mirror software images for installation) on one node and then load the images onto the remaining cluster nodes. This is the fastest way to install PowerHA® SystemMirror® Enterprise Edition for Metro Mirror.

PowerHA® SystemMirror® Enterprise Edition for Metro Mirror supports the Network Installation Management program and Alternate Disk Migration.

Installing from a hard disk

To install PowerHA® SystemMirror® Enterprise Edition for Metro Mirror software from your hard disk, you copy the software from the installation media to the hard disk prior to installation.

To copy the PowerHA® SystemMirror® Enterprise Edition for Metro Mirror software to your hard disk:

1. Place the PowerHA® SystemMirror® Enterprise Edition for Metro Mirror CD into the CD-ROM drive.
2. Enter `smit bffcreate`
The **Copy Software to Hard Disk for Future Installation** panel appears.
3. Enter the name of the CD-ROM drive in the **INPUT device / directory for software** field and press Enter. If you are unsure of the input device name, press F4 to list available devices. Select the correct drive and press Enter. That drive name appears in the **INPUT device / directory** field as the valid input device.
4. Press Enter to display the **Copy Software to Hard Disk for Future Installation** panel.
5. Enter field values as follows:

<i>Copy Software to Hard Disk for Future Installation fields</i>	
Field	Value
SOFTWARE name	Press F4 for a software listing. Install the images for PowerHA® SystemMirror® Enterprise Edition for Metro Mirror. For a list of the PowerHA® SystemMirror® Enterprise Edition for Metro Mirror images, see the section Contents of the installation media.
DIRECTORY for storing software	Change the value to the storage directory that all nodes using PowerHA® SystemMirror® Enterprise Edition for Metro Mirror can access.

6. Enter values for the other fields as appropriate for your site.
7. When you are satisfied with the entries, press Enter.
SMIT responds: **Are you sure?**
8. Press Enter again to copy the software.

Result

Installing PowerHA® SystemMirror® Enterprise Edition for Metro Mirror from the Hard Disk

After the PowerHA® SystemMirror® Enterprise Edition for Metro Mirror software is on your system, follow the instructions in the section Installing the software from the installation media to install the software.

Installing the software from the installation media

If you install the PowerHA® SystemMirror® Enterprise Edition for Metro Mirror software from the CD-ROM, install the software directly onto each cluster node.

To install the PowerHA® SystemMirror® Enterprise Edition for Metro Mirror software on a server node:

1. If you are installing directly from the installation media, insert the CD into the CD-ROM drive.
2. Enter `smit install_all`
SMIT displays the first **Install and Update from ALL Available Software** panel.
3. Enter the device name of the installation media or install directory in the **INPUT device / directory for software** field and press Enter.
If you are unsure about the input device name or about the install directory, press F4 to list the available devices. Then select the correct device or directory and press Enter. The correct value is entered into the **INPUT device / directory** field as the valid input device.
4. Enter field values as follows. Press F1 for help on any field.

You should use F4 to list the software before proceeding with the installation. This way you can install either the English, Japanese, or French message catalogs.

<i>INPUT device / directory fields</i>	
Field	Value
INPUT device / directory for software	This field shows the device or directory you specified earlier.
SOFTWARE to install	Press F4 for a software listing. In the software list, use the arrow keys to locate all software filesets associated with an image. Next press F7 to select either an image or a fileset. Then press Enter after making all selections. Your selections appear in this field.
PREVIEW only?	If set to yes , the preview option checks and verifies that installation prerequisites are met, for instance that required software is installed and sufficient disk space is available. Press F1 for details. When you are ready to perform the actual installation, set this field to no .
COMMIT software updates?	This field applies only when you install software updates (PTFs). See the F1 help for details.
SAVE replaced files?	This field applies only when you install software updates (PTFs). If you select no in response to commit software updates? select yes for this field. See F1 help for details.
AUTOMATICALLY install requisite software	Set this field to no if the prerequisite software for the latest version is already installed or if the OVERWRITE same or newer versions? field is set to yes ; otherwise, set this field to yes to install required software. See F1 help for details.
EXTEND filesystems if space needed?	Select yes if you have adequate hard disk space, no if you have limited space. See F1 help for details.
OVERWRITE same or newer versions?	For new installations, leave this field set to no . Set it to yes if you are reinstalling the software. If you set this field to yes , set the Automatically install requisite software field to no . See the F1 help for details.
VERIFY install and check file sizes?	Select yes if you want the system to perform some checks on the software you installed. See F1 help for details.
DETAILED output?	Select yes if you want a detailed log of all installation messages.
Process multiple volumes?	Select this option if you want to enable the processing of multiple-volume CDs. See F1 for information.

Field	Value
ACCEPT new license agreements?	Select yes for this item to proceed with installation. If you choose no , installation may stop with a warning that one or more filesets require software license agreements. You accept the license agreement only once for each node.
Preview new license agreements?	Select yes to view the text of the license agreements. The text appears in the current window in the language defined on your system.

5. When you are satisfied with the entries, press Enter. SMIT responds **Are you sure?**
6. Press Enter to install the software.
7. After you install the software on a node, reboot that node.

Result

Read the PowerHA® SystemMirror® Enterprise Edition for Metro Mirror Release Notes® in the /usr/es/sbin/cluster/release_notes_xd file for information that does not appear in the product documentation.

After you successfully install the software on each cluster node, you are ready to configure PowerHA® SystemMirror® Enterprise Edition for Metro Mirror for one of the following management types:

- Direct management
- DSCLI
- SVC PPRC

Upgrading to latest release of PowerHA® SystemMirror® Enterprise Edition

These topics discuss how to upgrade to the latest release of the PowerHA® SystemMirror® Enterprise Edition for Metro Mirror software.

Upgrading the PowerHA® SystemMirror® Enterprise Edition software

Before you upgrade to a new release of PowerHA® SystemMirror® Enterprise Edition for Metro Mirror, make sure you are familiar with the process for installing and configuring PowerHA® SystemMirror® Enterprise Edition.

See the section Installing the software from the installation media.

- Make sure that your system meets the installation prerequisites. See the section Installation prerequisites for PowerHA® SystemMirror® Enterprise Edition Metro Mirror.
- Ensure that each cluster node has its own PowerHA® SystemMirror® Enterprise Edition license.
- The PowerHA® SystemMirror® Enterprise Edition software uses 1 MB of disk space.
- Perform the installation process as the root user.

Before upgrading your cluster:

1. Archive any localized script and configuration files to prevent losing them during an upgrade.
2. Commit the installation (if it is applied but *not* committed) so that the PowerHA® SystemMirror® software can be installed over the existing version. To see if your configuration is already committed, enter: `ls1pp -h cluster.*`
3. If the word COMMIT is displayed under the Action header, continue to the next step. If not, run the **smit install_commit** utility before installing the latest version software. SMIT displays the **Commit Applied Software Updates (Remove Saved Files)** panel.
4. Enter field values as follows:

<i>Commit Applied Software Updates (Remove Saved Files) fields</i>	
Field	Value
SOFTWARE name	Press F4 for a software listing, and select all of the cluster.* filesets.
COMMIT old version if above version used it?	Set this field to yes .

Field	Value
EXTEND filesystem if space needed?	Set this field to yes .

5. Make a **mksysb** backup of each node's configuration.
If you restore a **mksysb** backup onto your system, you need to reset the SCSI IDs on the system.
6. Save any customized event information.

Result

Note: When upgrading do *not* leave the cluster at mixed versions for long periods. Mixed versions of the software in the cluster can impact availability within the cluster.

For information about installing the latest software, see the section Installing the software from the installation media. The PPRC filesets must be upgraded at the same time as the other cluster filesets.

Verifying the upgraded cluster definition

After the PowerHA® SystemMirror® Enterprise Edition for Metro Mirror software is installed on all nodes, and all nodes have been rebooted, you should verify the configuration. Verification provides errors or warnings to ensure that the cluster definition is the same on all nodes.

To verify the cluster:

1. From the command line, enter `smit sysmirror`.
2. From the SMIT interface, select **Cluster applications and resources > Verify and synchronise cluster configurations**, and press Enter.

Result

Verification verifies both PowerHA® SystemMirror® and PowerHA® SystemMirror® Enterprise Edition for Metro Mirror configuration.

Note: You cannot synchronize a mixed-version cluster. New functionality is available only when all nodes have been upgraded and the cluster has been synchronized. Do not expect commands like **clfindres** to supply correct information in a mixed cluster.

Recovering from a failed installation

When you install PowerHA® SystemMirror® Enterprise Edition, the **cl_convert** command runs automatically to convert the PowerHA® SystemMirror® configuration database from a previous release to that of the current release. If the installation fails, run **cl_convert** from the command line to convert the database.

In a failed conversion, run **cl_convert** using the **-F** flag.

To run a conversion utility requires:

- Root user privileges
- PowerHA® SystemMirror® Enterprise Edition version from which you are converting

The **cl_convert** utility logs conversion progress to the `/tmp/clconvert.log` file so that you can gauge conversion success. This log file is regenerated each time **cl_convert** or **clconvert_snapshot** runs.

Modifying previous cluster snapshots

After you upgrade your PowerHA® SystemMirror® Enterprise Edition software to the latest version, you may want to restore one or more of the previous version cluster snapshots you created using the Cluster Snapshot utility.

The default directory path for storage and retrieval of a snapshot is `/usr/es/sbin/cluster/snapshots`; however, you may have specified an alternate path using the `SNAPSHOTPATH` environment variable. Check in these locations before using the `/usr/es/sbin/cluster/conversion/clconvert_snapshot` utility to convert the snapshot.

The snapshot is based on your full PowerHA® SystemMirror® configuration - including the configuration to include PPRC replicated resources in a cluster. The **clconvert_snapshot** utility updates PowerHA® SystemMirror® configuration data with new information for the latest version.

To convert and apply a cluster snapshot, enter:

```
clconvert_snapshot -v version# -s snapshot_file_name
```

Where the **-s** flag is used with the snapshot file name you want to update or apply, and the **-v** flag is used with the version of the saved snapshot.

Addressing problems during the installation

If you experience problems during the installation, the installation program usually performs a cleanup process automatically.

If, for some reason, the cleanup is *not* performed after an unsuccessful installation:

1. Enter `smit install` to display the Installation and Maintenance menu.
2. Select **Install and Update Software**.
3. Select **Clean Up After a Interrupted Installation**.
4. Review the SMIT output (or examine the `/smit.log` file) for the cause of the interruption.
5. Fix any problems and repeat the installation process.

PowerHA® SystemMirror® Enterprise Edition for Metro Mirror with Direct management

These topics describe the planning, installation, and configuration tasks for setting up PowerHA® SystemMirror® Enterprise Edition for Metro Mirror with Direct management. Using this method, PowerHA® SystemMirror® directly manages PPRC pairs by communicating with the ESS Copy Services Server.

PowerHA® SystemMirror® Enterprise Edition for Metro Mirror with Direct management is the oldest type of PPRC support PowerHA® SystemMirror® provides. It has the simplest hardware configuration, but requires more work on the administrator's part to set up.

Briefly, PowerHA® SystemMirror® provides support for direct management of ESS PPRC by managing specified pair and path tasks as defined by the user on their ESS storage systems Copy Services Server (CSS). PowerHA® SystemMirror® provides monitoring, failover and fallback support for PPRC by issuing commands directly to the CSS via the ESS CLI interface.

Direct management can be used to support PPRC between ESS 800 storage systems.

Planning for Direct management

You should be familiar with the planning tasks for PowerHA® SystemMirror®.

You should have completed the planning steps in Metro Mirror general planning for PowerHA® SystemMirror® sites, at least.

To continue the plan for an PowerHA® SystemMirror® Enterprise Edition for Metro Mirror with Direct management environment:

- Plan for connections to the ESS
- Plan for ESS Copy Services
- Plan the PowerHA® SystemMirror® Enterprise Edition for Metro Mirror configuration:
 - Identify ESS systems to be included
 - Identify which resource groups will contain the PPRC replicated resources
- Identify PPRC replicated resources and provide information about the associated volumes
- (*Optional*) Plan for any user-specific PPRC tasks.

Planning prerequisites for Direct management

PowerHA® SystemMirror® Enterprise Edition PPRC with Direct management manages PPRC resources by communicating with the Copy Services Server (CSS) on ESS systems via the ESS CLI.

Accordingly, prior to configuring Direct management, ensure the following:

- The version of the ESS CLI shipped with the microcode for your storage system has been installed on all PowerHA® SystemMirror® cluster nodes.
- You have access to the ESS Copy Services Web Interface for the storage systems involved in your configuration.

You can only use enhanced concurrent volume groups with PPRC.

The ESS Copy Services Web Interface provides a configuration interface for setting up Direct management PPRC.


Sample configuration for Direct management

You can set up a mutual recovery configuration in which each site acts as a production site with the other site acting as an associated backup site.

Implementing a mutual recovery configuration requires:

- Two PowerHA® SystemMirror® sites (the same as for a single recovery configuration)
- Two resource groups.

The following figure shows a mutual recovery configuration with two resource groups in which Site_2 is the production site with Site_1 as the recovery site for Resource Group 2, and Site_1 is the production site with Site_2 as the recovery site for Resource Group 1. In Resource Group 1 the nodes n1 and n2 have a higher priority, and in Resource Group 2 the nodes n3 and n4 have a higher priority. The volume groups in both resource groups are included in PPRC replicated resources.



The order of the nodes in each resource group's nodelist indicates which site is considered to be the production site and which site is considered to be the backup site for that resource group. The site that includes the nodes with the highest priority is considered the production site for the resource group.

In the preceding example, the configuration for resource group 1 would be:

<i>Sample configuration for resource group 1</i>	
Field	Value
Ordered sitelist	Site_1, Site_2
Primary volume	volume100@ESS_1 (mirrored volume on ESS_1)
Secondary volume	volume200@ESS_2 (mirrored volume on ESS_2)
Ordered nodelist	n1, n2, n3, n4

In the preceding example, the configuration for resource group 2 would be:

Sample configuration for resource group 2	
Field	Value
Ordered sitelist	Site_2, Site_1
Primary volume	volume210@ESS_2 (mirrored volume on ESS_2)
Secondary volume	volume110@ESS_1 (mirrored volume on ESS_1)
Ordered nodelist	n3, n4, n1, n2

Planning integration with PPRC

Before you configure PowerHA® SystemMirror® Enterprise Edition for Metro Mirror, you configure PPRC on the ESS.

To configure PPRC, you set up:

- *PPRC paths*, ESCON® links that transfer mirrored data from one ESS to another
- *PPRC pairs*, volumes mirrored from one ESS to another.

The PPRC paths connect PPRC pairs.

Planning connections to the ESS

PowerHA® SystemMirror® Enterprise Edition for Metro Mirror communication with the ESS requires connections to each of the ESS subsystems through the ESSNet administrative network. This access lets PowerHA® SystemMirror® Enterprise Edition for Metro Mirror use the ESS Command Line Interface (CLI) to control the ESS subsystems during event processing.

Commands are sent to the ESS Copy Services Server software that runs on one of the ESS subsystems in the environment.

The PowerHA® SystemMirror® cluster nodes can either connect directly to the ESSNet or connect over a router or bridge. The interfaces that communicate with the ESSNet can be defined to PowerHA® SystemMirror®, but IP address takeover (IPAT) should *not* be used on those interfaces. If the interfaces are defined to PowerHA® SystemMirror®, *ensure that the nodes can gain access to ESSNet* at any time.

Note: The reliability of the connection between a node and the ESSNet directly affects the performance of PowerHA® SystemMirror® Enterprise Edition for Metro Mirror. A slow or unreliable network connection to ESSNet results in commands to initiate a failover being processed slowly or unreliably.

Planning Copy Services Server on ESS

Considerations for using Copy Service Server differ depending on the PPRC version.

PPRC Versions before 2.2

When you set up ESS Copy Services, you define one ESS cluster processor complex as the active Copy Services Server and one ESS cluster processor complex on another ESS as the backup Copy Services Server.

If the active Copy Services Server fails, a notification message is sent. Because PPRC versions prior to 2.2 do not support dual active Copy Services Servers, you manually reset the Copy Services Server to make the backup server assume the active role. *PowerHA® SystemMirror® cannot restart the Copy Services Server.*

Set up the active Copy Services Server at the recovery site.

PPRC Version 2.2 or later

When you set up ESS Copy Services, you define one ESS cluster processor complex in each ESS as a Copy Services Server. Because both Copy Services Servers are active at all times, you do not need to manually restart a Copy Service Server if one fails.

Configuring PPRC paths

When configuring PPRC paths that connect PPRC pairs to be included in a PPRC replicated resource, keep in mind that each ESCON® link can carry up to 64 paths at the same time. It also carries data in one direction at a time.

PowerHA® SystemMirror® Enterprise Edition for Metro Mirror requires that ESCON® links be available to carry data in both directions at the same time; therefore, you need a minimum of two physical links. To improve

throughput and to provide redundancy for the ESCON® cables and adapters, you should have at least four links. Note that the ESS manages the paths.

Planning for PPRC replicated resources

PPRC pairs are defined so that one volume in the pair resides on the ESS at one site, and the other volume in the pair resides on the ESS at the other site. In a PPRC pair, the volume that data is being written to is the source volume, and the volume that contains a mirrored copy of the data is the *target* volume.

The definition for a PPRC replicated resource contains the volume identifier and the name of the ESS for the source and the target volumes. PowerHA® SystemMirror® has knowledge of which volumes are mirrors of each other for each PPRC replicated resource.

To plan PPRC replicated resources for PowerHA® SystemMirror® Enterprise Edition for Metro Mirror, you should have a good understanding of the volume groups and volumes involved.

The terms primary and secondary are specific to each pair of PPRC-mirrored volumes as determined by the pair definition. In configurations that designate one site as a production site and the other as a backup site, the production site holds the primary volumes for the PPRC pairs, and the backup site holds the secondary volumes. Each PPRC pair is included in a resource group whose home node is located at the site that is primary for the PPRC pair.

In a mutual recovery configuration, in which nodes from both sites are active, each site contains primary volumes for some PPRC pairs and secondary volumes for others.

Note: PPRC copies the volume information, including the PVID, from one volume in a PPRC pair to the other. The volumes at both sites contain the same logical volumes and must therefore be imported with the same volume group name. This also allows single-name entries in a resource group definition.

Use ESS Copy Services to obtain information about the ESS configuration and the PPRC pairs that have been established.

Identifying volume groups

You use the volume group name when configuring PPRC tasks in ESS Copy Services and when you identify the volume group to an PowerHA® SystemMirror® node.

The volume group and logical volume names must be the same on all nodes in each site where the volume group can be brought online.

To identify volume group information on the ESS, complete the following steps:

1. Run the **lspv** command to view which volume group is associated with which hdisk.
2. Run the **rsList2105.sh** command to see which hdisks are associated with which serial numbers. Or, if you are using vpaths (as provided by an SDD driver) run the **lsvpcfg** command to see the information.

Planning tasks for the ESS

Tasks are a set of actions to be executed on an ESS. They allow you to automate a series of steps that otherwise would be performed through the ESS Copy Services Web Interface.

In a direct management environment, PowerHA® SystemMirror® Enterprise Edition for Metro Mirror relies on PPRC tasks for managing PPRC volume pairs and the associated PPRC paths. These tasks are used to establish and terminate PPRC pairs during failover and fallback of volumes. To access these tasks, nodes in an PowerHA® SystemMirror® cluster that supports PowerHA® SystemMirror® Enterprise Edition for Metro Mirror require access to ESSNet to execute commands.

Using user-specific task names

User-specific task names can be up to 16 characters long. If possible, base the task name on the naming conventions described in the Understanding task names topic.

Unless you have a configuration conflict (for example if you already have a task named with one of the recommended task names), you should name these tasks as listed in the section Configuring PPRC tasks.

Note: When you specify a user-specific task name for any task, you need to specify task names for all of the tasks (whether or not they are different from the recommended names) for each volume group.

You configure these tasks in ESS Copy Services and define them to PowerHA® SystemMirror® Enterprise Edition for Metro Mirror. For information about defining a group of user-specific task names to PowerHA® SystemMirror®, see the section Defining PPRC Tasks to PowerHA® SystemMirror® Enterprise Edition for Metro Mirror for Direct management.

Note: You should use the recommended task names, unless there is a very strong reason to do otherwise. If you do change names, ensure that the names are entered correctly in ESS Copy Services and in PowerHA® SystemMirror® Enterprise Edition.

Understanding task names

The recommended names for the PPRC tasks specify the volume group and PPRC tasks to manipulate the volume pairs in the volume group.

The recommended task names use the following naming conventions:

volume_group_action_object_direction_modifier_system

A task name may contain some or all of the components listed in the preceding syntax line.

The following table shows the values for the parts of the task name:

A table that displays the PPRC task names

<i>Task names</i>		
Task name	Short value	Full value
<i>volume_group</i>	The name of the volume group (less than 7 characters long)	
<i>_action</i>	The action to be executed:	
	Est Sus Ter Del	Establish Suspend Terminate Delete
<i>_object</i>	The pair or path for the volume group:	
	Pr Pt	Pair Path
<i>_direction</i>	The direction of the operation:	
	PS SP™	Primary to secondary Secondary to primary
<i>_modifier</i>	The type of copy:	
	NC FC SC FO FB F	No copy Full copy Sync changes Failover Failback Forced (only for the delete action)
<i>_system</i>	The type of system on which the action is being executed:	
	P S	Primary Secondary

Planning a PPRC resource group

If you have not done so, plan your resource groups.

Installing PowerHA® SystemMirror® Enterprise Edition for Metro Mirror for Direct management

These topics describe how to install PowerHA® SystemMirror® Enterprise Edition for Metro Mirror for Direct management.

Details for installing base PowerHA® SystemMirror® Enterprise Edition from the installation media are included in Installing PowerHA® SystemMirror® Enterprise Edition for Metro Mirror.

Note: At run time, make sure you have at least 9 MB of free space in the PowerHA® SystemMirror® log directory. PPRC commands use this directory.

Installation prerequisites for Direct management

There are some installation prerequisites for Direct management.

Before you install PowerHA® SystemMirror® Enterprise Edition Metro Mirror for Direct management, make sure that the following is true for each cluster node:

- The PowerHA® SystemMirror® Enterprise Edition software uses 1 MB of disk space in the **/usr** filesystem.
- You have root access to each node.

Software requirements for Direct management

There are some software requirements for PowerHA® SystemMirror® Enterprise Edition Metro Mirror.

The following software is required:

- AIX® level as stated in the support flash
- Minimum PowerHA® SystemMirror® version plus all current APARs
- Java™ Runtime Environment version appropriate for the ESS CLI version installed. (This will be tied to and should be available with the microcode level on the ESS systems.)
- ESS CLI as appropriate for your storage microcode (LIC) level:
The ESS Copy Services CLI software is closely tied to the ESS microcode level. CLI code is found on the MegaCDR that is included in the Customer Software packet.
 - IBM® 2105 Command Line Interface (**ibm2105cli.rte**)
or
 - IBM® 2105 Command Line Interface (**ibm2105esscli.rte**) (also included on the MegaCDR).

Depending on which version of the **ibm2105esscli** fileset ships with your level of ESS microcode, the files will be installed in one of three locations:

- a. **/usr/opt/ibm2105cli**
- b. **/usr/opt/ibm/ibm2105cli**
- c. **/opt/ibm/ESScli**

PowerHA® SystemMirror® Enterprise Edition assumes a default directory for the ESS CLI executables. Therefore, once the CLI executables are installed in one of the directories listed, you must create the following link:

```
/usr/opt/ibm2105cli -> < ESS cli installation location >
```

so that PowerHA® SystemMirror® Enterprise Edition can find the executables.

- (Optional but recommended) ESS microcode level vrmf 2.4.x.x. This version provides support for a dual active Copy Services Server.

Note: Install microcode vrmf 2.4.x.x on the ESS system to make support for a dual active Copy Services Server available. Otherwise, the Copy Services Server must be manually started on a backup Copy Services Server if the primary Copy Services Server is unavailable.

Installing filesets for Direct management

You need to install the necessary filesets for Direct management.

If you have not already done so, install the filesets listed in Contents of the installation media for Direct management.

Addressing problems during the installation

If you experience problems during the installation, the installation program usually performs a cleanup process automatically. If, for some reason, the cleanup is *not* performed after an unsuccessful installation, there are some steps that you can take.

These steps include:

1. Enter `smit install` to display the Installation and Maintenance menu.
2. Select **Install and Update Software**.
3. Select **Clean Up After an Interrupted Installation**.
4. Review the SMIT output (or examine the `/smit.log` file) for the cause of the interruption.
5. Fix any problems and repeat the installation process.

Configuring in a Direct management environment

These topics describe how to configure the ESS system to support PowerHA® SystemMirror® Enterprise Edition for Metro Mirror, and how to configure an PowerHA® SystemMirror® Enterprise Edition direct management (ESS CLI) environment.

Configuration prerequisites for Direct management

Before configuring Direct management, there are some prerequisite steps you need to take.

Before configuring PowerHA® SystemMirror® Enterprise Edition for Metro Mirror, ensure that:

- PPRC is configured and running on the ESS systems.
PowerHA® SystemMirror® Enterprise Edition for Metro Mirror supports System p® servers that can support ESS systems.
- You have a good understanding of PowerHA® SystemMirror® sites for PPRC replicated resources. For more information about sites, see the section PowerHA® SystemMirror® sites.
- The PowerHA® SystemMirror® Enterprise Edition for Metro Mirror software is installed on each cluster node that will support PowerHA® SystemMirror® Enterprise Edition for Metro Mirror.
For information about installing PowerHA® SystemMirror® Enterprise Edition for Metro Mirror, see Installing PowerHA® SystemMirror® Enterprise Edition for Metro Mirror.
- The PowerHA® SystemMirror® cluster is configured for:
 - Nodes
 - Networks and network interfaces
 - Service labels
 - Initial resource groups.

You can modify the attributes for a resource group later to accommodate PPRC replicated resources.

Configuration overview

Once you have installed PowerHA® SystemMirror® Enterprise Edition for Metro Mirror, you need to set up configuration.

To set up an PowerHA® SystemMirror® Enterprise Edition for Metro Mirror configuration:

1. On the ESS, configure PPRC support for PowerHA® SystemMirror® Enterprise Edition for Metro Mirror:
 - a. Configure PPRC tasks in ESS Copy Services
 - b. Import the volume groups on the nodes at the secondary site.
2. On cluster nodes, configure support for PowerHA® SystemMirror® Enterprise Edition for Metro Mirror:
 - a. Define PowerHA® SystemMirror® sites for PowerHA® SystemMirror® Enterprise Edition for Metro Mirror.
 - b. Define ESS systems to be included.
 - c. Define PPRC replicated resources.
 - d. Configure resource groups to include PPRC replicated resources.
 - e. (*Optional*) If user-specific task names are used, define them to PowerHA® SystemMirror®.
 - f. (*Optional*) If your cluster has a large number of volume groups, create LUN ID mapping files before startup as described in Improving the performance of volume groups.

Configuring PowerHA® SystemMirror® Enterprise Edition for Metro Mirror for Direct management support

Direct management configuration relies on *tasks* configured through ESS Copy Services and activated on the ESS system.

For information about tasks, see the section Planning tasks for the ESS.

Configure the ESS system in preparation for PowerHA® SystemMirror® Enterprise Edition for Metro Mirror before making configuration changes in PowerHA® SystemMirror® Enterprise Edition for Metro Mirror.

Configuring PPRC tasks

You create 24 PPRC tasks for each PPRC-protected volume group in the cluster. This section lists each of these tasks and lists the options in ESS Remote Copy Services that you use to set up the task.

When you create tasks, use the ESS Copy Services Web Interface. In this interface, serial numbers identify volumes. You select disks by serial number, and then create tasks that start with the volume group name.

Note: Use the task names recommended, unless there is a very strong reason to do otherwise. If you do change names, ensure that these names are entered correctly in ESS Copy Services and in PowerHA® SystemMirror® Enterprise Edition for Metro Mirror.

If you name a task other than the recommended name, see the section PPRC Tasks to PowerHA® SystemMirror® Enterprise Edition for Metro Mirror for Direct management.

Note: Enter the task names carefully. Direct management does *not* work if any of the task names are inaccurate.

To configure tasks for a volume group to use PowerHA® SystemMirror® Enterprise Edition for Metro Mirror:

1. Log into ESS Copy Services as a user with administrative privileges.
2. Use the ESS Copy Services Web Interface to define the ESS tasks needed to manage the paths (ESCON® links) that PPRC uses for the volume group.
Create the tasks as shown in the following list. Use the task names shown here, substituting the name of the volume group for *\$vgname*. This name can be up to seven characters long, for example vg1.

Note: Task names are case-sensitive.

For more information about volume group names, see (Identifying Volume Groups).

Task Name	Action	Option in ESS Copy Services to Define Task
<i>\$vgname</i> EstPtPS	Establish path(s) from the Logical Subsystem (LSS) in the primary ESS to the LSS in the secondary ESS.	Do not establish paths if they already exist
<i>\$vgname</i> DelPtPSF	Remove paths from the LSS in the primary ESS to the LSS in the secondary ESS.	Force removal of PPRC path even if pairs exist
<i>\$vgname</i> DelPtPS	Remove paths from the LSS in the primary ESS to the LSS in the secondary ESS.	
<i>\$vgname</i> EstPtSP	Establish paths from the LSS in the secondary ESS to the LSS in the primary ESS.	Do not establish paths if they already exist
<i>\$vgname</i> DelPtSP	Remove paths from the LSS in the secondary ESS to the LSS in the primary ESS.	
<i>\$vgname</i> DelPtSPF	Remove paths from the LSS in the secondary ESS to the LSS in the primary ESS.	Force removal of PPRC path even if pairs exist

3. Use the ESS Copy Services Web Interface to define the tasks needed to manage the pairs of volumes that are PPRC-protected. If you define multiple volumes in the volume group, you can use the **Multiple Selection** option on the ESS Copy Services Web Interface to define a single set of tasks to manage the entire volume group.

Create the tasks as shown in the following list. *Use the task names shown here*, substituting the name of the volume group for *\$vgname*. For more information about volume group names, see Identifying volume groups.

Task Name	Action	Option in ESS Copy Services to Define Task
<i>\$vgname</i> EstPrPSFC	Establish a PPRC relationship from the source volume(s) in the primary ESS to the target volume(s) in the secondary ESS.	Copy Entire Volume Permit Read from Secondary
<i>\$vgname</i> EstPrPSNC	Establish a PPRC relationship from the source volume(s) in the primary ESS to the target volume(s) in the secondary ESS.	Do Not Copy Volume Permit Read from Secondary
<i>\$vgname</i> EstPrPSSC	Establish a PPRC relationship from the source volume(s) in the primary ESS to the target volume(s) in the secondary ESS.	Copy Out-of-Sync Cylinders Only Permit Read from Secondary
<i>\$vgname</i> EstPrPSFO	Establish a PPRC relationship from the target volume(s) in the primary ESS to the source volume(s) in the secondary ESS.	PPRC Failover á
<i>\$vgname</i> EstPrPSFB	Establish a PPRC relationship from the target volume(s) in the primary ESS to the source volume(s) in the secondary ESS.	PPRC Failback Permit Read from Secondary
<i>\$vgname</i> SusPrPSP	Suspend a PPRC relationship between the source volume in the primary ESS and the target volume in the secondary ESS.	

Task Name	Action	Option in ESS Copy Services to Define Task
	This task is executed on the primary ESS.	
<i>\$vname SusPrPSS</i>	Suspend a PPRC relationship between the source volume in the primary ESS and the target volume in the secondary ESS. This task is executed on the secondary ESS.	
<i>\$vname TerPrPSP</i>	Terminate a PPRC relationship from the source volume in the primary ESS to the target volume in the secondary ESS. This task is executed on the primary ESS.	
<i>\$vname TerPrPSS</i>	Terminate a PPRC relationship from the source volume in the primary ESS to the target volume in the secondary ESS. This task is executed on the secondary ESS.	
<i>\$vname EstPrSPFC</i>	Establish a PPRC relationship from the source volume(s) in the secondary ESS to the target volume(s) in the primary ESS.	Copy Entire Volume Permit Read from Secondary
<i>\$vname EstPrSPNC</i>	Establish a PPRC relationship from the source volume(s) in the secondary ESS to the target volume(s) in the primary ESS.	Do Not Copy Volume Permit Read from Secondary
<i>\$vname EstPrSPSC</i>	Establish a PPRC relationship from the source volume(s) in the secondary ESS to the target volume(s) in the primary ESS.	Copy Out-of-Sync Cylinders Only Permit Read from Secondary
<i>\$vname EstPrSPFO</i>	Establish a PPRC relationship from the target volume(s) in the secondary ESS to the source volume(s) in the primary ESS.	PPRC Failover
<i>\$vname EstPrSPFB</i>	Establish a PPRC relationship from the target volume(s) in the secondary ESS to the source volume(s) in the primary ESS.	PPRC Failback Permit Read from Secondary
<i>\$vname TerPrSPS</i>	Terminate a PPRC relationship from the source volume(s) in the secondary ESS to the target volume(s) in the primary ESS. This task is executed on the secondary ESS.	
<i>\$vname TerPrSPP</i>	Terminate a PPRC relationship from the source volume(s) in the secondary ESS to the target volume(s) in the primary ESS. This task is executed on the primary ESS.	
<i>\$vname SusPrSPS</i>	Suspend a PPRC relationship between the source volume in the	

Task Name	Action	Option in ESS Copy Services to Define Task
	secondary ESS and the target volume in the primary ESS. This task is executed on the secondary ESS.	
<i>\$vgname</i> SusPrSPP	Suspend a PPRC relationship between the source volume in the secondary ESS and the target volume in the primary ESS. This task is executed on the primary ESS.	

- Run the ESS **rsExecuteTask.sh** command to verify that you can execute the tasks created in the previous step. Note that the command name is case-sensitive. The **rsExecuteTask.sh** command should return to a command prompt. If it does *not*, PowerHA® SystemMirror® Enterprise Edition for Metro Mirror does *not* work.

If the command does *not* return to a command prompt:

- Review task configuration and make corrections as needed.
- Run the **rsExecuteTask.sh** command again.

A host server IP address *not* being defined to the ESS causes a known configuration problem with the ESS. If this is present, the **rsExecuteTask.sh** command executes the ESS task, but the program does *not* return.

- Repeat steps 1 - 4 for each volume group that will use PowerHA® SystemMirror® Enterprise Edition for Metro Mirror.

Importing mirrored volume groups at the secondary site

If you have already created volume groups and are sure that volumes are mirroring correctly, skip this section. If you have not created the volume groups, complete the procedure in this section.

The secondary volume in a PPRC pair is not visible to the nodes at the secondary site when the PPRC relationship is active. As a result, use the following procedure to define the disks and volume groups for PowerHA® SystemMirror®.

To define disks and volume groups to PowerHA® SystemMirror®:

- Log into ESS Copy Services as a user with administrative privileges.
- Run the *\$vgname* **EstPtPS** task, where *\$vgname* is the name of the volume group. This establishes the path(s) from the primary volume to the secondary volume.
- Run the *\$vgname* **EstPrPSFC** task, where *\$vgname* is the name of the volume group. This establishes the PPRC pair(s) from the primary volume to the secondary volume and synchronizes the two copies. This may take awhile.
- Create the volume group using the AIX® Logical Volume Manager (LVM) on nodes on the active site.
- Run the *\$vgname* **TerPrPSP** task, where *\$vgname* is the name of the volume group. This terminates the PPRC pair(s).
- Run the *\$vgname* **DelPtPS** task, where *\$vgname* is the name of the volume group. This terminates the path(s).
- Run Configuration Manager (**cfgmgr**) on the backup host server to move ESS hdisks to the available state.
- Import the volume group to all nodes on the secondary site.

Result

The volume group can now be defined to PowerHA® SystemMirror® Enterprise Edition for Metro Mirror.

Configuring the PowerHA® SystemMirror® Enterprise Edition Metro Mirror Cluster for Direct management

Configure PowerHA® SystemMirror® Enterprise Edition for Metro Mirror after you configure tasks for PPRC on the ESS systems and define sites to support PPRC replicated resources in PowerHA® SystemMirror®.

Defining PPRC replicated resources

Most of the configuration options for PowerHA® SystemMirror® Enterprise Edition for Metro Mirror are available from the **Define PPRC Replicated Resources** panel.

To define PPRC replicated resources:

1. From the command line, enter `smit sysmirror`.
2. In SMIT, select **Cluster Applications and Resources > Resources > PPRC-Managed Replicated Resources Configuration > ESSCLI-Managed PPRC Replicated Resource Configuration**, and press Enter. From this panel you can:
 - Define ESS disk subsystems.
 - Define a PPRC replicated resource.
 - Define PPRC tasks.
 - Synchronize the PPRC configuration.
 - Verify the PPRC configuration.

Defining the ESS Disk Subsystems to PowerHA® SystemMirror®

You define the ESS subsystems included in the sites that support PowerHA® SystemMirror® Enterprise Edition for Metro Mirror to enable PowerHA® SystemMirror® to process failovers for PPRC replicated resources.

To define an ESS system to PowerHA® SystemMirror®:

1. From the command line, enter `smit sysmirror`.
2. In SMIT, select **Cluster Applications and Resources > Resources > PPRC-Managed Replicated Resources Configuration > ESSCLI-Managed PPRC Replicated Resource Configuration > ESS Disk Subsystem Configuration > Add an ESS Disk Subsystem**, and press Enter.
3. In the **Add an ESS Disk Subsystem** panel, enter field values as follows:

<i>Add an ESS Disk Subsystem fields</i>	
Field	Value
ESS Subsystem Name	The name that identifies the ESS subsystem to PowerHA® SystemMirror®. The name may consist of up to 32 alphanumeric characters, and may include underscores.
ESS site name	The name of the site where the ESS resides. The site is already defined in PowerHA® SystemMirror® for the site name to be available from the picklist.
ESS IP Address	The IP address, in dotted decimal notation, that the Copy Services Server for an ESS uses. <div style="border: 1px solid blue; padding: 5px; margin: 10px 0;"> The IP address specified here is the ESSNet address of the ESS at this site. </div>
ESS User ID	The user ID used to authenticate logging into the ESS.
ESS password	The password associated with the specified ESS User ID .

4. Press Enter.
5. Review the settings for the ESS subsystem. From the **Define ESS Disk Subsystem** panel, select **Change/Show an ESS Disk Subsystem** and select an ESS system.
6. Make changes to the configuration settings as needed.

Defining the PPRC pairs to PowerHA® SystemMirror®

You define the PPRC pairs, the primary volume to secondary volume mappings, to allow the PowerHA® SystemMirror® Enterprise Edition for Metro Mirror to manage them.

You must define one PPRC pair for each mirrored volume in the volume group.

To define a PPRC Pair to PowerHA® SystemMirror® Enterprise Edition for Metro Mirror:

1. From the command line, enter `smit sysmirror`.
2. In SMIT, select **Cluster Applications and Resources > Resources > PPRC-Managed Replicated Resources Configurations**, and press Enter.
3. In the **Add a PPRC Resource** panel, enter field values as follows:

<i>Add a PPRC Resource fields</i>	
Field	Value
PPRC Resource Name	The name of the PPRC volume pair. The name may consist of up to 32 alphanumeric characters, and may include underscores. For example, a PPRC replicated resource could be named <code>pprc4.1</code> , with <code>pprc</code> signifying the type of replicated resource, <code>4</code> identifying the volume group to which the PPRC-protected volume belongs, and <code>1</code> enumerating a PPRC-protected volume in the volume group. Another PPRC replicated resource that includes a different volume in the same volume group could be named <code>pprc4.2</code> .
PowerHA® SystemMirror® Sites	The two PowerHA® SystemMirror® sites, with the primary site for this PPRC pair listed first. The sites are already defined in PowerHA® SystemMirror®.
Primary Site ESS Info	The identifier for the volume in the primary ESS in the form: <i>volume_id @ ESS_name</i> where <i>volume_id</i> is the ESS volume ID of the disk, and <i>ess_name</i> is the name of the ESS subsystem containing this volume. The primary ESS is the ESS at the site that you chose as primary for this PPRC replicated resource.
Secondary Site ESS Info	The identifier for a volume in the secondary ESS in the form: <i>volume_id @ ESS_name</i> where <i>volume_id</i> is the ESS volume ID of the disk, and <i>ess_name</i> is the name of the ESS subsystem containing this volume. The secondary ESS is the ESS at the site that you chose as secondary for this PPRC replicated resource.

4. Press Enter.
5. Review the settings for the PPRC pairs. From the **Define a PPRC Replicated Resource** panel, select **Change/Show a PPRC Resource** and select a PPRC pair.
6. Make changes to the configuration settings as needed.

Defining the PPRC tasks to PowerHA® SystemMirror® Enterprise Edition for Metro Mirror for Direct management

You define PPRC tasks to PowerHA® SystemMirror® Enterprise Edition for Metro Mirror *only in cases where you named the PPRC tasks differently* than recommended. If you used the recommended naming convention, you can skip this section.

See [Configuring PPRC tasks](#).

Defining PPRC path tasks

If you specified another name for a PPRC path task in ESS Remote Copy Services, identify that task to PowerHA® SystemMirror® Enterprise Edition for Metro Mirror.

To define a user-specific PPRC path task name to PowerHA® SystemMirror® Enterprise Edition for Metro Mirror:

1. From the command line, enter `smit sysmirror`.

- In SMIT, select **Cluster Applications and Resources > Resources > PPRC-Managed Replicated Resources Configuration > ESSCLI-Managed PPRC Replicated Resource Configuration > Define PPRC Tasks > Define PPRC Path Tasks > Add a Group of PPRC Path Tasks**, and press Enter.
- In the **Add a Group of PPRC Path Tasks** panel, enter field values as follows:

<i>Add a Group of PPRC Path Tasks fields</i>	
Field	Value
Volume Group Name	The name of the volume group associated with the PPRC path task.
Establish Path Pri - Sec	The name of the task that establishes a PPRC path from the primary ESS to the secondary ESS as defined in ESS Copy Services.
Delete Path Pri - Sec	The name of the task that deletes a PPRC path from the primary ESS to the secondary ESS as defined in ESS Copy Services.
Delete Path Pri - Sec FORCED	The name of the task that deletes a PPRC path from the primary ESS to the secondary ESS with the Forced option as defined in ESS Copy Services.
Establish Path Sec - Pri	The name of the task that establishes a PPRC path from the secondary ESS to the primary ESS as defined in ESS Copy Services.
Delete Path Sec - Pri	The name of the task that deletes a PPRC path from the secondary ESS to the primary ESS as defined in ESS Copy Services.
Delete Path Sec - Pri FORCED	The name of the task that deletes a PPRC path from the secondary ESS to the primary ESS with the Forced option as defined in ESS Copy Services.

- Press Enter.
- Review the settings for the PPRC path tasks, from the **Define PPRC Path Tasks** panel, select **Change/ Show a Group of PPRC Path Tasks**.
- Make changes to the configuration settings as needed.

Defining PPRC pair tasks

If you specified another name for a PPRC pair task in ESS Remote Copy Services, identify that task to PowerHA® SystemMirror® Enterprise Edition for Metro Mirror.

Note: When you specify a user-specific task name, specify task names for all eighteen tasks (whether or not they are different from the recommended names) for each volume group.

To define a group of user-specific PPRC pair task names to PowerHA® SystemMirror® Enterprise Edition for Metro Mirror:

- From the command line, enter `smit sysmirror`.
- In SMIT, select **Cluster Applications and Resources > Resources > PPRC-Managed Replicated Resources Configuration > ESSCLI-Managed PPRC Replicated Resource Configuration > Define PPRC Tasks > Define PPRC Pair Tasks > Add a Group of PPRC Pair Tasks**, and press Enter.
- In the **Add a Group of PPRC Pair Tasks** panel, enter field values as follows:
A table to help you enter information about your environment.

<i>Add a Group of PPRC Pair Tasks fields</i>	
Field	Value
Volume Group Name	The name of the volume group associated with the PPRC pair task.
Direction of the operation: Primary to Secondary	
Establish Pair Pri - Sec NO COPY	The name of the task that establishes a PPRC pair from the primary ESS to the secondary ESS with the No Copy option set as defined in ESS Copy Services.

Field	Value
Establish Pair Pri - Sec FULL COPY	The name of the task that establishes a PPRC pair from the primary ESS to the secondary ESS with the Full Copy option set as defined in ESS Copy Services.
Establish Pair Pri - Sec RESYNC	The name of the task that establishes a PPRC pair from the primary ESS to the secondary ESS with the Copy Out-of-sync Cylinders Only option set as defined in ESS Copy Services.
Establish Pair Pri - Sec FAILOVER	The name of the task that establishes a PPRC pair from the primary ESS to the secondary ESS with the Failover option set as defined in ESS Copy Services.
Establish Pair Pri - Sec FAILBACK	The name of the task that establishes a PPRC pair from the primary ESS to the secondary ESS with the Failback option set as defined in ESS Copy Services.
Suspend Pair Pri - Sec on Pri	The name of the task executed on the primary ESS that suspends PPRC mirroring from the primary ESS to the secondary ESS as defined in ESS Copy Services.
Suspend Pair Pri - Sec on Sec	The name of the task executed on the secondary ESS that suspends PPRC mirroring from the primary ESS to the secondary ESS as defined in ESS Copy Services.
Terminate Pair Pri - Sec on Pri	The name of the task executed on the primary ESS that ends PPRC mirroring from primary ESS to secondary ESS as defined in ESS Copy Services.
Terminate Pair Pri - Sec on Sec	The name of the task executed on the secondary ESS that ends PPRC mirroring from primary ESS to secondary ESS as defined in ESS Copy Services.
Direction of the operation: Secondary to Primary	
Establish Pair Sec - Pri NO COPY	The name of the task that establishes a PPRC pair from the secondary ESS to the primary ESS with the No Copy option set as defined in ESS Copy Services.
Establish Pair Sec - Pri FULL COPY	The name of the task that establishes a PPRC pair from the secondary ESS to the primary ESS with the Full Copy option set as defined in ESS Copy Services.
Establish Pair Sec - Pri RESYNC	The name of the task that establishes a PPRC pair from the secondary ESS to the primary ESS with the Copy Out-of-sync Cylinders Only option set as defined in ESS Copy Services.
Establish Pair Sec - Pri FAILOVER	The name of the task that establishes a PPRC pair from the secondary ESS to the primary ESS with the Failover option set as defined in ESS Copy Services.
Establish Pair Sec - Pri FAILBACK	The name of the task that establishes a PPRC pair from the secondary ESS to the primary ESS with the Failback option set as defined in ESS Copy Services.
Suspend Pair Sec - Pri on Sec	The name of the task executed on the secondary ESS that suspends PPRC mirroring from the secondary ESS to the primary ESS as defined in ESS Copy Services.
Suspend Pair Sec - Pri on Pri	The name of the task executed on the primary ESS that suspends PPRC mirroring from the secondary ESS to the primary ESS as defined in ESS Copy Services.
Terminate Pair Sec - Pri on Sec	The name of the task executed on the secondary ESS that ends PPRC mirroring from the secondary ESS to the primary ESS as defined in ESS Copy Services.
Terminate Pair Sec - Pri on Pri	The name of the task executed on the primary ESS that ends PPRC mirroring from the secondary ESS to the primary ESS as defined in ESS Copy Services.

4. Press Enter.

5. Review the settings for the PPRC pair tasks, from the **Define PPRC Pair Tasks** panel, select **Change/Show a Group of PPRC Pair Tasks**.
6. Make changes to the configuration settings as needed.

Improving the performance of volume groups

During resource acquisition, PPRC automatically creates LUN ID mappings. If your cluster contains a large number of volume groups, you can save time by creating these mapping files manually before bringing up the cluster.

To create these mapping files, run the `/usr/es/sbin/cluster/pprc/utlis/cl_store_LUNPairs` command on each cluster node, and pass the names of your volume groups as parameters as shown in the following example:

```
cl_store_LUNPairs MyVg1 MyVg2 MyVg3
```

This populates the `/tmp` directory with files named **VolumePairs.VolumeGroupName**, where *VolumeGroupName* corresponds with the volume group names you specified as parameters (for example, MyVg1, MyVg2, and so on). Each file contains pairs of LUN IDs that resemble the following:

LUN IDs	
PRIMARY_LUN_ID	SECONDARY_LUN_ID
50122221	50122225
50222221	50222225

Verifying and synchronizing the PPRC configuration

The configuration changes that you have completed to this point need to be synchronized to the other cluster nodes.

