

IBM i
Version 7 Release 3

Availability
Implementing high availability



Note

Before using this information and the product it supports, read the information in [“Notices” on page 157.](#)

This edition applies to IBM i 7.3 (product number 5770-SS1) and to all subsequent releases and modifications until otherwise indicated in new editions. This version does not run on all reduced instruction set computer (RISC) models nor does it run on CISC models.

This document may contain references to Licensed Internal Code. Licensed Internal Code is Machine Code and is licensed to you under the terms of the IBM License Agreement for Machine Code.

© **Copyright International Business Machines Corporation 1998, 2019.**

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

Contents

- Implementing high availability.....1**
- What's new for IBM i 7.3.....1
- PDF file for Implementing high availability..... 3
- Installing IBM PowerHA SystemMirror for i licensed program.....3
- Uninstalling IBM PowerHA SystemMirror for i licensed program.....4
- Planning your high availability solution.....4
 - Cluster applications.....5
 - Planning clusters..... 12
 - Planning data resiliency..... 21
 - Planning environment resiliency.....39
 - Security planning for high availability.....40
- Implementing PowerHA..... 41
 - Configuring high availability infrastructure.....42
 - Configuring cluster administrative domains..... 63
 - Configuring independent disk pools..... 66
 - Configuring geographic mirroring..... 69
 - Configuring Metro Mirror..... 69
 - Configuring Global Mirror..... 70
 - Configuring switched logical units (LUNs)..... 71
 - Configuring a FlashCopy session..... 72
 - Configuring DS8000 Full System HyperSwap..... 72
 - Configuring DS8000 HyperSwap with independent auxiliary storage pools (IASPs)..... 73
- Managing PowerHA.....75
 - Scenarios: Managing high availability solutions..... 75
 - Managing clusters..... 79
 - Managing independent disk pools..... 118
 - Managing geographic mirroring..... 119
 - Managing Metro Mirror..... 124
 - Managing Global Mirror..... 127
 - Managing switched logical units (LUNs)..... 130
 - Managing the FlashCopy technology..... 130
 - Managing DS8000 Full System HyperSwap..... 133
 - Managing DS8000 HyperSwap with independent auxiliary storage pools (IASPs)..... 134
- Troubleshooting your high availability solution..... 141
 - Troubleshooting clusters..... 141
 - Troubleshooting geographic mirroring..... 151
 - Troubleshooting Metro Mirror, Global Mirror, and FlashCopy..... 151
 - Troubleshooting HyperSwap..... 151
- Related information for Implementing high availability.....153**
- Notices.....157**
- Programming interface information..... 158
- Trademarks..... 158
- Terms and conditions..... 159

Implementing high availability

For IBM® i environments, a task-based approach is used to implement high availability. The *task-based approach* allows you to design and build a customized high-availability solution for your business, using different interfaces for technologies that are related to high availability.

You are required to install the IBM PowerHA® SystemMirror® for i licensed program number (5770-HAS) on each system that is participating in high availability. The task-based approach uses the PowerHA graphical interface from which you can create and manage the cluster, cluster resource groups, cluster administrative domains, and independent ASPs. With the task-based approach, other interfaces, such as the CL commands, can also be used to manage the technologies within your solution.

Note: By using the code examples, you agree to the terms of the [Code license and disclaimer information](#).

What's new for IBM i 7.3

Read about this new topic collection for this release that contains previous information about Clusters and High Availability.

What's new as of December 2019

PowerHA has enhanced integrated recovery from a data-center outage in a HyperSwap with a LUN-level switching environment that enables HyperSwap protection to be restored with a single command.

Integration with DS8000 continues:

- Integration of PowerHA IASP-based replication with IBM Copy Services Manager for DS8000 Metro Mirror and DS8000 Global Mirror replication solutions.
- PowerHA has been enhanced to enable DS8000 HyperSwap configuration in conjunction with Global Mirror for regional outage protection.

SVC communications enhancements:

- Communication to the SVC has been enhanced for robustness in certain heavily loaded environments where the SVC communication fails with `ssh_exchange_identification: Connection closed by remote host`. PowerHA has implemented connection retries when this error is encountered.

PowerHA has improved control over changing the synchronization option for an administrative domain using the Change Cluster Administrative Domain (**CHGCAD**) command even when some nodes in the administrative domain are inactive.

PowerHA enhanced support for replication of encrypted passwords set with(using) the QSYSUPWD API. Previously, the QSYSUPWD API caused the user-affected profile to be marked as inconsistent within the administrative domain.

What's new as of October 2016

IBM PowerHA for i Standard Edition is enhanced to support a new representational state transfer (REST) server for advanced node failure detection. This function is provided through a new function PowerHA PTF for 7.2. The PTF numbers can be found on the [IBM PowerHA SystemMirror for i wiki](#).

Making use of the new REST support in enhanced commands requires IBM PowerHA for i Standard Edition to be installed with a valid license key and the new function PowerHA PTFs to be loaded and applied. Making use of the new REST support in enhanced APIs requires the new function base operating system PTFs to be loaded and applied.

The following commands are enhanced.

- [Add Cluster Monitor \(ADDCLUMON\)](#) command
- [Change Cluster Monitor \(CHGCLUMON\)](#) command

- [Remove Cluster Monitor \(RMVCLUMON\)](#) command
- [Display Cluster Information \(DSPCLUINF\)](#) command
- [Work Cluster \(WRKCLU\)](#) command

The following APIs are enhanced.

- [List Cluster Information](#) API
- [Add Cluster Monitor](#) API
- [Change Cluster Monitor](#) API
- [Remove Cluster Monitor](#) API

Enhanced 7.2 IBM PowerHA SystemMirror for i licensed program

IBM PowerHA SystemMirror for i licensed program is not being refreshed for 7.3. That means licensed program number (5770-HAS) 7.2 runs on either IBM i 7.3 or IBM i 7.2.

DS8000 HyperSwap with independent auxiliary storage pools (IASPs)

IBM PowerHA for i Enterprise Edition supports DS8000® HyperSwap® at an IASP level. HyperSwap functions independently for SYSBAS and IASPs. HyperSwap relationships can be configured for SYSBAS logical units only, IASP logical units only, or both. DS8000 HyperSwap with IASPs has been added to 7.2 IBM PowerHA SystemMirror for i licensed program.

To use DS8000 HyperSwap with IASPs requires IBM PowerHA for i Enterprise Edition be installed, with a valid license key. And the new function PowerHA PTF be loaded and applied. The PTF number can be found on the [IBM PowerHA SystemMirror for i wiki](#).

The following commands are enhanced or new to the IBM PowerHA SystemMirror for i licensed program.

- Enhanced [Add ASP Copy Description](#) command
- Enhanced [Change ASP Copy Description](#) command
- Enhanced [Retrieve ASP Copy Description](#) command
- Enhanced [Retrieve ASP Session](#) command
- New [Retrieve SVC Copy Description](#) command
- New [Retrieve SVC Session](#) command
- Enhanced [Change HyperSwap Status](#) command
- Enhanced [Display HyperSwap Status](#) command
- New [Work HyperSwap Status](#) command
- New [Add HA Configuration Description](#) command
- New [Change HA Configuration Description](#) command
- New [Display HA Configuration Description](#) command
- New [Remove HA Configuration Description](#) command
- New [Work HA Configuration Description](#) command
- Enhanced [Change Cluster Recovery](#) command

The following APIs are enhanced or new to the IBM PowerHA SystemMirror for i licensed program.


- Enhanced [Retrieve ASP Copy Information](#) API
- New [Retrieve HA Configuration Description](#) API
- New [Retrieve HA Status](#) API

Switched disks is no longer supported



Switching disks between partitions within the same CEC is no longer supported at IBM i 7.3.

PDF file for Implementing high availability

You can view and print a PDF file of this information about implementing high availability.

To view or download the PDF version of this document, select [Implementing high availability](#) .

You can view or download these related topic collection PDFs:


- [High availability overview](#)  contains the following topics:
 - Benefits of high availability
 - High availability criteria
 - Components of high availability
 - IBM PowerHA SystemMirror for i overview
- [High availability technologies](#)  contains the following topics:
 - IBM i Cluster technology
 - Advanced node failure detection
 - Cluster administrative domain
 - PowerHA data replication technologies
 - High availability management
 - Resource Monitoring and Control (RMC)

Saving PDF files

To save a PDF on your workstation for viewing or printing:

1. Right-click the PDF link in your browser.
2. Click the option that saves the PDF locally.
3. Navigate to the directory in which you want to save the PDF.
4. Click **Save**.

Downloading Adobe Reader

You need Adobe Reader installed on your system to view or print these PDFs. You can download a free copy from the [Adobe Web site](http://www.adobe.com/products/acrobat/readstep.html) (www.adobe.com/products/acrobat/readstep.html) .

Related reference

Related information for [Implementing high availability](#)

Product manuals, IBM Redbooks publications, websites, and other information center topic collections contain information that relates to the High availability topic collection. You also can find related information about implementing independent disk pools, PowerHA technologies, and disaster recovery. You can view or print any of the PDF files.

Installing IBM PowerHA SystemMirror for i licensed program

Before you can implement an IBM i high-availability solution, you must install the IBM PowerHA SystemMirror for i licensed program (5770-HAS) on each system that participates in high availability.

Before installing the IBM PowerHA SystemMirror for i licensed program, you should have completed the following installation requirements:

1. Install or upgrade to i 7.2 or i 7.3 operating system.
2. Install IBM i operating system Option 41 (HA Switchable Resources).

To install the IBM PowerHA for i licensed program, complete the following steps:

1. Enter GO LICPGM from a command line.
2. At the Work with Licensed Programs display, select option 11 (Install licensed programs).
3. Select Product 5770-HAS, option *BASE to install IBM PowerHA SystemMirror for i. Press enter.
4. At the Install Options display, type in the name of your installation device as requested. Press enter to start the installation.
5. Using asynchronous geographic mirroring, metro mirroring, global mirroring or DS8000 HyperSwap with IASPs requires IBM PowerHA for i Enterprise Edition (option 1) be installed. Select Product 5770-HAS, option 1 to install IBM PowerHA for i Enterprise Edition. Press enter.
6. Using PowerHA graphical interface, commands, synchronous geographic mirroring, switched logical units, or FlashCopy® requires IBM PowerHA for i Standard Edition (option 2) be installed. Select Product 5770-HAS, option 2 to install IBM PowerHA for i Standard Edition. Press enter.
7. Using DS8000 Full System HyperSwap requires IBM PowerHA for i Express® Edition (option 3) be installed. Select Product 5770-HAS, option 3 to install IBM PowerHA for i Express Edition. Press enter.

After successfully installing the IBM PowerHA SystemMirror for i licensed program, you need to restart the INETD server. For information about how to start INETD, see [“Starting the INETD server” on page 49](#).

Uninstalling IBM PowerHA SystemMirror for i licensed program

If you no longer want to use interfaces or functions associated with IBM PowerHA SystemMirror for i licensed program, you need to uninstall the product.

To uninstall the IBM PowerHA SystemMirror for i licensed program, complete these steps:

1. Type GO LICPGM and press enter. The Work with licensed programs menu appears.
2. Select option 12 (Delete licensed programs). The Delete licensed programs display appears.
3. Type 4 (Delete) in the Option column in front of 5770-HAS.
4. Press enter. You are shown the **Confirm Delete of Licensed Programs** display.
5. Press enter if your selections are correct. Press F12 to make corrections.
6. You receive confirmation messages from the system when the licensed programs are deleted.
7. Press F12 to return to the Work with licensed programs menu.

If you have difficulty deleting a licensed program following these steps, type **ENDSBS *ALL *IMMED** and press enter. Then proceed with step 1 again.

Planning your high availability solution

Before configuring an IBM i high-availability solution, proper planning is necessary to ensure all the requirements for the solution have been met.

Each high availability technology has minimum requirements that should be met before configuring a specific solution. In addition to these requirements, it is important to also determine which resources should be made resilient. These resources, such as applications, data, and devices, should be evaluated to determine whether they should be highly available. If they require high availability, it is important to make any necessary changes to the environment prior to configuring a solution for high availability. For example, you might have data that resides in SYSBAS, which should be highly available. Before configuring a solution, you should move that data to an independent disk pool. Applications might also require changes to enable high availability.

Cluster applications

Application resilience is one of the key elements in a clustered environment. If you are planning to write and use highly available applications in your cluster you should be aware that these applications have specific availability specifications.

By taking advantage of resilient applications in your cluster, an application can be restarted on a different cluster node without requiring you to reconfigure the clients. In addition, the data that is associated with the application will be available after switchover or failover. This means that the application user can experience minimal, or even seamless, interruption while the application and its data switch from the primary node to the backup node. The user does not need to know that the application and data have moved on the back end.

In order to achieve application resiliency in your cluster, applications that meet certain availability specifications must be used. Certain characteristics must be present in the application in order for it to be switchable, and therefore always available to the users of the application in the cluster. See [High Availability and Clusters](#) for details on these application traits. Because these requirements exist, you have the following options for using a switchable application in your cluster:

1. Purchase a cluster-enabled software application

Software products that are cluster-enabled meet certain high-availability requirements.

2. Write or change your own application to make it highly available

Independent software vendors and application programmers can customize applications to allow them to be switchable in a IBM i clustered environment.

Once you have a resilient application, it must be managed within your cluster.

Related information

[IBM PowerHA](#)

Identifying resilient applications

Not every application will give you the availability benefits of clustering.

An application must be resilient in order to take advantage of the switchover and failover capabilities provided by clustering. Application resilience allows the application to be restarted on the backup node without having to reconfigure the clients using the application. Therefore your application must meet certain requirements to take full advantage of the capabilities offered by clustering.

IBM i architecture for cluster-enabled applications

Additional end-user value is provided by any application that is highly available, recognizing applications that continue to be available in the event of an outage, planned or unplanned.

IBM i has provided an application resilience architecture that supports various degrees of highly available application. Applications on the high end of this spectrum demonstrate highly available characteristics, provide automation of the highly available environment, and are managed through high availability management interfaces.

These applications have the following characteristics:

- The application can switch over to a backup cluster node when the primary node becomes unavailable.
- The application defines the resilient environment in the Resilient Definition and Status Data Area to enable automatic configuration and activation of the application by a cluster management application.
- The application provides application resilience by means of an application CRG exit program to handle cluster related events, taking advantage of the capabilities of the IBM i cluster resource services.
- The application provides an application restart function that repositions the user to an application menu screen or beyond.

Applications that demonstrate more stringent availability and restart characteristics have the following characteristics:

- The application provides enhanced application resilience through more robust handling of cluster events (action codes) by the application CRG exit program.

- The application provides a greater level of application restart support. For host-centric applications, the user will be repositioned to a transaction boundary by commitment control or checkpoint functions. For client-centric applications, the user will experience a seamless failover with minimal service interruption.

Writing a highly available cluster application

A highly available application is one that can be resilient to a system outage in a clustered environment.

Several levels of application availability are possible:

1. If an application error occurs, the application restarts itself on the same node and corrects any potential cause for error (such as corrupt control data). You can view the application as though it had started for the first time.
2. The application performs some amount of checkpoint-restart processing. You can view the application as if it were close to the point of failure.
3. If a system outage occurs, the application is restarted on a backup server. You can view the application as though it had started for the first time.
4. If a system outage occurs, the application is restarted on a backup server and performs some amount of checkpoint-restart processing across the servers. You can view the application as if it were close to the point of failure.
5. If a system outage occurs, a coordinated failover of both the application and its associated data to another node or nodes in the cluster occurs. You can view the application as though it had started for the first time.
6. If a system outage occurs, a coordinated failover of both the application and its associated data to another node or nodes in the cluster occurs. The application performs some amount of checkpoint-restart processing across the servers. You can view the application as if it were close to the point of failure.

Note: In cases 1 through 4 above, you are responsible for recovering the data.

Making application programs resilient

Learn how to make application programs resilient.

A resilient application is expected to have the following characteristics:

- The application can be restarted on this node or another node
- The application is accessible to the client through the IP address
- The application is stateless or state information is known
- Data that is associated with the application is available after switchover

The three essential elements that make an application resilient to system outages in a clustered environment are:

The application itself

How tolerant is the application to errors or to system outages, and how transparently can the application restart itself?

The application can handle this through the use of clustering capabilities.

Associated data

When an outage occurs, does it affect the availability of any associated data?

You can store critical data in switched disks which allow data to remain available during an outage. Alternatively, a cluster middleware IBM Business Partner replication product that takes advantage of the clustering capabilities can handle this.

Control capabilities and administration

How easy is it to define the environment that supports the availability of the data and the application?

IBM PowerHA SystemMirror for i licensed program provides several interfaces to configure and manage high-availability solutions and technology. The IBM PowerHA SystemMirror for i licensed program provides the following interfaces:

PowerHA graphical interface

Use the graphical interface to easily configure, monitor, and manage your High Availability solution. For customers that are upgrading from a release before 7.2, it combines the simplicity of the High Availability Solutions Manager graphical interface with the flexibility of the Cluster Resource Services graphical interface in a single graphical interface.

PowerHA commands

These commands provide similar functions but are available through a command-line interface.

PowerHA APIs

These PowerHA APIs allow you to work with new function for independent disk pools.

In addition, you can also use a third-party cluster management interface that uses the clustering APIs and also combines resilient applications with resilient data can handle this.

Related information

[High availability management](#)

Restarting highly available cluster applications

To restart an application, the application needs to know its state at the time of the failover or switchover.

State information is application specific; therefore, the application must determine what information is needed. Without any state information, the application can be restarted on your PC. However, you must reestablish your position within the application.

Several methods are available to save application state information for the backup system. Each application needs to determine which method works best for it.

- The application can transfer all state information to the requesting client system. When a switchover or failover occurs, the application uses the stored state on the client to reestablish the state in the new server. This can be accomplished by using the Distribute Information API or Clustered Hash Table APIs.
- The application can replicate state information (such as job information and other control structures that are associated with the application) on a real-time basis. For every change in the structures, the application sends the change over to the backup system.
- The application can store pertinent state information that is associated with it in the exit program data portion of the cluster resource group for that application. This method assumes that a small amount of state information is required. You can use the `Change Cluster Resource Group (QcstChangeClusterResourceGroup)` API to do this.
- The application can store state information in a data object that is being replicated to the backup systems along with the application's data.
- The application can store state information in a data object contained in the IASP that also contains the application's data.
- The application can store the state information about the client.
- No state information is saved, and you need to perform the recovery.

Note: The amount of information that is required to be saved is lessened if the application uses some form of checkpoint-restart processing. State information is only saved at predetermined application checkpoints. A restart takes you back to the last known checkpoint which is similar to how database's commitment control processing works.

Calling a cluster resource group exit program

The cluster resource group exit program is called during different phases of a cluster environment.

This program establishes the environment necessary resiliency for resources within a cluster. The exit program is optional for a resilient device CRG but is required for the other CRG types. When a cluster resource group exit program is used, it is called on the occurrence of cluster-wide events, including the following:

- A node leaves the cluster unexpectedly
- A node leaves the cluster as a result of calling the `End Cluster Node (QcstEndClusterNode)` API or `Remove Cluster Node Entry (QcstRemoveClusterNodeEntry)` API

- The cluster is deleted as a result of calling the Delete Cluster (QcstDeleteCluster) API
- A node is activated by calling the Start Cluster Node (QcstStartClusterNode) API
- Communication with a partitioned node is re-established

The exit program completes the following processes:

- Runs in a named activation group or the caller's activation group (*CALLER).
- Ignores the restart parameter if the exit program has an unhandled exception or is canceled.
- Provides a cancel handler.

When a cluster resource group API is run, the exit program is called from a separate job with the user profile specified on the Create Cluster Resource Group (QcstCreateClusterResourceGroup) API. The separate job is automatically created by the API when the exit program is called. If the exit program for a data CRG is unsuccessful or ends abnormally, the cluster resource group exit program is called on all active nodes in the recovery domain by using an action code of Undo. This action code allows any unfinished activity to be backed out and the original state of the cluster resource group to be recovered.

Suppose an unsuccessful switchover occurs for a device CRG. After switching back all the devices, if all of the devices were varied-on successfully on the original primary node, clustering calls the exit program on the original primary node by using an action code of Start.

If the exit program for an application CRG is unsuccessful or ends abnormally, cluster resource services attempt to restart the application if the status of the CRG is active. The cluster resource group exit program is called by using an action code of Restart. If the application cannot be restarted in the specified maximum number of attempts, the cluster resource group exit program is called by using an action code of Failover. The restart count is reset only when the exit program is called by using an action code of start, which can be the result of a start CRG, a failover, or a switchover.

