Redbooks
ibm.com/redbooks

IBM PowerHA SystemMirror for i: Using Geographic Mirroring (Volume 4 of 4)

David Granum
Edward Grygierczyk
Sabine Jordan
Peter Mayhew
David Painter

Power Systems

IBM

Redbooks

International Technical Support Organization

IBM PowerHA SystemMirror for i: Using Geographic Mirrioring

June 2016

SG24-8401-00

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

First Edition (June 2016)

This edition applies to Version 7, Release 2 of IBM PowerHA SystemMirror for i (product number 5770-HAS).

Contents

Notices

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

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

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

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

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

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided “AS IS”, without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

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

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

DS8000®
FlashCopy®
IBM®
Power Systems™
PowerHA®
Redbooks®
Redbooks (logo) ®
Storwize®
SystemMirror®
Tivoli®

The following terms are trademarks of other companies:

Inc., and Inc. device are trademarks or registered trademarks of Kenexa, an IBM Company.

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

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

Preface

IBM® PowerHA® SystemMirror® for i is the IBM high-availability (HA), disk-based clustering solution for the IBM i operating system. When PowerHA for i is combined with IBM i clustering technology, PowerHA for i delivers a complete HA and disaster-recovery (DR) solution for business applications that are running in an IBM i environment. Use PowerHA for i to support HA capabilities with either native disk storage, IBM DS8000® storage servers, or IBM Storwize® storage servers.

This IBM Redbooks® publication helps you to install, tailor, and configure IBM PowerHA SystemMirror for i to use with geographic mirroring and native storage. This publication provides you with planning information to prepare to use the various PowerHA offerings with geographic mirroring with IBM i native storage. It also provides implementation and management information. It provides guidance about troubleshooting these solutions and identifies the documentation that you need to capture before you call IBM Support.

This book is part of a four-book set that gives you a complete understanding of PowerHA for i with native disk storage, IBM DS8000 storage servers, or IBM Storwize storage servers. The following IBM Redbooks publications are part of this PowerHA for i volume set:

- *IBM PowerHA SystemMirror for i: Preparation*, SG24-8400
- *IBM PowerHA SystemMirror for i: Using DS8000*, SG24-8403
- *IBM PowerHA SystemMirror for i: Using IBM Storwize*, SG24-8402

Important: The information that is presented in this volume set is for technical consultants, technical support staff, IT architects, and IT specialists who are responsible for providing HA and support for IBM i solutions. If you are new to HA, you need to first review the information that is presented in the first book of this volume set, *IBM PowerHA SystemMirror for i: Preparation (Volume 1 of 4)*, SG24-8400, to obtain a general understanding of clustering technology, independent auxiliary storage pools (IASPs), and the PowerHA architecture.

Authors

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

David Granum is an Advisory Software Engineer at IBM Rochester. He has been with the IBM i Global Support Center (iGSC) throughout his 18 years at IBM. His areas of expertise include PowerHA, performance, communications, and TCP/IP applications on the IBM i. He holds a Computer Science degree from South Dakota State University and a Masters degree in Telecommunications from Saint Mary's University of Minnesota.

Edward Grygierczyk is part of the IBM Lab Services team and based in Poland. He joined IBM in 2008, working for IBM STG Lab Services as a Consultant in IBM Power Systems™. Edward has worked in the IBM i area since 1991. He built a highly successful IBM Business Partner company in Poland before he joined IBM. Edward has a background in systems, availability, integration, and implementation. Edward was also an author of IBM i courses. Edward can be reached at edward.grygierczyk@pl.ibm.com.

Sabine Jordan is a Consulting IT Specialist, working in IBM Germany. She has worked as a Technical Specialist in the IBM i area for more than 20 years, specializing in HA. She has worked on IBM PowerHA SystemMirror for i implementations for both SAP and non-SAP environments that use geographic mirroring, IBM DS8000, and IBM SAN Volume Controller (SVC) remote copy services. Among these implementations, she created concepts for the design and implemented the entire project (cluster setup and application changes), in addition to performing customer education and testing. In addition, Sabine presents and delivers workshops (internally and externally) on IBM PowerHA SystemMirror for i and HA and DR.

Peter Mayhew is an IBM i Technical Services Consultant, working for IT Solutions Group, an IBM Business Partner. He has worked with midrange systems for almost 40 years, beginning as an IBM hardware service representative on the 3x systems, AS/400 CISC, and early RISC iSeries models. Leaving IBM after 23 years, Peter has maintained his skills on the IBM i, specializing in hardware and operating system installations, upgrades and migrations, and HA planning and implementation, supporting clients throughout the United States. Peter can be reached at Pete@itsolutions-group.com.

David Painter is part of the IBM Rochester Lab Services team and based in the United Kingdom. He joined IBM in 2006, working for IBM STG Lab Services as a Senior IT Specialist in IBM Power Systems, HA, and external storage. Before David joined IBM, he set up and built a highly successful IBM Business Partner company in the UK. David has a background in systems, availability, integration, and implementation. Today, David implements HA solutions across the world in clients that range from small and medium business (SMB) accounts to large enterprise accounts. David regularly teaches courses in IBM i HA and external storage for IBM i administrators. He also conducts many availability workshops and holds numerous certifications. David can be reached at davpaint@uk.ibm.com.

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

Philip Albu
Jenny Dervin
Steven Finnes
Scott Helt
Ben Rabe
Curt Schemmel
James Tilbury
IBM Rochester, IBM i Development and Support teams

Christian Aasland
Laural Bauer
Selwyn Dickey
IBM Systems Lab Services

Tracy Eidenschink
Cindy Mestad
Michael K Meyers
IBM i Performance and Scalability Services (PaSS) Center

Debbie Landon
Ann Lund
International Technical Support Organization

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

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

  ibm.com/redbooks

- Send your comments in an email to:

  redbooks@us.ibm.com

- Mail your comments to:

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

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

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

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

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

- Stay current on recent Redbooks publications with RSS Feeds:

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

1

Introduction to geographic mirroring

This chapter provides an introduction to IBM PowerHA SystemMirror for i with geographic mirroring.

The following topics are described in this chapter:

1.1 Value proposition

Today, businesses experience increasing demands for application availability, which requires that both small and large clients look for a solution that can help eliminate planned and unplanned downtime for their IT services.

Any unplanned outage can have severe implications, especially if the duration of the outage or recovery time exceeds business expectations. These implications vary and they can include loss of data, revenue, worker productivity, company reputation, and client loyalty. Companies that did not plan for the risk of an unplanned outage effectively, never completed their installation of an HA solution, or did not test a tape recovery plan are exposed to these negative business effects.

IBM PowerHA SystemMirror for i (PowerHA) offers a complete end-to-end integrated clustering solution for HA and disaster recovery (DR). PowerHA provides a data and application resiliency solution that is an integrated extension of the IBM i operating system and storage management architecture. Among other features is the design objective of providing application HA during both planned and unplanned outages.

PowerHA geographic mirroring offers a straightforward, cost-effective, HA solution for the small to mid-sized client. Typically, PowerHA geographic mirroring is used with internal disk storage, and it provides an alternative to solutions that require the additional configuration and management that are associated with an external storage device. In addition, if the systems are in different locations, it can provide protection in a site outage (disaster recovery).

For more information about IBM PowerHA SystemMirror for i, see the following website:

https://ibm.biz/Bd4JQx

1.2 Concept of geographic mirroring

Geographic mirroring refers to the IBM i host-based replication solution that is provided as a function of IBM PowerHA SystemMirror for i.

The basis of geographic mirroring is extending IBM i disk mirroring technology to a multiple systems environment. Geographic mirroring is managed by IBM i storage management so that replication is performed on a disk page segment (4k or 4224 bytes) basis. When a page of data is written, storage management automatically manages the replication of this page to the remote system.

Geographic mirroring requires a two-node clustered environment and uses data port services.

Data port services are provided by the System Licensed Internal Code (SLIC) to support the transfer of large volumes of data between a source node and a target node. The processes that keep the independent auxiliary storage pool (IASP) in synchronization run on the nodes that own the IASPs. This transport mechanism communicates over TCP/IP.

Figure 1-1 shows the concept of geographic mirroring.

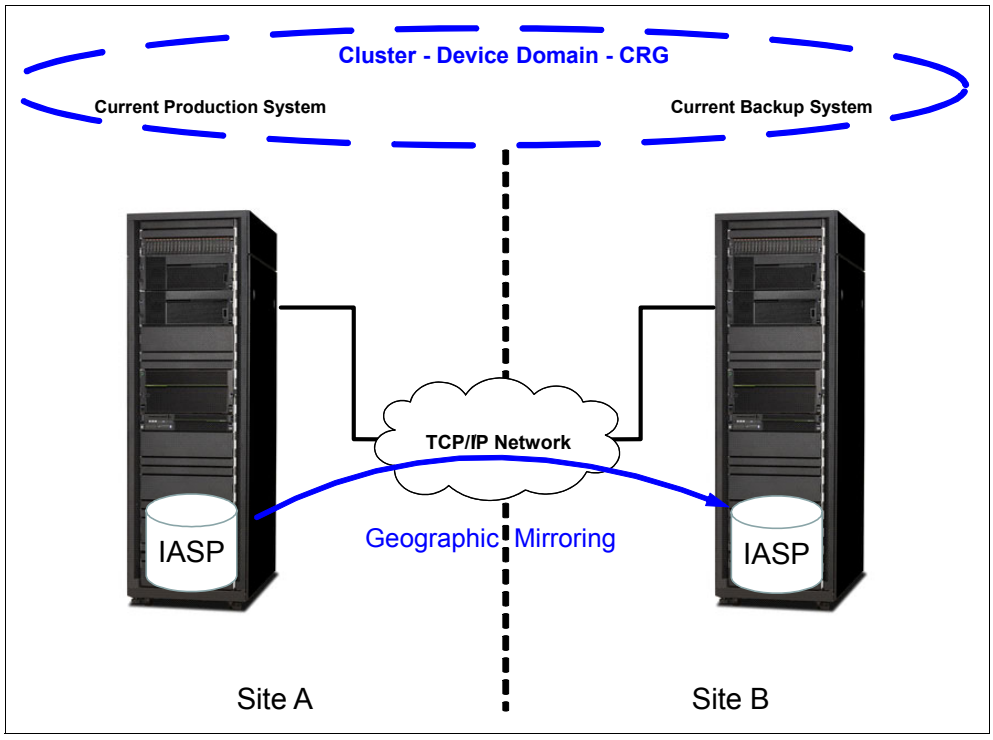

Figure 1-1 IBM PowerHA SystemMirror for i that uses geographic mirroring

For geographic mirroring, IASPs can be built on internal disk, external disk storage, or both because this technology is storage-agnostic.

The copy that is owned by the primary node is designated as the *production copy* and the copy that is owned by the backup node is the *mirror copy*. Users and applications can access the IASP on only the primary node while geographic mirroring is active. Changes that are made to the production copy (source system) are guaranteed by the geographic mirroring functionality to be made in the same order on the mirror copy (target system).

Geographic mirroring allows the production and mirrored copies to be on the same site for HA protection in a server failure. It is also possible to separate the two systems geographically for DR protection in a site-wide outage, if the communication link between the two sites has sufficient bandwidth.

Geographic mirroring is the only physical replication that can be used for IASPs that are built on the internal disk subsystems. It has a relatively low cost of implementation because it is a software-based solution and it does not have any particular storage requirements to meet for replication.

A possible disadvantage of this solution is that it uses the host resources, such as CPU and memory, which might affect the performance of other processes that are running on the system. Therefore, ensure that you plan for the additional overhead that is required for the replication.

1.3 Prerequisites and editions

Details of the prerequisites for IBM PowerHA SystemMirror for i and the available editions are described in the IBM Redbooks publication *IBM PowerHA SystemMirror for i: Preparation (Volume 1 of 4)*, SG24-8400.

For convenience, a brief summary is listed here. The following products must be installed on both nodes in the cluster:

- HA Switchable Resources, 5770-SS1 option 41
- IBM PowerHA SystemMirror for i, 5770-HAS *BASE
- PowerHA for i Standard Edition, 5770-HAS option 2
- If asynchronous mode operation is required, include PowerHA for i Enterprise Edition, 5770-HAS option 1

1.4 Modes of operation

Geographic mirroring uses two operational modes in conjunction to copy data from the source IASP to the target IASP:

- *Transmission delivery mode* refers to how the writes are delivered from the production copy system to the mirror copy system from a communications perspective.
- *Mirroring mode* refers to the method that is used to save the data on the mirror copy system.

Both of these modes can be configured as *synchronous* or *asynchronous* (although not all combinations are allowed).

If you use synchronous delivery mode, communication speed and throughput affect the application response time on the production system because the production system waits until a write operation at least reaches main memory on the backup system and the backup system sends a confirmation back to the production system before the local write to the IASP on the production system is considered finished.

PowerHA for i Enterprise Edition offers the asynchronous delivery mode of the changes to the remote IASP. In this case, the network delivery time does not affect application performance as significantly. However, the changes are not guaranteed to be delivered to the backup node in an unexpected failure of the primary node.

Synchronous or asynchronous mirroring mode determines whether the data needs to be written to the disk cache on the backup node to have a completed status on the primary node.

For more information about the modes of operation and the effect of various combinations, see 2.2, “Synchronous and asynchronous mode considerations” on page 6.

2

Planning

This chapter describes critical planning steps to consider when you implement IBM i host-based replication services.

Before you implement a geographic mirroring solution, consider the following items:

2.1 Storage options

Geographic mirroring is a host-based replication solution. Because the IBM i manages data transmission between a current production and current backup system, the type of storage that is used can be internal or external.

When the “geographic mirroring” term is used, many people immediately assume that the underlying storage that is used is internal storage. However, this assumption is not necessarily the case. Although geographic mirroring is the only IBM PowerHA SystemMirror for i replication option that is available when you use internal storage, geographic mirroring can also be used for an IBM i configuration that uses external storage. It can even be used with the external storage IBM FlashCopy® function.

2.2 Synchronous and asynchronous mode considerations

Geographic mirroring can be configured to use synchronous or asynchronous transmission delivery mode.

Depending on the transmission delivery mode, the mirroring mode can also be configured to be synchronous or asynchronous. *Transmission delivery mode* refers to how the writes are delivered from the production copy system to the mirror copy system from a communications perspective. *Mirroring mode* refers to the method of how the data is saved on the mirror copy system.

The rest of this section describes the various combinations of the synchronous and asynchronous transmission modes.

2.2.1 Synchronous transmission mode

Synchronous transmission can use either synchronous or asynchronous mirroring mode. In synchronous transmission mode, the job that is running on the production copy system must wait for writes to be completed on both the production copy system and the mirror copy system.

Synchronous transmission mode with synchronous mirroring mode

Synchronous transmission mode with synchronous mirroring mode specifies that both transmission delivery and mirroring mode are synchronous. Synchronous transmission mode with synchronous mirroring mode is sometimes referred to as *sync/sync mode*.

When geographic mirroring is active in synchronous mode (Figure 2-1 on page 7), the disk write operation is sent in parallel to the production copy’s disk subsystem (1) and the mirror copy system (2). These writes must complete to the disk (usually disk cache) on both the production copy system and the preferred target system. After the mirror copy system completes the write operation (3), the acknowledgment is then sent back to the production copy system (4).

In this configuration, the mirror copy independent auxiliary storage pool (IASP) is always eligible to become the production copy IASP because the order of writes is preserved on the mirror copy system. For this reason, this configuration is preferred, if possible. However, in many cases, the network infrastructure does not allow this configuration as a practical solution. For instance, on lower bandwidth or high latency communications links, it is not feasible to have to wait for the acknowledgment of the write being sent to the disk write cache, or even the memory, on the preferred target system.

Figure 2-1 shows geographic mirroring that is active in synchronous mode.

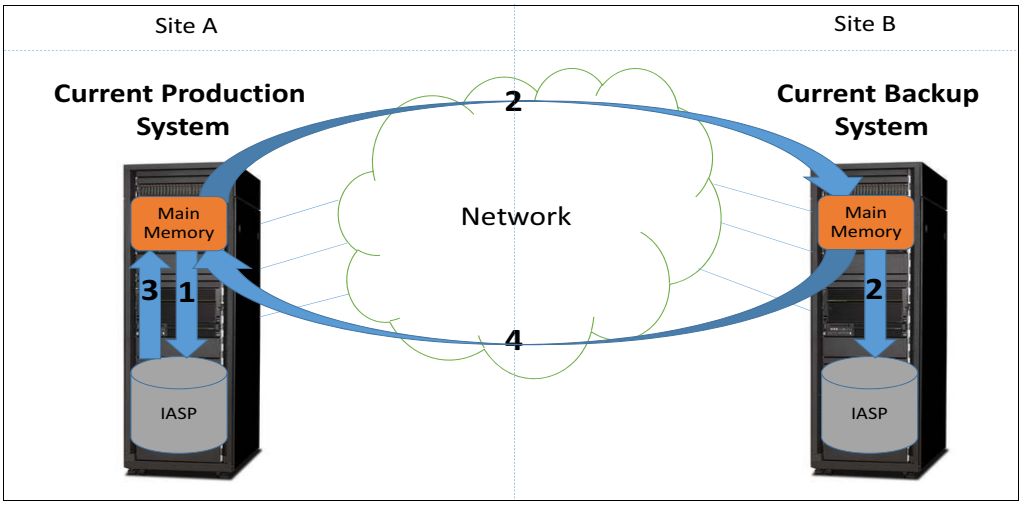

Figure 2-1 Synchronous transmission mode with synchronous mirroring mode

Synchronous transmission mode with asynchronous mirroring mode

Synchronous transmission mode with asynchronous mirroring mode is similar to sync/sync mode, with one difference. The communications method of waiting for write acknowledgment from the mirror copy system is still used in this case. However, when you use asynchronous mirroring mode, the acknowledgment is sent back from the mirror copy system when that data is in memory on the mirror copy system. (See operations 1 - 4 in Figure 2-2 on page 8.) This approach provides a faster acknowledgment because no waiting is necessary for the write to complete on the mirror copy system. The physical write operation (5) is performed later (asynchronously) to the disk on the mirror copy system. This approach is sometimes referred to as *sync/async mode*.

In this mode, the pending updates must be completed before the mirror copy can become the production copy. Performance might improve slightly on the production copy system during normal operation. However, switchover or failover times will be slightly longer because changes to the backup IASP will still reside in the main memory of the backup system and these changes must be written to disk before the IASP can be varied on.

Figure 2-2 shows synchronous transmission mode with asynchronous mirroring mode.

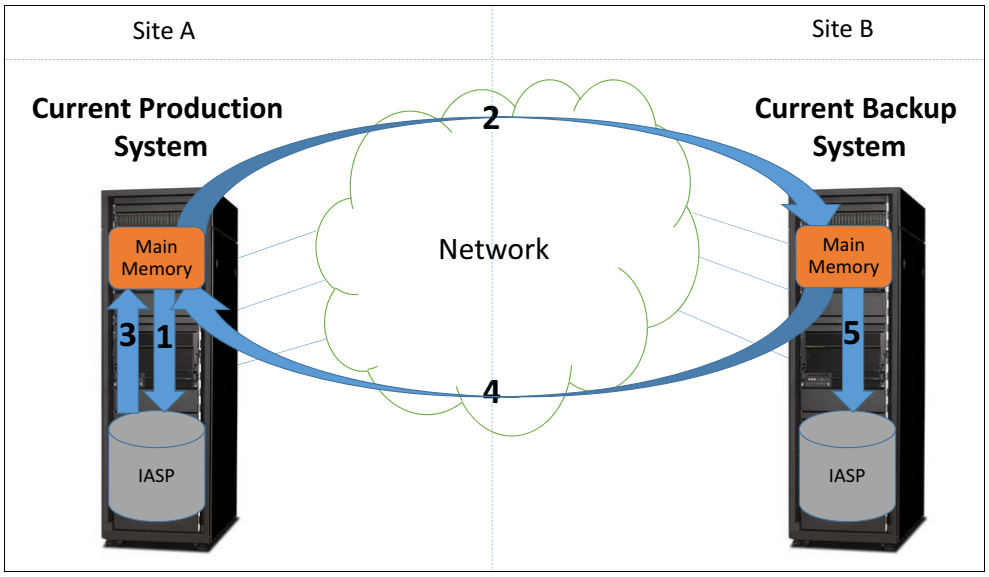

Figure 2-2 Synchronous transmission mode with asynchronous mirroring mode

2.2.2 Asynchronous transmission mode

Asynchronous transmission mode is available for environments in which synchronous mode is not a feasible solution. In asynchronous transmission mode, the job on the production copy system that performs the write to the IASP does not need to wait for the mirror copy write acknowledgment.

Asynchronous transmission mode with asynchronous delivery mode

Asynchronous transmission mode with asynchronous mirroring mode specifies that both transmission delivery and mirroring mode are asynchronous. This approach is sometimes referred to as *async/async mode*.

Asynchronous transmission delivery allows replication of the IASP beyond synchronous geographic mirroring limits. The write on disk operation does not wait until the operation is delivered to the mirror copy system.

Asynchronous transmission delivery requires asynchronous mirroring mode. It works by duplicating any changed IASP disk pages in the *BASE memory pool on the production copy system. Likely, page faulting will result in the *BASE pool as a result. These duplicated disk pages are sent asynchronously while the write order to the target (3) is preserved (Figure 2-3 on page 9). Therefore, at any time, the data on the target system (although not up-to-date) still represents a crash-consistent copy of the source system.

Figure 2-3 shows asynchronous transmission mode with asynchronous mirroring mode.

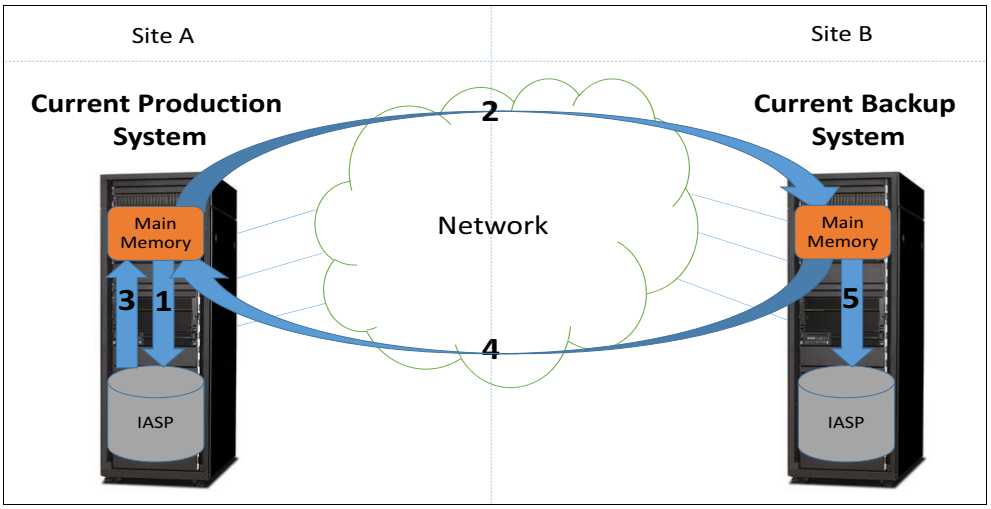

Figure 2-3 Asynchronous transmission mode with asynchronous mirroring mode

The asynchronous geographic mirroring option potentially affects performance to system resources, such as processor and memory. Communication lines with longer latency times might tie up additional memory resources for maintaining changed data. Therefore, the environment must be sized correctly in terms of both bandwidth/latency and system resources.

Communications and system performance considerations are described in 2.4, “Communication considerations for geographic mirroring” on page 12 and 2.5, “Determining bandwidth requirements” on page 14 for communications considerations. For system performance considerations, see 2.6, “System performance considerations” on page 21.

2.2.3 Geographic mirroring operations summary

The relationship between the operations in synchronous and asynchronous modes of operations are shown in Figure 2-4. The operations for each type of delivery are listed in Table 2-1.

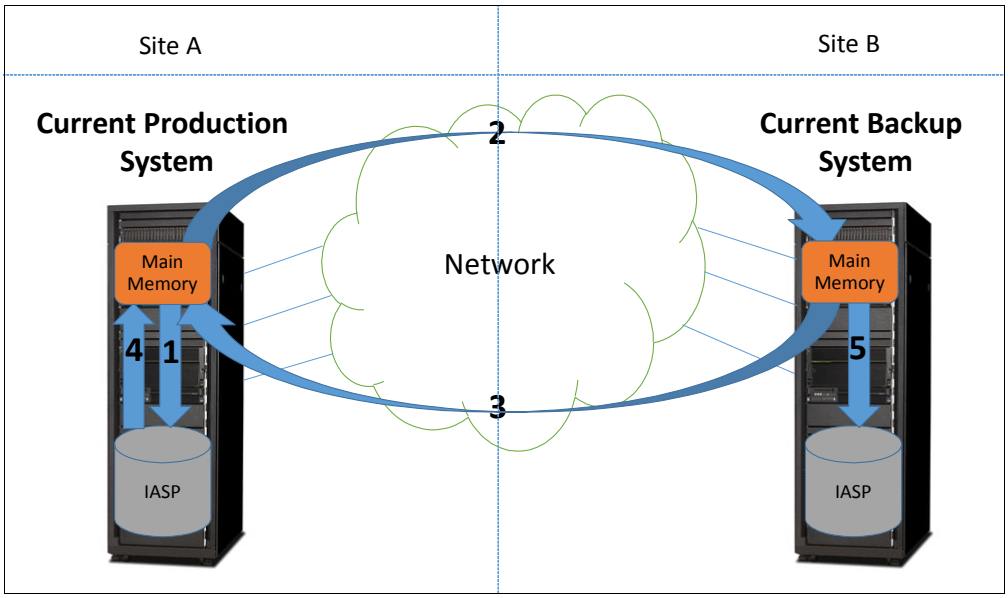

Figure 2-4 Geographic mirroring operations

Table 2-1 Geographic mirroring operations relationships

Transmission delivery	Mirroring mode	Local operations	Remote operations
*ASYNC	*ASYNC	1 and 4	2
*SYNC	*ASYNC	1 and 4	2, 3
*SYNC	*SYNC	1 and 4	2, 5, and 3

2.3 Other geographic mirroring considerations

This section describes the following additional topics that must be considered when you plan a geographic mirroring environment:

- 2.3.1, “Tracking space” on page 10
- 2.3.2, “Suspend timeout” on page 11

2.3.1 Tracking space

Tracking space enables geographic mirroring to track changed pages while it is in suspended status. With tracked changes, geographic mirroring can avoid full resynchronization after it resumes in many cases, which minimizes the exposure of time frames where no valid mirror copy is available. Tracking space is configured when geographic mirroring is configured or later by using the Change Auxiliary Storage Pool Session (**CHGASPSSN**) command.

Tracking space is allocated inside of the IASPs. The more tracking space that is specified, the more changes the system can track. The amount of space for tracking can be defined by the user up to 1% of the total IASP capacity.

When you specify tracking space size, the value refers to the percentage of total usable tracking space size. Therefore, if 100% is specified, 1% of the total IASP capacity is used as tracking space size.

For example, if the IASP is 100 GB, a maximum of 1 GB of storage space is used as the tracking space. If the tracking space parameter is set to 50%, 500 MB of storage space is used as tracking space.

Note: This tracking space does not contain any changed data. It merely holds information about what pages in the IASP were changed.

2.3.2 Suspend timeout

The *suspend timeout* in the ASP session specifies how long the application can wait when geographic mirroring cannot be performed. When an error, such as a failure of the communication link, prevents geographic mirroring from occurring, the production copy system waits and retries during the specified suspend timeout before it suspends geographic mirroring, which allows the application to continue.