Verifying configuration for PPRC replicated resources checks the configuration and reports on the following issues:

- PowerHA® SystemMirror® Enterprise Edition for Metro Mirror classes in the PowerHA® SystemMirror® Configuration Database are identical on all nodes.
- The PPRC Command Line Interface is installed correctly on each node.
- The PPRC volume groups are not defined in the **Concurrent Volume Groups** field for a resource group that spans on nodes across sites.
- Sites defined in the PowerHA® SystemMirror® disk subsystem exist in an PowerHA® SystemMirror® site.
- The IP addresses of the ESS systems exist and are reachable.
- The ESS systems for the PPRC replicated resources are defined as ESS disk subsystems to PowerHA® SystemMirror®.
- The two volumes of a PPRC pair are on different ESS systems and on different PowerHA® SystemMirror® sites.
- The volume IDs correspond to physical volumes defined to PowerHA® SystemMirror® cluster nodes.
- The PVIDs of the disks in a volume group at each end of a pair are the same.

Typically, you synchronize the PowerHA® SystemMirror® Enterprise Edition or Metro Mirror configuration with PowerHA® SystemMirror® cluster configuration.

You can also verify and synchronize only PowerHA® SystemMirror® Enterprise Edition for Metro Mirror configuration.

To verify and synchronize PowerHA® SystemMirror® Enterprise Edition for Metro Mirror configuration:

1. From the command line, enter `smit sysmirror`.
2. In SMIT, select **Cluster Applications and Resources > Resources > PPRC-Managed Replicated Resources Configuration**, and press Enter.
3. Select the type of PPRC resource.

4. Select **Verify PPRC Configuration**, and press Enter.

Result

To synchronize PowerHA® SystemMirror® Enterprise Edition for Metro Mirror changes from the command line interface, you can use the **cl_sync_pprc_config** command.

To verify PowerHA® SystemMirror® Enterprise Edition for Metro Mirror changes from the command line interface, you can use the **cl_verify_pprc_config** command.

Configuring resource groups

After defining PPRC replicated resources, you can add them to a resource group.

When you configure a resource group ensure that:

- The site policy is set to **Prefer Primary Site** or **Online on Either Site**.
- A **Startup** policy other than **Online on All Available Nodes** is specified.
- The **Resource Group Processing Ordering** is set to serial.

To add a PPRC replicated resource to a resource group:

1. From the command line, enter `smit sysmirror`.
2. In SMIT, select **Cluster applications and resources > Resource groups > Change/Show resources and attributes for a resource group**, and press Enter.
3. In the **Change/Show resources and attributes for a resource group** panel specify:
 - The name of the PPRC replicated resource in the **PPRC Replicated Resources** field.
 - The volume groups associated with the individual PPRC replicated resources.
4. Verify and synchronize the cluster.

Starting the cluster

After verifying and synchronizing the PowerHA® SystemMirror® Enterprise Edition for Metro Mirror configuration changes, start the PowerHA® SystemMirror® cluster.

All PPRC pairs must be in a **simplex-none-simplex-none** state at the initial cluster startup. This means that no relationship exists between the disk volumes at cluster startup. To view and modify the state of the PPRC pairs, use ESS Copy Services.

Changing PPRC replicated resources configuration

Use SMIT to change configuration for PPRC replicated resources.

Note: Changing resource configuration requires that PowerHA® SystemMirror® services be stopped on all nodes at both sites in the cluster.

Any configuration changes you make to any of the following components effects the others listed:

- Sites
- PPRC replicated resources
- Volumes
- Resource groups.

After you make a configuration change, verify and synchronize the configuration.

Changing the configuration for sites

You can change the configuration for sites.

To change site configuration in PowerHA® SystemMirror®:

1. From the command line, enter `smit sysmirror`.

2. From the SMIT interface, select **Cluster Nodes and Networks > Manage Sites > Change/Show a Site in Stretched Cluster**, and press Enter.

Changing the configuration for PPRC replicated resources

You can change the configuration for PPRC replicated resources.

To change the configuration for PPRC replicated resources:

1. From the command line, enter `smit sysmirror`.
2. In SMIT, select **Cluster Applications and Resources > Resources > PPRC-Managed Replicated Resources Configuration**, and press Enter.
From this panel, select:
 - **Define ESS Disk Subsystem**
For information about field values, see the section *Defining the ESS Disk Subsystems to PowerHA® SystemMirror®*.
 - **Define a PPRC Replicated Resource**
For information about field values, see the section *Defining the PPRC pairs to PowerHA® SystemMirror®*.
 - **Define PPRC Tasks**
For information about field values, see the section *Defining the PPRC Tasks to PowerHA® SystemMirror® Enterprise Edition for Metro Mirror for Direct management*.
 - **Synchronize PPRC Configuration**
 - **Verify PPRC Configuration**
3. After selecting a configuration option, select the **Change/Show** option for the value you want to change.

PowerHA® SystemMirror® Enterprise Edition for Metro Mirror with DSCLI management

These topics describe the planning, installation and configuration tasks for PowerHA® SystemMirror® Enterprise Edition for Metro Mirror with DSCLI management, from here on referred to as DSCLI management. DSCLI management simplifies how you manage PPRC replicated resources on IBM® TotalStorage® systems and how you can integrate PPRC replicated resources into an PowerHA® SystemMirror® configuration.

You are not required to define tasks on the ESS Web Interface when you use this management system.

Plan which resource groups will contain the DSCLI-managed PPRC replicated resources.

Overview of the DSCLI management system

DSCLI management provides the ability to execute Copy Services operations directly without relying on previously saved GUI tasks. The software dynamically manages DSCLI PPRC-controlled disks, providing a fully automated, highly available disaster recovery management solution.

The PowerHA® SystemMirror® interface is designed to communicate with the DSCLI so that once the basic PPRC environment is configured, PPRC relationships are created automatically and there is no need for manual access to the DSCLI.

Integration of DSCLI and PowerHA® SystemMirror® provides:

- Support for either **Prefer Primary Site** or **Online on Either Site** inter-site management policies.
- Flexible user-customizable resource group policies.
- Support for cluster verification and synchronization.
- Limited support for the PowerHA® SystemMirror® Cluster Single Point Of Control (C-SPOC). See *Installing the DSCLI management for PPRC Filesets*.
- Automatic failover/reintegration of server nodes attached to pairs of PPRC-protected disk subsystem within and across sites. See *Installing the DSCLI management for PPRC Filesets*.
- Management of PPRC for:

- Fallback/failback of PPRC paths and instances for automatic movement of PPRC-protected disks between PowerHA® SystemMirror® sites.
- Automatic fallback of PPRC-protected volume groups between nodes within a site. See Installing the DSCLI management for PPRC Filesets.

Using DSCLI allows PowerHA® SystemMirror® to:

- Automatically set up PPRC paths and instances that PowerHA® SystemMirror® will manage.
- Manage switching the direction of the PPRC relationships when a site failure occurs, so that the backup site is able to take control of the PowerHA® SystemMirror®-managed resource groups from the primary site.

Planning for DSCLI management

You should be familiar with the planning tasks for PowerHA® SystemMirror®.

Your environment must meet the following requirements:

- PowerHA® SystemMirror® sites have been planned.
- Basic DSCLI (and ESS CLI, if required for ESS storage,) support has been completely configured. Refer to the appropriate documentation on how to install and configure each.

To plan for a DSCLI management in an PowerHA® SystemMirror® cluster, complete the following tasks:

- Identify the Copy Services Servers (CSS) to be used.
- Identify the disk subsystems to be used in the cluster.
- Identify the vpaths to be used in the configuration, including the volume IDs for each that corresponds to the storage unit and LUN.
- Identify the PPRC Replicated Resources to be used
- Identify the Port Pairs to be used for PPRC Paths
- Identify Volume Pairs (LUNs)
- Identify the volume groups to be managed by PPRC Replicated Resources
- Plan which resource groups will contain the DSCLI-managed PPRC replicated resources (if not already done in the General Planning section).

Limitations and restrictions for DSCLI management

The current release of PowerHA® SystemMirror® Enterprise Edition for Metro Mirror with DSCLI management has some limitations and restrictions.

Check the IBM® Web site for the latest information on TotalStorage® models and PowerHA® SystemMirror® support.

<http://www.ibm.com/servers/storage/disk/index.html>

Refer to the README packaged with the DSCLI management filesets for the most up-to-date Limitations and Restrictions.

Volume group limitations

- A volume group must have the same volume major number across all cluster nodes. (It has been known to cause problems during cluster function time, and it is not guaranteed to be fixed during cluster verification.)
- Resource groups to be managed by PowerHA® SystemMirror® cannot contain volume groups with both PPRC-protected and non PPRC-protected disks. For example:
 - VALID: RG1 contains VG1 and VG2, both PPRC-protected disks.
 - INVALID: RG2 contains VG3 and VG4, VG3 is PPRC-protected, and VG4 is not protected.
 - INVALID: RG3 contains VG5, which includes both PPRC-protected and non-protected disks within the same volume group.

Managed resource limitations

Resource groups cannot manage both DSCLI and Direct management (ESS CLI)-managed PPRC resources simultaneously:

Note: ESS Storage resources (LSS and LUNs) are considered DSCLI PPRC Resources in this type of configuration because they are managed via the DSCLI interface and not the ESS CLI.

IBM® TotalStorage® Copy Services functions limitations

Only IBM® TotalStorage® Copy Services function Synchronous PPRC (Metro Mirror) is supported (No Global Copy or Global Mirror).

C-SPOC limitations

C-SPOC operations on nodes at the same site as the source volumes successfully perform all tasks supported in PowerHA® SystemMirror®.

C-SPOC operations will not succeed on nodes at the remote site (that contain the target volumes) for the following LVM operations:

- Creating or extending a volume group
- Operations that require nodes at the target site to write to the target volumes (for example, changing filesystem size, changing mount point, adding LVM mirrors) cause an error message in CSPOC. However, nodes on the same site as the source volumes can successfully perform these tasks. The changes are subsequently propagated to the other site via lazy update.
- For C-SPOC operations to work on all other LVM operations, it is highly recommended that you perform all C-SPOC operations when the cluster is active on all PowerHA® SystemMirror® nodes and the underlying SVC consistency groups are in a **consistent_synchronized** state.

Sample configuration for DSCLI management

You can set up a mutual recovery configuration in which each site acts as a production site with the other site acting as an associated backup site.

Implementing a mutual recovery configuration requires:

- Two PowerHA® SystemMirror® sites (the same as a single recovery configuration)
- Two resource groups.

A standard configuration includes two PowerHA® SystemMirror® sites, each comprised of nodes attached to two PPRC-managed groups of ESSes spread across both sites.

A CLI client (ESSNI client) must be installed on the PowerHA® SystemMirror® hosts. The ESSNI client is the interface between the application that intends to invoke PPRC commands and the ESSNI or HMC; it must be installed on all the ESSNI server nodes. PPRC services are invoked using the DSCLI within the ESSNI clients. The ESSNI client communicates with the ESSNI server.

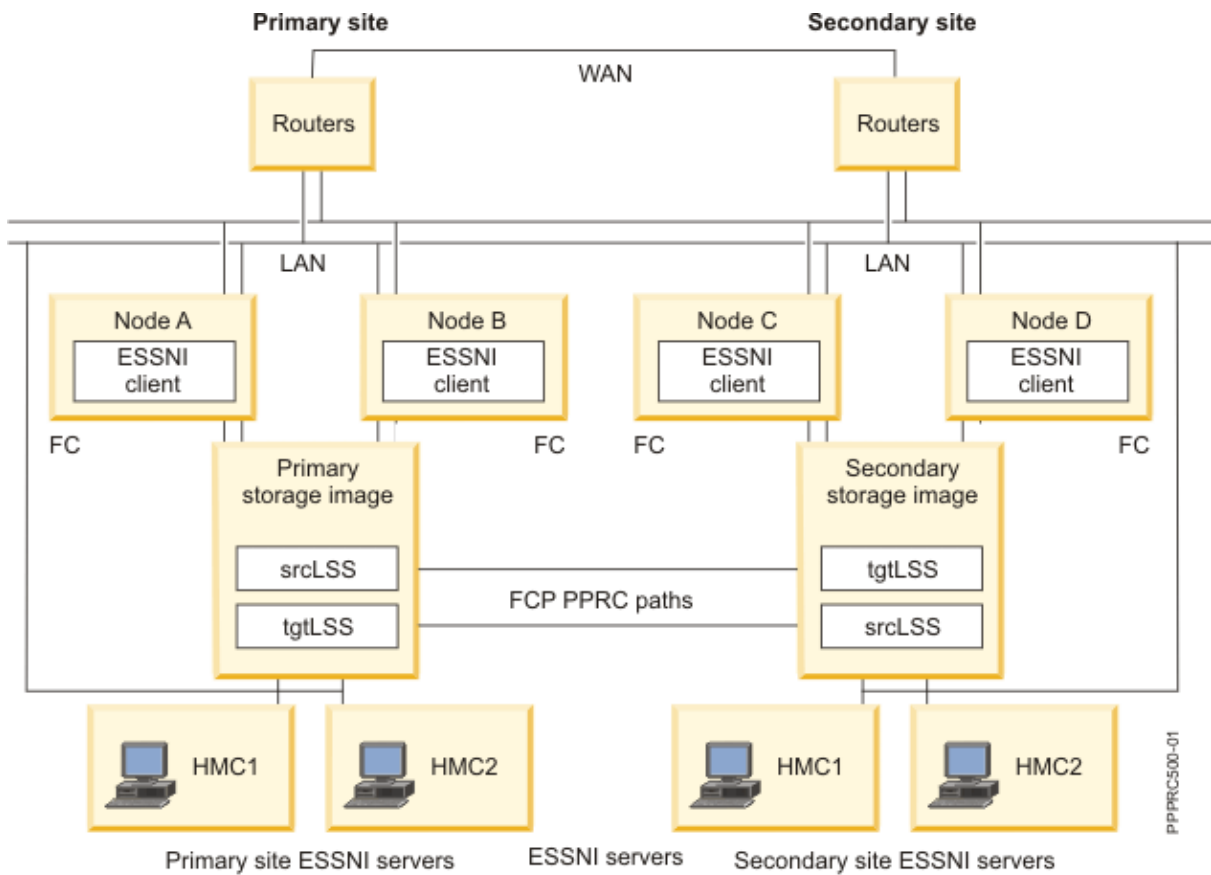
The ESSNI Server runs in an HMC for the ESS 2107 and on the management server for ESS 1750. It runs directly on the ESS clusters for the 2105. The ESS Server, in turn, communicates the CLI commands to the ESS disk controllers.

A PPRC replicated resource contains the ESS disk volume pairs information. The PowerHA® SystemMirror® resource group definition includes the volume groups built on the PPRC replicated volumes. PowerHA® SystemMirror® manages PPRC processing by dynamically executing the DSCLI commands. You no longer have to define tasks on the ESS Web Interface.

The example shows a typical implementation of two ESS model 2107s with PPRC in a four-node PowerHA® SystemMirror® geographic cluster. The cluster consists of four System p® nodes. Each ESS is connected to each node (server) via a SCSI or Fibre Channel connection. Two PPRC links (ESCON® or FC) between the ESSes provide the basic level of redundancy.

One link carries data from the source Logical Subsystem (LSS) on the Primary site to the target LSS on the Secondary site, and the other link carries data in the other direction (source is on the Secondary site, target is on the Primary site).

The configuration also includes point-to-point networks for heartbeating to connect the cluster nodes.



Example listing for DSCLI-managed configuration for mutual takeover

Here is an example of the configuration information for a DSCLI-managed mutual takeover configuration.

```

DS Subsystems
-----
m222:
  Cluster1 IP: 9.22.22.22
  Cluster2 IP: 9.44.44.44
  ESS Storage ID: IBM.2107-2222222
  Associated CS Server: m222h

m555:
  Cluster1 IP: 9.55.55.55
  Cluster2 IP: 9.77.77.77
  ESS Storage ID: IBM.2107-5555555
  Associated CS Server: m555h

Copy Services Servers
-----
m222h: IP address 9.112.112.2
m555h: IP address 9.115.115.2

Available IO ports
-----
m222: I0002, I0012
m555: I0005, I0015

Volumes (LUNs) to be used
-----
m222: 1200, 1201
m555: 1200, 1201

LSS to be used
-----
m222: 12    m555: 12

DSCLI-managed PPRC Replicated Resource for
Resource Group RG1
-----
PPRC Resource Name:   sample_res1
PowerHA SystemMirror Sites:      SiteA SiteB
Volume Pairs:         1200->1200
ESS Pair:             m222 m555

```

```
LSS Pair:          12 12
PPRC Type:         mmir
PRI-SEC PortPairIDs 0002->I0005
SEC-PRI PortPairIDs I0015->I0012
PPRC Link Type:    fcp
Volume Group       sample_VG1
```

DSCLI-managed PPRC Replicated Resource for

Resource Group RG2

```
-----
PPRC Resource Name:  sample_res2
PowerHA SystemMirror Sites:  SiteB SiteA
Volume Pairs:        1201->1201
ESS Pair:            m555 m222
LSS Pair:           12 12
PPRC Type:          mmir
PRI-SEC PortPairIDs I0005->I0002
SEC-PRI PortPairIDs I0012->I0015
PPRC Link Type:     fcp
Volume Group        sample_VG2
```

Note that the definitions for RG2 of volume pairs, ESS Pair, LSS Pair and the PortPairIDs are listed in reference to the source of the PPRC instance that will be used by the resource: RG2 is intended to come ONLINE with the PowerHA® SystemMirror® secondary site as the "source." RG1 is intended to come ONLINE with the PowerHA® SystemMirror® primary site as "source."

From this point, resource groups RG1 and RG2 will be configured to include the DSCLI-managed PPRC Replicated Resources sample_res1 and sample_res2 respectively.

Setting up volume groups and file systems on DSCLI-protected disks

Although *not* required, you should complete these steps prior to planning. These steps must be completed prior to the initial PowerHA® SystemMirror® verification to avoid verification errors.

1. Make sure that the hdisks and corresponding vpaths made available to your nodes are visible at those nodes. If they are *not*, and you can verify that the nodes have been cabled and configured correctly to make the vpaths available, reboot the node and run **cfgmgr** to make the disks viewable.
2. Based on the LUNs you have selected for a given PPRC relationship, determine which vpaths and hdisks correspond. Use the following utility on a node at both the primary and backup PowerHA® SystemMirror® sites:

```
/usr/sbin/lsvpcfg
```

It is likely (although *not* required) that the LUNS will be different at each site, and it is also likely (although *not* required) that the vpaths will be different on each node in each site.

The output will look something like this: (this output would be from an PowerHA® SystemMirror® node in siteA of the example configuration. Assume more than one type of storage unit is connected to this node.)

```
smithers) /usr/sbin/lsvpcfg
vpath12 (Avail ) 13AABKK1602 = hdisk14 (Avail ) hdisk44 (Avail )
hdisk74 (Avail ) hdisk104 (Avail )
vpath13 (Avail ) 13AABKK1603 = hdisk15 (Avail ) hdisk45 (Avail )
hdisk75 (Avail ) hdisk105 (Avail )
vpath14 (Avail ) 13AABKK1604 = hdisk16 (Avail ) hdisk46 (Avail )
hdisk76 (Avail ) hdisk106 (Avail )
vpath15 (Avail ) 22222221100 = hdisk17 (Avail ) hdisk47 (Avail )
hdisk77 (Avail ) hdisk107 (Avail )
vpath16 (Avail pv sample_VG1) 22222221200 = hdisk18 (Avail )
hdisk48 (Avail ) hdisk78 (Avail ) hdisk108 (Avail )
vpath17 (Avail pv sample_VG2) 22222221201 = hdisk19 (Avail )
hdisk49 (Avail ) hdisk79 (Avail ) hdisk109 (Avail )pilot> lshostvol.sh
```

The third column in this output corresponds to the storage system ID and LUN associated with the vpath listed in column one.

Examples

vpath12 (which has no volume group created on it yet) is on storage system IBM®.XXXX-13AABKK, LSS 16, LUN 002.

vpath17 (which has volume group sample_VG2 created on it) is on the same system, with a different LSS/LUN: IBM®.XXXX-222222, LSS 12, LUN 001.

To fill in the 'XXXX' from the system ID: either from your own documentation, or use

```
lsdev -Ccdisk | grep <hdisk associated with the vpath in question>
```

to display the underlying disk type: IBM® FC 1750 or IBM® FC 2107.

Example:

```
smithers) /usr/sbin/lsdev -Ccdisk | grep hdisk14
```

```
hdisk14 Available 2A-08-02      IBM FC 1750
```

You see that vpath12 is on IBM®.1750-13AABKK, LSS 16, LUN 002.

3. Create volume groups and file systems at the Primary PowerHA® SystemMirror® site.
 - a. On one of the nodes at the Primary PowerHA® SystemMirror® site, on the Vpaths) that correspond to the volume pairs for a given PPRC Relationship, set up the volume group(s) and file system(s) to be managed by PowerHA® SystemMirror®. Ensure that the Volume Major Numbers for the volume groups can be used on all PowerHA® SystemMirror® cluster nodes, and that the physical volume name for the filesystem can also be used across all PowerHA® SystemMirror® cluster nodes.

Use the `/usr/sbin/lvlstmajor` command on each node in the cluster to list the available volume major numbers and select a free number on all nodes.

- b. After successfully creating all necessary volume groups and file systems on the first node, import the data to all other nodes at the same site.
4. Use PPRC to mirror disk to the backup PowerHA® SystemMirror® site
5. Create a temporary PPRC relationship in order to copy the volume group/file set information to the remote disk.

Run these commands to set up the PPRC path and instance, and to copy over the local disk information (refer to the DSCLI documentation, or run `dscli help <command>` for more details):

- a. **mkpprcpath**

```
/opt/ibm/dscli/dscli -user <userid> -passwd <password> -hmc1  
<local hmc name> mkpprcpath -dev <local storage device ID> -remotedev  
<remote storage device ID> -srcless -tgtless -remotewwnn  
<WWNN> <local port>:<remote port>
```

Example, where IBM®.2107-222222 is at the Primary PowerHA® SystemMirror® site, IBM®.2107-555555 is at the Backup PowerHA® SystemMirror® site:

```
/opt/ibm/dscli/dscli -user sluggo -passwd batterup -hmc1  
m222h mkpprcpath -dev IBM.2107-222222 -remotedev IBM.2107-555555  
-srcless 12 -tgtless 13 -remotewwnn 6005076303FFC354 I0002:I0005
```

- b. **mkpprc**

```
/opt/ibm/dscli/dscli -user <userid> -passwd <password> -hmc1  
<local hmc name> mkpprc -dev <local storage device ID> -remotedev  
<remote storage device ID> -type <mmir| -mode full <local LUN>:<remote LUN>
```

Example:

```
/opt/ibm/dscli/dscli -user sluggo -passwd batterup -hmc1  
m222h mkpprc -dev IBM.2107-222222 -remotedev  
IBM.2107-555555 -type mmir -mode full 1002:1002
```

At this point, you should have a PPRC instance available, and in a copying state:

```
/opt/ibm/dscli/dscli -user <userid> -passwd <password> -hmc1
<local hmc name> lsprrc -dev <local storage device ID> -remotedev
<remote storage device ID> <local LUN>:<remote LUN>
```

c. **rmpprc**

Once the PPRC relationship has completed copying, delete the relationship (if you do *not*, then the Backup PowerHA® SystemMirror® site nodes will *not* have write access to the LUNs, and so you will *not* be able to import the new volume group(s):

```
/opt/ibm/dscli/dscli -user <userid> -passwd <password> -hmc1
<local hmc name> rmpprc -quiet -dev <local storage device ID>
-remotedev <remote storage device ID> <local LUN>:<remote LUN>
```

This step is necessary in order for the next LVM operations to complete successfully.

- d. Using SMIT or the command line on the backup PowerHA® SystemMirror® site (the site that is connected to the remote disk subsystem), import the volume group(s) created in step b.)

Result

At this point, the volume groups and filesystems necessary to configure PowerHA® SystemMirror® have been created.

Planning Primary and Secondary site layout for resource groups

In order for a resource group to come ONLINE correctly at cluster startup time, the definition of the Primary site must be the same for the PPRC replicated resource and the resource group that contains that resource.

Defining the Primary Site for a DSCLI-managed PPRC replicated resource

The order that data is entered into the SMIT panel fields when defining a PPRC Replicated Resource indicates which site will be primary. The following entries are order-sensitive. In all cases (unless otherwise specified), the first entry corresponds to information for the primary site.

- PowerHA® SystemMirror® Sites
- PPRC Volume Pairs

Example for a resource with multiple volume pairs:

```
[1200->1300 1201->1301]
```

Primary site LUNS are 1200 and 1201. Secondary site LUNS are 1300 and 1301

- ESS Pair
- LSS Pair
- Pri-Sec Port Pair IDs

For both Pri-Sec and Sec-Pri, the following format is correct when entering data in the SMIT panels:

```
[I0022->I0052 I0023->I0053]
```

where I0022 and I0023 correspond to Ports on the storage system directly connected to the site that will be primary for the PPRC Replicated Resource, and I0052 and I0053 are Ports on the storage system directly connected to the secondary site.

- Sec-Pri Port Pair IDs

If you were going to use the same Port pairs as in the previous example, the Sec-Pri listing would look like this:

```
[I0052->I0022 I0053->I0023]
```

where the Ports are on the same systems as described above.

Defining the Primary Site for an PowerHA® SystemMirror® resource group

When defining the Primary and Secondary site nodes on creating an PowerHA® SystemMirror® resource group, select nodes from the site where you want the resource group to come ONLINE during cluster startup as the Primary Site nodes.

Note that combining the **Online on Either Site** inter-site management policy with the **Online on First Available Node** startup policy will allow the resource group to come ONLINE on a node other than those defined as belonging to its defined Primary site.

Example:

```
PowerHA SystemMirror Site 1: node11, node12
PowerHA SystemMirror Site 2: node21, node22

To define an PowerHA SystemMirror Resource Group that will come up on Site 1:
Participating Nodes from Primary Site    [node11, node12]
Participating Nodes from Secondary Site[node21,node22]

To define an PowerHA SystemMirror Resource Group that will come up on Site 2:
Participating Nodes from Primary Site    [node21, node22]
Participating Nodes from Secondary Site[node11,node12]
```

Installing the DSCLI management for PPRC filesets

These topics describe how to install DSCLI management for PPRC filesets. You must be logged in as root to perform installation tasks.

Installing prerequisite software

Before installing PowerHA® SystemMirror® Enterprise Edition DSCLI management for PPRC (spprc filesets), the prerequisite software must be installed on cluster nodes.

Note: Check the README for updates on versions supported after the first release of this software.

1. The latest PowerHA® SystemMirror® Enterprise Edition version. (Refer to Installing PowerHA® SystemMirror® Enterprise Edition for Metro Mirror for more information)
2. IBM® Subsystem Device Driver (SDD) for the storage system you are using. Check the documentation and website for the currently approved SDD version to use with a given microcode version. Make sure the following filesets are installed:
 - a. `devices.fcp.disk.ibm.rte (1.0.0.0)`

Note: It may *not* be clear during installation of the SDD filesets that this first fileset is necessary, but it is critical for correct fallover behavior.

- b. `devices.sdd.**.rte (latest version)`
- c. `devices.ibm2105.rte (latest version)`

Note: This fileset provides connection scripts for all of the ESS and DS disk types, including ESS 800, DS 8000 and DS 6000.

3. DSCLI client software and other configuration specific prerequisites, as shipped with the microcode for the storage systems to be used in this cluster. Refer to the DSCLI documentation for more information on how to install and configure this software.
4. (*Optional*) If ESS storage systems are going to be included in your cluster) ESS CLI software is shipped with the microcode for the storage hardware. DSCLI management code expects the ESS CLI to be installed in the following (non-standard) directory

```
/opt/ibm/ibm2105cli
```

so you may have to create a link from the actual installation location to this location. Be mindful of this, as it will cause problems during both the verification and run-time for your cluster.

Installing the DSCLI management filesets

You need to install the necessary filesets for DSCLI management.

If you have not already done so, install the filesets listed in Contents of the installation media for DSCLI management.

Installation directories

All PowerHA® SystemMirror®-PPRC programs and scripts are located in specific directory and sub-directories.

These include:

```
/usr/es/sbin/cluster/pprc  
/usr/es/sbin/cluster/pprc/spprc
```

All DSCLI programs and scripts are located in the following directory and sub-directories:

```
/opt/ibm/dscli
```

All ESSCLI programs and scripts are expected to be located in the following directory:

```
/opt/ibm/ibm2105cli
```

Upgrading to the latest version of PowerHA® SystemMirror® Enterprise Edition for Metro Mirror

When upgrading from a previous version of PowerHA® SystemMirror® Enterprise Edition for Metro Mirror, you can choose to upgrade the base **pprc** filesets only, or you can add the **cluster.es.spprc** filesets (to add DSCLI management).

Note: Currently there is no migration path from eRCMF management or SVC management to DSCLI management. The statements here apply only to Direct management.

You can install the **cluster.es.spprc** filesets in your current base (Direct management) PPRC environment and continue to operate the current environment. This is possible because the base **pprc** and **spprc** configuration information is stored in different configuration databases (ODMs).

The **spprc** filesets are installed automatically if you already have a previous version of them on your system when you install via `smitty update_all`.

Configuring PowerHA® SystemMirror® Enterprise Edition for MetroMirror using DSCLI management

These topics explain how to configure the DSCLI management with PowerHA® SystemMirror®.

Configuration requirements

Before configuring PowerHA® SystemMirror® Enterprise Edition for Metro Mirror with the DSCLI management interface, ensure that:

- PPRC is configured and running on the storage systems.
- ESSNI client and server software is installed (DSCLI software on all PowerHA® SystemMirror® cluster nodes, for example).

- You have a good understanding of PowerHA® SystemMirror® sites for PPRC replicated resources. For more information about sites, see the section PowerHA® SystemMirror® Sites.
- Both the base PowerHA® SystemMirror® Enterprise Edition for Metro Mirror and DSCLI management filesets are installed on each cluster node.
- The PowerHA® SystemMirror® cluster is configured for:
 - Nodes
 - Sites
 - Networks and network interfaces
 - Service labels, application monitors, etc.
 - Initial resource groups

You can modify the attributes for a resource group later to accommodate PPRC replicated resources.

Steps for setting up the DSCLI management interface:

1. Configure PPRC-Managed Replicated Resources (use the SMIT panels at the bottom of the main PowerHA® SystemMirror® PPRC-Managed Replicated Resources menu):
 - a. Configure Copy Services Servers
 - b. Configure disk systems to be included
 - c. Configure DSCLI-managed PPRC replicated resources
2. Configure PowerHA® SystemMirror® resource groups to include PPRC-Managed replicated resources.

Configuring DSCLI-managed PPRC replicated resources

You should configure DSCLI-managed PPRC replicated resources using SMIT panel.

To define DSCLI-managed PPRC replicated resources:

1. From the command line, enter `smit sysmirror`.
2. In SMIT, select **Cluster Applications and Resources > Resources > PPRC-Managed Replication Resources Configuration > DSCLI-managed PPRC Replicated Resource Configuration**, and press Enter.

Result

From this panel you can:

- Configure Copy Services Server
- Configure DS ESS disk subsystems
- Configure DSCLI-managed PPRC replicated resources.

Configuring Copy Services Server

Configure Copy Service Server using the SMIT panel.

To configure the Copy Services Server:

1. From the command line, enter `smit sysmirror`.
2. In SMIT, select **Cluster Applications and Resources > Resources > PPRC-Managed Replication Resources Configuration > DSCLI-managed PPRC Replicated Resource Configuration > Copy Services Server Configuration > Add a Copy Services Server**, and press Enter.
3. In the **Add a Copy Services Server** panel, enter field values as follows:

<i>Add a Copy Services Server fields</i>	
Field	Value
CSS Subsystem Name	Name that identifies the Copy Service Server. The name may consist of up to 64 alphanumeric characters, and may include underscores.
CSS site name	Name of the PowerHA® SystemMirror® site where the CSS resides. The site must already be defined in PowerHA® SystemMirror® for the site name to be available from the picklist.

Field	Value
CLI Type	Select DSCLI if you are using the ESS 2107. Select ESSCLI if you are using the ESS 2105.
CSS IP Address	IP address, in dotted decimal notation, that the Copy Services Server uses. (This is different from the ESS IP address.)
CSS User ID	User ID used to authenticate logging into the CSS.
CSS password	Password associated with the specified CSS User ID .

4. Press Enter.
5. Repeat these steps for the CSS on the other site.

Defining the DS ESS Disk Subsystems to PowerHA® SystemMirror®

You define the DS ESS subsystems included in the sites that support PowerHA® SystemMirror® Enterprise Edition for MetroMirror to enable PowerHA® SystemMirror® to process failovers for PPRC replicated resources.

To define a DS ESS system to PowerHA® SystemMirror®:

1. From the command line, enter `smit sysmirror`.
2. In SMIT, select **Cluster Applications and Resources > Resources > PPRC-Managed Replication Resources Configuration > DSCLI-managed PPRC Replicated Resource Configuration > DS ESS Disk Subsystem Configuration > Add an ESS Disk Subsystem**, and press Enter.
3. In the **Add an ESS Disk Subsystem** panel, enter field values as follows:

<i>Add an ESS Disk Subsystem fields</i>	
Field	Value
ESS Subsystem Name	Name that identifies the ESS subsystem to PowerHA® SystemMirror®. The name may consist of up to 64 alphanumeric characters, and may include underscores.
ESS site name	Name of the site where the ESS resides. The site is already defined in PowerHA® SystemMirror® for the site name to be available from the picklist.
ESS Cluster 1 IP Address	IP address, in dotted decimal notation of the ESS or DS disk subsystem in cluster 1. The IP address specified here is the IP address of the DS/ESS at this site.
ESS Cluster 2 IP Address	IP address, in dotted decimal notation of the ESS or DS disk subsystem in cluster 2. The IP address specified here is the IP address of the DS/ESS at this site.
ESS User ID	User ID used to authenticate logging into the ESS, if available.
ESS password	Password associated with the specified ESS User ID , if available.
Full ESS Storage ID	Enter the fully qualified ESS storage image ID. This includes the manufacturer, device type, model and serial numbers (MTMS). The format is: manufacturer.type-model-serial number. For example: IBM.2107-921-75FA120
List of CS Servers	From the list, select the CSS that will manage the PPRC of this Disk Subsystem.

4. Press Enter.

5. Review the settings for the ESS subsystem.
From the **Configure an ESS Disk Subsystem** panel, select **Change/Show an ESS Disk Subsystem** and select an ESS system to view. Make changes if necessary and press Enter.
6. Repeat these steps to enter the information for the DS ESS at the second site.

Adding a DSCLI-managed PPRC replicated resource

You can add a DSCLI-managed PPRC replicated resource to your configuration.

To add a PPRC replicated resource:

1. From the command line, enter `smit sysmirror`.
2. In SMIT, select **Cluster Applications and Resources > Resources > PPRC-Managed Replication Resources Configuration > DSCLI-managed PPRC Replicated Resource Configuration > DSCLI-managed PPRC Replicated Resource Configuration > Add a PPRC Replicated Resource**, and press Enter.
3. In the **Add a Replicated Resource** panel, enter field values as follows:

<i>Add a Replicated Resource fields</i>	
Field	Value
PPRC Resource Name	Enter a name of the set of PPRC volume pairs that make up the PPRC replicated resource. Use no more than 64 alphanumeric characters and underscores.
PowerHA® SystemMirror® Sites	Enter the names of the PowerHA® SystemMirror® sites (already defined to PowerHA® SystemMirror®). Enter the name of the primary site followed by the secondary site.
PPRC Volume Pairs	List of PPRC volume pairs that are contained in this PPRC replicated resource. The format is: <div style="background-color: #f0f0f0; padding: 5px; margin: 5px 0;"> Primary Volume ID:Secondary Volume ID </div> <ul style="list-style-type: none"> All PPRC volume pairs in a PPRC replicated resource consist of volumes from the same LSS pair. Volumes in a PPRC replicated resource must be from the same volume group. A volume group can span more than one LSS. A volume group can span more than one ESS disk subsystem.
ESS pair	Set of ESSs associated with this PPRC resource. The first name in the list is the primary ESS and the second is the secondary ESS.
LSS Pair	Set of LSSs associated with this PPRC resource. The first name in the list is the primary LSS and the second is the secondary LSS.
PPRC Type	Indicates whether the PPRC volume relationships will be Metro Mirror mmir or Global Copy gep relationships. <ul style="list-style-type: none"> Metro Mirror maintains the PPRC relationship in a consistent manner. I/O write completion status is returned to the application once the updates are committed to the target ESS. Global Copy maintains the PPRC relationship in a non-synchronous manner. I/O write completion status is returned to the application once the updates are committed to the source ESS. Updates to the target volume are performed at a later point in time. The original order of updates is not strictly maintained.
Pri-Sec Port Pair ID	List of PPRC Path Port Pair IDs of PPRC Links between a primary LSS and secondary LSS. The source and target port must be a Fibre Channel / ESCON® I/O port that is configured for point-to-point or switch fabric topology. A PPRC Path Port Pair ID consists of two Port IDs, one designated as the source port and the other as the target port for a PPRC path. The first Port ID is the designated source Port. The second LSS ID is the designated target

Field	Value
	<p>Port. Separate the two Port IDs of a PPRC Path Port Pair ID with a colon and no whitespace.</p> <p>You can have up to 8 PPRC Path Port Pair IDs defined for each pair of LSSs. Use a white space to separate multiple PPRC Path Port Pair IDs.</p> <p>Example of three Port pairs:</p> <p>I1A10:I2A20<space>I1A11:I2A21<space>I1A12:I2A22</p>
Sec-Pri Port Pair IDs	<p>Same as above. The example here shows the reverse source and target order of the pri-sec pair IDs (12 > 11 instead of 11 > 12):</p> <p>I2A10:I1A20<space>I2A11:I1A21<space>I2A12:I1A22</p>
PPRC Link Type	<p>Select ESCON® or FCP depending on the connection you are using for the PPRC path.</p>
PPRC Critical Mode	<p>This option is used to write-protect the source volume. If the last path fails between the pairs, resulting in the inability to send information to the target, the source becomes write-protected. Current® updates and subsequent attempts to update the source will fail, with a unit check on S/390® or a check condition on SCSI. Values are ON (set critmod) or OFF (do not set critmod).</p> <p>The default is OFF.</p>
Volume Group	<p>Volume group that contains the PPRC volume pairs included in this PPRC replicated resource. The volume group can include volumes pairs from different LSSes as well as from different ESSes.</p>

4. Press Enter.
5. Repeat as necessary to define more DSCLI-managed PPRC replicated resources.

Configuring resource groups

After defining PPRC replicated resources, you can add them to a resource group.

When you configure a resource group ensure that:

- The site policy is set to **Prefer Primary Site** or **Online on Either Site**.
- A **Startup** policy other than **Online on All Available Nodes** is specified.

To add a PPRC replicated resource to a resource group:

1. From the command line, enter `smit sysmirror`.
2. In SMIT, select **Cluster Applications and Resources > Resource Groups > Change/Show Resources and Attributes for a Resource Group**, and press Enter.
3. In the **Change/Show Resources and Attributes for a Resource Group** panel specify:
 - The name of the PPRC replicated resource in the **PPRC Replicated Resources** field.
 - The volume groups associated with the individual PPRC replicated resources. The PPRC Replicated Resources entry is a picklist that displays the resource names created in the previous step. Make sure that the volume groups selected on the Resource Group configuration screen match the volume groups used in the PPRC Replicated Resource.
4. You must verify before synchronizing the cluster.

Verifying the DSCLI-managed PPRC configuration

Verifying the configuration for DSCLI-managed PPRC replicated resources checks the configuration.

It also reports on the following issues:

- SPPRC information in the PowerHA® SystemMirror® Configuration Database (ODM) is identical on all nodes.
- The DSCLI Command Line Interface is installed correctly on each node.

- The PPRC volume groups are not defined in the **Concurrent Volume Groups** field for a resource group that spans on nodes across sites.
- Sites are properly defined for the PowerHA® SystemMirror® and PPRC configuration.
- The IP addresses of the ESS systems exist and are reachable.
- The ESS systems for the PPRC replicated resources are defined as ESS disk subsystems to PowerHA® SystemMirror®.
- The two volumes of a PPRC pair are on different ESS systems and on different PowerHA® SystemMirror® sites.
- The volume IDs correspond to physical volumes defined to PowerHA® SystemMirror® cluster nodes.
- The PVIDs of the disks in a volume group at each end of a pair are the same.
- All PPRC volumes pairs in a PPRC replicated resource can only consist of volumes from the same LSS pair.
- Volumes in a PPRC replicated resource must be from the same volume group
- Correct PPRC links and their Port IDs have been defined
- The defined CLI path and the ESSNI client jar file exist on all the PowerHA® SystemMirror® servers
- The volume pairs have volumes that do exist on the ESSs defined to PowerHA® SystemMirror®.

To verify the PowerHA® SystemMirror® Enterprise Edition DSCLI-managed PPRC configuration:

1. From the command line, enter `smit sysmirror`.
2. In SMIT, select **Cluster Applications and Resources > Resources > Verify and Synchronize Cluster Configuration** and press Enter.
Since the cluster is inactive, the following options appear. Make sure to select **Verify** in the first field (not **both**, you do not synchronize the configuration yet):

<i>Verification and Synchronization fields</i>	
Field	Value
Verify Synchronize or Both	Select Verify only.
Automatically correct errors found during verification?	No is the default. PowerHA® SystemMirror® performs corrective actions without prompting you to perform any action. If you select Interactively , during verification you will be prompted when PowerHA® SystemMirror® finds a problem it can correct related to the following, for example: <ul style="list-style-type: none"> • Importing a volume group • Exporting and re-importing shared volume groups (mount points and filesystems issues) You then choose to have the action taken or not.
Force synchronization if verification fails?	No is the default. If you select Yes , cluster verification runs but verification errors are ignored and the cluster is synchronized. Use the default.
Verify changes only?	No is the default. (Run the full check on resource and topology configuration.) Use the default.
Logging	Standard is the default. You can also select Verbose . Verification messages are logged to <code>/var/hacmp/clverify/clverify.log</code> .

3. Press Enter. The verification output appears in the SMIT Command Status window.
4. If any error messages appear, make the necessary changes and run the verification procedure again. You may see Warnings if the configuration has a limitation on its availability; for example, only one interface per node per network is configured.

Synchronizing the cluster

Typically, you synchronize the PowerHA® SystemMirror® Enterprise Edition for Metro Mirror configuration with PowerHA® SystemMirror® cluster configuration.

To synchronize the PowerHA® SystemMirror® Enterprise Edition DSCLI-managed PPRC configuration:

1. From the command line, enter `smit sysmirror`.
2. In SMIT, select **Cluster Applications and Resources > Resources > Verify and Synchronize Cluster Configuration**, and press Enter.
3. The PowerHA® SystemMirror® Verification and Synchronization panel appears. Select either **both** or **Synchronize** in the first field and press Enter.

Result

The cluster is synchronized.

Starting the cluster

Verification runs automatically at cluster startup unless you turn this option off. After completing the steps above that set volume groups to show up only as vpaths, cluster verification will fail. To avoid this issue, set the **ignore verification errors?** field to **true** on the **Start Cluster Services** SMIT panel.

1. From the command line, enter `smit cl_admin`.
2. In SMIT, select **PowerHA® SystemMirror® Services > Start Cluster Services**, and press Enter.
3. Make the selections for the fields on this panel, setting the field **Ignore verification errors?** to **true**.
4. Press Enter to start cluster services.

Result

PowerHA® SystemMirror® will start up and manage the PPRC resources.

Changing the configuration for sites

Use SMIT to change configuration for PPRC replicated resources.

Note: Changing resource configuration requires that PowerHA® SystemMirror® services be stopped on all nodes at both sites in the cluster.

Configuration changes you make to any of the following components affects the others listed:

- Sites
- PPRC replicated resources
- Volumes
- Resource groups.

After you make a configuration change, verify and synchronize the configuration.

To change site configuration in PowerHA® SystemMirror®:

1. From the command line, enter `smit sysmirror`.
2. In SMIT, select **Cluster Nodes and Networks > Manage Sites > Change/Show a Site in a Stretched Cluster**, and press Enter.

Result

For information about field values, see the section Changing the configuration of DSCLI-managed PPRC replicated resources.

Changing the configuration of DSCLI-managed PPRC replicated resources

You can change the configuration of DSCLI-managed PPRC replicated resources using the SMIT panel.

To change or remove the configuration for DSCLI-managed PPRC replicated resources:

1. From the command line, enter `smit sysmirror`.
2. In SMIT, select **Cluster Applications and Resources > Resources > PPRC-Managed Replication Resources Configuration > DSCLI-managed PPRC Replicated Resource Configuration**, and press Enter. From this panel, select:

<i>Configure PPRC Replicated Resources fields</i>	
Field	Description
Configure Copy Services Server	For information about field values, see the section <i>Configuring Copy Services Server</i> .
Configure DS ESS Disk Subsystem	For information about field values, see the section <i>Defining the DS ESS Disk Subsystems to PowerHA® SystemMirror®</i> .
Configure a DSCLI-managed PPRC Replicated Resource	For information about field values, see the section <i>Adding a DSCLI-managed PPRC Replicated Resource</i> .

3. After selecting a configuration option, select the **Change/Show** option for the value you want to change, or the **Remove** option for the value you want to remove.

Configuring PPRC consistency groups with PowerHA® SystemMirror® Enterprise Edition for Metro Mirror with DSCLI management

These topics describe the planning, installation, and configuration tasks for maintaining the consistency of disk volumes as PPRC consistency groups within PowerHA® SystemMirror® resource groups.

Overview of consistency groups

When applications have one write that is dependent on the completion of another write, these applications are said to have dependent writes. Using dependent writes, these applications can manage the consistency of their data, so that a consistent state of the application data on disk is maintained if a failure occurs in the host machine, software, or the storage subsystem.

Common examples of application dependent writes are databases and their associated log files. Database data sets are related, with values and pointers from indexes to data. Databases have pointers inside the data sets, the catalog and directory data sets, and in the logs. Therefore, data integrity must always be kept across these components of the database.

In disaster situations, it is unlikely that the entire complex will fail at the same moment. Failures tend to be intermittent and gradual, and disaster can occur over many seconds, even minutes. Because some data might have been processed, and other data lost in this transition, data integrity on the secondary volumes is exposed. The mirrored data at the recovery site must be managed so that data consistency across all the volumes is preserved during an intermittent or gradual failure. For example, the situation where a fire starts in the data center. As the fire spreads, perhaps the adapters or the connectivity that is responsible for mirroring some of the data is damaged. If the storage system is able to continue operating, some transactions can continue mirroring while others cannot. This situation becomes a serious issue where dependent writes are concerned.

To maintain the consistency of data across multiple disk volumes at a backup location, the IBM® TotalStorage® disk subsystem Peer-to-Peer Remote Copy function supports the concept of a PPRC consistency group. Disk volumes in a PPRC relationship that are configured into a PPRC Consistency Group are maintained to ensure that a group of dependent updates made to the disk volumes at the primary location are made together as a unit on the disk volumes at the backup location to maintain data consistency. The PPRC Consistency Group attribute changes the behavior of volume pairs when an error occurs that affects any of the volumes in the group. Without the PPRC Consistency Group option, the DSS causes the volumes where the error is detected to enter a suspended state, which means that PPRC mirroring is suspended. This still allows updates to that volume. If the PPRC Consistency Group option is activated, the volume becomes suspended and additionally enters a "long busy" state, where updates are not possible in that state.

Rather than depending solely on the affected volume that remains in long busy state, the PowerHA® SystemMirror® Enterprise Edition initiates the freeze function to quickly suspend all mirroring between all volume pairs within the Consistency Group that is protected by the freeze. As a result, if you place all mirrored

pairs into the same Consistency Group, the consistency of dependent writes is protected on all volumes, LSSs, and disk systems.

As a result, if some disk volumes at the backup location cannot be updated, all are barred from update.

PowerHA® SystemMirror® Enterprise Edition for Metro Mirror provides support for the configuration of disk volumes as PPRC consistency groups within PowerHA® SystemMirror® resource groups.

For more information on PPRC consistency groups refer to:

- IBM® TotalStorage® Enterprise Storage Server® Implementing ESS Copy Services in Open Environments ITSO Redbook SG24-5757 (section 4.6)

Planning for PPRC consistency groups

You should be familiar with the planning tasks for PowerHA® SystemMirror®.

To use PowerHA® SystemMirror® Enterprise Edition for Metro Mirror with TotalStorage® disk volumes with the PPRC consistency group option, you define two PowerHA® SystemMirror® sites, each comprised of a number of AIX® servers attached to a group of PPRC-managed disk subsystems spread across both sites. PowerHA® SystemMirror® PPRC replicated resources are defined in the usual way to contain the ESS/DS disk volume pairs. Volume Groups are built on the PPRC-managed disk volumes and added to PowerHA® SystemMirror® Resource Groups. In addition, PPRC consistency groups are defined to consist of Logical Subsystem (LSS) disk pairs and the corresponding PPRC path information. All PPRC replicated resources that are part of a consistency group are required to be part of the same PowerHA® SystemMirror® resource group. This means that the PPRC volumes in a consistency group cannot span more than one PowerHA® SystemMirror® resource group. However, an PowerHA® SystemMirror® resource group can contain one or more PPRC consistency groups.