When the cluster resource group is started, the application CRG exit program called on the primary node is not to return control to cluster resource services until the application itself ends or an error occurs. After an application CRG is active, if cluster resource services must notify the application CRG exit program of some event, another instance of the exit program is started in a different job. Any action code other than Start or Restart is expected to be returned.

When a cluster resource group exit program is called, it is passed a set of parameters that identify the cluster event being processed, the current state of the cluster resources, and the expected state of the cluster resources.

For complete information about cluster resource group exit programs, including what information is passed to the exit program for each action code, see Cluster Resource Group Exit Program in the Cluster API documentation. Sample source code has been provided in the QUSRTOOL library which can be used as a basis for writing an exit program. See the TCSTAPPEXT member in the QATTSYSC file.

Application CRG considerations

An application cluster resource group manages application resiliency.

Managing application CRG takeover of IP addresses

You can manage application CRG takeover of IP addresses by using cluster resource services. You can also manage them manually.

You can manage the application takeover IP address that is associated with an application CRG in two ways. The easiest way, which is the default, is to let cluster resource services manage the takeover IP address. This method directs cluster resource services to create the takeover IP address on all nodes in the recovery domain, including nodes subsequently added to the recovery domain. When this method is selected, the takeover IP address cannot currently be defined on any node in the recovery domain.

The alternative way is to manage the takeover IP addresses yourself. This method directs cluster resource services to not take any steps to configure the takeover IP address; the user is responsible for configuration. You must add the takeover IP address on all nodes in the recovery domain (except on replicate nodes) before starting the cluster resource group. Any node to be added to the recovery domain of an active CRG must have the takeover IP address configured before being added.

Related concepts

Example: Application cluster resource group failover actions

This example shows how one failover scenario works. Other failover scenarios can work differently.

Multiple subnets

It is possible to have the application takeover IP address work across multiple subnets, although the default is to have all recovery domain nodes on the same subnet. To configure the application takeover IP address when the nodes in the recovery domain span multiple subnets, you need to enable the switchover environment.

Enabling application switchover across subnets with IPv4

Clustering, in general, requires that all cluster nodes in the recovery domain of an application cluster resource group reside on the same LAN (use the same subnet addressing). Cluster resource services supports a user configured takeover IP address when configuring application CRGs.

Address Resolution Protocol (ARP) is the network protocol that is used to switch the configured application takeover IP address from one node to another node in the recovery domain. To enable the application switchover across subnets, you need to use the virtual IP address support and the Routing Information Protocol (RIP) for IPv4.

The following manual configuration steps are required to enable the switchover environment. **This set of instructions must be done on all the nodes in the recovery domain, and repeated for the other nodes in the cluster that will become nodes in the recovery domain for the given application CRG.**

1. Select an IPv4 takeover IP address to be used by the application CRG.
 - To avoid confusion, this address should not overlap with any other existing addresses used by the cluster nodes or routers. For example, if choosing 19.19.19.19, ensure that 19.0.0.0 (19.19.0.0) are not routes known by the system routing tables.
 - Add the takeover interface (for example, 19.19.19.19). Create it with a line description of *VIRTUALIP, subnet mask of 255.255.255.255 (host route), maximum transmission unit of 1500 (any number in the range 576-16388), and autostart of *NO. This takeover address (for example, 19.19.19.19) must exist as a *VIRTUALIP address before identifying it as an associated local interface in next step. It does not, however, have to be active.
2. Associate the intended takeover IP address with one or both of the IP addresses that you specify to be used by cluster communications when you create the cluster or add a node to the cluster.
 - For example, this means that you make the 19.19.19.19 takeover address an associated local interface on the IP address for the cluster node. This must be done for each cluster address on each cluster node.
Note: The cluster addresses must be ended to accomplish this change under the Configure TCP/IP (CFGTCP) command.
3. Create the cluster and create any CRGs. For the application CRG, specify QcstUserCfgrTakeoverIpAddr for the **Configure takeover IP address** field. Do not start any application CRGs.
4. Using Configure TCP/IP applications (option 20) from the Configure TCP/IP menu, then Configure RouteD (option 2), then Change RouteD attributes (option 1), ensure that the Supply field is set to *YES. If not, set it to *YES. Then start or restart RouteD (RIP or RIP-2) on each cluster node.
 - NETSTAT option 3 shows the RouteD using a local port if currently running. RouteD must be running and advertising routes (ensure that the Supply field is set to *YES) on every cluster node in the CRG recovery domain.
5. Ensure that all the commercial routers in the network that interconnect the recovery domain LANs are accepting and advertising host routes for RIP.
 - This is not necessarily the default setting for routers. The language varies with router manufacturer, but the RIP interfaces settings should be set to send host routes and receive dynamic hosts.
 - This also applies to both the router interfaces that point to the systems as well as the router-to-router interfaces.

Note: Do not use an IBM i machine as the router in this configuration. Use a commercial router (IBM or otherwise) that is designed for routing purposes. IBM i routing cannot be configured to handle this function.

6. Manually activate the takeover address on one of the cluster nodes:

- a) Wait up to 5 minutes for RIP to propagate the routes.
- b) Ping the takeover address from all nodes in the CRG recovery domain and from selected clients on the LANs who will be using this address.
- c) Ensure the takeover address is ended again.

(Clustering will start the address on the specified primary node when the CRGs are started.)

7. Start the application CRGs.

- The takeover address is started by clustering on the specified, preferred node, and RIP advertises the routes throughout the recovery domain. RIP might take up to 5 minutes to update routes across the domain. The RIP function is independent from the start CRG function.

Important:

- If the above procedure is not followed for all cluster nodes in the application CRG recovery domain, the cluster hangs during the switchover process.
- Even though you do not perform a failover to replica nodes, it is a good idea to perform the procedure on the replica nodes in the event that they might be changed at a later date in time to become a backup.
- If you want to use multiple virtual IP addresses, then each one will require a separate application CRG and a separate IP address with which to be associated. This address may be another logical IP address on the same physical adapter or it may be another physical adapter altogether. Also, care must be taken to prevent ambiguities in the routing tables. This is best achieved by doing the following:
 - Add a *DFTRROUTE to the routing table for each virtual IP address.
 - To use multiple IP address use CFGTCP (option 2).
 - Set all parameters, including the next hop, the same to reach the router of choice; however, the Preferred binding interface should be set to the local system IP address that is associated with the virtual IP address that is represented by this route.

Enabling application switchover across subnets with IPv6

Clustering, in general, requires that all cluster nodes in the recovery domain of an application cluster resource group reside on the same LAN (use the same subnet addressing). Cluster resource services supports a user configured takeover IP address when configuring application CRGs.

Address Resolution Protocol (ARP) is the network protocol that is used to switch the configured application takeover IP address from one node to another node in the recovery domain. To enable the application switchover across subnets, you need to use the virtual IP address support and the Routing Information Protocol Next Generation (RIPng) for IPv6.

The following manual configuration steps are required to enable the switchover environment. **This set of instructions must be done on all the nodes in the recovery domain, and repeated for the other nodes in the cluster that will become nodes in the recovery domain for the given application CRG.**

1. Select an IPv6 takeover IP address to be used by the application CRG.

- To avoid confusion, this address should not overlap with any other existing addresses used by the cluster nodes or routers.
- It is recommended that this address is defined with a shorter IPv6 address prefix than any other IPv6 address that shares the same IPv6 prefix to ensure that the correct address is chosen for the source address in outbound packets.
- Add the takeover interface (for example, 2001:0DB8:1234::1. Create it with a line description of *VIRTUALIP, maximum transmission unit of 1500 (any number in the range 576-16388), and autostart of *NO.

2. Create the cluster and create any CRGs. For the application CRG, specify `QcstUserCfgrTakeoverIpAddr` for the **Configure takeover IP address** field. Do not start any application CRGs.
3. Use the Change RIP Attributes (CHGRIPA) command to set the RIPng attributes. Run the command: `CHGRIPA AUTOSTART(*YES) IP6COND(*NEVER) IP6ACPDFT(*NO) IP6SNDONLY(*VIRTUAL)`.
4. Ensure there is an IPv6 link-local address active on the system. An IPv6 link-local address starts with 'fe80:'.
5. Use the Add RIP Interface (ADDRIPIFC) command add a RIP interface used by the OMPROUTED server to advertise the virtual address used for the takeover IP address. For example, if fe80::1 is the active IPv6 link-local address, run the command: `ADDRIPIFC IFC('fe80::1') RCVDYNNET(*YES) SNDSTTRTE(*YES) SNDHOSTRTE(*YES) SNDONLY(*VIRTUAL)`.
6. Restart the OMPROUTED server using the following commands:
 - a) `ENDTCPSVR SERVER(*OMPROUTED) INSTANCE(*RIP)`
 - b) `STRTCPSVR SERVER(*OMPROUTED) INSTANCE(*RIP)`
7. Ensure that all the commercial routers in the network that interconnect the recovery domain LANs are accepting and advertising host routes for RIPng.
 - This is not necessarily the default setting for routers. The language varies with router manufacturer, but the RIPng interfaces settings should be set to send host routes and receive dynamic hosts.
 - This also applies to both the router interfaces that point to the systems as well as the router-to-router interfaces.

Note: Do not use an IBM i machine as the router in this configuration. Use a commercial router (IBM or otherwise) that is designed for routing purposes. IBM i routing cannot be configured to handle this function.
8. Manually activate the takeover address on one of the cluster nodes:
 - a) Wait up to 5 minutes for RIP to propagate the routes.
 - b) Ping the takeover address from all nodes in the CRG recovery domain and from selected clients on the LANs who will be using this address.
 - c) Ensure that the takeover address is ended again.
(Clustering will start the address on the specified primary node when the CRGs are started.)
9. Start the application CRGs.
 - The takeover address is started by clustering on the specified, preferred node, and RIPng advertises the routes throughout the recovery domain. RIPng might take up to 5 minutes to update routes across the domain. The RIPng function is independent from the start CRG function.

Important:

- If the above procedure is not followed for all cluster nodes in the application CRG recovery domain, the cluster hangs during the switchover process.
- Even though you do not perform a failover to replica nodes, it is a good idea to perform the procedure on the replica nodes in the event that they might be changed at a later date in time to become a backup.
- If you want to use multiple virtual IP addresses, then each one will require a separate application CRG and a separate IP address with which to be associated. This address may be another logical IP address on the same physical adapter or it may be another physical adapter altogether. Also, care must be taken to prevent ambiguities in the routing tables. This is best achieved by doing the following:
 - Add a *DFTRROUTE to the routing table for each virtual IP address.
 - To use multiple IP address use CFGTCP (option 2).
 - Set all parameters, including the next hop, the same to reach the router of choice; however, the Preferred binding interface should be set to the local system IP address that is associated with the virtual IP address that is represented by this route.

Example: Application cluster resource group failover actions

This example shows how one failover scenario works. Other failover scenarios can work differently.

The following happens when a cluster resource group for a resilient application fails over due to exceeding the retry limit or if the job is canceled:

- The cluster resource group exit program is called on all active nodes in the recovery domain for the CRG with an action code of failover. This indicates that cluster resource services is preparing to failover the application's point of access to the first backup.
- Cluster resource services ends the takeover Internet Protocol (IP) connection on the primary node.
- Cluster resource services starts the takeover IP address on the first backup (new primary) node.
- Cluster resource services submits a job that calls the cluster resource group exit program only on the new primary node with an action code of Start. This action restarts the application.

Related concepts

Managing application CRG takeover of IP addresses

You can manage application CRG takeover of IP addresses by using cluster resource services. You can also manage them manually.

Example: Application exit program

This code example contains an application cluster resource group exit program.

You can find this code example in the QUSRTOOL library.

Sample source code is provided in the QUSRTOOL library, which can be used as a basis for writing an exit program. See the TCSTAPPEXT member in the QATTSYSC file. To view the sample source code, issue one of the following commands:

- DSPPFM QUSRTOOL/QATTSYSC TCSTAPPEXT
- STRSEU SRCFILE(QUSRTOOL/QATTSYSC) SRCMBR(TCSTAPPEXT)
- WRKMBRPDM FILE(QUSRTOOL/QATTSYSC) MBR(TCSTAPPEXT), then place a 5 in front of the member name

If you receive CPF8056, EDT0213, or PDM0308 messages when running the commands, QUSRTOOL files need to be converted from save files to source physical files. Issue this command, CALL PGM(QUSRTOOL/UNPACKAGE) PARM('*ALL ' 1).

Note: By using the code examples, you agree to the terms of the [Code license and disclaimer information](#).

Planning clusters

Before implementing a high-availability solution, you must ensure that you met all prerequisites for clusters.

Hardware requirements for clusters

To implement a high-availability solution, you need to plan and configure a cluster. A cluster groups systems and resources in a high availability environment.

The following are the minimum hardware requirements for clusters:

- You will need at least two systems or partitions running IBM i. Cluster supports up to 128 systems within a cluster. Any System i[®] model that is capable of running IBM i V4R4M0, or later, is compatible for using clustering.
- External uninterruptible power supply or equivalent is recommended to protect from a sudden power loss which could result in a cluster partition.
- If you plan to use data resiliency technologies that require independent disk pools, you will also need to plan for hardware specific to your chosen data resiliency technology. You can also use different methods of disk protection to prevent failover from occurring should a protected disk fail.

Related concepts

Planning data resiliency

Data resilience is the ability for data to be available to users or applications. You can achieve data resiliency by using IBM i cluster technology with PowerHA technologies or logical replication technologies.

Related reference

[Planning checklist for clusters](#)

Complete the cluster configuration checklist to ensure that your environment is prepared properly before you begin to configure your cluster.

Related information

[Uninterruptible power supply](#)

[IP multicasting](#)

[Disk protection](#)

Software requirements for clusters

In order to use clustering, you must have the correct software and licenses.

1. Latest supported release of IBM i operating system installed.
2. TCP/IP Connectivity Utilities feature installed.
3. If you plan to use data resiliency technologies, there are additional requirements.
4. Option 41 (High Availability Switchable Resources) is required if you plan to use the following interfaces:
 - IBM PowerHA SystemMirror for i licensed program. This licensed program provides the following interfaces which require Option 41:
 - PowerHA graphical interface
 - PowerHA commands
 - PowerHA APIs
5. You can also use IBM Business Partner product or write your own high availability management application by using Cluster APIs.

Related concepts

[Determine site configuration](#)

PowerHA technologies provide several IBM i disaster recovery and high availability technologies.

[Planning data resiliency](#)

Data resilience is the ability for data to be available to users or applications. You can achieve data resiliency by using IBM i cluster technology with PowerHA technologies or logical replication technologies.

Related reference

[Planning checklist for clusters](#)

Complete the cluster configuration checklist to ensure that your environment is prepared properly before you begin to configure your cluster.

Related information

[Cluster APIs](#)

Communications requirements for clusters

Use any type of communications media in your clustering environment as long as it supports Internet Protocol (IP).

Cluster resource services uses TCP/IP and UDP/IP protocols to communicate between nodes. Local area networks (LANs), wide area networks (WANs), OptiConnect system area networks (SANs), or any combination of these connectivity devices are supported. Your choice should be based on the following factors:

- Volume of transactions
- Response time requirements

- Distance between the nodes
- Cost considerations

You can use these same considerations when determining the connection media to be used to connect primary and backup locations of resources. When planning your cluster, it is recommended that you designate one or more of your backup nodes in remote locations in order to survive a site loss disaster.

To avoid performance problems that might be caused by inadequate capacity, you need to evaluate the communication media that is used to handle the volumes of information that are sent from node to node. You can choose which physical media you prefer to use such as token ring, Ethernet, asynchronous transfer mode (ATM), SPD OptiConnect, or Virtual OptiConnect (a high-speed internal connection between logical partitions).

Related reference

Planning checklist for clusters

Complete the cluster configuration checklist to ensure that your environment is prepared properly before you begin to configure your cluster.

Dedicate a network for clusters

During normal operations, base clustering communication traffic is minimal. It is, however, highly recommended that you have redundant communication paths configured for each node in a cluster.

A redundant communications path means that you have two lines configured between two nodes in a cluster. If a failure on the first communication path occurs, the second communication path can take over to keep communications running between the nodes, thereby minimizing conditions that can put one or more nodes of the cluster into a cluster partition. One thing to consider when configuring these paths is that if both of your communications lines go into the same adapter on the system, these lines are still at risk if this single adapter fails. However, it should be noted that a cluster partition is not always avoidable. If your system experiences a power loss or if a hardware failure occurs, the cluster can become partitioned. By configuring two lines, you can dedicate one line for clustering traffic and the other line for the normal traffic and also for the backup line if the dedicated line for clustering goes down. The typical network-related cluster partition can best be avoided by configuring redundant communications paths between all nodes in the cluster.

Tips: Cluster communications

Consider these tips when you set up your communications paths.

- Be sure you have adequate bandwidth on your communication lines to handle the non cluster activity along with the clustering heartbeat function and continue to monitor for increased activity.
- For best reliability, do not configure a single communication path linking one or more nodes.
- Do not overburden the line that is responsible for ensuring that you are still communicating with a node.
- Eliminate as many single points of failure as possible, such as having two communication lines coming into a single adapter, same input-output processor (IOP), or same expansion unit.
- If you have an extremely high volume of data being passed over your communication lines, you may want to consider putting data replication and heartbeat monitoring on separate networks.
- User Datagram Protocol (UDP) multicast is the preferred protocol that the cluster communications infrastructure uses to send cluster management information between nodes in a cluster. When the physical media supports multicast capabilities, cluster communications uses the UDP multicast to send management messaging from a given node to all local cluster nodes that support the same subnet address. Messages that are sent to nodes on remote networks are always sent by using UDP point-to-point capabilities. Cluster communications does not rely on routing capability for multicast messages.
- The multicast traffic that supports cluster management messaging tends to fluctuate by nature. Depending on the number of nodes on a given LAN (that supports a common subnet address) and the complexity of the cluster management structure that is chosen by the cluster administrator, cluster-related multicast packets can easily exceed 40 packets per second. Fluctuations of this nature can have a negative effect on older networking equipment. One example is congestion problems on devices on the LAN that serve as Simple Network Management Protocol (SNMP) agents that need to evaluate every UDP multicast packet. Some of the earlier networking equipment does not have adequate bandwidth to

keep up with this type of traffic. You need to ensure that you or the network administrator has reviewed the capacity of the networks to handle UDP multicast traffic to make certain that clustering does not have a negative effect on the performance of the networks.

Performance planning for clusters

Since potentially significant differences exist in your communications environment, you have the capability to adjust variables that affect cluster communications to best match your environment.

The default values should normally be acceptable to most common environments. If your particular environment is not well suited for these defaults, you can tune cluster communications to better match your environment. Base-level tuning and advanced level tuning are available.

Base-level tuning

Base-level tuning allows you to set the tuning parameters to a predefined set of values identified for high, low, and normal timeout and messaging interval values. When the normal level is selected, the default values are used for cluster communications performance and configuration parameters. Selecting the low level causes clustering to increase the heartbeating interval and the various message timeout values. With fewer heartbeats and longer timeout values, the cluster is less sensitive to communications failures. Selecting the high level causes clustering to decrease the heartbeating interval and the various message timeout values. With more frequent heartbeats and shorter timeout values, the cluster is more sensitive to communications failures.

Advanced-level tuning

With advanced-level tuning, individual parameters can be tuned by using a predefined ranges of values. This allows more granular tuning to meet any special circumstances in your communications environment. If an advanced level of tuning is desired, it is recommended that you obtain help from IBM support personnel or equivalent. Setting the individual parameters incorrectly can easily result in decreased performance.

Tunable cluster communications parameters

The `Change Cluster Resource Services (QcstChgClusterResourceServices)` API enables some of the cluster topology services and cluster communications performance and configuration parameters to be tuned to better suit the many unique application and networking environments in which clustering occurs.

The **Change Cluster (CHGCLU)** command provides a base level of tuning, while the `QcstChgClusterResourceServices` API provides both base and advanced levels of tuning.

The `QcstChgClusterResourceServices` API and **Change Cluster Configuration (CHGCLUFCG)** command can be used to tune cluster performance and configuration. The API and command provide a base level of tuning support where the cluster will adjust to a predefined set of values identified for high, low, and normal timeout and messaging interval values. If an advanced level of tuning is desired, usually anticipated with the help of IBM support personnel, then individual parameters can be tuned through the use of the API over a predefined value range. Inappropriate changes to the individual parameters can easily lead to degraded cluster performance.

When and how to tune cluster parameters

The **CHGCLU** command and the `QcstChgClusterResourceServices` API provide for a fast path to setting cluster performance and configuration parameters without your needing to understand the details. This base level of tuning primarily affects the heartbeating sensitivity and the cluster message timeout values. The valid values for the base level of tuning support are:

1 (High Timeout Values/Less Frequent Heartbeats)

Adjustments are made to cluster communications to decrease the heartbeating frequency and increase the various message timeout values. With fewer heartbeats and longer timeout values, the cluster will be slower to respond (less sensitive) to communications failures.