The timeout value can be tuned by using the Change Auxiliary Storage Pool Session (**CHGASPSSN**) command as shown in Figure 2-5. The default value of the Suspend timeout (SSPTIMO) parameter is 120 seconds.

```
                      Change ASP Session (CHGASPSSN)

Type choices, press Enter.

Session  . . . . . . . . . . . . > GEOMIRROR     Name
Option . . . . . . . . . . . . . > *CHGATTR      *CHGATTR, *SUSPEND...
ASP copy:
  Preferred source . . . . . . . > S1CPYD        Name, *SAME
  Preferred target . . . . . . . > S2CPYD        Name, *SAME
              + for more values
Suspend timeout  . . . . . . . . > 120           60-3600, *SAME
Transmission delivery  . . . . . > *SYNC         *SAME, *SYNC, *ASYNC
Mirroring mode . . . . . . . . . > *SYNC         *SAME, *SYNC, *ASYNC
Synchronization priority . . . . > *MEDIUM       *SAME, *LOW, *MEDIUM, *HIGH
Tracking space . . . . . . . . . > 100           0-100, *SAME
FlashCopy type . . . . . . . . .   *SAME         *SAME, *COPY
Persistent relationship  . . . .   *SAME         *SAME, *YES, *NO




                                                                        Bottom
F3=Exit   F4=Prompt   F5=Refresh   F10=Additional parameters   F12=Cancel
F13=How to use this display        F24=More keys
```

Figure 2-5 Change ASP Session command

2.4 Communication considerations for geographic mirroring

When you implement an IBM i high-availability (HA) solution that uses geographic mirroring, it is critical that it includes a plan for adequate communication bandwidth so that it does not become a performance bottleneck. Both runtime replication and synchronization bandwidth requirements must be considered.

2.4.1 Network topology and configuration

Geographic mirroring can be used in environments over any distance. However, only business needs determine the latency that is acceptable for a certain application. Many factors affect communications latency. As a result, these factors might affect geographic mirroring performance.

Consider the following information about communications in a geographic mirroring environment:

- To provide consistent response time, geographic mirroring must have its own redundant communications lines. Without dedicated communication lines, contention might occur with other services or applications on the IBM i that use the same communication lines. Geographic mirroring supports up to four communications lines for dataport services.

  If the environment is configured with multiple lines, geographic mirroring distributes the load over multiple lines.

  Also, multiple communication lines among the nodes provide redundancy to the configuration.
- Geographic mirroring replication must be configured on separate interfaces and lines from the cluster heartbeat interface and line (the Internet Protocol (IP) address that is associated with each node in the cluster). If the same line is used, during periods of heavy geographic mirroring traffic, the cluster heartbeat messaging process can fail between nodes, causing a node to go into a Partition status.
- If the system is configured so that multiple applications or services require the use of the same communication line, part of this contention can be alleviated by implementing quality of service (QoS) through the TCP/IP functions of IBM i. The IBM i QoS solution enables the policies to request network priority and bandwidth for TCP/IP applications throughout the network. QoS can also be implemented with the network infrastructure devices.
- When you use multiple lines, the line speed setting needs to be the same for each line. If throughput differs, performance is gated by the slowest connection. For example, a configuration of lines that are configured for 100 megabits per second (Mbps) and 1 gigabit per second (Gbps) results in geographic mirroring throughput that is gated by the 100 Mbps link.
- Typically, the throughput for this traffic is determined by the slowest link in the network path between the sites, unless those sites are in the same physical location and local network. The decision between synchronous and asynchronous replication is largely determined by the throughput and latency between the two sites.
- The communications bandwidth must be sized for both the resynchronization process and normal runtime operations that occur in parallel. If geographic mirroring is suspended and a resynchronization must occur, the tracked changes function is performed. During this time, the high priority synchronization tasks are used on the production copy system, which might degrade application performance if the network bandwidth between the sites is saturated.

- If both the current production and current backup systems are connected to the same physical network and are on the same logical network, Schowler routes can be used to control the interfaces for this local network traffic.

  Schowler routing is a method for replacing direct routes in the IBM i IP routing table with network routes. This process allows the administrator to create host routes with preferred binding interfaces. These host routes allow control over which local interface is used when a system communicates with another system on the same logical network. These Schowler routes can be used for both heartbeat traffic and data port traffic. The following link provides more detailed information about this topic:

  https://ibm.biz/Bd4AcJ

- From an HA point of view, geographic mirroring interfaces that are associated with different Ethernet adapters are considered a preferred practice. The use of redundant switches and routers further improves the overall HA value of the environment. See Figure 2-6.

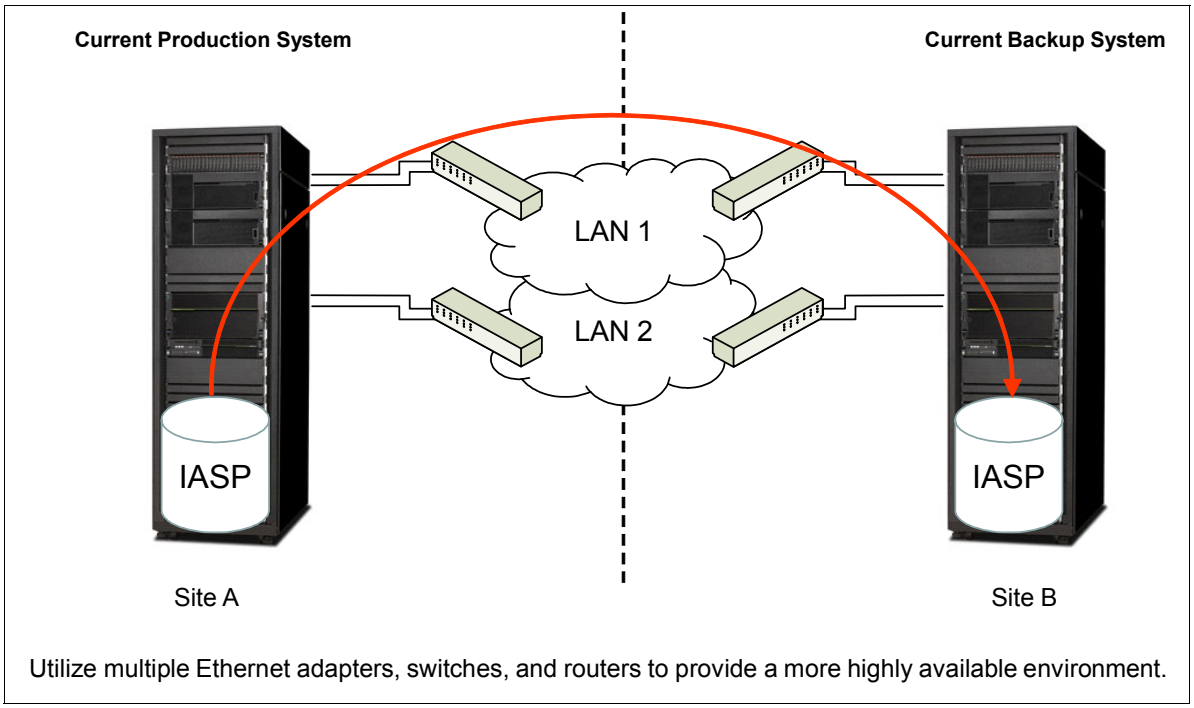

Figure 2-6 Preferred network configuration for geographic mirroring

2.4.2 Synchronization priority

One of the attributes that must be considered for geographic mirroring is synchronization priority.

The synchronization priority setting refers to a full or partial synchronization that is performed initially after you set up geographic mirroring, and after the geographic mirror session is suspended or detached. *Synchronization priority* determines the importance of how quickly geographic mirroring attempts to synchronize the current production copy of the IASP to the current backup copy of the IASP.

It is always best to synchronize the data as quickly as possible. However, potential drawbacks exist if you simply set the synchronization priority to *HIGH in all configurations.

The following list describes each of the synchronization priority settings:

- High

  High synchronization priority attempts to perform the synchronization as quickly as possible, without concern for system performance. No artificial delays are introduced into the transmission of the data.

- Medium

  Medium synchronization priority attempts to perform the synchronization with a balanced approach between transmission performance and system performance. An artificial delay of 1 second occurs for every 100 segments of data that is sent.

- Low

  Low synchronization priority attempts to perform synchronization and ensure that the affect on system resources and performance is minimal. An artificial delay of 1 second occurs for every 10 segments of data sent.

Important: While a synchronization is performed, other normal runtime write operations occur that can cause an additional load on the communication lines and system resources.

2.5 Determining bandwidth requirements

The bandwidth requirements for replicating both runtime disk operations and full synchronization operations must satisfy your business requirements.

2.5.1 Calculating full synchronization time bandwidth

When you plan for bandwidth requirements, you must determine the bandwidth that is required for a full synchronization.

To determine an approximation for the time that is needed for initial synchronization, take the total space that is used in the IASP times 8 to convert the number of bytes to bits (because throughput is normally measured in bits per second) divided by the effective communications capability of the chosen communications links.

In addition to the data in the IASP, the communications overhead must also be considered. Table 2-2 on page 15 includes a 15% overhead rate. The throughput of these links is largely the determining factor while the full synchronization completes.

Use the following formula to determine the time that is needed for initial synchronization, where *U* = IASP utilization in bytes and *C* = line speed in bits per second:

$(U \times 8 \times 1.15) / C$

Note: System performance factors can also play a significant role in synchronization time. These factors are described 2.6, “System performance considerations” on page 21.

Table 2-2 shows several examples of full synchronization time.

Table 2-2 Full synchronization time examples

IASP space used (GB)	Bandwidth (megabits per second (Mbps))	Time to full synchronization (hours)
2500	100	65.5
2500	1000	6.5
1000	1000	2.6
1000	100	26.1
1000	50	52.3
1000	10	261.5
500	25	6.5
500	10	130.7
500	5	261.4
250	5	130.7

During the time that is calculated in the last column of Table 2-2 for each environment, the environment is not switchable. Therefore, the business is vulnerable to an outage that might occur during this time frame. The importance of sufficient bandwidth to minimize this vulnerability is apparent from this table. Runtime disk writes also still occur during these synchronization times. Table 2-2 does not factor in the additional runtime writes. Calculating for that bandwidth is described in the following section, “Calculating runtime bandwidth”.

2.5.2 Calculating runtime bandwidth

Writes, which are geographically mirrored to the mirror copy system and that occur during normal runtime operations, to the IASP must also be considered when you determine bandwidth requirements. Runtime bandwidth can be calculated by sampling the collection services data on the production copy system.

Using Performance Data Investigator

The following steps assume that the Performance Data Investigator (PDI) Collection Services data was running for at least the previous five days. Also, the RETPERIOD value in the Configure Performance Collection (**CFGPFRCOL**) command is set to a minimum of 120 hours. To view historical performance data, Performance Data Investigator requires that all of the collections must be in a single library and no more than five collections can be graphed at the same time.

Use the following steps configure the environment for the data to be viewed in Performance Data Investigator:

1. On the IBM i command line, enter the Configure Performance Collection (**CFGPFRCOL**) command and press F4 for prompts. Document the Collection Library value, which is typically `QPFRDATA` or `QMPGDATA`. This example assumes that the collection services library is `QPFRDATA` for the remainder of these steps.
2. Create a library to copy the management collection objects to for this exercise. This library can be named anything. For this example, the library is named `BWDATA`:

   ```
   CRTLIB BWDATA
   ```

3. Obtain a list of all of the management collection objects in the performance data library by running the following Work with Objects Using PDM (**WRKOBJPDM**) command:

   ```
   WRKOBJPDM LIB(QPFRDATA) OBJ(*ALL) OBJTYPE(*MGTCOL)
   ```

   Note: The management collection object naming convention begins with the letter “Q”, which is followed by the Julian date and the time stamp. For example, Q247000032 is the management collection object from the 247th day of the year with the collection that started at 32 seconds after midnight on that day.

4. From the Work with Objects using PDM window, specify option `3` (Copy) for up to five collections. Most collection services configurations represent five days of system performance data. Press Enter.
5. From the Copy Objects window, specify the To library as the library that you created in step 2 (`BWDATA` in this example) and press Enter or F19 to submit the copying of the objects to batch. After the process completes, the collections are placed in the new library.
6. For the Performance Data Investigator tool to graph the collections in a single graph, the performance data files must be created. For *each* management collection object, run the following command, where `BWDATA` is the name of the library that you created and Qxxxxxxxxx is the name of the management collection object:

   ```
   CRTPFRDTA FROMMGTCOL(BWDATA/Qxxxxxxxxx)
   ```

7. Open a web browser on the workstation and specify the address or host name of the system for port 2001:

   ```
   http://<host_name_or_ip_address>:2001
   ```

8. When you are prompted, log in with a user that has *ALLOBJ special authority (Figure 2-7).

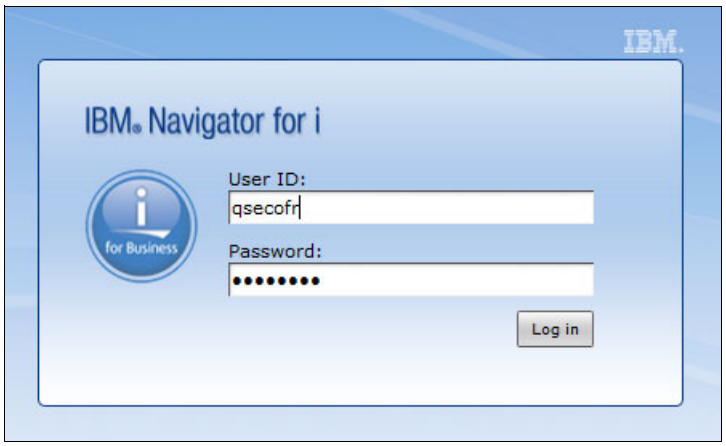

Figure 2-7 IBM Navigator for i sign-on panel

9. After you sign on, the Welcome window opens as shown in Figure 2-8.

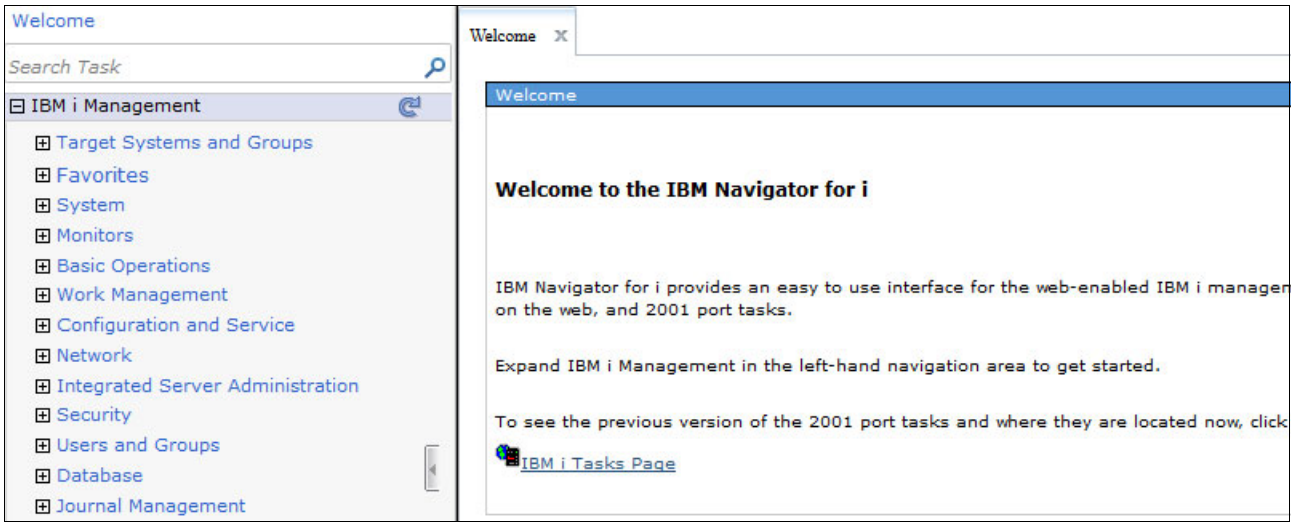

Figure 2-8 IBM Navigator for i Welcome window

10. To view the data in the new library in Performance Data Investigator, the collection table must be rebuilt. In the left pane, select **Performance** → **Manage Collections** as shown in Figure 2-9.

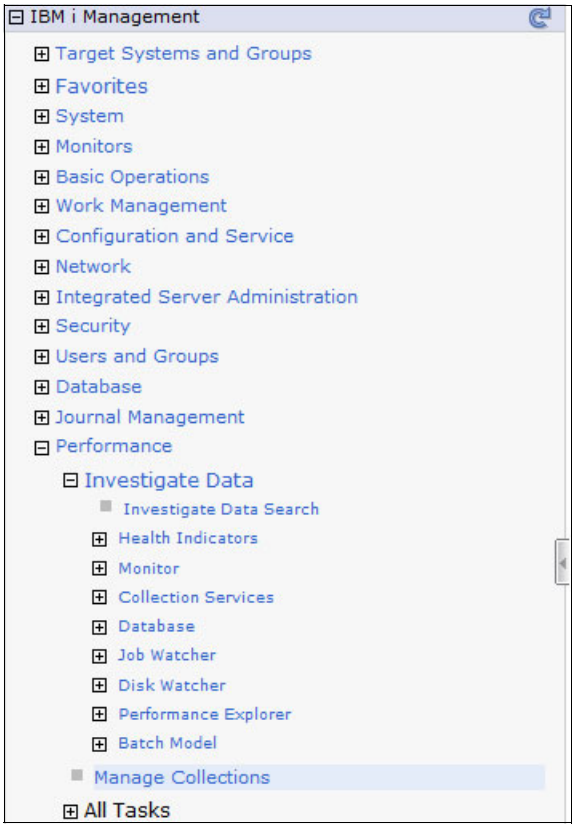

Figure 2-9 Performance Data Investigator, Manage Collections option

11. In the right pane, the list of collections is shown. Click **Actions** → **Maintain Collections** → **Rebuild Collection Table** as shown in Figure 2-10.

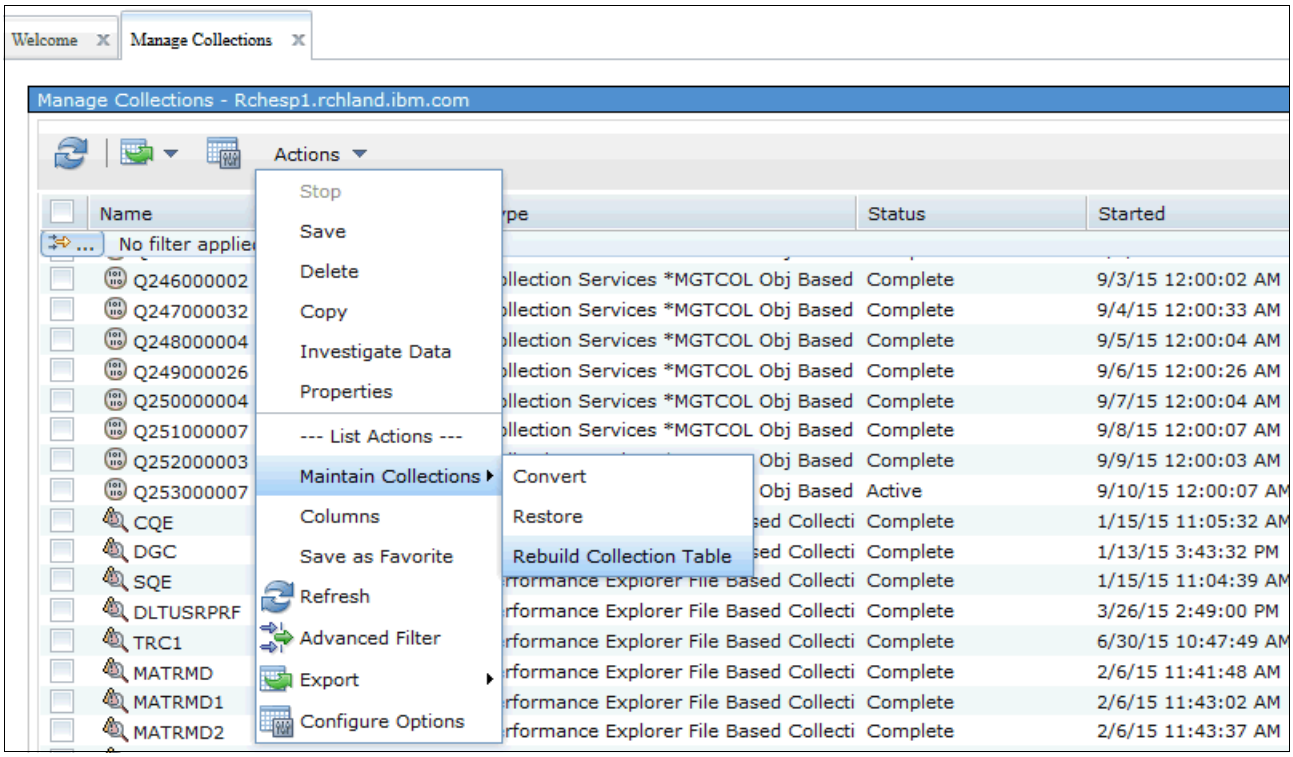

Figure 2-10 Performance Data Investigator, Rebuild Collection Table

12. The rebuild of the collection table takes a few moments to complete. After it completes, the Manage Collections list shows the collections in the library. In the left pane, select **Performance** → **Investigate Data** as shown in Figure 2-11.

Figure 2-11 Performance Data Investigator, Investigate Data

13. In the left pane, expand **Collection Services** → **Disk** → **Disk Throughput Overview for Disk Pools** as shown in Figure 2-12.

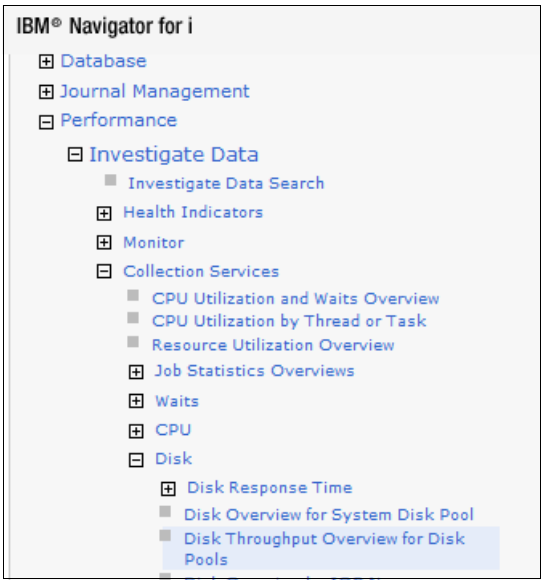

Figure 2-12 Performance Data Investigator, Collection Services - Disk

14. In the right pane, change the Collection Library to your library (`BWDATA` in this example) and the Collection Name to **All**, and click **Display** (Figure 2-13).

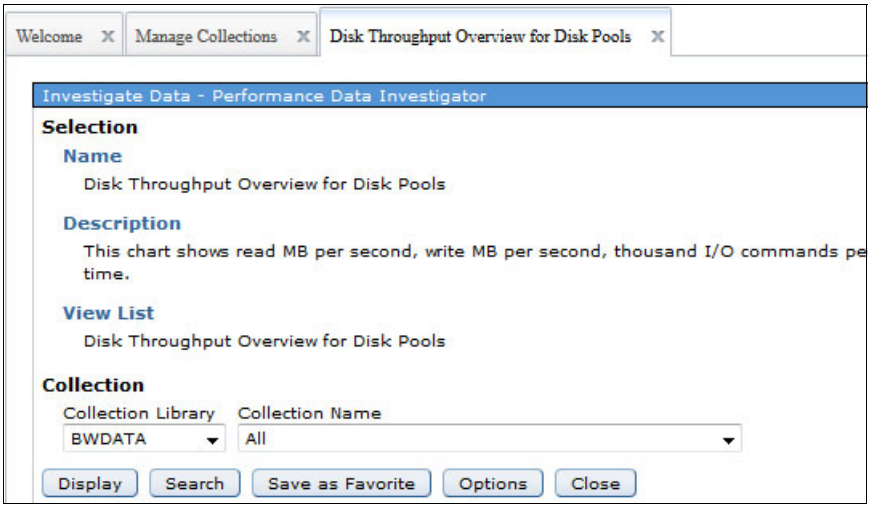

Figure 2-13 Performance Data Investigator, Disk Throughput Overview for Disk Pools collection criteria

15. The resulting graph is displayed. The Full Zoom Out tool icon that is shown in Figure 2-14 can be used to view the graph during the number of collections.

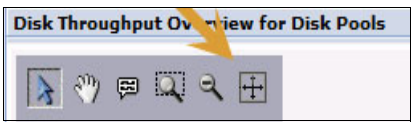

Figure 2-14 Performance Data Investigator, Full Zoom Out tool

16. The resulting graph that is shown in Figure 2-15 shows the IASP disk activity throughout the duration of the specified collections.

    In this example, only the data for a couple of days is available. However, it is a preferred practice for you to perform this entire process several times to view several weeks of data to view the amount of system IASP write activity.

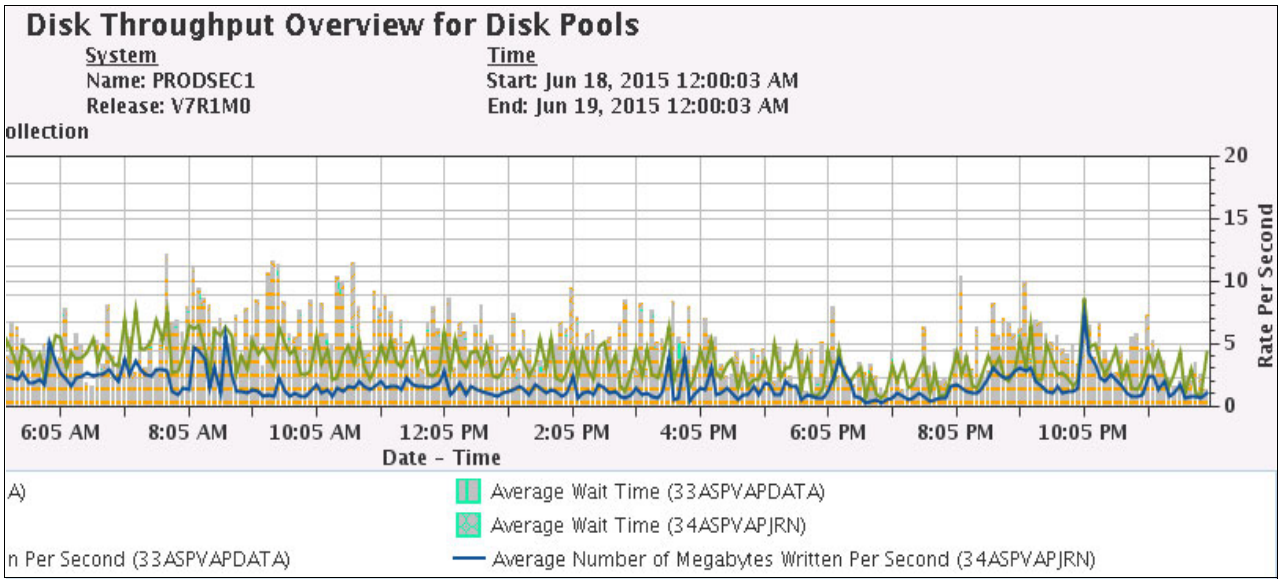

Figure 2-15 Performance Data Investigator Disk Throughput Overview for Disk Pools graph

On the graph that is shown in Figure 2-15, focus on the Average Number of Megabytes Written Per Second for the IASP to be used in the PowerHA environment.

If you did not migrate to an IASP yet, this tool can still be used to get a general idea of the number of megabytes that were written by looking at the Average Number of Megabytes Written Per Second for ASP 1. Review this graph for the peaks in the write rates to the IASP.