PowerHA® SystemMirror® maintains a common state of all consistency groups in a resource group - either all (or none) of them replicate data to the backup site. Cluster Administrators will most likely find the most convenient configuration where all the volume groups in a resource group are part of the same consistency group. Multiple consistency groups in the same resource group are appropriate when the configuration of paths and LSSes do not allow everything to be placed in the same consistency group. See the explanation associated with Figure 2 in Sample Configuration for an example of this.

The following prerequisites are required:

- PowerHA® SystemMirror® sites have been planned.
- Basic DSCLI support has been completely configured. Refer to the appropriate documentation on how to install and configure each.
- DSCLI management in an PowerHA® SystemMirror® cluster has been planned and the following tasks are completed:
 - Identify the Copy Services Servers (CSS) to be used.
 - Identify the disk subsystems to be used in the cluster.
 - Identify the vpaths (for configurations using SDD) to be used in the configuration, including the volume IDs for each that correspond to the storage unit and LUN.
 - Identify the PPRC Replicated Resources to be used.
 - Identify the Port Pairs to be used for PPRC Paths.
 - Identify Volume Pairs (LUNs).
 - Identify the volume groups to be managed by PPRC Replicated Resources.
- Plan PowerHA® SystemMirror® Resource Groups.
- Identify consistency groups that will be managed in a Resource Group.

Planning a resource group for PPRC consistency group

You need to plan your resource groups. In addition to basic resource group planning, it is assumed that resource group attributes, as well as the inter-site management policies for site support, has been planned.

You must also identify the following:

- The PPRC replicated resources to be added to the Resource Group.

- The Consistency Group that will be managed by the Resource Group.

Note: The freeze action is directed to all Consistency Groups in a Resource Group. Therefore, you should ensure that all consistency groups in a resource group have dependencies. Since the PPRC "freeze" action is performed at a resource group level, all disk resources in the same resource group are frozen if there is total loss of the PPRC link between one or two PPRC pairs in a resource group. This ensures that data of dependent applications are not mirrored to the remote site when data of one or two applications in the same resource group cannot be mirrored to the remote site.

If an application depends on more than one volume group, the volume groups might be defined to be part of the same consistency group or different consistency groups. However, it is recommended that they both be added to the same PowerHA® SystemMirror® Resource Group.

Sample consistency group configuration

These illustration display sample consistency group configuration.

Figure 1 illustrates a four-node, two site cluster on which PowerHA® SystemMirror® Enterprise Edition for Metro Mirror has been configured to provide a high availability disaster recovery solution. Site A and Site B are two geographically separated sites. Each site includes two nodes and a disk storage subsystem (labeled DSS A and DSS B). The DSS units are connected by redundant FCP links and use PPRC to replicate data from one storage unit to the other.

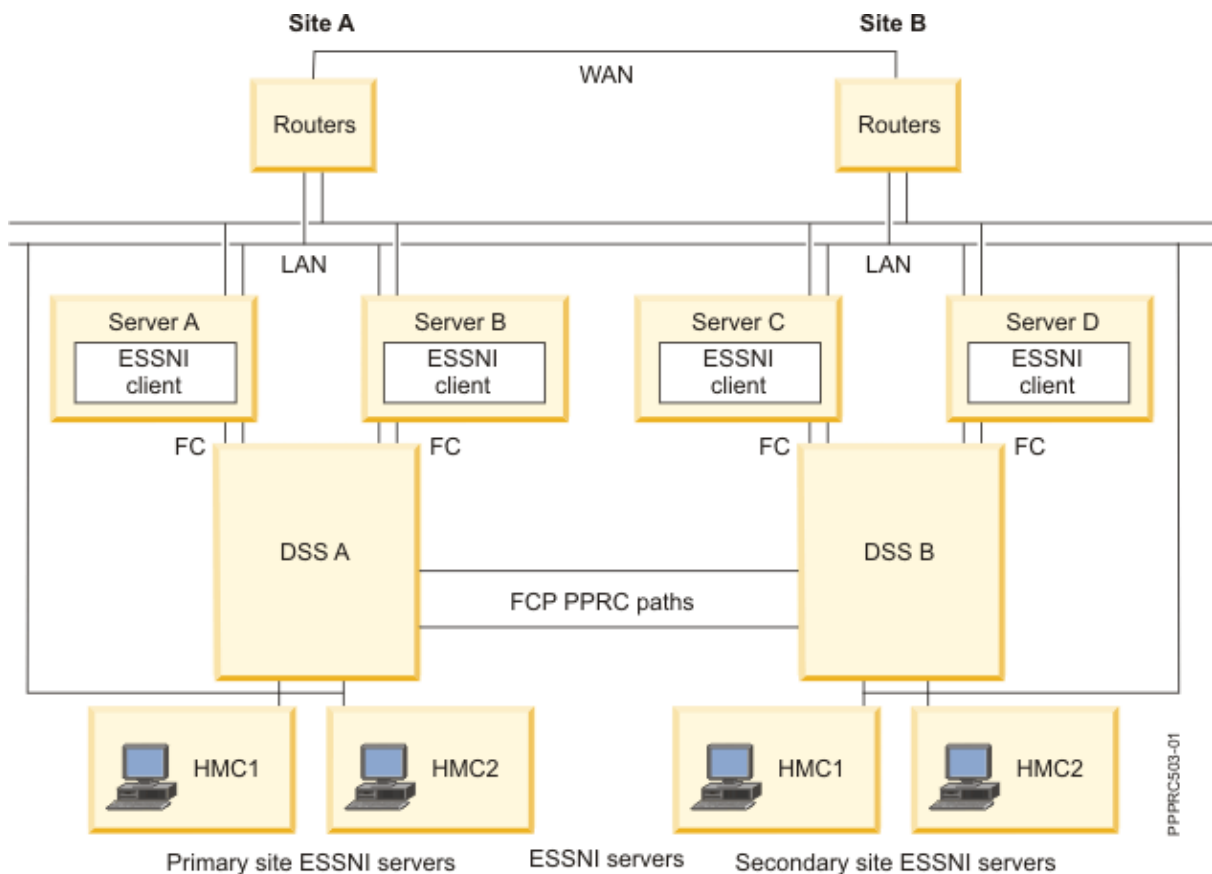
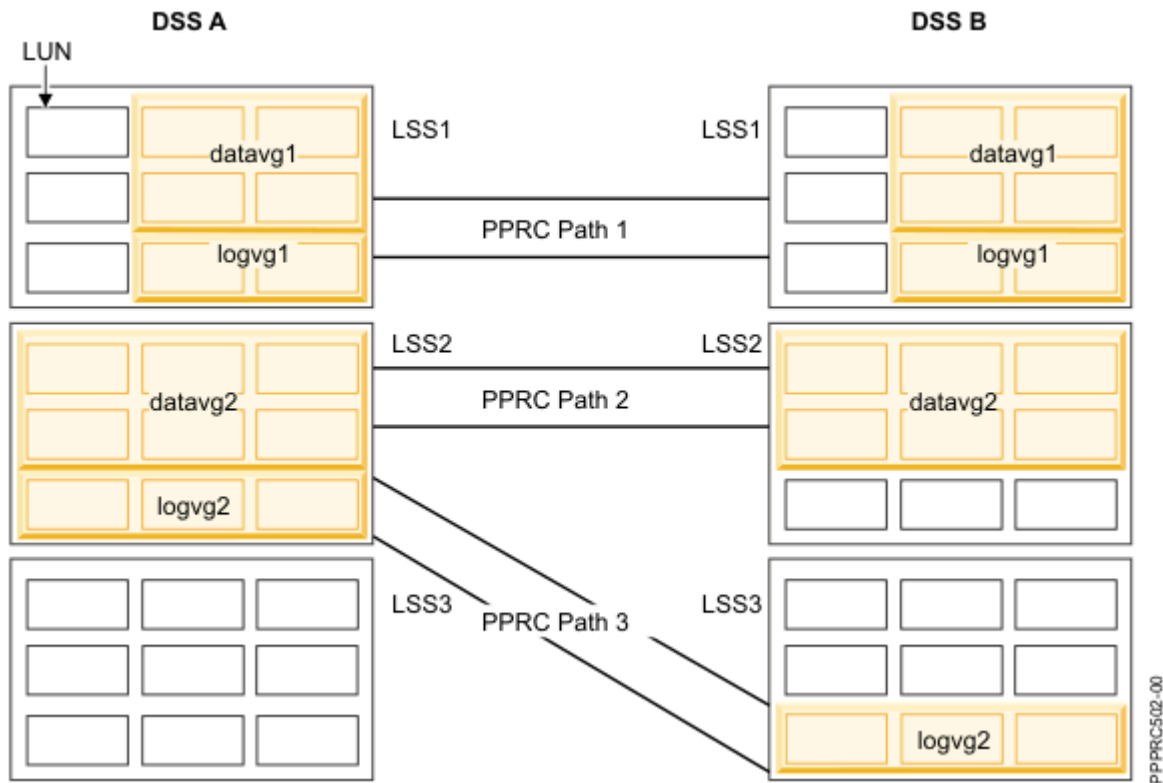


Figure 2 shows a cross-section of two disk subsystems, DSS A and DSS B, located at two geographically separated sites, siteA and siteB. As an example, 3 logical PPRC paths Path1, Path2 and Path3 have defined between the logical subsystems LSS1 on DSS A and LSS1 on DSS B.

A typical configuration, using Database as an example, would require the definition of two volume groups: one for the application data and the other for the log data. In this configuration, it would be required that both the application data (datavg1) and the log data (logvg1) be managed together in a consistency group. For example, a Volume Group (datavg1) that consists of four LUNs, carved from LSS1, on DSS A and a corresponding target Volume Group datavg1 with four LUNs from LSS1 on DSS B, has been defined. A set of four PPRC replicated resources (datapair1, datapair2, datapair3 and datapair4) needs to be defined over the four LUN pairs. Similarly,

logvg1 has been defined with source LUNs from LSS1 on DSS A and target LUNs from LSS1 on DSS B and the corresponding PPRC pair logpair1 defined over the LUN pairs.

To manage all the four PPRC pairs together in a consistency group, a consistency group CG1, is defined for the LSS pairs LSS1 on DSS A and LSS1 on DSS B over the PPRC path Path 1. Using Path1 as the path for this consistency group, whenever the path is established with the consistency group option enabled, all the PPRC pairs that are established between the source volumes of datavg1 and the target volumes of datavg1 are automatically put into a consistency group and associated with Path1. Therefore, all PPRC relationships in datavg1 and logvg1 are all part of the same PPRC consistency group.



On the other hand, CG2 and CG3 are created with the same source LSS (LSS2) but different target LSSes. The application data (datavg2) and log data (logvg2) both have their source volumes in the same LSS but have their target volumes in different LSSes. While datavg2 is associated with Path2, logvg2 is associated with Path3. They are, therefore, not part of the same PPRC consistency group. If an application depends on datavg2 and logvg2, they might be defined to be part of the same consistency group or different consistency group. However, it is recommended that they both be added to the same PowerHA® SystemMirror® Resource Group because the PPRC "freeze" action is performed at a resource group level. All disk resources in the same resource group are frozen if there is total loss of the PPRC link between any of the PPRC pairs in a resource group. This ensures that data of dependent applications are not mirrored to the remote site when data of one or two applications in the same resource group can not also be mirrored to the remote site.

Example configuration of consistency group CG1 described previously.

```

Add a PPRC Consistency Group

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                [Entry Fields]
* PPRC Consistency Group Name    [CG1]
* LSS Pair                       [01@DSS_A 01@DSS_B]
* Primary - Secondary Port Pair IDs [I000->I0101]
* Secondary - Primary Port Pair IDs [I0101-I000]
* Resource Group                  RG1

F1=Help    F2=Refresh    F3=Cancel    F4=List
  
```

F5=Reset	F6=Command	F7=Edit	F8=Image
F9=Shell	F10=Exit	Enter=Do	

Installing PowerHA® SystemMirror® Enterprise Edition for Metro Mirror consistency groups filesets

If you have not already done so, install the filesets according to instructions listed in the Installing PowerHA® SystemMirror® Enterprise Edition for Metro Mirror topic.

cluster.es.cgpprc.cmds

PowerHA® SystemMirror® Enterprise Edition for Metro Mirror Consistency Group commands

cluster.es.cgpprc.rte

PowerHA® SystemMirror® Enterprise Edition for Metro Mirror Consistency Group runtime commands

cluster.msg.en_US.cgpprc

PowerHA® SystemMirror® Enterprise Edition for Metro Mirror Consistency Group messages

Note: You can install the appropriate message set for your local.

Configuring PowerHA® SystemMirror® Enterprise Edition for Metro Mirror consistency groups

Configuration of Consistency Groups with PowerHA® SystemMirror® Enterprise Edition for Metro Mirror with DSCLI management is accomplished in three steps.

The disk storage configuration steps are accomplished through either WebSM or line mode commands. The PowerHA® SystemMirror® configuration steps are accomplished through a series of SMIT panels.

Configure the disk storage units to issue SNMP traps

The storage units must be configured to send traps to one or more of the nodes configured to receive them. You do this by using WebSM.

The precise method is hardware dependent; consult the publications associated with the storage subsystem.

Configure PowerHA® SystemMirror® to receive and handle SNMP traps

Before using the following procedure to enable SNMP Traps, you should ensure that any other consumer of SNMP traps - such as Netview or Tivoli® or other network management software - is not already running. Otherwise the PowerHA® SystemMirror® Cluster Information Daemon (clinfo) is not able to receive traps. Conversely, if PowerHA® SystemMirror® is configured to receive SNMP traps, no other management software is able to receive them.

PowerHA® SystemMirror® supports the receiving and handling of the following SNMP trap messages:

Generic Type = 6

Specific Types:

100

link degraded

101

link down

102

link up

200

LSS pair-consistency group error

201

Session consistency group error

202

LSS is suspended

PowerHA® SystemMirror® tests each SNMP trap received to ensure:

- It is from a valid storage unit (check by storage unit ID).
- It is from a storage unit that has been defined to PowerHA® SystemMirror®.
- It is from an LSS that has been previously configured into a PPRC Resource Group.

If PowerHA® SystemMirror® receives an SNMP trap that does not meet the previous criteria, it is logged, but otherwise ignored.

To enable SNMP Traps in PowerHA® SystemMirror® you must start the Cluster Information Daemon with SNMP Traps enabled, which provides Consistency Groups support. (If clinfo is already running, you must first stop it before using the following steps to restart it.) You do this on the Start Cluster Services SMIT panel (fast path: clstart):

```

Start Cluster Services

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                     [Entry Fields]
* Start now, on system restart or both      now          +
  Start Cluster Services on these nodes    [A2]          +
* Manage Resource Groups                    Automatically +
  BROADCAST message at startup?            true          +
  Startup Cluster Information Daemon?       false         +
  Ignore verification errors?              false         +
  Automatically correct errors found during Interactively +
  cluster start?

F1=Help      F2=Refresh    F3=Cancel    F4=List
F5=Reset     F6=Command    F7=Edit     F8=Image
F9=Shell     F10=Exit     Enter=Do
  
```

Place the cursor on the Startup Cluster Information Daemon? field and hit F4. This gives the following options:

```

Start Cluster Services

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                     [Entry Fields]
* Start now, on system restart or both      now          +
  Start Cluster Services on these nodes    [A2]          +
  BROADCAST message at startup?            true          +
  Startup Cluster Information Daemon?       false         +
  Ignore verification errors?              false         +
-----
|                                     Startup Cluster Information Daemon?                                     |
| Move cursor to desired item and press Enter.                                     |
| false                                                                              |
| true                                                                              |
| true with consistency group support                                             |
|                                                                              |
| F1=Help      F2=Refresh    F3=Cancel    F4=List                                     |
F1|  F8=Image    F10=Exit    Enter=Do                                     |
F5|  /=Find     n=Find Next                                     |
F9| -----
  
```

The options are:

- false - do not start cluster services

- true - start cluster services
- true with consistency group support - start cluster services with SNMP traps enabled

Select true with consistency group support and press Enter.

Configure PowerHA® SystemMirror® consistency group(s)

Before configuring a Consistency Group, you should have performed all the steps required to configure all the PPRC Replicated Resources that will be added to the consistency groups. The PPRC Replicated Resources should also be added to a Resource Group. After defining a PPRC replicated resources and added to a Resource Group, you can add them to a PPRC consistency Group.

To configure a PPRC consistency Group:

1. From the command line, enter `smit sysmirror`.
2. In SMIT, select **Cluster Applications and Resources > Resources > PPRC-Managed Replicated Resources Configuration > DSCLI-managed PPRC Replicated Resource Configuration > PPRC Consistency Groups Configuration**, and press Enter.
The PPRC Consistency Groups Configuration SMIT panel can also be accessed directly with the SMIT fast path: `def_consistgrp`.
3. In the Add a PPRC Consistency Group panel, enter the field values as follows:

```

DSCLI-managed Replicated Resource Configuration

Move cursor to desired item and press Enter.

Copy Services Server Configuration
DS ESS Disk Subsystem Configuration
DSCLI-managed PPRC Replicated Resource Configuration
PPRC Consistency Groups Configuration
Verify PPRC Configuration

F1=Help    F2=Refresh  F3=Cancel   F8=Image
F9=Shell   F10=Exit   Enter=Do

```

4. Select the PPRC Consistency Groups Configuration field to bring up the PPRC Consistency Groups Configuration SMIT panel (fast path: `def_consistgrp`):

```

PPRC Consistency Groups Configuration

Move cursor to desired item and press Enter.

Add a PPRC Consistency Group
Change/Show a PPRC Consistency Group
Remove a PPRC Consistency Group

F1=Help    F2=Refresh  F3=Cancel   F8=Image
F9=Shell   F10=Exit   Enter=Do

```

5. Select Add a PPRC Consistency Group to bring up the Add a PPRC Consistency Group SMIT panel (fast path: `claddconsistgrp.cmdhdr`):

```

Add a PPRC Consistency Group

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

* PPRC Consistency Group Name      [Entry Fields]
* LSS Pair                          []
* Primary - Secondary Port Pair IDs []
* Secondary - Primary Port Pair IDs []
* Resource Group                   []

```

F1=Help F2=Refresh F3=Cancel F4=List
 F5=Reset F6=Command F7=Edit F8=Image
 F9=Shell F10=Exit Enter=Do

6. Fill in the fields as follows:

<i>Add a PPRC Consistency Group fields</i>	
Field	Value
PPRC Consistency Group Name	Enter the name of the PPRC consistency group. PPRC Consistency Group is an alphanumeric string, up to 32 characters long. The name can be freely chosen and has no other restrictions.
LSS Pair	Enter the IDs of the LSS pair. This is the set of LSSs associated with this PPRC Consistency Group separated by a space. The first ID in the list is the LSS in the primary ESS/DSS, and the second is the LSS in the secondary ESS/DSS. Entries are of the format <code>lss_ID@ess_name</code> , where: <code>lss_ID</code> is the LSS ID and <code>ess_name</code> is the name of the ESS/DSS Disk Subsystem containing this LSS as listed in the disk subsystem definition. In smitty, you can use the F4 option to list the available LSSs on each disk subsystem. The selection screen is there only to aid you with this information. You must verify that any selections made through the F4 pick list are listed in the correct order and in the proper format.
Primary - Secondary Port Pair IDs	Enter the port pair IDs for establishing path(s) between the Primary and Secondary LSS. The primary port ID is listed first, followed by the secondary with the string <code>"->"</code> between them. You can list additional port pairs as a space separated list. These should match the port pairs in the resource group definition.
Secondary - Primary Port Pair IDs	Enter the port pair IDs for establishing path(s) between the Secondary and Primary LSS. The secondary port ID is listed first, followed by the primary with the string <code>"->"</code> between them. You can list additional port pairs as a space separated list. These should match the port pairs in the resource group definition.
Resource Group	Select the Resource Group that has PPRC replicated resources associated with this Consistency Group. This Resource Group must be defined and contain the PPRC replicated resources defined from this LSS pair. The resource group must be selected through the F4 pick list in SMIT.

7. Press Enter to apply the configuration settings.

Changing the configuration of a PPRC consistency group

You can change the configuration of a PPRC consistency group using SMIT.

To change or remove the configuration of an existing PPRC Consistency Group:

- From the command line, enter `smit sysmirror`.
- In SMIT, select **Cluster Applications and Resources Configuration > Resources Configuration > Configure PPRC-Managed Replicated Resources > DSCLI-managed PPRC Replicated Resource > PPRC Consistency Groups Configuration > Change/Show a PPRC Consistency Group** and press Enter.
- Select the PPRC consistency group you want to change or view.

```

PPRC Consistency Groups Configuration

Move cursor to desired item and press Enter.

Add a PPRC Consistency Group
Change/Show a PPRC Consistency Group
Remove a PPRC Consistency Group
  
```

```

| Select the PPRC Consistency Group to show or change
|
| Move cursor to desired item and press Enter.
|
|   CG1
|   CG2
|
| F1=Help      F2=Refresh      F3=Cancel
| F8=Image     F10=Exit       Enter=Do
F1 | / =Find      n=Find Next
F9 |
|-----|

```

4. Position the cursor to the consistency group to show or change and then press Enter to bring up the following SMIT panel:

```

Change / Show a PPRC Consistency Group

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

* PPRC Consistency Group Name      [Entry Fields]
New PPRC Consistency Group Name    CG1
* LSS Pair                          []
* Primary - Secondary Port Pair IDs [00@mtc111h 01@stand] +
* Secondary - Primary Port Pair IDs [I0233->I0002]
Resource Group                     [I0002->I0233] +
                                   RG1

F1=Help      F2=Refresh      F3=Cancel      F4=List
F5=Reset     F6=Command     F7=Edit        F8=Image
F9=Shell     F10=Exit       Enter=Do

```

5. Make the desired changes and then press Enter.
6. To remove a previously configured consistency group, select **Remove a PPRC Consistency Group** to bring up an inset panel from which to select the consistency group to remove:

```

PPRC Consistency Groups Configuration

Move cursor to desired item and press Enter.

Add a PPRC Consistency Group
Change/Show a PPRC Consistency Group
Remove a PPRC Consistency Group

|-----|
| Remove a PPRC Consistency Group
|
| Move cursor to desired item and press Enter.
|
|   CG1
|   CG2
|
| F1=Help      F2=Refresh      F3=Cancel
| F8=Image     F10=Exit       Enter=Do
F1 | / =Find      n=Find Next
F9 |
|-----|

```

7. Position the cursor to the consistency group to remove and then press Enter to bring up the following confirmation:

```

PPRC Consistency Groups Configuration

Move cursor to desired item and press Enter.

Add a PPRC Consistency Group

```

```

Change/Show a PPRC Consistency Group
Remove a PPRC Consistency Group

      ARE YOU SURE?

Continuing may delete information you may want
to keep. This is your last chance to stop
before continuing.
      Press Enter to continue.
      Press Cancel to return to the application.

F1 | F1=Help           F2=Refresh          F3=Cancel
F8 | F8=Image          F10=Exit            Enter=Do
F9 |

```

8. Press Enter to confirm the removal of the PPRC consistency group.

Verifying the PowerHA® SystemMirror® Enterprise Edition for Metro Mirror consistency groups configuration

The configuration of SNMP Traps and Consistency Groups is verified during PowerHA® SystemMirror® cluster verification.

You can also verify the Consistency Group configuration through the Verify PPRC Configuration option in the main DSCLI-managed Replication Resource Configuration SMIT Panel (fast path: pprc_ds), that appears as follows:

```

Verify PPRC Configuration

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                [Entry Fields]

Verify PPRC

F1=Help      F2=Refresh    F3=Cancel    F4=List
F5=Reset     F6=Command    F7=Edit      F8=Image
F9=Shell     F10=Exit      Enter=Do

```

Press Enter to begin PPRC configuration verification. Verify first checks the Copy Services Server, DS ESS Disk Subsystem and PPRC Replicated Resources configuration. Upon successful verification and the existence of Consistency Group configuration, verification of Consistency Groups takes place. Verification checks the following for each Consistency Group definition:

- LSS Pair: for both the source and target (LSS@Disk_Subsys) verify that the LSS exists on the DS ESS Disk Subsystem.
- Resource Group: verify that the specified Resource Group is a valid resource.
- Port Pair IDs: for both the primary and secondary list of Port Pair IDs, verify that each port exists on the respective primary and secondary disk subsystems.

Troubleshooting PPRC consistency groups

Use these tips to troubleshoot your PPRC consistency groups.

The following is an example of an Error Log entry made for an SNMP Trap:

```

LABEL:OPMSG
IDENTIFIER:      AA8AB241

Date/Time:Mon Jun 25 12:20:04 2007
Sequence Number: 16480
Machine Id:      0025A45C4C00
Node Id: regaa07
Class: 0
Type: TEMP

```

```
Resource Name:  SNMP_TRAP_NOTIF
```

```
Description  
OPERATOR NOTIFICATION
```

```
User Causes  
ERRLOGGER COMMAND
```

```
Recommended Actions  
REVIEW DETAILED DATA
```

```
Detail Data  
MESSAGE FROM ERRLOGGER COMMAND  
2007/06/25 10:17:53 PDT  
PPRC Links Up  
UNIT: Mnf Type-Mod SerialNm LS  
PRI:  IBM 2107-922 75-16231 00  
SEC:  IBM 1750-511 13-AA44A 00  
Path: Type PP  PLink SP  SLink RC  
1:    FIBRE 0101 XXXXXX 0001 XXXXXX OK  
2007/06/25 13:10:17 EDT  
PPRC Links Up  
UNIT: Mnf Type-Mod SerialNm LS  
PRI:  IBM 2105-800 13-22012 10  
SEC:  XXX 2105-XXX XX-16231 FD  
Path: Type PP  PLink SP  SLink RC  
1:    FIBRE 0004 XXXXXX 0101 XXXXXX OK
```

The following steps should be followed when an PowerHA® SystemMirror® cluster does not appear to exercise consistency group operations:

1. Check that the storage unit HMC is configured to send the SNMP traps to the IP address(es) of the PowerHA® SystemMirror® cluster node(s) that processes them. Normal best practice is to have the HMC configured to send SNMP traps to all PowerHA® SystemMirror® nodes. If that is not possible, at a minimum, it should send SNMP traps to two nodes, one at each site.
2. Check that the IP address of the copy services server that controls the storage subsystems is configured to PowerHA® SystemMirror®. This can be checked by doing smitty **pprc_def > DSCLI-managed PPRC Replicated Resource Configuration > Copy Services Server Configuration > Change / Show a Copy Services Server**.
3. Before starting clinfo with consistency group support, verify that there is no other consumer of SNMP traps running on the PowerHA® SystemMirror® nodes that process them. Run the following command:

```
Netstat -an | grep 162
```

There should not be any output from the command. If there is, it will look like the following:

```
udp40      0  *.162 *.*
```

Programs cannot share access to port 162.

4. Try pinging the nodes that are to receive the SNMP traps from the storage subsystem HMC, using the IP address configured in that HMC for the SNMP trap destination.
5. Try pinging the storage subsystem HMC, using the IP address defined for the copy services server, from the PowerHA® SystemMirror® cluster nodes that are to process SNMP traps.
6. Check the error log using `errpt -a` and see if there are any SNMP trap entries from the storage subsystem. An example of such an entry is:

```
LABEL:OPMSG  
IDENTIFIER:  AA8AB241  
  
Date/Time:Mon Jun 25 12:20:04 2007  
Sequence Number: 16480  
Machine Id:    0025A45C4C00  
Node Id:  regaa07  
Class: 0  
Type: TEMP  
Resource Name:  SNMP_TRAP_NOTIF  
  
Description  
OPERATOR NOTIFICATION
```


However, these operations cannot involve allocating additional storage that requires modification of the PPRC pairs. If you make configuration changes that allocate new storage or create logical units, first stop the associated PPRC volume pairs.

From C-SPOC, you cannot modify the configuration on nodes connected to target volumes for a PPRC pair.

Restarting ESS Copy Services Server on the backup CSS

An ESS system with a microcode level less than vrmf 2.2.x.x does not support dual active Copy Services Servers.

If the primary ESS Copy Services Server fails, you manually start the backup Copy Services Server as the active Copy Services Server for the ESS. You do this through the using the ESS Web Interface.

Detecting ESS failures

There is no simple way to detect the failure of an ESS subsystem in a direct management environment.

You can check for the loss of quorum for all the volume groups indicated by a LVM_SA_QUORCLOSE error. These errors are written to the AIX® error log. PowerHA® SystemMirror® checks whether the LVM_SA_QUORCLOSE error appeared in the AIX® error log file and if so, informs the Cluster Manager to selectively move the affected resource group. This PowerHA® SystemMirror® function is called selective fallover for volume group loss.

Note: If fallover does not occur, check that the LVM_SA_QUORCLOSE error appeared in the AIX® error log. When the AIX® error log buffer is full, new entries are discarded until space becomes available in the buffer and an error log entry informs you of this problem..

Certain storage subsystems have been found to generate large amounts of error log entries when they are experiencing problems. If the error log is not large enough, these error log entries can overrun the error log causing the LVM_SA_QUORCLOSE to be lost. To avoid this problem, configure the error log to have adequate space using the following command:

```
errdemon -B 1048576 -s 10485760
```

This sets the error log buffer size to 1MB and the log size to 10 MB.

Troubleshooting DSCLI-managed clusters

This section provides information that may help with troubleshooting DSCLI management clusters.

When a DSCLI management PowerHA® SystemMirror® Cluster is in a less than optimal functioning state, check the `/tmp/hacmp.out` file for any obvious issues, and then, if necessary, check the states of the PPRC instances and paths that underlay the volume groups to make sure they are write-accessible. Use the commands listed here to discover what the PPRC instances states are, as this will be a major indicator of the health of your volume groups.

If you encounter problems with your cluster that are *not* covered in this section or are *not* easily remedied, contact your IBM® Service Representative.

Common cluster startup problems

This section discusses some common cluster startup problems.

1. If a resource group will *not* stabilize on its primary node and site, it is probable that the PPRC Replicated Resource and Resource Group are defined to be primary on different sites. Refer to the Planning Primary and Secondary Site Layout for DSCLI-managed PPRC Replicated Resources in PowerHA® SystemMirror® Resource Groups. Make sure all the appropriate entries in the PPRC replicated resource are aligned in the correct direction, as described in this and the Sample Configuration section.
2. Another reason a resource group may *not* come ONLINE, may be unstable, or go into ERROR state is if there is trouble with the PPRC instance or path that the volume group and PPRC replicated resource being

managed requires. See the descriptions below for commands to use to display PPRC instance states and a description of normal working states.

3. Keep in mind the limitation that a volume group needs to have the same volume major number across all cluster nodes. This has been known to cause instability in a resource group when it initially comes online.
4. If the Pri-Sec and Sec-Pri port pairs are *not* declared correctly, this will lead to errors in creating the initial PPRC paths, which will keep the cluster from coming up correctly. In the `/tmp/hacmp.out` file, there will be an error message returned from the `mkpprcpath` call for the PPRC replicated resource associated with the resource group that will *not* come ONLINE.
5. In certain instances, disk reserves can be dropped on vpaths or hdisks, that DSCLI management is *not* able to break. These will keep a resource group from being able to come ONLINE because the disk that has a reserve on it will never be write accessible. See Other Handy AIX® Commands below for ideas on how to break disk reserves.

Handy DSCLI commands to know

Refer to the DSCLI documentation to become familiar with these management commands. Ideally you will *not* have to use the DSCLI interface except in troubleshooting situations.

n that case, however, the following commands are useful (all run through the `/opt/ibm/dscli/dscli` command):

- **lspprc**: list the current state of the PPRC instances (see here for valid and common operational states).
- **lspprcpath**: list the current state of the PPRC paths configured.
- **lsavailpprcport**: list the available PPRC port pairs between two storage systems.
- **failoverpprc**, **failbackpprc**: switch the direction of PPRC instances.
- **rmpprc**: remove an existing PPRC instance.
- **rpmkpprc**: create a new PPRC instances.
- **mkpprcpath**: create a new PPRC path.

Refer to the DSCLI documentation for detailed descriptions of these commands, and of others that will allow you to view details of the connected storage systems.

PPRC instance states

PPRC instance states exist during normal operation.

The states are described in pairs because if you poll the DSCLI at the primary site and the secondary site you will get different state information which reflects the I/O accessibility of the disk at that site, as well as the state of the PPRC instance.

Normal Working State

Primary state: Full Duplex

Secondary state: Target Full Duplex

This is the normal working state for a PPRC instance. In this state, the Primary is accessible for read and write I/O but the Secondary is *not* accessible for either.

Suspended State

Primary state: Full Duplex

Secondary state: Suspended

In this state, the Primary is accessible for read and write I/O and the Secondary is as well. Data integrity between the two sides of the PPRC instance is *not* ensured in this state.

In the context of PowerHA® SystemMirror® clusters, this PPRC instance state occurs when there has been a failure on the PowerHA® SystemMirror® site connected to the primary site, causing a failover to the secondary state.

Copy State

Primary and secondary state: Copy Pending

In this state, there is a full disk copy in process between one side of the PPRC instance and the other.

Handy AIX® commands

This section discusses some commonly used AIX® commands.

Refer to the AIX® documentation for more details on the following commands:

To discover whether a disk reserve exists on a given disk:

```
lquerypr -Vh /dev/<disk name>
```

To clear a disk reserve that exists on a given disk:

```
lquerypr -ch /dev/<disk name>
```

To list the existing vpath to storage volume ID (LUN + storage ID) mapping:

```
lsvpcfg
```

To list the existing vpaths on your system, which will display the volume group mapping:

```
lspv | grep vpath
```

To list the existing volume groups on your system:

```
lsvg
```

To list the existing file systems on your system:

```
lsfs
```

To list the available volume major numbers on your system:

```
lvlstmajor
```

Global Mirror

You can use Global Mirror on the IBM® DS8700 to implement asynchronous mirroring into your PowerHA® SystemMirror® configuration.

To implement an asynchronous PowerHA® SystemMirror® configuration with Global Mirror on DS8700 complete the following steps:

1. Plan for the deployment of DS8700 storage across two sites. This step involves understanding your application and middleware requirements and ensuring that you have those requirements covered in the context of a high availability PowerHA® SystemMirror® solution.
2. Set up and configure the DS8700 storage system.
3. Plan for the PowerHA® SystemMirror® implementation. This includes designing resource groups, policies, and scripts around your applications and mirrored storage resources.
4. Install and configure PowerHA® SystemMirror® to implement the settings you planned for in Step 3.
5. Test and verify the configuration is working correctly.

Note: This documentation focuses on Steps 3 and 4, but will mention and discuss the other steps at a high level.

Overview of Global Mirror

The IBM® DS8700 storage subsystem provides asynchronous mirroring using the Global Mirror feature. Global Mirror is implemented through a combination of FlashCopy® functions, Global-Copy functions, and Sessions functions to achieve consistency in the context of asynchronous replication.

FlashCopy®

A technology offered by all enterprise storage systems that allows users to create near-instantaneous copy of a physical volume. The instantaneous copy is handled by copying data in the background while allowing the source and target volumes to be read and written. While the background copying is in progress, all writes to the source volume are handled by the storage system by first copying the old block to the target volume. This process only occurs if the corresponding block on the target has not already been modified. Once the background copying completes, both source and target volumes, are exact copies of the last time the FlashCopy® command was issued.

Global Copy

A DS8000® technology that allows data to be mirrored to another storage unit in an asynchronous manner, without giving any guarantees about the order of data replication. Data can appear on the recovery site out-of-order. This is a good option for any application that wants to handle disaster recovery, but can tolerate out-of-order data.

Sessions

A mechanism by which a DS8000® Series storage system coordinates the replication process with the storage systems on the recovery site. When the Sessions function is used with the FlashCopy® function and the Global Copy function you can achieve asynchronous replication while maintaining consistency properties.

Planning for Global Mirror

To correctly implement Global Mirror you must plan correctly.

The following restrictions and limitations apply to Global Mirror:

- You must be using the IBM® DS8700 storage system. Global Mirror is not supported on earlier models of the DS8000® storage systems.
- If the volumes being replicated span across multiple IBM® DS8700 Logical Sub System, then the same Global Mirror Session Identifier must be used for all of them. The same Global Mirror Session Identifier must be used on both PowerHA® SystemMirror® sites.
- The Hardware Management Console (HMC) associated with the IBM® DS8700 storage system that is used with PowerHA® SystemMirror®, must be accessible using TCP/IP on all PowerHA® SystemMirror® nodes and both sites.
- The DS8000® configuration can not be changed using DARE, even though the DS8000® Global Mirror Replicated Resources or Mirror Groups can be included and excluded from a PowerHA® SystemMirror® Resource Group using DARE.
- PowerHA® SystemMirror® does not handle SNMP events in the context of DS8700 Global Mirror function. If there is a failure in the PPRC path, then in some cases, this can put the corresponding (active) Global Mirror session in a fatal state and PowerHA® SystemMirror® mirror would not know about it. To fix this problem, you must correct the PPRC path and then restart the corresponding resource groups on the appropriate nodes.

Planning the deployment of the IBM® DS8700 across two sites

The application data that you want to mirror using Global Mirror must already reside on a DS8700 Storage System.

If you do not have any application data, then follow the migration recommendations and steps described in the [DS8700 migration documentation](#).

To plan for the deployment of the IBM® DS8700 across two sites, complete the following steps:

1. Identify your production site for the application. A production site is the location where your application and its data would primarily reside and function.

2. Identify the volumes on the DS8700 Storage System that contain the application data that you want to be Highly Available with Disaster Recovery.
3. Identify the storage unit(s) and AIX® hosts that will be running on the recovery site.
4. Ensure that there are sufficient number of volumes and Fiber Channel ports available on the storage systems for the production site and recovery site. This allows a mirror path or PPRC path to be created between the storage units.
5. Verify that there is a FlashCopy® relationship established for each volume that resides on the production site and recovery site.

Planning the PowerHA® SystemMirror® implementation

To successfully implement PowerHA® SystemMirror® you must plan accordingly.

Before you implement PowerHA® SystemMirror® you must complete the following:

- Collect the following information for all the HMCs in your environment:
 - IP addresses
 - Login names and passwords
 - Associations with various storage units
- Verify that the `/opt/ibm/dscli/profile/dscli.profile` file does not contain the DSCLI password. If the DSCLI password is in the `/opt/ibm/dscli/profile/dscli.profile` file, any commands that are run by PowerHA® SystemMirror® fail.
- Verify that all the data volumes that need to be mirrored are visible to all relevant AIX® hosts. The DS8700 volumes should be appropriately zoned so that the FlashCopy® volumes are not visible to the PowerHA® SystemMirror® nodes.
- Ensure all HMCs are accessible using the TCP/IP network for all PowerHA® SystemMirror® nodes where you want to run Global Mirror.

Installing Global Mirror

PowerHA® SystemMirror® Enterprise Edition for AIX® file sets must be installed on all the AIX® nodes that you intend to run your PowerHA® SystemMirror® cluster.

Installation Prerequisites

Before installing PowerHA® SystemMirror® Enterprise Edition for AIX® for Global Mirror, make sure the following is true for each node in your cluster:

- You have root access to each node.

Software Requirements

The following software is required for PowerHA® SystemMirror® Enterprise Edition for AIX® Enterprise Edition for Global Mirror:

- Global Mirror functionality works with all the AIX® levels that are supported by PowerHA® SystemMirror® Standard Edition.
- PowerHA® SystemMirror® Enterprise Edition for AIX® 6.1, or later, with the APARs stated in the support flash and README.
- IBM® DS8700 microcode bundle must be 75.1.145.0, or later.
- Latest DS8000® DSCLI client interface must be installed on each PowerHA® SystemMirror® node.
- Add the pathname for the dscli client into the PATH for the root user on each PowerHA® SystemMirror® node.

Installation Steps

To install Global Mirror, complete the following steps:

1. Insert the PowerHA® SystemMirror® Enterprise Edition for AIX® installation media into the CD or DVD drive.
2. From the command line, enter `smitty install`.
3. From SMIT, select **Install and Update Software > Install Software**.
4. Select the following filesets from the list, **cluster.es.genxd.rte**, **cluster.es.genxd.cmds**, and **cluster.msg.genxd.xx_XX** where **xx_XX** is the message fileset for your language, and press Enter.

Configuring Global Mirror

After you have planned and installed Global Mirror you can start to configure it. You must ensure that the data volumes on both sites are exact replica of each other, otherwise you will not be able to import the volume group on the remote site.

Configuring the IBM® DS8700 to use Global Mirror

Before you can configure the IBM® DS8700 to use Global Mirror, data must already exist on the the IBM® DS8700.

To configuring the IBM® DS8700 to use Global Mirror, complete the following steps:

1. For each data volume that will participate in the Global Mirror on the production site, complete the following steps:
 - a. Identify a volume that will serve as its Global Copy Peer. This volume must reside on a DS8700 storage unit on the recovery site.
 - b. Depending on the DS8700 Logical Sub System (LSS) of the source and target volumes for the Global Copy, you want to create a PPRC path between the source and target LSS. A Global Copy path is a logical path created between two LSS that belong to different DS8700 Storage Systems. Global Copy paths are unidirectional, therefore a path must be created in both directions.

Each logical path between a given LSS pair, must be created on a different physical link. For redundancy and bandwidth, you should create multiple paths in the same direction.
 - c. Establish a Global Copy relationship between the volume on the production site storage system and its mirror volume on the recovery site.
 - d. On the production site for each Global Copy source volume, establish a persistent FlashCopy® relationship with the FlashCopy® source being the Global Copy volume.
 - e. On the recovery site for each Global Copy target volume, establish a persistent FlashCopy® relationship with the FlashCopy® source being the Global Copy volume.
2. On the production site, complete the following steps:
 - a. For each LSS that contains one or more volumes that are participating in the Global Copy relationships, create a Global Mirror Session. Verify that the Global Mirror Session Identifier matches for all the LSSs.
 - b. Add all the Global Copy volumes into the Global Mirror Session.
3. Repeat [Steps 2](#) for the recovery site. You must use the same Global Mirror Session Identifier that was used on the production site for the recovery site.
4. If the application has data volumes that are spread across multiple storage units on the production site, you must use the same Global Mirror Session Identifier on all the storage units.

Configuring PowerHA® SystemMirror® Enterprise Edition for Global Mirror

Once you have the IBM® DS8700 configured correctly, you can configure PowerHA® SystemMirror® Enterprise Edition for Global Mirror.

Before you start configuring PowerHA® SystemMirror® Enterprise Edition for Global Mirror, your environment must meet the following requirements:

- All PowerHA® SystemMirror® clusters are defined.
- All PowerHA® SystemMirror® nodes are defined.
- All PowerHA® SystemMirror® sites are defined.
- All PowerHA® SystemMirror® resource groups and associated resources are configured and working.

Configuring a Storage Agent

A Storage Agent is a generic name given by PowerHA® SystemMirror® for an entity such as the IBM® DS8000® Hardware Management Console (HMC). Storage Agents typically provide a one-point coordination point and often use TCP/IP as their transport for communication.

You must provide the IP address and authentication information that will be used to communicate with the HMC.

Adding a Storage Agent

To add a Storage Agent, complete the following steps:

1. From the command line enter, `smit sysmirror`.
2. From the SMIT interface, select **Cluster Applications and Resources > Resources > Configure DS8000 Global Mirror Resources > Configure Storage Agents > Add a Storage Agent** and press Enter.
3. Complete the following fields:

<i>Adding a Storage Agent fields</i>	
Fields	Values
Storage Agent Name	Enter the PowerHA® SystemMirror® name for this HMC.
IP Addresses	Select list of IP addresses for the HMC.
User ID	Enter the User ID that can access the HMC.
Password	Enter the password for the User ID that can access the HMC.

4. Verify all fields are correct and press Enter.

Changing an existing Storage Agent

To change an existing Storage Agent, complete the following steps:

1. From the command line enter, `smit sysmirror`.
2. From the SMIT interface, select **Cluster Applications and Resources > Resources > Configure DS8000 Global Mirror Resources > Configure Storage Agents > Change/Show Storage Agent** and press Enter.
3. From the list select the name of the Storage Agent that you want to change and press Enter.
4. Specify the changes you want to make in the fields.
5. Verify the changes are correct and press Enter.

Removing a Storage Agent

To remove a Storage Agent, complete the following steps:

1. From the command line enter, `smit sysmirror`.
2. From the SMIT interface, select **Cluster Applications and Resources > Resources > Configure DS8000 Global Mirror Resources > Configure Storage Agents > Remove Storage Agent** and press Enter.
3. From the list of Storage Agents select the name of the Storage Agent you want to remove and press Enter.

4. Press Enter to confirm the selected Storage Agent is the one you want to remove.

Configure a Storage System

A Storage System is a generic name given by PowerHA® SystemMirror® for an entity such as a DS8700 Storage Unit.

When using Global Mirror, you must associate one Storage Agent with each Storage System. You must provide the IBM® DS8700 system identifier for the Storage System. For example, IBM®.2107-75ABTV1 is a storage identifier for a DS8000® Storage System.

Adding a Storage System

To add a Storage System, complete the following steps:

1. From the command line enter, `smit sysmirror`.
2. From the SMIT interface, select **Cluster Applications and Resources > Resources > Configure DS8000 Global Mirror Resources > Configure Storage Systems > Add a Storage System** and press Enter.
3. Complete the following fields:

<i>Adding a Storage System fields</i>	
Fields	Values
Storage System Name	Enter the PowerHA® SystemMirror® name for the storage system. This name must be unique within the cluster definition.
Site Association	Enter the PowerHA® SystemMirror® site name.
Vendor Specific Identification	Enter the vendor specified unique identifies for the storage system.
WWNN	The World Wide Node Name for this storage system. Every storage system that supports FCP has a unique name.
Storage Agent Name	Press F4 to select the Name(s) of the Storage Agent that manage this Storage System from a list.

4. Verify all fields are correct and press Enter.

Changing a Storage System

To change a Storage System, complete the following steps:

1. From the command line enter, `smit sysmirror`.
2. From the SMIT interface, select **Cluster Applications and Resources > Resources > Configure DS8000 Global Mirror Resources > Configure Storage Systems > Change/Show Storage System** and press Enter.
3. From the list select the name of the Storage System that you want to change and press Enter.
4. Specify the changes you want to make in the fields.
5. Verify the changes are correct and press Enter.

Removing a Storage System

To remove a Storage System, complete the following steps:

1. From the command line enter, `smit sysmirror`.
2. From the SMIT interface, select **Cluster Applications and Resources > Resources > Configure DS8000 Global Mirror Resources > Configure Storage Systems > Remove Storage System** and press Enter.
3. From the list of Storage Systems select the name of the Storage System you want to remove and press Enter.
4. Press Enter to confirm the selected Storage System is the one you want to remove.

Configuring Mirror Group

A Mirror Group is a generic name given by PowerHA® SystemMirror® for a logical collection of volumes that have to be mirrored to another Storage System that resides on a remote site. A Global Mirror Session represents a Mirror Group.

Adding a Mirror Group

To add a Mirror Group, complete the following steps:

1. From the command line enter, `smit sysmirror`.
2. From the SMIT interface, select **Cluster Applications and Resources > Resources > Configure DS8000 Global Mirror Resources > Configure Mirror Groups > Add Mirror Group**, and press Enter.
3. Complete the following fields:

<i>Adding a Mirror Group fields</i>	
Fields	Values
Mirror Group Name	Enter the PowerHA® SystemMirror® for the replicated resource. This is what will be included into a PowerHA® SystemMirror® Resource Group.
Storage System Names	From the list select the Storage Systems on the production site that have data volumes forming this Mirror Group.
Vendor Specific Identifier	Enter the Global Mirror Sessions Identifier.
Maximum Coordination Time	Enter the maximum amount of time (in seconds) that the DS8700 can hold an I/O issued from the host for the purpose of Global Mirror operations.
Maximum Drain Time	Enter the maximum amount of time (in seconds) allowed for data to be drained before failing the current consistency group.
Consistency Group Interval Time	Enter the time (in seconds) to wait between the formation of two consistency groups.

4. Verify all fields are correct and press Enter.

Changing a Mirror Group

To change a mirror group, complete the following steps:

1. From the command line enter, `smit sysmirror`.
2. From the SMIT interface, select **Cluster Applications and Resources > Resources > Configure DS8000 Global Mirror Resources > Configure Mirror Groups > Configure Mirror Groups > Change/Show Mirror Group**, and press Enter.
3. From the list select the name of the Mirror Group that you want to change and press Enter.
4. Specify the changes you want to make in the fields.
5. Verify the changes are correct and press Enter.