2 (Default Values)

Normal default values are used for cluster communications performance and configuration parameters. This setting can be used to return all parameters to the original default values.

3 (Low Timeout Values/More Frequent Heartbeats)

Adjustments are made to cluster communications to decrease the heartbeating interval and decrease the various message timeout values. With more frequent heartbeats and shorter timeout values, the cluster is quicker to respond (more sensitive) to communications failures.

Example response times are shown in the following table for a heartbeat failure leading to a node partition:

Note: Times are specified in minutes:seconds format.

	1 (Less sensitive)			2 (Default)			3 (More sensitive)		
	Detection of Heartbeat at Problem	Analysis	Total	Detection of Heartbeat at Problem	Analysis	Total	Detection of Heartbeat at Problem	Analysis	Total
Single subnet	00:24	01:02	01:26	00:12	00:30	00:42	00:04	00:14	00:18
Multiple subnets	00:24	08:30	08:54	00:12	04:14	04:26	00:04	02:02	02:06

Depending on typical network loads and specific physical media being used, a cluster administrator might choose to adjust the heartbeating sensitivity and message timeout levels. For example, with a high speed high-reliability transport, such as OptiConnect with all systems in the cluster on a common OptiConnect bus, one might desire to establish a more sensitive environment to ensure quick detection leading to faster failover. Option 3 is chosen. If one were running on a heavily loaded 10 Mbs Ethernet bus and the default settings were leading to occasional partitions just due to network peak loads, option 1 could be chosen to reduce clustering sensitivity to the peak loads.

The Change Cluster Resource Services API also allows for tuning of specific individual parameters where the network environmental requirements present unique situations. For example, consider again a cluster with all nodes common on an OptiConnect bus. Performance of cluster messages can be greatly enhanced by setting the message fragment size parameter to the maximum 32,500 bytes to better match the OptiConnect maximum transmission unit (MTU) size than does the default 1,464 bytes. This reduces the overhead of fragmentation and reassembly of large messages. The benefit, of course, depends on the cluster applications and usage of cluster messaging resulting from those applications. Other parameters are defined in the API documentation and can be used to tune either the performance of cluster messaging or change the sensitivity of the cluster to partitioning.

Related reference

[QcstChgClusterResourceServices API](#)

Related information

[Change Cluster \(CHGCLU\) command](#)

Changing cluster resource services settings

The default values affecting message timeout and retry are set to account for most typical installations. However, it is possible to change these values to more closely match your communications environment.

The values can be adjusted in one of these ways:

- Set a general performance level that matches your environment.
- Set values for specific message tuning parameters for more specific adjustment

In the first method, the message traffic is adjusted to one of three communications levels. The normal level is the default and is described in detail in Heartbeat monitoring.

The second method should normally be done only under the advisement of an expert.

The Change Cluster Resource Services (QcstChgClusterResourceServices) API describes details on both methods.

Related reference

[QcstChgClusterResourceServices API](#)

Related information

[Heartbeat monitoring](#)

Planning multiple-release clusters

If you are creating a cluster that includes nodes at multiple cluster versions, then certain steps are required when you create the cluster.

By default, the current cluster version is set to the potential cluster version of the first node added to the cluster. This approach is appropriate if this node is at the lowest version level to be in the cluster. However, if this node is at a later version level, then you cannot add nodes with a lower version level. The alternative is to use the target cluster version value when you create a cluster to set the current cluster version to one less than the potential cluster version of the first node added to the cluster.

Note: If you are using the IBM PowerHA SystemMirror for i licensed program, it is required on all systems within the cluster.

For example, consider the case where a two-node cluster is to be created. The nodes for this cluster follow:

Node identifier	Release	Potential cluster version
Node A	V7R1	7
Node B	V7R2	8

If the cluster is to be created from Node B, care must be taken to indicate that this will be a mixed-release cluster. The target cluster version must be set to indicate that the nodes of the cluster will communicate at one less than the requesting node's potential node version.

Performance planning for clusters

When changes are made to a cluster, the overhead necessary to manage the cluster can be affected.

The only resources that clustering requires are those necessary to perform heartbeat monitoring, to manage the cluster resource groups and the cluster nodes, and to handle any messaging taking place between cluster resource groups and cluster nodes. After your clustering environment is operational, the only increase in overhead is if you make changes to the cluster.

During a normal operating environment, clustering activity should have minimal effect on your clustered systems.

Planning advanced node failure detection

Advanced node failure detection function can be used to reduce the number of failure scenarios which result in cluster partitions.

Before implementing advanced node failure detection, you must ensure that you have met all the prerequisites.

- To prevent cluster partitions when a cluster node has actually failed, a [Hardware Management Console \(HMC\) v7 or Virtual I/O Server \(VIOS\) partition on an Integrated Virtualization Manager \(IVM\) managed server](#) can be used.
- For each cluster node, determine the type of server that the cluster monitor will be using to monitor that node. If the node is managed either by an IVM or by an HMC at version V8R8.5.0 or earlier, then a common information model (CIM) server can be used. If the node is managed by an HMC at version V8R8.5.0, then a representational state transfer (REST) server can be used.

Note: For cluster nodes managed by HMCs at version V8R8.5.0, you may choose to use either a CIM server or a REST server. The REST server is recommended because CIM servers are not supported for later versions of HMC.

- Each node that is to have failures reported to it, will need to have a cluster monitor configured.

Hardware requirements for the advanced node failure detection

Advanced node failure detection feature can be used provided all the hardware requirements are met.

The following minimum hardware requirements are needed for the advanced node failure detection feature:

- At least two IBM i models or logical partitions
- Either Hardware Management Console (HMC) or a Virtual I/O Server (VIOS) on an Integrated Virtualization Manager (IVM) managed server

Software requirements for the advanced node failure detection

To use the advanced node failure detection function in an IBM i high-availability solution, the minimum software requirements should be met.

Each node planning to use the advanced node failure detection feature with common information model (CIM) servers has the following software requirements:

- 5770-SS1 Base operating system option 33 - Portable Application Solutions Environment
- 5770-SS1 Base operating system option 30 - Qshell
- 5733-SC1 - IBM Portable Utilities for IBM i
- 5733-SC1 option 1 - OpenSSH, OpenSSL, zlib
- 5770-UME IBM Universal Manageability Enablement
- 5770-HAS IBM PowerHA for i LP
- HMC version V8R8.5.0 or earlier. This is the last version of HMC to support the CIM server.

Note: IVM uses the CIM server.

Each node planning to use the advanced node failure detection feature with representational state transfer (REST) servers has the following software requirements:

- 5770-SS1 Base operating system option 3 - Extended Base Directory Support
- 5770-SS1 Base operating system option 33 - Portable Application Solutions Environment
- 5733-SC1 - IBM Portable Utilities for IBM i (Only required for initial configuration of a cluster monitor.)
- 5733-SC1 option 1 - OpenSSH, OpenSSL, zlib (Only required for initial configuration of a cluster monitor.)
- 5770-HAS IBM PowerHA for i LP
- HMC version V8R8.5.0 or later. This is the first version of HMC to support the REST server.
- PowerHA for i new function cluster monitor HMC REST support PTF

Planning checklist for clusters

Complete the cluster configuration checklist to ensure that your environment is prepared properly before you begin to configure your cluster.

TCP/IP requirements	
---	Start TCP/IP on every node you plan to include in the cluster using the Start TCP/IP (STRTCP) Command .
---	Configure the TCP loopback address (127.0.0.1) and verify that it shows a status of Active. Verify the TCP/IP loopback address by using the Work with TCP/IP Network Status (WRKTCPSTS) Command on every node in the cluster.

Table 1. TCP/IP configuration checklist for clusters (continued)

TCP/IP requirements	
---	Verify that the IP addresses used for clustering on a node have a status of Active. Use the Work with TCP/IP Network Status (WRKTCPSTS) Command to check the status of the IP addresses.
---	Verify that the Internet Daemon (INETD) server is active on all nodes in the cluster. If INETD server is not active, you need to start the INETD server. For information about how to start INETD server, see “Starting the INETD server” on page 49.
---	Verify that user profile for INETD, which is specified in /QIBM/ProdData/OS400/INETD/inetd.conf, does not have more than minimal authority. If this user profile has more than minimal authority, starting cluster node will fail. By default, QUSER is specified as the user profile for INETD.
---	Verify every cluster IP address on every node in the cluster can route to and send UDP datagrams to every other IP address in the cluster. If any cluster node uses an IPv4 address, then every node in the cluster needs to have an active IPv4 address (not necessarily configured as a Cluster IP address) that can route to and send TCP packets to that address. Also, if any cluster node uses an IPv6 address, then every node in the cluster needs to have an active IPv6 address (not necessarily configured as a Cluster IP address) that can route to and send TCP packets to that address. Use the PING command, specifying a local IP address, and the TRACEROUTE command, specifying UDP messages can be useful in determining if two IP addresses can communicate. PING and TRACEROUTE do not work between IPv4 and IPv6 addresses, or if a firewall is blocking PING and TRACEROUTE .
---	Verify that ports 5550 and 5551 are not being used by other applications. These ports are reserved for IBM clustering. Port usage can be viewed by using the Work with TCP/IP Network Status (WRKTCPSTS) command. Port 5550 is opened and is in a Listen state by clustering after INETD is started.

Table 2. Administrative domain checklist for clusters

Cluster resource services cluster interface considerations	
	Install IBM PowerHA SystemMirror for i (iHASM licensed program (5770-HAS). A valid license key must exist on all cluster nodes that will be in the high-availability solution.
---	Install <u>Option 41 (IBM i - HA Switchable Resources)</u> . A valid license key must exist on all cluster nodes that will be in the device domain.
---	Verify that all host servers are started by using the Start Host Server (STRHOSTSVR) Command: STRHOSTSVR SERVER(*ALL)

Table 3. Security configuration checklist for clusters

Security requirements	
---	Set the Allow Add to Cluster (ALWADDCLU) network attribute appropriately on the target node if you are trying to start a remote node. This should be set to *ANY or *RQSAUT depending on your environment. If this attribute is set to *RQSAUT, then IBM i option 34 (Digital Certificate Manager) and the CCA Cryptographic Service Provider (Option 35) must be installed. See <u>Enable a node to be added to a cluster</u> for details on setting the ALWADDCLU network attribute.
---	Enable the status of the user profile for INETD specified in /QIBM/ProdData/OS400/INETD/inetd.conf. It must not have *SECADM or *ALLOBJ special authorities. By default, QUSER is specified as the user profile for INETD.
---	Verify that the user profile that calls the cluster resource services APIs exists on all cluster nodes and has *IOSYSCFG authority.
---	Verify that the user profile to run the exit program for a cluster resource group (CRG) exists on all recovery domain nodes.

Table 4. Job configuration checklist for clusters

Job considerations	
---	Jobs can be submitted by the cluster resource services APIs to process requests. The jobs either run under the user profile to run the exit program specified when creating a cluster resource group, or under the user profile that requested the API (for varying on devices in resilient device CRGs only). Ensure that the subsystem that services the job queue associated with the user profile is configured as: *NOMAX for the number of jobs it can run from that job queue.
---	Jobs are submitted to the job queue specified by the job description that is obtained from the user profile defined for a CRG. The default job description causes the jobs to be sent to the QBATCH job queue. Because this job queue is used for many user jobs, the exit program job might not run in a timely fashion. Consider a unique job description with a unique user queue.
---	When exit program jobs are run, they use routing data from the job description to choose which main storage pool and run time attributes to use. The default values result in jobs that are run in a pool with other batch jobs that have a run priority of 50. Neither of these may produce the desired performance for exit program jobs. The subsystem initiating the exit program jobs (the same subsystem that is using the unique job queue) should assign the exit program jobs to a pool that is not used by other jobs initiated by the same subsystem or other subsystems. In addition, the exit program jobs should be assigned a run priority of 15 so that they run before almost all other user jobs.
---	Set the <u>QMLTTHDACN</u> system value to 1 or 2.

There are several software interfaces available for configuring and managing your cluster. One of these interfaces is PowerHA graphical interface. If you choose to use PowerHA, the following requirements must be satisfied.

Table 5. PowerHA configuration checklist for clusters

PowerHA graphical interface considerations	
---	Install IBM PowerHA SystemMirror for i licensed program. A valid license key must exist on all cluster nodes that will be in the high-availability solution.
---	Install <u>Option 41 (HA Switchable Resources)</u> . A valid license key must exist on all cluster nodes that will be in the device domain.
---	Verify that all host servers are started by using the Start Host Server (STRHOSTSVR) command: STRHOSTSVR SERVER(*ALL)
---	Verify that the Administration Server is started by using the Start TCP/IP Server (STRTCPSVR) command: STRTCPSVR SERVER(*HTTP) HTTPSVR(*ADMIN)

Table 6. Advanced node failure detection checklist for clusters

Advanced node failure detection considerations when using a CIM server or IVM	
---	Determine which cluster nodes are or can be managed with a Hardware Management Console (HMC) or a Virtual I/O Server (VIOS) partition on an Integrated Virtualization Manager (IVM) managed server
---	Determine which cluster node(s) are to receive messages when some other cluster node fails
---	On each cluster node that is to receive a message from an HMC or IVM, the following things must be done.
	Install base operating system option 33 - IBM Portable Application Solutions Environment for i
	Install 5733-SC1 - IBM Portable Utilities for i

Table 6. Advanced node failure detection checklist for clusters (continued)

Advanced node failure detection considerations when using a CIM server or IVM	
	Install 5733-SC1 option 1 - OpenSSH, OpenSSL, zlib
	Install 5770-UME - IBM Universal Manageability Enablement for i.
	Configure the enableAuthentication and sslClientVerificationMode properties for the 5770-UME product.
	Copy a digital certificate file from the HMC or VIOS and add it to an IBM i truststore.
	Start the *CIMOM server with STRTCPSVR *CIMOM CL command
	Configure the cluster monitor(s) with the ADDCLUMON CL command
Advanced node failure detection considerations when using a REST server	
---	Determine which cluster nodes are or can be managed with a Hardware Management Console (HMC) REST server
---	Determine which cluster node(s) are to receive messages when some other cluster node fails
---	On each cluster node that is to receive a message from an HMC, the following things must be done.
	Install base operating system option 3 - Extended Base Directory Support.
	Install base operating system option 33 - IBM Portable Application Solutions Environment for i
	Install 5733-SC1 - IBM Portable Utilities for i (Only required for initial configuration of a cluster monitor.)
	Install 5733-SC1 option 1 - OpenSSH, OpenSSL, zlib (Only required for initial configuration of a cluster monitor.)
	Copy a digital certificate file from the HMC and add it to an IBM i truststore.
	Configure the cluster monitor(s) with the ADDCLUMON CL command

Planning data resiliency

Data resiliency is the ability for data to be available to users or applications. You can achieve data resiliency by using IBM i cluster technology with PowerHA technologies or logical replication technologies.

For IBM i supported implementations of data resiliency, you have several choices of technologies. When these technologies are combined with IBM i cluster resource services, you can build a complete high-availability solution. These technologies can be categorized this way:

IBM i Independent disk pool technologies

These technologies are all based on IBM i implementation of independent disk pools. For high availability that uses independent disk pool technologies, it is required that all data that needs to be resilient be stored in an independent disk pool. In many cases, this requires migrating data to independent disk pools. This information assumes that migration of data has been completed.

The following IBM i supported technologies are based on independent disk pools:

- Switched logical units (LUNs)
- Geographic mirroring
- Metro Mirror
- Global Mirror
- HyperSwap with independent auxiliary storage pools (IASPs)

DS8000 Full System HyperSwap technology

DS8000 Full System HyperSwap is an IBM System Storage® technology, which has been integrated into PowerHA to provide a minimal downtime high availability solution for storage outages (seconds to minutes). Full System HyperSwap can also be combined with live partition mobility to define and implement affinity between the POWER® server hosting the IBM i logical partition and the storage server hosting the data.

Logical replication technologies

Logical replication is a journal-based technology, where data is replicated to another system in real time. Logical replication technologies use IBM i cluster resource services and journaling with IBM Business Partner applications. These solutions require a high availability business partner application to configure and manage the environment. This information does not provide specific requirements for these IBM Business Partner solutions. If you are implementing a logical replication solution for high availability, consult information related to application or contact a service representative.

Determine which data should be made resilient

Understand what types of data you should consider making resilient.

Determining which data you need to make resilient is similar to determining which kind of data you need to back up and save when you prepare a back up and recovery strategy for your systems. You need to determine which data in your environment is critical to keeping your business up and running.

For example, if you are running a business on the Web, your critical data can be:

- Today's orders
- Inventory
- Customer records

In general, information that does not change often or that you do not need to use on a daily basis probably does not need to be made resilient.

Determine site configuration

PowerHA technologies provide several IBM i disaster recovery and high availability technologies.

For PowerHA technologies, which involve two copies of the data, such as Geographic mirroring, Metro Mirror, Global Mirror and HyperSwap, there are considerations for planning the sites where the data will reside. One site is typically considered the production or source site. This site contains your production data, which is mirrored or copied to the remote site. The remote site, sometime referred to as a backup or target site, contains the mirrored copy of the production data. When a site-wide disaster occurs at the production site, the backup site resumes your business with the mirrored data. Before you configure one of these technologies, consider the following regarding your site plans.

Determine which sites will be production and backup sites

Assess the current hardware and software resources that are in place at each site to determine if there are any missing components that will be necessary for a cross-site mirroring solution.

Determine the distance between production and backup sites

Depending on your communication bandwidth and other factors, distance between sites can affect performance and latency in the mirroring technology you choose. Some cross-site mirroring technologies are better suited for sites that are at great distances, while others may have performance degradation.

For more information about disk pools, see [Planning your disk pools](#).

Related tasks

[Enabling and accessing disk units](#)

Planning for DS8000 Full System HyperSwap

You can use HyperSwap as a means to help reduce or eliminate outages due to storage and SAN-related outages in high availability environments that use the external storage units of the IBM Systems Storage DS8000 devices.

Hardware requirements for DS8000 Full System HyperSwap

To use the DS8000 Full System HyperSwap technology in an PowerHA solution, ensure that the minimum hardware requirements are met.

The following minimum hardware requirements are needed for DS8000 Full System HyperSwap:

- One System i model/logical partition with all storage provided by two IBM System Storage DS8000 external storage units that are attached to the system.
- Two supported IBM Systems Storage DS8000. Supported IBM System Storage units:
 - DS8700 or DS8800 (with release 6.3 SP7)
 - DS8870 with release 7.2
- Appropriate disk sizing for the system storage should be completed before any configuration. You need one set of disks on the source DS8000 unit, and an equivalent sized set of disks on the target DS8000 unit.

For information on storage technologies that are provided by IBM, see [“PowerHA supported storage servers”](#) on page 37.

Software requirements for DS8000 Full System HyperSwap

To use the DS8000 Full System HyperSwap technology in an IBM i high-availability solution, the minimum software requirements must be met.

The DS8000 Full System HyperSwap feature has the following minimum software requirements:


- IBM i model must be at IBM i 7.2.
- IBM PowerHA for i Express Edition licensed program must be installed.

Communications requirements for DS8000 Full System HyperSwap

To use the DS8000 HyperSwap technology, you must be using or planning to use a storage area network (SAN).

A SAN is a dedicated, centrally managed, secure information infrastructure that enables any-to-any interconnection between systems and storage systems. SAN connectivity is required for using IBM System Storage DS8000 external storage units.

IBM i product supports various SAN switches and directors. Refer to the IBM System Storage Interoperation Center (SSIC) website, for a complete list of supported adapters, switches and directors.

See the [IBM PowerHA SystemMirror for i wiki](#)  for the specific storage version numbers needed to work on IBM i.

At least one POWER model running IBM i attached to an IBM System Storage device. For a FlashCopy to a IBM i target system or partition and the FlashCopy target is not *NONE), both the FlashCopy source and FlashCopy target systems or partitions must be attached to the same external storage device.

Related information

[IBM System Storage Interoperation Center \(SSIC\)](#)

Performance requirements for DS8000 Full System HyperSwap

You should understand these performance considerations before configuring HyperSwap.

Before you use DS8000 HyperSwap, consider the following requirements and guidelines:

- The source and target volumes in a HyperSwap relationship must be the same storage type.
- The source and target logical volumes must be the same size.

- As with any synchronous replication technology, there are distance and bandwidth limitations that can affect performance.
- Similar to any system disk configuration, the number of disk units available to the application can have a significant effect on its performance.
- Putting more workload on a limited number of disk units may result in longer disk waits and ultimately longer response times to the application.