The Zoom Region tool in the toolbar can be used to view a specific range of time. This tool helps make the lines on the graph more readable. The write rates to the IASP during a particular time frame correspond to the amount of bandwidth that is needed for geographic mirroring traffic to sufficiently mirror the writes to the preferred target system during that same time frame.

Throughout the days of data that must be analyzed, peaks and valleys occur in the write rates to the IASP. If sufficient bandwidth is not available to geographic mirroring during any of these times, this process can result in performance issues on the preferred source system or an automatic suspension of geographic mirroring, or both.

The importance of the bandwidth analysis planning step cannot be stressed enough. In addition to the normal runtime disk writes, a full or partial synchronization process consumes additional network bandwidth and system resources.

2.6 System performance considerations

With geographic mirroring, IBM i performs the replication. It is important to consider performance when you plan to implement a geographic mirroring solution. Although asynchronous geographic mirroring allows more flexibility for the distance between systems, implications still result from undersizing the source, target, or the network infrastructure between the sites.

Minimizing the *latency* (that is, the time that the production system waits for the acknowledgment that the information was received on the target system) is key to good application performance.

The following sections describe how to size for good geographic mirroring performance.

2.6.1 General performance guidelines

When you implement geographic mirroring, various factors can influence the performance of the systems that are involved in the solution. To maximize the performance of the applications in a solution, several considerations affect your planning to maximize performance.

Two separate aspects must be considered when you size for a geographic mirroring environment:

- During the normal run time of the production environment, overhead will be added by geographic mirroring while the IBM i operating system sends disk writes to the target system.
- Consider the overhead and time that are required for the synchronization process. The synchronization process occurs when the target IASP is reconnected to the source IASP and changes are pushed from the current production system to the current backup system to make the two systems equivalent again.

The current backup system must be able to handle the production workload when a switchover or failover occurs. Size the target system to handle the production workload.

2.6.2 Production and backup system comparison

Geographic mirroring consumes resources on both the current production and the current backup systems. Optimal performance occurs when the current production and current backup systems are equivalent in terms of central processing unit (CPU), memory, and the disk subsystem, especially when you work with synchronous mirroring. An efficient communications link, as described in 2.4, “Communication considerations for geographic mirroring” on page 12, is assumed.

A system with a reduced CPU and memory configuration, such as the remnant system of an upgrade, can be used if you do not expect that it will satisfactorily run a full production workload. However, consider this option cautiously.

2.6.3 CPU considerations

Extra overhead exists on the CPU and memory when you perform geographic mirroring. Consider this overhead for both the source and the target systems. Geographic mirroring increases the CPU load to the system processors on both the system that owns the production copy of the IASP and the system that owns the mirror copy of the IASP. Sufficient CPU capacity must exist to handle this overhead.

As a rule, the partitions that are used to run geographic mirroring, on *both* the preferred source and the preferred target systems, need more than a partial processor. In a minimal CPU configuration, 5 - 20% CPU overhead can occur while you run geographic mirroring.

2.6.4 Memory considerations

Geographic mirroring requires extra memory in the machine pool during the synchronization process. For the optimal performance of geographic mirroring, particularly during synchronization, increase your machine pool size by at least the amount that is calculated by the following formula and use the Work with Shared Storage Pools (**WRKSHRPOOL**) command to set the machine pool size:

```
Extra machine pool size = 300 MB + (0.3 * number of disk arms in the IASP)
```

This extra memory is needed during the synchronization process on the system that owns the mirror copy of the IASP.

Important: The machine pool storage size must be large enough before you start the resynchronization. Otherwise, the increased memory is included after the synchronization process is in progress, and the synchronization process can take longer.

If the system value QPFRADJ is equal to 2 or 3**,** the system changes the storage pools automatically as needed based on the fault rate and priority settings within WRKSHRPOOL.

To prevent the performance adjuster function from reducing the machine pool size, take these steps:

1. Set the machine pool minimum size to the calculated amount (the current size plus the extra size for geographic mirroring from the formula) by using the **WRKSHRPOOL** command or the Change Shared Storage Pool (**CHGSHRPOOL**) command.
2. Set the Automatically adjust memory pools and activity levels (QPFRADJ) system value to zero, which prohibits the performance adjuster from changing the size of the machine pool.

Important: We consider it a preferred practice that you use the **WRKSHRPOOL** command to set the machine pool size to the calculated minimum. Disabling the Performance Auto Adjuster can affect performance in other areas of your environment.

An additional demand for memory occurs on the base pool on the current production system when the system is in asynchronous transmission delivery mode. The performance adjuster handles moving memory to the base pool to handle this additional page faulting if memory is available to move into this pool, as needed.

2.6.5 Disk subsystem considerations

Disk subsystem performance can affect overall geographic mirroring performance, especially when the disk subsystem is slower on the mirrored system. When mirroring is in synchronous mode, all write operations on the production copy are gated by the mirrored copy writes to disk. Therefore, a slow target disk subsystem can affect the source-side performance. You can minimize this effect on performance by running mirroring in asynchronous mode. Running in asynchronous mirroring mode alleviates the wait for the disk subsystem on the target side and sends back confirmation to the source side when the changed memory page is in memory on the target side.

System disk pool considerations

Similar to any system disk configuration, the number of disk units that are available to the application can affect its performance significantly. Adding workload on a limited number of disk units might result in longer disk waits and ultimately longer response times to the application. This additional workload is particularly important in relationship to temporary storage in a system that is configured with independent disk pools. All temporary storage (such as objects in the `QTEMP` library) is written to the SYSBAS disk pool. If your application does not use much temporary storage, you can function with fewer disk arms in the SYSBAS disk pool.

Although disk writes that are associated with production data will be on the IASP, disk activity will continue to be associated with the SYSBAS pool for the operating system and basic functions. As a start, use the guidelines that are shown in Table 2-3.

Table 2-3 Disk arms for (internal disk only) SYSBAS guidelines

Disk arms in IASPs	Arms for SYSBAS: Divide IASP arms by this number
Fewer than 20	3
20 - 40	4
Greater than 40	5

For example, if an IASP contains 10 drives, SYSBAS needs at least three drives. As another example, if an IASP contains 50 drives, SYSBAS needs at least 10 drives.

Note: Disk pool sizing depends on the application. These guidelines are only provided as a starting point. It is possible that fewer or more arms are required for a certain application environment. Understanding performance monitoring for the system's application environment is critical for sizing and capacity planning.

Monitoring the "percent busy" of the SYSBAS disks in the environment to ensure that the disk subsystem is sized correctly is always a good idea. In general, if this value exceeds 40% utilization, the size of the disk subsystem needs to be increased.

In an internal disk configuration environment, the disk that is assigned to the IASP must be placed on a separate I/O adapter from the SYSBAS disk to reduce any potential contention. Input/output adapter (IOA) cache is important because it provides greater data integrity and improved performance.

2.7 Backup planning for geographic mirroring

Before you implement HA that is based on geographic mirroring, you need to follow a backup strategy that adheres to your business needs.

If the current strategy is to end production applications and perform backups on the production system and your business has no requirement to change this strategy, you can continue to run backups as usual. This strategy ensures that the mirror copy is in sync for the longest period.

If you want backups from the target/mirror copy, geographic mirroring does not allow concurrent access to the mirror copy of the IASP. This rule has implications when you perform backups.

To back up the IASP data from the mirror copy system, run the following steps:

1. Quiesce the IASP on the current production copy system by using one of the following options:
   - The suspend option (*SUSPEND) of the Change Auxiliary Storage Pool Activity (**CHGASPACT**) command can be used:

     `CHGASPACT ASPDEV(<IASP_DEV_NAME>) OPTION(*SUSPEND)`

     The *SUSPEND option performs the following steps within the IASP:
     - Writes all changed pages from memory to disk.
     - Halts the initiation of new transactions and allows current transactions to reach the next commit boundary.
     - Halts the initiation of new non-transaction database writes and allows current writes to complete.
     - Writes all changed pages from memory to disk again.
   - The force write option (*FRCWRT) of the Change Auxiliary Storage Pool Activity (**CHGASPACT**) command can be used:

     `CHGASPACT ASPDEV(<IASP_DEV_NAME>) OPTION(*FRCWRT)`

     The *FRCWRT option writes only changed pages in memory at that time to disk, then it allows activity to the IASP to continue. Although this option is quicker than a suspend, it does not halt any transactions or wait for any commit boundaries to be reached.
2. Detach the mirror copy system IASP with tracking by using the following command. Tracking allows changes on the production system to be recorded so that those changes can be synchronized when the mirrored copy comes back online.

   `CHGASPSSN SSN(<SESSION_NAME>) OPTION(*DETACH) TRACK(*YES)`

   Using TRACK(*YES) requires a partial data resynchronization between the production and mirrored copies only. A detach without tracking requires a full synchronization to occur.
3. If the **CHGASPACT** command with the *SUSPEND option was used, the session can be resumed by using the following command:

   `CHGASPACT ASPDEV(<IASP_DEV_NAME>) OPTION(*RESUME)`
4. Vary on the detached copy of the IASP on the mirror copy system:

   `VRYCFG CFGOBJ(IASP_DEV) CFGTYPE(*DEV) STATUS(*ON)`
5. Perform the backup procedure.

6. Vary off the IASP on the mirror copy system:

   ```
   VRYCFG CFGOBJ(IASP_DEV) CFGTYPE(*DEV) STATUS(*OFF)
   ```

7. Reattach the mirror copy IASP to the original production host:

   ```
   CHGASPSSN SSN(<SESSION_NAME>) OPTION(*REATTACH)
   ```

While the IASP is detached on the mirror copy system, the changes on the production system are being tracked, but they are not transmitted to the mirror copy until it becomes available after the backup completes.

Important: The mirror copy system and production system will not be synchronized until all tracked changes are transmitted. During this synchronization time, no failover or switchover can occur.

To minimize synchronization time, which in turn limits exposure, the preferred practice is to always use tracking when you suspend or detach the mirror copy IASP.

3

Implementing geographic mirroring

This chapter provides detailed steps that show how to configure geographic mirroring on a two-node cluster. These steps are documented by using IBM i control language (CL) commands.

The following topics are described in this chapter:

3.1 Gathering information

We suggest that you complete a configuration worksheet with the basic required parameters to configure geographic mirroring before you continue. A blank example worksheet is provided in Appendix B, “Worksheet for configuring geographic mirroring” on page 99.

3.2 Configuring geographic mirroring

A geographic mirroring solution requires a cluster with two cluster nodes, an independent auxiliary storage pool (IASP), and optionally an administrative domain. These components make up the basic framework of geographic mirroring.

This section provides an example of setting up a two-node geographic mirroring high availability (HA) environment. This example assumes that the IASP was created already on the primary system. For details about how to create a cluster framework, see the IBM Redbooks publication, *IBM PowerHA SystemMirror for i: Preparation (Volume 1 of 4)*, SG24-8400.

Important: An IASP with the same name must *not* exist on the secondary or backup system.

The cluster resource group (CRG) dataport interfaces and Internet Protocol (IP) addresses must be configured and active before you continue. You can configure up to four interfaces on each node. Although it is not required, it is a preferred practice for a minimum of two redundant interfaces to be used in a production environment. It is important to ensure that every dataport address on a node can communicate with at least one dataport address on every other node in the CRG.

The values that are shown in Table 3-1 are used in the configuration example in this chapter. The settings for your configuration will differ.

Table 3-1 Configuration values that are used in the example in this chapter

Parameter	Keyword and type	Value	Description	Commands where parameter is used
IASP	ASPDEV CHAR(10)	IASPHA	IASP name or device description.	CFGDEVASP
Cluster	CLUSTER CHAR(10)	PWRHACLU	Cluster name.	CRTCLU ADDCLUMON ADDDEVDMN CRTCAD ADDCADMRE
Node identifier (primary)	NODE CHAR(8)	ITSO1NOD	Primary or source node name.	CRTCLU ADDCLUMON ADDDEVDMNE CRTCAD
Node identifier (backup)	NODE CHAR(8)	ITSO2NOD	Backup or target node name.	CRTCLU ADDCLUMON ADDDEVDMNE CRTCAD

Parameter	Keyword and type	Value	Description	Commands where parameter is used
Node identifier (additional)	NODE CHAR(8)	N/A	Additional node name, if required.	CRTCLU ADDCLUMON ADDDEVDMNE CRTCAD
IP address (primary node)	Dotted IP	192.168.80.172	Cluster IP address for the primary node. It can have 1 or 2 addresses.	CRTCLU
IP address (backup node)	Dotted IP	192.168.80.182	Cluster IP address for the backup node. It can have 1 or 2 addresses.	CRTCLU ADDCLUNODE
IP address (additional node)	Dotted IP	N/A	Cluster IP address for the additional node. It can have 1 or 2 addresses. Not applicable for all configurations.	CRTCLU ADDCLUNODE
Device domain	DEVDMN CHAR(10)	PWRHADMN	Device domain name.	ADDDEVDMNE
Cluster administrative domain	ADMDMN CHAR(10)	PWRHACAD	Cluster administrative domain name.	CRTCAD ADDCADMRE
Cluster resource group	CRG CHAR(10)	PWRHACRG	Cluster resource group (CRG) name.	CRTCRG
Site name (primary)	SRCSITE CHAR(8)	ITSO1STE	Primary node site name.	CRTCRG
Site name (backup)	TGTSITE CHAR(8)	ITSO2STE	Backup node site name.	CRTCRG
Dataport IP address (primary node)	Dotted IP	192.168.80.173 192.168.80.174	1 - 4 numeric IP addresses that are used for dataport services (primary).	CRTCRG
Dataport IP address (backup node)	Dotted IP	192.168.80.183 192.168.80.184	1 - 4 numeric IP addresses that are used for dataport services (backup).	CRTCRG
Session	SSN CHAR(10)	PWRHASSN	Auxiliary storage pool (ASP) session name.	CFGGEOMIR
Source ASP copy description	CHAR(10)	ITSO1CPY	Copy description on the primary node.	CFGGEOMIR
Backup ASP copy description	CHAR(10)	ITSO2CPY	Copy description on the backup node.	CFGGEOMIR

Run the following steps to configure geographic mirroring:

1. Create the IASP device description on the backup node as shown in Figure 3-1. The Device description and Resource name parameters must exactly match the IASP name that was used in the cluster framework configuration. For details about how to create a cluster framework, see the IBM Redbooks publication, *IBM PowerHA SystemMirror for i: Preparation (Volume 1 of 4)*, SG24-8400.

   Important: If the IASP description was added to the cluster administrative domain previously, the device description exists and this step is not required.

```
                    Create Device Desc (ASP) (CRTDEVASP)

Type choices, press Enter.

Device description . . . . . > IASPHA        Name
Resource name  . . . . . . . > IASPHA        Name
Relational database  . . . . .   *GEN
Message queue  . . . . . . . .   *SYSOPR       Name
  Library  . . . . . . . . . .                 Name, *LIBL, *CURLIB
Text 'description' . . . . . .   My GeoMir IASP

                                                                    Bottom
F3=Exit   F4=Prompt   F5=Refresh   F10=Additional parameters   F12=Cancel
F13=How to use this display        F24=More keys
```

Figure 3-1 Creating the IASP device description on the backup system

2. Create the cluster resource group (CRG) on the source node by using the **CRTCRG** command as shown in Figure 3-2. The CRG defines the resources that are involved in the availability relationship that you are creating.

```
                    Create Cluster Resource Group (CRTCRG)

Type choices, press Enter.

Cluster  . . . . . . . . . . . . > PWRHACLU      Name
Cluster resource group . . . . . > PWRHACRG      Name
Cluster resource group type  . . > *DEV          *DATA, *APP, *DEV, *PEER
CRG exit program . . . . . . . . > *NONE         Name, *NONE
  Library  . . . . . . . . . . .                 Name
User profile . . . . . . . . . . > *NONE         Name, *NONE
Recovery domain node list:
  Node identifier  . . . . . . . > ITSO1NOD      Name
  Node role  . . . . . . . . . . > *PRIMARY      *CRGTYPE, *PRIMARY...
  Backup sequence number . . . .   *LAST         1-127, *LAST
  Site name  . . . . . . . . . . > ITSO1STE      Name, *NONE
  Data port IP address . . . . . > '192.168.80.173'
                                   '192.168.80.174'

  Node identifier  . . . . . . . > ITSO2NOD      Name
  Node role  . . . . . . . . . . > *BACKUP       *CRGTYPE, *PRIMARY...
  Backup sequence number . . . . > 1             1-127, *LAST
  Site name  . . . . . . . . . . > ITSO2STE      Name, *NONE
  Data port IP address . . . . . > '192.168.80.183'
                                   '192.168.80.184'


Exit program format name . . . .   EXTPO100      EXTPO100, EXTPO101...
 Exit program data . . . . . . . *NONE

  Distribute info user queue . . .   *NONE         Name, *NONE
  Library  . . . . . . . . . . .                 Name
Job  . . . . . . . . . . . . . .   *JOBD         Name, *JOBD, *CRG
Configuration object list:
  Configuration object . . . . .   IASPHA        Name, *NONE
  Configuration object type  . .   *DEVD         *DEVD, *CTLD, *LIND, *NWSD
  Configuration object online  .   *ONLINE       *OFFLINE, *ONLINE, *PRIMARY
  Server takeover IP address . .   192.168.80.170

              + for more values
Text description . . . . . . . .   *BLANK

Failover message queue . . . . .   *NONE         Name, *NONE
  Library  . . . . . . . . . . .                 Name

F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
```

Figure 3-2 Creating a cluster resource group (multiple displays)

In Figure 3-2 on page 31, you can see the various parameters that can be entered for the **CRTCRG** command. Not all of the parameters are required. The parameters are described:

- CLUSTER: The name of your cluster.
- CRG: The name of the new cluster resource group (CRG) that you are creating. It must be unique within the cluster.
- CRGTYPE: Multiple types of CRGs exist. However, for geographic mirroring, you need to create a *DEV because you are managing switchable devices.
- Optional: EXITPGM: You can define a program to call whenever a cluster event occurs. If you define this program, you must also specify a user profile for it to run under. Usually, this value is left as *NONE. For more information, see the Cluster Resource Group Exit Program topic in the IBM i 7.2 Knowledge Center:

  http://www.ibm.com/support/knowledgecenter/ssw_ibm_i_72/apis/clrgexit.htm?lang=en
- USRPRF: This parameter is required if an exit program is specified. Otherwise, it must be *NONE.
- RCYDMN: This parameter defines the cluster nodes that are involved with this switchable resource. You need to specify the cluster node names, the node role (one *PRIMARY and one *BACKUP for geographic mirroring). All of the dataport IP addresses on one node must be able to communicate with all dataport IP addresses on the other node.
- EXITPGMFMT: This parameter indicates the parameter format to pass to the exit program (if specified).
- EXITPGMDTA: This short piece of data will be passed to the exit program.
- CRGMSGUSRQ: If you want to distribute cluster-wide information through this cluster, you must specify a USRQ object name here.
- JOB: This parameter is the name of the JOB that is used for any actions by the CRG.
- CFGOBJ: The IASP must be specified as a configuration object to link it to the CRG. Setting the object online parameter to *ONLINE automatically varies on the IASP at the completion of a switch.
- FLVMSGQ: A separate failover message queue, failover wait time, and failover action can be designated for the CRG. If *NONE is specified, these values default to the settings that were configured for the cluster.

For more information about the parameters for the **CRTCRG** command, see the Create Cluster Resource Group (CRTCRG) topic in the IBM i 7.2 Knowledge Center:

https://www.ibm.com/support/knowledgecenter/ssw_ibm_i_72/cl/crtcrg.htm?lang=en

3. Run the Configure Geographic Mirroring (**CFGGEOMIR**) command on the source node as shown in Figure 3-3. The IASP on the source node must be varied off to run this command.

```
                  Configure Geographic Mirror (CFGGEOMIR)

Type choices, press Enter.

ASP device . . . . . . . . . . . . > IASPHA        Name
Action . . . . . . . . . . . . . . > *CREATE       *CREATE, *DELETE
Source site  . . . . . . . . . . .   *             Name, *
Target site  . . . . . . . . . . .   *             Name, *
Session  . . . . . . . . . . . . . > PWRHASSN      Name, *NONE
  Source ASP copy description  . >   ITSO1CPY     Name
  Target ASP copy description  . >   ITSO2CPY     Name
Transmission delivery  . . . . .   *SYNC          *SYNC, *ASYNC
Disk units . . . . . . . . . . .   *SELECT        Name, *SELECT
               + for more values

                           Additional Parameters

Confirm  . . . . . . . . . . . .   *YES           *YES, *NO
Cluster  . . . . . . . . . . . .   *             Name, *
Cluster resource group . . . . .   *             Name, *


Mirroring mode . . . . . . . . .   *SYNC          *SYNC, *ASYNC
Synchronization priority . . . .   *HIGH          *HIGH, *MEDIUM, *LOW
Suspend timeout  . . . . . . . .   120            60-3600
Tracking space . . . . . . . . .   100            0-100

F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
```

Figure 3-3 Configuring geographic mirroring

When you use the **CFGGEOMIR** command, you need to specify the IASP name that you are making highly available and the *CREATE action.

You are not required to specify the site names that you configured previously in the CRG because only two sites are in the CRG recovery domain.

Although the session name is optional, it is usual to specify a name here. This mirroring session name links the source and target copy descriptions that you specify. If these names do not exist, the command creates them for you.

This example shows the Transmission delivery mode of *SYNC (synchronous) and the Mirroring mode of *SYNC. These names can be changed now based on the node configuration or later by using the Change Auxiliary Storage Pool Session (**CHGASPSSN**) command. You can also change the synchronization priority, timeouts, and tracking space later.

The tracking space is a percentage of the maximum available space and not a percentage of the IASP. Unless a specific reason requires you to choose a certain value, 100 is a good choice.

Specifying *SELECT for Disk units presents a panel of all non-configured disk units on the target site (backup node) as shown in Figure 3-4. Use option 1 to select the disk units that you want from the list. The total capacity must be 95% or greater of the source site (or primary) copy.

```
                        Select Non-Configured Disk Units

ASP device . . . . . :   IASPHA        Mirror copy node . . . :   ITSO2NOD
Prod copy capacity . . :   111864      Mirror copy capacity . :   0
Prod disk units  . . . :   3           Mirror disk units  . . :   0

Type options, press Enter.
  1=Select

     Resource
Opt  Name        Serial Number    Type Model  Capacity  Rank  Eligible
 1   DD002       Y4XSQ7YJMG86     6B22 0050      37287  002      Yes
 1   DD005       YHYY7N8D5ARX     6B22 0050      37287  002      Yes
 1   DD006       YEH7YY7PAUTB     6B22 0050      37287  002      Yes

                                                                  Bottom
F1=Help F9=Calculate Selection F11=View 2 F12=Cancel
```

Figure 3-4 Selection window of non-configured units for the IASP

4. Optional: If either the source or target nodes (or both) use the IBM Storage Area Network (SAN) Volume Controller (SVC) external storage server or the IBM Storwize external storage server and if you want to use IBM FlashCopy to copy the IASP, you must remove the ASP copy description for that node and add an SVC copy description by using the following steps:

   a. End the ASP session by using the following End Auxiliary Storage Pool Session (**ENDASPSSN**) command. This command also deletes the session configuration.

      ```
      ENDASPSSN SSN(PWRHASSN)
      ```

   b. Remove the ASP copy description by using the following Remove Auxiliary Storage Pool Copy Description (**RMVASPCPYD**) command:

      ```
      RMVASPCPYD ASPCPY(ITSO1CPY)
      ```

   c. Add the correct copy description by using either the Add Auxiliary Storage Pool Copy Description (**ADDASPCPYD**) command or the Add SAN Volume Controller Auxiliary Storage Pool Copy Description (**ADDSVCCPYD**) command.

An SVC example is shown in Figure 3-5.

Note: The Storage host (SVCHOST) and Virtual disk range (VRTDSKRNG) parameters must be obtained from the person who is responsible for configuring the external storage device.

```
                  Add SVC ASP Copy Description (ADDSVCCPYD)

Type choices, press Enter.

ASP copy . . . . . . . . . . . . . ASPCPY       > ITSO1CPYF
ASP device . . . . . . . . . . . . ASPDEV       > IASPHA
Cluster resource group . . . . . . CRG          > PWRHACRG
Cluster resource group site  . . . SITE         > ITSO1STE
Node identifier  . . . . . . . . . NODE         > ITSO1NOD
Storage host:                      SVCHOST
  User name  . . . . . . . . . .                _________
  Secure shell key file  . . . .                __________

  Internet address . . . . . . .                ____________________________

Virtual disk range:                VRTDSKRNG
  Range start  . . . . . . . . .                _______
  Range end  . . . . . . . . . .                _______
  Host identifier  . . . . . . .                  *ALL
                          + for more values
                          + for more values
Device domain  . . . . . . . . . . DEVDMN         *
Recovery domain:                   RCYDMN
  Cluster node . . . . . . . . .                  *NONE
  Host identifier  . . . . . . .
                          + for more values
                          + for more values
                                                                        More...
 F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
```

Figure 3-5 Adding the SVC copy description

d. Start (create) the ASP session by using the Start ASP Session (**STRASPSSN**) command, as shown in Figure 3-6.

The Session name is a unique name. Because this name is for geographic mirroring, the Session type is *GEOMIR. You then need to specify the source and target copy descriptions. Keep the default settings for the other parameters.

```
                        Start ASP Session (STRASPSSN)

Type choices, press Enter.

Session  . . . . . . . . . . . . > PWRHASSN      Name
Session type . . . . . . . . . . > *GEOMIR       *GEOMIR, *METROMIR...
ASP copy:
  Preferred source . . . . . . . > ITSO1CPY      Name
  Preferred target . . . . . . . > ITSO2CPY      Name
              + for more values
Device domain  . . . . . . . . .   *             Name, *
Transmission delivery  . . . . .   *CFG          *CFG, *SYNC, *ASYNC



                                                                        Bottom
F3=Exit   F4=Prompt   F5=Refresh   F10=Additional parameters   F12=Cancel
F13=How to use this display        F24=More keys
```

Figure 3-6 Starting an ASP session after you change a copy description

5. Start the CRG by using the Start Cluster Resource Group (**STRCRG**) command as shown in Figure 3-7.

```
                     Start Cluster Resource Group (STRCRG)

Type choices, press Enter.

Cluster  . . . . . . . . . . . .   PWRHACLU      Name
Cluster resource group . . . . .   PWRHACRG      Name
Exit program data  . . . . . . .   *SAME



                                                                        Bottom
 F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
```

Figure 3-7 Starting the cluster resource group

6. Vary on the IASP on the primary node by using the **VRYCFG** command. This command also starts the initial synchronization of the source and target IASPs. Depending on the size of the IASP and the bandwidth between the nodes, the time to complete the command might be extensive.

```
VRYCFG CFGOBJ(IASPHA) CFGTYPE(*DEV) STATUS(*ON)
```

7. Run the Display ASP Session (**DSPASPSSN**) command to display the status of the mirroring configuration as shown in Figure 3-8. The state of the production role shows AVAILABLE and the state of mirror role shows ACTIVE. Geographic mirroring is functioning correctly and it is in sync.

Note: The Display ASP Session panel can also be displayed by using the Work with Cluster (**WRKCLU**) command, selecting option 10, and then using option 25.