Removing a Mirror Group

To remove a Mirror Group, complete the following steps:

1. From the command line enter, `smit sysmirror`.
2. From the SMIT interface, select **Cluster Applications and Resources > Resources > Configure DS8000 Global Mirror Resources > Configure Mirror Groups > Configure Mirror Groups > Remove Mirror Group**, and press Enter.
3. From the list of Mirror Groups select the name of the Mirror Group you want to remove and press Enter.
4. Press Enter to confirm the selected Mirror Group is the one you want to remove.

Configuring Resource Groups

After configuring the Mirror Group, include the Mirror Group into the desired PowerHA® SystemMirror® Resource Group.

Configure resource groups following the procedures in the [Configuring PowerHA® SystemMirror® resource groups](#) topic. Make sure that you understand site support for PowerHA® SystemMirror® resource groups.

When you configure a resource group you must meet the following requirements:

- The site policy is set to Prefer Primary Site or Online on Either Site.
- A startup policy other than Online on All Available Nodes is specified.
- The Resource Group Processing Ordering is set to Serial.

To add a Global Mirror replicated resource to a resource group, complete the following steps:

1. From the command line enter, `smit sysmirror`.
2. From the SMIT interface, select **Cluster Applications and Resources > Resources Groups > Change/Show Resources and Attributes for a Resource Group**, and press Enter.
3. Enter the following:
 - a. The name of the mirror groups in the Global Mirror Replicated Resources field.
 - b. The name of the volume groups associated with the individual Global Mirror replicated resources in the Volume Groups field.

The volume group names must be listed in the same order as the DS8700 mirror group names.

4. Verify and synchronize the cluster.

SVC replication resources

These topics present information for planning, installing, and configuring an PowerHA® SystemMirror® Enterprise Edition Metro Mirror for SAN Volume Controller (SVC) cluster.

SVC virtualizes storage, assists in managing Peer-to-Peer Remote Copy (PPRC) services on TotalStorage® systems, and simplifies how you can integrate PPRC replicated resources into an PowerHA® SystemMirror® configuration. SVC management allows SVC PPRC resources to be managed by PowerHA® SystemMirror® with a minimum of user configuration.

Overview of SVC management

The IBM® TotalStorage® SAN Volume Controller (SVC) is a virtualization appliance solution that maps virtualized volumes visible to hosts and applications to physical volumes on storage devices.

It is an in-band implementation that minimizes the dependency on unique hardware and software, telecopying the storage functions expected in a SAN environment from the storage subsystems and managing storage resources. In addition, it provides advanced copy services for data migration and business continuity similar to the Peer to Peer Remote Copy function in the ESS and DS4000® (formerly named the FASTT) Storage subsystems. Since the copy services operate on the virtual volumes, dramatically simpler replication configurations can be created using the SAN Volume Controller, rather than replicating each physical volume in the managed storage pool.

PowerHA® SystemMirror® Enterprise Edition for Metro Mirror with SVC management enhances PPRC's ability to provide a fully automated, highly available disaster recovery management solution by taking advantage of SVC's ability to provide virtual disks derived from varied disk subsystems. The PowerHA® SystemMirror® interface is designed so that once the basic SVC environment is configured, PPRC relationships are created automatically; no additional access to the SVC interface is needed.

PowerHA® SystemMirror® configuration information required for SVC PPRC resource management is simple; it is easily acquired from the existing SVC configuration.

The integration of PowerHA® SystemMirror® and SVC PPRC provides:

- PowerHA® SystemMirror® management of SVC PPRC for automatic failover/reintegration of SVC PPRC-protected virtual disks between PowerHA® SystemMirror® sites
- Support for user-defined policy-based resource groups
- Support for the following Inter-Site Management Policies for resource groups:
 - Prefer Primary Site
 - Online on Either Site.
- Support for the Subsystem Device Driver (SDD) (for the 2145 (SVC) system)
- Support for cluster verification and synchronization
- Flexible user-customizable resource group movement policies
- Limited support for C-SPOC (PowerHA® SystemMirror® Single Point of Control utility)
- Automatic failover/reintegration of server nodes attached to SVC-provided virtual disk pairs within sites
- Management of automatic failover of PPRC-protected volume groups between nodes within a site
- SVC management to switch the SVC PPRC relationships over so that the backup site can take control of the PowerHA® SystemMirror®-managed resource groups from the primary site in case of a site failure
- SVC PPRC Command Line Interface or GUI to manually manage SVC PPRC consistency groups and relationships.

It is important to understand how the terms primary and auxiliary are used in the PowerHA® SystemMirror® Enterprise Edition Metro Mirror for the SVC environment. In general, primary and auxiliary refer to the SVC virtual disks that are on either end of a SVC PPRC link. Primary and secondary refer to the PowerHA® SystemMirror® sites that host the resource groups that manage the SVC PPRC replicated resources that contain those SVC PPRC links.

The terms primary and auxiliary can also refer to the SVC clusters themselves. In general, the primary SVC cluster is connected to the PowerHA® SystemMirror® production site, and the auxiliary SVC cluster is connected to the PowerHA® SystemMirror® backup and recovery site.

Planning for SVC management

You should be familiar with the planning tasks for PowerHA® SystemMirror®.

- PowerHA® SystemMirror® sites have been planned.
- Basic SVC and SVC PPRC support has been completely configured. Refer to the SVC documentation on how to install and configure SVC and SVC PPRC support.
- SVC Clusters and native PPRC support on those SVC Clusters has already been configured. To plan for SVC-managed PPRC replicated resources in an PowerHA® SystemMirror® cluster, complete the following planning tasks:
 - Identify the SVC Clusters and which PowerHA® SystemMirror® sites they belong to.
 - Identify SVC relationships.
 - Identify the volume groups associated with SVC virtual disks (VDisks) to be used in the relationships.
 - Identify SVC consistency groups.
 - Identify the relationships to be used in the consistency groups.
Note: The term 'Consistency Group Name' can be, in this instance, interchanged with the term 'PPRC Replicated Resource Name.' In the context of SVC, they are the same.
- Plan which resource groups will contain the SVC-managed PPRC replicated resources.

Limitations and restrictions for PowerHA® SystemMirror® Enterprise Edition Metro Mirror for SVC

The current release of PowerHA® SystemMirror® Enterprise Edition SVC PPRC has some limitations and restrictions.

These include:

- Although SVC Host Name Aliases are arbitrary, for PowerHA® SystemMirror® Enterprise Edition SVC PPRC, they must match the node names. This ensures that the **ssh** commands used to execute SVC tasks are completed for the correct nodes.
- PowerHA® SystemMirror® uses commands to communicate with the SVC PPRC cluster that require **ssh**. Therefore, some version of **ssh** must be installed and configured on all SVC PPRC cluster nodes.
- Resource Groups to be managed by PowerHA® SystemMirror® cannot contain volume groups with both SVC PPRC-protected and non SVC-PPRC protected disks.
For example:
 - VALID: RG1 contains VG1 and VG2, both PPRC-protected disks.
 - INVALID: RG2 contains VG3 and VG4, VG3 is PPRC-protected, and VG4 is not.
 - INVALID: RG3 contains VG5, which includes both PPRC-protected and non-protected disks within the same volume group.
- The Host Aliases on your SVC Clusters must match the node names (short names) used in PowerHA® SystemMirror® to define each cluster node.
- SVC PPRC does not support concurrent resource groups that are configured with a startup policy of Online on All Available Nodes (OOAN).
- You cannot use C-SPOC for the following LVM operations to configure nodes at the remote site (that contain the target volumes):
 - Creating or extending a volume group
 - Operations that require nodes at the target site to write to the target volumes (for example, changing filesystem size, changing mount point, adding LVM mirrors) cause an error message in CSPOC. This includes functions such as changing filesystem size, changing mount points, and adding LVM mirrors. However, nodes on the same site as the source volumes can successfully perform these tasks. The changes are subsequently propagated to the other site via lazy update.

For C-SPOC operations to work on all other LVM operations, it is highly recommended that you perform all C-SPOC operations when the cluster is active on all PowerHA® SystemMirror® nodes and the underlying SVC consistency groups are in a *consistent_synchronized* state.

Sample configuration for SVC management

You can set up a mutual recovery configuration in which each site acts as a production site with the other site acting as an associated backup site.

PowerHA® SystemMirror® support the SVC feature Change Volumes for Global Mirror. If you want to use this feature, you must configure and manage it outside of PowerHA® SystemMirror®. For more information about the SVC feature Change Volumes for Global Mirror, see [Low Bandwidth Global Mirror \(SVC & V7000\)](#).

Implementing a mutual recovery configuration requires:

- Two PowerHA® SystemMirror® sites (the same as a single recovery configuration)
- At least two resource groups.

Basic SVC PPRC cluster configuration and operation

The infrastructure we consider here is a wide-area cluster system whereby the set of local cluster nodes are directly attached to local storage systems but with no access to the geographically separated remote storage systems.

In a typical setup with AIX® server nodes in a wide-area PowerHA® SystemMirror® cluster configuration, the server nodes at each PowerHA® SystemMirror® cluster site are connected directly to the local SVC Clusters. Two or more inter-cluster links are required between the two SVC clusters. The PowerHA® SystemMirror® nodes at the same geographical site access the same shared volume groups, but the nodes at each site access them from different physical volumes.

The SVC PPRC maintains separate identical local copies of the application data on two separate back end storage subsystems. Several virtual disks are mirrored via primary SVC cluster PPRC from the primary site to the standby site. When node or site failure occurs, access to disk resources are *not* passed from one node to another. Instead, all highly available applications are restarted at the standby site using data copy on secondary volumes. Under normal operation, the application is active on a server at the production site, and all updates to the application data are automatically replicated to the backup disk subsystem by the SVC PPRC framework. PPRC protects the backup copy from inadvertent modification. When a total site failure occurs, the application is restarted on a backup server at the remote site.

Prior to restarting the application, which, in this context, means *not* only the application that end users interact with, but also all dependent data base software or other middleware, PowerHA® SystemMirror® Enterprise Edition initiates SVC actions to ensure that the backup disk volumes are in the appropriate state to allow application access.

PowerHA® SystemMirror® support for geographic clustering is based on the concept of *replicated resources*. These are defined as a resource type that has a primary and secondary instance, corresponding to the source and target of data copies that are replicated across two locations. In SVC PPRC, a PPRC consistency group, comprised of a list of SVC relationships defined to SVC and associated with an PowerHA® SystemMirror® geographically separated cluster definition and defined to PowerHA® SystemMirror® is referred to as an *SVC PPRC replicated resource*. The definition for an SVC PPRC replicated resource contains the VDisk name and the PowerHA® SystemMirror® volume groups associated with the VDisk pairs. SVC recognizes which volumes mirror each other for each PPRC replicated resource.

In the example shown here, the PowerHA® SystemMirror® production site includes:

- Server A and Server B
- The ESS labeled Primary ESS
- A SAN Volume Controller connected by Fibre Channel PPRC links to the SVC at the secondary site and two connections to the Primary ESS.

The PowerHA® SystemMirror® recovery site includes:

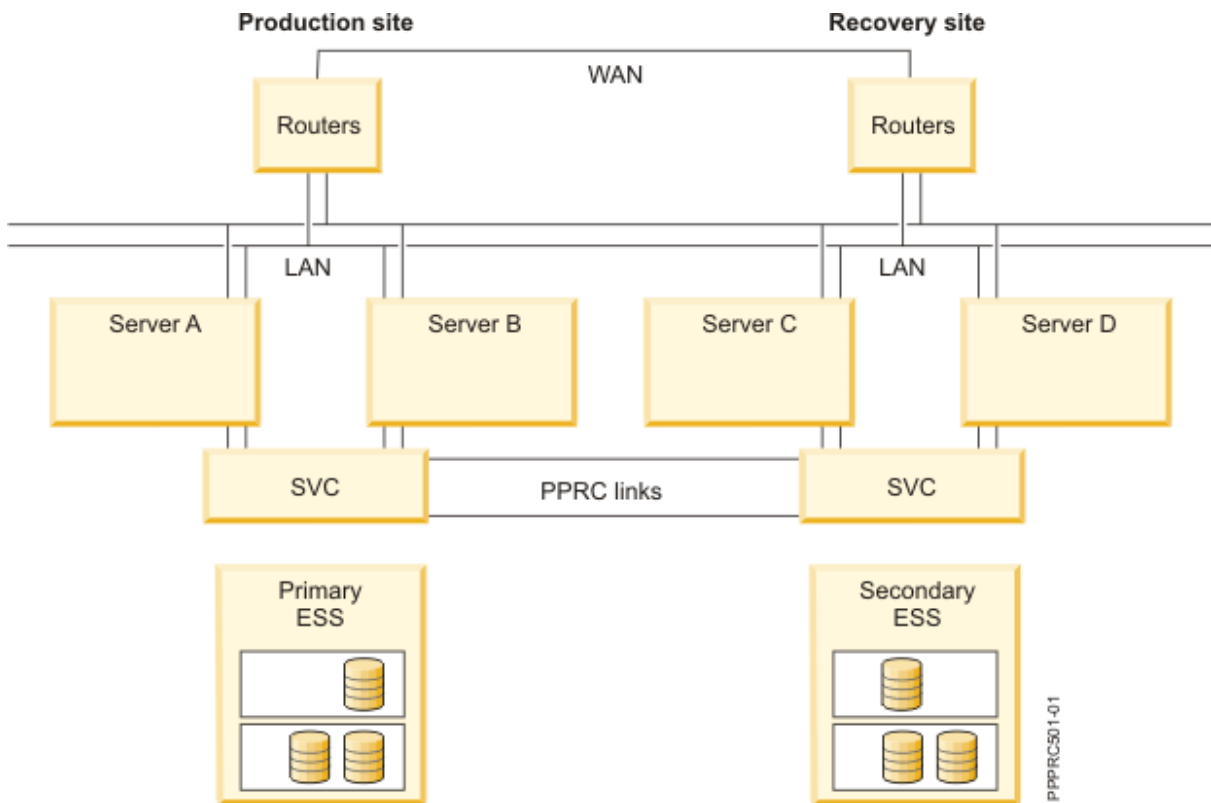
- Server C and Server D
- The ESS labeled Secondary ESS
- A SAN Volume Controller connected by Fibre Channel PPRC links to the SVC at the primary site and two connections to the Primary ESS.

The PowerHA® SystemMirror® resource group contains:

- The four server nodes
- IP networks to connect the server nodes
- One or more shared volume groups
- The PPRC replicated resource associated with the volumes in the volume group.

The configuration also includes point-to-point networks for heartbeating to connect the cluster nodes.

Note that all nodes also require connections to the SVC. At least one storage controller is connected to each SVC cluster via a SCSI or Fiber Channel connection.



Example PowerHA® SystemMirror® Enterprise Edition Metro Mirror for SVC mutual takeover configuration

This example lists the information for a two-node SVC PPRC cluster configured for mutual takeover.

```

SVC Clusters
-----
svc_C11, svc_C12

SVC PPRC Relationships Configuration
-----
Relationship Name = sample_rel1
Primary Vdisk info = volume_id1@svc_C11
Auxiliary Vdisk info = volume_id2@svc_C12

SVC PPRC Relationships Configuration
-----
Relationship Name = sample_rel2
Primary Vdisk info = volume_id3@svc_C11
Auxiliary Vdisk info = volume_id4@svc_C12
SVC PPRC-Replicated Resource Configuration
for Resource Group RG1
-----

SVC PPRC Consistency Group Name = CG1
Primary SVC Cluster Name = svc_C11
Auxiliary SVC Cluster Name = svc_C12
List of Relationships = sample_rel1, sample_rel2

**
SVC PPRC Relationships Configuration
for Consistency Group CG3
-----
Relationship Name = sample_rel3
Primary Vdisk info = volume_id5@svc_C12
Auxiliary Vdisk info = volume_id6@svc_C11

SVC PPRC-Replicated Resource Configuration
for Resource Group RG2
-----
SVC PPRC Consistency Group Name = CG1
Primary SVC Cluster Name = svc_C12

```

```
Auxiliary SVC Cluster Name = svc_C11
List of Relationships = sample_re13
```

From this point, resource groups RG1 and RG2 are configured to include the SVC PPRC-Replicated Resources CG1 and CG2 respectively.

Note: Since nodes located at each site treat the pairs differently in some situations, using PPRC pairs in resource groups that have no site preference could lead to unpredictable results.

Planning PPRC-SVC management support

Before you configure SVC management option, you configure SVC itself to manage PPRC mirrored volumes.

Verify that your SVC configuration meets the following requirements:

- Makes a consistent point-in-time copy of data available at all times
- Has virtual disks (vDisks) made available from both SVC Clusters to be used in the PowerHA® SystemMirror® configuration.

For information about configuring SVC, see the appropriate SVC documentation. See also the Limitations and Restrictions for PowerHA® SystemMirror® Enterprise Edition Metro Mirror for SVC for any naming convention restrictions.

Setting up volume groups and filesets on SVC PPRC-protected disks

As part of planning your PPRC SVC environment, decide what vDisks will be used to support volume groups and file systems on the PowerHA® SystemMirror® cluster nodes.

Regardless of whether you simply plan the vDisks at this time or perform the steps listed below, you must complete these steps prior to the initial PowerHA® SystemMirror® SVC PPRC verification step. That step creates the final SVC PPRC relationships.

This procedure assumes that the hdisks and corresponding vpaths made available to your nodes are visible at those nodes. If they are *not*, and you can verify that SVC has been configured correctly to make the virtual disks (vDisks) available, reboot the node and run **cfgmgr** to make the disks viewable.

Discover AIX® vpaths associated with SVC vDisks

You can use SVC vDisks to discover AIX® vpaths.

1. Select vDisks from your SVC configuration to be used in SVC PPRC relationships (that will be grouped in consistency groups) that will be managed by PowerHA® SystemMirror®.
2. Based on the vDisks you have chosen to use in a given SVC PPRC relationship, find out which vpaths and hdisks correspond.
 - a. On an PowerHA® SystemMirror® node, run:

```
lsdev -Ccdisk | more
```

From this output, you can discover which hdisks are associated with the SAN Volume Controller:

```
HAnode1> lspv
...
hdisk32 Available 81-08-02 SAN Volume Controller Device
hdisk33 Available 81-08-02 SAN Volume Controller Device
hdisk34 Available 81-08-02 SAN Volume Controller Device
hdisk35 Available 81-08-02 SAN Volume Controller Device
hdisk36 Available 81-08-02 SAN Volume Controller Device
hdisk37 Available 81-08-02 SAN Volume Controller Device
vpath0 Available Data Path Optimizer Pseudo Device Driver
vpath1 Available Data Path Optimizer Pseudo Device Driver
```

vpath2	Available	Data Path Optimizer	Pseudo Device Driver
vpath3	Available	Data Path Optimizer	Pseudo Device Driver

In the example listing above, hdisk32 - hdisk37 are all SAN Volume Controller Devices.

- b. Then run one of the following commands:
 'lsvpcfg | grep <hdisk>', to list the vpath for a specific hdisk
 'lsvpcfg | grep vpath', to list information for all vpaths.

SVC-related vpaths have an output similar to this:

```
[HAnode1][usr/es/sbin/cluster/svcpprc/utils]> lsvpcfg | grep vpath
vpath2 (Avail pv ) 6005076801898009980000000000012A = hdisk4 (Avail )
hdisk11 (Avail ) hdisk24 (Avail ) hdisk31 (Avail )

vpath3 (Avail pv ) 60050768018980099800000000000129 = hdisk5 (Avail )
hdisk12 (Avail ) hdisk25 (Avail ) hdisk32 (Avail )
```

- c. Using the volume ID number from **lsvpcfg**:

```
(ie: 6005076801898009980000000000012A )
```

Run the following command to get any associated vDisk value. (If there is doubt as to the output, run the command without the trailing 'grep' and 'cut'):

```
ssh admin@<SVC Cluster IP address> svcinfo lshostvdiskmap -delim : |
grep <SVC volume ID value> | cut -f5 -d":"
```

Example:

```
HAnode1> ssh admin@9.22.22.22 svcinfo lshostvdiskmap -delim : |grep
6005076801898009980000000000012C |cut -f5 -d":"

vDisk1
HAnode1>
```

The value returned should be the name given to the vDisk during its creation on the SVC system.

Setting up volume group and filesystem for PowerHA® SystemMirror® management

On the vpath(s) that corresponds to the master vDisk(s) for a given SVC PPRC Relationship, set up the volume group and filesystem to be managed by PowerHA® SystemMirror®. Ensure that the Volume Major Number for the volume group can be used on all PowerHA® SystemMirror® cluster nodes, and that the physical volume name for the file system can also be used across all PowerHA® SystemMirror® cluster nodes.

1. Use the **lvlstmajor** command on each cluster node to display the available Volume Major Numbers.
2. Vary off (using **varyoffvg**) the newly created volume group and import it to all nodes in the local PowerHA® SystemMirror® site.
3. Create a temporary SVC PPRC relationship to copy the volume group/file set information to the auxiliary vDisk. Run the following commands. (Refer to the SVC CLI documentation for more details.)

```
ssh admin@<master SVC Cluster IP> svctask mkrcrelationship -master
<vDisk_name>-aux <vDisk_name> -cluster <Aux Cluster name> -name
<relationship name>

ssh admin@<master SVC Cluster IP> svctask starttrrelationship
<relationship name>
```

At this point, wait until the relationship moves from inconsistent_copying to consistent_synchronised state. Check the state by running:

```
ssh admin@master<master SVC Cluster IP> svcinfo lsrelationship
[relationship name]
```

4. Once the SVC PPRC relationship has completed copying, delete the relationship: `ssh admin@<master SVC Cluster IP> svctask rmrcrelationship <relationship name>`
This step is necessary in order for the next LVM operations to complete successfully.
5. Using SMIT or the command line on the Backup PowerHA® SystemMirror® site (the site that is connected to the Auxiliary SVC Cluster,) import the volume group(s) created in step 2c.
6. Ensure that on ALL cluster nodes, the AUTO VARYON feature for volume groups is set to NO: On each node run: `chvg -a 'n' -Q 'y' <volume group name>`
PowerHA® SystemMirror® will attempt to auto-correct this during verification but will *not* be able to do so in the case of remote PPRC. (If done at this point, it will save you time later.)

Result

At this point, the volume groups and file systems necessary to configure PowerHA® SystemMirror® have been created.

Installing PowerHA® SystemMirror® Enterprise Edition Metro Mirror for SVC

This section describes how to install PowerHA® SystemMirror® Enterprise Edition Metro Mirror for SVC. You must be logged in as root to perform installation tasks.

Installing prerequisite software

Before installing PowerHA® SystemMirror® Enterprise Edition Metro Mirror for SVC, the following software must be installed on cluster nodes:

1. openssh version 4.6.1 or higher (+ license), for access to the SVC interface.
2. Base PowerHA® SystemMirror® Enterprise Edition at required minimum level (cluster.es.server.rte) or higher.
3. Make sure the recommended versions of the following filesets are installed:
 - a. `devices.fcp.disk.ibm.rte (1.0.0.0)`
 - b. `devices.sddpcm.**.rte ()`

All these required filesets might vary depend upon SVC SDD version. Hence, refer to the SVC documentation and website for the prereq filesets information.

It may *not* be clear during installation of the **sdd** filesets that this first fileset is necessary, but it is critical in order to ensure that failover works correctly.

Result

Installing PowerHA® SystemMirror® Enterprise Edition Metro Mirror for SVC filesets

To install SVC filesets for PowerHA® SystemMirror® Enterprise Edition, complete the following steps:

1. Insert the PowerHA® SystemMirror® Enterprise Edition for AIX® DVD into the DVD drive.
2. From the command line, enter `smit installp`.
3. From the System Management Interface Tool (SMIT) interface, select **Install and Update Software** > **Install Software**, and press Enter.
4. Select the following filesets and press Enter.
 - `cluster.es.svcpprc`
 - `cluster.es.genxd`

Configuring PowerHA® SystemMirror® Enterprise Edition Metro Mirror for SVC

After you install SVC following the procedures outlined in its documentation, follow the steps to configure the cluster.

These steps include:

1. Check SVC and PowerHA® SystemMirror® configuration requirements.
2. Configure PowerHA® SystemMirror® and SVC PPRC.
3. Verify the PowerHA® SystemMirror® and SVC PPRC configuration.
4. Synchronize the cluster.
5. Start the cluster.

SVC and PowerHA® SystemMirror® configuration requirements

Make sure the configuration requirements are met.

These requirements include:

1. Configure **ssh** for each node that communicates with the SVC clusters. Public or private key pairs must be created on the nodes for the configured user and distributed to the SAN Volume Controller (SVC) clusters for PowerHA® SystemMirror® Enterprise Edition Metro Mirror for SVC to function. If the consistency groups and SVC Peer-to-Peer Remote Copy (PPRC) relationships are already created in the SVC storage, a user with privileges to perform all PPRC operations can access the SVC storage. Otherwise, the user must have administrator privilege to access the SVC storage.

The SVC PPRC replicated methods that are used with PowerHA® SystemMirror® Enterprise Edition rely heavily on remote commands executed by using **ssh**. It must be possible, for example, to remotely run a command via **ssh** on all cluster nodes that are part of a resource group that contains a PPRC replicated resource. Private/public key pairs must be installed on both the nodes that access the SVC clusters, and the SVC clusters themselves in order for the SVC PPRC commands to work correctly.

See the SVC documentation for directions on how to complete this task.

2. Make sure the Host Aliases on your SVC Clusters match the node names (short names) used in PowerHA® SystemMirror®. To get a list of the current host to vDisk maps on the primary and auxiliary SVC clusters, use the SVC CLI commands:

```
ssh admin@<ip_for_PRIMARY_SVC_Cluster> svcinfo lshostvdiskmap | more
ssh admin@<ip_for_AUXILIARY_SVC_Cluster> svcinfo lshostvdiskmap | more
```

Result

Check the host names listed there against your PowerHA® SystemMirror® node names. If they differ, then refer to the SVC documentation on how to change the names to match, either via the SVC CLI or GUI interface.

Adding an SVC cluster configuration

You can use the SMIT interface to add an SAN Volume Controller (SVC) cluster configuration.

To add an SVC cluster configuration, complete the following steps:

1. From the command line, enter `smit sysmirror`.
2. In SMIT, select **Cluster Applications and Resources > Resources > Configure SVC PPRC-Replicated Resources > SVC Clusters Definition to PowerHA® SystemMirror® > Add an SVC Cluster**, and press Enter:
3. Enter field values as follows:

Add an SVC Cluster fields	
Field	Value
SVC Cluster Name	Enter the same name used by SVC. This name cannot be more than 20 alphanumeric characters and underscores.
SVC Cluster Role	Select Primary or Auxiliary .
PowerHA® SystemMirror® site	Select the PowerHA® SystemMirror® site associated with this SVC cluster.
SVC Cluster IP Address	IP address of this cluster.

Field	Value
SVC Cluster Second IP Address	If you are using SVC 5.1, or later, you can enter a second IP address for this cluster. If the IP address that is specified in the SVC Cluster IP Address field is not available, the IP address you specify in this field is used.
Remote SVC Partner	Name of the IBM® SAN Volume Controller (SVC) cluster that is hosting vDisks from the other side of the SVC Peer-to-Peer Remote Copy (PPRC) link.
User Name	User name that is used to access IBM SVC cluster by using the Secured Shell (SSH) protocol. The user name must be configured to use an SSH key-based authentication.

4. Press Enter.

Adding SVC PPRC relationships

After you add a SAN Volume Controller (SVC) cluster, you can add an SVC Peer-to-Peer Remote Copy (PPRC) relationship. In PowerHA® SystemMirror® Version 7.2.4, or later, adding SVC PPRC relationships is optional. If you are not specifying the SVC PPRC relationship details, PowerHA® SystemMirror® automatically discovers the relationship details from the SVC storage for the configured SVC PPRC consistency groups.

To add an SVC PPRC relationship, complete the following steps:

1. From the command line, enter `smit sysmirror`.
2. Select **Cluster Applications and Resources > Resources > Configure SVC PPRC-Replicated Resources > SVC PPRC Relationships Definition > Add an SVC PPRC Relationship**, and press Enter.
3. Enter field values as follows:

<i>Add an SVC PPRC Relationship field</i>	
Field	Value
Relationship Name	The name that is used by both SVC and PowerHA® SystemMirror® for configuration of SVC PPRC Relationships. Use no more than 20 alphanumeric characters and underscores.
Master VDisk Info	The Primary and Auxiliary VDisk names use the format <code>vdisk_name@svc_cluster_name</code> . See the section (PPRC-protected disks) for steps on how to discover what virtual disks (VDisks) correspond to the vpaths used for your configuration. The Master VDisk is the disk that resides at the primary site for the resource group the SVC PPRC Relationship will be a part of.
Auxiliary VDisk Info	The Auxiliary VDisk is the disk at the backup site for the resource group the SVC PPRC Relationship will be a part of.

4. Press Enter when you have finished the definition.
5. Create as many SVC PPRC Relationships that are required to manage all the vpaths used for volume groups.

Adding SVC PPRC replicated resources

You can use the SMIT interface to add SVC PPRC replicated resources.

To add an SVC PPRC replicated resource, complete the following steps:

1. From the command line, enter `smit sysmirror`.
2. From the SMIT interface, select **Cluster Applications and Resources > Resources > Configure SVC PPRC-Replicated Resources > SVC PPRC-Replicated Resource Configuration > Add an SVC PPRC Resource**, and press Enter.
3. Enter field values as follows:

<i>Add an SVC PPRC Resource fields</i>	
Field	Value
SVC PPRC Consistency Group Name	The name to be used by SVC and also to be used in the resource group configuration. Use no more than 20 alphanumeric characters and underscores.
Master SVC Cluster Name	Name of the Primary cluster is the SVC Cluster connected to the PowerHA® SystemMirror® Primary Site
Auxiliary SVC Cluster Name	Name of the SVC Cluster connected to the PowerHA® SystemMirror® Backup/Recovery Site
List of Relationships	List of names of the SVC PPRC relationships
Copy Type	Global Mirror processing provides a long-distance remote copy solution across two sites by using asynchronous technology. To use Global Mirroring, select GLOBAL . Metro Mirror functions offer a synchronous remote copy option that constantly updates a secondary copy of a volume to match changes that are made to a source volume. To use Metro Mirroring, select METRO .
PowerHA® SystemMirror® Recovery Action	If a site fallover occurs, specify the recovery policy that is used by PowerHA® SystemMirror®. Select AUTOMATED to automatically use the recovery policy. Select MANUAL to use manual intervention when a site fallover occurs.

4. Press Enter.

Adding SVC PPRC consistency groups to PowerHA® SystemMirror® resource groups

Add the SVC PPRC Consistency Groups created in the previous step to PowerHA® SystemMirror® resource groups.

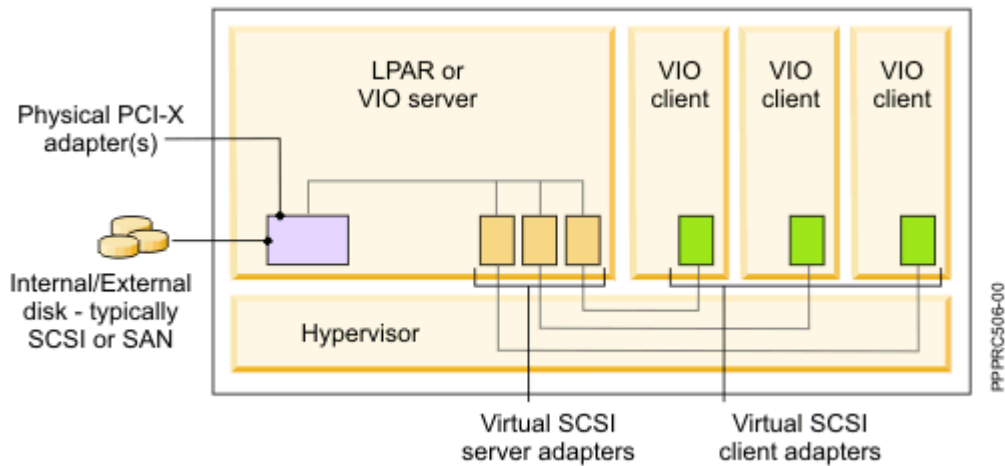
1. From the command line, enter `smit sysmirror`.
2. From the SMIT interface, select **Cluster Applications and Resources > Resource Groups**, and press Enter. Select **Change/Show** or **Add a Resource Group** depending on whether you are working with an existing resource group or creating a new resource group.

The **SVC PPRC Replicated Resources** entry appears at the bottom of the page in SMIT. This entry is a picklist that displays the resource names created in the previous step. Make sure that the volume groups selected on the Resource Group configuration screen match the vDisks used for the SVC PPRC Replicated Resource.

3. Press Enter.

Setting up PowerHA® SystemMirror® Enterprise Edition for SVC Metro Mirror on AIX® Virtual I/O clients

Virtual I/O Server allows a machine to be divided into LPARs, with each LPAR running a different OS image allowing the sharing of physical resources between the LPARs including virtual SCSI and virtual networking. The VIO Server owns real PCI adapters (Ethernet, SCSI, or SAN), but lets other LPARs share them remotely using the built-in Hypervisor services. These other LPARs are called Virtual I/O client partitions or VIO clients. And because they don't need real physical disks or real physical Ethernet adapters to run, they can be created quickly and inexpensively.



As an example, in previous Figure 1, the VIO Server has a few disks that could be SCSI or fiber channel storage area network (SAN) disks. The VIO clients use the VIO client device driver just as they would a regular local device disk to communicate with the matching server VIO device driver, then the VIO Server actually does the disk transfers on behalf of the VIO client.

When PowerHA® SystemMirror® Enterprise Edition for SVC Metro Mirror is configured on VIO clients the SVC clusters are not directly attached to the VIO clients, therefore the normally SCSI query commands can not be used to extract the necessary SVC vdisk information. Even though no special configuration steps are required to define the SVC PPRC resource to PowerHA® SystemMirror® Enterprise Edition, it is required that the following procedures have been performed already before verification can succeed. If this has not been done perform the following actions before proceeding to perform the usual SVC PPRC configuration steps.

These steps assume that the disks subsystems have already been physically attached to the SVC clusters. All the necessary SCSI server adapters have been created on the servers and virtual client SCSI adapters have been mapped to client partitions.

1. On the SVC clusters, perform these actions:
 - a. Identify the managed disks MDisks using the **svcinfolsmdisk** command
 - b. Identify or create the managed disk groups MDiskgrp using **svcinfolsmdiskgrp** or the **svctask mkmdiskgrp** command
 - c. Identify or Create the Virtual Disks using the **svcinfolsvdisk** or the **svctask mkvdisk** command
 - d. Map the VDisks to the VIO Servers as hosts using the **svctask mkvdiskhostmap**.
2. On the VIO servers, perform these actions
To get access to the regular AIX® command line interface on the VIO server it is recommended that you run **oem_setup_env**.
 - a. Run **cfgmgr**
 - b. Use **odmget -q "id=unique_id" CuAt** to identify the hdisks/vpaths mapped to the SVC vdisks on the servers
 - c. Select the disk to export by running **lsdev** to show the virtual SCSI server adapters that could be used for mapping with a physical disk
 - d. Run the **mkvdev** command using the appropriate hdisk#s respectively to create the Virtual Target Device (this commands maps the LUNs to the virtual I/O clients:


```
$ mkvdev -vdev hdisk# -vadapter vhost# -dev vhdisk#
```
 - e. Map the VDisks to the VIO Servers as hosts using the **svctask mkvdiskhostmap**.
3. On the VIO clients
After the **mkvdev** command runs successfully on the VIO server, the LUNs will be exported to the VIO clients.
 - a. Run **lsdev -Cc disk** to identify the LUN info on the client
 - b. Run **cl_vpath_to_vdisk** command to identify the SVC vdisk – LUN mapping on the VIO clients.

Verifying and synchronizing the PowerHA® SystemMirror® Enterprise Edition cluster

At this point, provided you have configured all the other aspects of your resource group (volume groups being the most important), verify the cluster and SVC PPRC configuration.

Take the following steps:

1. Verify the PowerHA® SystemMirror® cluster configuration (using SMIT).
2. Verify the SVC PPRC configuration (using CLI).
3. Synchronize the PowerHA® SystemMirror® Cluster configuration (using SMIT).

Verifying the PowerHA® SystemMirror® cluster configuration

You can verify the PowerHA® SystemMirror® cluster configuration using SMIT.

Be careful to run verification only. You must use the **Cluster Applications and Resources** SMIT path to get the correct menu.

1. From the command line, enter `smit sysmirror`.
2. Select **Cluster Applications and Resources > Resources > Verify and Synchronize Cluster Configuration** and press Enter.

Enter the field values as follows:

Verification and Synchronization fields	
Field	Value
Verify Synchronize or Both	Select Verify .
Automatically correct errors found during verification?	Select Yes .
Force synchronization if verification fails?	Select No (the default.)
Verify changes only?	Select no to run all verification checks that apply to the current cluster configuration.
Logging	Standard is the default. You can also select Verbose . Verification messages are logged to <code>/var/hacmp/clverify/clverify.log</code> .

3. Press Enter. The output from the verification is displayed in the SMIT Command Status window.
4. If you receive error messages, make the necessary changes and run the verification procedure again. You may see Warnings if the configuration has a limitation on its availability; for example, only one interface per node per network is configured.

Verifying the SVC PPRC configuration

You can verify the SVC PPRC configuration using the CLI.

To run the SVC PPRC verification, execute the following command:

```
/usr/es/sbin/cluster/svcpprc/utlis/cl_verify_svcpprc_config
```

If any configuration errors appear, go back into the PowerHA® SystemMirror® SMIT panels and correct the errors, then re-run this script.

All the SVC PPRC Relationships and Consistency Groups named in the previous steps are created during this step. If there were any errors during execution of this script, the SVC PPRC consistency groups and relationship may *not* be created.

To verify that the configuration exists, run the following two commands against either of the SVC clusters:

```
ssh admin@<SVC Cluster IP> svcinfo lsrcrelationship
ssh admin@<SVC Cluster IP> svcinfo lsrconsistgrp
```

Synchronizing the PowerHA® SystemMirror® cluster configuration

You can propagate the new SVC PPRC configuration information (and possibly PowerHA® SystemMirror® site information) across your PowerHA® SystemMirror® cluster.

Follow these steps:

1. From the command line, enter `smit sysmirror`.
2. Select **Cluster Applications and Resources > Verify can Synchronize Cluster Configuration**, and press Enter.

Enter the field values as follows:

Verification and Synchronization fields	
Field	Value
Verify Synchronize or Both	Select Synchronize .
Automatically correct errors found during verification?	Select No (the default.)
Force synchronization if verification fails?	Select No (the default.)
Verify changes only?	Select no to run all verification checks that apply to the current cluster configuration.
Logging	Standard is the default. You can also select Verbose . Verification messages are logged to <code>/var/hacmp/clverify/clverify.log</code> .

3. Press Enter. The cluster is synchronized. The output is displayed in the SMIT Command Status window.

Starting the PowerHA® SystemMirror® Enterprise Edition SVC PPRC cluster

Your PowerHA® SystemMirror® cluster is now ready to be started.

PowerHA® SystemMirror® will manage the SVC PPRC replicated resources you have just defined. PowerHA® SystemMirror® starts these SVC PPRC resources automatically, so there should be no need to manually manage the SVC resources from this point, under normal PowerHA® SystemMirror® cluster operation.

Changing an PowerHA® SystemMirror® Enterprise Edition for Metro Mirror SVC configuration

Using SMIT, you can change, show, or remove cluster configurations, PPRC relationships, and resources.

Be sure to update the PowerHA® SystemMirror® resource group information and synchronize the cluster after making changes.

Changing an SVC cluster configuration

You can use the SMIT interface to the change an SAN Volume Controller (SVC) cluster configuration.

To change the SVC cluster, complete the following steps:

1. From the command line, enter `smit sysmirror`.
2. In SMIT, select **Cluster Applications and Resources > Resources > Configure SVC-PPRC Replicated Resources > SVC Clusters Definition to PowerHA® SystemMirror® > Change/Show an SVC Cluster**, and press Enter:
3. Select the SVC cluster to change from the picklist, and press Enter.
4. Enter field values as follows:

<i>Change/Show an SVC Cluster fields</i>	
Field	Value
SVC Cluster Name	The current name of the SVC cluster.
New SVC Cluster Name	Enter the name of the SVC Cluster Use no more than 20 characters.
SVC Cluster Role	Select Primary or Auxiliary .
PowerHA® SystemMirror® site	Select the PowerHA® SystemMirror® site associated with this SVC cluster. The site must already be defined in PowerHA® SystemMirror®.
SVC Cluster IP Address	IP address of this cluster.
SVC Cluster Second IP Address	If you are using SVC 5.1, or later, you can enter a second IP address for this cluster. If the IP address that is specified in the SVC Cluster IP Address field is not available, the IP address you specify in this field is used.
Remote SVC Partner	The name of the SVC Cluster that is hosting vDisks from the other side of the SVC PPRC link.
User Name	User name that is used to access IBM SVC cluster by using the Secured Shell (SSH) protocol. The user name must be configured to use an SSH key-based authentication.

5. Press Enter.

Removing an SVC cluster configuration

You can remove an SVC cluster configuration.

To remove the SVC cluster definition:

1. From the command line, enter `smit sysmirror`.
2. In SMIT, select **Cluster Applications and Resources > Resources > Configure SVC-PPRC Replicated Resources > SVC Clusters Definition to PowerHA® SystemMirror® > Remove an SVC Cluster**, and press Enter:
3. Select the SVC cluster to remove from the picklist and press Enter.
4. Confirm that you want to remove this SVC cluster definition by pressing Enter again.

Changing SVC PPRC relationships

After you define an SVC PPRC relationship, you can change it.

To change the SVC PPRC relationships:

1. From the command line, enter `smit sysmirror`.
2. Select **Cluster Applications and Resources > Resources > Configure SVC-PPRC Replicated Resources > SVC PPRC Relationships Definition > Change/Show an SVC PPRC Relationship**, and press Enter.
3. Select the SVC PPRC relationship to change and press Enter.
4. Enter field values as follows:

<i>Change/Show an SVC PPRC Relationship fields</i>	
Field	Value
Relationship Name	The current relationship name displays here.
New Relationship Name	The new name to be used by both SVC and PowerHA® SystemMirror® for configuration of SVC PPRC Relationships. Use no more than 20 characters.
Primary VDisk Info	The Primary and Auxiliary VDisk names use the format: <code>vdisk_name@svc_cluster_name</code> See the section Setting up volume groups and filesets on SVC PPRC-protected disks for steps on how to discover what virtual disks (VDisks) correspond to the vpaths used for your

Field	Value
	configuration. The Primary VDisk is the disk that resides at the primary site for the resource group the SVC PPRC Relationship will be a part of.
Auxiliary VDisk Info	The Auxiliary VDisk is the disk at the backup site for the resource group the SVC PPRC Relationship will be a part of.

5. Press Enter when you have finished the definition.
6. Create as many SVC PPRC Relationships as will be necessary to manage all the vpaths used for volume groups that will be managed by HAMCP.

Removing SVC PPRC relationships

You can remove an SVC PPRC relationship.

To remove the SVC PPRC relationships:

1. From the command line, enter `smit sysmirror`.
2. Select **Cluster Applications and Resources > Resources > Configure SVC-PPRC Replicated Resources > SVC PPRC Relationships Definition > Remove an SVC PPRC Relationship**, and press Enter.
3. Select the SVC PPRC relationship to remove and press Enter.
4. Confirm that you want to remove this SVC PPRC relationship by pressing Enter again.

Changing SVC PPRC resources

You can use the SMIT interface to change existing SVC PPRC resources.

To change the SVC PPRC resources, complete the following steps:

1. From the command line, enter `smit sysmirror`.
2. Select **Cluster Applications and Resources > Resources > Configure SVC-PPRC Replicated Resources > SVC PPRC-Replicated Resource Configuration > Change/Show an SVC PPRC Resource**, and press Enter.
3. Select the SVC PPRC resource to change and press Enter.
4. Enter field values as follows:

<i>Change/Show an SVC PPRC Resource fields</i>	
Field	Value
SVC PPRC Consistency Group Name	The current name that is used by SVC and also used in the Resource Group Configuration.
New SVC PPRC Consistency Group Name	The new name that is used by SVC and also to be used in the Resource Group Configuration.
Master SVC Cluster Name	The name of the Primary cluster is the SVC Cluster connected to the PowerHA® SystemMirror® Primary Site.
Auxiliary SVC Cluster Name	The name of the SVC Cluster connected to the PowerHA® SystemMirror® Backup/Recovery Site.
List of Relationships	A list of names of the SVC PPRC relationships.
Copy Type	Global Mirror processing provides a long-distance remote copy solution across two sites by using asynchronous technology. To use Global Mirroring, select GLOBAL . Metro Mirror functions offer a synchronous remote copy option that constantly updates a secondary copy of a volume to match changes that are made to a source volume. To use Metro Mirroring, select METRO .
PowerHA® SystemMirror® Recovery Action	If a site failover occurs, specify the recovery policy that is used by PowerHA® SystemMirror®. Select AUTOMATED to

Field	Value
	automatically use the recovery policy. Select MANUAL to use manual intervention when a site fallover occurs.

5. Press Enter.

Removing SVC PPRC resources

You can remove SVC PPRC resources.

To remove the SVC PPRC resources:

1. From the command line, enter `smit sysmirror`.
2. Select **Cluster Applications and Resources > Resources > Configure SVC-PPRC Replicated Resources > SVC PPRC-Replicated Resource Configuration > Remove an SVC PPRC Resource**, and press Enter.
3. Select the SVC PPRC resource to remove and press Enter.
4. Press Enter again to confirm this action.

Troubleshooting PowerHA® SystemMirror® Enterprise Edition for Metro Mirror for SVC

These topics provide information that might help you with troubleshooting SVC PPRC clusters.

SVC PPRC states

The following PPRC Volume states are possible, for either Consistency Groups or PPRC Relationships:

- `inconsistent_stopped`
In this state, the Primary is accessible for read and write I/O but the Secondary is *not* accessible for either. A copy process needs to be started to make the Secondary Consistent.
- `inconsistent_copying`
In this state, the Primary is accessible for read and write I/O but the Secondary is *not* accessible for either. This state is entered after a Start command is issued to an InconsistentStopped Relationship or Consistency Group. It is also entered when a Forced Start is issued to an Idling or ConsistentStopped Relationship or Consistency Group. A background copy process will run which will copy data from the Primary to the Secondary virtual disk.
- `consistent_stopped`
In this state, the Secondary contains a Consistent image but it might be out-of-date with respect to the Primary.
- `consistent_synchronized`
In this state, the Primary VDisk is accessible for read and write I/O. The Secondary VDisk is accessible for read-only I/O. Writes that are sent to the Primary VDisk are sent to both Primary and Secondary VDIs. Either good completion must be received for both writes, or the write must be failed to the host, or a state transition out of ConsistentSynchronized must take place before a write is completed to the host.
- `idling`
Both Primary and Auxiliary disks are operating in the Primary role. Consequently both are accessible for write I/O. In this state the Relationship or Consistency Group will accept a 'Start' command. Remote Copy maintains a record of regions on each disk which have received write I/O while Idling. This will be used to determine what areas need to be copied following a 'Start' command.
- `idling_disconnected`
The Virtual Disk(s) in this half of the Relationship or Consistency Group are all in the Primary role and will accept read or write I/O. No configuration activity is possible (except for deletes or stops) until the Relationship becomes connected again. At that point, the Relationship will transition to a Connected state.
- `inconsistent_disconnected`

The Virtual Disks in this half of the Relationship or Consistency Group are all in the Secondary role and will *not* accept read or write I/O. No configuration activity except for deletes is permitted until the Relationship becomes connected again.

- `consistent_disconnected`
The VDisks in this half of the Relationship or Consistency Group are all in the Secondary role and will accept read I/O but *not* write I/O. This state is entered from `ConsistentSynchronized` or `ConsistentStopped` when the Secondary side of a Relationship becomes disconnected.
- `empty`
This state only applies to Consistency Groups. It is the state of a Consistency Group which has no Relationships and hence no other state information to show. It is entered when a Consistency Group is first created. It is exited when the first Relationship is added to the Consistency Group at which point the state of the Relationship becomes the state of the Consistency Group.

Viewing the state of consistency groups

To view the current state of the consistency groups PowerHA® SystemMirror® is managing (via the configured resource groups) execute the following command on any of the PowerHA® SystemMirror® cluster nodes:

```
ssh admin@<SVC cluster IP address> svcinfo lsrrconsistgrp
```

where <SVC Cluster IP Address> is for either SVC cluster being used by PowerHA® SystemMirror®.

Viewing SVC clusters

You can use commands to view information about SVC clusters.

These commands are stored in the `/usr/es/sbin/cluster/svcpprc/cmds` directory.

cllssvc command

List SVC cluster information.

```
cllssvc [-n < svccluster_name >] [-c]
```

The command lists information about all SVC clusters in the PowerHA® SystemMirror® configuration or a specific SVC cluster. If no SVC is specified, all SVC clusters defined will be listed. If a specific SVC cluster is provided via the `-n` flag, information about this SVC only will be displayed. The `-c` flag displays information in a colon-delimited format.