Planning geographic mirroring

Geographic mirroring is a sub-function of cross-site mirroring. This technology provides disaster recovery and high availability in IBM i environments.

Hardware requirements for geographic mirroring

If you plan to use geographic mirroring for IBM i high availability, ensure that the minimum hardware requirements are met.

- All independent disk pool hardware requirements must be met.
- At least two IBM i models, which can be separated geographically, are required.
- At least two sets of disks at each site that are roughly of similar capacity are required.
- A separate storage pool for jobs using geographic mirrored independent disk pools should be configured. Performing geographic mirroring from the main storage pool can cause the system to hang under extreme load conditions.
- Geographic mirroring is performed when the disk pool is available. When geographic mirroring is being performed, the system value for the time of day (QTIME) should not be changed.
- Communications requirements for independent disk pools are critical because they affect throughput.
- Geographic mirroring traffic is dispersed round robin across the potentially multiple communication lines available to it. It is recommended that if multiple lines are provided for geographic mirroring, that those lines be of similar speed and capacity.
- It is recommended that a separate communication line be used for the clustering heartbeat to prevent contention with the geographic mirroring traffic.

Software requirements for geographic mirroring

If you plan to use geographic mirroring as part of an IBM i high availability solution, the following software is required.

- To use advanced features of geographic mirroring, IBM PowerHA SystemMirror for i license program must be installed.
- To use new and enhanced functions and features of this technology, it is recommended that you install the most current release and version of the operating system on each system or logical partition that is participating in a high-availability solution that is based on this technology. If the production system and backup system are at different operating system releases, it is required that the backup system be at the more current release. This scenario can be used to upgrade a production environment to the next release.

Note: Once the IASP has been switched and varied on to the newer release system, it cannot be switched back until the new target system has also been upgraded.

- To perform some of the disk management tasks necessary to implement independent disk pools, use the Configure Device ASP (CFGDEVASP) command, or IBM Navigator for i.
- You need to install IBM i Option 41 HA Switchable Resources. Option 41 gives you the capability to switch independent disk pools between systems. To switch an independent disk pool between systems, the systems must be members of a cluster and the independent switched disk must be associated with a device cluster resource group in that cluster. Option 41 is also required for working with high availability management interfaces, which are provided as part of the IBM PowerHA SystemMirror for i licensed program.

Related information

[High Availability and Clusters](#)

Communications requirements for geographic mirroring

When you are implementing an IBM i high-availability solution that uses geographic mirroring, you should plan communication lines so that geographic mirroring traffic does not adversely affect system performance.

The following is recommended:

- Geographic mirroring can generate heavy communications traffic. If geographic mirroring shares the same IP connection with another application, for example clustering, then geographic mirroring might be suspended, which results in synchronization. Likewise, clustering response might be unacceptable, which results in partitioned nodes. Geographic mirroring should have its own dedicated communication lines. Without its own communication line, geographic mirroring can contend with other applications that use the same communication line and affect user network performance and throughput. This also includes the ability to negatively affect cluster heartbeat monitoring, resulting in a cluster partition state. Therefore, it is recommended that you have dedicated communication lines for both geographic mirroring and clusters. Geographic mirroring supports up to four communications lines.

Geographic mirroring distributes changes over multiple lines for optimal performance. The data is sent on each of the configured communication lines in turn, from 1 to 4, over and over again. Four communication lines allow for the highest performance, but you can obtain relatively good performance with two lines.

If you use more than one communication line between the nodes for geographic mirroring, it is best to separate those lines into different subnets, so that the usage of those lines is balanced on both systems.

- If your configuration is such that multiple applications or services require the use of the same communication line, some of these problems can be alleviated by implementing Quality of Service (QoS) through the TCP/IP functions of IBM i. The IBM i quality of service (QoS) solution enables the policies to request network priority and bandwidth for TCP/IP applications throughout the network.
- Ensure that throughput for each data port connection matches. This means that the speed and connection type should be the same for all connections between system pairs. If throughput is different, performance will be gated by the slowest connection.
- Consider the delivery method for a geographic mirroring ASP session. Before 7.1, the mirroring uses synchronous communication between the production and mirror copy systems. This delivery method is best for low latency environments. In 7.1, asynchronous support was added, which means asynchronous communications is used between the production and mirror copy systems. This delivery method is best for high latency environments. This delivery method will consume more system resources on the production copy node than synchronous delivery.
- Consider configuring a virtual private network for TCP/IP connections for the following advantages:
 - Security of data transmission by encrypting the data
 - Increased reliability of data transmission by sending greater redundancy

Related reference

[Quality of Service \(QoS\)](#)

Journal planning for geographic mirroring

When implementing high availability based on IBM i geographic mirroring, you should plan for journal management.

Journal management prevents transactions from being lost if your system ends abnormally. When you journal an object, the system keeps a record of the changes you make to that object. Regardless of the high availability solution that you implement, journaling is considered a best practice to prevent data loss during abnormal system outages.

Related information

[Journal management](#)

Backup planning for geographic mirroring

Before implementing high availability based on geographic mirroring, you should understand and plan a backup strategy within this environment.

Before configuring any high-availability solution, assess your current backup strategy and make appropriate changes if necessary. Geographic mirroring does not allow concurrent access to the mirror copy of the independent disk pool, which has implications to performing remote backups. If you want to back up from the geographically mirrored copy, you must quiesce mirroring on the production system and suspend the mirrored copy with tracking enabled. Tracking allows for changes on the production to be tracked so that they can be synchronized when the mirrored copy comes back online. Then you must vary on the suspended "mirror" copy of the independent disk pool, perform the backup procedure, vary off "the suspended mirror copy" and then resume the independent disk pool to the original production host. This process only requires "partial data resynchronization" between the production and mirrored copies.

Your system is running exposed while doing the backups and when synchronization is occurring. It is also recommended that you suspend mirroring with tracking enabled, which speeds up the synchronization process. Synchronization is also required for any persistent transmission interruption, such as the loss of all communication paths between the source and target systems for an extended period of time. You can also use redundant communication paths to help eliminate some of those risks associated with a communication failure.

It is recommended that you should also use geographic mirroring in at least a three system, or logical partitions, where the production copy of the independent disk pool can be switched to another system at the same site that can maintain geographic mirroring.

Related concepts

Scenario: [Performing backups in geographic mirroring environment](#)

This scenario provides an overview of tasks that are necessary when performing a remote backup in a PowerHA solution that uses geographic mirroring.

Performance planning for geographic mirroring

When implementing a geographic mirroring solution, you need to understand and plan your environment to minimize potential effects on performance.

A variety of factors can influence the performance of geographic mirroring. The following factors provide general planning considerations for maximizing performance in a geographic mirroring environment:

CPU considerations

Geographic mirroring increases the CPU load, so there must be sufficient excess CPU capacity. You might require additional processors to increase CPU capacity. As a general rule, the partitions you are using to run geographic mirroring need more than a partial processor. In a minimal CPU configuration, you can potentially see 5 - 20% CPU overhead while running geographic mirroring. If your backup system has fewer processors in comparison to your production system and there are many write operations, CPU overhead might be noticeable and affect performance.

Base pool size considerations

If asynchronous delivery transmission is used for geographic mirroring, it may be necessary to also increase the amount of storage in the base pool of the system. The amount to increase the base pool by depends primarily on the amount of latency which occurs due to the distance between the two systems. Larger amounts of latency will require larger amounts of the base pool.

Machine pool size considerations

For optimal performance of geographic mirroring, particularly during synchronization, increase your machine pool size by at least the amount given by the following formula:

- The amount of extra machine pool storage is: $300 \text{ MB} + .3\text{MB} \times \text{the number of disk ARMs in the independent disk pool}$. The following examples show the additional machine pool storage needed for independent disk pools with 90 disk ARMs and a 180 disk ARMs, respectively:

- $300 + (.3 \times 90 \text{ ARMs}) = 327 \text{ MB}$ of additional machine pool storage
- $300 + (.3 \times 180 \text{ ARMs}) = 354 \text{ MB}$ of additional machine pool storage

The extra machine pool storage is required on all nodes in the cluster resource group (CRG) so that the target nodes have sufficient storage in case of switchover or failover. As always, the more disk units in the independent disk pool, the better the performance should be, as more things can be done in parallel.

To prevent the performance adjuster function from reducing the machine pool size, you should do one of the following:

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 Work with Shared Storage Pools (WRKSHRPOOL) command or Change Shared Storage Pool (CHGSHRPOOL) command.

Note: It is recommended to use this option with the Work with Shared Storage Pools (WRKSHRPOOL) option.
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.

Disk unit considerations

Disk unit and IOA performance can affect overall geographic mirroring performance. This is especially true when the disk subsystem is slower on the mirrored system. When geographic mirroring is in a synchronous mirroring 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 geographic mirroring in asynchronous mirroring mode. Running in asynchronous mirroring mode alleviates the wait for the disk subsystem on the target side, and sends confirmation back 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 available to the application can have a significant affect on its performance. Putting additional workload on a limited number of disk units might result in longer disk waits and ultimately longer response times to the application. This is particularly important when it comes to temporary storage in a system configured with independent disk pools. All temporary storage is written to the SYSBAS disk pool. If your application does not use a lot of temporary storage, then you can get by with fewer disk arms in the SYSBAS disk pool. You must also remember that the operating system and basic functions occur in the SYSBAS disk pool. This is also true for the mirror copy system since, in particular, the TCP messages that are sent to the mirror copy can potentially page in the system asp.

Network configuration considerations

Network cabling and configuration can potentially impact geographic mirroring performance. In addition to ensuring that network addressing is set up in different subnets for each set of data port IP addresses, network cabling and configuration should also be set up in the same manner.

Planning switched logical units (LUNs)

A single copy of the data is maintained on the logical units within the IBM® System Storage™ storage unit.

When an outage occurs on the primary node, access to the data on the switchable logical units switches to a designated backup node. Additionally, independent disk pools can be used in a cross-site mirroring (XSM) environment. This allows a mirror copy of the independent disk pool to be maintained on a system that is (optionally) geographically distant from the originating site for availability or protection purposes.

Careful planning is required if you plan to take advantage of switchable logical units residing on independent disk pools or cross-site mirroring (XSM).

You should also evaluate your current system disk configuration to determine if additional disk units may be necessary. Similar to any system disk configuration, the number of disk units available to the

application can have a significant affect on its performance. Putting additional workload on a limited number of disk units might result in longer disk waits and ultimately longer response times to the application. This is particularly important when it comes to temporary storage in a system configured with independent disk pools. All temporary storage is written to the SYSBAS disk pool. If your application does not use much temporary storage, then you can get by with fewer disk arms in the SYSBAS disk pool. You must also remember that the operating system and basic functions occur in the SYSBAS disk pool.

Before you can use IBM Navigator for i to perform any disk management tasks, such as creating an independent disk pool, you need to set up the proper authorizations for dedicated service tools (DST).

Hardware requirements for switched logical units

To configure and manage an PowerHA solution that uses switched logical units technology, you should ensure that the minimum hardware requirements are met.

The following minimum hardware requirements are recommended:

- At least two IBM i partitions or systems that are separated geographically with at least one IBM System Storage external storage unit that is attached to each system. The IBM System Storage external storage units are supported on all System i models that support Fibre Channel attachment for external storage.
- PowerHA supports switched logical units (LUNs) on DS8000, SAN Volume Controller, and Storwize®. For specific Fibre Channel adapters supported, see [IBM System Storage Interoperation Center \(SSIC\)](#).
- Appropriate disk sizing for the system storage should be completed before any configuration. With switched logical units you have one set of disks.

For more information on the storage technologies that are provided by IBM i, see [PowerHA supported storage servers](#).

Related information

[PowerHA SystemMirror for IBM i Cookbook](#)

[IBM System Storage DS8000 Information Center-> Copy Services](#)

[IBM i and IBM Storwize Family: A Practical Guide to Usage Scenarios](#)

[IBM SAN Volume Controller Information Center](#)

Software requirements for switched logical units

Before configuring an IBM i high-availability solution that uses switched logical units (LUNs), ensure that the minimum software requirements have been met.

Switched logical units have the following minimum software requirements:

- Each IBM i model with in the high-availability solution must be running at least IBM i 7.1 for use with the IBM PowerHA SystemMirror for i licensed program.
- IBM PowerHA SystemMirror for i licensed program installed on each system participating in the high-availability solution that uses switched logical units.
- You need to install IBM i Option 41 HA Switchable Resources. Option 41 gives you the capability to switch independent disk pools between systems. To switch an independent disk pool between systems, the systems must be members of a cluster and the independent switched disk must be associated with a device cluster resource group in that cluster. Option 41 is also required for working with high availability management interfaces, which are provided as part of the IBM PowerHA SystemMirror for i licensed program.
- For DS8000 System Storage Devices, the IBM PowerHA for i licensed program also requires storage command-line interface (DSCLI). DSCLI is required software for all the IBM System Storage solutions, such as FlashCopy technology, Metro Mirror, Global Mirror, there is a requirement to have DSCLI installed each of the systems or partitions participating in the high availability solution, which uses these storage solutions. DSCLI has these additional software requirements:
 - Java™ Version 1.4
 - Option 35 (CCA Cryptographic Service Provider) installed on each system or partition
- For SAN Volume Controller (SVC), or Storwize System Storage Devices, the IBM PowerHA SystemMirror for i licensed program also requires the following:

- Portable Utilities for IBM i (5733-SC1)
- Ensure that the latest PTF are installed.

For more information on the storage technologies that are provided by IBM i, see [PowerHA supported storage servers](#).

Communications requirement for Switched logical units

Before configuring a PowerHA solution that uses switched logical units, ensures that the minimum communication requirements have been met.

To use the Switched logical unit technology, you must be using or planning to use a storage area network (SAN).

A SAN is a dedicated, centrally managed, secure information infrastructure that enables any-to-any interconnection between systems and storage systems. SAN connectivity is required for using IBM System Storage.

The following are the minimum communication requirements for a PowerHA solution that uses switched logical units:

- The System i product supports various SAN switches and directors. Refer to the IBM System Storage Interoperation Center (SSIC) website, for a complete list of supported switches and directors.
- In addition, taking advantage of multipath I/O is highly recommended in order to improve overall resiliency and performance. Multipath I/O provides the ability to have multiple Fibre Channel devices that are configured to the same logical disk units within the storage. When correctly configured this allows single devices, or I/O enclosures to fail without losing connections to the disk units. Multipath also provides performance benefits by spreading workloads across all available connections (paths). Each connection for a multipath disk unit function independently. Several connections provide improved resiliency by allowing disk storage to be used even if a single path fails.
- Switched logical units require at least one TCP/IP communications interface between the systems in the cluster.

Planning Metro Mirror

IBM PowerHA SystemMirror for i supports Metro Mirror, which provides high availability and disaster recovery. To effectively configure and manage a high availability solution that uses this technology, proper planning is required.

Related information

[Guidelines and recommendations for using Copy Services functions with DS8000](#)

Hardware requirements for Metro Mirror

To configure and manage an PowerHA solution that uses Metro Mirror technology, you should ensure that the minimum hardware requirements are met.

The following minimum hardware requirements are recommended:

- At least two IBM i partitions or systems that are separated geographically with at least one IBM System Storage external storage unit that is attached to each system. The IBM System Storage external storage units are supported on all System i models that support Fibre Channel attachment for external storage.
- PowerHA supports Metro Mirror on DS8000, SAN Volume Controller, and Storwize. For specific Fibre Channel adapters supported, see [IBM System Storage Interoperation Center \(SSIC\)](#).
- Appropriate disk sizing for the system storage should be completed before any configuration. You need one set of disk for the source, and an equal set of disk units for the target.

For more information on the storage technologies that are provided by IBM i, see [PowerHA supported storage servers](#).

Related information

[PowerHA SystemMirror for IBM i Cookbook](#)

[IBM System Storage DS8000 Information Center-> Copy Services](#)

[IBM i and IBM Storwize Family: A Practical Guide to Usage Scenarios](#)

Software requirements for Metro Mirror

Before configuring an IBM i high-availability solution that uses Metro Mirror, ensure that the minimum software requirements have been met.

Metro Mirror has the following minimum software requirements:

- Each IBM i model within the high-availability solution must be running at least IBM i V6R1 for use with the IBM PowerHA SystemMirror for i licensed program.
Note: For prior releases, you can still use the IBM Advanced Copy Services for PowerHA on i, which is an offering from Lab Services, to work with IBM System Storage solutions.
- IBM PowerHA SystemMirror for i licensed program installed on each system participating in the high-availability solution that uses Metro Mirror.
- You need to install IBM i Option 41 HA Switchable Resources. Option 41 gives you the capability to switch independent disk pools between systems. To switch an independent disk pool between systems, the systems must be members of a cluster and the independent switched disk must be associated with a device cluster resource group in that cluster. Option 41 is also required for working with high availability management interfaces, which are provided as part of the IBM PowerHA SystemMirror for i licensed program.
- For DS8000 System Storage Devices, the IBM PowerHA for i licensed program also requires storage command-line interface (DSCLI). DSCLI is required software for all the IBM System Storage solutions, such as FlashCopy technology, Metro Mirror, Global Mirror, there is a requirement to have DSCLI installed on each of the systems or partitions participating in the high availability solution, which uses these storage solutions. DSCLI has these additional software requirements:
 - Java Version 1.4
 - Option 35 (CCA Cryptographic Service Provider) installed on each system or partition
- For SAN Volume Controller (SVC), or Storwize System Storage Devices, the IBM PowerHA SystemMirror for i licensed program also requires the following:
 - Portable Utilities for IBM i (5733-SC1)
- Ensure that the latest PTF have been installed.

Related information

[PowerHA SystemMirror for IBM i Cookbook](#)

[IBM System Storage DS8000 Information Center-> Copy Services](#)

[IBM i and IBM Storwize Family: A Practical Guide to Usage Scenarios](#)

Communications requirement for Metro Mirror

Before configuring PowerHA solution that uses Metro Mirror, ensure that the minimum communication requirements have been met.

To use the Metro Mirror technology, you must be using or planning to use a storage area network (SAN).

A SAN is a dedicated, centrally managed, secure information infrastructure that enables any-to-any interconnection between systems and storage systems. SAN connectivity is required for using IBM System Storage.

The following are the minimum communication requirements for a PowerHA solution that uses Metro Mirror:

- The System i product supports a variety of SAN switches and directors. Refer to the IBM System Storage Interoperation Center (SSIC) website, for a complete list of supported switches and directors.
- In addition, taking advantage of multipath I/O is highly recommended in order to improve overall resiliency and performance. Multipath I/O provides the ability to have multiple Fibre Channel devices that are configured to the same logical disk units within the storage. When correctly configured this allows single devices, or I/O enclosures to fail without losing connections to the disk units. Multipath also provides performance benefits by spreading workloads across all available connections (paths).

Each connection for a multipath disk unit functions independently. Several connections provide improved resiliency by allowing disk storage to be used even if a single path fails.

Related concepts

[PowerHA supported storage servers](#)

IBM System Storage provides enhanced storage capabilities.

Related reference

[Storage area network \(SAN\) Web site](#)

Related information

[IBM SAN Volume Controller Information Center](#)

Journal planning for Metro Mirror

Journaling is important for increasing recovery time for all high availability solutions. In the case of IBM System Storage based technologies, such as Metro Mirror, it is vital that journaling be used to force write operations to external storage units, since mirroring of data occurs outside of IBM i storage.

Journal management prevents transactions from being lost if your system ends abnormally. When you journal an object, the system keeps a record of the changes you make to that object. Regardless of the high availability solution that you implement, journaling is considered a best practice to prevent data loss during abnormal system outages.

Related information

[Journal management](#)

Backup planning for Metro Mirror

With Metro Mirror, you can use the FlashCopy feature to create a copy of data stored in IBM System Storage external storage units.

FlashCopy operations provide the ability to create point-in-time copies. As soon as the FlashCopy operation is processed, both the source and target volumes are available for application use. The FlashCopy feature can be used with other IBM System Storage technologies, such as Metro Mirror and Global Mirror, to create consistent, point-in-time copy of data at a remote site, which then can be backed up with your standard backup procedures. You should complete the following before implementing the FlashCopy technology:

- Identify the source volumes and target volumes for FlashCopy relationships. You should select FlashCopy target volumes in different ranks for better performance.
- Understand FlashCopy data consistency considerations. There are environments where data is stored in system memory cache and written to disk at some later time. To avoid these types of restart actions, ensure that all data that is related to the FlashCopy source volume has been written to disk before you perform the FlashCopy operation.
- You can use an existing Metro Mirror source volume as a FlashCopy target volume. This allows you to create a point-in-time copy using a target volume of a FlashCopy pair and then mirror that data to a source Metro Mirror volume at a remote location.