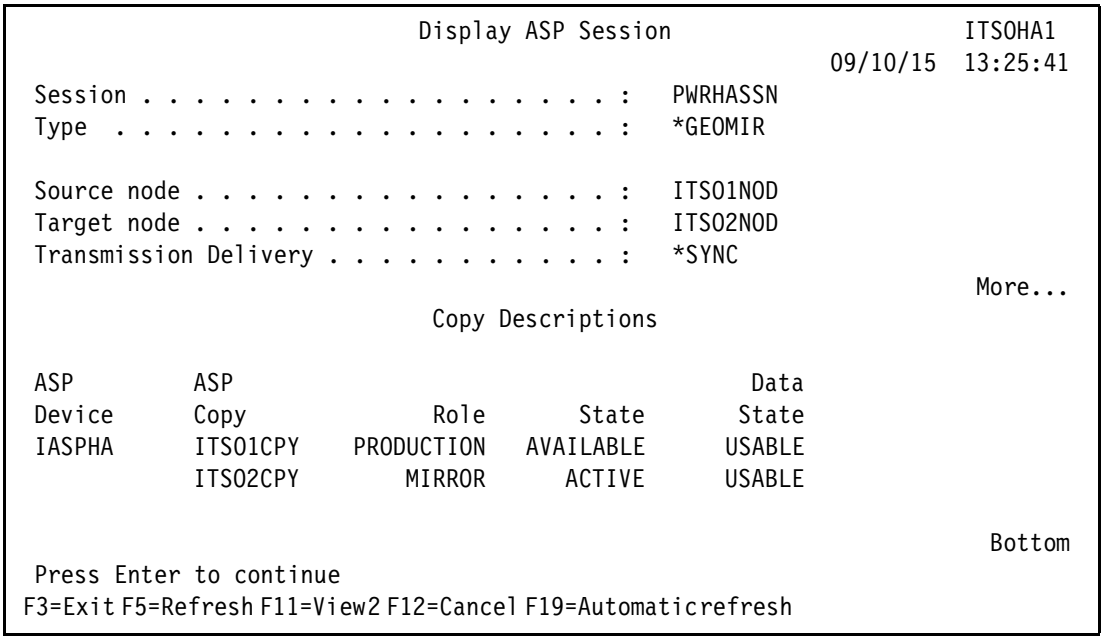

Figure 3-8 Displaying an ASP session that shows an active PowerHA configuration

The setup for geographic mirroring is complete.

4

Monitoring and managing IBM PowerHA SystemMirror for i

This chapter describes methods that help you monitor and manage your IBM PowerHA SystemMirror for i in a geographic mirroring environment.

The following topics are described in this chapter:

Note: All of the sample control language (CL) programs that are shown in this chapter are available as a save file. For more information about how to access and download these save files, see Appendix C, “Additional material” on page 103.

4.1 Monitoring your environment

Monitoring of your environment is probably the most important task in a high-availability (HA) environment.

Historically, it was considered sufficient to check your environment once a day. However, if you check only once a day, you only learn that your availability solution is unusable at the time you run the check, and that your environment was unusable for anything from a minute to 24 hours ago. If you check only once a day and find a failure, on average you must assume that the solution was not available for you to use for an average of 12 hours. Putting that in perspective, your availability solution is only usable approximately 50% of the time.

This situation is easily prevented by implementing an automated, proactive monitoring regime. By using tools and automation to monitor on a regular and automatic basis, you are alerted when situations need a resolution. Therefore, you understand your ability to switch if a problem occurs with your production partition.

Important: It is imperative that you implement effective and complete monitoring processes for your specific environment. These monitoring tools must be as resilient as the systems that they monitor. Failing to monitor correctly can lead to problems later.

Consider many areas when you monitor your environment, including the following areas:

- 4.1.1, “QSYSOPR message queue” on page 40
- 4.1.2, “Monitoring the status of your replication” on page 41
- 4.1.3, “Monitoring the status of cluster nodes” on page 45
- 4.1.4, “Monitoring the status of the cluster administrative domain” on page 48
- 4.1.5, “Monitoring the cluster resource group status” on page 50
- 4.1.6, “Monitoring the status of the independent auxiliary storage pool” on page 53
- 4.1.7, “Other monitoring considerations” on page 57

Different environments have different characteristics and different potential points of failure. Ensure that any monitoring processes that you implement can identify any of these potential points of failure to give you time to make informed decisions about the resolution of any issues that arise. What you do not want is to discover that you have a problem when you need to switch in a hurry.

Note: The example control language (CL) programs that are provided in this chapter are designed for use with an IBM PowerHA SystemMirror for i (PowerHA) environment. If you have an IBM Systems Lab Services Toolkit, they have other monitoring methods that complement or replace the basic PowerHA methods. These methods are identified as part of the installation service of the toolkit.

4.1.1 QSYSOPR message queue

The first place to always look for information about the health of your replication solution is the QSYSOPR message queue. PowerHA periodically checks the health of the environment and sends messages to the QSYSOPR message queue to indicate any problems that it identifies.

4.1.2 Monitoring the status of your replication

The most important area to monitor is the actual replication. It is imperative that any issues here are identified as early as possible. If the replicated copy of data is unusable, the whole solution is put at risk.

Traditionally, on a logical replication solution, you monitor the status of your replication once a day. With the hardware solutions, it is a simpler process, so you can monitor the status of the replication more frequently. However, do not run it *too* frequently.

Note: A good policy might be to monitor 2 or 3 times a day. You might monitor first in the morning, in the middle of the day, and before major batch processes are scheduled to run. However, your environment will have its own requirements. Do not be tempted to run every half hour, which is excessive.

To monitor for any replication problems, you can write a simple CL program that checks the current state of the replication with the Retrieve Auxiliary Storage Pool Session (**RTVASPSSN**) command and alert your system operators by sending a message to the QSYSOPR message queue or by using the Send SMTP Email (**SNDSMTPEMM**) command to send an email message to an email recipient.

Example 4-1 shows a sample CL program to assist you in building your own monitoring tool for replication sessions. This CL program is available as a save file. For more information about how to access and download this save file, see Appendix C, “Additional material” on page 103.

Tip: If you want to use the sample code that is shown in Example 4-1, you must remember to specify the name of an appropriate email address if you want to use that method of notification.

Example 4-1 Sample CL program to monitor replication status

```
/*****************************************************************************/
/* MONASPSSN - Monitor the replication status                                */
/*                                                                           */
/*             This program is intended to provide a simple way of           */
/*             monitoring the replication status                             */
/*                                                                           */
/*             This is achieved by using the RTVASPSSN command to obtain     */
/*             information about the replication environment and based       */
/*             upon the type of environment it will identify the status of   */
/*             the environment.                                              */
/*                                                                           */
/*             If the replication is not in a fully consistent status, it    */
/*             will send a message to the system operator and also send an   */
/*             email to the named user to alert the administrators to the    */
/*             fact that they should investigate the problems.               */
/*                                                                           */
/*            ============================================================  */
/*           I NOTE:  This sample code assumes that there are only 2 (TWO)  I */
/*           I copy descriptions involved in the session.                   I */
/*           I                                                              I */
/*           I        It is possible that more could exist in which case    I */
/*           I this sample code would need to be reworked to allow for this I */
```

```
/*           I possibility.                                                 I */
/*            ============================================================  */
/*****************************************************************************/
/* Written by  : David Painter                                               */
/* Date Written: Sep 3, 2015                                                 */
/* Date Changed: None                                                        */
/*****************************************************************************/
pgm
             DCL        VAR(&SSN) TYPE(*CHAR) LEN(10)VALUE('GEOMIR    ')
             DCL        VAR(&TYPE) TYPE(*CHAR) LEN(10)
             DCL        VAR(&ASPCPYLST) TYPE(*CHAR) LEN(4472)
             DCL        VAR(&DELIVERY) TYPE(*CHAR) LEN(8)
             DCL        VAR(&MODE) TYPE(*CHAR) LEN(8)
             DCL        VAR(&INTRANSIT) TYPE(*DEC) LEN(10 0)
             DCL        VAR(&TIMEOUT) TYPE(*DEC) LEN(4 0)
             DCL        VAR(&PRIORITY) TYPE(*CHAR) LEN(8)
             DCL        VAR(&TRACKSPACE) TYPE(*DEC) LEN(3 0)
             DCL        VAR(&PERSISTENT) TYPE(*CHAR) LEN(8)
             DCL        VAR(&FLASHTYPE) TYPE(*CHAR) LEN(8)
             DCL        VAR(&ASPCPY) TYPE(*CHAR) LEN(31232)

/* Variables used to extract the copy information */
             DCL        VAR(&OFFIASP1) TYPE(*INT) STG(*DEFINED) +
                          LEN(4) DEFVAR(&ASPCPY 1)
             DCL        VAR(&LENIASP) TYPE(*INT) STG(*DEFINED) +
                          LEN(4) DEFVAR(&ASPCPY 5)
             DCL        VAR(&NUMIASP) TYPE(*INT) STG(*DEFINED) +
                          LEN(4) DEFVAR(&ASPCPY 9)

/* Variables for Copy information 1 */
             DCL        VAR(&ASPCPY1) TYPE(*CHAR) LEN(70)
             DCL        VAR(&ASPCPYD1) TYPE(*CHAR) STG(*DEFINED) +
                          LEN(10) DEFVAR(&ASPCPY1 1)
             DCL        VAR(&ASPDEV1) TYPE(*CHAR) STG(*DEFINED) +
                          LEN(10) DEFVAR(&ASPCPY1 11)
             DCL        VAR(&ASPROLE1) TYPE(*CHAR) STG(*DEFINED) +
                          LEN(10) DEFVAR(&ASPCPY1 21)
             DCL        VAR(&ASPSTATE1) TYPE(*CHAR) STG(*DEFINED) +
                          LEN(10) DEFVAR(&ASPCPY1 31)
             DCL        VAR(&ASPDSTATE1) TYPE(*CHAR) STG(*DEFINED) +
                          LEN(10) DEFVAR(&ASPCPY1 41)
             DCL        VAR(&ASPNODE1) TYPE(*CHAR) STG(*DEFINED) +
                          LEN(8) DEFVAR(&ASPCPY1 51)
             DCL        VAR(&ASPPCTTRK1) TYPE(*int) STG(*DEFINED) +
                          LEN(4) DEFVAR(&ASPCPY1 59)
             DCL        VAR(&ASPOOS1) TYPE(*int) STG(*DEFINED) +
                          LEN(4) DEFVAR(&ASPCPY1 63)
             DCL        VAR(&ASPSYNCP1) TYPE(*INT) STG(*DEFINED) +
                          LEN(4) DEFVAR(&ASPCPY1 67)

/* Variables for Copy information 2 */
             DCL        VAR(&ASPCPY2) TYPE(*CHAR) LEN(70)
             DCL        VAR(&ASPCPYD2) TYPE(*CHAR) STG(*DEFINED) +
                          LEN(10) DEFVAR(&ASPCPY2 1)
             DCL        VAR(&ASPDEV2) TYPE(*CHAR) STG(*DEFINED) +
```

```
                    LEN(10) DEFVAR(&ASPCPY2 11)
             DCL    VAR(&ASPROLE2) TYPE(*CHAR) STG(*DEFINED) +
                    LEN(10) DEFVAR(&ASPCPY2 21)
             DCL    VAR(&ASPSTATE2) TYPE(*CHAR) STG(*DEFINED) +
                    LEN(10) DEFVAR(&ASPCPY2 31)
             DCL    VAR(&ASPDSTATE2) TYPE(*CHAR) STG(*DEFINED) +
                    LEN(10) DEFVAR(&ASPCPY2 41)
             DCL    VAR(&ASPNODE2) TYPE(*CHAR) STG(*DEFINED) +
                    LEN(8) DEFVAR(&ASPCPY2 51)
             DCL    VAR(&ASPPCTTRK2) TYPE(*int) STG(*DEFINED) +
                    LEN(4) DEFVAR(&ASPCPY2 59)
             DCL    VAR(&ASPOOS2) TYPE(*INT) STG(*DEFINED) +
                    LEN(4) DEFVAR(&ASPCPY2 63)
             DCL    VAR(&ASPSYNCP2) TYPE(*INT) STG(*DEFINED) +
                    LEN(4) DEFVAR(&ASPCPY2 67)

/* Work variables */
             DCL    VAR(&X) TYPE(*DEC) LEN(5 0)

/* Step 1 - Use RTVASPSSN to obtain the session data */
             RTVASPSSN  SSN(&SSN) TYPE(&TYPE) ASPCPYLST(&ASPCPYLST) +
                          DELIVERY(&DELIVERY) MODE(&MODE) +
                          INTRANSIT(&INTRANSIT) TIMEOUT(&TIMEOUT) +
                          PRIORITY(&PRIORITY) +
                          TRACKSPACE(&TRACKSPACE) +
                          PERSISTENT(&PERSISTENT) +
                          FLASHTYPE(&FLASHTYPE) ASPCPY(&ASPCPY)
             MONMSG     MSGID(HAE0000 CPF0000) EXEC(DO)
             SNDPGMMSG  MSGID(CPF9898) MSGF(QCPFMSG) MSGDTA('An +
                          error occurred with the MONASPSSN command, +
                          check joblog for more information') +
                          MSGTYPE(*ESCAPE)
             RETURN
             ENDDO

/* Step 2 - If there are more than 2 copy descriptions involved issue a warning */
             IF         COND(&NUMIASP *NE 2) THEN(DO)
             SNDPGMMSG  MSGID(CPF9898) MSGF(QCPFMSG) MSGDTA('There +
                          are more than 2 copy descriptions +
                          involved, this sample is invalid') +
                          MSGTYPE(*ESCAPE)
             RETURN
             ENDDO

/* Step 3 - Extract the two data about the source and target ASP copies */
             CHGVAR     VAR(&X) VALUE(&OFFIASP1 + 1)
             CHGVAR     VAR(&ASPCPY1) VALUE(%SST(&ASPCPY &X &LENIASP))
             CHGVAR     VAR(&X) VALUE(&X + &LENIASP)
             CHGVAR     VAR(&ASPCPY2) VALUE(%SST(&ASPCPY &X &LENIASP))

/* Step 4 - Based upon the type of session, calculate what is a success state */
/* FlashCopy */

             IF         COND(&TYPE *EQ '*FLASHCOPY') THEN(DO)
             SNDPGMMSG  MSGID(CPF9898) MSGF(QCPFMSG) MSGDTA('A +
```

```
                    flashcopy session is not supported for +
                    this function') MSGTYPE(*ESCAPE)
          RETURN
          ENDDO

/* Geographic Mirroring */
          IF        COND(&TYPE *EQ '*GEOMIR') THEN(DO)
          IF        COND((&ASPROLE1 *EQ 'MIRROR') *AND +
                    ((&ASPDSTATE1 *EQ 'USABLE') *OR +
                    (&ASPDSTATE1 *EQ 'INSYNC'))) THEN(DO)
          SNDPGMMSG MSGID(CPF9898) MSGF(QCPFMSG) MSGDTA('The +
                    IASP replication is in the correct status +
                    for this setup') MSGTYPE(*COMP)
          RETURN
          ENDDO
          IF        COND((&ASPROLE2 *EQ 'MIRROR') *AND +
                    ((&ASPDSTATE2 *EQ 'USABLE') *OR +
                    (&ASPDSTATE2 *EQ 'INSYNC'))) THEN(DO)
          SNDPGMMSG MSGID(CPF9898) MSGF(QCPFMSG) MSGDTA('The +
                    IASP replication is in the correct status +
                    for this setup') MSGTYPE(*COMP)
          RETURN
          ENDDO
          ENDDO

/* Metro Mirror */
          IF        COND(&TYPE *EQ '*METROMIR') THEN(DO)
          IF        COND((&ASPROLE1 *EQ 'TARGET') *AND +
                    (&ASPSTATE1 *NE 'ACTIVE') *AND +
                    (&ASPSTATE1 *NE 'AVAILABLE')) THEN(DO)
          SNDPGMMSG MSGID(CPF9898) MSGF(QCPFMSG) MSGDTA('The +
                    IASP replication is in the correct status +
                    for this setup') MSGTYPE(*COMP)
          RETURN
          ENDDO
          IF        COND((&ASPROLE2 *EQ 'TARGET') *AND +
                    (&ASPSTATE2 *NE 'ACTIVE') *AND +
                    (&ASPSTATE2 *NE 'AVAILABLE')) THEN(DO)
          SNDPGMMSG MSGID(CPF9898) MSGF(QCPFMSG) MSGDTA('The +
                    IASP replication is in the correct status +
                    for this setup') MSGTYPE(*COMP)
          RETURN
          ENDDO
          ENDDO

/* Global Mirror */
          IF        COND(&TYPE *EQ '*GLOBALMIR') THEN(DO)
          IF        COND((&ASPROLE1 *EQ 'TARGET') *AND +
                    (&ASPSTATE1 *NE 'ACTIVE') *AND +
                    (&ASPSTATE1 *NE 'AVAILABLE')) THEN(DO)
          SNDPGMMSG MSGID(CPF9898) MSGF(QCPFMSG) MSGDTA('The +
                    IASP replication is in the correct status +
                    for this setup') MSGTYPE(*COMP)
          RETURN
          ENDDO
```

```
            IF         COND((&ASPROLE2 *EQ 'TARGET') *AND +
                         (&ASPSTATE2 *NE 'ACTIVE') *AND +
                         (&ASPSTATE2 *NE 'AVAILABLE')) THEN(DO)
            SNDPGMMSG  MSGID(CPF9898) MSGF(QCPFMSG) MSGDTA('The +
                         IASP replication is in the correct status +
                         for this setup') MSGTYPE(*COMP)
            RETURN
            ENDDO
            ENDDO

/* Step 5 - Send an email to an admin user if there is an error */
            SNDSMTPEMM RCP((SYSTEMADMINISTRATOR@YOURCOMPANY)) +
                         SUBJECT('IASP replication problems have +
                         been found.') NOTE('Some problems were +
                         found when checking the IASP +
                         replication.  You should use the +
                         DSPASPSSN command to review the situation +
                         and take any necessary actions to correct +
                         the errors.')
            MONMSG     MSGID(TCP5090 TCP5092)

/* Step 6 - Send a message to QSYSOPR if there is an error */
            SNDPGMMSG  MSGID(CPF9898) MSGF(QCPFMSG) MSGDTA('Errors +
                         were found with the IASP replication, +
                         please review') TOMSGQ(*SYSOPR)

/* Step 7 - Send an escape message to indicate there is an error */
            SNDPGMMSG  MSGID(CPF9898) MSGF(QCPFMSG) MSGDTA('Errors +
                         were found with the IASP replication, +
                         please review') MSGTYPE(*ESCAPE)
            RETURN
endpgm
```

4.1.3 Monitoring the status of cluster nodes

For PowerHA to function, you need to ensure that the cluster nodes are ACTIVE under normal conditions. Therefore, you must monitor to verify that the cluster nodes are ACTIVE under normal conditions.

When you verify the status of the cluster nodes, you must perform this verification from the perspective of a single node. You retrieve the local node status and then from the local node you also obtain the remote node status. This process is important because if you fail to perform this process correctly, it can result in you thinking that the cluster nodes are both active when in reality they are in a PARTITION status.

Figure 4-1 shows an example of a partitioned cluster. If you check on the partition LONDON, it states that it is ACTIVE, and if you check on the partition FRANKFURT, it also states that it is ACTIVE. Only by performing the check on a single partition can you see that the cluster is in a PARTITION state.

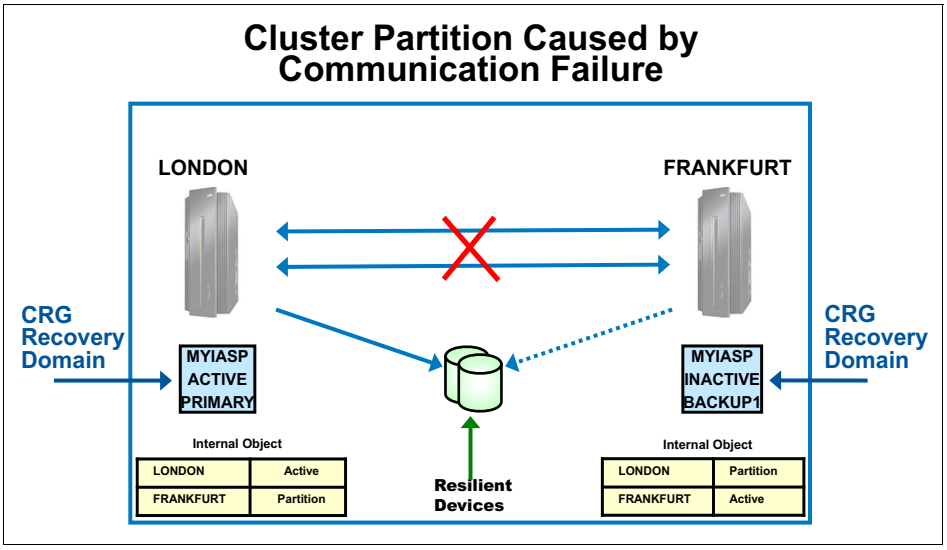

Figure 4-1 A partitioned cluster

Important: Do not use remote commands to query other nodes in the cluster. You will not obtain the correct information.

Example 4-2 shows a sample CL program that can assist you in building your own monitoring tool for the cluster node status. This CL program is available as a save file. For more information about how to access and download this save file, see Appendix C, “Additional material” on page 103.

Example 4-2 Sample CL program to monitor the cluster node status

```
/*****************************************************************************/
/* MONNODSTS - Monitor the cluster node status                               */
/*                                                                           */
/*             This program is intended to provide a simple way of           */
/*             monitoring the status of the cluster nodes                    */
/*                                                                           */
/*             This is achieved by using system APIs to obtain the current   */
/*             status of the nodes, it should be run on a single node only   */
/*             to avoid missing a partition condition.                       */
/*****************************************************************************/
/* Written by  : David Painter                                               */
/* Date Written: Sep 3, 2015                                                 */
/* Date Changed: None                                                        */
/*                                                                           */
/* Required service programs: QHASM/QHAAPI, QCSTCTL1                         */
/*****************************************************************************/
pgm
             DCL        VAR(&CLUNAME) TYPE(*CHAR) LEN(10)

/* API Variables */
```

```
             DCL        VAR(&APIRCVR) TYPE(*CHAR) LEN(18)
             DCL        VAR(&APIRCVRLEN) TYPE(*INT) LEN(4) VALUE(18)
             DCL        VAR(&APIUSRSPC) TYPE(*CHAR) LEN(20)

/* User space header fields and pointers */
             DCL        VAR(&PTR_USRSPC) TYPE(*PTR)
             DCL        VAR(&HEADER) TYPE(*CHAR) STG(*BASED) +
                          LEN(192) BASPTR(&PTR_USRSPC)
             DCL        VAR(&OFFLIST) TYPE(*UINT) STG(*DEFINED) +
                          LEN(4) DEFVAR(&HEADER 125)
             DCL        VAR(&NUMLIST) TYPE(*UINT) STG(*DEFINED) +
                          LEN(4) DEFVAR(&HEADER 133)
             DCL        VAR(&LENENTRY) TYPE(*UINT) STG(*DEFINED) +
                          LEN(4) DEFVAR(&HEADER 137)

/* User space list data fields and pointers */
             DCL        VAR(&PTR_LIST) TYPE(*PTR)
             DCL        VAR(&LISTDATA) TYPE(*CHAR) STG(*BASED) +
                          LEN(512) BASPTR(&PTR_LIST)
             DCL        VAR(&LENNODE) TYPE(*UINT) STG(*DEFINED) +
                          LEN(4) DEFVAR(&LISTDATA 1)
             DCL        VAR(&NODE) TYPE(*CHAR) STG(*DEFINED) LEN(8) +
                          DEFVAR(&LISTDATA 5)
             DCL        VAR(&NODESTS) TYPE(*UINT) STG(*DEFINED) LEN(4) +
                          DEFVAR(&LISTDATA 13)

/* Work variables */
             DCL        VAR(&ERRFLG) TYPE(*CHAR) LEN(1) VALUE('N')
             DCL        VAR(&X) TYPE(*UINT) LEN(4)

/* Step 1 - Use the QhaRetrieveHAInfo API to obtain the cluster name */
             CALLPRC    PRC('QhaRetrieveHAInfo') PARM((&APIRCVR) +
                          (&APIRCVRLEN) ('RHAI0100') (X'00000000'))
             CHGVAR     VAR(&CLUNAME) VALUE(%SST(&APIRCVR 9 10))

/* Step 2 - Use the QcstListClusterInfo API to obtain the cluster node status */
             CHGVAR     VAR(&APIUSRSPC) VALUE('CLUAPI    QGPL      ')
             CALL       PGM(QUSCRTUS) PARM(&APIUSRSPC 'SAMPLE    ' +
                          X'00000001' X'00' '*ALL      ' 'Temporary +
                          user space                          ')
             MONMSG     MSGID(CPF9870)
             CALLPRC    PRC('QcstListClusterInfo') PARM((&APIUSRSPC) +
                          (&CLUNAME) ('LCTI0100') ('*ALL    ') +
                          (X'00000000'))

/* Step 3 - Get a pointer to the USRSPC, and then get pointer to the first list
data */
             CALL       PGM(QUSPTRUS) PARM(&APIUSRSPC &PTR_USRSPC)
             CHGVAR     VAR(&PTR_LIST) VALUE(&PTR_USRSPC)
             CHGVAR     VAR(%OFFSET(&PTR_LIST)) VALUE(&OFFLIST)
             CHGVAR     VAR(&X) VALUE(1)

/* Step 4 - If this node is not active, then there is an error */
 TESTSTS:    IF         COND(&NODESTS *NE 2) THEN(DO)
             SNDPGMMSG  MSGID(CPF9898) MSGF(QCPFMSG) MSGDTA('Node: ' +
```

```
                        *CAT &NODE *TCAT ' is not active') +
                        MSGTYPE(*DIAG)
             CHGVAR     VAR(&ERRFLG) VALUE('Y')
             ENDDO

/* Step 5 - Have we processed all the entries,  if not increment and test again */
             IF         COND(&X *LT &NUMLIST) THEN(DO)
             CHGVAR     VAR(&X) VALUE(&X + 1)
             CHGVAR     VAR(%OFFSET(&PTR_LIST)) +
                        VALUE(%OFFSET(&PTR_LIST) + &LENNODE)
             GOTO       CMDLBL(TESTSTS)
             ENDDO

/* Step 6 - Send a message to the user if there is no error */
             IF         COND(&ERRFLG *NE 'Y') THEN(DO)
             SNDPGMMSG  MSGID(CPF9898) MSGF(QCPFMSG) MSGDTA('All +
                        cluster nodes are active') MSGTYPE(*COMP)
             RETURN
             ENDDO

/* Step 7 - Send an email to an admin user if there is an error */
             SNDSMTPEMM RCP((SYSTEMADMINISTRATOR@YOURCOMPANY)) +
                        SUBJECT('Not all cluster nodes are +
                        active, please investigate.') NOTE('Not +
                        all of the cluster nodes are active, +
                        while this could be deliberate it should +
                        be investigated to ensure that there is +
                        no exposure to the availability or +
                        replication solution.)')
             MONMSG     MSGID(TCP5090 TCP5092)

/* Step 6 - Send an message to QSYSOPR if there is an error */
             SNDPGMMSG  MSGID(CPF9898) MSGF(QCPFMSG) MSGDTA('Not all +
                        cluster nodes are active, please +
                        investigate') TOMSGQ(*SYSOPR)

/* Step 7 - Send an escape message to indicate there is an error */
             SNDPGMMSG  MSGID(CPF9898) MSGF(QCPFMSG) MSGDTA('Not all +
                        cluster nodes are active, please +
                        investigate') MSGTYPE(*ESCAPE)
             RETURN
endpgm
```

4.1.4 Monitoring the status of the cluster administrative domain

The *cluster administrative domain* is used to monitor multiple types of resources. Although it is reliable, sometimes the nodes that are involved are out of sync. For example, if you deleted a user profile and forgot to remove the monitored resource entry from the cluster administrative domain, an inconsistency results. Another example is if you add a user with a specific job description that does not exist on one of the other nodes.