Sample output

```
[/usr/es/sbin/cluster/svcpprc/cmds]> cllssvc
svc9A
svc78

[/usr/es/sbin/cluster/svcpprc/cmds]> cllssvc -n svc9A
svccluster_name svccluster_role sitename cluster_ip r_partner
svc9A AuxiliaryVancouver 9.114.230.93 svc78

[/usr/es/sbin/cluster/svcpprc/cmds]> cllssvc -n svc9A -c
#SVCNAME:ROLE:SITENAME:IPADDR:RPARTNER
svc9A:Auxiliary:Vancouver:9.114.230.93:svc78
```

cllssvcpprc command

List information about all SVC PPRC resources or a specific SVC PPRC resource.

```
cllssvcpprc [-n < svcpprc_consistencygrp >] [-c] [-a] [-h]
```

If no resource name is specified, the names of all PPRC resources defined will be listed. If the **-a** flag is provided, full information about all PPRC resources is displayed. If a specific resource is provided via the **-n** flag, information about this resource only will be displayed. The **-c** flag displays information in a colon-delimited format. The **-h** flag turns off the display of column headers.

Sample output

```
[/usr/es/sbin/cluster/svcpprc/cmds]> cllssvcpprc<
HASVC1

[/usr/es/sbin/cluster/svcpprc/cmds]> cllssvcpprc -n HASVC1
svcpprc_consistencygrp PrimaryCluster AuxiliaryCluster relationships
HASVC1          svc78 svc9A svcrel1

[/usr/es/sbin/cluster/svcpprc/cmds]> cllssvcpprc -n HASVC1 -ca
#NAME:PRIMARY:AUXILIARY:RELATIONSHIPS
HASVC1:svc78:svc9A:svcrel1
```

cllsrelationship command

List information about all SVC PPRC relationships or a specific PPRC relationship.

```
cllsrelationship [-n <relationship_name>] [-c] [-a] [-h]
```

If no resource name is specified, the names of all PPRC resources defined will be listed. If the **-a** flag is provided, full information about all PPRC relationship is displayed. If a specific relationship is provided via the **-n** flag, information about this relationship only will be displayed. The **-c** flag displays information in a colon-delimited format. The **-h** flag turns off the display of column headers.

Sample output

```
[/usr/es/sbin/cluster/svcpprc/cmds]> cllsrelationship
svcrel1

[/usr/es/sbin/cluster/svcpprc/cmds]> cllsrelationship -n svcrel1
relationship_name PrimaryVdisk_info AuxiliaryVdisk_info
svcrel1 c48f1rp06_075@svc78 c48f2rp08_095@svc9A

[/usr/es/sbin/cluster/svcpprc/cmds]> cllsrelationship -n svcrel1 -c
#RELATIONSHIP:PRIMARYVDISK:AUXVDISK
svcrel1:c48f1rp06_075@svc78:c48f2rp08_095@svc9A
```

cl_verify_svcpprc_config command

Verifies the SVC definition in the PowerHA® SystemMirror® configuration.

After the successful verification of the SVC configurations, it establishes all the SVC relationships defined to PowerHA® SystemMirror® on the SVC clusters and adds them to the corresponding consistency groups.

Configuring PowerHA® SystemMirror® Enterprise Edition for SVC policy based replication

After you install SVC, follow the steps to configure the enterprise edition for policy based replication.

These steps include:

1. Check the SAN Volume Controller (SVC) policy and PowerHA® SystemMirror® configuration requirements.
2. Configure PowerHA® SystemMirror® and SVC Policy based replication.
3. Verify the PowerHA® SystemMirror® and SVC Policy based replication
4. Synchronize the cluster.
5. Start the cluster.

- Ensure that you follow the SVC and PowerHA® SystemMirror® configuration requirements. For more information, see [“SVC and PowerHA SystemMirror configuration requirements” on page 82.](#)
- You can use the SMIT interface to add an SVC policy configuration. For more information, see [“Adding an SVC policy configuration” on page 93.](#)
- After you have configured all the other requirements of the resource group, verify the cluster and the SVC policy configuration.

Adding an SVC policy configuration

You can use the SMIT interface to add an SVC policy configuration.

You can use the SMIT interface to add an SAN Volume Controller (SVC) Policy configuration. The policy based configuration is available in the PowerHA® SystemMirror® Enterprise Version 7.2.9, or later. It supports asynchronous replication between two systems and provides a better consistency between all the volumes and clusters in a particular volume group.

To add an SVC policy configuration, complete the following steps:

1. From the command line, run the `smit hacmp` command.
2. In SMIT, select **Cluster Applications and Resources > Resources > Configure SVC PPRC-Replicated Resources > SVC Policy based Replication Configuration > Add an SVC Policy based replication Resource**, and press Enter:
3. Enter field values as follows:

<i>Add an SVC policy fields</i>	
Field	Value
SVC Replication Policy	Enter the name for the SVC policy.
Policy Replication Volume Group Name	Enter the volume group configured in storage.
Primary SVC Cluster Name	Select the name associated with the SVC cluster.
Auxiliary SVC Cluster Name	Select the secondary name associated with the SVC cluster.
Topology	The default value is 2-site-async-dr.

4. Press Enter.

SVC Policy Based Replication resources

These topics present information for planning, installing, and configuring PowerHA® SystemMirror Enterprise Edition for SAN Volume Controller (SVC) cluster. Policy based replication is the successor of the Remote Copy technology to provide replication services for IBM Storage FlashSystem, IBM SAN Volume Controller, and IBM Storage Virtualize version 8.5.2, and later.

SVC virtualizes storage, assists in managing Policy based replication (PBR) services on TotalStorage® systems, and simplifies how you can integrate PBR replicated resources into an PowerHA® SystemMirror® configuration. SVC management allows SVC PBR resources to be managed by PowerHA® SystemMirror® with a minimum of user configuration.

Overview of SVC policy-based replication management

The IBM® TotalStorage® SAN Volume Controller (SVC) is a virtualization appliance that maps virtualized volumes that are available to hosts and applications to physical volumes on the storage devices.

You can use policy-based replication (PBR) services with remote copy to maintain a copy of the data at a remote site that contains another IBM Storage Virtualize system. The remote site is known as the disaster recovery (DR) system. The site in which the hosts access the data is known as the production system.

You can get the information to configure and integrate PowerHA® SystemMirror® for the SVC PBR resource management from the existing SVC configuration resources.

The integration of PowerHA® SystemMirror® and SVC PBR provides:

- Automatic failover or integration of the SVC PBR protected virtual disks among various PowerHA® SystemMirror® sites
- Automatic failover or integration of server nodes that are attached to the SVC virtual disks in the PowerHA® SystemMirror® sites
- Automatic failover of volume groups secured by policy-based replication among nodes in a PowerHA® SystemMirror® site
- User-defined resource groups that are policy-based
- Inter-site management policies for the resource groups:
 - Prefer primary site
 - Online on either site
- Verification and synchronization for resource group clusters
- Flexible and customizable policies for resource groups
- PowerHA SystemMirror Single Point of Control utility support
- SVC management to switch the SVC PBR replication modes so that the backup site can manage resource groups from the primary site during a site failure
- SVC PBR command-line interface to manually manage the SVC PBR configuration
- SVC PBR disaster recovery support for site disks that are not accessible for the read operation
- **lsvdisk** storage command to verify that the disks are mapped to nodes by comparing the Universally Unique Identifier (UUID) strings.

The following terms are used in the PowerHA® SystemMirror® Enterprise Edition for AIX® for the SVC environment:

- Primary and auxiliary: Refer to the SVC virtual disks in the SVC PBR links.
- Primary and secondary: Refer to the PowerHA® SystemMirror® sites that host the resource groups. These resource groups manage the SVC policy-based replicated data in the SVC PBR links.

The terms primary and auxiliary are also associated with SVC clusters. The primary SVC cluster is connected to the PowerHA® SystemMirror® production site. The auxiliary SVC cluster is connected to the PowerHA® SystemMirror® backup or recovery site.

Configuring PowerHA® SystemMirror® Enterprise Edition policy based replication for SVC flash storage

You can configure PowerHA® SystemMirror® Enterprise Edition policy based replication for SVC flash storage.

These steps include:

1. Check the SAN Volume Controller (SVC) policy and PowerHA® SystemMirror® configuration requirements for flash storage.
2. Configure the PowerHA® SystemMirror® and SVC cluster configuration for flash storage.
3. Verify the PowerHA® SystemMirror® Enterprise Edition SVC cluster configuration.
4. Synchronize the cluster.
5. Start the cluster.

SVC and PowerHA® SystemMirror® configuration requirements for flash storage

Make sure the configuration requirements are met.

These requirements include:

1. Configure the **Secured Shell (ssh)** key for each node that communicates with the SVC clusters. Public or private key pairs must be created on the nodes for the configured user and distributed to the SAN Volume Controller (SVC) clusters for PowerHA® SystemMirror® Enterprise Edition for SVC to function.

If the production and recovery relationships are already created in the SVC storage, the user who has the privileges to perform all policy based replication (PBR) operations can access the SVC storage. Otherwise, the user needs to have administrator privilege rights to access the SVC storage.

The SVC policy based replication methods that are used with PowerHA® SystemMirror® Enterprise Edition depend a lot on remote commands executed by using the **ssh** key. It must be possible, for example, to remotely run a command via the **ssh** key on all cluster nodes that are part of a resource group that contains a PBR replicated resource. Private or public key pairs must be installed on both the nodes that access the SVC clusters. The SVC clusters can also operate themselves for the SVC PBR commands to work properly.

See the SVC documentation for directions on how to complete this task.

2. Make sure the Host Aliases on your SVC Clusters match the node names (short names) used in PowerHA® SystemMirror®. To get a list of the current host to vDisk maps on the primary and auxiliary SVC clusters, use the following SVC CLI commands:

```
ssh admin<ip_for_PRIMARY_SVC_Cluster> svcinfo lshostvdiskmap | more
ssh admin<ip_for_AUXILIARY_SVC_Cluster> svcinfo lshostvdiskmap | more
```

Result

Check the host names listed there against your PowerHA® SystemMirror® node names. If they differ, then refer to the SVC documentation on how to change the names to match, either via the SVC CLI or GUI interface.

Note: A policy defined as a replicated resource can also be attached to a resource group. While a storage policy can be associated with multiple SVC storage volume groups, all PowerHA® SystemMirror® volume groups added to this resource group should be mapped to a single SVC storage volume group.

Adding an SVC cluster configuration for Flash Storage

You can use the System Management Interface Tool (SMIT) interface to add a SAN Volume Controller (SVC) cluster configuration for flash storage.

To add an SVC cluster configuration for flash storage, complete the following steps:

1. From the command line, enter `smit sysmirror`.
2. In SMIT, select **Cluster Applications and Resources > Resources > Configure SVC PPRC-Replicated Resources > SVC Clusters Definition to PowerHA® SystemMirror® > Add an SVC Cluster**, and press Enter:
3. Enter field values as follows:

<i>Add an SVC Cluster fields</i>	
Field	Value
SVC Cluster Name	Enter the same name used by SVC. This name cannot be more than 20 alphanumeric characters and underscores.
SVC Cluster Role	Select Primary or Auxiliary .
PowerHA® SystemMirror® site	Select the PowerHA® SystemMirror® site associated with this SVC cluster.
SVC Cluster IP Address	IP address of this cluster.
SVC Cluster Second IP Address	If you are using SVC version 5.1, or later, you can enter a second IP address for this cluster. If the IP address that is specified in the SVC Cluster IP Address field is not available, the IP address you specify in this field is used.
Remote SVC Partner	Name of the IBM® SAN Volume Controller (SVC) cluster that is hosting vDisks from the other side of the SVC Policy Based Replication (PBR) link.
User Name	User name that is used to access IBM SVC cluster by using the Secured Shell (SSH) protocol. The user name must be configured to use an SSH key-based authentication.

4. Press Enter.

Adding SVC PBR replicated resources

You can use the SMIT interface to add an SVC replicated resource configuration.

You can use the SMIT interface to add an SAN Volume Controller (SVC) Policy configuration. The policy based configuration is available in the PowerHA® SystemMirror® version 7.2.9, or later. It supports replication between two systems that are not synchronized and provides better consistency across all volumes and clusters in a particular volume group.

1. From the command line, run the `smit hacmp` command.
2. From the SMIT interface, select **Cluster Applications and Resources > Resources > Configure SVC PPRC-Replicated Resources > SVC Policy based Replication Configuration > Add an SVC Policy based replication Resource**, and press Enter.
3. Enter field values as follows:

<i>Add an SVC policy fields</i>	
Field	Value
SVC Replication Policy	Name of the SVC policy. The policy name should match with the SVC storage policy.
Policy Replication Volume Group Name	The name of the volume group configured in storage.
Primary SVC Cluster Name	The name associated with the SVC cluster.
Auxiliary SVC Cluster Name	The secondary name associated with the SVC cluster.
Topology	Defines that two site locations are involved and asynchronous replication is performed for the disaster recovery process. Default value is 2-site-async-dr.

4. Press Enter.
A policy defined as a replicated resource can be attached to a resource group. The same storage policy that is available in the storage can be associated with multiple SVC storage volume groups. All PowerHA® SystemMirror® volume groups connected to a resource group should be mapped to a single SVC storage volume group.

Adding SVC PBR policy as replicated resource to PowerHA® SystemMirror® resource groups

After you add the SVC replicated resource configuration, add the SVC PBR resource to PowerHA® SystemMirror® resource groups.

1. From the command line, enter `smit sysmirror`.
2. From the SMIT interface, select **Cluster Applications and Resources > Resource Groups**, and press Enter.
3. Select **Change/Show**
The **SVC Policy Based Replication Resources** entry appears at the bottom of the page in SMIT. This entry is a pick list that displays the replicated resource names created in the previous step. Make sure that the volume groups selected on the Resource Group configuration screen match the vDisks used for the SVC storage volume group and associated with the PBR replication resource policy.
4. Press Enter.

You can only attach one replication resource policy to a resource group.

Verifying and synchronizing the PowerHA® SystemMirror® Enterprise Edition cluster for flash storage

After you have configured all the other aspects of your resource group (volume groups being the most important), verify the cluster and SVC PBR configuration.

Take the following steps:

1. Verify the PowerHA® SystemMirror® cluster configuration (using SMIT).
2. Verify the SVC PPRC configuration (using CLI).
3. Synchronize the PowerHA® SystemMirror® Cluster configuration (using SMIT).
 - You can verify the PowerHA SystemMirror cluster configuration using SMIT. For more information see, [“Verifying the PowerHA SystemMirror cluster configuration” on page 86](#).
 - You can verify the SVC PBR configuration using the CLI. For more information, see [“Verifying the SVC PBR configuration” on page 97](#).
 - You can propagate the new SVC PBR configuration information (and PowerHA® SystemMirror® site information) across your PowerHA® SystemMirror® cluster. For more information, see [“Synchronizing the PowerHA SystemMirror cluster configuration for flash storage” on page 98](#).

Verifying the PowerHA® SystemMirror® cluster configuration

You can verify the PowerHA® SystemMirror® cluster configuration using SMIT.

Be careful to run verification only. You must use the **Cluster Applications and Resources** SMIT path to get the correct menu.

1. From the command line, enter `smit sysmirror`.
2. Select **Cluster Applications and Resources > Resources > Verify and Synchronize Cluster Configuration** and press Enter.

Enter the field values as follows:

<i>Verification and Synchronization fields</i>	
Field	Value
Verify Synchronize or Both	Select Verify .
Automatically correct errors found during verification?	Select Yes .
Force synchronization if verification fails?	Select No (the default.)
Verify changes only?	Select no to run all verification checks that apply to the current cluster configuration.
Logging	Standard is the default. You can also select Verbose . Verification messages are logged to <code>/var/hacmp/clverify/clverify.log</code> .

3. Press Enter. The output from the verification is displayed in the SMIT Command Status window.
4. If you receive error messages, make the necessary changes and run the verification procedure again. You may see Warnings if the configuration has a limitation on its availability; for example, only one interface per node per network is configured.

Verifying the SVC PBR configuration

You can verify the SVC PBR configuration using the CLI.

To run the SVC PBR verification, execute the following command:

```
/usr/es/sbin/cluster/svcpprc/utils/cl_verify_svcpbr_config
```

If any configuration errors appear, go back into the PowerHA® SystemMirror® SMIT panels and correct the errors, then re-run this script.

Verification includes whether all PowerHA® SystemMirror® volume group disks are associated with single SVC volume group in storage for a specific resource group, and attached to the policy in the SVC storage. Make sure that the replicated resource is created from this policy. If there are any errors during execution of this script, the SVC PBR verification is not successful.

To verify that the configuration exists, run the following two commands against either of the SVC clusters:

```
ssh admin@<SVC Cluster IP> svcinfo lsvolumegroupreplication
ssh admin@<SVC Cluster IP> svcinfo lsreplicationpolicy
```

Synchronizing the PowerHA® SystemMirror® cluster configuration for flash storage

You can propagate the new SVC PBR configuration information (and possibly PowerHA® SystemMirror® site information) across your PowerHA® SystemMirror® cluster.

Follow these steps:

1. From the command line, enter `smit sysmirror`.
2. Select **Cluster Applications and Resources > Verify can Synchronize Cluster Configuration**, and press Enter.
Enter the field values as follows:

<i>Verification and Synchronization fields</i>	
Field	Value
Verify Synchronize or Both	Select Synchronize .
Automatically correct errors found during verification?	Select No (the default.)
Force synchronization if verification fails?	Select No (the default.)
Verify changes only?	Select no to run all verification checks that apply to the current cluster configuration.
Logging	Standard is the default. You can also select Verbose . Verification messages are logged to <code>/var/hacmp/clverify/clverify.log</code> .

3. Press Enter. The cluster is synchronized. The output is displayed in the SMIT Command Status window.

Starting the PowerHA® SystemMirror® Enterprise Edition SVC PBR cluster

Your PowerHA® SystemMirror® cluster is now ready to start.

PowerHA® SystemMirror® will manage the SVC PBR replicated resources you have just defined. PowerHA® SystemMirror® starts these SVC PBR resources automatically, so there should be no need to manually manage the SVC resources under the normal PowerHA® SystemMirror® cluster operation.

Changing an PowerHA® SystemMirror Enterprise Edition for SVC PBR configuration

Using SMIT, you can change, show, or remove cluster configurations, and PBR resources.

Be sure to update the PowerHA® SystemMirror® resource group information and synchronize the cluster after making changes.

Changing an SVC cluster configuration

You can use the SMIT interface to change an SAN Volume Controller (SVC) cluster configuration.

To change the SVC cluster, complete the following steps:

1. From the command line, enter `smit sysmirror`.
2. In SMIT, select **Cluster Applications and Resources > Resources > Configure SVC-PPRC Replicated Resources > SVC Clusters Definition to PowerHA® SystemMirror® > Change/Show an SVC Cluster**, and press Enter:
3. Select the SVC cluster to change from the picklist, and press Enter.
4. Enter field values as follows:

<i>Change/Show an SVC Cluster fields</i>	
Field	Value
SVC Cluster Name	The current name of the SVC cluster.
New SVC Cluster Name	Enter the name of the SVC Cluster Use no more than 20 characters.
SVC Cluster Role	Select Primary or Auxiliary .
PowerHA® SystemMirror® site	Select the PowerHA® SystemMirror® site associated with this SVC cluster. The site must already be defined in PowerHA® SystemMirror®.
SVC Cluster IP Address	IP address of this cluster.
SVC Cluster Second IP Address	If you are using SVC 5.1, or later, you can enter a second IP address for this cluster. If the IP address that is specified in the SVC Cluster IP Address field is not available, the IP address you specify in this field is used.
Remote SVC Partner	The name of the SVC Cluster that is hosting vDisks from the other side of the SVC PPRC link.
User Name	User name that is used to access IBM SVC cluster by using the Secured Shell (SSH) protocol. The user name must be configured to use an SSH key-based authentication.

5. Press Enter.

Removing an SVC cluster configuration

You can remove an SVC cluster configuration.

To remove the SVC cluster definition:

1. From the command line, enter `smit sysmirror`.
2. In SMIT, select **Cluster Applications and Resources > Resources > Configure SVC-PPRC Replicated Resources > SVC Clusters Definition to PowerHA® SystemMirror® > Remove an SVC Cluster**, and press Enter:
3. Select the SVC cluster to remove from the picklist and press Enter.
4. Confirm that you want to remove this SVC cluster definition by pressing Enter again.

Changing SVC PBR resources

You can use the SMIT interface to change existing SVC PBR resources.

To change the SVC PBR resources, complete the following steps:

1. From the command line, enter `smit sysmirror`.
2. Select **Cluster Applications and Resources > Resources > Configure SVC-PPRC Replicated Resources > SVC Policy based Replication Configuration > Change / Show SVC Policy based replication Resource**, and press Enter.
3. Select the SVC PBR resource to change and press Enter.

4. Enter field values as follows:

<i>Change/Show an SVC PBR Resource fields</i>	
Field	Value
SVC Replication Policy	The current Policy that is used by SVC and also used in the resource group configuration.
New SVC Replication Policy	The new policy name that is used by SVC and also to be used in the resource group configuration.
Policy replication Volume group Name	The new volume group name that is used by SVC and also to be used in the resource group configuration.
Primary SVC Cluster Name	The name of the Primary cluster is the SVC Cluster connected to the PowerHA® SystemMirror® primary site.
Auxiliary SVC Cluster Name	The name of the SVC Cluster connected to the PowerHA® SystemMirror® backup or recovery site.
Topology	The list of the topology related to the cluster in the site.

5. Press Enter.
A policy defined as a replicated resource can be attached to a resource group. While a storage policy can be associated with multiple SVC storage volume groups, all PowerHA® SystemMirror® volume groups connected to the resource group should be mapped to a single SVC storage volume group.

Removing SVC PBR resources

You can remove SVC Policy Based Replication (PBR) resources.

To remove the SVC PBR resources:

1. From the command line, enter `smit sysmirror`.
2. Select **Cluster Applications and Resources > Resources > Configure SVC-PPRC Replicated Resources > SVC Policy Based Replication Resource Configuration > Remove an SVC Policy Based Replication Resource**, and press Enter.
3. Select the SVC PBR resource to remove and press Enter.
4. Press Enter again to confirm this action.
Deleting the SVC PBR resource will not modify any configuration in the SVC storage site.

Troubleshooting PowerHA® SystemMirror® Enterprise Edition for PBR for SVC

These topics provide information that might help you with troubleshooting SAN Volume Controller (SVC) Policy Based Replication (PBR) clusters.

SVC PBR states

The following PBR Volume states are possible, for PBR Relationships:

- **Production**
In this state, the volumes in the group are accessible for host I/O, and configuration changes are allowed. This state is the source for replication to any recovery copies. By definition, a production volume group always contains an up-to-date version of application data for host access. Adding the resource group where the storage is in production will not modify any storage states.
- **Recovery**
In this state, the volumes in the group can receive only replication I/O, and volumes act as a target for replication. They cannot be used for host I/O. Most configuration changes are not permitted and must be performed on the production site. Adding the resource group where the storage is in recovery state depends on the link status. If the storage is not in suspended state, you should change the replication state to production and activate the resource group.
- **Independent**

In this state, the independent access is enabled for the volume group. During a shutdown, each part of the volume group is accessible for the host I/O. It permits configuration changes, and replication is suspended in this state. You can add the resource group where the storage is in Independent state to production and activate the resource group.

Viewing the state of volume groups

To view the current state of the PBR resource PowerHA® SystemMirror® in managing the volume group (through the configured resource groups), execute the following command on any of the PowerHA® SystemMirror® cluster nodes:

```
ssh admin@<SVC cluster IP address> svcinfo lsvolumegroupreplication vg_id
```

where <SVC Cluster IP Address> is for either SVC cluster being used by PowerHA® SystemMirror® and the `vg_id` is the SVC volumegroup ID.

Disaster recovery management with EMC SRDF replication

PowerHA® SystemMirror® implements disaster recovery with SRDF storage devices by using the EMC supplied Symmetrix command-line interface (SYMCLI).

PowerHA® SystemMirror® uses SYMCLI to automatically manage and monitor the status of SRDF devices and SRDF links. PowerHA® SystemMirror® can also use SYMCLI to automatically manage the switching of the direction for the SRDF relationships when a site failover occurs. If a primary site fails, the backup site is able to take control of the resource that PowerHA® SystemMirror® manages from the primary site.

PowerHA® SystemMirror® uses SYMCLI to discover the SRDF deployment and capture any missing replication-related information from the administrator. Using this information, PowerHA® SystemMirror® manages the replicated devices as part of the highly available environment for the applications and resource groups.

To enable the PowerHA® SystemMirror® for HADR of EMC storage devices, complete the following steps:

1. Plan the storage deployment and mirroring functions that are necessary for your environment. This step is related to the applications and middleware that are deployed in the environment, that is managed by PowerHA® SystemMirror®.
2. Use the EMC tools to configure and deploy the storage devices.

Note: Verify whether the gatekeeper devices are mapped and masked to a single host and are not shared for concurrent I/O across hosts. For more information about gatekeeper devices, go to the EMC Enterprise Single Sign-on web page (<http://support.emc.com>) and search for *emc25597*.

3. Use PowerHA® SystemMirror® interfaces to discover the storage devices that are deployed.
4. Define PowerHA® SystemMirror® policies for the applications or resource groups that are using the mirrored storage devices.

Note: Complete [steps 1](#) and [step 2](#) before you configure PowerHA® SystemMirror® for HADR management of EMC storage devices.

EMC SRDF concepts

PowerHA® SystemMirror® implements disaster recovery with SRDF storage devices by using an EMC-supplied Symmetrix command-line interface (SYMCLI).

PowerHA® SystemMirror® uses SYMCLI to automatically manage SRDF links and to switch the direction of the SRDF relationships when a site failover occurs. If a primary site fails, the backup site controls the PowerHA® SystemMirror® managed resource groups from the primary site.

PowerHA® SystemMirror® uses these interfaces to discover and integrate the replicated storage from an EMC storage device into the High Availability and Disaster Recovery (HADR) framework of PowerHA® SystemMirror®.

You can use this configuration to manage the HADR framework of your applications that are using the mirrored storage devices.

Integration of SRDF and PowerHA® SystemMirror® provides the following benefits:

- Support for either Prefer Primary Site or Online on Either Site inter-site management policies.
- Flexible user-customizable resource group policies.
- Support for cluster verification and synchronization.
- Limited support for the PowerHA® SystemMirror® Cluster Single Point Of Control (C-SPOC).
- Automatic failover and reintegration of server nodes that are attached to pairs of SRDF disk subsystems within sites and across sites.

EMC SRDF modes

PowerHA® SystemMirror® supports SRDF replicated resources in synchronous mode (SRDF/S) and in asynchronous mode (SRDF/A).

SRDF synchronous replication

In the synchronous mode, the EMC storage device responds to the host that issued a write operation to the source of the composite group after the EMC storage, which contains the target of the composite group, acknowledges that it has received and checked the data.

SRDF asynchronous replication

In the asynchronous mode, the EMC storage device provides a consistent point-in-time image on the target of the composite group, which is a short period of time behind the source of the composite group. Asynchronous mode is managed in sessions. Asynchronous mode transfers data in predefined timed cycles or in delta sets to ensure that data at the remote target of the composite group site is in the dependent write consistent state.

EMC SRDF consistency groups

All Symmetrix Remote Data Facility (SRDF) operations in PowerHA® SystemMirror® Enterprise Edition clusters are performed on a composite group and not on individual device groups. The composite group that is enabled for consistency is known as the *consistency group*.

The consistency groups operate in unison to preserve the integrity and dependent write consistency of a database distributed across multiple arrays. Consistency for a SRDF replicated resource is maintained at the composite group level on the EMC storage device.

The following table displays the valid states for a composite group.
Displays the valid states for a composite group.

<i>Valid composite group states</i>	
State	Description
SyncInProgress	A synchronization is currently in progress between the source of the composite group and the target of the composite group. There are existing invalid tracks between the two pairs and the logical link between both sides of an RDF pair is up.
Synchronized	This state is applicable only to the synchronous mirroring mode. The source of the composite group and the target of the composite group are currently in a synchronized state. The same content exists on the target of the composite group as the source of the composite group. There are no invalid tracks between the two pairs.
Split	The source of the composite group and the target of the composite group are currently Ready to their hosts, but the link is Not Ready or Write Disabled.

State	Description
Failed Over	The source of the composite group is currently Not Ready or Write Disabled and operations are failed over to the target of the composite group.
R1 Updated	The source of the composite group is currently Not Ready or Write Disabled to the host, there are no local invalid tracks on the source of the composite group side, and the link is Ready or Write Disabled.
R1 UpdInProg	The source of the composite group is currently Not Ready or Write Disabled to the host, there are invalid local tracks on the source side, and the link is Ready or Write Disabled.
Suspended	The RDF links have been suspended and are Not Ready or Write Disabled. If the source of the composite group is Ready while the links are suspended, any I/O will accumulate as invalid tracks owed to the target of the composite group.
Partitioned	The SYMAPI is currently unable to communicate through the corresponding RDF path to the remote Symmetrix.
Mixed	A composite SYMAPI device group that is in an RDF pair state. Different SRDF pair states exist within a device group.
Invalid	This is the default state when no other SRDF state applies. The combination of the source of the composite group, the target of the composite group, and RDF link states and statuses do not match any other pair state. This state might occur if there is a problem at the disk director level.
Consistent	The target source of the composite group for the SRDF/A capable devices are in a consistent state. Consistent state signifies the normal state of operation for device pairs that are operating in asynchronous mode. This state is applicable only to asynchronous mirroring.
Transmit Idle	This state is applicable only to asynchronous mirroring. The SRDF/A session cannot move data in the transmit cycle across the link because the link is not available.

EMC SRDF/Metro SmartDR configuration

PowerHA® SystemMirror® uses Symmetrix command-line interface (SYMCLI) to discover and integrate the replicated storage from an EMC storage device into the High Availability and Disaster Recovery (HADR) framework. You can use this configuration to manage the HADR framework for applications that are using mirrored storage devices. In addition to traditional EMC Symmetrix Remote Data Facility (SRDF) configuration, EMC provides SRDF/Metro SmartDR configuration that uses three SRDF/Metro storage devices. EMC SRDF/Metro SmartDR configuration is a two-region HADR framework, which integrates SRDF/Metro and SRDF/Async replicated resources.

The EMC SRDF/Metro SmartDR configuration adds the following capabilities to SRDF/Metro replicated resources:

- Integrates SRDF/Metro (Metro) and SRDF/Async (SRDF/A) replicated resources by enabling HADR framework for a SRDF/Metro session.
- Replicate data to a single SRDF/Metro storage device by coupling the SRDF/A sessions on each side of a SRDF/Metro resource pair.

EMC SRDF/Metro SmartDR configuration has the following considerations:

- Witness configuration is required for all EMC SRDF/Metro SmartDR configurations.
- Only a single SRDF/A session can send data to the disaster recovery site.
- If you switch an application data transfer to another site, ensure that a copy of dependent writes of the application data is maintained on the disaster recovery site.

Planning for SRDF replication

Before you use SRDF replication for disaster recovery, you must plan to implement SRDF management within an PowerHA® SystemMirror® environment.

Review the following information before you implement SDRF replication in your environment:

- Configure PowerHA® SystemMirror® Enterprise Edition 7.1.2, or later, across all sites.
- Identify SRDF mirroring relationships between storage devices.
- Configure the EMC Solution Enabler software on all nodes in the cluster.
- Export the consistency groups from nodes that are configured correctly and import the configuration to other nodes in the cluster.
- Configure the various mirroring relationships by using the EMC Solution Enabler software.

Note: Disks are grouped together as necessary for consistency device groups, or composite groups are configured, and SRDF pairs are created. You cannot use PowerHA® SystemMirror® to create SRDF pairs.

Review the following information to plan for SRDF replicated resources in an PowerHA® SystemMirror® cluster:

- Identify the composite groups or device groups that must be included in the PowerHA® SystemMirror® resource groups.
- Identify the SRDF pairs and establish the SRDF relationship in PowerHA® SystemMirror®.
- Identify the volume sets, associated volume groups, and the volume group associations with the applications and middleware software that you are deploying.
- Identify associations between applications and the resource groups.
- Identify which resource groups contain the SRDF replicated resources.

Limitations and restrictions of SRDF management

Before you implement SRDF management, you must be aware of the limitations and restrictions.

PowerHA® SystemMirror® Enterprise Edition for mirroring with SRDF has the following limitations and restrictions:

- Multi-Hop configuration is not supported.
- Mirroring to BCV devices is not supported.
- Concurrent RDF configuration is not supported.
- PowerHA® SystemMirror® Enterprise Edition does not support EMC storages that are indirectly connected to the AIX® host system through a host computer.
- PowerHA® SystemMirror® does not trap SNMP notification events for VMAX storage. If the SRDF link goes down when the cluster is up, and then the link is repaired, you must manually resync the pairs.
- The SRDF pair must be created outside of the cluster control. You must create the SRDF pairs before you start cluster services.
- EMC PowerPath disks are not supported for use as a tiebreaker disk.
- PowerHA® SystemMirror® Enterprise Edition does not correct the SRDF pairs or resynchronize the pairs if the pair states are in the invalid state.

Note: A pair in an invalid state might lead to data corruption if PowerHA® SystemMirror® Enterprise Edition tries to recover the state. A pair can be in an invalid state if there is a problem at the disk director level.

- Resource groups that are managed by PowerHA® SystemMirror® cannot contain volume groups with disks that are SRDF-protected and disks that are not SRDF-protected.
For example:
 - Correct example: RG1 contains VG1 and VG2, VG1 and VG2 are both SRDF-protected disks.
 - Incorrect example: RG2 contains VG3 and VG4, where VG3 is an SRDF-protected disk, and VG4 is not an SRDF-protected disk.
- Use C-SPOC for the following LVM operations to configure nodes at the remote site that contain the target volumes:
 - Creating a volume group.
 - Operations that require nodes at the target site to write to the target volumes (for example, changing file system size, changing the mount point, and adding LVM mirrors) cause an error message in C-SPOC. However, nodes on the same site as the source volumes can successfully perform these tasks. The changes are then propagated to the other site by using the lazy update function.

Note: For C-SPOC operations to work on all other LVM operations, you must perform all C-SPOC operations when the cluster is active on all PowerHA® SystemMirror® nodes, and the underlying SRDF pairs are in a synchronized state.

Example configuration for SRDF management

This example assumes that your environment consists of a mutual recovery configuration in which a production site stores different applications and services, and a recovery site acts as an associated backup site.

The following configuration is required to implement a mutual recovery configuration:

- Two PowerHA® SystemMirror® sites, which are the same as a single recovery configuration
- Two resource groups, with each resource group configured to use mirrored storage

In PowerHA® SystemMirror®, you can configure two sites and a resource group that includes nodes, SRDF-managed replicated resources, and the volume groups that are associated with the replicated resources.

The SYMCLI client is installed on the PowerHA® SystemMirror® hosts and on each of the PowerHA® SystemMirror® managed nodes. The EMC SRDF operations are completed by using the SYMCLI software on a AIX® server. You can use the SYMCLI software on the AIX® server to display EMC SRDF information and run SRDF operations from the command line or with a script file.

An SRDF replicated resource contains the EMC disk volume pairs (composite group) information. The PowerHA® SystemMirror® resource group definition includes the volume groups that are built on SRDF replicated volumes. PowerHA® SystemMirror® manages SRDF processing by dynamically running the SYMCLI commands.

Example for traditional SRDF configuration

The following image displays a typical implementation of two EMC VMAX storage with SRDF in a four-node PowerHA® SystemMirror® geographic cluster. The cluster consists of four IBM® Power Systems™ servers. Each storage device is connected to a node (Server A, Server B, Server C, and Server D) through SCSI or Fibre Channel links. The SRDF links are established between the Primary storage and the Secondary storage. The configuration also includes point-to-point networks for heartbeating to connect the cluster nodes together.

The order of the nodes in each resource group's node list indicates which site is considered as the production site and which site is considered as the backup site for that resource group.

In the following figure the configuration for resource group 1 is as follows:

Ordered site list

Production site, Recovery site

Ordered node list

Server A, Server B, Server C, and Server D

The data mirroring goes from Primary storage, which is owned by the Production site, to the Secondary storage, which is owned by the Secondary site.




Figure 1: Sample configuration of resource groups for SRDF-managed replicated resources

Example for EMC SRDF/Metro SmartDR configuration

Figure 2 displays a typical implementation of EMC SRDF/Metro SmartDR configuration with three EMC VMAX storage devices in a four-node PowerHA® SystemMirror® geographic cluster. The primary storage consists of two VMAX storage devices that are connected to production site nodes. The secondary storage is connected to recovery site nodes.

SRDF/Metro links are established between the primary storage and the secondary storage. The SRDF/Metro SmartDR configuration also includes point-to-point networks to connect the cluster nodes for heartbeating.

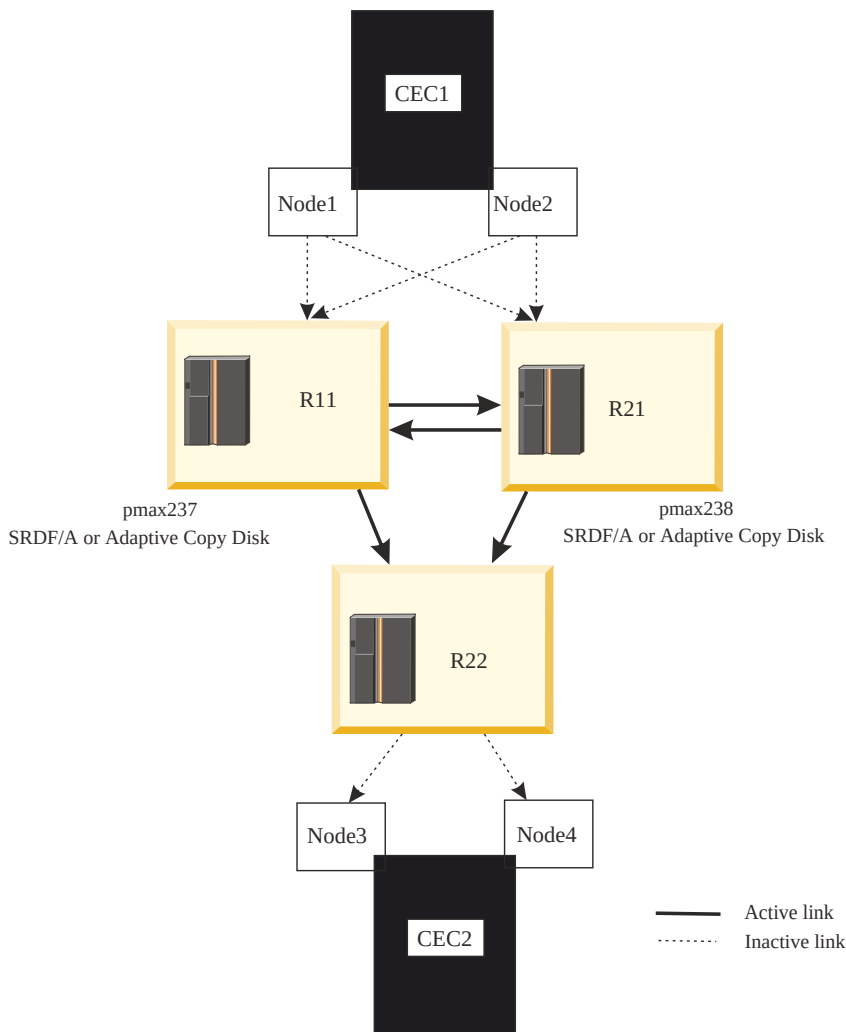


Figure 2: Sample configuration of resource groups for SRDF/Metro SmartDR replicated resources

Setting up volume groups and file systems on SRDF protected disks

As part of planning your PowerHA® SystemMirror® Enterprise Edition SRDF environment, you must decide what disks are used to support volume groups and file systems on the PowerHA® SystemMirror® cluster nodes.

Note: If EMC hdisks are not available to your nodes, you can verify that EMC storage is configured correctly to make the EMC disks available by rebooting the node and by running the `cfgmgr` command to make the disks viewable.

You cannot use PowerHA® SystemMirror® Enterprise Edition to create SRDF replication pairs. You must use the EMC Solution Enabler interfaces to create the SRDF pairs before you can use PowerHA® SystemMirror® Enterprise Edition.

To discover composite groups, devices groups, and AIX® disks that are associated with EMC storage devices, complete the following steps:

1. From the command line, enter the following command to list all host-visible Symmetrix devices:

```
# symdev list pd
```

Example of the output follows:

```
Symmetrix ID: 000190100304
```

Device Name	Directors	Device

		Cap

Sym	Physical	SA	:P	DA	:IT	Config	Attribute	Sts	(MB)
0026	/dev/rhdiskpower48	04A:1		16D:D5		RDF1+Mir	Grp'd	RW	4314
0027	/dev/rhdiskpower49	04A:1		01D:D4		RDF1+Mir	Grp'd	RW	4314
0028	/dev/rhdiskpower50	04A:1		01A:C8		RDF1+Mir	Grp'd	RW	4314
0029	/dev/rhdiskpower51	04A:1		16A:C9		RDF1+Mir	Grp'd	RW	4314
002A	/dev/rhdiskpower52	04A:1		01C:C6		RDF1+Mir	N/Grp'd	RW	4314
002B	/dev/rhdiskpower53	04A:1		16C:C7		RDF1+Mir	N/Grp'd	RW	4314
0F5C	/dev/rhdisk56	04A:1		01C:CA		2-Way Mir	N/Grp'd	VCM RW	6
0F9A	/dev/rhdiskpower0	04A:1		16C:CB		2-Way Mir	N/Grp'd	RW	6
0F9B	/dev/rhdiskpower1	04A:1		16A:CB		2-Way Mir	N/Grp'd	RW	6

- From the list, select the devices that stores the application data.
- Identify the AIX® disks for the Symmetrix device you selected by running the following command:

```
# powermt display dev=Symmetrix_device_name
Pseudo name=Symmetrix_device_name
```

In the following example, the Symmetrix device maps to hard disks hdisk105, hdisk50, hdisk120, and hdisk175 through the two I/O paths, fscsi1 and fscsi0:

```
Symmetrix ID=000190100304
Logical device ID=0026
state=alive; policy=Symm0pt; priority=0; queued-I/Os=0
=====
----- Host -----
### HW Path I/O Paths - Stor - -- I/O Path - -- Stats ---
-----
1 fscsi1 hdisk105 FA 13aB active alive 0 0
0 fscsi0 hdisk120 FA 4aB active alive 0 0
0 fscsi0 hdisk175 FA 13aB active alive 0 0
1 fscsi1 hdisk50 FA 4aB active alive 0 0
-----
```

- Create a volume group with the physical volumes of the Symmetrix device. For example, when you create these volume groups you are creating file systems and logical volumes. You must configure volume groups by using Logical Volume Manager (LVM) on the primary site.
- Create composite groups and device groups on the EMC storage and add Symmetrix devices to these groups.
- Establish full mirroring between the SRDF pairs so that the entire LVM metadata and other data are copied over to the mirrored volumes.
- Configure the SRDF replicated resources in PowerHA® SystemMirror® Enterprise Edition.
- Split the SRDF by using the **split** command to access the mirrored data on the secondary site.
- Import the volume groups on the secondary site by specifying the same volume groups name and major number that is used on the primary site for the mirrored storage device.
- Varyon the volume groups and mount the file systems.

Installing SRDF filesets

To perform any installation task, you must be logged in to the system as a root user.

Software prerequisite

Before you install SRDF filesets, the following software must be installed on all nodes in the cluster:

- Your environment must have one of the following version of the AIX® operating system installed:
 - IBM® AIX® 6.1 with Technology Level 8, or later
 - IBM® AIX® 7.1 with Technology Level 2, or later
- PowerHA® SystemMirror® Enterprise Edition at the same version.
- EMC Solution Enabler for AIX® software.
- SYMCLI software for AIX®. The fileset level for SYMCLI.SYMCLI.rte must be 7.0.0.0, or later.

- You must have SYMCLI Version 9.2.0.0, or later for EMC SRDF/Metro SmartDR configuration.
- Java™ SE Runtime Environment Version 1.6.0, or later.

Note: The default installation directory `/usr/java[0-9]*/jre/bin` is required to work with SRDF filesets.

Installing SRDF filesets

The following filesets are included on the PowerHA® SystemMirror® Enterprise Edition media and must be installed for SRDF management to work:

cluster.es.sr.cmds

PowerHA® SystemMirror® Enterprise Edition for SRDF commands

cluster.es.sr.rte

PowerHA® SystemMirror® Enterprise Edition for SRDF runtime commands

cluster.msg.en_US.sr

PowerHA® SystemMirror® Enterprise Edition for SRDF English messages

Complete the following steps to install the PowerHA® SystemMirror® SRDF filesets:

1. From the command line, enter `smit install`.
2. From the SMIT interface, select **Install and Update Software** › **Install Software** and press Enter.
3. Enter the path to the SRDF filesets that you want to install.

Note: All PowerHA® SystemMirror® SRDF filesets are installed in the `/usr/es/sbin/cluster/sr` directory.

Configuring SRDF replicated resources

After you install SYMCLI you must configure the cluster.

SRDF and PowerHA® SystemMirror® configuration requirements

You must verify that your environment meets specific requirements before you can configure PowerHA® SystemMirror® Enterprise Edition to work with the SRDF.

Before configuring PowerHA® SystemMirror® Enterprise Edition to work with SRDF, verify that the following criteria is implemented in your environment:

- The SRDF is configured and running on the storage systems.
- The SYMCLI software is installed.
- You have a good understanding of how PowerHA® SystemMirror® sites and SRDF replicated resources work.
- The base PowerHA® SystemMirror® Enterprise Edition for Metro Mirror and SRDF management filesets are installed on each node in the cluster.
- The PowerHA® SystemMirror® cluster has the following functions configured:
 - Nodes
 - Sites
 - Networks and network interfaces
 - Service labels and application monitors
 - Initial resource groups

Note: You can modify the attributes for a resource group later to accommodate SRDF replicated resources.

Adding an SRDF replicated resource

You must add an SRDF replicated resource before you can add a them to a resource group.

To add an SRDF replicated resource, complete the following steps:

Other:

- A device group configuration, without a composite group definition, is supported only in PowerHA® SystemMirror® 7.1.3, or later.
- An EMC SRDF/Metro SmartDR configuration is supported only in PowerHA® SystemMirror® 7.2.5 with Service Pack 1, or later.

1. From the command line, enter `smit sysmirror`.
2. From the SMIT interface, select **Cluster Applications and Resources > Resources > Configure EMC SRDF Replicated Resources > Add EMC SRDF Replicated Resource**, and press Enter.
3. Enter values for each of the following fields and press Enter:

EMC SRDF Replicated Resource Name

The name of the SRDF replicated resource that is defined to PowerHA® SystemMirror® as a composite group or a device group on the EMC storage device. You can choose an existing composite group or device group. You can also define a new group by listing the device groups that are included in the composite group or by listing the device group in the replicated resource.

EMC SRDF Mode { *ASYNC* / *SYNC* }

The SRDF mirroring mode. You can enter **SYNC** (synchronous mirroring) or **ASYNC** (asynchronous mirroring).

If you select **SmartDR** as the SRDF Group Type, the default mode **ASYNC** is used for SRDF mirroring. For more information about SRDF Group Type, see [EMC SRDF modes](#).

Device Groups

A list of device groups that is included in the preceding composite group or a device group in the replicated resource.

If you select **SmartDR** as the SRDF Group Type, do not enter any value for this field. For more information about SRDF Group Type, see [EMC SRDF consistency groups](#).

Recovery Action { *AUTO* / *MANUAL* }

If a site failover occurs, specify the disaster recovery policy that is used by the PowerHA® SystemMirror® resource group. Enter **Manual** if you want manual intervention when site failover occurs. Enter **Automated** if you do not want manual intervention when site failover occurs.

If you specify **Manual**, it does not indicate that a manual intervention is required for all failover scenarios. There are some conditions, such as cluster partition, in which an automatic failover from one site to another, can cause potential data divergence and integrity issues. If PowerHA® SystemMirror® detects the potential issue for such a case, and if the recovery action associated

with the mirror group is set to **Manual**, PowerHA® SystemMirror® does not run an automatic failover.

Resync Action { *AUTO* / *MANUAL* }

If a site failover occurs, specify how you want to resynchronize the data between the sites after the failed site is brought back online. Enter **AUTO** if you want to automatically resynchronize the data between the sites as soon as the failed site is online. Enter **MANUAL** if you want to manually resynchronize the data between the sites. The default value for this field is **AUTO**.