Performance planning for Metro Mirror

You should understand these performance considerations prior to configuring Metro Mirror.

Before you use Metro Mirror, consider the following requirements and guidelines:

- The source and target volumes in a Metro Mirror relationship must be the same storage type.
- The source and target logical volumes must be the same size or the target must be larger in size.
- For Metro Mirror environments, distribute the work loads by not directing all updates to a small set of common volumes on a single target storage unit. The performance impact at the target site storage unit adversely affects the performance at the source site.
- Similar to any system disk configuration, the number of disk units available to the application can have a significant affect on its performance. Putting additional workload on a limited number of disk units might result in longer disk waits and ultimately longer response times to the application. This is particularly important when it comes to temporary storage in a system configured with independent

disk pools. All temporary storage is written to the SYSBAS disk pool. If your application does not use a lot of temporary storage, then you can get by with fewer disk arms in the SYSBAS disk pool. You must also remember that the operating system and basic functions occur in the SYSBAS disk pool.

Related information

[Guidelines and recommendations for using Copy Services functions with DS8000](#)

Planning Global Mirror

IBM PowerHA SystemMirror for i supports Global Mirror, which provides high availability and disaster recovery in environments that use external storage solutions. To effectively configure and manage high availability that uses this technology, proper planning is required.

Related information

[Guidelines and recommendations for using Copy Services functions with DS8000](#)

Hardware requirements for Global Mirror

To configure and manage an PowerHA solution that uses Global Mirror technology, you should ensure that the minimum hardware requirements are met.

The following minimum hardware requirements should be met for Global Mirror:

- At least two IBM i partitions or systems that are separated geographically with at least one IBM System Storage external storage unit that is attached to each system. The IBM System Storage external storage units are supported on all System i models that support Fibre Channel attachment for external storage.
- PowerHA supports Global Mirror on DS8000, SAN Volume Controller, and Storwize. For specific Fibre Channel adapters supported, see [IBM System Storage Interoperation Center \(SSIC\)](#).
- Appropriate disk sizing for the system storage should be completed before any configuration. You need one set of disk for the source, and an equal set of disk units for the target and another set for each consistency copy.

For more information on the storage technologies that are provided by IBM i, see [PowerHA supported storage servers](#).

Related information

[PowerHA SystemMirror for IBM i Cookbook](#)

[IBM System Storage DS8000 Information Center-> Copy Services](#)

[IBM i and IBM Storwize Family: A Practical Guide to Usage Scenarios](#)

[IBM SAN Volume Controller Information Center](#)

Software requirements for Global Mirror

Before configuring a PowerHA solution that uses Global Mirror, ensure that the minimum software requirements have been met.

Global Mirror has the following minimum software requirements:

- Each IBM i model with in the high-availability solution must be running at least IBM i V6R1 for use with the IBM PowerHA SystemMirror for i licensed program.
Note: For prior releases, you can still use the IBM Advanced Copy Services for PowerHA on i, which is an offering from Lab Services, to work with IBM System Storage solutions. If you are using Global Mirror on multiple platforms, or if you want to implement Global Mirror on multiple IBM i partitions, you can also use the IBM Advanced Copy Services for PowerHA on i.
- IBM PowerHA SystemMirror for i licensed program must be installed on each system participating in the high-availability solution that use Global Mirror.
- For DS8000 System Storage Devices, the IBM PowerHA for i licensed program also requires storage command-line interface (DSCLI). DSCLI is required software for all the IBM System Storage solutions, such as FlashCopy technology, Metro Mirror, Global Mirror, there is a requirement to have DSCLI installed each of the systems or partitions participating in the high availability solution, which uses these storage solutions. DSCLI has these additional software requirements:
 - Java Version 1.4

- Option 35 (CCA Cryptographic Service Provider) installed on each system or partition
- For SAN Volume Controller (SVC), or Storwize System Storage Devices, the IBM PowerHA SystemMirror for i licensed program also requires the following:
 - Portable Utilities for IBM i (5733-SC1)
- Ensure that the latest PTF have been installed.

Related information

[PowerHA SystemMirror for IBM i Cookbook](#)

[IBM System Storage DS8000 Information Center-> Copy Services](#)

[IBM i and IBM Storwize Family: A Practical Guide to Usage Scenarios](#)

Communications requirement for Global Mirror

Before configuring a PowerHA solution that uses Global Mirror, you should ensure that the minimum communication requirements have been met.

To use the Global Mirror technology you must be using or planning to use a storage area network (SAN).

A SAN is a dedicated, centrally managed, secure information infrastructure that enables any-to-any interconnection between systems and storage systems. SAN connectivity is required for using IBM System Storage.

The following are the minimum communication requirements for a PowerHA solution that use Global Mirror:

- The System i product supports a variety of SAN switches and directors. Refer to the IBM System Storage Interoperation Center (SSIC) website, for a complete list of supported switches and directors.
- In addition, taking advantage of multipath I/O is highly recommended in order to improve overall resiliency and performance. Multipath I/O provides the ability to have multiple Fibre Channel devices that are configured to the same logical disk units within the storage. When correctly configured this allows single devices, or I/O enclosures to fail without losing connections to the disk units. Multipath also provides performance benefits by spreading workloads across all available connections (paths). Each connection for a multipath disk unit functions independently. Several connections provide improved resiliency by allowing disk storage to be used even if a single path fails.

Related concepts

[PowerHA supported storage servers](#)

IBM System Storage provides enhanced storage capabilities.

Related reference

[Storage area network \(SAN\) Web site](#)

Related information

[IBM SAN Volume Controller Information Center](#)

Journal planning for Global Mirror

Journaling is important for decreasing recovery time for all high availability solutions. In the case of IBM System Storage based technologies, such as Global Mirror, journaling forces write operations to external storage units, which is necessary because data mirroring occurs outside of IBM i storage.

Journal management prevents transactions from being lost if your system ends abnormally. When you journal an object, the system keeps a record of the changes you make to that object. Regardless of the high availability solution that you implement, journaling is considered a best practice to prevent data loss during abnormal system outages.

Related information

[Journal management](#)

Backup planning for Global Mirror

When using Global Mirror technology within your high-availability solution, you can use the FlashCopy feature to create point-in-time copy of your data.

FlashCopy operations provide the ability to create point-in-time copies. As soon as the FlashCopy operation is processed, both the source and target volumes are available for application use. The FlashCopy feature can be used with other IBM System Storage technologies, such as Metro Mirror and Global Mirror, to create consistent, point-in-time copy of data at a remote site, which then can be backed up with your standard backup procedures. You should complete the following before implementing the FlashCopy technology:

- Identify the source volumes and target volumes for FlashCopy relationships. You should select FlashCopy target volumes in different ranks for better performance.
- Understand FlashCopy data consistency considerations. There are environments where data is stored in system memory cache and written to disk at some later time. To avoid these types of restart actions, ensure that all data that is related to the FlashCopy source volume has been written to disk before you perform the FlashCopy operation.

Performance planning for Global Mirror

You should understand these performance considerations before configuring Global Mirror.

Before you use Global Mirror, consider these performance guidelines:

- The source and target volumes in a Global Mirror relationship must be the same storage type.
- The source and target logical volumes must be the same size or the target must be the same size or the target must be larger in size.
- Similar to any system disk configuration, the number of disk units available to the application can have a significant affect on its performance. Putting additional workload on a limited number of disk units might result in longer disk waits and ultimately longer response times to the application. This is particularly important when it comes to temporary storage in a system configured with independent disk pools. All temporary storage is written to the SYSBAS disk pool. If your application does not use a lot of temporary storage, then you can get by with fewer disk arms in the SYSBAS disk pool. You must also remember that the operating system and basic functions occur in the SYSBAS disk pool.

Related information

[Guidelines and recommendations for using Copy Services functions with DS8000](#)

Planning the FlashCopy feature

You can use the FlashCopy feature as a means to reduce your backup window in PowerHA environments that use the external storage units of the IBM System Storage. Before using the FlashCopy feature, ensure that the minimum requirements have been met.

Hardware requirements for the FlashCopy feature

To use the FlashCopy technology in an PowerHA solution, ensure that the minimum hardware requirements are met.

The following minimum hardware requirements are needed for the FlashCopy feature:

- At least one IBM i attached to an IBM System Storage device. For a FlashCopy to a IBM i target system or partition (Where the FlashCopy target is not *NONE), both the FlashCopy source and FlashCopy target systems or partitions must be attached to the same external storage device.
- PowerHA supports FlashCopy on DS8000, SAN Volume Controller, and Storwize. For specific Fibre Channel adapters supported, see [IBM System Storage Interoperation Center \(SSIC\)](#).
- Appropriate disk sizing for the system storage should be completed before any configuration.
 - For FlashCopy with DS8000 System Storage Devices, you need one set of disks for the source, an equal set of disk units for the target, and another set for each consistency group. Space efficient FlashCopy (FlashCopy SE) can be used to reduce this disk requirement.
 - For FlashCopy with SAN Volume Controller (SVC), or Storwize System Storage Devices, you need one set of disks for the source, and an equal set of disk units for the target.

For more information about the storage technologies that are provided by IBM, see [“PowerHA supported storage servers”](#) on page 37.

Related concepts

[PowerHA supported storage servers](#)

IBM System Storage provides enhanced storage capabilities.

Related information

[PowerHA SystemMirror for IBM i Cookbook](#)

[IBM System Storage DS8000 Information Center-> Copy Services](#)

[IBM i and IBM Storwize Family: A Practical Guide to Usage Scenarios](#)

[IBM SAN Volume Controller Information Center](#)

Software requirements for the FlashCopy feature

To use the FlashCopy technology in an PowerHA solution, the minimum software requirements should be met.

The FlashCopy feature has the following minimum software requirements:

- Each IBM i model with in the PowerHA solution must be running at least IBM(r) iV6R1 for use with the IBM PowerHA SystemMirror for i licensed program.

Note: For prior releases, you can still use the IBM Advanced Copy Services for PowerHA on i, which is an offering from Lab Services, to work with IBM System Storage solutions.

- IBM PowerHA SystemMirror for i installed on each system.
- Ensure that the latest PTF have been installed.

Communications requirements for the FlashCopy feature

To use the FlashCopy technology in a PowerHA solution, ensure that the minimum communications requirements are met.

The following minimum communication requirements should be met for the FlashCopy feature:

- At least two IBM i partitions or systems that have at least one IBM(r) System Storage(r) external storage unit that is attached to each system. The IBM(r) System Storage(r) external storage units are supported on all System i(r) models that support Fibre Channel attachment for external storage.
- PowerHA supports Metro Mirror on DS8000, SAN Volume Controller, and Storwize. For specific Fibre Channel adapters supported, see [IBM System Storage Interoperation Center \(SSIC\)](#).

For more information on the storage technologies that are provided by IBM i, see [“PowerHA supported storage servers”](#) on page 37.

Related concepts

[PowerHA supported storage servers](#)

IBM System Storage provides enhanced storage capabilities.

Planning DS8000 HyperSwap with independent auxiliary storage pools (IASPs)

When planning a HyperSwap configuration with a PowerHA logical unit (LUN) switching environment and or a live partition mobility environment, the following requirements must be met.

Hardware requirements for DS8000 HyperSwap with independent auxiliary storage pools (IASPs)

To use the DS8000 HyperSwap with IASPs technology in a PowerHA solution, ensure that the following minimum hardware requirements are met.

- Two IBM i logical partitions (on the same or separate POWER servers) with all storage provided by two IBM System Storage DS8000 external storage units that are attached to the systems.
- Two supported IBM Systems Storage DS8000. Supported IBM System Storage units:
 - DS8700 or DS8800 (with release 6.3 SP7)
 - DS8870 with release 7.2

- Appropriate disk sizing for the system storage should be completed before any configuration. You need one set of disks on the source DS8870 unit, and an equivalent sized set of disks on the target DS8870 unit. To protect both partitions from storage outages, you also need extra disk to replicate SYSBAS from each system to the other. For example, if you have partitions A and B, the storage for partition A should be enough to store SYSBAS A, IASP A, and SYSBAS B (copy). The storage for partition B should be enough to store SYSBAS B, IASP B (copy), and SYSBAS A (copy).

For information on storage technologies that are provided by IBM, see [“PowerHA supported storage servers”](#) on page 37.

Software requirements for DS8000 HyperSwap with independent auxiliary storage pools (IASPs)

To use the DS8000 HyperSwap technology with IASPs in an IBM i high-availability solution, the minimum software requirements must be met.


The IASP-based HyperSwap has the following minimum software requirements:

- Both IBM i partitions must be at IBM i 7.2 at Technology Refresh (TR) 4 or at IBM i 7.3.
- IBM PowerHA for i Enterprise Edition licensed program must be installed on both partitions, with the latest HA Group PTF installed.
- The PowerHA product requires IBM i Option 41 HA Switchable Resources as a prerequisite.
- For PowerHA to interact with the DS8870 System Storage device, the storage command-line interface (DSCLI) is required. DSCLI has these additional software requirements:
 - Java Version 1.4 (or higher)
 - Option 35 (CCA Cryptographic Service Provider) installed on each system or partition
- Before you can use IBM Navigator for i to perform any disk management tasks you need to set up the proper authorizations for dedicated service tools (DST).

Communications requirements for DS8000 HyperSwap with independent auxiliary storage pools (IASPs)

To use the DS8000 HyperSwap technology, you must be using or planning to use a storage area network (SAN).

A SAN is a dedicated, centrally managed, secure information infrastructure that enables any-to-any interconnection between systems and storage systems. SAN connectivity is required for using IBM System Storage DS8000 external storage units.

IBM i product supports various SAN switches and directors. Refer to the [IBM System Storage Interoperation Center \(SSIC\)](#)  website for a complete list of supported adapters, switches, and directors.

Taking advantage of multipath I/O is highly recommended in order to improve overall resiliency and performance. Multipath I/O provides the ability to have multiple Fibre Channel devices that are configured to the same logical disk units within the storage. When correctly configured, this allows single devices, or I/O enclosures to fail without losing connections to the disk units. Multipath I/O also provides performance benefits by spreading workloads across all available connections (paths). Each connection for a multipath disk unit functions independently. Several connections provide improved resiliency by allowing disk storage to be used even if a single path fails.

Journal planning for DS8000 HyperSwap with independent auxiliary storage pools (IASPs)

Journaling is important for improved recovery time for all high availability solutions, and is highly recommended. In the case of IBM System Storage based technologies, such as HyperSwap, it is vital that journaling be used to force write operations to external storage units, since mirroring of data occurs outside of IBM i storage.

Journal management prevents transactions from being lost if your system ends abnormally. When you journal an object, the system keeps a record of the changes you make to that object. Regardless of the high availability solution that you implement, journaling is considered a best practice to prevent data loss during abnormal system outages.

Backup planning for DS8000 HyperSwap with independent auxiliary storage pools (IASPs)

With HyperSwap, you can use the FlashCopy feature to create a copy of data that is stored in IBM System Storage external storage units.

FlashCopy operations provide the ability to create point-in-time copies. As soon as the FlashCopy operation is processed, both the source and target volumes are available for application use. The FlashCopy feature can be used with other IBM System Storage technologies, such as Metro Mirror and Global Mirror, to create consistent, point-in-time copy of data at a remote site, which then can be backed up with your standard backup procedures. You should complete the following before implementing the FlashCopy technology:

- Identify the source volumes and target volumes for FlashCopy relationships. You should select FlashCopy target volumes in different ranks for better performance.
- Understand FlashCopy data consistency considerations. There are environments where data is stored in system memory cache and written to disk at some later time. To avoid these types of restart actions, ensure that all data that is related to the FlashCopy source volume has been written to disk before you perform the FlashCopy operation.
- You can use an existing HyperSwap source volume as a FlashCopy target volume. Select this option to create a point-in-time copy that uses a target volume of a FlashCopy pair and then mirror that data to a source HyperSwap volume at a remote location.

Performance requirements for DS8000 HyperSwap with independent auxiliary storage pools (IASPs)

You should understand these performance considerations before configuring HyperSwap.



Before you use HyperSwap, consider the following requirements and guidelines:

- The source and target volumes in a HyperSwap Metro Mirror relationship must be the same storage type.
- The source and target logical volumes must be the same size.
- As with any synchronous replication technology, there are distance and bandwidth limitations that can affect performance.
- Similar to any system disk configuration, the number of disk units available to the application can have a significant affect on its performance. Putting more workload on a limited number of disk units may result in longer disk waits and ultimately longer response times to the application. This is important when it comes to temporary storage in a system configured with IASPs. All temporary storage is written to the SYSBAS disk pool. If your application does not use a lot of temporary storage, then you can get by with fewer disk arms in the SYSBAS disk pool. Remember, the operating system and basic functions occur in the SYSBAS disk pool.

PowerHA supported storage servers

IBM System Storage provides enhanced storage capabilities.

PowerHA supports the Metro Mirror, Global Mirror, LUN switching, and FlashCopy technologies with the following storage servers: DS8000 storage family, SAN Volume Controller (SVC), and the IBM Storwize models. For more information on specific implementations, see the following guides.

- [IBM i and IBM Storwize Family: A Practical Guide to Usage Scenarios](#) 
- [IBM i and IBM System Storage: A Guide to Implementing External Disks on IBM i](#) 

Depending on the IBM System Storage Device, the IBM PowerHA SystemMirror for i licensed program also requires:

- For DS8000 System Storage Devices, the IBM PowerHA for i licensed program also requires storage command-line interface (DSCLI). DSCLI is required for all of the IBM DS8000 Series System Storage solutions, such as FlashCopy, Metro Mirror, and Global Mirror.
- There is a requirement to have DSCLI installed on each of the systems or partitions participating in the high availability solution, which uses these storage solutions.
- DSCLI has these additional software requirements:

- Java Version 1.4
- Option 35 (CCA Cryptographic Service Provider) installed on each system or partition
- For SAN Volume Controller (SVC), or Storwize System Storage Devices, the IBM PowerHA for i licensed program also requires:
 - Portable Utilities for IBM i (5733-SC1)

Related concepts

Communications requirement for Metro Mirror

Before configuring PowerHA solution that uses Metro Mirror, ensure that the minimum communication requirements have been met.

Communications requirement for Global Mirror

Before configuring a PowerHA solution that uses Global Mirror, you should ensure that the minimum communication requirements have been met.

Managing switched logical units (LUNs)

Switched logical units are independent disk pools created from logical units created in an IBM System Storage that have been configured as part of a device cluster resource group (CRG).

Hardware requirements for the FlashCopy feature

To use the FlashCopy technology in an PowerHA solution, ensure that the minimum hardware requirements are met.

Communications requirements for the FlashCopy feature

To use the FlashCopy technology in a PowerHA solution, ensure that the minimum communications requirements are met.

Scenario: Global Mirror

This scenario describes an PowerHA solution that is based on external storage and provides disaster recovery and high availability for storage systems that are separated by great distances. Global Mirror is an IBM Systems Storage solution that copies data asynchronously from the storage unit at the production site to the storage unit at the backup site. In this way, data remains consistent at the backup site.

Scenario: Metro Mirror

This scenario describes a PowerHA solution, which is based on external storage and provides disaster recovery and high availability for storage systems, which are separated by short distances. Metro Mirror is an IBM System Storage solution that copies data synchronously from the storage unit at the production site to the storage unit at the backup site. In this way data remains consistent at the backup site.

Scenario: Performing a FlashCopy function

In this example, an administrator wants to perform a backup from the remote copy of data that is stored in an external storage unit at the backup site. Using the FlashCopy function available with IBM Storage Solutions, the administrator reduces his backup time considerably.

Related tasks

Configuring a FlashCopy session

For IBM i high-availability environments that use IBM System Storage technology, you can configure a FlashCopy session to create a point-in-time copy of data.

Configuring switched logical units (LUNs)

Switched logical units is an independent disk pool that is controlled by a device cluster resource group and can be switched between nodes within a cluster. For IBM i high availability solutions that use switchable logical units with IBM® System Storage™, you need to configure an ASP Copy Description that defines the host identifiers and volume groups for each cluster node in which the ASP device can be switched.

Configuring Metro Mirror

For IBM i high availability solutions that use IBM System Storage Metro Mirror technology, you need to configure a session between the IBM i machine and IBM System Storage external storage units that have Metro Mirror configured. In IBM i, Metro Mirror sessions do not set up the mirroring on the external storage units, but rather sets up a relationship between the IBM i systems and the existing Metro Mirror configuration on external storage units.