To monitor for these types of problems, you can write a simple CL program to monitor for any entries that are not CONSISTENT and alert your system operators or monitoring tools as your business needs dictate.

You can monitor the cluster administrative domain by using the Print Cluster Administrative Domain Monitored Resource Entry (**PRTCADMRE**) command to produce an output file of monitored resource entries that are not in a global consistent state. Then, you can discover how many entries are in the file and if the number is not zero, you can alert the system operators by sending a message to the QSYSOPR message queue or by using the Send SMTP Email (**SNDSMTPEMM**) command to send an email message to an email recipient.

Note: If you use the sample code that is shown in Example 4-3, you must remember to specify the name of the cluster administration domain and an appropriate email address if you plan to use that method of notification.

Example 4-3 shows a sample CL program that can assist you in building your own monitoring tool for the cluster administrative domain. This CL program is available as a save file. For more information about how to access and download this save file, see Appendix C, “Additional material” on page 103.

Example 4-3 Sample CL program to monitor the cluster administrative domain

```
/*****************************************************************************/
/* MONADMDMN - Monitor the cluster administrative domain                     */
/*                                                                           */
/*             This program is intended to provide a simple way of           */
/*             monitoring the Cluster Administrative Domain                  */
/*                                                                           */
/*             This is achieved by simply generating a file of Monitored     */
/*             Resource Entries that are not in a consistent state           */
/*                                                                           */
/*             If this file has any records in it then there are issues so a */
/*             message is sent to the QSYSOPR message queue and an email     */
/*             is sent to a named user as well.  This alerts the admin team  */
/*             that they should investigate the problems.                    */
/*                                                                           */
/*****************************************************************************/
/* Written by  : David Painter                                               */
/* Date Written: Sep 2, 2015                                                 */
/* Date Changed: None                                                        */
/*****************************************************************************/
pgm
             DCL        VAR(&NBRCURRCD) TYPE(*DEC) LEN(10 0)

/* Step 1 - Use the PRTCADMRE to generate a file of errors */
             PRTCADMRE  ADMDMN(ADMDMN) RSCGLBSTS(*ADDED *ENDED +
                          *FAILED *INCONSISTENT *PENDING) +
                          DETAIL(*RSC) OUTPUT(*OUTFILE) +
                          OUTFILE(QTEMP/ADMDMNERR)
             MONMSG     MSGID(HAE0000 CPF0000) EXEC(DO)
             SNDPGMMSG  MSGID(CPF9898) MSGF(QCPFMSG) MSGDTA('An +
                          error occurred with the MONADMDMN command, +
                          check joblog for more information') +
                          MSGTYPE(*ESCAPE)
             RETURN
             ENDDO

/* Step 2 - Retrieve the number of records in the error file */
             RTVMBRD    FILE(QTEMP/ADMDMNERR) NBRCURRCD(&NBRCURRCD)
```

```
/* Step 3 - If the number of records is 0 then terminate normally */
             IF         COND(&NBRCURRCD *EQ 0) THEN(DO)
             SNDPGMMSG  MSGID(CPF9898) MSGF(QCPFMSG) MSGDTA('The +
                          Cluster Administrative Domain is all +
                          synchronized') MSGTYPE(*COMP)
             RETURN
             ENDDO

/* Step 4 - Send an email to an admin user if there is an error */
             SNDSMTPEMM RCP((SYSTEMADMINISTRATOR@YOURCOMPANY)) +
                          SUBJECT('Cluster Administrative Domain +
                          errors have been found.') NOTE('Some +
                          errors were found when checking the +
                          Cluster Administrative Domain.  You +
                          should use the WRKCADMRE command to +
                          review the situation and take any +
                          necessary actions to correct the errors.')
             MONMSG     MSGID(TCP5090 TCP5092)

/* Step 5 - Send a message to QSYSOPR if there is an error */
             SNDPGMMSG  MSGID(CPF9898) MSGF(QCPFMSG) MSGDTA('Errors +
                          were found in the Cluster Administrative +
                          Domain, please review') TOMSGQ(*SYSOPR)

/* Step 6 - Send an escape message to indicate there is an error */
             SNDPGMMSG  MSGID(CPF9898) MSGF(QCPFMSG) MSGDTA('Errors +
                          were found in the Cluster Administrative +
                          Domain, please review') MSGTYPE(*ESCAPE)
             RETURN
endpgm
```

4.1.5 Monitoring the cluster resource group status

Use cluster resource groups (CRGs) to define replication characteristics for the PowerHA configuration. They are also effectively an “ON/OFF” switch to instruct PowerHA if you expect it to take charge in a failover situation. Therefore, it is important that your CRGs are ACTIVE under normal operations.

To monitor for these problems, you can write a simple CL program to alert your system operators or monitoring tools when CRGs are not in an ACTIVE status, as your business needs dictate.

You can monitor the status of CRGs by using system application programming interface (API) programs to provide the status.

Note: If you use the sample code that is in Example 4-4, you must remember to specify the name of an appropriate email address if you plan to use that method of notification.

Example 4-4 shows a sample CL program that can assist you in building your own monitoring tool for the CRG status. This CL program is available as a save file. For more information about how to access and download this save file, see Appendix C, “Additional material” on page 103.

Example 4-4 Sample CL program to monitor the CRG status

```
/*****************************************************************************/
/* MONCRGSTS - Monitor the cluster resource group status                     */
/*                                                                           */
/*             This program is intended to provide a simple way of           */
/*             monitoring the status of the cluster resource groups          */
/*                                                                           */
/*             This is achieved by using system APIs to obtain the current   */
/*             status of the CRGs.                                           */
/*                                                                           */
/*****************************************************************************/
/* Written by  : David Painter                                               */
/* Date Written: Sep 9, 2015                                                 */
/* Date Changed: None                                                        */
/*                                                                           */
/* Required service programs: QCSTCRG3                                       */
/*****************************************************************************/
pgm
             DCL        VAR(&CLUNAME) TYPE(*CHAR) LEN(10)

/* API Variables */
             DCL        VAR(&APIRCVR) TYPE(*CHAR) LEN(18)
             DCL        VAR(&APIRCVRLEN) TYPE(*INT) LEN(4) VALUE(18)
             DCL        VAR(&APIUSRSPC) TYPE(*CHAR) LEN(20)

/* User space header fields and pointers */
             DCL        VAR(&PTR_USRSPC) TYPE(*PTR)
             DCL        VAR(&HEADER) TYPE(*CHAR) STG(*BASED) +
                          LEN(192) BASPTR(&PTR_USRSPC)
             DCL        VAR(&OFFLIST) TYPE(*UINT) STG(*DEFINED) +
                          LEN(4) DEFVAR(&HEADER 125)
             DCL        VAR(&NUMLIST) TYPE(*UINT) STG(*DEFINED) +
                          LEN(4) DEFVAR(&HEADER 133)
             DCL        VAR(&LENENTRY) TYPE(*UINT) STG(*DEFINED) +
                          LEN(4) DEFVAR(&HEADER 137)

/* User space list data fields and pointers */
             DCL        VAR(&PTR_LIST) TYPE(*PTR)
             DCL        VAR(&LISTDATA) TYPE(*CHAR) STG(*BASED) +
                          LEN(512) BASPTR(&PTR_LIST)
             DCL        VAR(&CRGNAME) TYPE(*CHAR) STG(*DEFINED) +
                          LEN(10) DEFVAR(&LISTDATA 1)
             DCL        VAR(&CRGTYPE) TYPE(*UINT) STG(*DEFINED) +
                          LEN(2) DEFVAR(&LISTDATA 11)
             DCL        VAR(&CRGSTS) TYPE(*UINT) STG(*DEFINED) LEN(4) +
                          DEFVAR(&LISTDATA 13)
             DCL        VAR(&CRGPRI) TYPE(*CHAR) STG(*DEFINED) LEN(8) +
                          DEFVAR(&LISTDATA 17)
             DCL        VAR(&CRGAPPID) TYPE(*CHAR) STG(*DEFINED) +
                          LEN(2) DEFVAR(&LISTDATA 25)
```

```
/* Work variables */
             DCL        VAR(&ERRFLG) TYPE(*CHAR) LEN(1) VALUE('N')
             DCL        VAR(&X) TYPE(*UINT) LEN(4)

/* Step 1 - Use the QhaRetrieveHAInfo API to obtain the cluster name */
             CALLPRC    PRC('QhaRetrieveHAInfo') PARM((&APIRCVR) +
                          (&APIRCVRLEN) ('RHAI0100') (X'00000000'))
             CHGVAR     VAR(&CLUNAME) VALUE(%SST(&APIRCVR 9 10))

/* Step 2 - Use the QcstListClusterResourceGroups API to obtain the cluster node
status */
             CHGVAR     VAR(&APIUSRSPC) VALUE('CLUAPI    QGPL      ')
             CALL       PGM(QUSCRTUS) PARM(&APIUSRSPC 'SAMPLE    ' +
                          X'00000001' X'00' '*ALL      ' 'Temporary +
                          user space                              ')
             MONMSG     MSGID(CPF9870)
             CALLPRC    PRC('QcstListClusterResourceGroups') +
                          PARM((&APIUSRSPC) ('CRGL0100') (&CLUNAME) +
                          (X'00000000'))

/* Step 3 - Get a pointer to the USRSPC, and then get pointer to the first list
data */
             CALL       PGM(QUSPTRUS) PARM(&APIUSRSPC &PTR_USRSPC)
             CHGVAR     VAR(&PTR_LIST) VALUE(&PTR_USRSPC)
             CHGVAR     VAR(%OFFSET(&PTR_LIST)) VALUE(&OFFLIST)
             CHGVAR     VAR(&X) VALUE(1)

/* Step 4 - If this CRG is not active, then there is an error */
 TESTSTS:    IF         COND(&CRGSTS *NE 10) THEN(DO)
             SNDPGMMSG  MSGID(CPF9898) MSGF(QCPFMSG) MSGDTA('CRG: ' +
                          *CAT &CRGNAME *TCAT ' is not active') +
                          MSGTYPE(*DIAG)
             CHGVAR     VAR(&ERRFLG) VALUE('Y')
             ENDDO

/* Step 5 - Have we processed all the entries,  if not increment and test again */
             IF         COND(&X *LT &NUMLIST) THEN(DO)
             CHGVAR     VAR(&X) VALUE(&X + 1)
             CHGVAR     VAR(%OFFSET(&PTR_LIST)) +
                          VALUE(%OFFSET(&PTR_LIST) + &LENENTRY)
             GOTO       CMDLBL(TESTSTS)
             ENDDO

/* Step 6 - Send an email to an admin user if there is an error */
             IF         COND(&ERRFLG *NE 'Y') THEN(DO)
             SNDPGMMSG  MSGID(CPF9898) MSGF(QCPFMSG) MSGDTA('All +
                          cluster resource groups are active') +
                          MSGTYPE(*COMP)
             RETURN
             ENDDO

/* Step 7 - Send an email to an admin user if there is an error */
             SNDSMTPEMM RCP((SYSTEMADMINISTRATOR@YOURCOMPANY)) +
                          SUBJECT('Not all cluster resource groups +
```

```
                          are active, please review.') NOTE('Not +
                          all of the cluster resource groups are +
                          active, while this could be deliberate it +
                          should be investigated to ensure that +
                          there is no exposure to the availability +
                          or replication solution.)')
             MONMSG     MSGID(TCP5090 TCP5092)

/* Step 6 - Send an message to QSYSOPR if there is an error */
             SNDPGMMSG  MSGID(CPF9898) MSGF(QCPFMSG) MSGDTA('Not all +
                          cluster resource groups are active, +
                          please reviews') TOMSGQ(*SYSOPR)

/* Step 7 - Send an escape message to indicate there is an error */
             SNDPGMMSG  MSGID(CPF9898) MSGF(QCPFMSG) MSGDTA('Not all +
                          cluster resource groups are active, +
                          please review') MSGTYPE(*ESCAPE)
             RETURN
endpgm
```

4.1.6 Monitoring the status of the independent auxiliary storage pool

One important consideration for the independent auxiliary storage pool (IASP) is that it cannot overflow into SYSBAS. Therefore, you need to review the current utilization regularly.

Although you can set a threshold within System Service Tools (SST), this setting notifies you only when the threshold is exceeded. You can also write a program to use the system API **QYASPOL**, which can provide you with the total and current capacity. It is also possible to obtain the number of logical unit numbers (LUNs) that are involved in the IASP configuration.

By using the information from the system API, it is possible to raise an alert when the utilization grows by more than a certain percentage since the last time the check was run, or if the number of disks in the IASP changed, which might require a corresponding change in the PowerHA configuration.

Note: If you use the sample code that is in Example 4-5, you must remember to specify the name of an appropriate email address if you plan to use that method of notification.

Example 4-5 shows a sample CL program that can assist you in building your own monitoring tool to verify the IASP utilization status. This CL program is available as a save file. For more information about how to access and download this save file, see Appendix C, “Additional material” on page 103.

Example 4-5 Sample CL program to verify the IASP utilization status

```
/*****************************************************************************/
/* MONIASPSTS- Monitor the IASP status                                       */
/*                                                                           */
/*             This program is intended to provide a simple way of           */
/*             monitoring the status of the IASP, specifically the level     */
/*             of capacity utilization, and number of disks.
/*                                                                           */
/*             This is achieved by using system APIs to obtain information   */
/*             about the IASP and to alert the operations should any         */
```

```
/*             potential problems exist.                                    */
/*                                                                          */
/*             If problems are identified, then the program                 */
/*             will send a message to the system operator and also send an   */
/*             email to the named user to alert the administrators to the    */
/*             fact that they should investigate the problems.              */
/*                                                                          */
/*            ============================================================  */
/*           I NOTE:  This sample code assumes that there are only 2 (TWO)  I */
/*           I copy descriptions involved in the session. It also assumes   I */
/*           I that only a single IASP exists. It is possible that more     I */
/*           I could exist in which case this sample code would need to be  I */
/*           I be reworked to allow for this possibility                    I */
/*            ============================================================  */
/****************************************************************************/
/* Written by  : David Painter                                              */
/* Date Written: Sep 10, 2015                                               */
/* Date Changed: None                                                       */
/****************************************************************************/
pgm
/* Variables for growth calculation, Note this is the percentage growth that
triggers an event */
             DCL        VAR(&GROWTH) TYPE(*DEC) LEN(5 1) VALUE(10.0)

/* QYASPOL variables */
             DCL        VAR(&YASPRCVR) TYPE(*CHAR) LEN(148)
             DCL        VAR(&YASPRCVLEN) TYPE(*UINT) LEN(4) VALUE(148)
             DCL        VAR(&YASPLINF) TYPE(*CHAR) LEN(80)
             DCL        VAR(&YASPREC) TYPE(*INT) LEN(4) VALUE(-1)
             DCL        VAR(&YASPFTRNUM) TYPE(*INT) LEN(4) VALUE(1)
             DCL        VAR(&YASPFTR) TYPE(*CHAR) LEN(16)
             DCL        VAR(&YASPFTRSIZ) TYPE(*UINT) STG(*DEFINED) +
                          LEN(4) DEFVAR(&YASPFTR 1)
             DCL        VAR(&YASPFTRKEY) TYPE(*UINT) STG(*DEFINED) +
                          LEN(4) DEFVAR(&YASPFTR 5)
             DCL        VAR(&YASPFTRDSZ) TYPE(*UINT) STG(*DEFINED) +
                          LEN(4) DEFVAR(&YASPFTR 9)
             DCL        VAR(&YASPFTRDTA) TYPE(*INT) STG(*DEFINED) +
                          LEN(4) DEFVAR(&YASPFTR 13)

/* Break out variables for &YASPRCVR */
             DCL        VAR(&ASPNUM) TYPE(*UINT) STG(*DEFINED) +
                          LEN(4) DEFVAR(&YASPRCVR 1)
             DCL        VAR(&DSKNUM) TYPE(*UINT) STG(*DEFINED) +
                          LEN(4) DEFVAR(&YASPRCVR 5)
             DCL        VAR(&ASPCAPT) TYPE(*UINT) STG(*DEFINED) +
                          LEN(4) DEFVAR(&YASPRCVR 9)
             DCL        VAR(&ASPCAPTA) TYPE(*UINT) STG(*DEFINED) +
                          LEN(4) DEFVAR(&YASPRCVR 13)
             DCL        VAR(&ASPCAPP) TYPE(*UINT) STG(*DEFINED) +
                          LEN(4) DEFVAR(&YASPRCVR 17)
             DCL        VAR(&ASPCAPPA) TYPE(*UINT) STG(*DEFINED) +
                          LEN(4) DEFVAR(&YASPRCVR 21)
             DCL        VAR(&ASPCAPU) TYPE(*UINT) STG(*DEFINED) +
                          LEN(4) DEFVAR(&YASPRCVR 25)
```

```
             DCL        VAR(&ASPCAPUA) TYPE(*UINT) STG(*DEFINED) +
                          LEN(4) DEFVAR(&YASPRCVR 29)
             DCL        VAR(&ASPSTGTHR) TYPE(*UINT) STG(*DEFINED) +
                          LEN(4) DEFVAR(&YASPRCVR 61)

/* Utilization variables */
             DCL        VAR(&UTILPCT) TYPE(*DEC) LEN(5 1)  /* Percent Util*/
             DCL        VAR(&UTILPCTD) TYPE(*DEC) LEN(5 1) /* Difference in Util*/
             DCL        VAR(&UTILPCTL) TYPE(*DEC) LEN(5 1) /* Last Percent Util */
             DCL        VAR(&IASPUTIL) TYPE(*CHAR) LEN(32)
             DCL        VAR(&LSTUSED) TYPE(*UINT) STG(*DEFINED) +
                          LEN(4) DEFVAR(&IASPUTIL 1)
             DCL        VAR(&LSTCAP) TYPE(*UINT) STG(*DEFINED) +
                          LEN(4) DEFVAR(&IASPUTIL 5)
             DCL        VAR(&LSTDAT) TYPE(*DEC) STG(*DEFINED) LEN(6 +
                          0) DEFVAR(&IASPUTIL 9)
             DCL        VAR(&LSTTIM) TYPE(*DEC) STG(*DEFINED) +
                          LEN(6 0) DEFVAR(&IASPUTIL 15)
             DCL        VAR(&LSTDSK) TYPE(*UINT) STG(*DEFINED) +
                          LEN(4) DEFVAR(&IASPUTIL 21)
             DCL        VAR(&CURUSED) TYPE(*UINT) LEN(4)
             DCL        VAR(&QDATE) TYPE(*CHAR) LEN(6)
             DCL        VAR(&QTIME) TYPE(*CHAR) LEN(6)
             DCL        VAR(&ERRDSK) TYPE(*CHAR) LEN(1) VALUE('N')
             DCL        VAR(&ERRUTIL) TYPE(*CHAR) LEN(1) VALUE('N')

/* Step 1 - Use the QYASPOL YASP0200 to obtain the IASP data */
             CHGVAR     VAR(&YASPFTRSIZ) VALUE(16)
             CHGVAR     VAR(&YASPFTRKEY) VALUE(1)
             CHGVAR     VAR(&YASPFTRDSZ) VALUE(4)
             CHGVAR     VAR(&YASPFTRDTA) VALUE(-3)
             CALL       PGM(QYASPOL) PARM(&YASPRCVR &YASPRCVLEN +
                          &YASPLINF &YASPREC &YASPFTRNUM &YASPFTR +
                          'YASP0200' X'00000000')

/* Step 2 - Calculate the current percentage capacity utilization */
             CHGVAR     VAR(&CURUSED) VALUE(&ASPCAPT - &ASPCAPTA)
             CHGVAR     VAR(&UTILPCT) VALUE(100.0 * (&CURUSED / +
                          &ASPCAPT))

/* Step 3 - Retrieve the previous capacity utilization percentage */
             RTVSYSVAL  SYSVAL(QDATE) RTNVAR(&QDATE)
             RTVSYSVAL  SYSVAL(QTIME) RTNVAR(&QTIME)
             RTVDTAARA  DTAARA(MONIASPSTS *ALL) RTNVAR(&IASPUTIL)
             MONMSG     MSGID(CPF1015) EXEC(DO) /* *DTAARA does not +
                          exist, yet */
             CRTDTAARA  DTAARA(MONIASPSTS) TYPE(*CHAR) LEN(32)
             CHGVAR     VAR(&LSTUSED) VALUE(&ASPCAPT - &ASPCAPTA)
             CHGVAR     VAR(&LSTCAP) VALUE(&ASPCAPT)
             CHGVAR     VAR(&LSTDAT) VALUE(000101)
             CHGVAR     VAR(&LSTTIM) VALUE(000000)
             CHGVAR     VAR(&LSTDSK) VALUE(&DSKNUM)
             ENDDO
/* Step 4 - If the growth of the storage is an increase on the last utilization,*/
/* then error. The logic here is that if the actual storage used has increased  */
```

```
/* by more than the desired percentage since last run, then an alert is raised. */
/* If the actual capacity has changed. then the process is bypassed.            */
/* Check if installed capacity has changed */
             IF         COND(&ASPCAPT *NE &LSTCAP) THEN(DO)
             GOTO       CMDLBL(CHKDSK)
             ENDDO
             IF         COND((((&CURUSED - &LSTUSED) / &LSTUSED) * +
                          100) *LT &GROWTH) THEN(DO)
             GOTO       CMDLBL(CHKDSK)
             ENDDO
             CHGVAR     VAR(&ERRUTIL) VALUE('Y')

/* Step 5 - Has the number of disks changed? */
 CHKDSK:
             IF         COND(&LSTDSK *NE &DSKNUM) THEN(CHGVAR +
                          VAR(&ERRDSK) VALUE('Y'))

/* Step 6 - Send an email to an admin user if there is an error */
             IF         COND(&ERRDSK *EQ 'Y') THEN(DO)
             SNDSMTPEMM RCP((SYSTEMADMINISTRATOR@YOURCOMPANY)) +
                          SUBJECT('The IASP disk config has +
                          changed, verify PowerHA config.') +
                          NOTE('The configuration of your IASP has +
                          changed, please verify that your PowerHA +
                          configuration has been modified for these +
                          changes.')
             MONMSG     MSGID(TCP5090 TCP5092)
             ENDDO
             IF         COND(&ERRUTIL *EQ 'Y') THEN(DO)
             SNDSMTPEMM RCP((SYSTEMADMINISTRATOR@YOURCOMPANY)) +
                          SUBJECT('The IASP utilization has grown +
                          by more than expected.') NOTE('The IASP +
                          utilization has grown by more than the +
                          growth percentage specified in the +
                          monitor program.  You should investigate +
                          the situation to ensure appropriate +
                          actions are taken.')
             MONMSG     MSGID(TCP5090 TCP5092)
             ENDDO

/* Step 7 - Send a message to QSYSOPR if there is an error */
             IF         COND(&ERRDSK *EQ 'Y') THEN(DO)
             SNDPGMMSG  MSGID(CPF9898) MSGF(QCPFMSG) MSGDTA('The +
                          disk configuration has changed for the +
                          IASP, Please review PowerHA settings') +
                          TOMSGQ(*SYSOPR)
             ENDDO
             IF         COND(&ERRUTIL *EQ 'Y') THEN(DO)
             SNDPGMMSG  MSGID(CPF9898) MSGF(QCPFMSG) MSGDTA('The +
                          IASP utilization has grown by more than +
                          the expected growth percentage,  please +
                          review') TOMSGQ(*SYSOPR)
             ENDDO

/* Step 8 - Update data area with the current information */
```

```
             CHGVAR     VAR(&LSTUSED) VALUE(&ASPCAPT - &ASPCAPTA)
             CHGVAR     VAR(&LSTCAP) VALUE(&ASPCAPT)
             CHGVAR     VAR(&LSTDAT) VALUE(&QDATE)
             CHGVAR     VAR(&LSTTIM) VALUE(&QTIME)
             CHGVAR     VAR(&LSTDSK) VALUE(&DSKNUM)
             CHGDTAARA  DTAARA(MONIASPSTS *ALL) VALUE(&IASPUTIL)

/* Step 9 - Send an escape message to indicate there is an error */
             IF         COND((&ERRDSK *EQ 'Y') *OR (&ERRUTIL *EQ +
                          'Y')) THEN(DO)
             SNDPGMMSG  MSGID(CPF9898) MSGF(QCPFMSG) MSGDTA('Errors +
                          were found with the IASP utilization or +
                          configuration') MSGTYPE(*ESCAPE)
             ENDDO
             RETURN
endpgm
```

4.1.7 Other monitoring considerations

In addition to the monitoring that is described previously in this chapter, you need to investigate other methods of monitoring. The following list shows typical areas to monitor:

- Network switches
- Power service
- Storage area network (SAN) switches
- Storage subsystems

Any of these areas can cause system failures and affect your ability to use the systems.

Often, these areas (and other areas) can be monitored by the use of Simple Network Management Protocol (SNMP) monitoring software. SNMP is an industry standard that allows a device to raise a *trap*, which is a notification that is broadcast on your network. An SNMP monitor can then receive the trap and decode it, which allows the monitor to raise alerts in real time.

For example, a storage subsystem typically raises a trap if it loses connectivity to a remote storage subsystem or if a physical component fails internally.

4.2 Managing your environment

Although little management is required for a PowerHA environment, you need to consider several areas. The following areas are covered in this section:

4.2.1 Verifying the health of your replication solution

Sometimes, you want to look at the overall health of your replication solution. With a PowerHA environment, you can look at the PowerHA graphical user interface (GUI) as shown in Figure 4-2.

As you can see in Figure 4-2, multiple indicators show the health of your replication. At the top of the window, the local node shows a green check mark, which indicates that PowerHA is operational on this node. Further down, you can see that the CRG also has a green check mark, which indicates that the CRG is active. You can see the green arrow in the list of nodes, which in this example shows that the replication is working and replicating from CTCIHA4A to CTCIHA4D.

If any box is not green, a problem exists and you must investigate.

Tip: It is useful to hover the mouse over the color indicator. A pop-up message appears and describes the indicated condition.

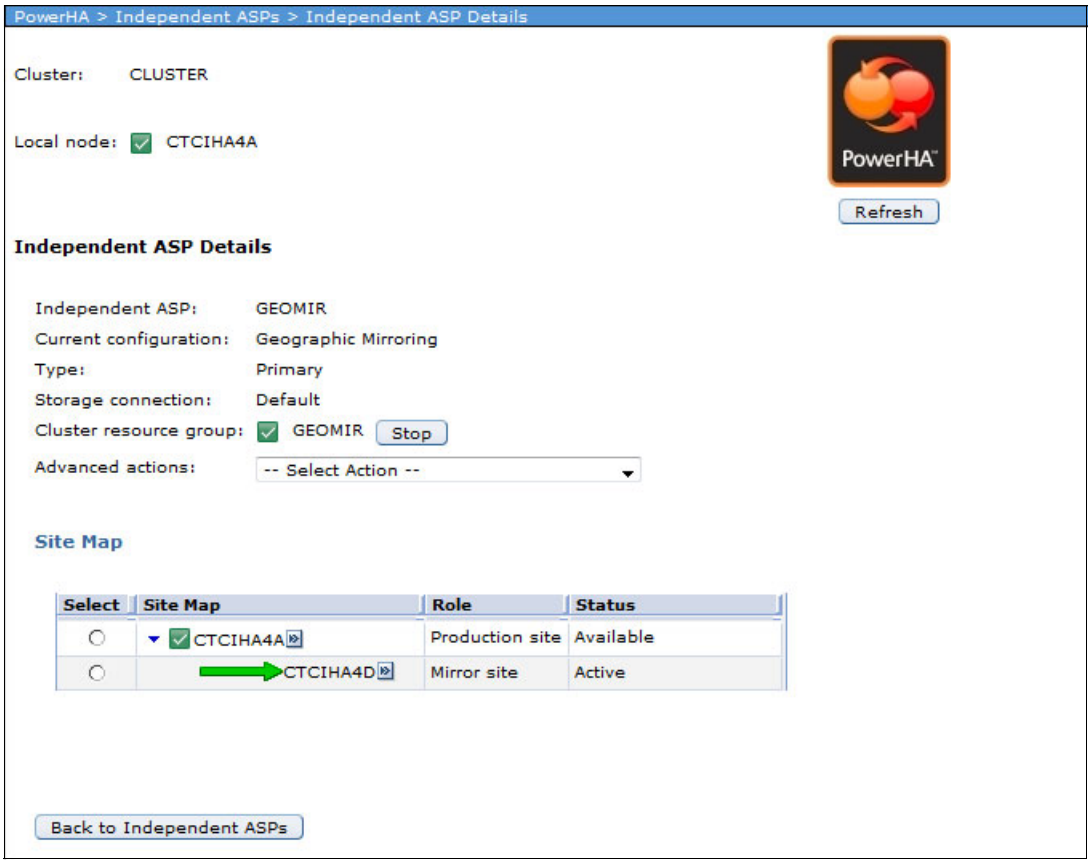

Figure 4-2 PowerHA GUI view of the health of the replication

4.2.2 Adding users to the systems

When you need to add users to the systems, you must also add them to the cluster administrative domain.