Consistency Enabled { *YES* / *NO* }

Enter **Yes** to enable consistency for the SRDF replicated resource, which is defined as the composite group on the EMC storage device. Enter **No** to disable consistency for the SRDF replicated resource.

If you select **SmartDR** as the SRDF Group Type, the default value **YES** is used for this field. For more information about SRDF Group Type, see [EMC SRDF consistency groups](#).

SRDF Group Type { *CompGroup* / *DevGroup* / *SmartDR* }

If devices are managed as a composite group, select **CompGroup**. If devices are managed as a device group, select **DevGroup**. If devices are managed as a Metro/SmartDR environment, select **SmartDR**. The following fields are mandatory for Metro/SmartDR environment:

SmartDR Environment Name

Enter the environment name for storage that is used during Metro/SmartDR configuration. The environment name is case-sensitive and can be up to 16 characters long. It can include letters, numbers, the dash (-) and the underscore (_) characters.

Symmetrix Array ID for MetroR1

Enter the array ID for MetroR1 storage. This is a numeric value. For example, 000197902237

Symmetrix Array ID for MetroR2

Enter the array ID for MetroR2 storage. This is a numeric value. For example, 000197902238

Symmetrix Array ID for DR

Enter the array ID for Disaster Recovery (DR) storage. This is a numeric value. For example, 000197902239

Adding a SRDF replicated resource to a PowerHA® SystemMirror® resource group

The SRDF replicated resources must be added to a resource group before they can be managed by PowerHA® SystemMirror®.

To add a SRDF replicated resource to a PowerHA® SystemMirror® resource group, complete the following steps:

1. From the command line, enter `smit sysmirror`.
2. From the SMIT interface, select **Cluster Applications and Resources > Resource Groups > Change/Show Resources and Attributes for a Resource Group**, and press Enter.
3. In the **EMC SRDF Replicated Resource** field select the SRDF replicated resource, and press Enter.

The resource group that you select must be the same resource group you used when you created the SRDF replicated resource.

4. If the EMC SRDF replicated resource is of SmartDR type, select **Never Fallback** as the Resource Group Fallback Policy. From the SMIT interface, select **Cluster Applications and Resources > Resource Groups > Change/Show Nodes and Policies for a Resource Group** and press Enter.

This step is optional if the EMC replicated resource is not of SmartDR type.

5. Verify and synchronize the cluster.

Verifying the SRDF configuration

You must verify the SRDF configuration from the command-line interface.

To verify the SRDF configuration, complete the following steps:

1. From the command line, run the `/usr/es/sbin/cluster/sr/utlils/cl_verify_sr_config` command.
2. If the verification is successful, the following message is displayed:

```
PowerHA SystemMirror support for EMC SRDF configuration verified successfully.  
Status=0
```

If the verification is not successful, fix any error messages and run the `cl_verify_sr_config` command again.

Prerequisites for verifying the EMC SRDF/Metro SmartDR configuration

Before you verify the EMC SRDF/Metro SmartDR configuration, you must ensure that the following prerequisites are complete:

- You must have an installation of Symmetrix command-line interface (SYMCLI) software for AIX® Version 9.2.0.0, or later.
- The service state of SRDF/Metro SmartDR environment must be Normal. If any storage disks that are used in pairs are unmapped or if witness configuration is used to maintain appropriate cluster state, the service state of SRDF/Metro SmartDR environment can be any other state but not the Normal state.
- You must not use previously used storage disk-pairs or SRDF/Metro SmartDR environment in EMC SRDF/Metro SmartDR configuration.
- If the resource group have SRDF/Metro storage devices as replicated resources, you must use **Never Fallback** as the fallback policy for PowerHA® SystemMirror® resource groups.

Starting the SRDF cluster

After you configure and synchronize SRDF replicated resources, you can start the PowerHA® SystemMirror® Enterprise Edition cluster. PowerHA® SystemMirror® Enterprise Edition automatically manages the SRDF replicated resources that you defined. Therefore, you do not have to manually manage the SRDF replicated resources.

To start the PowerHA® SystemMirror® Enterprise Edition cluster, complete the following steps:

1. From the command line, enter `smit cl_admin`.
2. From the C-SPOC interface, select **PowerHA® SystemMirror® Services > Start Cluster Services**.
3. Complete all fields, and press **Enter** to start the cluster services.

Verification runs automatically at cluster startup. You can turn off this option, but it is not recommended.

Monitoring EMC SRDF/Metro SmartDR configuration

During a failover event, the Disaster Recovery (DR) storage site becomes the primary site and the SRDF/Metro SmartDR link moves to the `Failed Over` state. In this scenario, until a failback event is triggered, EMC SRDF/Metro SmartDR resources fails to replicate the data back to the SRDF/Metro SmartDR storage site from the DR

storage site. Based on the number of invalid tracks that are generated during failback process, EMC SRDF/Metro SmartDR replication might get delayed.

During EMC SRDF/Metro SmartDR configuration, PowerHA® SystemMirror® automatically adds EMC SRDF/Metro SmartDR application monitor when a SRDF/Metro SmartDR resource is added as a replicated resource to a resource group. To ensure proper data replication, the EMC application monitor regularly runs the Symmetrix command-line interface (SYMCLI) **update** command.

The following scripts are used for EMC SRDF Metro/SmartDR configuration:

- /usr/es/sbin/cluster/sr/utlils/cl_smartdr_start
- /usr/es/sbin/cluster/sr/utlils/cl_smartdr_stop
- /usr/es/sbin/cluster/sr/utlils/cl_smartdr_monitor

EMC SRDF/Metro SmartDR application controller and monitor are created by using default values. Sample data for EMC SRDF/Metro SmartDR application controller and monitor is shown below:

```
# clmgr query appctl cl_app_smartdr
NAME="cl_app_smartdr"
MONITORS="cl_mon_smartdr"
STARTSCRIPT="/usr/es/sbin/cluster/sr/utlils/cl_smartdr_start"
STOPSCRIPT="/usr/es/sbin/cluster/sr/utlils/cl_smartdr_stop"
CPU_USAGE_MONITOR=""
PROCESS_TO_MONITOR_CPU_USAGE=""
CPU_USAGE_MONITOR_INTERVAL="0"
STARTUP_MODE="background"
RESOURCE_GROUP="RG1"
MONITORING_STATUS="OFFLINE"
RAW_MONITORING_STATUS="OFFLINE"

# clmgr query appmon cl_mon_smartdr
NAME="cl_mon_smartdr"
APPLICATIONS="cl_app_smartdr"
TYPE="user"
MODE="longrunning"
MONITORMETHOD="/usr/es/sbin/cluster/sr/utlils/cl_smartdr_monitor"
MONITORINTERVAL="300"
HUNG SIGNAL="9"
STABILIZATION="300"
RESTARTCOUNT="0"
RESTARTINTERVAL="0"
FAILUREACTION="notify"
NOTIFYMETHOD=""
CLEANUPMETHOD=""
RESTARTMETHOD=""
MONITORRETRYCOUNT=""
AMLOGGING=""
```

Note: To change default values of an application monitor, see [Changing the configuration of an application monitor](#).

Changing SRDF replicated resources

You can change the SRDF replicated resources by using the SMIT interface.

Complete the following steps to change the SRDF replicated resources:

1. From the command line, enter `smit sysmirror`.
2. From the SMIT interface, select **Cluster Applications and Resources > Resources > Configure EMC SRDF-replicated resources > Change/Show EMC SRDF replicated resource**, and press Enter.
3. Select the SRDF replicated resource that you want to change from the list and press Enter.
4. Make your changes to the fields, and press Enter.
5. Verify and synchronize the cluster.

Removing a SRDF replicated resource

You can remove a SRDF replicated resource by using the SMIT interface.

To remove a SRDF replicated resource, complete the following steps:

1. From the command line, enter `smit sysmirror`.
2. From the SMIT interface, select **Cluster Applications and Resources > Resources > Configure EMC SRDF-replicated resources > Remove EMC SRDF replicated resource**, and press Enter.
3. Select the SRDF replicated resource that you want to remove from the list and press Enter.
4. Confirm that you want to remove this SRDF replicated resources and press Enter.
5. Verify and synchronize the cluster.

Hitachi Truecopy/HUR replication resources

PowerHA® SystemMirror® Enterprise Edition Version 7.1.2 with Service Pack 2, or later, supports High Availability Disaster Recovery support for third-party vendor storage from Hitachi.

Hitachi Storage system supports short distance replication and long-distance replication through Truecopy synchronous and asynchronous replication with Hitachi Universal Replicator (HUR) technologies. PowerHA® SystemMirror® Enterprise Edition enables integrated discovery and management of mirrored resources for management of Highly Available (HA) resource groups.

PowerHA® SystemMirror® Enterprise Edition enablement for High-Availability and Disaster Recovery (HADR) of Hitachi mirrored storage involves the following steps:

1. Plan the storage deployment and replication necessary for your environment. This process is related to the applications and middleware being deployed in the environment which would eventually be HA managed by PowerHA® SystemMirror® Enterprise Edition.
2. Use the provided storage configuration tools provide by Hitachi to configure the storage devices you defined in step 1 and deploy them.
3. Use PowerHA® SystemMirror® Enterprise Edition interfaces to discover the deployed storage devices and define the HA policies for the applications or resource groups that are consuming the mirrored storage.

Overview of Truecopy/HUR Management

Hitachi Truecopy/HUR storage management uses Command Control Interface (CCI) operations from the AIX® operating system and PowerHA® SystemMirror® Enterprise Edition environment.

PowerHA® SystemMirror® Enterprise Edition uses these interfaces to discover and integrate the Hitachi Storage replicated storage into the HADR framework of PowerHA® SystemMirror® Enterprise Edition. This integration enables you to manage HADR for applications consuming the mirrored storage.

Integration of Truecopy/HUR and PowerHA® SystemMirror® Enterprise Edition provides the following:

- Support for Prefer Primary Site or Online on Either Site inter-site management policies.
- Flexible user-customizable resource group policies.
- Support for cluster verification and synchronization.
- Limited support for the PowerHA® SystemMirror® Enterprise Edition Cluster Single Point Of Control (C-SPOC).
- Automatic failover and reintegration of server nodes attached to pairs of Truecopy/HUR disk subsystem within sites and across sites.
- Automatic management for Truecopy/HUR links.
- Management for switching the direction of the Truecopy/HUR relationships when a site failure occurs. This process allows the backup site to take control of the PowerHA® SystemMirror® Enterprise Edition managed resource groups from the primary site.

Planning for Truecopy/HUR Management

To successfully use Truecopy/HUR Management, you must correctly plan for its implementation.

Before you continue, your environment must meet the following requirements:

- PowerHA® SystemMirror® Enterprise Edition sites are already correctly configured.
- Truecopy/HUR support is already correctly configured.

Limitations for Truecopy/HUR Management

To implement Truecopy/HUR Management correctly, you must know its limitations.

PowerHA® SystemMirror® Enterprise Edition for mirroring with Truecopy/HUR has the following limitations:

- Only fence level NEVER is supported for synchronous mirroring.
- Only HUR is supported for asynchronous mirroring.
- The dev_name must map to a logical devices and the dev_group should be defined under HORCM_LDEV section in the horcm.conf file.
- The PowerHA® SystemMirror® Enterprise Edition Truecopy/HUR solution uses dev_group for any basic operation. For example, pairresync/pairevtwait/horctakeover. If there are several dev_names in a dev_group, the dev_group must be consistency enabled.
- PowerHA® SystemMirror® Enterprise Edition does not trap SNMP notification events for Truecopy/HUR storage. If a Truecopy link goes down when the cluster is up and later the link is repaired you need to manually resync the pairs.
- The creation of pairs is done outside the cluster control. You must create the pairs before starting cluster services.
- Resource groups that are managed by PowerHA® SystemMirror® Enterprise Edition cannot contain volume groups with both Truecopy/HUR-protected and non-Truecopy/HUR-protected disks. A Resource group must have either a Truecopy protected disk or a HUR protected disk.
- All nodes in the PowerHA® SystemMirror® Enterprise Edition cluster must use same horcm instance.
- You cannot use C-SPOC for the following LVM operations to configure nodes at the remote site that contain the target volume:
 - Creating a volume group.
 - Operations that require nodes at the target site to write to the target volumes. For example, changing file system size, changing mount point, or adding LVM mirrors cause an error message in CSPOC. However, nodes on the same site as the source volumes can successfully perform these tasks. The changes are then propagated to the other site using lazy update.

Note: For C-SPOC operations to work on all other LVM operations, it is highly recommended that you perform all C-SPOC operations when the cluster is active on all PowerHA® SystemMirror® Enterprise Edition nodes and the underlying Truecopy/HUR PAIRs are in a PAIR state.

Example Truecopy/HUR Management configuration

You can set up a mutual recovery configuration in which each site acts as a production site, for different applications and services, with the other site acting as an associated backup site.


To implement a mutual recovery configuration you must have the following:

- Two PowerHA® SystemMirror® Enterprise Edition sites (the same as a single recovery configuration)
- Two resource groups

In PowerHA® SystemMirror® Enterprise Edition you can configure two sites and a resource group that includes nodes, Truecopy/HUR-managed replicated resources, and the volume groups associated with the replicated

resources. A Truecopy/HUR replicated resource contains the Hitachi disk volume pairs (dev_group) information. The PowerHA® SystemMirror® Enterprise Edition resource group definition includes the volume groups built on top of the Truecopy/HUR replicated volumes. PowerHA® SystemMirror® Enterprise Edition manages Truecopy/HUR processing by dynamically executing the Command Control Interface (CCI) commands.

The following image shows a typical implementation of two Hitachi Storage devices with Truecopy/HUR in a four-node PowerHA® SystemMirror® Enterprise Edition geographic cluster. The cluster consists of four nodes. Each storage device is connected to each node using a SCSI or Fibre Channel connection. Truecopy/HUR links are set between Primary Hitachi Storage and Secondary Hitachi Storage. The configuration also includes point-to-point networks for heartbeating to connect the cluster nodes.



Planning worksheets for Truecopy/HUR Management

You can use planning worksheets to start planning for your Truecopy/HUR Management implementation.

For information about configuring Truecopy/HUR, see the Truecopy/HUR Configuration Guide that is maintained by Hitachi.

Sample Truecopy/HUR Replicate Resource Worksheet

The following are definitions to help you complete the worksheet:

Truecopy/HUR Resource Name

The name of the Hitachi Truecopy/HUR Resource. A Resource Name is a user-defined name that can have one or more device groups with the same fence level and copy mode. The name can include up to 128 alphanumeric characters and underscores.

Truecopy/HUR Mode

Defines the Truecopy/HUR relationship mode. For synchronous Truecopy only NEVER fence level is supported. If the dev_groups is managed by Truecopy, the mode is SYNC. If the dev_group is managed by HUR, the mode is ASYNC.

Device Groups

One or more dev_groups that are defined in the configuration file. The dev_group with the same COPY mode can be merged together as device groups.

Horctakeover Timeout Value

The -t option for the **horcrtakeover** command. Use a 3short timeout for horctakeover of TrueCopy Synchronous. A short timeout or long timeout can be used for horctakeover of the Hitachi Universall Replicator (HUR). It is your decision what timeout option you want to use. For example, if returning to production is more important to you than having the most current data in a failure use a short timeout. If

ensuring that you lose as little data as possible in a disaster, even if that means the recovery time is elongated, use a long timeout.

Horcm Instance

The horcm instance that you used. For example, you use the horcm0.conf file then the value is **horcm0**. If you use the horcm.conf file then the value is **horcm**. All nodes in the PowerHA® SystemMirror® Enterprise Edition cluster must use same horcm instance. For example, if node1 uses the horcm0.conf file then the same instance (horcm0.conf) must be used by all other nodes in the cluster.

Pairevwait Timeout Value

The -t option for the **pairevwait** command. Specifies the interval of monitoring a status specified using the -s option and the time-out period in units of 1 second. Use along timeout for any pairevwait monitoring a pairevwait. The default value is 3600 seconds.

Use the following worksheet to record the Truecopy/HUR replicated resources that go in the PowerHA® SystemMirror® Enterprise Edition resource groups:

<i>SYNCHRONOUS Truecopy</i>					
Truecopy Resource Name	Truecopy Mode	Device Groups	Horctakeover Timeout Value	Horcm Instance	Pairevwait Timeout Value
TRU_RES1	SYNC	Oradb1 Oradb2 Oradb3	300	horcm0	3600
TRU_RES2	SYNC	Oradb5 Oradb6	300	horcm0	3600
<i>ASYNCHRONOUS HUR</i>					
Truecopy Resource Name	Truecopy Mode	Device Groups	Horctakeover Timeout Value	Horcm Instance	Pairevwait Timeout Value
TRU_RES3	ASYNCHRONOUS	Oradb11 Oradb12 Oradb13	3600	horcm0	3600
TRU_RES3	ASYNCHRONOUS	Oradb15 Oradb16	3600	horcm0	3600
<i>Blank worksheet</i>					
Truecopy Resource Name	Truecopy Mode	Device Groups	Horctakeover Timeout Value	Horcm Instance	Pairevwait Timeout Value

Configuring volume groups and filesystems on Truecopy/HUR protected disks

As part of planning your PowerHA® SystemMirror® Enterprise Edition Truecopy/HUR environment, decide what disks will be used to support volume groups and file systems on the PowerHA® SystemMirror® Enterprise Edition cluster nodes.

You must make sure that the Hitachi hdisks are made available to your nodes before you continue the configuration process. If the Hitachi hdisk are not available to your nodes you can reboot the nodes and run the **cfgmgr** command.

Note: PowerHA® SystemMirror® Enterprise Edition does not create replication pairs using the Hitachi interfaces. You must use the Hitachi Storage interfaces to create the same replicated pairs before using PowerHA® SystemMirror® Enterprise Edition to achieve HA/DR solution. For information about setting up

Truecopy/HUR pairs, see the Hitachi Command Control Interface (CCI) User and Reference Guide that is maintained by Hitachi.

Discovering dev_groups and AIX® disks associated with Hitachi storage

You need to identify the correct mapping between the Hitachi device groups and the corresponding AIX® hdisks.

Complete the following step to discover dev_groups and AIX® disks associated with Hitachi storage:

1. On the PowerHA® SystemMirror® Enterprise Edition nodes select the disks that are Hitachi disks and will be used in Truecopy/HUR relationships by running the **lsdev** command. In the following example, hdisk4 and hdisk5 are the Hitachi disks.

```
# lsdev -C -c disk | grep hdisk | inqraid
hdisk4 -> [SQ] CL2-A Ser = 45306 LDEV = 256 [Hitachi ] [OPEN-V]
        HORC = S-VOL HOMRCF[MU#0 = SMPL MU#1 = SMPL MU#2 = SMPL]
hdisk5 -> [SQ] CL2-A Ser = 45306 LDEV = 257 [Hitachi ] [OPEN-V]
        HORC = S-VOL HOMRCF[MU#0 = SMPL MU#1 = SMPL MU#2 = SMPL]
```

2. On the PowerHA® SystemMirror® Enterprise Edition nodes find the hdisk that will be managed by the RG (Resource Group) and get the LDEV mapping by running the **lsdev** command. In the following example, hdisk4 and hdisk5 need to be managed by a HA/DR solution, and the LDEV numbers are 256 for hdisk4 and 257 for hdisk5.

```
# lsdev -C -c disk | grep hdisk | raidscan -find
DEVICE_FILE      UID      S/F  PORT  TARG      LUN  SERIAL
LDEV  PRODUCT_ID
hdisk4          0        F    CL2-A  0         0    45306
 256    OPEN-V
hdisk5          0        F    CL2-A  0         1    45306
 257    OPEN-V
```

3. Use the **HORCM LDEV** section in the horcm2.conf file to identify the dev_group that will be managed by PowerHA® SystemMirror® Enterprise Edition. For example if hdisk4 and hdisk5 will be part of the resource group then identify the dev_group from the Horcm configuration file. In the following example, dev_group VG01 (LDEV 256) and VG02 (LDEV 257) need to be managed by the PowerHA® SystemMirror® Enterprise Edition Truecopy/HUR solution.

```
HORCM_LDEV
#dev_group      dev_name      Serial#      CU:LDEV(LDEV#)      MU#
VG01            oradb1        45306        256                 0
VG02            oradb2        45306        257                 0
```

4. Verify that the PAIRs are established by running the **pairvolchk** command or the **pairedisplay** command on the dev_group.

```
# pairvolchk -g VG01 -IH2
pairvolchk : Volstat is P-VOL.[status = PAIR fence = NEVER MINAP = 3 ]
```

Note: If pairs are not yet established, you must create the pairs. For instructions about creating pairs, see the CCI/Storage Navigator documentation that is maintained by Hitachi. The following example is displayed when pairs are not created.

```
# pairvolchk -g VG02 -IH2
pairvolchk : Volstat is SMPL.[status = SMPL]
```

5. Verify that dev_group is consistency enabled by running the **pairvolchk** command and CTGID is in the output. You can also verify if the dev_group is managed by Truecopy SYNC or HUR. If the fence value is NEVER, the dev_group is managed by Truecopy SYNC. If the fence value is ASYNC, the dev_group is managed by HUR.

```
# pairvolchk -g VG01 -IH2
pairvolchk : Volstat is P-VOL.[status = PAIR fence = NEVER MINAP = 3 CTGID = 2]
```

Setting up volume groups and filesystems on Truecopy/HUR-protected disks

After you have identified the hdisks and dev_groups that will be managed by PowerHA® SystemMirror® Enterprise Edition you need to set up the volume groups and filesystems.

Complete the following steps to set up volume groups and filesystems on Truecopy/HUR-protected disks:

1. On the hdisks that corresponds to the dev_group, set up the volume group and file system to be managed by PowerHA® SystemMirror® Enterprise Edition. Verify that the Volume Major Number for the volume group can be used on all PowerHA® SystemMirror® Enterprise Edition cluster nodes by running the **lvfstmajor** command on each cluster node. You must also verify that the physical volume name for the filesystem can also be used across all PowerHA® SystemMirror® Enterprise Edition cluster nodes.
2. Vary off the newly created volume group by running the **varyoffvg** command, and import it to all nodes in the local PowerHA® SystemMirror® Enterprise Edition site.
3. Run the **pairresync** command on the local disks and make sure that they are in the PAIR state. This process verifies that the local disk information has been copied to the remote storage.

```
#pairresync -g VG01 -IH2
#pairresync -g VG02 -IH2
```

4. Once the Truecopy relationship has completed copying, run the **pairsplit** command to SPLIT the relationship. If you do not complete this step future LVM operations will not complete successfully.
5. On the node at the backup PowerHA® SystemMirror® Enterprise Edition site, remove the paired Hitachi hdisk by running the **rmdev** command. For example, the hdisk4 at primary site is paired with hdisks9 at remote site and hdisks5 is paired with hdisk10.

```
#rmdev -d -l hdisk9
#rmdev -d -l hdisk10
```

6. Run the **cfgmgr** on the backup node.

```
#cfgmgr
```

7. Using SMIT or the command line on the backup PowerHA® SystemMirror® Enterprise Edition site (the site that is connected to the secondary Hitachi Storage) import the volume groups that you created in step 1.
8. Resync the pairs you split steps 4 by running the **pairresync** command.

```
#pairresync -g VG01 -IH2
#pairresync -g VG02 -IH2
```

9. Verify that on all cluster nodes, the AUTO VARYON feature for volume groups is set to NO by running the **chvg** command.

```
chvg -a 'n' -Q 'y' <volume group name here>
```

Note: PowerHA® SystemMirror® Enterprise Edition attempts to automatically set the AUTO VARYON to NO during verification, except in the case of remote Truecopy/HUR.

Installing Truecopy management filesets

You must be logged in as root to perform installation tasks. You must verify that the base filesets are installed on cluster nodes.

For instructions about installing the base fileset, see [“Installing PowerHA SystemMirror Enterprise Edition for Metro Mirror”](#) on page 15.

Installing prerequisite software

Before installing PowerHA® SystemMirror® Enterprise Edition: Truecopy/HUR Management, the following software must be installed on cluster nodes:

- You must have one of the following operating system levels installed:
 - IBM® AIX® 6.1 with Technology Level 8 with Service Pack 1, or later.
 - IBM® AIX® 7.1 with Technology Level 2 with Service Pack 1, or later.
- PowerHA® SystemMirror® Enterprise Edition Version 7.1.2 with Service Pack 2, or later.
- Command Control Interface (CCI) software for AIX®. For more information about CCI, see Hitachi Command Control Interface (CCI) User and Reference Guide that is maintained by Hitachi.

Installing Truecopy/HUR Management filesets

The **cluster.es.tc** fileset is included on the PowerHA® SystemMirror® Enterprise Edition media and need to be installed for Truecopy Management. You also must select the message fileset for your locales. For example, if you wanted to install the US English fileset you would select the following:

```
cluster.msg.en_US.tc    POWERHA TRUECOPY/HUR Messages - U.S. English
```

All PowerHA® SystemMirror® Enterprise Edition Truecopy/HUR filesets are installed in the `/usr/es/sbin/cluster/tc` directory.

Configuring Truecopy/HUR resources

After you install Truecopy/HUR you can configure the cluster resources.

Checking Truecopy/HUR and PowerHA® SystemMirror® Enterprise Edition configuration requirements

Before you can configure PowerHA® SystemMirror® Enterprise Edition for Truecopy/HUR interfaces, you must verify that all your environment settings are correct.

The following criteria must be met before you can configure PowerHA® SystemMirror® Enterprise Edition for Truecopy/HUR interfaces:

1. Truecopy/HUR is configured and running on the storage system
2. Command Control Interface (CCI) software is installed. For example, CCI software is installed on PowerHA® SystemMirror® Enterprise Edition cluster nodes.
3. Both the base PowerHA® SystemMirror® Enterprise Edition for Metro Mirror and Truecopy/HUR Management filesets are installed on each cluster node.
4. The PowerHA® SystemMirror® Enterprise Edition cluster is configured for nodes, sites, networks, network interfaces, service labels, and application monitors.

You can modify the attributes for a resource group later to accommodate Truecopy/HUR replicated resources.

Configuring Truecopy/HUR-Managed replicated resource

You must define Truecopy/HUR-managed replicated resources before adding them to resource groups.

To define a Truecopy/HUR-managed replicated resource, complete the following steps:

1. From the command line, enter `smit sysmirror`.
2. In SMIT select, **Cluster Applications and Resources > Resources > Configure Hitachi Truecopy Replicated Resources > Add Hitachi Truecopy/HUR Replicated Resource**, and press Enter.
3. Complete the available fields, and press Enter.

AFTER COMPLETING THE TASK Repeat the steps as necessary to define additional Truecopy/HUR-Managed replicated resources.

Adding Truecopy/HUR replicated resources to PowerHA® SystemMirror® Enterprise Edition resource groups

After you have configured the Truecopy/HUR replicated resources you must add them to the PowerHA® SystemMirror® Enterprise Edition resource groups.

To add a Truecopy replicated resource to a PowerHA® SystemMirror® Enterprise Edition resource group, complete the following steps:

1. From the command line, enter `smit sysmirror`.
2. In SMIT select, **Cluster Applications and Resources > Resource Groups**. Depending on whether you are working with an existing resource group or creating a resource group. Make sure that the volume groups selected on the Resource Group configuration screen match the volume groups used in the Truecopy/HUR Replicated Resource.
 - a. If you are creating a resource group select, **Change/Show Resources and Attributes for a Resource Group**.
 - b. If you are adding a resource group select, **Add a Resource Group**.
3. In the **Truecopy Replicated Resources** field verify that the volume group you specify is the same volume group that you specified on the Resource Group configuration panel.
4. Press Enter.

AFTER COMPLETING THE TASK Next you must verify the cluster.

Verifying the PowerHA® SystemMirror® Enterprise Edition cluster configuration

You can use the SMIT interface to verify the configuration of a PowerHA® SystemMirror® Enterprise Edition cluster.

To run the verification process for the cluster configuration, complete the following steps:

1. From the command line, enter `smit sysmirror`.
2. In SMIT select, **Custom Cluster Configuration > Verify and Synchronize Cluster Configuration (Advanced)** and press Enter.
3. Enter the following values for the available fields:

<i>Verify and Synchronize Cluster Configuration</i>	
Fields	Values
Verify Synchronize or Both	Verify
Include custom verification library checks	Yes
Automatically correct errors found during verification?	Yes
Force synchronization if verification fails?	No
Verify changes only?	No
Logging	Standard

Verification messages are logged to the `/var/hacmp/clverify/clverify.log` file.

4. Press Enter. The output from the verification process is displayed in the SMIT Command Status window.

If you receive error messages, make the necessary changes and run the verification procedure again. You might see warning messages if the configuration has a limitation on its availability. For example, only one interface per node per network is currently configured.

Verifying the Truecopy/HUR configuration

Before you can synchronize the newly created cluster, you must verify the Truecopy/HUR configuration.

To complete Truecopy/HUR verification run the following command:

```
/usr/es/sbin/cluster/tc/Utils/cl_verify_tc_config
```

Correct any configuration errors that appear, then run the script again.

Synchronizing the PowerHA® SystemMirror® Enterprise Edition cluster configuration

You must verify the PowerHA® SystemMirror® Enterprise Edition cluster and the Truecopy/HUR configuration before you can synchronize the cluster.

To propagate the new Truecopy/HUR configuration information and possibly PowerHA® SystemMirror® Enterprise Edition site information across your PowerHA® SystemMirror® Enterprise Edition cluster, complete the following steps:

1. From the command line, enter `smit sysmirror`.
2. In SMIT select, **Custom Cluster Configuration > Verify and Synchronize Cluster Configuration (Advanced)** and press Enter.
3. In the Verify Synchronize or Both field, enter **Synchronize**.
4. In the Automatically correct errors found during verification field, enter **No**.
5. Press Enter. The output is displayed in the SMIT Command Status window.

Starting the PowerHA® SystemMirror® Enterprise Edition Truecopy/HUR cluster

PowerHA® SystemMirror® Enterprise Edition manages the Truecopy/HUR replicated resources you have defined for the cluster.

PowerHA® SystemMirror® Enterprise Edition manages these Truecopy/HUR resources automatically, so there is no need to manually manage the Truecopy/HUR resources.

To start a Truecopy/HUR cluster services, complete the following steps:

Note: Perform the following only after configuring and synchronizing a cluster.

1. From the command line, enter `smit cl_admin`.
2. From SMIT, select **Manage Services > Start Cluster Services**, and press Enter.
3. Complete all fields, and press **Enter** to start the cluster services.

Changing Truecopy/HUR-Managed replicated resource

Using SMIT you can change Truecopy/HUR replicated resources.

To change Truecopy/HUR Replicated Resources, complete the following steps:

1. From the command line, enter `smit sysmirror`.
2. In SMIT select, **Cluster Applications and Resources > Resources > Configure Hitachi Truecopy/HUR Replicated Resources > Change/Show Hitachi Truecopy/HUR Replicated Resources**, and press Enter.
3. Update the PowerHA® SystemMirror® Enterprise Edition resource group information and synchronize the cluster.

Removing Truecopy/HUR-Managed replicated resource

Using SMIT you can remove Truecopy/HUR replicated resources.

To remove Truecopy/HUR Replicated Resources, complete the following steps:

1. From the command line, enter `smit sysmirror`.
2. In SMIT, select **Cluster Applications and Resources > Resources > Configure Hitachi Truecopy/HUR Replicated Resources > Remove Hitachi Truecopy/HUR Replicated Resources**, and press Enter.
3. Update the PowerHA® SystemMirror® Enterprise Edition resource group information and synchronize the cluster.

XIV® Remote Mirror replication resources

To use mirroring with the XIV® Remote Mirror management, you must be running PowerHA® SystemMirror® Enterprise Edition 7.1.2, or later in your environment. By using the XIV® Remote Mirror management, you can deploy your PowerHA® SystemMirror® cluster to be integrated with XIV® Remote Mirror replicated resources.

PowerHA® SystemMirror® discovers the XIV® deployment, and captures any missing replication-related information from the administrator. By using this captured information, PowerHA® SystemMirror® manages the XIV® replicated storage as part of the High Availability management environment for the applications and resource groups.

PowerHA® SystemMirror® supports high availability disaster recovery (HADR) for storage from the IBM® XIV® Storage Systems. XIV® Storage Systems that are capable of replication, support short distance and long-distance mirroring through synchronous and asynchronous mirroring functions.

PowerHA® SystemMirror® enables integrated discovery and management of mirrored storage resources for the high availability management of resource groups.

To use PowerHA® SystemMirror® for HADR of XIV® mirrored storage, you must complete the following steps:

1. Plan for the deployment of XIV® storage across two sites. You must understand your application and middleware requirements and verify that these requirements are covered in the context of a high availability PowerHA® SystemMirror® solution.
2. Use the following XIV® configuration tools to set up and configure the XIV® Storage System:
 - XIV® Storage System command-line interface (XCLI)
 - XIV® Storage Management GUI (XIVGUI)
3. Plan for the PowerHA® SystemMirror® implementation. You should design resource groups, policies, and scripts around your applications and mirrored storage resources.
4. Install and configure PowerHA® SystemMirror® to implement the settings that you planned for in [Step 3](#).
5. Test and verify that the configuration is working correctly.

Note: This documentation focuses on [step 3](#) and [step 4](#), but does mention and discuss the other steps at a high level. You must complete [step 1](#) and [step 2](#) before configuring PowerHA® SystemMirror® for HADR management of XIV® storage.

XIV® Remote Mirror concepts

The XIV® Remote Mirror function of the IBM® XIV® Storage System enables a real-time copy between two or more storage systems over Fibre Channel or iSCSI links. This function provides a method to protect data from site failures.

The purpose of mirroring is to create a set of consistent data that can be used by production applications during problems with production volumes or for other purposes. XIV® remote mirroring is application and operating system independent, and does not require server processor cycle usage. XIV® remote mirroring can be a synchronous copy or an asynchronous copy.

Mirroring, whether synchronous or asynchronous, requires two or more XIV® systems. The source and target of the asynchronous mirroring can reside on the same site and form a local mirroring, or the source and target can reside on different sites and enable a disaster recovery plan.

A peer volume or a peer consistency group is assigned either a primary or a backup role when the mirror is defined. By default, in a new mirror definition, the location of the primary designates the primary system, and

the backup designates the secondary system. A mirror must have one primary system and one secondary system. The actual function of the peer is based on its role.

XIV® enables a set of remote mirrors to be grouped into a consistency group. When using synchronous mirroring or asynchronous mirroring, the consistency groups handle many remote mirror pairs as a group to make mirrored volumes consistent. Consistency groups simplify the handling of many remote volume pairs because you do not have to manage the remote volume pairs individually.

Planning for XIV® Remote Mirror

Before you start using the XIV® Remote Mirror management, you must plan to implement the XIV® Remote Mirror management within a PowerHA® SystemMirror® environment.

The following restrictions and limitations apply to Remote Mirror:

- Remote mirror is only supported on IBM® XIV® Storage Systems 1.2.0, or later.
- The XIV® configuration cannot be changed by using dynamic automatic reconfiguration (DARE), even though the XIV® Remote Mirror replicated resources or mirror groups can be included and excluded from a PowerHA® SystemMirror® resource group by using DARE.
- PowerHA® SystemMirror® does not handle Simple Network Management Protocol (SNMP) events in the context of the XIV® Remote Mirror function. If there is a failure in the XIV® links, in some cases, this can put the corresponding Remote Mirror in a failed state and PowerHA® SystemMirror® might not be notified. To fix this problem, you must correct the links and then restart the corresponding resource groups on the appropriate nodes.

Planning the IBM® XIV® Storage System deployment

Before you can deploy an IBM® XIV® Storage System, you must plan accordingly.

To plan for the deployment of the IBM® XIV® across two sites, complete the following steps:

1. Identify your production site for the application. A production site is the location where your application and its data primarily reside and function.
2. Identify the volumes on the XIV® Storage System that contain the application data that you want to be high availability disaster recovery (HADR).
3. Identify the storage units and the AIX® hosts that will be running on the recovery site.
4. Verify that there are sufficient number of volumes and Fiber Channel ports that are available on the storage systems for the production site and recovery site. You can create mirror links between the storage units.

Planning the PowerHA® SystemMirror® implementation

To successfully implement PowerHA® SystemMirror® within your IBM® XIV® Storage Systems environment, you must plan accordingly.

Before you implement PowerHA® SystemMirror® within your IBM® XIV® Storage Systems environment, you must complete the following tasks:

- Collect the following information for all the XIV®s in your environment:
 - IP addresses
 - Login names and their corresponding passwords
 - Associations with various storage units
- Verify that all data volumes that need to be mirrored are visible to all relevant AIX® hosts.
- Verify that all XIV®s are accessible by using the TCP/IP network for all PowerHA® SystemMirror® nodes where you want to run XIV® Remote Mirror.

Installing XIV® Remote Mirror

PowerHA® SystemMirror® Enterprise Edition for AIX® file sets must be installed on all the AIX® nodes on which you run PowerHA® SystemMirror® clusters.

Software requirements

You must have the following software installed to use XIV® Remote Mirror:

- PowerHA® SystemMirror® Enterprise Edition for AIX® 7.1.2, or later, and the required version of the AIX® operating system and Reliable Scalable Cluster Technology (RSCT).
- IBM® XIV® microcode bundle 10.2.4, or later.
- IBM® XIV® command-line interface (XCLI) for AIX® 2.4.4, or later, on each PowerHA® SystemMirror® node. The (XCLI) provides a mechanism for issuing commands to manage and maintain the XIV® systems.

Note: Before you install XCLI, you must have the appropriate AIX® operating system release installed.

Installing XIV® Remote Mirror

You must have root access to each node before installing PowerHA® SystemMirror® Enterprise Edition support for XIV® Remote Mirror.

To install XIV® Remote Mirror, complete the following steps:

1. Insert the PowerHA® SystemMirror® Enterprise Edition for AIX® media into the DVD or CD-ROM drive.
2. From the command line, enter `smitty install`.
3. From SMIT, select **Install and Update Software > Install**.
4. Select the following file sets from the list:
 - `cluster.es.genxd.rte`
 - `cluster.es.genxd.cmds`
 - `cluster.msg.genxd.xx_XX`, where, `xx_XX` is the message fileset for your language
5. Press Enter to complete the installation.

Installing the XCLI

For the installation requirements to install XCLI and a list of available packages, see *XCLI Release Notes*®.

To download and install the IBM® XIV® Storage System command-line interface (XCLI) on the XCLI client system, see [Installing the XCLI](#) in the IBM® XIV® Storage System Information Center.

Note: PowerHA® SystemMirror® uses the XCLI commands from the `/opt/xiv/XIVGUI/xcli` directory. Hence you must install the XCLI in the same path, otherwise PowerHA® SystemMirror® throws error.

Configuring XIV® Remote Mirror

After you have planned and installed XIV® Remote Mirror you can start to configure it. You must verify that the data volumes on both sites are a exact replica of each other, otherwise you cannot import the volume groups on the remote site.

Configuring XIV® to use Remote Mirror

Before you can configure the IBM® XIV® to use Remote Mirror, data must already exist on the IBM® XIV® system.

To configure the IBM® XIV® to use Remote Mirror, complete the following steps:

1. Define the XIV® mirroring target.
To connect two XIV® systems for remote mirroring, each system must be defined to be a mirroring target of the other. An XIV® mirroring target is an XIV® system with volumes that receive data copied through XIV® remote mirroring. To define an XIV® mirroring target for an XIV® system, name the target and specify whether Fibre Channel or iSCSI protocol is used to copy the data.

2. Set the maximum initialization and synchronization rates.
With the XIV® system you can specify a user specific maximum rate (in MBps) for remote mirroring coupling initialization, and a different user specific maximum rate for resynchronization.

The initialization rate and resynchronization rate are specified for each mirroring target. The maximum initialization rate must be less than or equal to the maximum synchronization job rate (asynchronous mirroring only), which must be less than or equal to the maximum resynchronization rate.

The default settings follow:

- Maximum initialization rate: 100 MBps
- Maximum synchronization job: 300 MBps
- Maximum resynchronization rate: 300 MBps

3. Connect the XIV® mirroring ports.
After you define the remote mirroring targets, you must make one-to-one connections (iSCSI or Fibre Channel) between ports on each XIV® system.

4. Define the XIV® Volume mirror coupling and peers.
After the mirroring targets are defined, a coupling or mirror might be defined, which creates a mirroring relationship between two peers.

For each data volume that participates in the Remote Mirror on the production site, you must identify a volume that serves as its Remote Mirror Peer. This volume must reside on an XIV® storage unit at the recovery site.

5. Activate XIV® mirror couplings.
When an XIV® mirror coupling is activated, all actual data existing on the primary is copied to the backup. This process is referred to as initialization. XIV® Remote Mirroring copies volume identification information (that is, physical volume ID or PVID) and any actual data on the volumes. Space that has not been used is not copied.
6. Create and add the volume mirror couplings to consistency group mirror coupling.
When an XIV® consistency group mirror coupling is created, the volumes included in the consistency group must not have any application data on them so that there is no data movement and the initialization process is fast.

After a volume mirror coupling has completed the initialization, the primary volume can be added to a mirrored consistency group in the same storage pool. With each mirroring type there are certain additional constraints, such as same role, target, and schedule. The backup volume is automatically added to the consistency group on the remote XIV® system. One or more additional mirrored volumes can be added to a mirrored consistency group at a later time in the same way.

Note: In a consistency group all volumes have the same role. Consistency groups are handled as a single entity. For example, in asynchronous mirroring, a delay in replicating a single volume affects the status of the entire consistency group.

Configuring a storage agent

A *storage agent* is a generic name given by PowerHA® SystemMirror® for an entity that provides a coordination point for storage management and often uses TCP/IP as the communication interface.

For PowerHA® SystemMirror® to manage the XIV® remote mirrors, a storage agent must be defined. This is the XIV® management interface, which resides in the XIV® controller. You must provide the IP address and authentication information that is used to communicate with the XIV® storage agent.

Adding a storage agent

Before you can use a storage agent, you must add it to your environment.

To add a storage agent, complete the following steps:

1. From the command line, enter `smitty sysmirror`.
2. From the SMIT interface, select **Cluster Applications and Resources > Resources > Configure XIV Remote Mirror Resources > Configure Storage Agents > Add a Storage Agent**, and press Enter.
3. Enter the following information for the fields:

<i>Add a Storage Agent SMIT fields</i>	
Fields	Values
Storage Agent Name	Enter the PowerHA® SystemMirror® name for the storage agent. This name must be unique within the cluster definition.
IP Addresses	Select the IP address for the Storage Agent from the list.
User ID	Enter the User ID that can access the XIV® system.
Password	Enter the password for the User ID that can access the XIV® system.

4. Verify all fields are correct and press Enter.

Modifying an existing storage agent

You can make changes to any of the properties of an existing storage agent that you have already added to your environment.

To modify any of the properties for a storage agent, complete the following steps:

1. From the command line, enter `smitty sysmirror`.
2. From the SMIT interface, select **Cluster Applications and Resources > Resources > Configure XIV Remote Mirror Resources > Configure Storage Agents > Change/Show Storage Agent**, and press Enter.
3. From the list select the name of the storage agent that you want to change and press Enter.
4. Enter the changes you want to make in the fields.
5. Verify the changes are correct and press Enter.

Removing a storage agent

If you do not want a storage agent to be included in your environment, you can remove it.

To remove a storage system, complete the following steps:

1. From the command line, enter `smitty sysmirror`.
2. From the SMIT interface, select **Cluster Applications and Resources > Resources > Configure XIV Remote Mirror Resources > Configure Storage Agents > Remove a Storage Agent**, and press Enter.
3. From the list of storage agents select the name of the storage agent that you want to remove and press Enter.
4. Press Enter to confirm that the selected storage system is the one you want to remove.

Configuring a storage system

A *storage system* is a generic term that is used by PowerHA® SystemMirror® for an entity such as a DS8000® storage unit.

When using XIV® Remote Mirror, you must provide the IBM® XIV® system identifier for the Storage System. For example, MX90012 is a storage identifier for an XIV® Storage System. This information can be added, changed, or removed through the Add, Change, or Remove a Storage System SMIT menus.

Adding a storage system

Before you can start using your storage system, you must add it to your environment.

To add a storage system, complete the following steps:

1. From the command line, enter `smitty sysmirror`.
2. From the SMIT interface, select **Cluster Applications and Resources > Resources > Configure XIV Remote Mirror Resources > Configure Storage Systems > Add a Storage System**, and press Enter.
3. Enter the following information for the fields:

<i>Add a Storage System fields</i>	
Fields	Values
Storage System Name	Enter the PowerHA® SystemMirror® name for the storage system. This name must be unique within the cluster definition.
Storage Agent Name	Enter the name for the storage agent that manage this storage system.
Site Association	Enter the vendor specific unique identifier for the storage system.
Vendor Specific Identification	Enter the unique vendor specific ID for this storage system. For example, MX90012.
User ID	Enter the user ID that can access the XIV® system.
Password	Enter the password for the user ID that can access the XIV® system.

4. Verify that all fields are correct and press Enter.

Modifying a storage system

You can make changes to any of the properties of an existing storage system that you have already added to your environment.

To modify the properties for a storage system, complete the following steps:

1. From the command line, enter `smitty sysmirror`.
2. From the SMIT interface, select **Cluster Applications and Resources > Resources > Configure XIV Remote Mirror Resources > Configure Storage Systems > Change/Show Storage System**, and press Enter.
3. From the list select the name of the storage system that you want to change and press Enter.
4. Enter the changes you want to make in the fields.
5. Verify that the changes are correct and press Enter.

Removing a storage system

If you do not want a storage system to be included in your environment, you can remove it.

To remove a storage system, complete the following steps:

1. From the command line, enter `smitty sysmirror`.
2. From the SMIT interface, select **Cluster Applications and Resources > Resources > Configure XIV Remote Mirror Resources > Configure Storage Systems > Remove Storage System**, and press Enter.

- From the list of storage systems, select the name of the storage system that you want to remove and press Enter.
- Press Enter to confirm that the selected storage system is the one you want to remove.

Configuring mirror groups

A *mirror group* is a generic name given by PowerHA® SystemMirror® for a logical collection of volumes that have to be mirrored to another storage system that resides on a remote site. An XIV® consistency group represents a mirror group.

Adding a mirror group

Before you can start using your mirror group, you must add it to your environment.

To add a mirror group, complete the following steps:

- From the command line, enter `smitty sysmirror`.
- From the SMIT interface, select **Cluster Applications and Resources > Resources > Configure XIV Remote Mirror Resources > Configure Mirror Groups > Add Mirror Group**, and press Enter.
- Enter the following information for the fields:

<i>Add Mirror Group SMIT fields</i>	
Fields	Values
Mirror Group Name	Enter the PowerHA® SystemMirror® name for the replicated resource. This name is included in a PowerHA® SystemMirror® resource group.
Storage System Name	From the list, select the storage system on the production site that have data volumes, which forms the mirror group.
Vendor Specific Identifier	Enter the XIV® consistency group identifier.
Recovery Action	<p>If a site fail over occurs, specify the disaster recovery policy that will be taken by the PowerHA® SystemMirror® resource group. Enter Manual if you want manual intervention when site failover occurs.</p> <p>Enter Automated if you do not want manual intervention when site failover occurs.</p> <div style="border: 1px solid blue; padding: 5px; margin-top: 10px;"> <p>If you specify Manual, that does not indicate that a manual intervention is required for all failover scenarios. There are some conditions, such as cluster partition, in which doing an automatic failover from one site to another, can cause potential data divergence and integrity issues. If PowerHA® SystemMirror® detects the potential for such a case, and if the recovery action associated with the mirror group is set to Manual, PowerHA® SystemMirror® does not execute an automatic failover.</p> </div>

- Verify that all fields are correct and press Enter.

Modifying a mirror group

You can make changes to any of the properties of an existing mirror group that you have already added to your environment.

To modify any of the properties for a mirror group, complete the following steps:

1. From the command line, enter `smitty sysmirror`.
2. From the SMIT interface, select **Cluster Applications and Resources > Resources > Configure XIV Remote Mirror Resources > Configure Mirror Groups > Change/Show a Mirror Groups**, and press Enter.
3. From the list, select the name of the mirror group that you want to change and press Enter.
4. Enter the changes that you want to make in the fields.
5. Verify that the changes are correct and press Enter.

Removing a mirror group

If you do not want a mirror group to be included in your environment, you can remove it.

To remove a mirror group, complete the following steps:

1. From the command line, enter `smitty sysmirror`.
2. From the SMIT interface, select **Cluster Applications and Resources > Resources > Configure XIV Remote Mirror Resources > Configure Mirror Groups > Remove Mirror Groups**, and press Enter.
3. From the list, of mirror groups select the name of the mirror group that you want to remove and press Enter.
4. Press Enter to confirm the selected mirror group is the one you want to remove.

Configuring resource groups

After you configure the mirror group, you should include the mirror group into the desired PowerHA® SystemMirror® resource group.

Before you configure resource groups, you must have a basic understanding of site support for PowerHA® SystemMirror® resource groups. For more information about site support for PowerHA® SystemMirror® resource groups, see [Planning resource groups](#).