Configuring Global Mirror

For IBM i high availability solutions that use IBM System Storage Global Mirror technology, you need to configure a session between the IBM i machine and IBM System Storage external storage units that have Global Mirror configured. In IBM i, Global Mirror sessions do not set up the mirroring on the external storage units, but rather sets up a relationship between the IBM i systems and the existing Global Mirror configuration on external storage units.

Planning environment resiliency

Environment resiliency ensures that your objects and attributes remain consistent among resources defined in the high-availability environment. You need to identify which resources require a consistent environment to function properly and create a cluster administrative domain that will ensure that these resource attributes remain consistent in your high availability solution.

Planning for a cluster administrative domain

The cluster administrative domain requires planning to manage resources that are synchronized among nodes within a cluster administrative domain. In order to ensure that an application will run consistently on any system in a high-availability environment, all resources that affect the behavior of the application need to be identified, as well as the cluster nodes where the application will run, or where application data might reside.

A cluster administrator can create a cluster administrative domain and add monitored resources that are synchronized among nodes. The IBM i cluster provides a list of system resources that can be synchronized by a cluster administrative domain, represented by monitored resource entries (MREs).

When designing a cluster administrative domain, you should answer the following questions:

What nodes will be included in the cluster administrative domain?

You should determine what nodes in a cluster are to be managed by the cluster administrative domain. These are the cluster nodes representing the systems where an application can run or where the application data is stored, and that require a consistent operational environment. Nodes cannot be in multiple cluster administrative domains. For example, if you have four nodes in a cluster (Node A, Node B, Node C and Node D), Nodes A and B can be in one cluster administrative domain and Nodes C and D can be in another. However you cannot have Nodes B and C in a third cluster administrative domain and still have them in their original cluster administrative domain.

What will be the naming convention for cluster administrative domains?

Depending on the complexity and size of your clustered environment, you might want to establish a standard naming convention for peer CRGs and cluster administrative domains. Since a peer CRG is created when you create a cluster administrative domain, you will want to differentiate other peer CRGs from those that represent cluster administrative domains. For example, peer CRGs that represent cluster administrative domains can be named *ADMDMN1*, *ADMDMN2*, and so forth, while other peer CRGs can be named *PEER1*. You can also use the [List Cluster Resource Group Information \(QcstListClusterResourceGroupIn\) API](#) to determine whether the peer CRG is used as a cluster administrative domain. A peer CRG which represents a cluster administrative domain can be identified by its application identifier, which is QIBM.AdminDomain.

Planning monitored resources entries (MRE)

Monitored resources are IBM i objects that can be defined within a cluster administrative domain. These resources need to remain consistent across the systems in a high-availability environment otherwise during an outage applications might not perform as expected. You should plan which supported resources within your environment should be monitored.

You need to determine which system resources need to be synchronized. You can select attributes for each of these resources to customize what is synchronized. Applications that run on multiple nodes might need specific environment variables to run properly. In addition data that spans several nodes might also require certain user profiles to be accessed. Be aware of the operational requirements for your applications and data before you determine what resources need to be managed by a cluster administrative domain.

Security planning for high availability

Prior to configuring your high-availability solution, you should reassess the current security strategies in your environment and make any appropriate changes to facilitate high availability.

Distributing cluster-wide information

Learn about the security implications of using and managing cluster-wide information.

The `Distribute Information (QcstDistributeInformation)` API can be used to send messages from one node in a cluster resource group recovery domain to other nodes in that recovery domain. This can be useful in exit program processing. However, it should be noted that there is no encryption of that information. Secure information should not be sent with this mechanism unless you are using a secure network.

Non persistent data can be shared and replicated between cluster nodes by using the Clustered Hash Table APIs. The data is stored in non persistent storage. This means the data is retained only until the cluster node is no longer part of the clustered hash table. These APIs can only be used from a cluster node that is defined in the clustered hash table domain. The cluster node must be active in the cluster.

Other information distributed by using cluster messaging is similarly not secured. This includes the low level cluster messaging. When changes are made to the exit program data, there is no encryption of the message containing that data.

Considerations for using clusters with firewalls

If you are using clustering in a network that uses firewalls, you should be aware of some limitations and requirements.

If you are using clustering with a firewall, you need to give each node the ability to send outbound messages to and receive inbound messages from other cluster nodes. An opening in the firewall must exist for each cluster address on each node to communicate with every cluster address on every other node. IP packets traveling across a network can be of various types of traffic. Clustering uses ping, which is type ICMP, and also uses UDP and TCP. When you configure a firewall, you can filter traffic based on the type. For clustering to work the firewall needs to allow traffic of ICMP, UDP and TCP. Outbound traffic can be sent on any port and inbound traffic is received on ports 5550 and 5551.

In addition, if you are making use of advanced node failure detection, any cluster node that is to receive failure messages from a Hardware Management Console (HMC) or a Virtual I/O Server (VIOS) on an Integrated Virtualization Manager (IVM) managed server must be able to communicate with that HMC or VIOS partition. The cluster node will send to the HMC or VIOS on the IP address that is associated with the HMC's or VIOS domain name and to port 5989. The cluster node will receive from the HMC or VIOS on the IP address that is associated with the cluster node's system name and on port 5989.

Maintaining user profiles on all nodes

You can use two mechanisms for maintaining user profiles on all nodes within a cluster.

In a high-availability environment, a user profile is considered to be the same across systems if the profile names are the same. The name is the unique identifier in the cluster. However, a user profile also contains a user identification number (UID) and group identification number (GID). To reduce the amount of internal processing that occurs during a switchover, where the independent disk pool is made unavailable on one system and then made available on a different system, the UID and GID values should be synchronized across the recovery domain for the device CRG. Administrative Domain can be used to synchronize user profiles, including the UID and GID values, across the cluster.

One mechanism is to create a cluster administrative domain to monitor shared resources across nodes in a cluster. A cluster administrative domain can monitor several types of resources in addition to user profiles, providing easy management of resources that are shared across nodes. When user profiles are updated, changes are propagated automatically to other nodes if the cluster administrative domain is active. If the cluster administrative domain is not active, the changes are propagated after the cluster administrative domain is activated. This method is recommended because it automatically maintains user profiles with a high-availability environment.

With the second mechanism, administrators can also use Management Central in IBM Navigator for i to perform functions across multiple systems and groups of systems. This support includes some common

user-administration tasks that operators need to perform across the multiple systems in their cluster. With Management Central, you can perform user profile functions against groups of systems. The administrator can specify a post-propagation command to be run on the target systems when creating a user profile.

Important:

- If you plan to share user profiles that use password synchronization within a cluster, you must set the Retain Server Security (QRETSVRSEC) system value to 1.
- If you change QRETSVRSEC to 0 after you add a monitored resource entry (MRE) for a user profile, and then change a password (if password is being monitored), the global status for the MRE is set to Inconsistent. The MRE is marked as unusable. Any changes made to the user profile after this change are not synchronized. To recover from this problem, change QRETSVRSEC to 1, remove the MRE, and add the MRE again.

Related tasks

Creating a cluster administrative domain

In a high-availability solution, the cluster administrative domain provides the mechanism that keeps resources synchronized across systems and partitions within a cluster.

Implementing PowerHA

The task-based approach to configuring and managing IBM i high availability allows you to configure and manage a customized high-availability solution that is based on your business needs. Graphical and command-line interfaces are used for configuring and managing your high-availability solution.

The task-based approach gives the knowledgeable user the means to customize and implement a personalized solution.

PowerHA graphical interface

The PowerHA graphical interface allows you to configure and manage cluster technologies, which are integral to a high-availability solution. To use this interface, the PowerHA licensed program must be installed. With this interface you can perform the following functions from any node in the cluster:

- Create and manage a cluster
- Create and manage nodes
- Create and manage cluster resource groups
- Create and manage cluster administrative domains
- Create and manage monitored resources
- Create and manage independent ASPs
- Configure and manage geographic mirroring
- Configure and manage Metro Mirror
- Configure and manage Global Mirror
- Configure and manage FlashCopy
- Monitor the status of your high availability solution
- Perform cluster related operations, such as switchovers for planned outages

Disk Management interface

The Disk Management interface allows you to configure and manage independent disk pools which are necessary when implementing several data resiliency technologies. Depending on the type of data resiliency technology that is implemented, installation requirements might be necessary to use some of these functions:

- Create a disk pool

- Make a disk pool available
- Make a disk pool unavailable
- Configure geographic mirroring

PowerHA command-line interface

The command-line interface allows you to perform many different high availability tasks with CL commands. For each cluster-related task, the corresponding CL command has been identified.

PowerHA application programming interface

These application programming interfaces (APIs) allow you to work with PowerHA version, and retrieve PowerHA related information.

Related information

[PowerHA commands](#)

[PowerHA APIs](#)

Configuring high availability infrastructure

Before you can configure a high-availability solution in your IBM i environment, ensure that you have completed the appropriate planning and understand your resources and goals for high availability and disaster recovery. Use configuration scenarios for high availability and tasks that are associated with high availability technologies to create your own high-availability solution.

Scenarios: Configuring high availability

Configuration scenarios provide examples of different PowerHA environments and step-by-step configuration tasks to help you implement a high-availability solution that is based on your needs and resiliency requirements.

These scenarios contain descriptions of the business objectives for high availability and provide a graphic that illustrates the resources within the high-availability solution. Each solution example contains step-by-step instructions to set up and test high availability. However, this information does not cover all configuration cases and additional testing may be required to verify high availability.

For additional details on specific environments, see the following Redbooks:

- PowerHA SystemMirror for IBM i Cookbook
- IBM i and IBM Storwize Family: A Practical Guide to Usage Scenarios
- Simple Configuration Example for Storwize V7000 FlashCopy and PowerHA SystemMirror for i

Related information

[PowerHA SystemMirror for IBM i Cookbook](#)

[IBM System Storage DS8000 Information Center-> Copy Services](#)

[IBM i and IBM Storwize Family: A Practical Guide to Usage Scenarios](#)

[Simple Configuration Example for Storwize V7000 FlashCopy and PowerHA SystemMirror for i](#)

Scenario: Geographic mirroring

This scenario describes an IBM i high-availability solution that uses geographic mirroring in a two node cluster. This solution provides both disaster recovery and high availability.

Overview

Geographic mirroring is a PowerHA technology where data is mirrored to a copy of the independent disk pool at the remote location. This solution provides disaster recovery in the event of a site-wide outage on the production system (System 1). In that situation, failover to the backup site (System 2) occurs, where operations can continue on the mirrored copy of the data. This solution provides a simple and less expensive alternative to external storage-based solutions, such as IBM System Storage Global Mirror and

Metro Mirror. However, geographic mirroring does not offer all the performance options that the external storage solutions provide.

Objectives

This solution has the following advantages:


- Provides availability for your business resources during planned outages
- Provides availability for business resources during unplanned outages
- Provides availability for business resources during a disaster
- Enables data to remain current and may not need to be synchronized

This solution has the following restrictions:

- There is no concurrent access to the disk pool. Only the production copy is normally accessed. However, you can detach the mirror copy for offline processing of a second copy of the data.
- Potentially affects performance because increased central processing unit (CPU) is required to support geographic mirroring
- Consider using redundant communication paths and adequate bandwidth

Details

The following graphic illustrates this solution:



Configuration steps

Complete the following steps to configure the high availability technologies that are associated with this scenario using the PowerHA graphical interface:

1. [Complete planning checklist for cluster](#)
2. [Create a cluster](#)
3. [Create a cluster administrative domain](#)
4. [Start a cluster administrative domain](#)
5. [Create an independent disk pool](#)
6. [Add monitored resource entries](#)
7. [Make the independent ASP highly available](#)
8. [Configure geographic mirroring](#)
9. [Vary on the independent ASP](#)
10. [Perform a switchover to test your high-availability solution](#)

Complete the following steps to configure the high availability technologies that are associated with this scenario using commands:

1. [Complete planning checklist for clusters](#)
2. [Create a cluster](#)
3. [Add nodes](#)
4. [Start nodes](#)
5. [Add nodes to device domain](#)


6. [Create a cluster administrative domain](#)
7. [Start cluster administrative domain](#)
8. Create an independent disk pool using [Configure Device ASP](#).
9. [Add monitor resource entries](#)
10. [Create device CRG](#)
11. [Start a device CRG](#)
12. Use [Vary Configuration](#) to make the disk pool available.
13. Perform a [switchover](#) to test the configuration.

Scenario: Metro Mirror

This scenario describes a PowerHA solution, which is based on external storage and provides disaster recovery and high availability for storage systems, which are separated by short distances. Metro Mirror is an IBM System Storage solution that copies data synchronously from the storage unit at the production site to the storage unit at the backup site. In this way data remains consistent at the backup site.

Overview

The Metro Mirror solution provides a high availability and disaster recovery by using external storage units within a metropolitan area. The independent disk pool is replicated between the external storage devices to provide availability for both planned and unplanned outages. When Metro Mirror receives a host update to the production volume, it completes the corresponding update to the backup volume. Metro Mirror supports a maximum distance of 300 km (186 mi). Delays in response times for Metro Mirror are proportional to the distance between the volumes.

This scenario covers the configuration of IBM-supplied IBM i high availability technology and does not provide installation or configuration instructions regarding IBM System Storage DS8000 series. This information assumes that an IBM System Storage solution is already in place before the IBM i high availability configuration. For installation and configuration information about DS8000, see [Copy Services in the IBM System Storage DS8000 Information Center](#) .

A similar scenario can also be accomplished with other system storage technologies. For more information on storage technologies that support Metro Mirror, see [“PowerHA supported storage servers” on page 37](#).

Objectives

This solution has the following advantages:


- Replication is entirely managed by the external storage unit, thus no IBM i CPU is used. Replication continues in the storage unit even when the system experiences a system-level failure.
- Availability for business resources during planned or unplanned outages, which include maintenance outages or software/PTF related outages as well as disaster recovery.
- I/O remains consistent and does not need to be synchronized
- Fast recovery times when used with journaling. Journaling recovers data more quickly in the event of an unplanned outage or failover. Journaling forces data changes to disk where the mirroring is occurring. If you do not use journaling, you might lose data that is in memory. Journaling provides recovery of these data-level transactions and helps with the recovery times.
- The ability to use the FlashCopy function on the source or target side of Metro Mirror.

This solution has the following restrictions:

- Requires external storage hardware
- Consider using redundant communication paths and adequate bandwidth
- There is no concurrent access to the disk pool

Details

The following graphic illustrates this solution:



Configuration steps

Complete the following steps to configure the high availability technologies that are associated with this scenario using the PowerHA graphical interface:

1. [Complete planning checklist for clusters](#)
2. [Create a cluster of two active nodes](#)
3. [Create a cluster administrative domain](#)
4. [Start a cluster administrative domain](#)
5. [Create an independent disk pool](#)
6. [Add monitored resource entries](#)
7. [Make the independent ASP highly available](#)
8. [Configure Metro Mirror](#)
9. [Make disk pool available](#)
10. [Perform a switchover to test your high-availability solution](#)

Complete the following steps to configure the high availability technologies that are associated with this scenario using other interfaces:

1. [Complete planning checklist for clusters](#)
2. [Create a cluster](#)
3. [Add nodes](#)
4. [Start nodes](#)
5. [Add nodes to device domain](#)
6. [Create a cluster administrative domain](#)
7. [Start a cluster administrative domain](#)
8. [Create an independent disk pool using Configure Device ASP](#).
9. [Add monitored resource entries](#)
10. [Create a device CRG](#)
11. [Start a device CRG](#)
12. [Configure Metro Mirror session](#)
13. [Make disk pool available](#)
14. [Perform a switchover to test your high-availability solution.](#)

Related concepts

[PowerHA supported storage servers](#)


IBM System Storage provides enhanced storage capabilities.

Scenario: Global Mirror

This scenario describes an PowerHA solution that is based on external storage and provides disaster recovery and high availability for storage systems that are separated by great distances. Global Mirror is an IBM Systems Storage solution that copies data asynchronously from the storage unit at the production site to the storage unit at the backup site. In this way, data remains consistent at the backup site.

Overview

The Global Mirror solution provides a disaster recovery solution by using external storage units across long distances. The independent disk pool is replicated between the external storage devices to provide availability for both planned and unplanned outages.

This scenario covers the configuration of IBM-supplied IBM i high availability technology and does not provide installation or configuration instructions regarding IBM System Storage DS8000 series. This information assumes that an IBM System Storage solution is already in place before the IBM i high availability configuration. For installation and configuration information about DS8000, see [Copy Services in the IBM System Storage DS8000 Information Center](#) .

A similar scenario can also be accomplished with other system storage technologies. For more information on storage technologies that support Global Mirror, see [“PowerHA supported storage servers”](#) on page 37.

Objectives

Cross-site mirroring with Global Mirror provides the following advantages:


- Replication is entirely managed by the external storage unit, thus no IBM i CPU is used. Replication continues in the storage unit even when the system experiences a system-level failure.
- Availability for business resources during planned or unplanned outages, which include maintenance outages or software/PTF related outages as well as disaster recovery.
- Fast recovery times when used with journaling. Journaling recovers data more quickly in the event of an unplanned outage or failover. Journaling forces data changes to disk where the mirroring is occurring. If you do not use journaling, you might lose data that is in memory. Journaling provides recovery of these data-level transactions and helps with the recovery times.
- The ability to use the FlashCopy function on the source or target side of Global Mirror.

This solution has the following restrictions:

- The solution requires IBM System Storage DS8000 server hardware.
- To achieve acceptable performance, consider using redundant communication paths and adequate bandwidth.
- There is no concurrent access to the disk pool.
- A consistency group is required for the Global Mirror target copy. A consistency group is not required for the Global Mirror source copy, but it is highly recommended.
- Reverse replication occurs automatically on a switchover only if the new target has a consistency group. Reverse replication never occurs automatically on a failover.
- When reverse replication does not occur on a switchover or failover, the configuration will consist of two source copies.
 - If the desired target copy node has a consistency group, then a reattach operation will convert it to a target copy and automatically initiate replication.
 - If the desired target copy node does not have a consistency group, then recovery requires manual intervention with the System Storage DS8000 Storage Manager interface to initiate replication and synchronize the current source and target.

Details

The following graphic illustrates this solution:



Configuration steps

Complete the following steps to configure the high availability technologies that are associated with this scenario using the PowerHA graphical interface:

1. [Complete planning checklist for cluster](#)
2. [Create a cluster](#)
3. [Create a cluster administrative domain](#)
4. [Start a cluster administrative domain](#)
5. [Create an independent disk pool](#)
6. [Add monitored resource entries](#)
7. [Make the independent ASP highly available](#)
8. [Configure Global Mirror](#)
9. [Vary on the independent ASP](#)
10. [Perform a switchover to test your high-availability solution](#)

Complete the following steps to configure the high availability technologies associated with this scenario using other interfaces:

1. [Complete planning checklist for cluster](#)
2. [Create a cluster](#)
3. [Add a nodes](#)
4. [Start a node](#)
5. [Add nodes to a device domain](#)
6. [Create cluster administrative domain](#)
7. [Start a cluster administrative domain](#)
8. [Create an independent disk pool using Configure Device ASP](#).
9. [Add ASP copy description](#)
10. [Add monitored resource entries](#)
11. [Create a device CRG](#)

12. [Start a device CRG](#)
13. [Start ASP session](#)
14. [Configure global mirror session](#)
15. [Make disk pool available](#)
16. Perform a [switchover](#) to test your high-availability solution

Related concepts

[PowerHA supported storage servers](#)

IBM System Storage provides enhanced storage capabilities.

Setting up TCP/IP for high availability

Because cluster resource services uses only IP to communicate with other cluster nodes, which are systems or logical partitions within a high availability environment, all cluster nodes must be IP-reachable, which means that you must have IP interfaces configured to connect the nodes in your cluster.

IP addresses must be set up either manually by the network administrator in the TCP/IP routing tables on each cluster node or they might be generated by routing protocols running on the routers in the network. This TCP/IP routing table is the map that clustering uses to find each node; therefore, each node must have its own unique IP address.

Each node can have up to two IP addresses assigned to it. These addresses must not be changed in any way by other network communications applications. Be sure when you assign each address that you take into account which address uses which kind of communication line. If you have a preference for using a specific type of communication media, make sure that you configure the first IP address by using your preferred media. The first IP address is treated preferentially by the reliable message function and heartbeat monitoring. Every cluster IP address on every node must be able to reach every other IP address in the cluster. If any cluster node uses an IPv4 address, then every node in the cluster needs to have an active IPv4 address (not necessarily configured as a cluster IP address) that can route to and send TCP packets to that address. Also, if any cluster node uses an IPv6 address, then every node in the cluster needs to have an active IPv6 address (not necessarily configured as a cluster IP address) that can route to and send TCP packets to that address. One way to verify that one address can reach another address is if you can ping and use a UDP message trace route in both directions; however, PING and TRACEROUTE do not work between IPv4 and IPv6 addresses, or if a firewall is blocking them.