Depending on how often you add users, you can use either of the following options:

- Use the Add Cluster Admin Domain MRE (**ADDCADMRE**) command to add a user to the cluster administrative domain.

  You need to write a simple CL program that performs the **ADDCADMRE** command for you automatically.

- Use the system-supplied exit point on the Create User Profile (**CRTUSRPRF**) command to automatically add the user to the list of monitored resource entries (MREs).

  The exit point is called QIBM_QSY_CRT_PROFILE. The documentation for this exit point is in the IBM Knowledge Center:

  http://www.ibm.com/support/knowledgecenter/ssw_ibm_i_72/apis/XCRTUP.htm?lang=en

  Important: When you write a program for the QIBM_QSY_CRT_PROFILE exit point and then when you create a user profile on other nodes in the partition, that profile also calls the exit program. Therefore, include a check to determine whether the user that is running the program is QSYS. If yes, terminate with no further action.

4.2.3 Adding network printers and other SYSBAS objects to the systems

Many objects in SYSBAS can be replicated with the cluster administrative domain. How you manage these various objects varies, depending on many factors.

For devices, such as network printers, which are not created frequently, adding them manually after you create them is often the sensible approach. For other objects, you might use a program that creates them in a specific way, in which case, it makes sense to change that program to include the Add Cluster Admin Domain MRE (**ADDCADMRE**) command.

Important: Whichever route you choose, certain objects “happen to be the same” and need to be managed under change control, rather than “must be the same” and need to be replicated.

4.2.4 Powering down your system

When you power down a cluster node, the cluster considers the node as failed unless certain steps are followed. These steps differ slightly, depending on whether the system is the primary node or the backup node.

Follow the correct sequence to ensure that mirroring and clustering end correctly (in addition to any existing application power-down procedures) to avoid a failover:

1. If the system is the backup node, the source node (production) is not affected:
   a. Suspend mirroring from the primary node.
   b. End the backup cluster node.
   c. The backup node can now be powered down.
2. If the system is the primary node and this outage is planned where production applications do not need to be available, a switch to the backup node is not required:
   a. Perform any existing application shutdown procedures.
   b. Vary off the IASP.
   c. Suspend mirroring from the primary node.
   d. End the backup cluster node.
   e. End the cluster administrative domain, the CRG, and the primary cluster node.
   f. The primary node can now be powered down.
3. If the system is the primary node and this outage is planned, but production applications need to remain available, a switch to the backup node is required:
   a. Perform any existing application shutdown procedures.
   b. Vary off the IASP.
   c. Issue the Change Cluster Resource Group Primary (**CHGCRGPRI**) command. Upon completion, verify that mirroring is active and in sync.
   d. Suspend mirroring and end the current target node.
   e. The primary node (now the mirroring target) can be powered down now.

For more information about switching operations for a geographic mirroring environment, see Chapter 5, “Switching a geographic mirroring environment” on page 67.

4.2.5 Upgrading your operating system

One of the benefits of PowerHA is the ability to upgrade your operating system and to apply program temporary fixes (PTFs) or technology refreshes (TRs) with less impact to production than without PowerHA. Two scenarios are involved:

- PTFs and technology refreshes
- Operating system upgrades

Although both scenarios benefit from the use of PowerHA, the way that you manage the upgrades differs slightly.

PTFs and technology refreshes

With the PTFs and technology refreshes scenario, you can simply install the updates to the current nonproduction node. After you correctly install the updates, you can then schedule a switch at a convenient time so that you can use the new updates. If you experience a problem, you can simply switch back again until you resolve the problem.

After you are comfortable that the updates work well, you can perform the updates on the remaining node.

Important: If you use geographic mirroring and the updates that you want to apply require an IPL, suspend the replication before you start the process.

Operating system upgrades

The operating system upgrades scenario is slightly more complex because after an IASP is updated to a new operating system level, it cannot be used on an older-level partition. Therefore, it is necessary to use a different approach to the upgrade.

You can upgrade the operating system and install the associated PTFs on the nonproduction node without affecting production. Instead of switching to use the new code level, vary off the IASP and then issue a detach of the replication with the following Change Auxiliary Storage Pool Session (**CHGASPSSN**) command, where *sessionname* is the name of the replication session:

```
CHGASPSSN SSN(sessionname) OPTION(*DETACH)
```

This command makes the replication target usable. You can vary on the IASP on the newly upgraded node and test.

The overall process consists of the following steps:

1. Vary off the production IASP copy.
2. Issue the **CHGASPSSN** command with **OPTION(*DETACH)**.
3. Vary on the IASP copy on the upgraded system and test it.

   Restriction: If you use geographic mirroring, you must complete the operating system upgrade on the original production system before you attempt to restart the replication.

4. When you are ready to go into production, vary off the upgraded copy.
5. Issue an End Cluster Resource Group (**ENDCRG**) command.
6. Issue a Change Cluster Resource Group (**CHGCRG *CHGCUR**) command and change the primary and backup nodes.
7. Issue a Start Cluster Resource Group (**STRCRG**) command.
8. Issue a **CHGASPSSN** with **OPTION(*REATTACH)** command from the upgraded node to the old production node.
9. Vary on the IASP on the upgraded node and replication restarts. In certain conditions, you might need to resume the replication.

Important: After you upgrade the operating system and apply technology refreshes or PowerHA group PTFs, you must check that the actual cluster version and PowerHA version match the new potential cluster version and PowerHA version. Use the Change Cluster Version (**CHGCLUVER**) command, which can increment versions or modification levels if all of the nodes can operate at the new level. Failure to follow this process can result in an inability to start your cluster nodes and use PowerHA.

4.2.6 Adding disks to an IASP in a PowerHA environment

Important: You must follow a particular sequence when you add logical unit numbers (LUNs) to a PowerHA environment.

Although normally no data loss occurs during this sequence of steps, it is a good idea to have a current backup of the IASP data in place before you begin these steps.

Also, the IBM i system performs more efficiently with more LUNs of equivalent sizes. Consider this effect on disk performance when you add capacity to your IASP.

Important: No replication can be in place for the new LUNs when you start this procedure.

Run the following steps to add disks to an IASP in a PowerHA environment:

1. Use the Work with Cluster (**WRKCLU**) command to ensure that all nodes are active in the cluster.
2. Add the LUNs to the mirror copy (current backup system) IASP by using the following steps:
   a. Run **STRSST** and sign on.
   b. Select option 3 to work with disk units.
   c. Select option 2 to work with the disk configuration.
   d. Select option 2 to add units to ASPs.
   e. Select option 3 to add units to existing ASPs.
   f. Select the non-configured disk units to add by placing the ASP number on the specified line. If a LUN in this set of LUNs was used before by this system or another system, you receive a warning message. Press F10 to continue to use these LUNs.
   g. Press F10 to Add and Balance.
3. After the Add and Balance completes, perform a manual switchover that promotes the mirror copy (current backup system) IASP to the production copy (current production system) IASP.
4. Add the disk units to the *new mirror copy* (previously the production copy) IASP by using the following steps:
   a. Run **STRSST** and sign on.
   b. Select option 3 to work with disk units.
   c. Select option 2 to work with the disk configuration.
   d. Select option 2 to add units to ASPs.
   e. Select option 3 to add units to existing ASPs.
   f. Select the non-configured disk units to add by placing the ASP number on the specified line. If a LUN in this set of LUNs was used before by this system or another system, you receive a warning message. Press F10 to continue to use these LUNs.
   g. Press F10 to Add and Balance.
5. After the Add and Balance completes, perform a manual switchover that promotes the mirror copy IASP back to the original production copy IASP.

Now, both the current production system and the current backup system have additional capacity in the IASP and the geographic mirroring of the data is still occurring.

4.2.7 Removing disks from an IASP in a PowerHA environment

Be careful when you remove LUNs from a PowerHA environment because a particular sequence must be followed.

Run the following steps to remove LUNs from a PowerHA environment:

1. Use the **WRKCLU** command to ensure that all nodes are active in the cluster.
2. Remove the LUNs from the IASP on the production node only:
   a. Run **STRSST** and sign on.
   b. Select option 3 to work with disk units.
   c. Select option 2 to work with the disk configuration.
   d. Select option 12 to work with removing units from the configuration.
   e. Select option 3 to remove units from the configuration.
   f. Select the disk unit or disk units to remove by placing a 4 in front of the disk units to remove and press Enter.
   g. On the Confirm Remove Disk Units panel, take a moment to verify that the units that are shown are the correct units to remove, and then press Enter to confirm the choice.
   h. Now, the disk units are removed from the IASP.
   i. Exit SST.
3. Perform a manual switchover that promotes the mirror copy (current backup system) IASP to the production copy (current production system) IASP.
4. Remove the disk units from the *new mirror copy* (previously the production copy) IASP by using these steps:
   a. Run **STRSST** and sign on.
   b. Select option 3 to work with disk units.
   c. Select option 2 to work with the disk configuration.
   d. Select option 12 to work with removing units from the configuration.
   e. Select option 3 to remove units from the configuration.
   f. Select the disk unit or disk units to remove by placing a 4 in front of the disk units to remove and press Enter.
   g. On the Confirm Continuation panel, press Enter to continue.
   h. On the Confirm Remove Disk Units panel, press Enter to confirm the removal of the disk units.
   i. On the Function Status window, the completion percentage is shown. After the completion percentage reaches 100%, a confirmation message is posted that indicates that the selected units were removed successfully.
   j. Exit SST.

4.2.8 Server replacement

Depending on the PowerHA SystemMirror solution that you implemented, replacing one or two of the servers in your HA environment involves a number of steps. The steps are described in this section.

When you replace the backup system, including its disks, in an environment that uses geographic mirroring with internal disks, you cannot simply perform a save and restore operation.

Run the following steps:

1. To preserve your old backup system if the migration to the new backup system fails, perform a detach of the ASP session by using the following command:

   ```
   CHGASPSSN OPTION(*DETACH)
   ```

2. Power down the old backup system.
3. *Your production IASP needs to be varied off to perform this step*. Unconfigure geographic mirroring from the production system by using either the IBM System Navigator for i or the Configure Geographic Mirror (**CFGGEOMIR**) command with the ***DELETE** action. This action results in an error message that states that the backup system cannot be found. You can instruct the system to ignore this status message and to proceed with the unconfiguration.
4. Start the new backup system. The system must have non-configured drives that are available to become the new backup IASP. You must ensure that clustering works between the production system and the new backup system. It might be necessary to remove and add the backup cluster node to and from the cluster and recovery domain.
5. *Your production IASP needs to be varied off to perform this step.* Configure geographic mirroring from the production system to the new backup system by using either the GUI interfaces or the **CFGGEOMIR** command.
6. When the configuration of geographic mirroring is finished, vary on the production IASP and ensure that geographic mirroring is active. A full resynchronization is required.

Exchanging the production system without first unconfiguring geographic mirroring and reconfiguring it afterward is also not possible. Consider the use of the Change Cluster Resource Group Primary (**CHGCRGPRI**) command to switch over to the backup system and then follow the steps that are described above.

To replace only the *backup* server in a geographic mirroring environment that uses external storage, ensure that you suspend geographic mirroring from the production site first. Then, power down the old backup server, attach the new server to the existing external storage, and restart the new backup server. Finally, resume geographic mirroring to run a partial synchronization.

When you replace only the *production* server in a geographic mirroring environment that uses external storage, either switch to the backup system first or ensure that you end your production system correctly (by varying off the IASP, running **ENDCRG**, and running **ENDCLUNOD**, before you run the **PWRDWNSYS** command) before you exchange the server hardware.

4.2.9 Unconfiguring geographic mirroring

In a certain situation, IBM Support might direct you to unconfigure geographic mirroring to perform a specific activity, or geographic mirroring might not be needed in your environment any longer. This section documents the process for unconfiguring your geographic mirroring environment.

Run the following steps:

1. Run the Display Cluster Information (**DSPCLUINF**) command for the cluster from each node and press Enter after the initial DSPCLUINF panel is displayed. This action provides the status of all cluster nodes from each cluster node. All cluster nodes must be active before you continue.
2. Vary off the IASP by using the **VRYCFG** command.
3. End the CRG by using the **ENDCRG** command.
4. Unconfigure geographic mirroring by using the Configure Geographic Mirroring (**CFGGEOMIR**) command with the ***DELETE** action:

   ```
   CFGGEOMIR ASPDEV(<iasp_device_name>) ACTION(*DELETE)
   ```

   Press F16 to confirm the command.
5. The message in Figure 4-3 is posted to your session if the command was successful.

```
Configure Geographic Mirror *DELETE request completed.
```

Figure 4-3 Unconfigure geographic mirroring success message

Unconfiguring the geographic mirroring process does not delete the IASP or the CRG object from either node. The production copy system's copy description also remains. If these objects are no longer needed, they can be manually deleted from the systems.

5

Switching a geographic mirroring environment

This chapter describes the geographic mirroring synchronization process and shows how to perform a planned switch of a geographic mirroring environment between production and backup nodes. In addition, the procedures to recover from an unplanned failover are covered.

This chapter describes the following topics:

5.1 Synchronization

When geographic mirroring is resumed after a suspend or detach, the mirror copy is resynchronized with the production copy. The production copy can function normally during synchronization, but performance might be affected negatively.

During synchronization, the contents of the mirror copy are unusable, and it cannot become the production copy. If the independent disk pool is made unavailable during the synchronization process, synchronization resumes where it left off when the independent auxiliary storage pool (IASP) is made available again.

The message CPI095D "`Cross-site Mirroring (XSM) synchronization for IASP`" is sent to the QSYSOPR message queue every 15 minutes to indicate the progress of the synchronization.

Two types of synchronization exist:

- Full synchronization

  Full synchronization indicates that a complete synchronization takes place. Changes to the production copy were not tracked to apply to the synchronization, or the data status of the mirror copy was not determined. First, a full synchronization deletes all data in the backup IASP. Then, the full synchronization copies the current data from the production IASP to the backup IASP.

- Partial synchronization

  Partial synchronization indicates that changes to the production copy and mirror copy were tracked while geographic mirroring was suspended or detached. This tracking can shorten the synchronization time considerably because a complete synchronization is unnecessary.

  In this case, when the mirror copy is reattached and geographic mirroring is resumed, only tracked changes need to be synchronized. Changes that are made on the production copy (since the suspend or detach was performed) are sent to the mirror copy, and any changes that were made on the mirror copy are overwritten with the original data from the production copy of the IASP.

  Important: Any changes that are made on the mirror copy while it is detached are undone, and any tracked changes from the production copy are applied.

  Message CPI095D indicates the type of synchronization.

Figure 5-1 shows the message details.

```
                        Additional Message Information

Message ID . . . . . :   CPI095D       Severity . . . . . :   80
Message type . . . . . :   Information
Date sent  . . . . . :   09/25/15      Time sent  . . . . . :   09:32:52

Message . . . . :   Cross-site Mirroring (XSM) synchronization for IASP 33 is
  0% complete.
Cause . . . . . :   Mirror copy Independent Auxiliary Storage Pool (IASP) 33
  on the target system with clustering node ID ITSO1NOD is being synchronized
  with the production copy IASP 33 on the source system with clustering node
  ID ITSO2NOD.
  The synchronization process is 0 percent complete.
    If the percent complete is 0, synchronization has started recently.
    If the percent complete is 100, synchronization has completed.
    If the percent complete is less than 100, then synchronization is still
  active. The data on the mirror copy is not usable while synchronization is
  active and the mirror copy is not available for switchover or failover.

  The synchronization is of type 1. The synchronization types and their
    meanings are as follows:
    1 - The synchronization being performed is a synchronization of tracked
  changes.
    2 - The synchronization being performed is a synchronization of all data.

                                                                        Bottom
Press Enter to continue.

 F3=Exit   F6=Print   F9=Display message details   F12=Cancel
F21=Select assistance level
```

Figure 5-1 Message CPI095D that shows the synchronization status

5.2 Planned switch

Normally, a *planned switch* is performed when you want to perform a maintenance action on your current primary system or as part of a regular role swap. This section guides you through a planned switch operation by using the PowerHA graphical user interface (GUI) or by using Work with Cluster (**WRKCLU**) command menu options.

Important: A successful switch depends on the PowerHA switch process. Also, a successful switch depends on business applications that can run on the backup node and users that can access those applications.

The following topics are described in this section:

- 5.2.1, “Performing a planned switch by using the PowerHA GUI” on page 70
- 5.2.2, “Performing the switch operation by using control language commands” on page 75
- 5.2.3, “Additional considerations when you use Storwize or DS8000 and FlashCopy” on page 79

5.2.1 Performing a planned switch by using the PowerHA GUI

You must end all applications or jobs that use the IASP on the current production node before you perform a planned switch. Otherwise, these jobs are ended abnormally.

Any jobs that use the IASP can be shown by running the Work with ASP Jobs (**WRKASPJOB**) command as shown in Figure 5-2.

Note: No option is available to show IASP jobs by using the PowerHA GUI.

```
                              Work with ASP Jobs                      ITSOHA1
                                                              09/12/15  10:07:35
Type options, press Enter.
  4=End job   5=Work with   7=Send message

Opt     ASP/Job          User           Number     Type        Status
 _      IASPHA
 _        DSP01          PWRHAUSR       004656     INT          RUN


                                                                        Bottom
Parameters or command
===>
F3=Exit F4=Prompt F5=Refresh F9=Retrieve F12=Cancel
```

Figure 5-2 WRKASPJOB command that shows an active job that uses the IASP

Run the following steps to perform a planned switch by using the PowerHA GUI:

1. From a web browser, open the Navigator for IBM i GUI by entering the following URL and logging in. This step can be performed on either cluster node:

   `http://<system ip address>:2001`

2. On the Navigator for IBM i Welcome window, select **PowerHA**. Figure 5-3 shows the main status display of the PowerHA GUI.

   All entries must show a green check mark for their status before you proceed. If your display does not show green check marks for all entries, you need to investigate and resolve any issues first. Chapter 6, “Troubleshooting and collecting data for problem determination” on page 87 can help assist you with this investigation.

   Select **Cluster Resource Groups**.

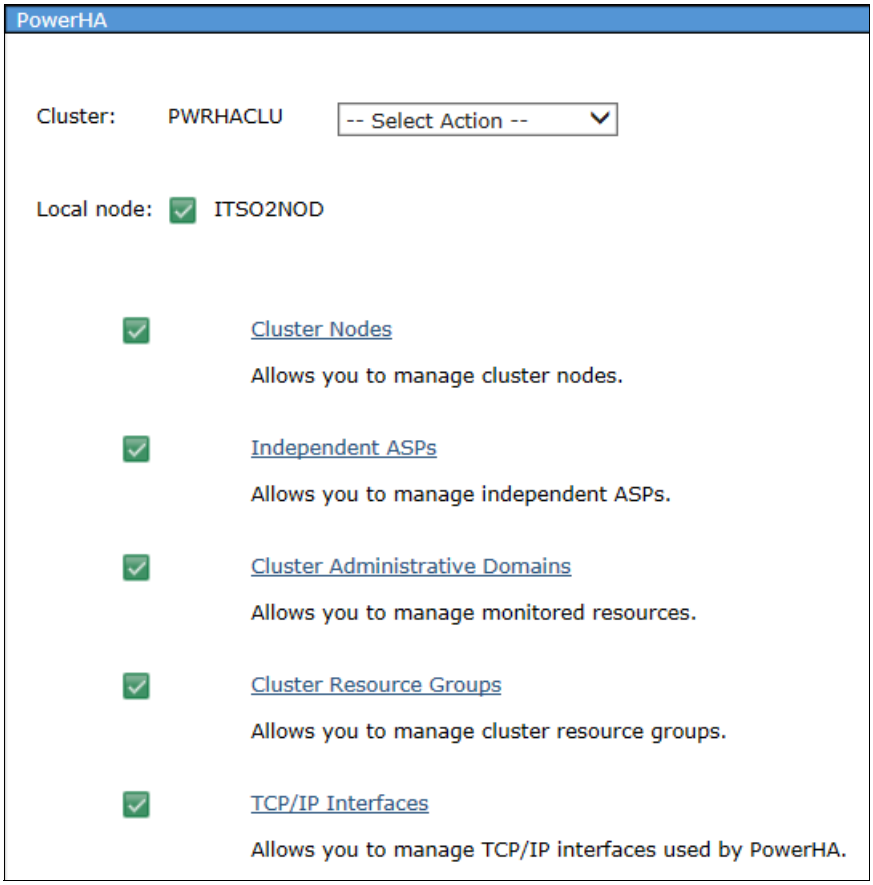

Figure 5-3 Main PowerHA status display

3. To perform a switch, select **Switch** from the menu (drop-down menu) of the cluster resource group (CRG) as shown in Figure 5-4.

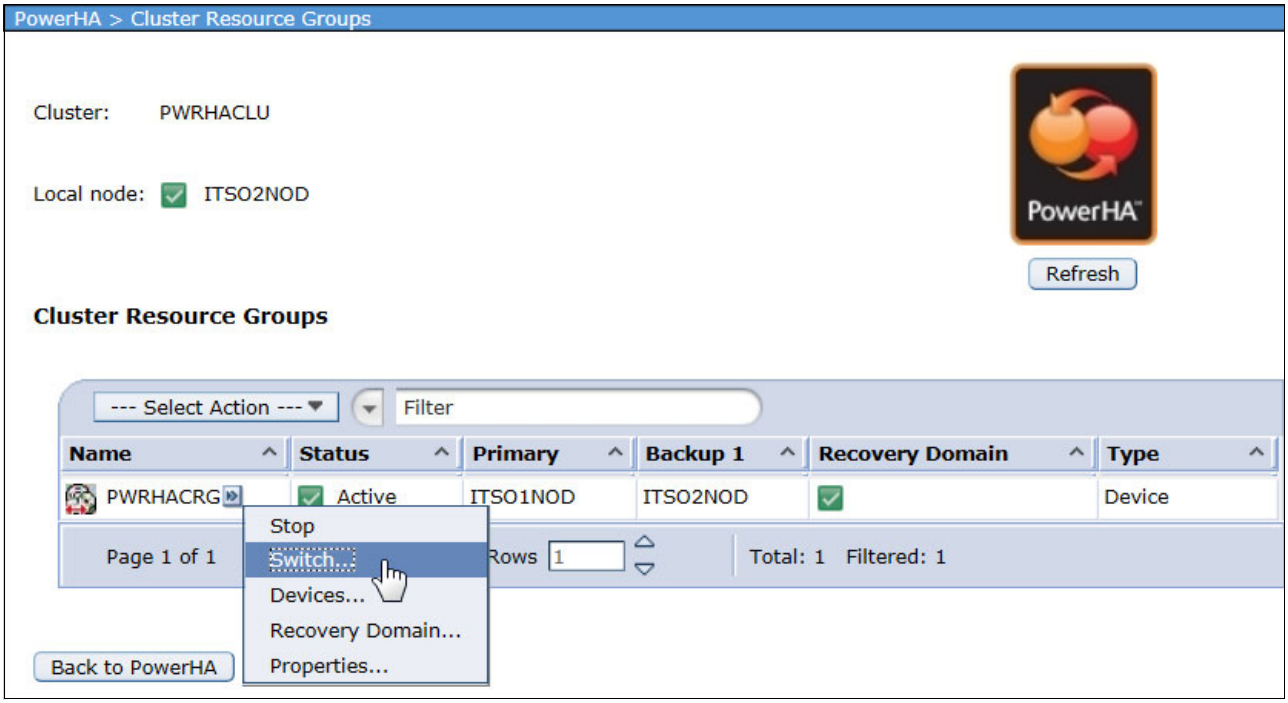

Figure 5-4 CluSter Resource Group GUI display that shows the Switch option

4. As shown in Figure 5-5, PowerHA shows you a preview of how the nodes and their roles in the CRG look before and after a switch.

   Verify that the roles of the nodes are as expected and click **OK** to proceed with the switch.

   The switchover process performs the following functions:

   a. The switchover process varies off the IASP on the current primary cluster node if the IASP is not varied off.

   b. If a server takeover Internet Protocol (IP) address is defined in the CRG, this IP interface is ended on the current primary cluster node.

   c. The recovery domain information in the CRG is updated to show the new current primary and backup nodes.

   d. On the new primary node, the IASP is varied on if this action is configured in the CRG device entry.

   e. On the new primary node, the server takeover IP address is started if a server takeover IP address is defined in the CRG device entry.

   f. Reverse replication is started.