When you configure a resource group you must meet the following requirements:

- Set the site policy to **Prefer Primary Site** or **Online on Either Site**.
- Specify a startup policy other than **Online on All Available Nodes**.
- Set the **Resource Group Processing Ordering** field to **Serial**.

Adding a remote mirror to a resource group

Before you can start using your remote mirror, you must add it to an existing resource group in your environment.

To add a remote mirror resource that is replicated to a resource group, complete the following steps:

1. From the command line, enter `smitty sysmirror`.
2. From the SMIT interface, select **Cluster Applications and Resources > Resource Groups > Change/Show Resources and Attributes for a Resource Group** and press Enter.
3. Enter the name of the mirror groups in the **DS8000(GM)/XIV Replicated Resources** field.
4. Enter the name of the volume groups associated with the individual remote mirror resources that are replicated in the **Volume Groups** field.

The volume group names must be listed in the same order as the XIV® mirror group names.

5. Verify and synchronize the cluster.

HyperSwap® for PowerHA® SystemMirror®

The HyperSwap® function in PowerHA® SystemMirror® Enterprise Edition 7.1.2, or later, provides for continuous availability against storage errors. HyperSwap® is based on storage-based synchronous replication. HyperSwap® technology enables the host to transparently switch an applications I/O operation to the auxiliary volumes, provided physical connectivity exists between the host and the auxiliary storage subsystem.

The HyperSwap® function in PowerHA® SystemMirror® supports the following capabilities within your environment:

- Eliminates primary disk subsystems as the single point of failure.
- Provides maintenance for storage devices without any application downtime.
- Provides migration from an old storage device to a new storage system.

HyperSwap® for PowerHA® SystemMirror® concepts

The HyperSwap® function in PowerHA® SystemMirror® Enterprise Edition 7.1.2, or later, enhances application availability for storage errors by using IBM® DS8000® metro mirroring. If you use the HyperSwap® function in your environment, your applications stay online even if errors occur on the primary storage because PowerHA® SystemMirror® 7.1.2, or later, transparently routes the application I/O to an auxiliary storage system.

The HyperSwap® function uses a model of communication, which is called in-band, that sends the control commands to a storage system through the same communication channel as the I/O for the disk. The HyperSwap® function supports the following types of configurations:

Traditional Metro Mirror Peer-to-Peer Remote Copy (PPRC)

The primary volume group is only visible in the primary site and the auxiliary volume group is only visible in the auxiliary site.

HyperSwap®

The primary and auxiliary volume group are visible from the same node in the cluster.

You typically configure the HyperSwap® function to be used in the following environments:

Single node environment

A single compute node is connected to two storage systems that are in two sites. This HyperSwap® configuration is ideal to protect your environment against simple storage failures in your environment.

Multiple site environment

A cluster has multiple nodes that are spread across two sites. This HyperSwap® configuration provides high availability and disaster recovery for your environment.

Mirror groups in HyperSwap® for PowerHA® SystemMirror® represent a container of disks and have the following characteristics:

- Mirror group contain information about the disk pairs across the site. This information is used to configure mirroring between the sites.
- Mirror groups can contact a set of logical volume manager (LVM) volume groups and a set of raw disks that are not managed by the AIX® operating system.
- All the disks devices that are associated with the LVM volume groups and raw disks that are part of a mirror group are configured for consistency. For example, the IBM® DS8800 views a mirror group as one entity regarding consistency management during replication.
- The following types of mirror groups are supported:

User mirror group

Represents the middleware-related disk devices. The HyperSwap® function is prioritized internally by PowerHA® SystemMirror® and is considered low priority.

System mirror group

Represents critical set of disks for system operation, such as, rootvg disks and paging space disks. These types of mirror groups are used for mirroring a copy of data that is not used by any other node or site other than the node that host these disks.

Repository mirror group

Represents the cluster repository disks of that are used by Cluster Aware AIX (CAA).

Swap timeout for HyperSwap®

Swap time is the amount of I/O delay time in seconds that PowerHA® SystemMirror® causes while performing a HyperSwap® operation on a mirror group. The swap timeout value is specific for each mirror group in a cluster.

There are different swap timeout values for planned HyperSwap® and unplanned HyperSwap®. The swap timeout value for a planned HyperSwap® is 120 seconds, and it cannot be changed. The swap timeout value for an unplanned HyperSwap® is 0 - 180 seconds.

To determine the swap timeout value to use for an unplanned HyperSwap®, consider the following factors about your environment:

1. Number of nodes where the application is hosted. The greater the number of nodes means that more information is being shared.
2. Network latency and application network usage.
3. Number of disks that are used by the application.
4. I/O response time requirements for the application.

Planned HyperSwap® for PowerHA® SystemMirror®

A planned HyperSwap® occurs when you initiate a HyperSwap® from the primary storage subsystem to the auxiliary storage subsystem.

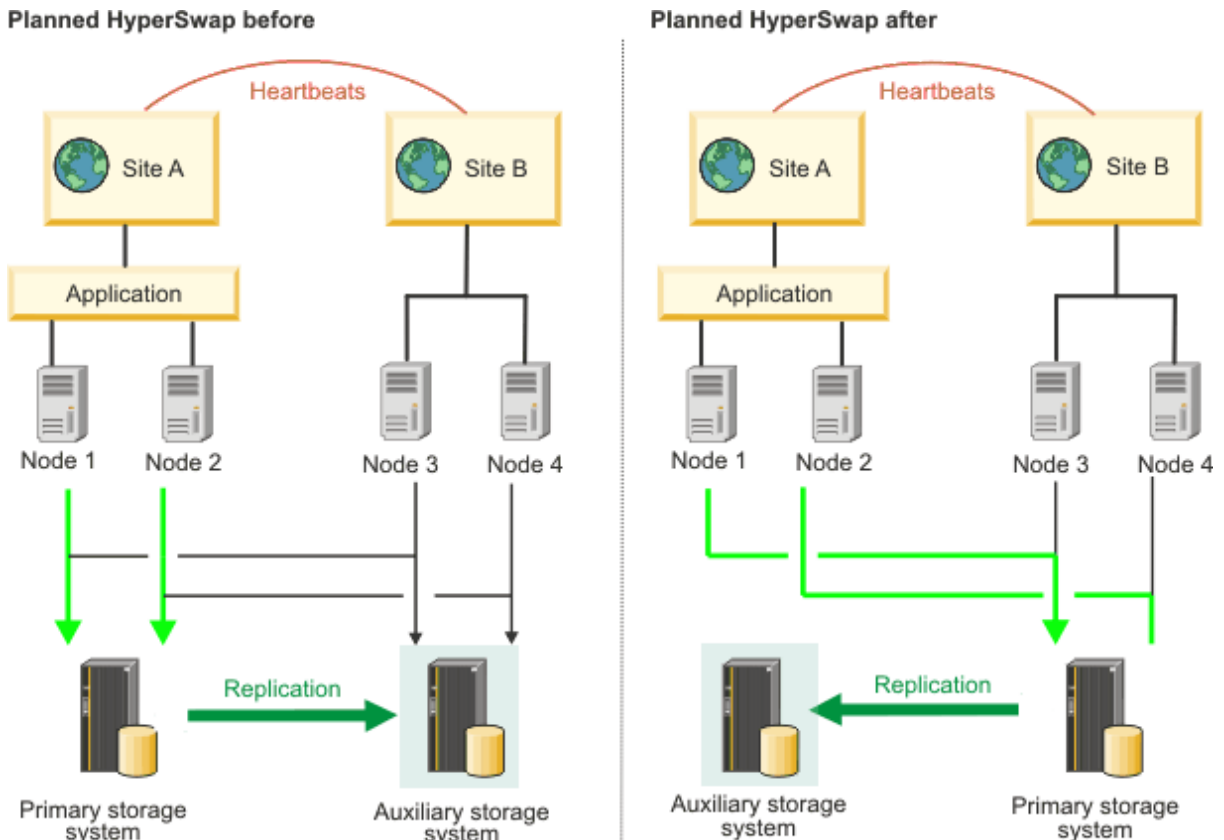
During a planned HyperSwap®, I/O activity for an application stops after coordination occurs across the host in the cluster. The application I/O is switched to the auxiliary storage subsystem and the application I/O activity continues to function as normal.

A planned HyperSwap® is ideal when you perform maintenance on the primary storage subsystem, or when you migrate from an old storage subsystem to a new storage subsystem.

The following figure shows a cluster configuration using PowerHA® SystemMirror® Enterprise Edition for AIX® that has the following characteristics:

- Two sites called Site A and Site B.
- Two nodes for each site for a total of four nodes.
- A concurrent application, for example a DB2® application that is active on Node 1 and Node 2.
- Application disks are replicated by using IBM® DS8800 metro mirroring.
- All four nodes can access both instances of the application disks that are being replicated.

The following figure shows the changes in your environment when a failure occurs and your sites are configured for a planned HyperSwap®. The primary storage system on Site A is changed to the auxiliary storage system because the application is running on Node 1 and Node 2 can access the storage system on Site B as shown in following figure. Therefore, the application that is running on Site A now stores data on the primary storage system at Site B.



pprc008-2

Figure 3: Planned HyperSwap® configuration

Unplanned HyperSwap® for PowerHA® SystemMirror®

An unplanned HyperSwap® occurs when a primary storage system fails, and the operating system detects and reacts by performing a failover. During the failover, the application I/O on the primary storage system is transparently redirected to a auxiliary storage system and the application I/O continues to run.

During the HyperSwap® process, when the applications are being redirected to an auxiliary storage system, the application I/O is temporarily suspended.

If an unplanned HyperSwap® does not complete successfully, the application I/O fails and a resource group failover event starts based on the site policy. You cannot define a failover event in a site policy for concurrent resource groups.

There are multiple scenarios when an unplanned HyperSwap® can occur.

Scenario: Unplanned HyperSwap® configuration for node access failure

In this scenario, a node loses access to the primary storage system and the redirects application I/O to an auxiliary storage system.

The following figure shows a cluster configuration in the PowerHA® SystemMirror® Enterprise Edition for AIX® environment with the following characteristics:

- Two sites called Site A and Site B.
- Two nodes for each site, which amounts to a total of four nodes. Node 1 and Node 2 have access to each storage system at Site A and Site B.
- A concurrent application, for example a DB2® application that is active on Node 1 and Node 2.
- Application disks that are replicated by using IBM® DS8800 metro mirroring.

In the following figure Node 1 and Node 2 lose access to the primary storage system on Site A, and an unplanned HyperSwap® occurs. To correct this problem, PowerHA® SystemMirror® validates whether the nodes that host the application (Node 1 and Node 2) can access the auxiliary storage system on Site B. The HyperSwap® function automatically redirects the application I/O from Node 1 and Node 2 to the storage system on Site B. The storage system on Site B becomes the primary storage system.

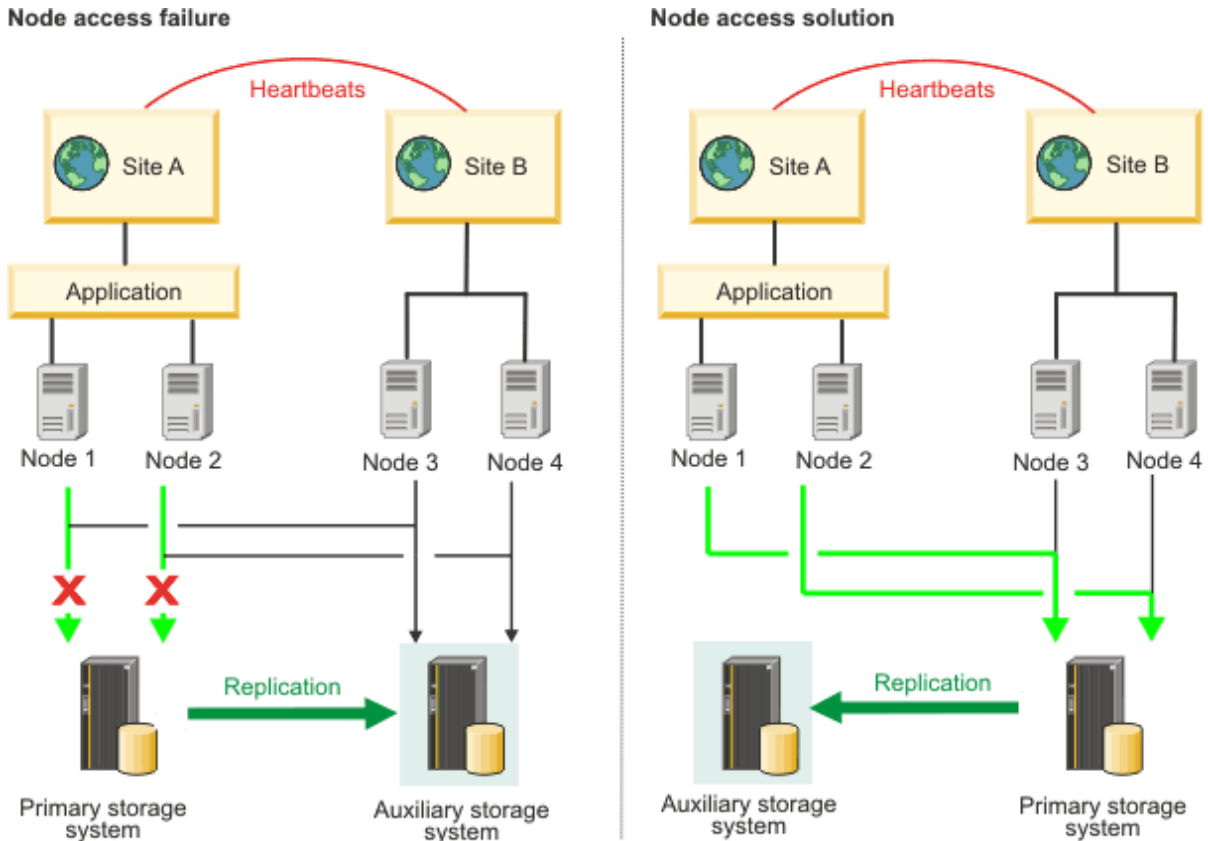


Figure 4: Unplanned HyperSwap® configuration for node access failure

Scenario: Unplanned HyperSwap® configuration for site partition failure

In this scenario, the heartbeat connection between two sites fails. Replication between both sites does not occur until the heartbeat connection is fixed and working again.

The following figure is for a cluster configuration in the PowerHA® SystemMirror® Enterprise Edition for AIX® environment with the following characteristics:

- Two sites called Site A and Site B.
- Two nodes for each site for a total of four nodes. All nodes have access to each storage system at Site A and Site B.
- A concurrent application, for example a DB2® application that is active on Node 1 and Node 2.
- Application disks are replicated by using IBM® DS8800 metro mirroring.

The following figure describes the effect of an automated failover for applications in a site partition. The nodes for the cluster in this scenario cannot communicate with each other.

Node 1 and Node 2 can communicate with each other, but they cannot communicate with Node 3 and Node 4. Therefore, Node 1 and Node 2 consider Node 3 and Node 4 to be offline (and vice versa). The application running on Node 1 and Node 2 is configured for automated failover, and the nodes at each Site A and Site B attempt to bring the application online. The application is now running independently on both sites and can cause complex problems, such as data corruption between the two sites.

pprc09-2

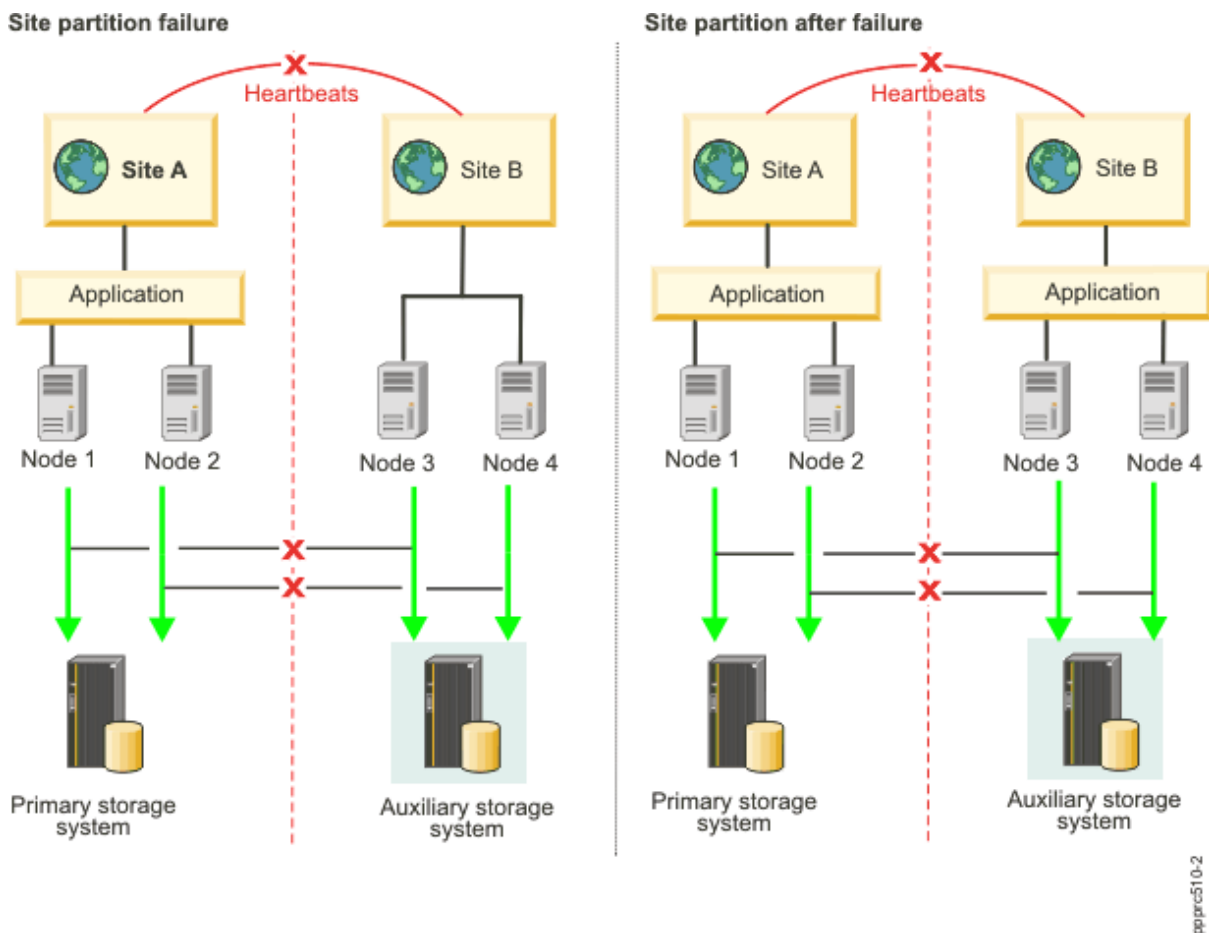


Figure 5: Unplanned HyperSwap® configuration for site failure

You can use the following configuration options in your environment to minimize or eliminate unplanned site failure:

Use split and merge policies

If a failover occurs, you can specify the type of policy you want PowerHA® SystemMirror® to implement. The policies can reduce site partition during a failover. The policies that are combined with the HyperSwap® function can reduce the possibility of data corruption. Thus, if your cluster uses the HyperSwap® function you should configure split and merge policies, so that you can automatically recover data during a failover. The following types of policies are supported by the HyperSwap® function:

- Split policy
- Merge policy
- Split and merge policy

Recovery action

You can configure your cluster to use either the automatic or manual options for mirror group recovery during a failover.

Starting resource groups and mirror groups

When you use PowerHA® SystemMirror® to start resource groups and mirror groups, you must verify that the selected resource or mirror group is in the auxiliary storage system. If you select the auxiliary storage system, the mirror group or resource group does not start and goes into an error state because the primary storage system cannot be accessed. Selecting the correct resource and mirror group minimizes data corruption if your cluster splits during a failover.

Swapping resource groups and mirror groups

When a failure occurs on the primary volume group, an automatic swap occurs. However, the swap does not depend on a cluster split or a recovery option. The swapping process ensures that the storage systems that contain the resource groups and mirror groups are accessible, so that a swap can occur.

Fallover process for resource groups and mirror groups

When failover occurs for a node or site that contains resource groups and mirror groups, PowerHA® SystemMirror® identifies if another node or site is accessible. If the node or site is not accessible, the resource groups and mirror groups is not started on the auxiliary storage system unless the cluster is configured with a split or merge policy.

Scenario: Unplanned HyperSwap® configuration for suspended storage systems

In this scenario, the storage systems on both sites are not synchronizing and replication is not occurring.

The following figure is for a cluster configuration in the PowerHA® SystemMirror® Enterprise Edition for AIX® environment with the following characteristics:

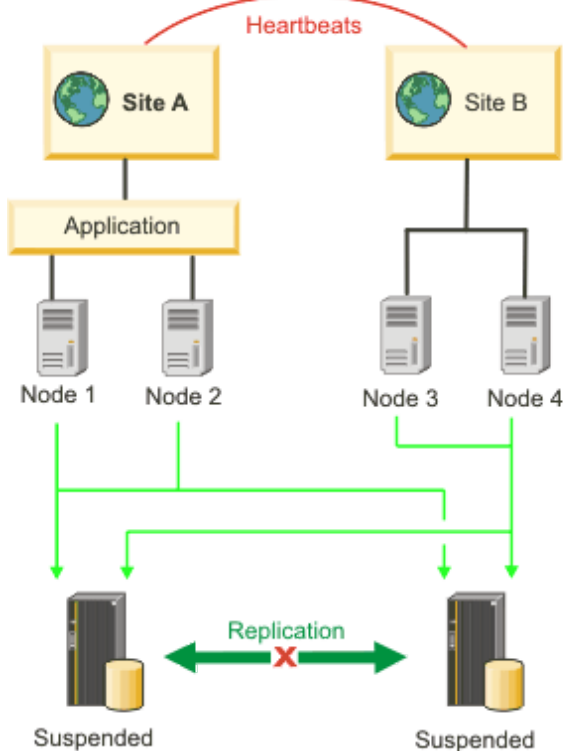
- Two sites called Site A and Site B.
- Two nodes for each site for a total of four nodes. All nodes have access to each storage system at Site A and Site B.
- The storage systems for both sites are in a suspended state.
- Application disks are replicated by using IBM® DS8800 metro mirroring.

In the following figure, the storage systems for Site A and Site B are in a suspended state. HyperSwap® for PowerHA® SystemMirror® cannot determine where the latest data is available because the storage systems are in a suspended state. There is no automatic recovery solution to fix the problem.

To correct the problem identified in this scenario, complete the following steps:

1. Perform a failback recovery operation on the storage system.
2. From the command line, enter `smit sysmirror`.
3. From the SMIT interface, select **System Management (C-SPOC) > Resource Group and Applications > Bring Resource Group Online**, and press Enter.

Suspended storage system failure



Suspended storage system solution




Figure 6: Unplanned HyperSwap® configuration for suspended storage systems

pprcsf12-2

HyperSwap® configuration for a system mirror group

In a HyperSwap® configuration for a system mirror group, only one node can access the rootvg disk that is being mirrored to both the primary storage system and auxiliary storage system. The rootvg disk contains all the information about booting, configuring, and the file system that is specific to a node.

In the following figure, the rootvg disk is mirrored from the primary storage system to the auxiliary storage system. A node must not have access to another node's rootvg disk. For example, Node 2 cannot have access to the rootvg disk for Node 1. You must set up zoning correctly in your network environment to verify that a node does not have access to another node's rootvg disk.




Figure 7: HyperSwap® configuration for a system mirror group

HyperSwap® configuration for a repository disk

For linked clusters, repository disks are shared by nodes that are at the same site, and each site has its own repository disk. Sites based on a stretched cluster share a common repository disk with all the nodes in the cluster.

In the following figure, Node 1 and Node 2 can access the repository disk on Site A and the replicated repository disk on Site B. If any type of failure occurs where Node 1 and Node 2 cannot access the repository disk in the primary storage system, the nodes can access the repository disk in the auxiliary storage system. You must configure your network zone correctly to ensure that only nodes for a specific site have access the repository disk that belongs to that site. For example, Node 3 and Node 4 on Site B do not have access to the repository disk in either storage system because it belongs to Node 1 and Node 2, which are part of Site A.

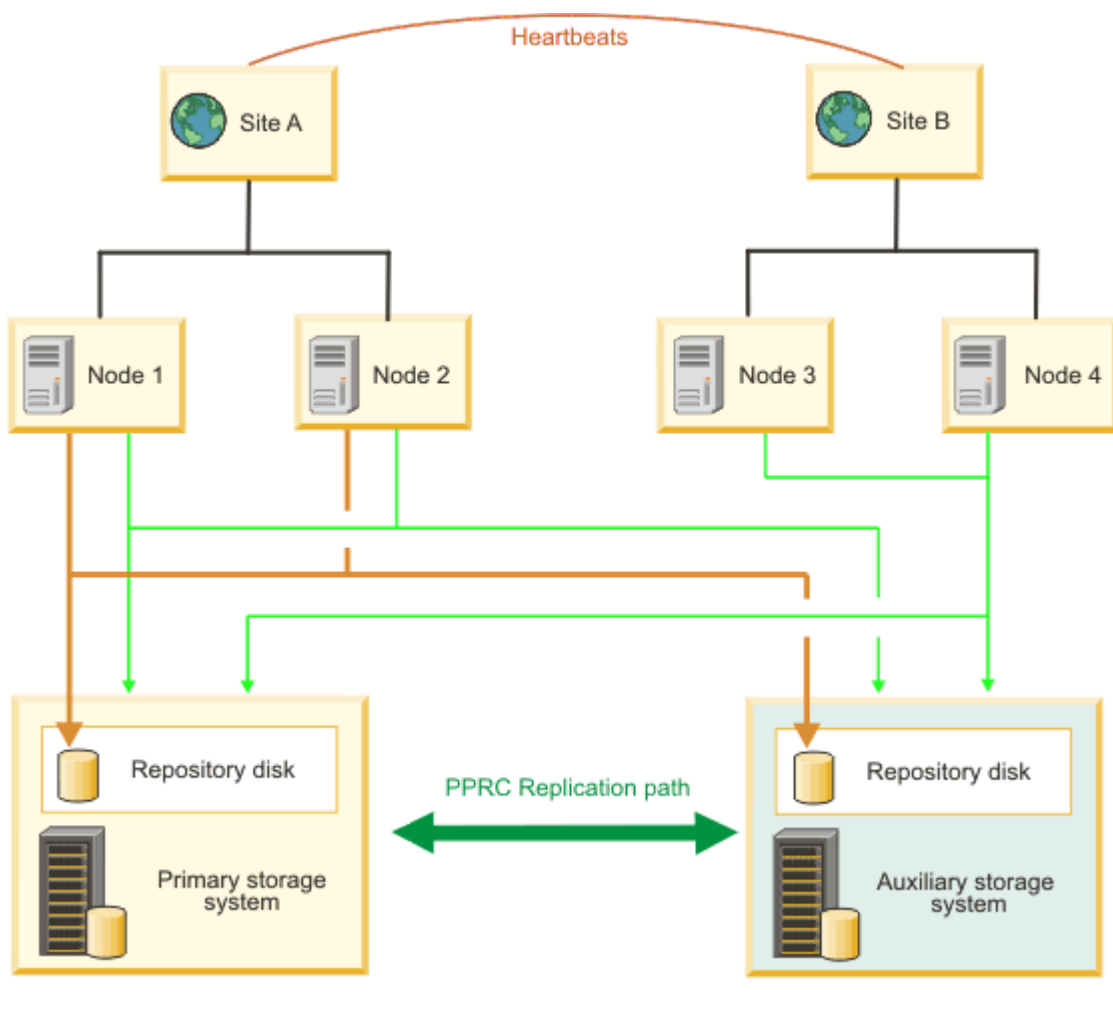


Figure 8: HyperSwap® configuration for a repository disk

HyperSwap® configuration for active-active workloads

A PowerHA® SystemMirror® cluster that has two sites with specific applications (parallel database) and uses the HyperSwap® function can operate concurrently with a shared disk.

Your application workloads in this cluster configuration can be balanced between two clusters with access to shared data, while application data integrity is maintained across the sites through Metro Mirror and HyperSwap® functions. If a cluster failure or site failure occurs, the application workloads can establish a connection to access the data by using a different site. This function is called active-active. Thus, the HyperSwap® function extends the capability to build an active-active data center solution.

Active-active workloads across sites

In an active-active configuration that supports application workloads across sites, all nodes have applications that are online and that perform I/O functions concurrently to a storage system. In this configuration, PowerHA® SystemMirror® performs swap and resync functions.

The following figure shows an active-active configuration that supports application workloads across sites that have the following characteristics:

- Two sites that are called Site A and Site B.
- Two nodes at each site, which amounts to a total of four nodes.
- All nodes have mirror groups that are configured to use the HyperSwap® function.
- All nodes have application workloads that are online and can access the Primary storage system.

In the following figure, if a failover occurs for the primary storage system at Site A, the I/O for the storage subsystem moves to the auxiliary storage system at Site B. The auxiliary storage system becomes the primary

storage system. When the storage system at Site A recovers and is brought back online, it automatically becomes the auxiliary storage system.

Before




Figure 9: Active-active configuration across sites

Active-active workloads within a site

In an active-active configuration that supports application workloads within a site, all nodes in a specific site can have applications that are online and that perform I/O concurrently to a storage system. In this configuration, PowerHA® SystemMirror® performs swap and resync functions.

The following figure shows an active-active configuration that supports application workloads that has a site with the following characteristics:

- Two sites that are called Site A and Site B.
- Two nodes at each site, which amounts to a total of four nodes.
- All nodes have mirror groups that are configured to use the HyperSwap® function or the traditional in-band Metro Mirror function.
- The application is operating only on Node 1 and Node 2.

In the following figure, the application workload moves from Site A to Site B if failure occurs at both Node 1 and Node 2. If Node 1 fails, the application continues to operate on Node 2 and vice versa. If both Node 1 and Node 2 fail, the application workload is moved to Node 3 and Node 4 at Site B. Also, the auxiliary storage system at Side B becomes the primary storage system. When the storage system at Site A recovers and is brought back online, it automatically becomes the auxiliary storage system.




Figure 10: Active-active configuration within a site

HyperSwap® resync function

PowerHA® SystemMirror® attempts to automatically resync cluster data during a failover for Peer-to-Peer Remote Copy (PPRC) replication technologies.

In PowerHA® SystemMirror® 7.1.3, or later, if failover occurs for a mirror group, you can use the HyperSwap® function to automatically resync data without distributing applications when the mirror group is brought back online. If failover occurs for your replication paths between storage systems, the HyperSwap® function can automatically re-create the replication path before the storage systems are brought online.

In the following figure, the primary storage system has a failure and is offline. The application I/O from Node 1 and Node 2 cannot access the primary storage system.




Figure 11: Primary storage system is offline

In the following figure, the application I/O resumes on the auxiliary storage system. The auxiliary storage system at Site B becomes the primary storage system. The previously primary storage system at Site A is in a suspended state. The replication paths between Site A and Site B are broken.

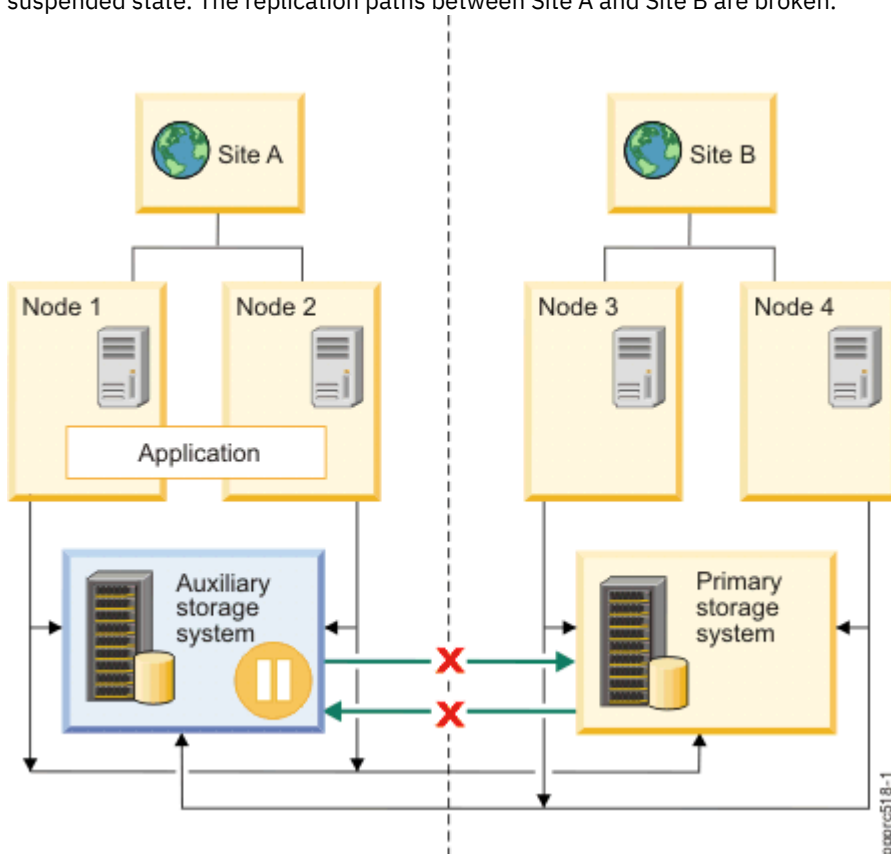


Figure 12: Replication paths are broken

In the following figure, the auxiliary storage system recovers and is brought back online. The HyperSwap® function automatically establishes the replication paths between the primary and auxiliary storage systems. The mirroring process between both Site A and Site B resumes.

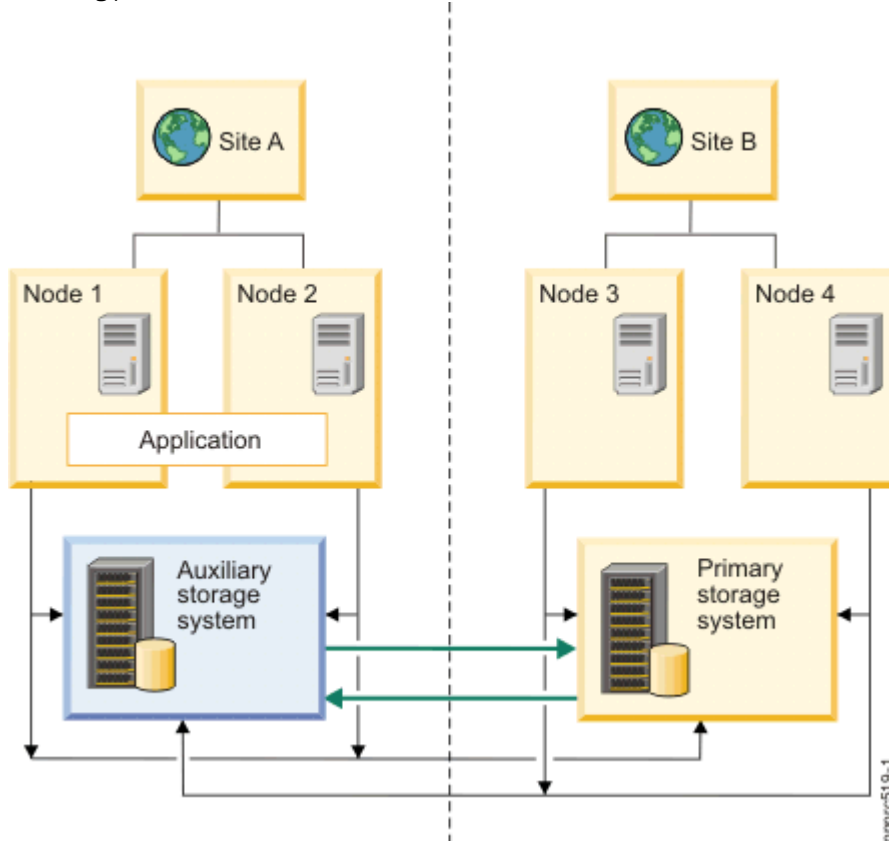


Figure 13: Mirroring resumes between Site A and Site B

Planning for HyperSwap® for PowerHA® SystemMirror®

Before you can install HyperSwap® for PowerHA® SystemMirror®, you must verify that your environment is working correctly.

To correctly plan for a HyperSwap® configuration, review the following information:

- Storage level Peer-to-Peer Remote Copy (PPRC) relationships and PPRC paths must be defined before you configure HyperSwap® for PowerHA® SystemMirror® or In-band PPRC Metro Mirror for PowerHA® SystemMirror®.
- HyperSwap® for PowerHA® SystemMirror® is only supported on IBM® DS8800, or later, storage systems.
- Concurrent workloads across sites, such as Oracle Real Application Clusters (RAC), and concurrent resource groups are supported in stretched clusters and linked clusters that are using HyperSwap® enabled mirror groups. PowerHA® SystemMirror® does not support concurrent resource groups for other PPRC replications solutions.
- You can configure HyperSwap® enabled mirror groups for automatic resynchronization when a replication failure occurs. Error logs are created when the replication failure occurs, and the log file can be used to identify the cause of the failure.
- To maintain consistency group schematics, suspended operations on a storage device must function on the entire logical subsystem (LSS). For example, if a single IBM® DS8800 LSS contains PPRC volumes from more than one application and if one of the replication connections breaks, all PPRC paths are removed. If the applications are not managed by PowerHA® SystemMirror®, the PPRC paths must be manually re-created or refreshed after the replication connection is reestablished.
- Refresh a mirror group in the cluster after any storage-level PPRC configuration changes or if a PPRC path is broken. When you refresh a mirror group, you are re-creating the paths for the mirror group and refreshing any existing configuration changes. Do not run any HyperSwap® functions when you are refreshing a mirror group.

- To add a node to a mirror group, you must perform configuration operations from a node where all disks are accessible.
- Metro Mirror (in-band) functions, including HyperSwap®, are supported in Virtual I/O Server (VIOS) configurations by the N-Port ID Virtualization (NPIV) method of disk management.
- Metro Mirror (in-band) functions, including HyperSwap®, are not supported by the virtual SCSI (VSCSI) method of disk management.
- To use Live Partition Mobility (LPM), you must bring the resource group that contains the mirror group into an unmanaged state by using the C-SPOC utility to stop cluster services with the **Unmanage Resource Groups** option. After you complete the LPM configuration process, you must bring the resource group back online by using SMIT. This process brings all mirror groups and resource groups back online.
- Disk replication relationships must adhere to a one-to-one relationship between the underlying LSS. An LSS that is already part of a mirror group, cannot be part of another mirror group.
- Repository disks require that you specify an alternate disk or a disk that is not configured to use the HyperSwap® function when you set HyperSwap® property to **Disable**.
- SCSI reservations are not supported for devices that use the HyperSwap® function.
- You must verify and synchronize the cluster when you change the cluster configuration. If you change the mirror group configuration while cluster services are active (DARE), those changes might be interpreted as failures, which result in unwanted cluster events. You must disable the HyperSwap® function before you change any settings in an active cluster environment.
- A node in an unmanaged state does not respond or perform any swap or resync functions. However, the I/O to the node continues. Thus, having a node in an unmanaged state is ideal for performing LPM on the node. After you complete the LPM process, you can move the node from unmanaged to online state by using SMIT. If an unplanned event occurs while the node is in an unmanaged state, the node is moved to a halt state.

Fileset requirements

To run HyperSwap® for PowerHA® SystemMirror®, you must install the following filesets on all nodes in the cluster:

<i>HyperSwap® fileset requirements for PowerHA® SystemMirror®</i>		
Fileset name	Minimum fileset level for AIX® Version 7.1	Minimum fileset level for AIX® Version 6.1
cluster.es.genxd.cmds	7.1.2.0	6.1.8.0
devices.common.IBM.storwork.rte	7.1.2.0	6.1.8.0
devices.common.IBM.mpio.rte	7.1.2.0	6.1.8.0
devices.fcp.disk.rte	7.1.2.0	6.1.8.0

HyperSwap® support for DS8800 Metro Mirroring

IBM® DS8800 provides Metro Mirror (in-band) communication capabilities, which allow PowerHA® SystemMirror® to manage DS8800 replication without using the Hardware Management Console (HMC).

Using the Inband communication model, PowerHA® SystemMirror® can send replication commands with specific Small Computer System Interface (SCSI) commands. The SCSI commands are sent through the same Fibre Channel that is used to send I/O requests to the storage device. This process greatly reduces the configuration complexities and failover times.

HyperSwap® support for raw disks

Some applications, such as DB2®, do not require a file system because they perform functions like logging, tracking data, and caching. These types of applications perform better when using raw disk I/O instead of using file I/O.

Normally, any raw disks that PowerHA® SystemMirror® manages require a physical volume ID (PVID). However, with HyperSwap® for PowerHA® SystemMirror® you can add raw disks without defining a PVID. This support includes raw disks that are shared in a single configuration or raw disks that replicate across sites.

HyperSwap® configuration settings PowerHA® SystemMirror® 7.1 Enterprise Edition for AIX®

Before you configure HyperSwap®, you must know the optimal settings to use.

The following table displays the optimal settings to use for a HyperSwap® configuration.

<i>HyperSwap® configuration settings</i>			
Name	Component	Value	Description
dyntrk	Protocol driver (fscsi)	Enabled	Provides transparent I/O recovery if the N_port ID for a device is changed. For example, if the Fibre Channel moves from one switch port to another port. This option is handled at the host bus adapter (HBA) level.
fc_err_recov	Protocol driver (fscsi)	fast_fail	Detects Fibre Channel problems between the switch and the storage device.
hcheck_interval	Disk driver (hdisk)	60	Sets the time interval when a health check request is sent to the storage device. The default setting is 60 seconds.

HyperSwap® zoning requirements

To use the HyperSwap® function in your cluster, you must configure zones correctly for system mirror groups, repository mirror groups, and user mirror groups.

The zoning requirements for specific types of mirror groups follow:

System mirror group

System mirror groups are local to a specific node. Therefore, the Peer-to-Peer Remote Copy (PPRC) disks that are associated with a system mirror group are only visible to the single node in the cluster. For example, if there are two sites (Site A and Site B) with one node at each site (Node A and Node B), you can create a system mirror group in Node A at Site A. In this example, you must verify that the disks in the system mirror group are only visible to Node A. The PPRC peers for these disks are only visible to Node A in Site A, even though the PPRC peer disk is in Site B.

Repository mirror group

The disks in a repository mirror group are accessible by the nodes that are listed in the associated site property list for the repository mirror group. You must verify that the HyperSwap® disk associated with the repository mirror group is accessible to all the nodes that are specified in the site property list.

User mirror group that is enabled for the HyperSwap® function

User mirror groups are associated to a set of one or more nodes in the cluster. The associated nodes can span across sites, and can have primary disks and auxiliary disks visible to each node in the cluster.

User mirror group for traditional Metro Mirror PPRC with in-band communication

These user mirror groups are not configured to use the HyperSwap® function. The associated nodes can span across sites, and can have primary disks visible to primary site nodes and auxiliary disks visible to auxiliary site nodes.

Installing HyperSwap® for PowerHA® SystemMirror®

Before installing HyperSwap® for PowerHA® SystemMirror®, you must verify that each user has root access to each node in the cluster. Your systems must also be running IBM® AIX® 6.1 with Technology Level 8, or later, or IBM® AIX® 7.1 with Technology Level 2, or later.

To install HyperSwap® for PowerHA® SystemMirror®, complete the following steps:

1. Insert the *PowerHA® SystemMirror® Enterprise Edition for AIX®* DVD into the DVD drive.
2. From the command line, enter `smit installp`.
3. From the System Management Interface Tool (SMIT) interface, select **Install and Update Software** > **Install Software**, and press Enter.
4. Select the `cluster.es.genxd.cmds` fileset and press Enter.

Configuring HyperSwap® for PowerHA® SystemMirror®

Before you can configure HyperSwap® for PowerHA® SystemMirror®, your IBM® DS8800 must be configured correctly.

Verify that your environment meets the following requirements before you configure HyperSwap® for PowerHA® SystemMirror®:

- A PowerHA® SystemMirror® cluster is defined.
- All PowerHA® SystemMirror® nodes are defined.
- All PowerHA® SystemMirror® sites are defined.
- All PowerHA® SystemMirror® resource groups and associated resources are configured and working properly.

Configuring a storage system for HyperSwap®

A *storage system* is a generic term that is used by PowerHA® SystemMirror® for an entity such as an IBM® DS8800 storage unit.

Configuring a storage system to use in-band communication for HyperSwap®

To use the HyperSwap® function, your storage system must use in-band communication. You might have to use the DS command-line interface (DSCLI) or the graphical management interface to configure a storage system to use in-band communication through a SCSI connection.

To configure a storage system for in-band communication, you must create a host profile on each storage system and associate the host profile with all PowerHA® SystemMirror® nodes in the cluster.

To configure a storage system to use in-band communication, complete the following steps:

1. Verify that the primary disk and auxiliary disk in the storage systems are correctly zoned to all PowerHA® SystemMirror® nodes.
2. Identify the host connection identifier that is associated with PowerHA® SystemMirror® nodes by entering the following command:

```
lshostconnect -dev storage system ID
```

3. For each node in the cluster, enter the following command to change the host profile to *IBM pSeries - AIX with Powerswap support*:

```
chhostconect -profile "IBM pSeries - AIX with Powerswap support" host connection ID
```

4. Define Peer-to-Peer Remote Copy (PPRC) paths and PPRC relationships for each disk that is part of the storage system that is used as a replicated resource. The PPRC paths are bidirectional. Thus, PPRC paths must be created from both directions. For example, there must be a PPRC path from site A to site B and from site B to site A. You must have configured PPRC paths before you start the mirror group.

5. Enter the following command to verify that both the primary and auxiliary disks are replicating correctly:

```
lsppic -l primary volume ID
```

The command displays various fields. The **State** field of the primary disk must display **Full-Duplex**. The value for the **Out-of-Sync Tracks** field must be **0**.

Configuring a PowerHA® SystemMirror® node to use in-band communication for HyperSwap®

After you configured a storage system to use in-band communication for, you must configure each node in the cluster to use in-band communication for HyperSwap®.

To configure a PowerHA® SystemMirror® cluster to use in-band communication, complete the following steps on each node in the cluster:

1. Enter the following command to configure all disks that are part of the storage system to use the AIX Path Control Module (PCM) driver:

```
manage_disk_drivers -d device -o AIX_AAPCM
```

2. Restart the node by entering the **shutdown -r** command.
3. SCSI reservations are not supported for disks that are used in a HyperSwap® mirror group. Enter the following command to verify that no disk reservations are set:

```
devsrsv -c query -l hdisk_name
```

The command returns the following information:

```
ODM Reservation Policy      : NO RESERVE
Device Reservation State    : NO RESERVE
```

4. Enter the following command to change the disk reservation policy to *no_reserve*:

```
chdev -a reserve_policy=no_reserve -l hdisk_number
```

5. Enter the following command to transform a disk that is associated with the primary disk to transparently route its I/O to the auxiliary disk:

```
chdev -a san_rep_cfg=migrate_disk -l hdisk_name -U
```

For example, if the primary disk is *hdisk1* and the auxiliary disk is *hdisk2*, you would run the following command:

```
chdev -a san_rep_cfg=migrate_disk -l hdisk1 -U
```

When you run this command the auxiliary disk *hdisk2* moves to a Defined state and all application I/O for *hdisk1* is transparently routed to the auxiliary disk.

Enter the following command to rollback *hdisk1* from being enabled by using the HyperSwap® function:

```
chdev -a san_rep_cfg=revert_disk -l hdisk1 -U
```

After you enter this command, both the primary and secondary disks are displayed for the node.

6. Enter the following command to verify that the disks for the storage system are configured correctly:

```
lsattr -E0 -a san_rep_cfg,san_rep_device,reserve_policy -l hdisk_number
```

The command returns following information:

```
#san_rep_cfg:san_rep_device:reserve_policy  
migrate_disk:yes:no_reserve
```

The value for `san_rep_cfg` attribute might be displayed as `none`, which is acceptable. However, the value for the `san_rep_device` attribute must be set to `yes`.

7. Use the SMIT interface to import volume groups on all PowerHA® SystemMirror® nodes by completing the following steps:
 - a. From the command line, enter `smit sysmirror`.
 - b. From SMIT interface, select **System Managemnet (C-SPOC) > Storage > Volume Groups > Import a Volume Group**, and press Enter.
 - c. From the list, select the volume group that you want to import, and press Enter.

Configuring a PowerHA® SystemMirror® node to use in-band communication for traditional Metro Mirror PPRC

After you configured a storage system to use in-band communication for, you must configure each node in the cluster to use in-band communication for traditional Metro Mirror Peer-to-Peer Remote Copy (PPRC).