Note: You need to be sure that the loopback address (127.0.0.1) is active for clustering. This address, which is used to send any messages back to the local node, is normally active by default. However, if it has been ended by mistake, cluster messaging cannot function until this address has been restarted.

Setting TCP/IP configuration attributes

To enable cluster resource services, certain attribute settings are required in the TCP/IP configuration of your network.

You must set these attributes before you can add any node to a cluster:

- Set IP datagram forwarding to *YES by using the **CHGTCPA (Change TCP/IP Attributes)** [command](#) if you plan to use a System i product as the router to communicate with other networks and you have no other routing protocols running on that server.
- Set the INETD server to START. See [“Starting the INETD server”](#) on page 49 for information about starting the INETD server.
- Set User Datagram Protocol (UDP) CHECKSUM to *YES using the **CHGTCPA (Change TCP/IP Attributes)** command.
- Set MCAST forwarding to *YES if you are using bridges to connect your token ring networks.
- If you are using OptiConnect for IBM i to communicate between cluster nodes, start the QSOC subsystem by specifying STRSBS(QSOC/QSOC).

Starting the INETD server

The Internet Daemon (INETD) server must be started in order for a node to be added or started, as well as for merge partition processing.

It is recommended that the INETD server always be running in your cluster.

You can start the INETD server through IBM Navigator for i by completing the following steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. In the navigation tree, expand **i5/OS Management** and select **Network**.
4. On the Network page, select **TCP/IP Servers**. A list of available TCP/IP servers is displayed.
5. From the list, select **INETD**.
6. From the **Select Action** menu, select **Start**. The status of the server changes to **Started**.

Alternatively, you can start the INETD server by using the Start TCP/IP Server (STRTCPSVR) command and specifying the SERVER(*INETD) parameter. When the INETD server is started, a User QTCP (QTOGINTD) job is present in the Active Jobs list on the node.

Related reference

[STRTCPSVR \(Start TCP/IP Server\) command](#)

Configuring clusters

Any IBM i implementation of high availability requires a configured cluster to control and manage resilient resources. When used with other data resiliency technologies, such as switched disk, cross-site mirroring, or logical replication, cluster technology provides the key infrastructure that is necessary for high-availability solutions.

Cluster resource services provides a set of integrated services that maintain cluster topology, perform heartbeating monitoring, and allow creation and administration of cluster configuration and cluster resource groups. Cluster resource services also provides reliable messaging functions that keep track of each node in the cluster and ensures that all nodes have consistent information about the state of cluster resources. The Cluster Resource Service graphical user interface, which is part of the IBM PowerHA SystemMirror for i (iHASM) licensed program number (5770-HAS), allows you to configure and manage clusters within your high-availability solution. In addition, the licensed program also provides a set of control language (CL) commands that will allow you to work with cluster configurations.

There is also application program interfaces (APIs) and facilities that can be used by application providers or customers to enhance their application availability.

In addition to these IBM technologies, high availability business partners provide applications that use clusters with logical replication technology.

Creating a cluster

To create a cluster, you need to include at least one node in the cluster and you must have access to at least one of the nodes that will be in the cluster.

If only one node is specified, it must be the system that you are currently accessing. For a complete list of requirements for creating clusters, see the [“Planning checklist for clusters”](#) on page 18.

If you will be using switchable devices in your cluster or by using cross-site mirroring technologies to configure a high-availability solution, there are additional requirements. See [Scenarios: Configuring high availability solutions](#) for several configuration examples of high-availability solutions which use these technologies. Each scenario provides step-by-step configuration tasks and an overview of outage coverage this solution provides. You can use these examples to configure your high-availability solution or customize them to suit your needs.

Use the following steps to create a cluster using the PowerHA graphical interface:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.

4. On the **PowerHA** page, select **Create a new cluster**.
5. Follow the instructions in the Create Cluster wizard to create the cluster.

After you create the new cluster, the Welcome page changes to display the name of the cluster at the top of the page. The Welcome page lists several tasks for working with clusters.

After you have created a cluster you need to add any additional nodes and create CRGs.

Related information

[Create Cluster \(CRTCLU\) command](#)

[Create Cluster \(QcstCreateCluster\) API](#)

Enabling nodes to be added to a cluster

Before you can add a node to a cluster, you need to set a value for the Allow add to cluster (ALWADDCLU) network attribute.

Use the **Change Network Attributes (CHGNETA)** command on any server that you want to set up as a cluster node. The **CHGNETA** command changes the network attributes of a system. The ALWADDCLU network attribute specifies whether a node allows another system to add it as a node in a cluster.

Note: You must have *IOSYSCFG authority to change the network attribute ALWADDCLU.

Possible values follow:

***SAME**

The value does not change. The system is shipped with a value of *NONE.

***NONE**

No other system can add this system as a node in a cluster.

***ANY**

Any other system can add this system as a node in a cluster.

***RQSAUT**

Any other system can add this system as a node in a cluster only after the cluster add request has been authenticated.

The ALWADDCLU network attribute is checked to see if the node that is being added is allowed to be part of the cluster and whether to validate the cluster request through the use of X.509 digital certificates. A *digital certificate* is a form of personal identification that can be verified electronically. If validation is required, the requesting node and the node that is being added must have the following installed on the systems:

- IBM i Option 34 (Digital Certificate Manager)
- IBM i Option 35 (CCA Cryptographic Service Provider)

When *RQSAUT is selected for the ALWADDCLU, the certificate authority trust list for the IBM i cluster security server application must be correctly set up. The server application identifier is QIBM_QCST_CLUSTER_SECURITY. At a minimum, add certificate authorities for those nodes that you allow to join the cluster.

Adding nodes

The IBM i PowerHA graphical interface allows you to create a cluster with multiple nodes. After the cluster has been created you may add additional nodes through an active node in the cluster. A cluster can contain up to 128 nodes.

Use the following steps to add a node to an existing cluster using the PowerHA graphical interface:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Nodes**.
5. On the **Cluster Nodes** tab, select **Add Cluster node ...** from the **Select Action** menu.
6. Specify the required information and click **OK**.

Starting nodes

Starting a cluster node activates clustering and cluster resource services on a node in an IBM i high availability environment.

A node can start itself and is able to rejoin the current active cluster, provided it can find an active node in the cluster.

To start clustering on a node using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Nodes**.
5. On the **Cluster Nodes** page, select **Start** from the context menu of the node you want to start.

Adding a node to a device domain

A device domain is a subset of nodes in a cluster that shares device resources.

If you are implementing a high-availability solution that contains independent disk pools-based technologies, such as switched disk or cross-site mirroring, you must define the node as a member of a device domain. After you add the node to a device domain, you can create a device cluster resource group (CRG) that defines the recovery domain for the cluster. All nodes that will be in the recovery domain for a device CRG must be in the same device domain. A cluster node can belong to only one device domain.

To create and manage device domains, you must have PowerHA Option 41 (HA Switchable Resources) installed. A valid license key must exist on all cluster nodes in the device domain.

The PowerHA graphical interface simplifies device domain management by ensuring that the required nodes are in a device domain when a device CRG is created or when a node is added to a device CRG.

To add a node to a device domain using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click **Cluster Nodes**.
5. On the **Cluster Nodes** page, select **Properties ...** from the context menu of the node you want to add to a device domain.
6. Click **Edit** in the General section of the **Properties** page.
7. Specify the name of the device domain to which you want to add the node in the **Device Domain** field and click **Save**.

Creating cluster resource groups (CRGs)

Cluster resource groups (CRGs) manage high availability resources, such as applications, data, and devices. Each CRG type manages the particular type of resource in a high-availability environment.

The PowerHA graphical interface allows you to create different CRGs for management of your high availability resources. Each CRG type can be used separately or in conjunction with other CRGs. For example, you may have a stand-alone business application that requires high availability. After you have enabled the application for high availability, you can create CRGs to help manage availability for that application.

If you want only an application, not its data to be available in the event of an outage, you can create an application CRG. However, if you want to have both the data and application available, you can store both within an independent disk pool, which you can define in a device CRG. If an outage occurs, the entire independent disk pool is switched to a backup node, making both the application and its data available.

Creating application CRGs

If you have applications in your high-availability solution that you want to be highly available, you can create an application cluster resource group (CRG) to manage failovers for that application.

You can specify to allow an active takeover IP address when you create the application CRG. When you start an application CRG that allows for an active takeover IP address, the CRG is allowed to start.

To create an application CRG using the PowerHA graphical interface, complete the following steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Resource Groups**.
5. On the **Cluster Resource Group** page, select **Create Cluster Resource Group ...** from the **Select Action** menu.
6. Follow the instructions in the **Create Cluster Resource Group** wizard to create an application CRG.

Creating data CRGs

Data cluster resource groups (CRGs) are primarily used with logical replication applications, which are provided by several high availability Business Partners. If you are implementing a high-availability solution based on logical replication you can create a data CRG to assist the replication of data between primary and backup nodes.

To create a data CRG PowerHA graphical interface, complete the following steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Resource Groups**.
5. On the **Cluster Resource Group** page, select **Create Cluster Resource Group ...** from the **Select Action** menu.
6. Follow the instructions in the Create Cluster Resource Group wizard to create a data CRG.

Creating device CRGs

A device cluster resource group (CRG) is made up of a pool of hardware resources that can be switched as an entity. To create switchable devices within a high-availability solution, the nodes that use these devices need to be a part of a device CRG.

Prior to creating a device CRG, add all nodes that will share a switchable resource to a device domain. The PowerHA graphical interface simplifies this by ensuring that the required nodes are in a device domain when creating a device CRG.

To create a device CRG using the PowerHA graphical interface, complete the following steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Resource Groups**.
5. On the **Cluster Resource Group** page, select **Create Cluster Resource Group ...** from the **Select Action** menu.
6. Follow the instructions in the **Create Cluster Resource Group** wizard to create a device CRG.

Creating peer CRGs

You can create a peer CRG to define node roles in load-balancing environments.

To create a peer CRG in a cluster using thePowerHA graphical interface , complete the following steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.

3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Resource Groups**.
5. On the **Cluster Resource Group** page, select **Create Cluster Resource Group ...** from the **Select Action** menu.
6. Follow the instructions in the Create Cluster Resource Group wizard to create a peer CRG.

Starting a CRG

Starting a cluster resource group (CRG) enables resilience for the CRG.

To start a CRG using the PowerHA graphical interface, complete the following steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Resource Groups**.
5. On the **Cluster Resource Group** page, select **Start** from the context menu of the CRG you want to start.

Related information

[Start Cluster Resource Group \(STRCRG\) command](#)

[Create Cluster Resource Group \(QcstCreateClusterResourceGroup\) API](#)

Specifying message queues

You can either specify a cluster message queue or a failover message queue. These message queues help you determine causes of failures in your PowerHA environment.

A cluster message queue is used for cluster-level messages and provides one message, which controls all cluster resource groups (CRGs) failing over to a specific node. A failover message queue is used for CRG-level messages and provides one message for each CRG that is failing over.

Specifying a cluster message queue

Note: You can also configure a cluster to use a cluster message queue by specifying the message queue while running the Create Cluster wizard.

To specify a cluster message queue using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, select **Properties ...** from the **Select Action** menu.
5. Click **Edit** in the Advanced section of the **Properties** page.
6. Specify the cluster message queue information in the **Cluster Message Queue** field and click **Save**.

Specifying a failover message queue

To specify a failover message queue using the PowerHA graphical interface, complete these steps:

1. In a web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** in the IBM Navigator for i window.
4. On the **PowerHA** page, click **Cluster Resource Groups**.
5. On the **Cluster Resource Group** page, select **Properties ...** from the context menu of the CRG for which you want a failover message queue.
6. Click **Edit** in the Advanced section of the **Properties** page.
7. Specify the failover message queue information in the **Failover Message Queue** field and click **Save**.

Related information

[Create cluster](#)

[Create cluster resource group](#)

[Change cluster resource group](#)

Performing switchovers

Switchovers can be performed to test the high availability solution or to handle a planned outage for the primary node, such as a backup operation or scheduled system maintenance.

Performing a manual switchover causes the current primary node to switch over to the first backup node. The recovery domain of the cluster resource group defines these roles. When a switchover occurs, the roles of the nodes that are currently defined in the recovery domain change such that:

- The current primary node is assigned the role of last active backup.
- The current first backup is assigned the role of primary node.
- Subsequent backups are moved up one in the order of backups.

A switchover is only allowed on application, data, and device CRGs that have a status of Active.

Note: If you are performing a switchover on a device CRG, you should synchronize the user profile name, UID, and GID for performance reasons. Cluster administrative domain simplifies synchronization of user profiles.

To perform a switchover using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Click **PowerHA** from the IBM Navigator for i window.
3. On the **PowerHA** page, click on **Cluster Resource Groups**.
4. On the **Cluster Resource Groups** page, select **Switchover ...** from the context menu of the CRG that you want to switchover.
5. Click **OK** on the confirmation panel.

The selected cluster resource group is now switched to the backup node. The Status column is updated with the new node name.

There is another way to perform a switchover of a CRG that is managing an Independent ASP using the PowerHA graphical interface:

To perform a switchover using the PowerHA[®] graphical interface, follow these steps:

1. In a web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, select **Details ...** from the context menu of the Independent ASP managed by the device CRG that you want to switchover.
6. On the **Independent ASPs Details** page, select **Switch within Site...** or **Make Production Copy...** from the context menu of the Independent ASP, depending upon which kind of switchover you want to perform.
7. Click **OK** on the confirmation panel.

Related concepts

[Cluster administrative domain](#)

Related tasks

[Configuring cluster administrative domains](#)

In a high-availability environment, it is necessary that the application and operational environment remain consistent among the nodes that participate in high availability. Cluster administrative domain is

the PowerHA implementation of environment resiliency and ensures that the operational environment remains the consistent across nodes.

Related information

[Change Cluster Resource Group Primary \(CHGCRGPRI\) command](#)
[Initiate Switchover \(QcstInitiateSwitchOver\) API](#)

Configuring nodes

Nodes are systems or logical partitions that are participating in an IBM i high availability solution.

There are several tasks related to node configuration. When you use the Create Cluster wizard, you can configure a simple two-node cluster. You can add additional nodes up to a total of 128. Depending on the technologies that comprise your high-availability solution, additional node configuration tasks might be required.

Starting nodes

Starting a cluster node activates clustering and cluster resource services on a node in an IBM i high availability environment.

A node can start itself and is able to rejoin the current active cluster, provided it can find an active node in the cluster.

To start clustering on a node using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Nodes**.
5. On the **Cluster Nodes** page, select **Start** from the context menu of the node you want to start.

Related information

[Start Cluster Node \(STRCLUNOD\)\) command](#)
[Start Cluster Node \(QcstStartClusterNode\) API](#)

Enabling nodes to be added to a cluster

Before you can add a node to a cluster, you need to set a value for the Allow add to cluster (ALWADDCLU) network attribute.

Use the **Change Network Attributes (CHGNETA)** command on any server that you want to set up as a cluster node. The **CHGNETA** command changes the network attributes of a system. The ALWADDCLU network attribute specifies whether a node allows another system to add it as a node in a cluster.

Note: You must have *IOSYSCFG authority to change the network attribute ALWADDCLU.

Possible values follow:

***SAME**

The value does not change. The system is shipped with a value of *NONE.

***NONE**

No other system can add this system as a node in a cluster.

***ANY**

Any other system can add this system as a node in a cluster.

***RQSAUT**

Any other system can add this system as a node in a cluster only after the cluster add request has been authenticated.

The ALWADDCLU network attribute is checked to see if the node that is being added is allowed to be part of the cluster and whether to validate the cluster request through the use of X.509 digital certificates. A *digital certificate* is a form of personal identification that can be verified electronically. If validation is required, the requesting node and the node that is being added must have the following installed on the systems:

- IBM i Option 34 (Digital Certificate Manager)
- IBM i Option 35 (CCA Cryptographic Service Provider)

When *RQSAUT is selected for the ALWADDCLU, the certificate authority trust list for the IBM i cluster security server application must be correctly set up. The server application identifier is QIBM_QCST_CLUSTER_SECURITY. At a minimum, add certificate authorities for those nodes that you allow to join the cluster.

Adding nodes

The IBM i PowerHA graphical interface allows you to create a cluster with multiple nodes. After the cluster has been created you may add additional nodes through an active node in the cluster. A cluster can contain up to 128 nodes.

Use the following steps to add a node to an existing cluster using the PowerHA graphical interface:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Nodes**.
5. On the **Cluster Nodes** tab, select **Add Cluster node ...** from the **Select Action** menu.
6. Specify the required information and click **OK**.

Related information

Add Cluster Node Entry (ADDCLUNODE) command

[Add Cluster Node Entry \(QcstAddClusterNodeEntry\) API](#)

Adding a node to a device domain

A device domain is a subset of nodes in a cluster that shares device resources.

If you are implementing a high-availability solution that contains independent disk pools-based technologies, such as switched disk or cross-site mirroring, you must define the node as a member of a device domain. After you add the node to a device domain, you can create a device cluster resource group (CRG) that defines the recovery domain for the cluster. All nodes that will be in the recovery domain for a device CRG must be in the same device domain. A cluster node can belong to only one device domain.

To create and manage device domains, you must have PowerHA Option 41 (HA Switchable Resources) installed. A valid license key must exist on all cluster nodes in the device domain.

The PowerHA graphical interface simplifies device domain management by ensuring that the required nodes are in a device domain when a device CRG is created or when a node is added to a device CRG.

To add a node to a device domain using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click **Cluster Nodes**.
5. On the **Cluster Nodes** page, select **Properties ...** from the context menu of the node you want to add to a device domain.
6. Click **Edit** in the General section of the **Properties** page.
7. Specify the name of the device domain to which you want to add the node in the **Device Domain** field and click **Save**.

Related information

Add Device Domain Entry (ADDDEVDMNE) command

[Add Device Domain Entry \(QcstAddDeviceDomainEntry\) API](#)

Configuring advanced node failure detection

Advanced node failure detection can be used to prevent cluster partitions when a cluster node has actually failed. A Hardware Management Console (HMC) or a Virtual I/O Server (VIOS) partition on an Integrated Virtualization Manager (IVM) managed server can be used.

In this example, a HMC is being used to manage two different IBM systems. For example, the HMC can power up each system or configure logical partitions on each system. In addition, the HMC is monitoring the state of each system and logical partitions on each system. Assume that each system is a cluster node and cluster resource services is monitoring a heartbeat between the two cluster nodes.

With the advanced node failure detection function, cluster resource services can be configured to make use of the HMC. For example, Node A can be configured to have a cluster monitor that uses the HMC. Whenever HMC detects that Node B fails (either the system or the logical partition for Node B), it notifies cluster resource services on Node A of the failure. Cluster resource services on Node A then marks Node B as failed and perform failover processing rather than partitioning the cluster.

Likewise, Node B can also be configured to have a cluster monitor. In this example, then, a failure of either Node A or Node B would result in a notification from the HMC to the other node.

Advanced node failure detection implementation differs depending on the method of connection to the HMC server. Select your method of connection from the list:

- [Configuring advanced node failure detection on hardware management console \(HMC\) with REST server](#)
- [Configuring advanced node failure detection in a Virtual I/O Server \(VIOS\) on an Integrated Virtualization Manager \(IVM\) managed server](#)
- [Configuring advanced node failure detection on hardware management console \(HMC\) with CIM server.](#)

Follow the instructions in the link to configure advanced node failure detection in your network.

Configuring advanced node failure detection on hardware management console (HMC) with CIM server

A Hardware Management Console (HMC) can be used with advance node failure detection to prevent cluster partitions when a cluster node has actually failed.

Prior to configuration, verify the *CIMOM TCP and the *SSHD TCP servers are running.

1. To ensure the *CIMOM TCP server is running on your IBM i. Look for the QUMECIMOM job within the QSYSWRK subsystem to determine if it is running. If the server is not running, start it using the command **STRTCPSVR *CIMOM**.
2. Ensure the *SSHD TCP server is running on your IBM i. At the command line display, enter **STRTCPSVR *SSHD**. To start the *SSHD server, set the QSHRMEMCTL system value is set to 1.

Note: You must have access to the HMC either through the physical monitor and keyboard or remotely through a configured SSH client. You cannot access the HMC with telnet or web interface.

The *CIMOM TCP server must be configured and started on each cluster node that has a cluster monitor configured on it. The default configuration of the *CIMOM server that is provided by the installation of the 5770-UME LP must be changed so that the IBM i system can communicate with the CIM server.