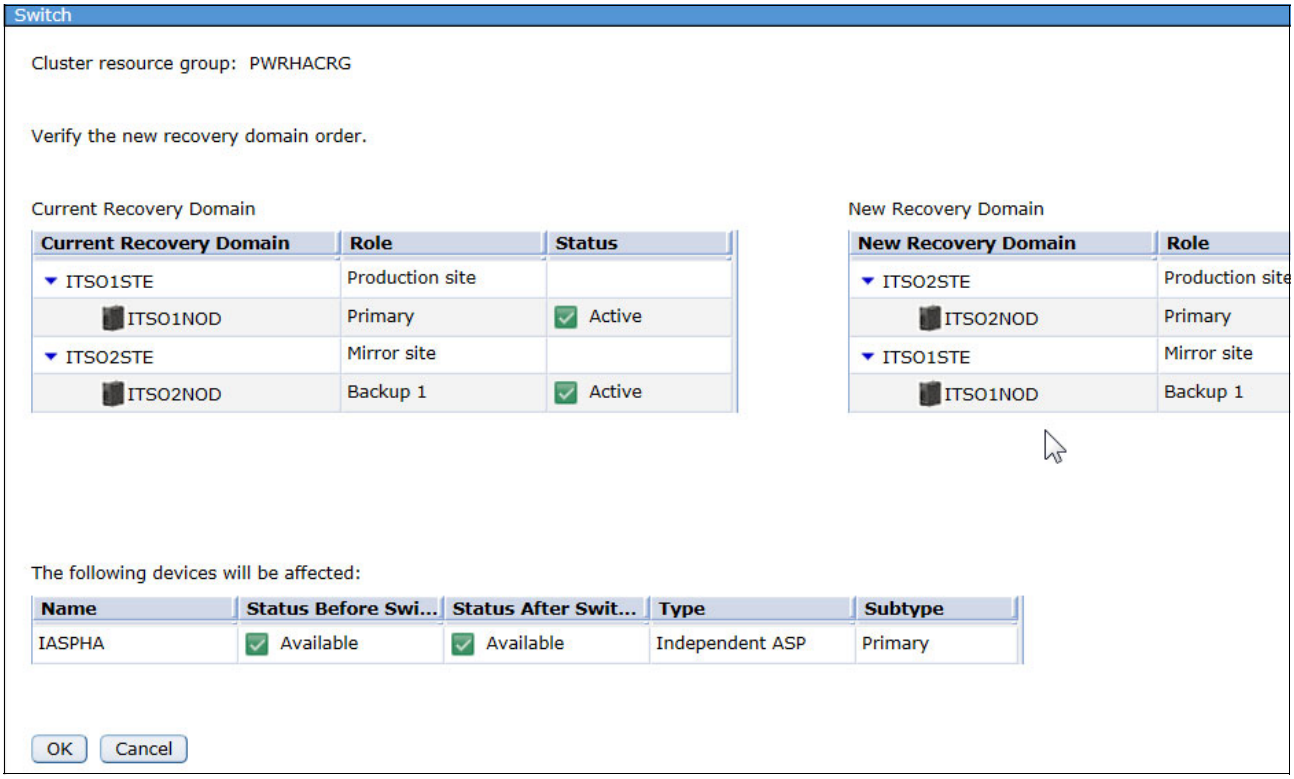

Figure 5-5 Recovery domain switch preview before and after a switch

5. As shown in Figure 5-6, the status is displayed and updated automatically as the switch progresses.

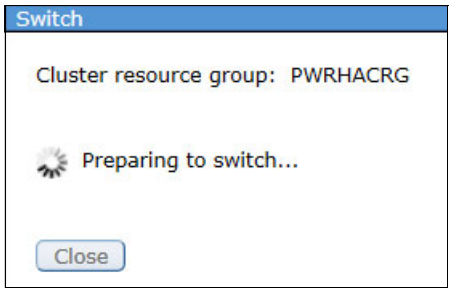

Figure 5-6 Switch progress status

6. When the switch completes, the status window is updated as shown in Figure 5-7. Click **Close**.

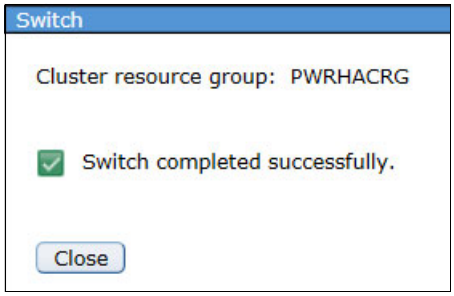

Figure 5-7 Switch completion status

7. On the Cluster Resource Groups display, click **Refresh** to view the updated role status for all of the nodes in the recovery domain as shown in Figure 5-8.

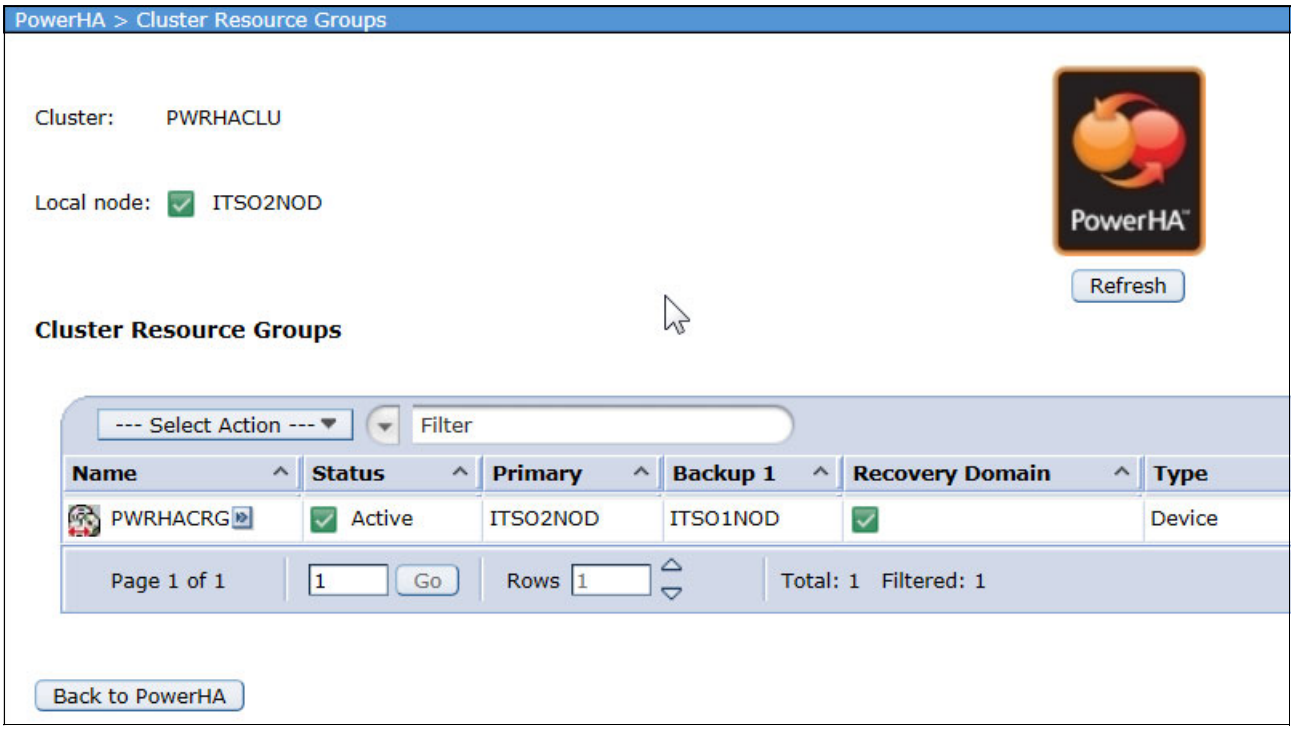

Figure 5-8 Switch verification window

5.2.2 Performing the switch operation by using control language commands

You must end all applications or jobs that use the IASP on the source node before a planned switch is performed. Otherwise, these jobs are ended abnormally.

Note: Ensure that the interactive job that invokes the switch operation is not using the IASP. The interactive job is ended; however, the switch will continue to completion.

Any jobs that use the IASP can be shown by running the Work with ASP Jobs (**WRKASPJOB**) command as shown in Figure 5-9.

```
                              Work with ASP Jobs                      ITSOHA1
                                                                 09/12/15  10:07:35
Type options, press Enter.
  4=End job   5=Work with   7=Send message

Opt     ASP/Job          User           Number     Type        Status
 _      IASPHA
 _        DSP01          QSECOFR        004656     INT          RUN

                                                                          Bottom
Parameters or command
===>
F3=Exit F4=Prompt F5=Refresh F9=Retrieve F12=Cancel
```

Figure 5-9 WRKASPJOB command that shows an active job that uses the IASP

Run the following steps to perform a planned switch by using IBM i control language (CL) commands:

1. From the Work with Cluster (**WRKCLU**) menu that is shown in Figure 5-10, select option 9 to work with cluster resource groups.

```
                            Work with Cluster
                                                          System:   ITSOHA1
 Cluster  . . . . . . . . . . . . . . . . :   PWRHACLU

 Select one of the following:

      1. Display cluster information
      2. Display cluster configuration information

      6. Work with cluster nodes
      7. Work with device domains
      8. Work with administrative domains
      9. Work with cluster resource groups
     10. Work with ASP copy descriptions

     20. Dump cluster trace

 Selection or command
 ===> _
F1=Help F3=Exit F4=Prompt F9=Retrieve F12=Cancel
```

Figure 5-10 Work with Cluster menu

2. On the Work with Cluster Resource Groups display (Figure 5-11), specify option 3 (Change primary) to change the cluster resource group (CRG) primary node.

```
                      Work with Cluster Resource Groups

 Consistent information in cluster  . . . :   Yes

 Type options, press Enter.
   1=Create   2=Change   3=Change primary          4=Delete   5=Display
   6=Recovery domain     7=Configuration objects   8=Start    9=End
   20=Dump trace

            Cluster                                                Primary
 Opt     Resource Group     Type        Status                     Node

  _        PWRHACRG         *DEV        Active                     ITSO1NOD


                                                                        Bottom
 Parameters for options 1, 2, 3, 8, 9 and 20 or command
 ===>
 F1=Help   F3=Exit   F4=Prompt   F5=Refresh   F9=Retrieve   F12=Cancel
F13=Work with cluster menu
```

Figure 5-11 Work with Cluster Resource Groups display

3. On the Change CRG Primary (**CHGCRGPRI**) display, press Enter to continue (Figure 5-12).

```
                      Change CRG Primary (CHGCRGPRI)

Type choices, press Enter.

Cluster  . . . . . . . . . . . . > PWRHACLU      Name
Cluster resource group . . . . . > PWRHACRG      Name
Exit program data  . . . . . . .   *SAME


                                                                      Bottom
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
```

Figure 5-12 Change CRG Primary display

4. The display shows a status of inhibited while the switch takes place. After the switch is finished, the display in Figure 5-13 is shown.

Note: The primary node is changed.

```
                    Work with Cluster Resource Groups

Consistent information in cluster  . . . :   Yes

Type options, press Enter.
  1=Create   2=Change   3=Change primary          4=Delete   5=Display
  6=Recovery domain     7=Configuration objects   8=Start    9=End
  20=Dump trace

           Cluster                                                Primary
Opt     Resource Group     Type        Status                    Node

 _        PWRHACRG         *DEV        Active                    ITSO2NOD


                                                                      Bottom
Parameters for options 1, 2, 3, 8, 9 and 20 or command
===>
F1=Help   F3=Exit   F4=Prompt   F5=Refresh   F9=Retrieve   F12=Cancel
F13=Work with cluster menu
```

Figure 5-13 Work with Cluster Resource Groups display after a switch completes

5. You can also run the Display ASP Session **DSPASPSSN SSN(PWRHASSN)** command to confirm that the nodes and ASP copies were reversed as shown in Figure 5-14.

```
                            Display ASP Session                     ITSOHA1
                                                          09/10/15  14:03:04
 Session . . . . . . . . . . . . . . . . . . :   PWRHASSN
 Type  . . . . . . . . . . . . . . . . . . . :   *GEOMIR

 Source node . . . . . . . . . . . . . . . . :   ITSO2NOD
 Target node . . . . . . . . . . . . . . . . :   ITSO1NOD
 Transmission Delivery . . . . . . . . . . . :   *SYNC
 Mirroring Mode  . . . . . . . . . . . . . . :   *SYNC
 Suspend timeout . . . . . . . . . . . . . . :   120
 Synchronization priority  . . . . . . . . . :   *HIGH

                             Copy Descriptions

ASP         ASP                                        Data
Device      Copy              Role       State       State
IASPHA      ITSO2CPY    PRODUCTION   AVAILABLE      USABLE
            ITSO1CPY        MIRROR      ACTIVE      USABLE


                                                                      Bottom
 Press Enter to continue
F3=Exit F5=Refresh F11=View2 F12=Cancel F19=Automaticrefresh
```

Figure 5-14 DSPASPSSN command shows the role reversal after a switch

6. Pressing F11 (View2) shows the secondary window as shown in Figure 5-15.

```
                            Display ASP Session                     ITSOHA1
                                                          09/19/15  20:33:12
 Session . . . . . . . . . . . . . . . . . . :   PWRHASSN
 Type  . . . . . . . . . . . . . . . . . . . :   *GEOMIR

 Source node . . . . . . . . . . . . . . . . :   ITSO2NOD
 Target node . . . . . . . . . . . . . . . . :   ITSO1NOD
 Transmission Delivery . . . . . . . . . . . :   *SYNC
                                                                     More...
                             Copy Descriptions

 ASP         ASP                     Total data  Synchronization
 Device      Copy             State  out of sync     progress
 IASPHA      ITSO2CPY        ACTIVE            0        0%
             ITSO1CPY        ACTIVE            0        0%




                                                                      Bottom
 Press Enter to continue

 F3=Exit   F5=Refresh   F11=View 1   F12=Cancel   F19=Automatic refresh
```

Figure 5-15 Secondary window for the DSPASPSSN command

5.2.3 Additional considerations when you use Storwize or DS8000 and FlashCopy

Storwize and DS8000 external storage can be configured to present logical unit numbers (LUNs) for an IASP in a geographic mirroring environment.

The procedure to perform a planned switch is the same as the procedure that is described in this section. However, if FlashCopy is used, additional considerations about the FlashCopy procedure might apply after a switch is performed. For more information, see *IBM PowerHA SystemMirror for i: Using DS8000 (Volume 2 of 4)*, SG24-8403, and *IBM PowerHA SystemMirror for i: Using IBM Storwize (Volume 3 of 4)*, SG24-8402.

5.3 Unplanned site failover

A *failover* occurs when the source node fails and the backup node takes over. The default failover procedures depend on the cluster failover wait time and failover default action settings.

Important: As with a planned switch, the success of a failover operation also depends on your previous testing and verification to show that the business applications can run on the backup node and users can access those applications.

We use the following failure scenarios to explain failover:

- A primary node failure triggers an automatic failover. This scenario can occur either by a panic message that is sent to the backup node or cluster monitors in place.
- A primary node or cluster communications failure results in node Partition status.

Each scenario requires different failover and recovery actions, which we describe in the following sections:

- 5.3.1, “Primary node failure that triggers an automatic failover event” on page 79
- 5.3.2, “Sudden primary node or cluster communications failure” on page 80

5.3.1 Primary node failure that triggers an automatic failover event

An unplanned automatic failover event can be triggered by a panic message that is sent by the primary node due to ending a cluster node, ending Transmission Control Protocol (TCP), ending all subsystems, or powering down the partition or system. The failover can also be triggered by a power state change event that is sent by the Hardware Management Console (HMC) server for a partition failure.

For an automatic failover event, a CPABB02 “`Cluster resource groups are failing over to node` *backup-node*`.  (C G)`” inquiry message is sent to the cluster or CRG message queue on the backup node if a failover message queue is defined for either the cluster or the CRG. If no failover message queue is defined, the failover starts immediately without posting any message.

The cluster parameters or Failover Wait Time (FLVWAITTIM) and Failover Default Action (FLVDFTACT) determine the next actions. With the default settings of FLVWAITTIM=`*NOWAIT` and FLVDFTACT=`*PROCEED`, an automatic failover begins immediately.

Setting Failover Wait Time parameter as a duration in minutes or *NOMAX allows a user on the backup node to respond to the CPABB02 inquiry message to either proceed with the failover or cancel the failover. The Failover Default Action parameter determines whether PowerHA proceeds with the failover or cancels the failover processing after the specified failover wait time expires and no response to the inquiry message was entered.

Note: Regardless of the cluster parameter settings, the primary IASP is taken offline by PowerHA for a failover event.

5.3.2 Sudden primary node or cluster communications failure

In the absence of advanced node failure detection, a sudden primary node failure or simply a cluster communications failure can result in a condition that is known as a *cluster partition*. The cluster partition condition means that the backup node cannot determine the status of the production node reliably. In this situation, an automatic failover does not occur, regardless of the cluster failover settings. In fact, a failover is not possible unless you perform additional steps.

First, determine whether the primary node or source node is still in operation. If the production workload can continue, you do not need to fail over to the backup node. For a cluster communications failure only, message ID CPDB715 is issued on the primary node. Mirroring continues and after communications are restored, the cluster software reconnects the cluster nodes automatically.

Otherwise, if the primary node is no longer responsive or available, the status of that node must be changed from Partition to Failed. A status of Partition or Failed allows a failover to proceed to the backup node by running the following steps:

1. Log on to the backup system and verify the node status. If the primary node is in a Partition status, a manual switch or failover cannot be performed. Figure 5-16 shows the ITSO1NOD (primary/source) node in a Partition status by using Work with Cluster Nodes (**WRKCLU**) menu option 6.

```
                          Work with Cluster Nodes

Local node . . . . . . . . . . . . . . . :   ITSO2NOD
Consistent information in cluster  . . . :   Yes

Type options, press Enter.
  1=Add   2=Change   4=Remove   5=Display more details   6=Work with monitors
  8=Start   9=End    20=Dump trace

Opt     Node         Status         Device Domain

 _      ITSO1NOD     Partition       PWRHADMN
 _      ITSO2NOD     Active          PWRHADMN

                                                                        Bottom
Parameters for options 1, 2, 9 and 20 or command
===>
F1=Help   F3=Exit     F4=Prompt    F5=Refresh   F9=Retrieve
F11=Order by status   F12=Cancel   F13=Work with cluster menu
```

Figure 5-16 Work with Cluster Nodes display that shows a Partition condition

2. The status of the primary node must be changed to Failed by using the Change Cluster Node Entry (**CHGCLUNODE**) command as shown in Figure 5-17. This command also varies off the IASP on the target node.

```
                    Change Cluster Node Entry (CHGCLUNODE)

Type choices, press Enter.

Cluster  . . . . . . . . . . . . > PWRHACLU      Name
Node identifier  . . . . . . . . > ITSO1NOD      Name
Option . . . . . . . . . . . . . > *CHGSTS       *ADDIFC, *RMVIFC, *CHGIFC...



                                                                        Bottom
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
```

Figure 5-17 Change Cluster Node Entry command

3. Figure 5-18 shows the node status changed to Failed. The **CHGCLUNODE** command also triggers a cluster failover without a failover inquiry message, but it still requires the user to vary on the IASP on the new primary/source node and start the takeover IP interface.

```
                          Work with Cluster Nodes

 Local node . . . . . . . . . . . . . . . . :   ITSO2NOD
 Consistent information in cluster  . . . :   Yes

 Type options, press Enter.
   1=Add   2=Change   4=Remove   5=Display more details   6=Work with monitors
   8=Start   9=End    20=Dump trace

 Opt     Node         Status         Device Domain

  _      ITSO1NOD     Failed          PWRHADMN
  _      ITSO2NOD     Active          PWRHADMN


                                                                          Bottom
 Parameters for options 1, 2, 9 and 20 or command
 ===>
 F1=Help   F3=Exit     F4=Prompt    F5=Refresh   F9=Retrieve
F11=Order by status F12=Cancel F13=Work with cluster menu
```

Figure 5-18 Work with Cluster Nodes display that shows a failed node

4. Vary on the IASP on the backup (now source) node. The status changes to AVAILABLE as shown in Figure 5-19. Also, start the takeover address, if required.

```
                       Work with Configuration Status                ITSOHA2
                                                          09/14/15  16:06:57
 Position to  . . . . .               Starting characters

 Type options, press Enter.
   1=Vary on   2=Vary off   5=Work with job   8=Work with description
   9=Display mode status    13=Work with APPN status...

 Opt  Description       Status               -------------Job--------------
  _   IASPHA            AVAILABLE









                                                                          Bottom
 Parameters or command
 ===>
 F3=Exit   F4=Prompt   F12=Cancel   F23=More options   F24=More keys
Vary on completed for device IASPHA.
```

Figure 5-19 IASP configuration status on the new source node

5. Run the following Display ASP Session (**DSPASPSSN**) command to verify that the node ASP copy shows as AVAILABLE. See Figure 5-20. The backup node now has the production copy.

```
DSPASPSSN SSN(PWRHASSN)
```

Important: The terms can be confusing here. After a switch/failover, the “backup” node (which refers to a physical system/partition) will be the “source” node and contain the “production” copy of the IASP when it is shown in the DSPASPSSN display.

```
                              Display ASP Session                    ITSOHA2
                                                          09/15/15  09:54:34
 Session . . . . . . . . . . . . . . . . . . . :   PWRHASSN
 Type  . . . . . . . . . . . . . . . . . . . . :   *GEOMIR
 Source node . . . . . . . . . . . . . . . . . :   ITSO2NOD
 Target node . . . . . . . . . . . . . . . . . :   ITSO1NOD
 Transmission Delivery . . . . . . . . . . . . :   *SYNC

                               Copy Descriptions
ASP         ASP                                          Data
 Device      Copy              Role       State        State
 IASPHA      ITSO2CPY    PRODUCTION   AVAILABLE      USABLE
             ITSO1CPY        MIRROR   SUSPENDED      USABLE


                                                                      Bottom
 Press Enter to continue
F3=Exit F5=Refresh F11=View2 F12=Cancel F19=Automaticrefresh
```

Figure 5-20 Display ASP Session on the backup node

6. After the original primary system is repaired and started, ensure that the IASP on that node is in a varied-off status and that the cluster node is Inactive. Start the repaired node from the current production node and verify that all nodes show Active as shown in Figure 5-21.

```
                          Work with Cluster Nodes

 Local node . . . . . . . . . . . . . . . . :   ITSO2NOD
 Consistent information in cluster  . . . :   Yes

 Type options, press Enter.
   1=Add   2=Change   4=Remove   5=Display more details   6=Work with monitors
   8=Start   9=End    20=Dump trace

 Opt     Node         Status         Device Domain

  _      ITSO1NOD     Active          PWRHADMN
  _      ITSO2NOD     Active          PWRHADMN


Parameters for options 1, 2, 9 and 20 or command
 ===>
 F1=Help   F3=Exit     F4=Prompt    F5=Refresh   F9=Retrieve
F11=Order by status F12=Cancel F13=Work with cluster menu
```

Figure 5-21 Work with Cluster Nodes display that shows that both nodes are active

7. Start the CRG if it is not started by using the following command:

```
STRCRG CUSTER(PWRHACLU) CRG(PWRHACRG)
```

8. Use the Work with Cluster (**WRKCLU**) menu that is shown in Figure 5-10 on page 76, and select option 10 to work with ASP copy descriptions. On Figure 5-22, enter option 22 to change session on either ASP's Opt line and press F4.

```
                      Work with ASP Copy Descriptions               ITSOHA2
                                                                 09/15/15  11:46:39
 Device domain  . . . . . . . . . . . . . :   PWRHADMN

 Type options, press Enter.
   1=Add copy         2=Change copy       4=Remove copy    5=Display copy
   21=Start session   22=Change session   24=End session   25=Display session

         ASP            ASP            ASP               Session
 Opt     Device         Copy           Session             Type

  _      IASPHA         ITSO1CPY       PWRHASSN          *GEOMIR
  22     IASPHA         ITSO2CPY       PWRHASSN          *GEOMIR



                                                                          Bottom
 Parameters or command
 ===>
F3=Exit F4=Prompt F5=Refresh F9=Retrieve F12=Cancel
```

Figure 5-22 Work with ASP Copy Descriptions display

9. On the Change ASP Session display, enter the *RESUME option as shown in Figure 5-23.

```
                      Change ASP Session (CHGASPSSN)

Type choices, press Enter.

Session  . . . . . . . . . . . . > PWRHASSN      Name
Option . . . . . . . . . . . . . > *RESUME       *CHGATTR, *SUSPEND...
Device domain  . . . . . . . . .   *             Name, *
ASP device . . . . . . . . . . .   IASPHA        Name, *ALL
               + for more values


                                                                      Bottom
F3=Exit   F4=Prompt   F5=Refresh   F10=Additional parameters   F12=Cancel
F13=How to use this display        F24=More keys
```

Figure 5-23 Change ASP session command to resume mirroring

10. Refresh the display to show that the status changed to RESUMING as shown in Figure 5-24.

    Pressing F11 (View2) displays the progress and the amount of data that is out of sync. After an unplanned failover, a full resynchronization of the IASP likely will be required. The full resynchronization can take a long time.

    Important: The backup copy is unusable until the resume process is finished.

```
                           Display ASP Session                    ITSOHA2
                                                          09/15/15  11:50:37
Session . . . . . . . . . . . . . . . . . . :   PWRHASSN
Type  . . . . . . . . . . . . . . . . . . . :   *GEOMIR

Source node . . . . . . . . . . . . . . . . :   ITSO2NOD
Target node . . . . . . . . . . . . . . . . :   ITSO1NOD
Transmission Delivery . . . . . . . . . . . :   *SYNC
                                                                     More...
                             Copy Descriptions

ASP         ASP                                        Data
Device      Copy              Role       State       State
IASPHA      ITSO2CPY    PRODUCTION   AVAILABLE      USABLE
            ITSO1CPY        MIRROR    RESUMING    UNUSABLE




                                                                      Bottom
Press Enter to continue
F3=Exit F5=Refresh F11=View2 F12=Cancel F19=Automatic refresh
```

Figure 5-24 ASP copy in resuming status

The resume operation can also be initiated from the PowerHA GUI under the Independent ASP Details section as shown in Figure 5-25.

Note: It is considered a preferred practice to run the Reclaim Storage (**RCLSTG**) command on SYSBAS on a failed production node after you perform a failover and before you resume mirroring. Also, you need to schedule an **RCLSTG** of the IASP at the earliest convenience, preferably on the backup node before you switch back to “preferred production”.

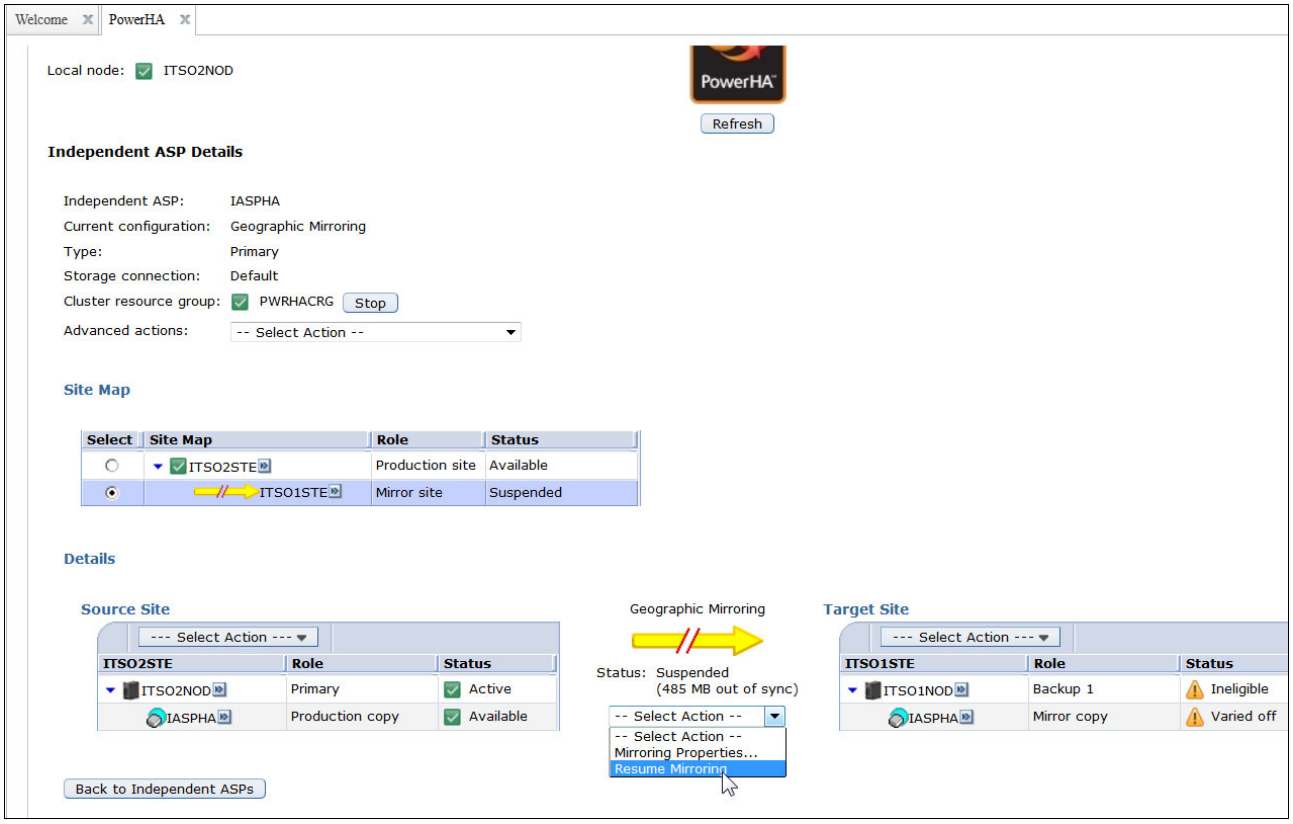

Figure 5-25 GUI display that shows Resume Mirroring

6

Troubleshooting and collecting data for problem determination

This chapter provides troubleshooting methods and the requirements for the documentation that you need to gather. Also, it provides support to help you report a problem to IBM. This chapter describes the following topics:

For a complete description of IBM PowerHA SystemMirror for i troubleshooting procedures, see Chapter 6, “Troubleshooting and collecting data for problem determination” in the IBM Redbooks publication, *IBM PowerHA SystemMirror for i: Preparation (Volume 1 of 4)*, SG24-8400.