To configure a PowerHA® SystemMirror® cluster to use in-band communication for traditional Metro Mirror PPRC, complete the following steps on each node in the cluster:

1. Enter the following command to configure all disks that are part of the storage system to use the AIX Path Control Module (PCM) driver:

```
manage_disk_drivers -d device -o AIX_AAPCM
```

2. Restart the node by entering the `shutdown -x` command.
3. SCSI reservations are not supported for disks that are used in a HyperSwap® mirror group. Enter the following command to verify that no disk reservations are set:

```
devfsrv -c query -l hdisk_name
```

The command returns the following information:

```
ODM Reservation Policy      : NO RESERVE  
Device Reservation State    : NO RESERVE
```

4. Enter the following command to change the disk reservation policy to `no_reserve`:

```
chdev -a reserve_policy=no_reserve -l hdisk_number
```

5. Enter the following command to change the attributes of the primary disk for all nodes in the primary site:

```
chdev -a san_rep_cfg=migrate_disk -l hdisk_name -U
```

Where `hdisk_name` is the name of the primary disk.

6. Enter the following command to change the attributes of the auxiliary disk for all nodes in the auxiliary site:

```
chdev -a san_rep_cfg=migrate_disk -l hdisk_name -U
```

Where `hdisk_name` is the name of the auxiliary disk.

7. Enter the following command to verify that the disks in the storage system are configured correctly:

```
lsattr -E0 -a san_rep_cfg,san_rep_device,reserve_policy -l hdisk_number
```

The command returns the following information:

```
#san_rep_cfg:san_rep_device:reserve_policy
migrate_disk:detected:no_reserve
```

The value for `san_rep_cfg` might display `none`, which is acceptable. However, the `san_rep_device` value must be detected or `yes`.

Adding a storage system for HyperSwap®

Before you can start using your storage system, you must add it to your environment.

To add a storage system, complete the following steps:

1. From the command line, enter `smit sysmirror`.
2. From the SMIT interface, select **Cluster Application and Resources > Resources > Configure DS8800-Metro Mirror (In-band) Resources > Configure Storage Systems > Add a Storage System**, and press Enter.
3. Enter the following information for the fields:
The fields for adding a storage system for HyperSwap®

<i>Add storage system fields</i>	
Field	Description
Storage system name	Enter the PowerHA® SystemMirror® name for the storage system. This name must be unique within the cluster definition, and cannot be more than 32 characters.
Site association	Press F4 to select an available site name from the list.
Vendor specific identifier	Press F4 to select the vendor-specific identifier from the list of storage systems that are connected to the local node.
WWNN	Press F4 to select the worldwide node name from a list of storage systems that are connected to the local node.

4. Verify that all fields are correct and press Enter.

Changing a storage system for HyperSwap®

You can change any of the properties of an existing storage system that you already added to your environment.

To change a storage system, complete the following steps:

1. From the command line, enter `smit sysmirror`.
2. From the SMIT interface, select **Cluster Application and Resources > Resources > Configure DS8800-Metro Mirror (In-band) Resources > Configure Storage Systems > Change/Show a Storage System**, and press Enter.
3. From the list, select the name of the storage system that you want to change and press Enter.
4. Enter the following information for the fields:
Change storage system fields

<i>Change storage system fields</i>	
Field	Description
Storage system name	Displays the current name for the PowerHA® SystemMirror® storage system. You cannot change this field.

Field	Description
New storage system name	Enter the new name for the PowerHA® SystemMirror® storage system. This name must be unique within the cluster definition. You cannot change this field in dynamic automatic reconfiguration (DARE) or when the underlying mirror group is being used. The name can have a maximum of 32 characters.
Site association	Press F4 to select an available site name from a list. You cannot change this field in DARE or when the underlying mirror group is being used. This field automatically discovers the changed site name from the PowerHA® SystemMirror®
Vendors identifier	Enter the unique vendor-specific ID for this storage system. For example, IBM®.2107-75ABTV1 is a storage identifier for a DS8800 storage system. Press F4 to display a list of the storage systems that are connected to the local node. You cannot change this field in DARE or when the underlying mirror group is being used.
WWNN	Press F4 to select the worldwide node name from a list of storage systems that are connected to the local node. You cannot change this field in DARE or when the underlying mirror group is being used.

5. Verify that all fields are correct and press Enter.

Removing a storage system for HyperSwap®

If you do not want a storage system to be included in your environment, you can remove it.

To remove a storage system, complete the following steps:

1. From the command line, enter `smit sysmirror`.
2. From the SMIT interface, select **Cluster Application and Resources > Resources > Configure DS8800-Metro Mirror (In-band) Resources > Configure Storage Systems > Remove a Storage System**, and press Enter.
3. Enter the storage system name you want to remove, and press Enter.
4. Press Enter to confirm that the selected storage system is the one you want to remove.

Configuring user mirror groups for HyperSwap®

You can add and change user mirror groups by using HyperSwap® functions.

Adding a user mirror group for HyperSwap®

Before you can start using a user mirror group, you must add it to your environment.

Review the following information if you are adding a user mirror group to a single node configuration that uses the HyperSwap® function:

- You cannot create a user mirror group if the HyperSwap® field is set to **Disable**.
- You cannot change the HyperSwap® field after the mirror group is added.
- The resource group must contain a single node in the node list. Otherwise, the cluster verification process fails.
- The **Recovery action** field must be set to **Manual**.

- You should use the tie breaker option for the split and merge policies.
- You can have up to 64 volume groups in a mirror group.
- You can have up to 1024 raw disks in a mirror group.

To add a user mirror group, complete the following steps:

1. From the command line, enter `smit sysmirror`.
2. From the SMIT interface, select **Cluster Application and Resources > Resources > Configure DS8800-Metro Mirror (In-band) Resources > Configure Mirror Groups > Add a Mirror Group > User**, and press Enter.
3. Enter the following information for the fields:
Adding user mirror group fields

<i>Adding user mirror group fields</i>	
Fields	Description
Mirror group name	Enter the name for the mirror group.
Volume groups	Select the volume groups that you want to include in the mirror group.
Raw disks	Select the raw disks that you want to include in the mirror group.
HyperSwap®	Select Enable to use HyperSwap® for the mirror group. Select Disable , the default value, to not use HyperSwap® for the mirror group and to use a traditional Metro Mirror PPRC with in-band communication. Verify that your environment has the correct zoning configurations before you change this field from Enable to Disable or from Disable to Enable . You cannot change this field in dynamic automatic reconfiguration (DARE).
Consistency group	Select Enable , the default, to use the consistency group property for the mirror group. Select Disable to not use the consistency group property for the mirror group.
Unplanned HyperSwap® timeout (in seconds)	Enter the timeout value for unplanned HyperSwap®. This value represents how long a connection remains unavailable before an unplanned HyperSwap® site failover occurs. The default value is 60.
HyperSwap® priority	Select Low , Medium , or High as the HyperSwap® priority. The default setting is Medium .
Recovery action	Specify the disaster recovery policy that PowerHA® SystemMirror® will use in case of a node or site failover. Select Automatic to use a split or merge policy that will automatically resync the node or site. Select Manual if you do not want use a policy and to use the recommend action displayed in the <code>/var/hacmp/log/hacmp.out</code> file to correct the failover.
Re-sync action	Select Automatic to automatically resynchronize the node or site if the replicated volumes or network paths recover from the failover. Select Manual to display the recommended action for HyperSwap® enabled mirror groups by using the <code>exrpt</code> command, and for HyperSwap® disabled mirror groups the recommended action is displayed in the <code>/var/hacmp/log/hacmp.out</code> file.

- Verify that all fields are correct and press Enter.

Changing a user mirror group for HyperSwap®

You can change any of the properties of an existing user mirror group that you already added to your environment.

If you change settings during dynamic reconfiguration (DARE), you must verify and synchronize the cluster. If a failure occurs during the verification and synchronization process, you must manually change the settings back to the previous values and run the verification and synchronization process again.

To change a user mirror group, complete the following steps:

- From the command line, enter `smit sysmirror`.
- From the SMIT interface, select **Cluster Application and Resources > Resources > Configure DS8800-Metro Mirror (In-band) Resources > Configure Mirror Groups > Change/Show a Mirror Group > User**, and press Enter.
- Enter the following information for the fields:

<i>Changing user mirror group fields</i>	
Field	Description
Mirror group name	Enter the name for the mirror group.
New mirror group name	Enter the new name that you want to use for the mirror group. This new name can have a maximum of 32 characters.
Volume groups	Select the volume groups that you want to include in the mirror group.
Raw disks	Select the raw disks that you want to include in the mirror group.
Associated storage systems	Displays the name of the underlying storage systems on which the mirror group exists. You cannot edit this field because the displayed information is automatically discovered based on the disks and volume groups that are added to the mirror group.
HyperSwap®	Select Enable to use HyperSwap® for the mirror group. Select Disable , the default value, to not use HyperSwap® for the mirror group and to use a traditional Metro Mirror PPRC with in-band communication. Verify that your environment has the correct zoning configurations before you change this field from Enable to Disable or from Disable to Enable . You cannot change this field in dynamic automatic reconfiguration (DARE).
Consistency group	Select Enable , the default, to use the consistency group property for the mirror group. Select Disable to not use the consistency group property for the mirror group.
Unplanned HyperSwap® timeout (in seconds)	Enter the timeout value for unplanned HyperSwap®. This value represents how long a connection remains unavailable before an unplanned HyperSwap® site failover occurs. The default value is 60.
HyperSwap® priority	Select Low , Medium , or High as the HyperSwap® priority. The default setting is Medium .
Recovery action	Specify the disaster recovery policy that PowerHA® SystemMirror® will use in case of a node or site failover. Select Automatic to use a split or merge policy that will automatically resync

Field	Description
	the node or site. Select Manual if you do not want use a policy and to use the recommend action displayed in the <code>/var/hacmp/log/hacmp.out</code> file to correct the failover.
Re-sync action	Select Automatic to automatically resynchronize the node or site if the replicated volumes or network paths recover from the failover. Select Manual to display the recommended action for HyperSwap® enabled mirror groups by using the errpt command, and for HyperSwap® disabled mirror groups the recommended action is displayed in the <code>/var/hacmp/log/hacmp.out</code> file.

4. Verify that all fields are correct and press Enter.

Support for Metro Global Mirror configuration

PowerHA® SystemMirror® Enterprise Edition 7.2.7, or later, provides DS8800 Metro Global Mirror (MGM) support in integration with Copy Services Manager (CSM).

Before you use the DS8800 MGM support function, you must enable HyperSwap® configuration between the primary storage and secondary storage and you must configure Global Mirror relationship with the third storage. In PowerHA® SystemMirror®, HyperSwap® is responsible for a failover from the primary storage to secondary storage. However, the recovery operation must be performed manually by the user by using CSM. The Copy Services Manager handles resynchronization of the Metro and Global Mirror relationships and also handles the recovery of Global Mirror copy.

To use the DS8800 MGM support function, perform the following steps:

1. Acquire the CSM license and three storages.
2. Set values of Recovery action and Re-sync action field to **Manual**. For more information, see [“Changing a user mirror group for HyperSwap” on page 151](#).
3. Set the value of the AUTO_RESYNC attribute to false in the `/usr/es/sbin/cluster/xd_generic/xd_ds8k_mm/ds8k_inband_mm.cfg` config file.

Configuring a system mirror group for HyperSwap®

You can add or change a system mirror group by using HyperSwap® functions.

Adding a system mirror group for HyperSwap®

Before you can start using a system mirror group, you must add it to your environment.

To add a system mirror group, complete the following steps:

1. From the command line, enter `smit sysmirror`.
2. From the SMIT interface, select **Cluster Application and Resources > Resources > Configure DS8800-Metro Mirror (In-band) Resources > Configure Mirror Groups > Add a Mirror Group > System**, and press Enter.
3. Enter the following information for the fields:

<i>Adding system mirror group fields</i>	
Field	Description
Mirror group name	Enter the name for the mirror group.
Volume groups	Select the volume groups that you want to include in the mirror group.

Field	Description
Raw disks	Select the raw disks that you want to include in the mirror group.
Node name	Displays the name of the local node. This field is automatically populated and cannot be changed.
HyperSwap®	This field must be set to Enabled .
Consistency group	This field must be set to Enabled .
Unplanned HyperSwap® timeout (in seconds)	Enter the timeout value for unplanned HyperSwap®. This value represents how long a connection remains unavailable before an unplanned HyperSwap® site failover occurs. The default value is 60.
HyperSwap® priority	Select Low , Medium , or High as the HyperSwap® priority. The default setting is High .

- Verify that all fields are correct and press Enter.

Result

Note: The resynchronization function occurs automatically for system mirror groups.

Changing a system mirror group for HyperSwap®

You can change the properties of an existing system mirror group that you already added to your environment.

To change a system mirror group, complete the following steps:

- From the command line, enter `smit sysmirror`.
- From the SMIT interface, select **Cluster Application and Resources > Resources > Configure DS8800-Metro Mirror (In-band) Resources > Configure Mirror Groups > Change/Show a Mirror Group > System**, and press Enter.
- Enter the following information for the fields:

<i>Changing system mirror group</i>	
Fields	Description
Mirror group name	Enter the name for the mirror group.
New mirror group name	Enter the new name that you want to use for the mirror group. This new name can have a maximum of 32 characters.
Volume groups	Select the volume groups that you want to include in the mirror group.
Raw disks	Select the raw disks that you want to include in the mirror group.
Associated storage systems	Displays the name of the underlying storage systems on which the mirror group exists. You cannot edit this field because the displayed information is automatically discovered based on the disks and volume groups that are added to the mirror group.
Node name	Displays the name of the local node. This field is automatically populated and cannot be changed.
HyperSwap®	You cannot change this field.
Consistency group	You cannot change this field.

Fields	Description
Unplanned HyperSwap® timeout (in seconds)	Enter the timeout value for unplanned HyperSwap®. This value represents how long a connection remains unavailable before an unplanned HyperSwap® site failover occurs. The default value is 60.
HyperSwap® priority	Select Low , Medium , or High as the HyperSwap® priority. The default setting is High .

4. Verify that all fields are correct and press Enter.

Configuring a cluster repository mirror group for HyperSwap®

You can add or change a cluster repository mirror group by using HyperSwap® functions.

Adding a cluster repository mirror group for HyperSwap®

Before you can start using a cluster repository mirror group, you must add it to your environment.

PowerHA® SystemMirror® implements the HyperSwap® function by using Cluster Aware AIX® (CAA). CAA requires that a repository disk is defined in your environment to create a cluster across multiple sites. You must first configure the cluster to function across multiple sites before you implement the HyperSwap® function.

To configure the HyperSwap® function with a repository disk, you must identify a new disk pair that uses the HyperSwap® function.

To add a cluster repository mirror group, complete the following steps:

1. From the command line, enter `smit sysmirror`.
2. From the SMIT interface, select **Cluster Application and Resources > Resources > Configure DS8800-Metro Mirror (In-band) Resources > Configure Mirror Groups > Add a Mirror Group > Cluster Repository**, and press Enter.
3. Enter the following information for the fields:

<i>Adding cluster repository mirror group fields</i>	
Field	Description
Mirror group name	Enter the name for the mirror group.
Site name	Select the site that is used as the cluster repository disk and contains the mirror group. If you select a stretched cluster, both site names are displayed. If you select a linked cluster, the local site name is displayed.
Non-HyperSwap® disk	Select the name of the non-HyperSwap® disk that is used as the cluster repository disk and contains the mirror group. If the repository disk is on a non-HyperSwap® disk select that disk. This field is optional if the repository disk is already on the HyperSwap® enabled disk.
HyperSwap® disk	Select the name of the Peer-to-Peer Remote Copy (PPRC) disk that is used as the cluster repository and contains the mirror group.
HyperSwap®	This field must be set to Enabled .
Consistency group	This field must be set to Enabled .
Unplanned HyperSwap® timeout (in seconds)	Enter the timeout value for unplanned HyperSwap®. This value represents how long a connection remains unavailable before an unplanned HyperSwap® site failover occurs. The default value is 60.

Field	Description
HyperSwap® priority	Select Low , Medium , or High as the HyperSwap® priority. The default setting is High .
Re-sync action	Select Automatic to automatically resynchronize the node or site if the replicated volumes or network paths recover from the failover. Select Manual to display the recommended action for HyperSwap® enabled mirror groups by using the errpt command, and for HyperSwap® disabled mirror groups the recommended action is displayed in the <code>/var/hacmp/Log/hacmp.out</code> file.

- Verify that all fields are correct and press Enter.

Result

Note: During the cluster verification and synchronization process, if a repository disk is not configured to use the HyperSwap® function it moves to a disk that is configured to use the HyperSwap® function.

Changing a cluster repository mirror group for HyperSwap®

You can change any of the properties of an existing cluster repository mirror group that you already added to your environment.

To change a cluster repository mirror group, complete the following steps:

- From the command line, enter `smit sysmirror`.
- From the SMIT interface, select **Cluster Application and Resources > Resources > Configure DS8800-Metro Mirror (In-band) Resources > Configure Mirror Groups > Change/Show a Mirror Group > Cluster Repository**, and press Enter.
- Enter the following information for the fields:

<i>Changing cluster repository mirror group fields</i>	
Field	Description
Mirror group name	Enter the name for the mirror group.
New mirror group name	Enter the new name that you want to use for the mirror group. This new name can have a maximum of 32 characters.
Site name	Select the site that is used as the cluster repository disk and contains the mirror group. If you select a stretched cluster, both site names are displayed. If you select a linked cluster, the local site name is displayed.
Non-HyperSwap® disk	Select the name of the non-HyperSwap® disk that is used as the cluster repository disk and contains the mirror group. If the repository disk is on a non-HyperSwap® disk select that disk. This field is optional if the repository disk is already on the HyperSwap® enabled disk.
HyperSwap® disk	Select the name of the Peer-to-Peer Remote Copy (PPRC) disk that is used as the cluster repository and contains the mirror group.
Associated storage systems	Displays the name of the underlying storage systems on which the mirror group exists. You cannot edit this field because the displayed information is automatically discovered based on

Field	Description
	the disks and volume groups that are added to the mirror group.
HyperSwap®	You cannot change this field.
Consistency group	You cannot change this field.
Unplanned HyperSwap® timeout (in seconds)	Enter the timeout value for unplanned HyperSwap®. This value represents how long a connection remains unavailable before an unplanned HyperSwap® site failover occurs. The default value is 60.
HyperSwap® priority	Select Low , Medium , or High as the HyperSwap® priority. The default setting is High .
Re-sync action	Select Automatic to automatically resynchronize the node or site if the replicated volumes or network paths recover from the failover. Select Manual to display the recommended action for HyperSwap® enabled mirror groups by using the errpt command, and for HyperSwap® disabled mirror groups the recommended action is displayed in the <code>/var/hacmp/log/hacmp.out</code> file.

4. Verify that all fields are correct and press Enter.

Removing a mirror group for HyperSwap®

You can remove a user, system, or cluster repository mirror group by using SMIT.

To remove a user, system, or cluster repository mirror group, complete the following steps:

1. From the command line, enter `smit sysmirror`.
2. From the SMIT interface, select **Cluster Application and Resources > Resources > Configure DS8800-Metro Mirror (In-band) Resources > Configure Mirror Groups > Remove a Mirror Group**, and press Enter.
3. Select the mirror group you want to remove and press Enter.
4. Press Enter to confirm that the selected mirror group is the one you want to remove.

AFTER COMPLETING THE TASK

Note: After you remove a mirror group, the HyperSwap® configuration for the storage systems or volume groups is also removed from the PowerHA® SystemMirror® database.

Configuring a single node to use the HyperSwap® function

You can configure a single node to use the HyperSwap® function in a PowerHA® SystemMirror® cluster.

You can use the HyperSwap® function in a single node configuration by using a minimal amount of hardware. However, your environment must have a primary storage system and an auxiliary storage system.

In a single node configuration that uses the HyperSwap® function, you can only have one mirror group for the associated node. To create a mirror group in this configuration, your cluster cannot be configured to use multiple sites. You can use the HyperSwap® function to change a mirror group that is associated with a single node.

To configure a single node to use the HyperSwap® function, complete the following steps:

1. From the command line, enter `smit sysmirror`.
2. From the SMIT interface, select **Cluster Application and Resources > Resources > Configure DS8800-Metro Mirror (In-band) Resources > Cluster Wide HyperSwap® Policies**, and press Enter.
3. From **Single node HyperSwap®** list, select **Enabled**.

4. Verify that all fields are correct and press Enter.

AFTER COMPLETING THE TASK

Note: You cannot convert the single node configuration that uses the HyperSwap® function to a multisite configuration. If you want to create a cluster with multiple sites, you must first remove the existing single node configuration from the cluster.

Managing mirror groups for HyperSwap®

After you add a mirror group to the cluster, you can manage it by using the **System Management (C-SPOC)** panel in SMIT.

Managing user mirror groups for HyperSwap®

User mirror groups contain disks that are used for PowerHA® SystemMirror® application workloads. User mirror groups are only active on a node if the corresponding resource group is also active.

To manage user mirror groups, complete the following steps:

1. From the command line, enter `smit cspoc`.
2. From the SMIT interface, select **Storage > Manage Mirror Groups > User Mirror Groups**, and press Enter.
3. Complete the following fields, and press Enter.

<i>User mirror group fields</i>	
Field	Description
Mirror groups	Enter the name or names of one or more mirror groups for which you want to change the type of operation.
Operation	Press F4 and select from the following options: Swap Select this option to perform a planned HyperSwap® on the selected mirror group. You can use this option only if the mirror group is enabled for HyperSwap® and the mirror group is part of a resource group that is currently online. Refresh Select this option to rediscover the latest replication path from the underlying storage for the selected mirror group. Show Active Path Select this option to display the currently active site and storage system for the selected mirror group. You can use this option only if the mirror group is enabled for HyperSwap® and the mirror group is part of a resource group that is currently online.

4. Verify that all fields are correct and press Enter.

Managing system mirror groups for HyperSwap®

System mirror groups are a group of disks and volume groups that are critical to a single node that is running the AIX® operating system. An example of a system mirror group is a rootvg mirror group or a mirror group that contains the paging space.

A system mirror group is only active on a single node; therefore, any HyperSwap® request does not require any coordination with other nodes in the cluster.

To manage a system mirror group, complete the following steps:

1. From the command line, enter **smit cspoc**.
2. From the SMIT interface, select **Storage > Manage Mirror Groups > System Mirror Group**, and press Enter.
3. Complete the following fields, and press Enter.

<i>System mirror group fields</i>	
Field	Description
Mirror groups	Enter the name or names of one or more mirror groups for which you want to change the type of operation.
Node name	Enter the name of one or more nodes that you want to change to a different storage system.
Operation	<p>Press F4 and select from the following options:</p> <p>Swap</p> <p>Select this option to perform a planned HyperSwap® on the selected mirror group. You can use this option only if the mirror group is enabled for HyperSwap® and the mirror group is part of a resource group that is currently online.</p> <p>Refresh</p> <p>Select this option to rediscover the latest replication path from the underlying storage for the selected mirror group.</p> <p>Show Active Path</p> <p>Select this option to display the currently active site and storage system for the selected mirror group. You can use this option only if the mirror group is enabled for HyperSwap® and the mirror group is part of a resource group that is currently online.</p>

4. Verify that all fields are correct and press Enter.

Managing cluster repository mirror groups for HyperSwap®

Cluster repository mirror groups contain disks that are used for the Cluster Aware AIX® (CAA) cluster repository. These disks are shared between multiple nodes, but they cannot be shared across sites.

To manage a cluster repository mirror group, complete the following steps:

1. From the command line, enter **smit cspoc**.
2. From the SMIT interface, select **Storage > Manage Mirror Groups > Cluster Repository Mirror Group**, and press Enter.
3. Complete the following fields, and press Enter.

<i>Cluster repository mirror group fields</i>	
Field	Description
Mirror groups	Enter the name or names of one or more mirror groups for which you want to change the type of operation.
Operation	<p>Press F4 and select from the following options:</p> <p>Swap</p> <p>Select this option to perform a planned HyperSwap® on the selected mirror group. You can use this option only if the mirror group is enabled for HyperSwap® and the mirror group is part of a resource group that is currently online.</p> <p>Refresh</p> <p>Select this option to rediscover the latest replication path from the underlying storage for the selected mirror group.</p> <p>Show Active Path</p> <p>Select this option to display the currently active site and storage system for the selected mirror group. You can use this option only if the mirror group is enabled for HyperSwap® and the mirror group is part of a resource group that is currently online.</p>

4. Verify that all fields are correct and press Enter.

Managing mirror groups by nodes for HyperSwap®

You can manage mirrors groups that belong to different nodes in the cluster.

To manage mirror groups that belong to different nodes, complete the following steps:

1. From the command line, enter `smi t cspoc`.
2. From the SMIT interface, select **Storage > Manage Mirror Groups > Manage Mirror Groups by Nodes**, and press Enter.
3. Complete the following fields, and press Enter.

<i>Mirror groups by nodes fields</i>	
Field	Description
Node name	Enter the name of one or more nodes that you want to change to a different storage system.
Include system mirror groups	Select yes to include system mirror groups. If you select yes , the system mirror groups must be active on the selected nodes. Select no to not include system mirror groups.
Operation	<p>Press F4 and select from the following options:</p> <p>Swap</p> <p>Select this option to perform a planned HyperSwap® on the selected mirror group. You can use this option only if the mirror group is enabled for HyperSwap® and the mirror group is part of a resource group that is currently online.</p>

Field	Description
	<p>Refresh</p> <p>Select this option to rediscover the latest replication path from the underlying storage for the selected mirror group.</p> <p>Show Active Path</p> <p>Select this option to display the currently active site and storage system for the selected mirror group. You can use this option only if the mirror group is enabled for HyperSwap® and the mirror group is part of a resource group that is currently online.</p>

4. Verify that all fields are correct and press Enter.

Managing mirror groups by site for HyperSwap®

You can manage all mirrors groups that belong to a site.

To manage user mirror groups that belong to a site, complete the following steps:

1. From the command line, enter `smit cspoc`.
2. From the SMIT interface, select **Storage > Manage Mirror Groups > Manage Mirror Groups by Sites**, and press Enter.
3. Complete the following fields, and press Enter.

<i>Mirror groups by nodes fields</i>	
Field	Description
Site name	Enter the site name for the mirror groups you want manage. If you want to manage a stretched cluster, you must enter more than one site name. If you want to manage a linked cluster, you must manage each site separately.
Include system mirror groups	Select yes to include system mirror groups. If you select yes , the system mirror groups must be active on the selected site. Select no to not include system mirror groups.
Include cluster repository mirror groups	Select yes to include cluster repository mirror groups. If you select yes , the cluster repository mirror groups must be active on the selected site. Select no to not include cluster repository mirror groups.
Operation	<p>Press F4 and select from the following options:</p> <p>Swap</p> <p>Select this option to perform a planned HyperSwap® on the selected mirror group. You can use this option only if the mirror group is enabled for HyperSwap® and the mirror group is part of a resource group that is currently online.</p> <p>Refresh</p> <p>Select this option to rediscover the latest replication path from the underlying storage for the selected mirror group.</p> <p>Show Active Path</p>

Field	Description
	Select this option to display the currently active site and storage system for the selected mirror group. You can use this option only if the mirror group is enabled for HyperSwap® and the mirror group is part of a resource group that is currently online.

- Verify that all fields are correct and press Enter.

Troubleshooting HyperSwap® configurations

To successfully troubleshoot a HyperSwap® configuration, you must know where the log files are located and the commands that you can use to fix common problems.

HyperSwap® log files

The following table identifies the different log files that you can use to troubleshoot various HyperSwap® problems.

Log files to troubleshoot HyperSwap® problems

<i>Log files to troubleshoot HyperSwap® problems</i>		
Name	Location	Description
hacmp.out	/var/hacmp/log/hacmp.out	Displays messages that are related to detection, migration, termination, and execution of a PowerHA® SystemMirror® cluster for an application.
clutils.log	/var/hacmp/log/clutils.log	Displays the results of the automatic verification that runs on a specified PowerHA® SystemMirror® cluster node every 24 hours.
clxd_debug.log	/var/hacmp/xd/log/clxd_debug.log	Displays information about the clxd daemon.

HyperSwap® configurations use kernel extension. Thus, you can view error or warning messages from this kernel extensions by using the **errpt** command. The following example shows HyperSwap® configuration when you run the **errpt** command:

```
# errpt
IDENTIFIER  TIMESTAMP    T C RESOURCE_NAME DESCRIPTION
F6D1B11B   0705061213  P S phake          PROGRAM INTERRUPT
F6D1B11B   0705061213  P S phake          PROGRAM INTERRUPT
F6D1B11B   0705061213  P S phake          PROGRAM INTERRUPT
F6D1B11B   0705061213  P S phake          PROGRAM INTERRUPT
F6D1B11B   0705061213  P S phake          PROGRAM INTERRUPT
F6D1B11B   0705061213  P S phake          PROGRAM INTERRUPT
F6D1B11B   0705061213  P S phake          PROGRAM INTERRUPT
63B1A1E6   0705061213  I H pha_1065458178 PPRC Replication Path Recovered
DCB47997   0705060813  T H hdisk71        DISK OPERATION ERROR
4BD7BBF6   0705060813  T H pha_1065458178 PPRC Replication Path Failed
BFCFD000   0705060813  T H hdisk71        PPRC Device Suspended
```

You can also configure kernel extension to create debug logs in the `/etc/syslog.conf` file by completing the following steps:

- In the `/etc/syslog.conf` file, add the following line:

```
kernel.debug /var/hacmp/xd/log/syslog.phake rotate size 500k files 7
```

- Create a file that is called `syslog.phake` in the `/var/hacmp/xd/log` directory.

- Refresh the **syslogd** daemon.

Note: The debug logs are also logged in the console.

Troubleshooting common problems

The following common problems might occur when are configuring the HyperSwap® function:

- After a system restarts, the **clxd** daemon is not operating. To troubleshoot this problem, you must start the **clxd** daemon by completing the following steps:

- From the command, run the following command:

```
lssrc -s clxd
```

- From the command line, run the following command:

```
tail -f /var/hacmp/xd/log/clxd_debug.log
```

- From the command line, run the following command:

```
lspprc -Ao
```

- From the command line, run the following command:

```
startsrc -s clxd
```

If the **clxd** daemon is still not starting, contact IBM® support.

- The following error message is displayed when you are creating a mirror group:

```
Failed to changed XD Mirror Group 'User_RD_MG2'. errno=22 msg=No Storage System name added
OR PPRC path not established properly
OR HyperSwap is not enabled for disks. Check clxd_debug.log for details
```

This message means that you did not correctly configure the disks that you are using for the mirror group that uses the HyperSwap® function. To troubleshoot this problem, review the information in the [“Configuring a storage system for HyperSwap” on page 145](#) topic.

- The following error message is displayed when you are configuring a repository mirror group and HyperSwap® moves the mirror group to a disk that is configured to use the HyperSwap® function:

```
ERROR: Failed to run 'chrepos'.
Follow recommended user actions to recover from here
(if needed in case caavg_private has not been seen in some or all nodes).
```

To troubleshoot this problem, you can use the **chrepos -r** command on the disk.

- The disks for the mirror group are in a halt state. You can run the **lspprc -A** command to display the status of your disks. The following example output is displayed when you run the **lspprc -A** command:

#	hdisk#	PPRC state	Primary path group ID	Secondary path group ID	Primary Storage WWNN	Secondary Storage WWNN
	hdisk104	Active	-1	0	500507630affc16b	5005076308ffc6d4
	hdisk14	Halt	0(s)	1	500507630affc16b	5005076308ffc6d4
	hdisk15	Halt	0(s)	1	500507630affc16b	5005076308ffc6d4
	hdisk16	Active	0(s)	-1	500507630affc16b	5005076308ffc6d4
	hdisk17	Active	0(s)	-1	500507630affc16b	5005076308ffc6d4
	hdisk18	Active	0(s)	1	500507630affc16b	5005076308ffc6d4

To move a disk from the halt state to an active state, you must reboot the node or bring the resource group for the associated disk online.

Failover rehearsal

Failover rehearsal provides mirroring status for different storage devices in a multi-site cluster.

PowerHA® SystemMirror® 7.2.2, or later, provides you with the following information when you move resource to other sites:

- Storage mirroring availability between storage sites
- Direction of mirroring between sites
- Types of mirroring
- Status of mirroring

Note: Failover rehearsal feature is available only in SAN Volume Controller (SVC) storage devices in a linked cluster environment.

```
clmgr manage cluster rehearsal \  
[ RESOURCE_GROUP="<RG#1> [,<RG#2>,..." ] ]
```

Example

To check the mirroring status of storage devices included in RG1 and RG2 resource groups, enter the following command:

```
clmgr manage cluster rehearsal RESOURCE_GROUP="RG1 RG2"
```

Configuring Fibre Channel disks

The `fc_err_recov` parameter has a default value of `delayed_fail` that causes extensive retries in response to problems caused while communicating with the storage subsystem. For the `fc_err_recov` parameter, you can specify the delay in the response time by a value of 10 minute increments before an unrecoverable error is recognized. You can use the `delayed_fail` value if there are not any alternative communication paths to the storage subsystem, or if the storage subsystem holds a rootvg disk, a paging space disk, or other system critical information.

The `fast_fail` value for the `fc_err_recov` parameter causes fewer retries and more timely recognition of an unrecoverable error by the system than the `delayed_fail` value. Depending on the problem, the `fast_fail` value allows the AIX® MPIO device driver, the AIX® Logical Volume Manager (LVM), or PowerHA® SystemMirror® to react quickly to a communication problem with the storage subsystem. You can use the `fast_fail` value if the configuration is likely to experience solid failures, such as severing of the Fibre Channel links or loss of power to the storage subsystem.

Before you change the values for the `fc_err_recov` parameter, you must determine the appropriate Fibre Channel SCSI device or devices by completing the following steps:

1. To list all the disk and the owning volume groups that are known to the subsystem, run the **lspv** command. The names of the disks in the volume groups on the storage subsystem are managed by PowerHA® SystemMirror®.
2. To determine the name of the Fibre Channel SCSI adapter that is used to communicate with the devices, run the **lsdev -C -l <disk_name> -F parent** command for each disk.

To change the values for the `fc_err_recov` parameter by using the SMIT interface, complete the following steps:

1. From the command line, enter **smit devices**.
2. From the SMIT interface, select **FC adapter > FC SCSI protocol device > Change/Show characteristics of an FC SCSI device**, and press Enter.
3. Select the Fibre Channel SCSI device that you want to change.
4. Change the **FC Fabric Event Error RECOVERY Policy** field to the desired value, and press Enter.
5. In the **Apply change to DATABASE only** field, enter **yes**.

6. Reboot the node.

Troubleshooting manual recovery for storage subsystems

PowerHA® SystemMirror® Enterprise Edition Version 7.1.2, or later, supports various storage subsystems that provide high availability for applications and services by monitoring for failures and implementing an automatic recovery for the failure. The storage subsystems use various replication technologies to manage the replication of data between a primary and auxiliary data center.

If the storage subsystem is online and available, PowerHA® SystemMirror® Enterprise Edition 7.1.2, or later, can automatically manage the replicated data during the failover and fallback. However, the following scenarios explain in what circumstances PowerHA® SystemMirror® Enterprise Edition does not automatically manage the replicated data and when manual intervention is required:

- PowerHA® SystemMirror® Enterprise Edition cannot determine the status of the storage subsystem, storage links, or device groups. In this scenario, PowerHA® SystemMirror® Enterprise Edition stops the cluster event processing and displays the corrective actions in the `/var/hacmp/log/hacmp.out` log file. To troubleshoot storage subsystem problems, review the information in the RECOMMENDED USER ACTIONS section in the `/var/hacmp/log/hacmp.out` log file.
When the storage subsystem is brought back online, you must manually resume cluster event processing by selecting **Problem Determination Tools > Recover PowerHA® SystemMirror® From Script Failure** from the SMIT interface.
- A failover occurs for a partitioned cluster across different sites. The primary and auxiliary partitions begin to write data to a local storage subsystem. When the primary partition recovers and the storage links are brought back online, you must determine whether the data from the two sites can be merged or if one site's data can replace the other site's data. In this scenario, you do not want PowerHA® SystemMirror® Enterprise Edition to use the automatic recovery function.
To configure PowerHA® SystemMirror® Enterprise Edition to use manual recovery, complete the following steps:
 - a. From the command line, enter `smit sysmirror`.
 - b. In the SMIT interface, select **Cluster Applications and Resources > Resources**.
 - c. Select the storage subsystem that you want to configure for manual recovery.
 - d. From the **Recovery Action** field, select **MANUAL**.
- If an outage affects all the mirror links between the source site and target site, IBM FlashSystem® A9000 or IBM® XIV® Storage System on the primary storage might not fail over to the secondary storage. In this scenario, the mirror consistency group relationship is still active, but the `mirror_switch_roles` command fails. If you want the mirror consistency group to fail over to the secondary storage, you must manually perform the following steps:
 - a. Deactivate the mirror consistency group relationship on the primary storage by running the following command:

```
mirror_deactivate -y cg=cgname
```
 - b. On the secondary storage, change the role of the consistency group to Primary by running the following command:

```
mirror_change_role -y cg=cgname role=Master
```
 - c. On the primary storage, change the role of the consistency group to Slave by running the following command:

```
mirror_change_role -y cg=cgname role=Slave
```

Note: When you change the role of the consistency group on the secondary storage, the volume group on the secondary storage can be in the Varied ON state as the mirror volumes on the secondary storage are no longer in read only mode. When you run the `mirror_change_role` command on the primary storage, a time delay occurs because the I/

O activity is broken on the host. To avoid the time delay, stop the disk I/O activity on the host before you run the **mirror_change_role** command.

Other: The input and output syntax of command line interface (CLI) commands uses the legacy terminology of "Master", "SMaster", and "Slave" volumes, which in any documentation except the CLI reference, are referred to as "Primary", "Secondary", and "Tertiary". This inconsistency is a necessary compromise, required to avoid changes to older CLI commands that are in customer use, and also to keep the CLI terminology consistent across the board. The new terminology helps emphasize the commonality between the more recent functions of Multi-site HA/DR, high availability (HyperSwap), and the disaster recovery (Synchronous and Asynchronous mirroring) ones. It is used outside the CLI reference, where broader concepts can be explained.

Notices

This information was developed for products and services offered in the US.

IBM® may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM® representative for information on the products and services currently available in your area. Any reference to an IBM® product, program, or service is not intended to state or imply that only that IBM® product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM® intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM® may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

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

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM® may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM® product and use of those websites is at your own risk.

IBM® may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

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 Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

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.

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

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM® has not tested those products and cannot confirm the

accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

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

All IBM® prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM®, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM®, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM® shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (your company name) (year).

Portions of this code are derived from IBM® Corp. Sample Programs.

© Copyright IBM® Corp. _enter the year or years_.

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 the 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> 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](http://www.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](http://www.ibm.com/legal/copytrade.shtml) at www.ibm.com/legal/copytrade.shtml.

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

Index

- active-active workloads [138](#)
- active-active workloads across sites [138](#)
- active-active workloads within a site [139](#)
- configuration settings [144](#)
- Configuring [145](#)
- Configuring a cluster repository mirror group [154](#)
- configuring a single node [156](#)
- Configuring a storage system [145](#)
- Configuring system mirror groups [152](#)
- Configuring user mirror groups [149](#)
- in-band communication [145](#), [146](#)
- Installing [145](#)
- Managing cluster repository mirror groups [158](#)
- Managing mirror groups by nodes [159](#)
- Managing mirror groups by sites [160](#)
- Managing system mirror groups [158](#)
- Managing user mirror groups [157](#)
- Planned [132](#)
- Planning [142](#)
- PowerHA SystemMirror [131](#)
- Removing a mirror group [156](#)
- resync function [140](#)
- support for DS8800 [143](#)
- support for raw disks [143](#)
- SVC policy based replication management
 - overview [93](#)
- timeout [132](#)
- troubleshooting [161](#)
- Unplanned [133](#)
- zoning requirements [144](#)

A

adding

- DSCLI management
 - replicated resource [49](#)

C

CD-ROM [17](#)

changing

- consistency group
 - site configuration [60](#)
- Direct management
 - replicated resource configuration [38](#), [37](#)
 - site configuration [37](#)
- DSCLI management
 - replicated resource configuration [52](#)
 - site configuration [52](#)
- SVC management
 - relationships [88](#)
 - resources [89](#), [99](#)

cl_verify_svcpprc_config command

- SVC management [92](#)

cllsrelationship command

- SVC management [92](#)

cllssvc command

- SVC management [91](#)

cllssvcpprc command

- SVC management [91](#)

cluster

configuring

- Direct Management [32](#)
- SVC Management [82](#), [93](#)
- SVC Management for flash storage [95](#)

for Metro Mirror [11](#)

starting

- Direct management [37](#)
- DSCLI management [52](#)
- SVC management [87](#)

SVC management

- viewing [91](#)

configuring

consistency group [57](#)

Direct management [27](#)

Direct management

- cluster [32](#)
- overview [28](#)
- PPRC paths [23](#)
- PPRC tasks [28](#)
- prerequisites [27](#)
- resource groups [37](#)
- support [28](#)

DSCLI management [46](#)

DSCLI management

- replicated resource [47](#)
- resource groups [50](#)

PBR for SVC flash storage [94](#)

SVC management [81](#), [92](#)

SVC management

- cluster [82](#), [93](#)
- cluster for flash storage [95](#)

consistency group

changing

- site configuration [60](#)

configuring [57](#)

installing [57](#)

overview [53](#)

planning [54](#)

planning resource group [54](#)

sample configuration [55](#)

troubleshooting

- configuration [62](#)

verifying

- configuration [62](#)

Copy Services Server

defining

- DSCLI managing [47](#)

planning for Direct management [23](#)

D

defining

Direct management

ESS disk subsystems 32

PPRC pairs 32

PPRC tasks 33

replicated resource 32

DSCLI management

Copy Services Server 47

ESS disk subsystems 48

SVC management

pbr replicated resource configuration 96

replicated resource configuration 83

Direct management

changing replicated resource configuration 38, 37

changing site configuration 37

configuration overview 28

configuring 27

configuring cluster 32

configuring PPRC paths 23

configuring PPRC tasks 28

configuring prerequisites 27

configuring resource groups 37

configuring support 28

defining ESS disk subsystems 32

defining PPRC pair 32

defining PPRC task 33

defining replicated resource 32

for Metro Mirror 21

improving volume group performance 36

installation troubleshooting 27

installing 26

installing prerequisites 26

installing software requirements 26

planning 21

planning connections 23

planning Copy Services Server 23

planning integration 23

planning PPRC replicated resource 24

planning PPRC resource group 25

planning prerequisites 22

planning tasks for ESS 24

sample configuration 22

starting the cluster 37

synchronizing PPRC configuration 36

verifying PPRC configuration 36

volume groups 24

disk subsystem 9

DSCLI management

adding replicated resource 49

changing replicated resource configuration 52

changing site configuration 52

changing site configuration

consistency group 60

configuring 46, 81, 92

configuring replicated resource 47

configuring resource groups 50

consistency group

configuring 57

defining Copy Services Server 47

defining ESS disk subsystems 48

for Metro Mirror 38

installation directories 46

installing 45

installing consistency group 57

installing software requirements 45

overview 11, 38

overview

consistency group 53

planning 39

planning

consistency groups 54

planning prerequisites 39

planning resource group 44

planning resource group

consistency group 54

sample configuration 40

sample configuration

consistency group 55

starting the cluster 52

synchronizing PPRC configuration 52

troubleshooting

consistency group configuration 62

upgrading 46

verifying configuration

consistency group 62

verifying PPRC configuration 50

volume groups 42

E

ESS

planning for Direct management 24

ESS disk subsystems

defining

Direct managing 32

DSCLI managing 48

F

fallback 13

fallover 13

for Metro Mirror

cluster 11

Direct management 21

DSCLI management 38

DSCLI management

overview 38

fallback 13

fallover 13

installation media 16

installing 15

installing from CD-ROM 17

installing from hard disk 17

installing from installation server 17

- installing prerequisites 15
- maintaining 64
- management types 13
- planning 13
- planning overview 14
- recovering from failed installation 20
- SVC management**
 - overview 74
- troubleshooting 65
- upgrading 19
- upgrading software 19
- verifying upgrade 20

H

- hard disk 17

I

- improving
 - Direct management**
 - volume group performance 36
- installation directories
 - DSCLI management 46
- installation server 17
- installing
 - CD-ROM 17
 - consistency group 57
 - Direct management 26
 - Direct management**
 - prerequisites 26
 - software requirements 26
 - troubleshooting 27
 - DSCLI management 45
 - DSCLI management**
 - directories 46
 - software requirements 45
 - for Metro Mirror 15
 - hard disk 17
 - installation media 16
 - installation server 17
 - prerequisites for for Metro Mirror 15
 - recovering from failed installation 20
 - SVC management 81
 - troubleshoot 21

M

- maintaining
 - for Metro Mirror 64
- Metro Mirror PPRC**
 - in-band communication 147
- mirroring 9
- modifying
 - previous snapshots 20

O

- overview 8
- overview

- consistency group 53
- Direct management**
 - configuration 28
- DSCLI management 38
- planning 14
- pprc 10
- SVC management 74
- SVC policy based replication management 93

P

- planning
 - consistency group**
 - resource group 54
 - consistency groups 54
 - Direct management 21
 - Direct management**
 - connections 23
 - Copy Services Server 23
 - integration 23
 - PPRC replicated resource 24
 - PPRC resource group 25
 - prerequisites 22
 - sample configuration 22
 - tasks for ESS 24
 - volume groups 24
 - DSCLI management 39
 - DSCLI management**
 - prerequisites 39
 - resource group 44
 - sample configuration 40
 - volume groups 42
 - for Metro Mirror 13
 - overview for for Metro Mirror 14
 - resource group 14
 - site 14
 - SVC management 75
 - SVC management**
 - prerequisites 76
 - sample configuration 76
 - support 79
 - volume groups 79
- PowerHA SystemMirror**
 - 131
- PPRC pair**
 - defining
 - Direct management 32
- PPRC paths**
 - configuring for Direct management 23
- PPRC replicated resource 12
- PPRC replicated resource**
 - planning for Direct management 24
- PPRC resource group**
 - planning for Direct management 25
- PPRC task**
 - defining
 - Direct management 33
- PPRC tasks**
 - configuring for Direct management 28

prerequisites 8

R

recovering

from failed installation 20

relationship

SVC management

changing 88

removing 89

removing

SVC management 88

SVC management

relationships 89

resources 90, 100

replicated resource

adding

DSCLI management 49

changing configuration

Direct management 38, 37

DSCLI management 52

configuring

DSCLI management 47

defining

Direct management 32

defining configuration

SVC management 83, 96

Direct management 12

DSCLI-managed 12

SVC-managed 12

resource

SVC management

changing 89, 99

removing 90, 100

resource group 12

resource group

configuring

Direct management 37

DSCLI management 50

planning 14

planning for consistency group 54

planning for DSCLI management 44

S

sample configuration

consistency group 55

Direct management 22

DSCLI management 40

SVC management 76

site

changing configuration

consistency group 60

Direct management 37

DSCLI management 52

planning 14

snapshots

modifying previous 20

software requirements

Direct management 26

DSCLI management 45

starting

Direct management

cluster 37

DSCLI management

cluster 52

SVC management

cluster 87

SVC management

changing

relationships 88

resources 89, 99

cl_verify_svcpprc_config command 92

cllsrelationship command 92

cllssvc command 91

cllssvcpprc command 91

configuring 94

configuring cluster 82, 93, 95

defining pbr replicated resource configuration 96

defining replicated resource configuration 83

installing 81

overview 11, 74

planning 75

planning

support 79

volume groups 79

planning prerequisites 76

removing 88

removing

relationships 89

resources 90, 100

sample configuration 76

starting the cluster 87

synchronizing PBR configuration 97

synchronizing PPRC configuration 86

troubleshooting 90, 100

verifying PBR configuration 97

verifying PPRC configuration 86

viewing

clusters 91

SVC pbr management

overview 93

synchronizing

Direct management

PPRC configuration 36

DSCLI management

PPRC configuration 52

SVC management

PBR configuration 97

PPRC configuration 86

T

troubleshooting

consistency group 62

Direct management installation 27

for Metro Mirror 65

installation 21

SVC management 90, 100

U

upgrading

DSCLI management [46](#)

for Metro Mirror [19](#)

software for for Metro Mirror [19](#)

verifying [20](#)

V

verifying

consistency group

configuration [62](#)

Direct management

PPRC configuration [36](#)

DSCLI management

PPRC configuration [50](#)

SVC management

PBR configuration [97](#)

PPRC configuration [86](#)

upgrade [20](#)

viewing

SVC management

clusters [91](#)

volume group

Direct management [24](#)

DSCLI management [42](#)

improving

Direct management [36](#)

SVC management [79](#)

© Copyright International Business Machines Corporation 2017, 2025

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