6.1 Reducing the likelihood of errors

The best approach is to avoid PowerHA problems, if possible. Although avoiding all problems and errors is not plausible, specific actions can help you to minimize the number of issues that you experience when you work with PowerHA.

6.1.1 Verifying the PTF level for the system and PowerHA SystemMirror for i

As with any product, it is important that you establish a fix strategy to help you minimize the possibility of encountering code defects that were already discovered. Ensure that the System Licensed Internal Code (SLIC) and operating system code are up-to-date with the PowerHA SystemMirror for i product.

For more information about the group program temporary fix (PTF) level and recommended fixes for PowerHA, see chapter 6, “Troubleshooting and collecting data for problem determination” in the IBM Redbooks publication, *IBM PowerHA SystemMirror for i: Preparation (Volume 1 of 4)*, SG24-8400.

6.1.2 Configuring and managing the PowerHA environment

This book, and the other PowerHA books in this volume series, provide steps that were tested and used when we configured and managed PowerHA on IBM i systems. Chapter 3, “Implementing geographic mirroring” on page 27 provides information about implementing the PowerHA environment in a geographic mirroring environment. Chapter 4, “Monitoring and managing IBM PowerHA SystemMirror for i” on page 39 provides information about managing the PowerHA environment.

The following links point to other IBM documentation that can help you configure and manage your PowerHA environment:

- IBM i 7.2 Knowledge Center:

  http://www.ibm.com/support/knowledgecenter/ssw_ibm_i_72/rzahg/ic-homepage.htm
- IBM PowerHA SystemMirror for i DeveloperWorks website:

  https://ibm.biz/Bd4ub8
- IBM Support portal:

  https://ibm.biz/BdXqvs

6.2 Common tools and commands used for troubleshooting

This section briefly describes the following tools and commands that can be used for troubleshooting a geographic mirroring environment:

- 6.2.1, “MustGather Data Collector Tool (QMGTOOLS)” on page 89
- 6.2.2, “GEOSTAT” on page 89
- 6.2.3, “ASMINFO” on page 91
- 6.2.4, “DSMINFO” on page 92

6.2.1 MustGather Data Collector Tool (QMGTOOLS)

The MustGather Data Collector Tool (QMGTOOLS) is a suite of debug tools to gather data for various issues. This tool is the primary tool to debug geographic mirroring issues. We advise that anyone who sets up PowerHA on IBM i needs to familiar with the data that is collected by this tool. This knowledge will help you to minimize the time that is spent resolving a PowerHA issue.

The installation and data collection capabilities of QMGTOOLS are described in detail in Chapter 6, “Troubleshooting and collecting data for problem determination” in the IBM Redbooks publication, *IBM PowerHA SystemMirror for i: Preparation (Volume 1 of 4)*, SG24-8400.

6.2.2 GEOSTAT

You use the **GEOSTAT** tool to determine the current state of geographic mirroring.

The following steps show you how to dump out this information for the system:

1. From the command line, run Start System Service Tools (**STRSST**) and press Enter.
2. Sign in with a service tool profile and password that has authority to Display/Alter/Dump in SST.
3. Select option `1` - Start a service tool, and press Enter.
4. Select option `4` - Display/Alter/Dump, and press Enter.
5. To dump to a spooled file, select option `2` - Dump to printer, and press Enter.
6. Select option `2` - Licensed Internal Code (LIC) data, and press Enter.
7. Select option `14` - Advanced analysis, and press Enter.
8. On the Select Advanced Analysis Command panel, type option `1` (Select) next to the top blank line under the Command column. In the blank line, type `GEOSTAT`, and press Enter.
9. In the Options field, type `-ALL`, and press Enter.
10. Exit out of SST and type the Work with Job (**WRKJOB**) command and press Enter. From the Work with Job display, select option `4` to Work with spooled files for that job. The spool file that is named QSYSPRT contains the **GEOSTAT** data.

Figure 6-1 shows a display of the QSYSPRT spooled file.

```
*...+....1....+....2....+....3....+....4....+....5....+....6....+....7.
     DISPLAY/ALTER/DUMP
RUNNING MACRO: GEOSTAT                          -ALL
PROCESSING CLUSTER NODE: RCHESP1
CRG ESPCRG     IS ACTIVE -- PRIMARY NODE
PROCESSING IASP: 33 ESPIASP    AVAILABLE  PRIMARY ASP
IASP INFORMATION:
     PHYSICAL COPY ID:                  0XE7
     COPY ROLE:                         0XD7  PRODUCTION
     REMOTEMIRRORERRORRECOVERYPOLICY:      2  SUSPEND
     REMOTEMIRRORENCRYPTION:               1  NO
     REMOTEMIRRORTRACKRESOURCES:      172277  PAGES
     CONNECTION ID:                      E8
     MIRRORCOPYSTATE:                     1  ACTIVE
     MIRRORCOPYDATASTATE:                 2  USABLE
     REMOTEMIRRORDELIVERY:                0  SYNC
     REMOTEMIRRORPERFORMANCEMODE:         1  SYNC
     REMOTEMIRRORSYNCPRIORITY:           10  HIGH
     REMOTEMIRRORERRORRECOVERYTIMEOUT:    2  MINUTES
     REMOTEMIRRORAUTORESUME:              0
     SYNCSTATUS:                          1  SYNC IS NOT REQUIRED
     SYNCTYPE:                            1  SYNC IS NOT REQUIRED
------------------------------------------------------------------------------
F3=Exit   F12=Cancel   F19=Left   F20=Right   F24=More keys
```

Figure 6-1 QSYSPRT spooled file that shows the results of the GEOSTAT command

From the information that is shown in Figure 6-1, you can see the current state of the independent auxiliary storage pool (IASP) and the cluster resource group (CRG). The following key information can be determined from this output to assist you in troubleshooting:

- CRG: Shows the status of the device CRG. In this example, it is ACTIVE.
- Processing IASP: Shows the IASP number and name. In this example, it is 33 ESPIASP.
- Copy Role: Specifies whether this node is the Current Production (CP) or Current Backup (CB) (Mirror Copy) node. In this example, it is PRODUCTION.
- MirrorCopyState: Specifies the state of geographic mirroring. In this example, it is ACTIVE.
- MirrorCopyDataState: Specifies the state of the IASP data on the CB system. In this example, it is USABLE.
- RemoteMirrorDelivery: Specifies the transmission delivery mode for this geographic mirroring session. In this example, it is SYNC.
- RemoteMirrorPerformanceMode: Specifies the mirroring mode for this geographic mirroring session. In this example, it is SYNC.
- RemoteMirrorSyncPriority: Specifies the synchronization priority for this geographic mirroring session if a synchronization is needed. In this example, it is HIGH.
- SyncStatus: Specifies whether a synchronization is required. In this example, it is SYNC IS NOT REQUIRED.
- SyncType: Specifies whether a partial or full synchronization is required if SyncStatus shows SYNC IS REQUIRED. In this example, the value is SYNC IS NOT REQUIRED.

The **GEOSTAT** output in Figure 6-1 on page 90 is from a CP system where geographic mirroring is in a normal operating state. The **GEOSTAT** output in Figure 6-2 shows a geographic mirroring environment that was suspended and is now in a resuming state.

In this case, the MirrorCopyState shows RESUMING. While the session resumes, the MirrorCopyDataState always shows a status of UNUSABLE. The **GEOSTAT** macro is the tool to use to provide the best information about the synchronization status.

Note: The **ASMINFO** output is also collected by QMGTOOLS.

```
DISPLAY/ALTER/DUMP            GEOSTAT
RUNNING MACRO: GEOSTAT                          -ALL
PROCESSING CLUSTER NODE: RCHESP1
CRG ESPCRG     IS ACTIVE -- PRIMARY NODE
PROCESSING IASP: 33 ESPIASP    AVAILABLE  PRIMARY ASP
IASP INFORMATION:
     PHYSICAL COPY ID:                  0XE7
     COPY ROLE:                         0XD7  PRODUCTION
     REMOTEMIRRORERRORRECOVERYPOLICY:      2  SUSPEND
     REMOTEMIRRORENCRYPTION:               1  NO
     REMOTEMIRRORTRACKRESOURCES:      172277  PAGES
     CONNECTION ID:                      E8
     MIRRORCOPYSTATE:                     3  RESUMING
     MIRRORCOPYDATASTATE:                 3  UNUSABLE
     REMOTEMIRRORDELIVERY:                0  SYNC
     REMOTEMIRRORPERFORMANCEMODE:         1  SYNC
     REMOTEMIRRORSYNCPRIORITY:           10  HIGH
     REMOTEMIRRORERRORRECOVERYTIMEOUT:    2  MINUTES
     REMOTEMIRRORAUTORESUME:              0
     SYNCSTATUS:                          4  PERFORMING PARTIAL SYNC ONLY
     SYNCTYPE:                            2  PARTIAL SYNC IS REQUIRED
     SYNCHRONIZATION IS IN THE SYNCHRONIZATION STAGE WITH 68 PERCENT COMPLETE
     TRACKING SPACE:                      0% USED.
     DATA TRACKED CONN ID E8:        670010 PAGES.
```

Figure 6-2 GEOSTAT that shows a resuming IASP session

6.2.3 ASMINFO

The **ASMINFO** native macro is used to collect information about the geographic mirroring environment. This output provides historical information for IBM Support and development teams about reasons that the session was suspended and the duration of the “suspend to resume” time.

In addition, this data provides time stamps when geographic mirroring went into and out of flow control. The flow control is a condition where geographic mirroring temporarily pauses the sending of writes from the CP system to the CB system due to networking performance issues or CB performance side issues.

To collect this data, run the following steps:

1. From the command line, run Start System Service Tools (**STRSST**) and press Enter.
2. Sign in with a service tool profile and password that has authority to Display/Alter/Dump in SST.

3. Select option 1 - Start a service tool, and press Enter.
4. Select option 4 - Display/Alter/Dump, and press Enter.
5. To dump to a spooled file, select option 2 - Dump to printer, and press the Enter.
6. Select option 2 - Licensed Internal Code (LIC) data, and press the Enter.
7. Select option 14 - Advanced analysis, and press Enter.
8. On the Select Advanced Analysis Command panel, type option 1 (Select) next to the top blank line under the Command column. In the blank line, type ASMINFO, and press Enter.
9. In the Options field, type the following information where *nnn* is the IASP number and press Enter:

   ```
   -ASP nnn -r 50000
   ```

10. Exit out of SST and type the Work with Job (**WRKJOB**) command and press Enter. From the Work with Job display, select option 4 to Work with spooled files for that job. The spool file that is named QPCSMPRT contains the **ASMINFO** data. Provide this data to IBM Support for analysis.

Note: The **ASMINFO** output is also collected by QMGTOOLS.

6.2.4 DSMINFO

The **DSMINFO** macro provides information for various IASP (and SYSBAS) issues. IBM Support might request that **DSMINFO** output is collected for various IASP issues, such as detach and reattach issues, IASP **VRYCFG** issues, and others.

This data can be collected by running the following steps:

1. From the operating system command line, run Start System Service Tools (**STRSST**) and press Enter.
2. Sign in with a service tool profile and password that has authority to Display/Alter/Dump in SST.
3. Select option 1 - Start a service tool, and press Enter.
4. Select option 4 - Display/Alter/Dump, and press Enter.
5. To dump to a spooled file, select option 2 - Dump to printer, and press the Enter.
6. Select option 2 - Licensed Internal Code (LIC) data, and press Enter.
7. Select option 14 - Advanced analysis, and press Enter.
8. On the Select Advanced Analysis Command panel, type option 1 (Select) next to the top blank line under the Command column. In the blank line, type **DSMINFO**, and press Enter. Leave the Options field blank.
9. Exit out of SST and type the Work with Job (**WRKJOB**) command and press Enter. From the Work with Job display, select option 4 to Work with spooled files for that job. One spool file that is named QPCSMPRT contains the **DSMINFO** data. Provide this data to IBM Support for analysis.

Note: The **DSMINFO** data is also collected by QMGTOOLS.

6.3 Methods for troubleshooting

This section provides general items that help when you troubleshoot geographic mirroring, several common issues that can occur with geographic mirroring, and how to resolve those issues.

6.3.1 Ensuring that job descriptions have the correct logging levels

Several user profiles are used to perform various cluster and PowerHA actions. It is important when you troubleshoot that you ensure that the job descriptions for these user profiles were configured with the correct logging levels. The user profiles and associated job descriptions are listed in Table 6-1.

Table 6-1 PowerHA user profiles and associated job descriptions

User profile	Job description
QSYS	QGPL/QDFTJOBD
QCLUSTER	QSYS/QCSTJOBD
QHAUSRPRF	QHASM/QHAJOBD
QLPAR (IBM Copy Services Manager and Full System Copy Services Manager toolkits only)	QGPL/QLPARJOBD
Current interactive user	Display by running a WRKJOB, option 2.

To check each of these job descriptions, run the Display Job Description (**DSPJOBD**) command for each one. Each job description's message logging value must be set to the following value:

- Message logging level: 4
- Message logging severity: 00
- Message logging text: *SECLVL

If the LOG value of the job descriptions in Table 6-1 are not set to these values, use the Change Job Description (**CHGJOBD**) command to change them.

Note: If the job description is changed, the active jobs that use that job description do not pick up the new values until the job is restarted.

6.3.2 Using the command line interface when you re-create problems

When you re-create problems, it is best to use control language (CL) commands rather than menu options. CL commands allow easier analysis and better logging within the job log where the action is taken.

For example, if you experience an issue when you try to work with ASP copy descriptions, it is best to run the Work with Auxiliary Storage Pool Copy Descriptions (**WRKASPCPYD**) command, rather than to perform a Work with Cluster (**WRKCLU**) command and use option 10. This approach provides better debug information to the user, and to IBM Support, if necessary.

6.4 Common return codes when you manage geographic mirroring

When you are working in a geographic mirroring environment, certain actions can cause various outcomes that result in a CPFBA48 message with a wide range of reason codes. Due to the large number of reason codes that can occur for all known issues that result in CPFBA48, it is not possible to specify all of them within the message text.

The following document provides information about each of these possible return codes, their meaning, and possible recovery procedures:

http://www.ibm.com/support/docview.wss?uid=nas8N1020895

A

PowerHA Tools for IBM i

This appendix describes the PowerHA Tools for IBM i offerings and services that are available from IBM Systems Lab Services.

The PowerHA Tools for IBM i complement and extend the PowerHA and IBM storage capabilities for high availability (HA) and disaster recovery (DR).

The PowerHA Tools for IBM i provide the following benefits:

- Helps reduce business risk and improve resiliency for critical applications.
- Simplifies setup and automation of HA, DR, and backup solutions.
- Reduces the cost of maintaining and regular testing of an HA/DR environment.
- Facilitates flexible deployment options for single or multi-site protection.
- Assures consistent deployment by using preferred practices and experienced consultants.

For more information about PowerHA Tools for IBM i, see the following IBM Systems Lab Services and Training website:

http://www.ibm.com/systems/services/labservices

PowerHA Tools for IBM i

Table A-1 lists PowerHA Tools for IBM i that are available from IBM Systems Lab Services.

Table A-1 PowerHA Tools for IBM i

PowerHA Tools for IBM i	Capability	Benefit	DS8000	Storwize	Internal
Smart Assist for PowerHA on IBM i	Provides operator commands and scripts to supplement the PowerHA installation and ongoing operations for independent auxiliary storage pool (IASP)-enabled applications.	Simplifies deployment and ongoing management of HA for critical IBM i applications.	X	X	X
IASP Copy Services Manager (Automated recovery with faster IASP-level vary on with no system IPL)					
FlashCopy	Automates FlashCopy of IASP for daily offline backup with seamless Backup, Recovery, and Media Services (BRMS) integration.	Increases application availability by reducing or eliminating backup window for routine daily backups.	X	X	
Logical unit number (LUN)-level switching	Simplifies deployment and automates switching an IASP between IBM i cluster nodes in one data center.	Enables a business continuity manager to provide a simple, single-site HA solution.	X[1]	[2]	
Metro Mirror or Global Mirror	Simplifies initial deployment and automates ongoing server and storage management of two-site Metro Mirror or Global Mirror HA or DR solutions. Requires IASP-enabled applications.	Enables a business continuity manager to provide the seamless operation of integrated server and storage operations for two-site HA and DR.	X		
Metro Global Mirror (MGM)	Extends PowerHA functionality to provide a three-site server or storage replication solution that contains Metro Mirror for HA with Global Mirror for DR. Requires IASP-enabled applications and IBM Tivoli® Productivity Center for Replication.	Enables a business continuity manager to further lower business risk and maximize business resilience for critical applications that require three-site HA and DR protection.	X		
Full System Copy Services Manager (Automated recovery, and requires a full-system IPL)					
FlashCopy	Automates full-system FlashCopy for daily offline backup with integrated support for BRMS without IASP-enabled applications.	Increases application availability by reducing or eliminating the backup window for online daily backups. Enables an entry solution while you plan the IASP enablement.	X	X	
Metro Mirror or Global Mirror	Simplifies initial deployment and automates ongoing server and storage management of two-site Metro Mirror or Global Mirror HA or DR solutions without IASP-enabled applications.	Enables a business continuity manager to provide seamless operation of integrated server and storage operations for HA and DR. Enables an entry solution while an IASP enablement is planned.	X	X[3]	

[1] DS8000 support is available with PowerHA Tools for IBM i V6.1 or earlier and included in PowerHA SystemMirror V7.1.

[2] V7000 support is included with PowerHA V7.1 TR6.

[3] Storwize Full System Replication requires Full System Replication for PowerHA.

IBM Lab Services Offerings for PowerHA for IBM i

Table A-2 lists the PowerHA for IBM i service offerings that are available from IBM Systems Lab Services.

Table A-2 PowerHA for IBM i service offerings

PowerHA for IBM i service offering	Description
IBM i High Availability Architecture and Design Workshop	An experienced IBM i consultant conducts a planning and design workshop to review solutions and alternatives to meet HA and DR and backup and recovery requirements. The consultant provides an architecture and implementation plan to meet these requirements.
PowerHA for IBM i Bandwidth Analysis	An experienced IBM i consultant reviews network bandwidth requirements for implementing storage data replication. IBM reviews I/O data patterns and provides a bandwidth estimate to build into the business and project plan for clients who are deploying PowerHA for IBM i.
IBM i Independent Auxiliary Storage Pool (IASP) Workshop	An experienced IBM i consultant provides jumpstart services for migrating applications into an IASP. Training includes enabling applications for IASPs, clustering techniques, and managing PowerHA and HA and DR solution options with IASPs.
PowerHA for IBM i Implementation Services	An experienced IBM consultant provides services to implement an HA/DR solution for IBM Power Systems servers with IBM Storage. Depending on specific business requirements, the end-to-end solution implementation can include a combination of PowerHA for IBM i and PowerHA Tools for IBM i, and appropriate storage software, such as Metro Mirror, Global Mirror, or FlashCopy.

B

Worksheet for configuring geographic mirroring

This appendix provides a worksheet to assist with the PowerHA geographic mirroring configuration as described in Chapter 3, “Implementing geographic mirroring” on page 27.

The worksheet in Table B-1 on page 100 can be used for configuring geographic mirroring, and it is a handy reference when completing the commands.

We suggest that you not use “primary” and “backup” as part of the naming conventions because these names can cause confusion after a switch is performed. After you configure geographic mirroring, names cannot be changed unless you delete the entire configuration and start over.

For cluster communications, which are also known as the *heartbeat*, a minimum configuration of one IP address is required. However, for redundancy, it is considered a preferred practice to use a shared Ethernet adapter, a virtual IP address across two Ethernet adapters, or two IP addresses with two Ethernet adapters. In addition, all cluster IP addresses on a specific node must be able to communicate with all cluster IP addresses on all nodes in the cluster.

For cluster resource group (CRG) dataport communications, a minimum configuration of one IP address is required. For redundancy and improved throughput, up to four Ethernet addresses can be used. These Ethernet addresses can be configured on a shared Ethernet adapter, as virtual IP addresses across multiple Ethernet adapters, or even multiple IP addresses on different subnets. No matter which configuration you choose, each dataport address on a certain node must be able to communicate with at least one address on all other nodes in the CRG.

Table B-1 shows a worksheet for you to use to configure geographic mirroring for your environment.

Table B-1 Worksheet for configuring geographic mirroring

Parameter	Keyword and type	Value	Description	Commands where parameter is used
IASP	ASPDEV CHAR(10)		Independent auxiliary storage pool (IASP) name or device description.	CFGDEVASP
Cluster	CLUSTER CHAR(10)		Cluster name.	CRTCLU ADDCLUMON ADDDEVDMN CRTCAD ADDCADMRE
Node identifier (primary)	NODE CHAR(8)		Primary or source node name.	CRTCLU ADDCLUMON ADDDEVDMNE CRTCAD
Node identifier (backup)	NODE CHAR(8)		Backup or target node name.	CRTCLU ADDCLUMON ADDDEVDMNE CRTCAD
Node identifier (additional)	NODE CHAR(8)		Additional node name, if required.	CRTCLU ADDCLUMON ADDDEVDMNE CRTCAD
IP address (primary node)	Dotted IP		Cluster IP address for the primary node. It can have 1 or 2 addresses.	CRTCLU
IP address (backup node)	Dotted IP		Cluster IP address for the backup node. It can have 1 or 2 addresses.	CRTCLU ADDCLUNODE
IP address (additional node)	Dotted IP		Cluster IP address for the additional node. It can have 1 or 2 addresses. It is not applicable for all configurations.	CRTCLU ADDCLUNODE
Device domain	DEVDMN CHAR(10)		Device domain name.	ADDDEVDMNE
Cluster administrative domain	ADMDMN CHAR(10)		Cluster administrative domain name.	CRTCAD ADDCADMRE
Cluster resource group	CRG CHAR(10)		Cluster resource group (CRG) name.	CRTCRG
Site name (primary)	SRCSITE CHAR(8)		Primary node site name.	CRTCRG
Site name (backup)	TGTSITE CHAR(8)		Backup node site name.	CRTCRG

Parameter	Keyword and type	Value	Description	Commands where parameter is used
Dataport IP address (primary node)	dotted IP		One to four numeric IP addresses that are used for dataport services (primary).	CRTCRG
Dataport IP address (backup node)	Dotted IP		One to four numeric IP addresses that are used for dataport services (backup).	CRTCRG
Session	SSN CHAR(10)		Auxiliary storage pool (ASP) session name.	CFGGEOMIR
Source ASP copy description	CHAR(10)		Copy description on the primary node.	CFGGEOMIR
Backup ASP copy description	CHAR(10)		Copy description on the backup node.	CFGGEOMIR

Several of the parameters that are listed in Table B-1 on page 100 are described in more detail:

- IASP: The name of the IASP device that contains the initial source copy of the Independent Auxiliary Storage Pool (IASP). For more information about creating the IASP, see the IBM Redbooks publication, *IBM PowerHA SystemMirror for i: Preparation (Volume 1 of 4)*, SG24-8400.
- Cluster: A meaningful name up to 10 characters that identifies the unique cluster. The nodes are added to the cluster.
- Node identifier: A meaningful name up to eight characters that represents the system or partition. A minimum of two nodes must be defined. The primary or source node normally contains the production or source IASP. The backup or target node normally contains the backup or target IASP.
- IP address: One or two IP addresses are configured on each node for cluster communications. For a production environment, we consider it a preferred practice (but it is not required) to use two dedicated, redundant Ethernet ports with the cluster IP addresses.

  Reminder: All cluster IP addresses on a specific node must be able to communicate with all cluster IP addresses on any other node in the cluster.

- Device domain: A meaningful name up to 10 characters that identifies the device domain within the cluster. The nodes and IASP belong to the device domain.
- Cluster administrative domain: A meaningful name up to 10 characters that identifies the administrative domain for the nodes of the cluster.
- Cluster resource group: A meaningful name up to 10 characters that identifies the recovery domain that consists of the IASP, cluster nodes, and administrative domain.
- Site name: A meaningful name up to eight characters that identifies a node within a recovery domain. Each node must have a unique site name.

- Dataport IP address: One to four IP addresses are configured on each node for dataport services (which is the transmission of replicated data from the source node to the target node). For a production environment, we consider it a preferred practice (but not required) to use at least two dedicated, redundant Ethernet ports.

  Reminder: Every dataport IP address on a specific node must be able to communicate with at least one IP address on any other node in the recovery domain.

- Session: A meaningful name up to 10 characters, which identifies the ASP session, that links the ASP copy descriptions.
- Copy description: A meaningful name up to 10 characters that identifies the copy of the node IASP. Each ASP copy in the session must have a unique copy description.

C

Additional material

This book refers to additional material that can be downloaded from the Internet as described in the following sections.

Locating the web material

The web material associated with this book is available in softcopy on the Internet from the IBM Redbooks web server. Point your web browser at this website:

`ftp://www.redbooks.ibm.com/redbooks/SG248401`

Alternatively, you can go to the IBM Redbooks website:

ibm.com/redbooks

Select **Additional materials** and open the directory that corresponds with the IBM Redbooks form number, SG248401.

Using the web material

The additional web material that accompanies this book includes the following IBM i save files:

File name	*Description*
MONASPSSN.mbr	Sample CL program to monitor replication status
MONNODSTS.mbr	Sample CL program to monitor the cluster node status
MONADMDMN.mbr	Sample CL program to monitor the cluster administrative domain
MONCRGSTS.mbr	Sample CL program to monitor the cluster resource group status
MONIASPSTS.mbr	Sample CL program to verify the IASP utilization status

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- *IBM PowerHA SystemMirror for i: Preparation (Volume 1 of 4)*, SG24-8400
- *IBM PowerHA SystemMirror for i: Using DS8000 (Volume 2 of 4)*, SG24-8403
- *IBM PowerHA SystemMirror for i: Using IBM Storwize (Volume 3 of 4)*, SG24-8402

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Online resources

These websites are also relevant as further information sources:

- IBM PowerHA SystemMirror for i website:

  https://ibm.biz/Bd4JQx
- Schowler Routes on the IBM i:

  https://ibm.biz/Bd4AcJ
- IBM i 7.2 Knowledge Center:

  http://www.ibm.com/support/knowledgecenter/ssw_ibm_i_72/rzahg/ic-homepage.htm
- IBM PowerHA SystemMirror for i DeveloperWorks website:

  https://ibm.biz/Bd4ub8
- IBM Support portal:

  https://ibm.biz/BdXqvs
- Geographic Mirroring: Listing of Most Common Return Codes for MSGCPFBA48

  http://www.ibm.com/support/docview.wss?uid=nas8N1020895

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Redbooks

IBM PowerHA SystemMirror for i: Using Geographic Mirrioring

(0.2”spine)
0.17”<->0.473”
90<->249 pages

SG24-8401-00

ISBN 0738441716

Printed in U.S.A.