Use these steps to assist you in setting up the HMC to monitor for node failures the cluster nodes to be monitored must register with the HMC through a CIM server on one of the nodes. To register the CIM and cluster require digital certification. The HMC must be used to copy this file to the cluster node with the following steps:

1. Access the HMC or connect to it with a Secure Shell (SSH) session.
2. At the HMC or in the SSH, locate the security certification file to share with your IBM i cluster node. The file is similar to, /etc/Pegasus/server.pem QSECOFR@LP0236A:/server_name.pem. Prepare to copy the file to your server with the **SCP** command. Before initiating this secure copy, change the location name from LP0236A to the name of your IBM i system. Change the file server_name.pem to the name of your HMC, for example, yourHMC.pem.
3. At the HMC or in the SSH, copy to your IBM i cluster node your modified file using the secure copy **SCP** command: **scp /etc/Pegasus/server.pem QSECOFR@YOUR_IBM_i_system:/yourHMC.pem**.

You must have a home directory associated with your profile on the IBM i. For example: if using the QSECOFR profile as the profile running the **SCP** command, you need a /home/QSECOFR directory in the integrated file system on the IBM i. Verify that you have the directory created in the correct profile.

4. Sign off the HMC or close your SSH session.

With the digital certificate on your IBM i cluster node, follow this procedure to enter the file into the truststore:

5. Sign on your IBM i system and open the command line display.
6. In the command line display, enter **call qp2term** to enter the PASE shell environment.
7. Locate the HMC digital certificate: *yourHMC.pem/QOpenSys/QIBM/UserData/UME/Pegasus/ssl/truststore/yourHMC.pem*.
8. Add this digital certificate to your IBM i system truststore with the **MOVE** command: **mv/QOpenSys/QIBM/ProdData/UME/Pegasus/bin/cimtrust -a -U QSECOFR -f /QOpenSys/QIBM/UserData/UME/Pegasus/ssl/truststore/yourHMC.pem -T s**.
9. Press F3 to exit the PASE environment.
10. On the command line display enter **ENDTCPSVR *CIMOM** to end the CIM server.

To configure the *CIMOM server to communicate with the IBM I, change the enableAuthentication and sslClientVerificationMode security settings following these steps:

11. Restart the CIM server and pick up the new certificate, enter **STRTCPSVR *CIMOM** in the command line.
12. In the command line display, enter **call qp2term** to start PASE shell and run the **CIMCONFIG** command.
13. Enter */QOpenSys/QIBM/ProdData/UME/Pegasus/bin/cimconfig -s enableAuthentication=false -p*.
14. Enter */QOpenSys/QIBM/ProdData/UME/Pegasus/bin/cimconfig -s sslClientVerificationMode=optional -p*.

These two alterations change the security configuration attributes, permitting the IBM I to communicate with the CIM server. See [Authentication on CIMOM](#) for more information about sslClientVerificationMode attribute.

15. Exit the PASE shell with F3 and end the *CIMOM server using **ENDTCPSVR *CIMOM**.
16. Restart the *CIMOM server again from the command line with the **STRTCPSVR *CIMOM**.

Installing a new version of software on the HMC partition may generate a new certificate which will then cause communication between the HMC partition and the cluster node to fail producing error CPFBBBCB with error code 4. If this occurs, add the new digital certificate to the truststore on the nodes which have that HMC or VIOS partition configured in a cluster monitor.

You are ready to perform the cluster configuration sequence using either the **ADDCLUMON** command or with the IBM Navigator for i. Follow the instructions in the topic, [Add a cluster monitor to a node](#). For additional information about the **ADDCLUMON**, see the [Add Cluster Monitor \(ADDCLUMON\)](#) command in the Knowledge Center.

Configuring advanced node failure detection on hardware management console (HMC) with REST server

A Hardware Management Console (HMC) can be used with advanced node failure detection to prevent cluster partitions when a cluster node has actually failed.

1. Using HMC with a Representational state transfer (REST) server requires HMC version V8R8.5.0 or greater.
2. The Add cluster monitor (**ADDCLUMON**) command must be used with the representational state transfer (REST) server. The PowerHA® graphical interface only supports the Common Informational Model (CIM) server for the cluster monitor.
3. Check the QSSLPCL system value. Verify that it is set correctly for the release currently running.

Note: An incorrect value in QSSLPCL may result in a CPFBBBCB diagnostic message with reason code 4.

These steps guide you through obtaining the digital certificate of your HMC, storing it and referencing it to allow advanced node failure detection for the cluster node.

Important: This guide describes steps making use of features of both HMC and of the Digital Certificate Manager. Changes to either of these products may cause portions of this guide to become invalid. If you suspect such changes are preventing you from following the steps outlined in this guide successfully, contact your technical support provider.

Begin by extracting the digital certificates for the HMC and copying them to the IBM i system in the cluster node with these steps:

1. Sign on your IBM i system and open the command line display.
2. In the command line display, enter **CALL QP2TERM** to enter the PASE shell environment.
3. Retrieve the digital certificates from the HMC with this command:

```
openssl s_client -connect HMC_name:443 < /dev/null | awk '/BEGIN/,/END/{ if(/BEGIN/){a++;  
out="HMC_name"a".pem"; print > out}}
```

Replace *HMC_name* with the name of your system's HMC. This copies the certificates into files named *HMC_name1.pem* ... *HMC_nameN.pem*, where *N* is the number of certificates copied from your system's HMC.

4. Press F3 to exit the **QP2TERM** environment.

Create a certificate store to hold the digital certificates by following these steps:

5. Open the IBM Navigator for i and click **Internet Configurations**.
6. On the **Internet Configurations** page, click **Digital Certificate Manager**.
You need to enter your user profile and password.
7. In the **Digital Certificate Manager** page, click **Create New Certificate Store**.
8. In the page that appears, you should have an option for ***SYSTEM**. Make sure that the button is selected and click **Continue**.
If the ***SYSTEM** option is not there, you already have a ***SYSTEM** store created. Skip forward to step 12.
9. Select **No - Do not create a certificate in the certificate store**.
10. Create a password for the ***SYSTEM** store and click **Continue**.

The password is case-sensitive. It is recommended not to use special characters. This password is not attached to a user profile and it will not lock you out of the system after too many retries.

You have successfully created the ***SYSTEM** store.

Select the ***SYSTEM** certificate store by following these steps:

11. Click **Select a Certificate Store** and select the ***SYSTEM** option, click **continue**.
12. Sign in with the password for the certificate store and click **Continue**, then **Manage Certificates**.

Import the HMC certificates into the security store.

13. Select **Import certificate** and click **Continue**.

If your HMC has only one certificate, perform these steps for that certificate. If your HMC has multiple certificates, perform these steps for each certificate except the first certificate (*HMC_name1.pem*), starting with the last certificate and moving backwards through the list of certificates. For example, if there are three certificates: *HMC_name1.pem*, *HMC_name2.pem*, and *HMC_name3.pem*, perform these steps for *HMC_name3.pem* first, then for *HMC_name2.pem*.

14. Select **Certificate Authority (CA)** and click **Continue**.
15. Enter the path name of the certificate you want to import. For example, the path and file name may be */HMC_name1.pem*. Click **Continue**.

After importing the certificates, sign on to your IBM i and use the command line to run the Add cluster monitor (**ADDCLUMON**) command to run the cluster configuration steps. For additional information about **ADDCLUMON**, see the [Add Cluster Monitor \(ADDCLUMON\)](#) command in the Knowledge Center.

Configuring advanced node failure detection in a Virtual I/O Server (VIOS) on an Integrated Virtualization Manager (IVM) managed server

A Virtual I/O Server (VIOS) partition on an Integrated Virtualization Manager (IVM) managed server can be used with advanced node failure detection to prevent cluster partitions when a cluster node has actually failed.

Prior to configuration, verify the *CIMOM TCP and the *SSHD TCP servers are running.

1. To ensure the *CIMOM TCP server is running on your IBM i. Look for the QUMECIMOM job within the QSYSWRK subsystem to determine if it is running. If the server is not running, start it using the command **STRTCPSVR *CIMOM**.
2. Ensure the *SSHD TCP server is running on your IBM i. At the command line display enter: **STRTCPSVR *SSHD** to start the *SSHD server. Verify the **QSHRMEMCTL** system value is set to 1.

The representational state transfer (REST) server is not supported by IVM.

To implement advanced node failure detection in a VIOS partition on an IVM managed server begin by obtaining and copying the digital certificate to your IBM i system. The certificate requires further configuration before adding detection monitors.

Use these steps to help assist you in setting up the VIOS partition to monitor for node failures. To register the cluster node you need to obtain the digital certificate and have the VIOS copy the file to the IBM i:

1. Open a telnet session and sign on to the VIOS partition.
2. Switch to a non-restricted shell by entering `oem_setup_env`.
3. Locate the security certification file to share with your IBM i cluster node. The file is similar to, `/usr/bin/scp /opt/freeware/cimom/pegasus/etc/cert.pem QSECOFR@system-name:/server.pem`. Prepare to copy the file to your server by changing `system-name` to the name of your IBM i system. Change the file `server.pem` to the name of your VIOS partition, for example, `yourVIOS.pem`.

Your file will resemble: `/usr/bin/scp /opt/freeware/cimom/pegasus/etc/cert.pem QSECOFR@YOUR_IBM_i_system:/yourVIOS.pem`

4. Copy to your IBM i cluster node your modified file using the secure copy **SCP** command: `/usr/bin/scp /opt/freeware/cimom/pegasus/etc/cert.pem QSECOFR@YOUR_IBM_i_system:/yourVIOS.pem`.
5. Start the CIMOM server running in the VIOS partition by entering **startnetsvc cimserver**.
6. Sign off of the VIOS partition.

With the digital certificate on your IBM i cluster node, follow this procedure to enter the file into the truststore:

7. Sign on your IBM i system and open the command line display.
8. In the command line display, enter **call qp2term** to enter the PASE shell environment.
9. Locate the VIOS digital certificate: `yourVIOS.pem` and add this digital certificate to your IBM i system trusts with the **MOVE** command: `mv yourVIOS.pem/QOpenSys/QIBM/ProdData/UME/Pegasus/bin/cimtrust -a -U QSECOFR -fyourVIOS.pem -T s`.
10. Press F3 to exit the PASE environment.
11. On the command line display enter **ENDTCPSVR *CIMOM** to end the CIM server.

To configure the *CIMOM server to communicate with the IBM I, change the enableAuthentication and sslClientVerificationMode security settings following these steps:

12. Restart the CIM server and pick up the new certificate with **STRTCPSVR *CIMOM** in the command line.
13. At the command line display, enter **call qp2term** to start PASE shell and run the **CIMCONFIG** command.
14. Enter `/QOpenSys/QIBM/ProdData/UME/Pegasus/bin/cimconfig -s enableAuthentication=false -p`.
15. Enter `/QOpenSys/QIBM/ProdData/UME/Pegasus/bin/cimconfig -s sslClientVerificationMode=optional -p`.

These two alterations change the security configuration attributes, permitting the IBM I to communicate with the CIM server. See [Authentication on CIMOM](#) for more information about `sslClientVerificationMode` attribute.

16. Exit the PASE shell with F3 and end the *CIMOM server using **ENDTCPSVR *CIMOM**.

17. Restart the *CIMOM server again from the command line with the **STRTCPSVR *CIMOM**.

Installing a new version of software on the HMC partition may generate a new certificate which will then cause communication between the HMC partition and the cluster node to fail producing error CPFBBBCB with error code 4. If this occurs, add the new digital certificate to the truststore on the nodes which have that HMC or VIOS partition configured in a cluster monitor.

You are ready to perform the cluster configuration sequence using either the **ADDCLUMON** command or with the IBM Navigator for i. Follow the instructions in the topic, [Add a cluster monitor to a node](#). For additional information about the **ADDCLUMON**, see the [Add Cluster Monitor \(ADDCLUMON\)](#) command in the Knowledge Center.

Configuring CRGs

Cluster resource groups (CRGs) manage resources within an IBM i high availability environment. Several tasks enable the management of high availability resources through CRGs.

Starting a CRG

Starting a cluster resource group (CRG) enables resilience for the CRG.

To start a CRG using the PowerHA graphical interface, complete the following steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Resource Groups**.
5. On the **Cluster Resource Group** page, select **Start** from the context menu of the CRG you want to start.

Related information

[Start Cluster Resource Group \(STRCRG\) command](#)

[Create Cluster Resource Group \(QcstCreateClusterResourceGroup\) API](#)

Creating cluster resource groups (CRGs)

Cluster resource groups (CRGs) manage high availability resources, such as applications, data, and devices. Each CRG type manages the particular type of resource in a high-availability environment.

The PowerHA graphical interface allows you to create different CRGs for management of your high availability resources. Each CRG type can be used separately or in conjunction with other CRGs. For example, you may have a stand-alone business application that requires high availability. After you have enabled the application for high availability, you can create CRGs to help manage availability for that application.

If you want only an application, not its data to be available in the event of an outage, you can create an application CRG. However, if you want to have both the data and application available, you can store both within an independent disk pool, which you can define in a device CRG. If an outage occurs, the entire independent disk pool is switched to a backup node, making both the application and its data available.

Creating application CRGs

If you have applications in your high-availability solution that you want to be highly available, you can create an application cluster resource group (CRG) to manage failovers for that application.

You can specify to allow an active takeover IP address when you create the application CRG. When you start an application CRG that allows for an active takeover IP address, the CRG is allowed to start.

To create an application CRG using the PowerHA graphical interface, complete the following steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.

2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Resource Groups**.
5. On the **Cluster Resource Group** page, select **Create Cluster Resource Group ...** from the **Select Action** menu.
6. Follow the instructions in the **Create Cluster Resource Group** wizard to create an application CRG.

Related information

[Create Cluster Resource Group \(CRTCRG\) command](#)

[Create Cluster Resource Group \(QcstCreateClusterResourceGroup\) API](#)

Creating data CRGs

Data cluster resource groups (CRGs) are primarily used with logical replication applications, which are provided by several high availability Business Partners. If you are implementing a high-availability solution based on logical replication you can create a data CRG to assist the replication of data between primary and backup nodes.

To create a data CRG PowerHA graphical interface, complete the following steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Resource Groups**.
5. On the **Cluster Resource Group** page, select **Create Cluster Resource Group ...** from the **Select Action** menu.
6. Follow the instructions in the Create Cluster Resource Group wizard to create a data CRG.

Related information

[Create Cluster Resource Group \(CRTCRG\) command](#)

[Create Cluster Resource Group \(QcstCreateClusterResourceGroup\) API](#)

Creating device CRGs

A device cluster resource group (CRG) is made up of a pool of hardware resources that can be switched as an entity. To create switchable devices within a high-availability solution, the nodes that use these devices need to be a part of a device CRG.

Prior to creating a device CRG, add all nodes that will share a switchable resource to a device domain. The PowerHA graphical interface simplifies this by ensuring that the required nodes are in a device domain when creating a device CRG.

To create a device CRG using the PowerHA graphical interface, complete the following steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Resource Groups**.
5. On the **Cluster Resource Group** page, select **Create Cluster Resource Group ...** from the **Select Action** menu.
6. Follow the instructions in the **Create Cluster Resource Group** wizard to create a device CRG.

Related information

[Create Cluster Resource Group \(CRTCRG\) command](#)

[Create Cluster Resource Group \(QcstCreateClusterResourceGroup\) API](#)

Creating peer CRGs

You can create a peer CRG to define node roles in load-balancing environments.

To create a peer CRG in a cluster using the PowerHA graphical interface, complete the following steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Resource Groups**.
5. On the **Cluster Resource Group** page, select **Create Cluster Resource Group ...** from the **Select Action** menu.
6. Follow the instructions in the Create Cluster Resource Group wizard to create a peer CRG.

Related information

[Create Cluster Resource Group \(CRTCRG\) command](#)

[Create Cluster Resource Group \(QcstCreateClusterResourceGroup\) API](#)

Configuring cluster administrative domains

In a high-availability environment, it is necessary that the application and operational environment remain consistent among the nodes that participate in high availability. Cluster administrative domain is the PowerHA implementation of environment resiliency and ensures that the operational environment remains the consistent across nodes.

Creating a cluster administrative domain

In a high-availability solution, the cluster administrative domain provides the mechanism that keeps resources synchronized across systems and partitions within a cluster.

To create the cluster administrative domain, a user must have *IOSYSCFG authority and authority to the QCLUSTER user profile. To manage a cluster administrative domain, a user must be authorized to the CRG that represents the cluster administrative domain, the QCLUSTER user profile, and cluster resource group commands.

To create a cluster administrative domain using the PowerHA graphical interface, complete the following steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Administrative Domains**.
5. On the **Cluster Administrative Domain** page, select **Create Administrative Domain ...** from the Select Action menu.
6. Specify the cluster administrative domain and click **OK**.

Related concepts

[Maintaining user profiles on all nodes](#)

You can use two mechanisms for maintaining user profiles on all nodes within a cluster.

Related information

[Create Cluster Administrative Domain \(CRTCAD\) command](#)

[Create Cluster Administrative Domain \(QcstCrtClusterAdminDomain\) API](#)

Adding a node to the cluster administrative domain

You can add additional nodes to a cluster administrative domain within a high-availability solution.

Before adding a node to a cluster administrative domain, ensure that node is also part of the cluster in which the cluster administrative domain resides. If it is not, you cannot add the node to the cluster administrative domain. The cluster administrative domain does not have to be active, but the resources will just not be made consistent until it is active.

When you add a node to the administrative domain, the MREs from the domain are copied to the node being added. If the monitored resource does not exist on the new node, it is created by the cluster administrative domain. If the monitored resource already exists on the node being added, it is synchronized with the rest of the cluster administrative domain if the domain is active. That is, the values

of the attributes for each monitored resource on the node that is joining are changed to match the global values for the monitored resources in the active domain.

To add a node to a cluster administrative domain using the PowerHA graphical interface, follow these steps:

1. Log on to the system with your user profile and password.
2. Click **PowerHA** from the IBM Navigator for i window.
3. On the **PowerHA** page, click on **Cluster Administrative Domains**.
4. On the **Administrative Domains** page, select **Domain Nodes ...** from the context menu of the cluster administrative domain to which you want to add a node.
5. On the **Domain Nodes** page, select **Add Domain Node ...** from the Select Action menu.
6. Select the node you want to add and click **OK**.

Related information

[Add Cluster Administrative Domain Node Entry \(ADDCADNODE\) command](#)

[Add Node To Recovery Domain \(QcstAddNodeToRcvyDomain\) API](#)

Starting a cluster administrative domain

Cluster administrative domains provide environment resiliency for resources within an IBM i high-availability solution.

When the cluster administrative domain is started, any change made to any monitored resource while the cluster administrative domain was ending is propagated to all active nodes in the cluster administrative domain.

To start a cluster administrative domain using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Administrative Domains**.
5. On the **Cluster Administrative Domains** page, select **Start** from the context menu of the cluster administrative domain you want to start.

The Status column shows that the cluster administrative domain is started.

Related concepts

[Synchronization of monitored resource](#)

Synchronization of monitored resources occurs when monitored resources are changed on nodes which have been defined in the cluster administrative domain.

Related information

[Start Cluster Administrative Domain \(STRCAD\) command](#)

Synchronization of monitored resource

Synchronization of monitored resources occurs when monitored resources are changed on nodes which have been defined in the cluster administrative domain.

During this synchronization process, the cluster administrative domain attempts to change each resource with attributes whose values do not match its global values, unless there is a pending change for that resource. Any pending change is distributed to all active nodes in the domain and applied to each affected resource on each node. When the pending changes are distributed, the global value is changed and the global status of each affected resource is changed to *consistent* or *inconsistent*, depending on the outcome of the change operation for the resource on each node. If the affected resource is changed successfully on every active node in the domain, the global status for that resource is *consistent*. If the change operation failed on any node, the global status is set to *inconsistent*.

If changes are made to the same resource from multiple nodes while the cluster administrative domain is inactive, all of the changes are propagated to all of the active nodes as part of the synchronization process when the domain is started. Although all pending changes are processed during the activation of the

cluster administrative domain, there is no guaranteed order in which the changes are processed. If you make changes to a single resource from multiple cluster nodes while the cluster administrative domains are inactive, there is no guaranteed order to the processing of the changes during activation.

If a node joins an inactive cluster administrative domain (that is, the node is started while the cluster administrative domain is ended), the monitored resources are not resynchronized until the cluster administrative domain is started.

Note: The cluster administrative domain and its associated exit program are IBM-supplied objects. They should not be changed with the QcstChangeClusterResourceGroup API or the Change Cluster Resource Group (CHGCRG) command, or unpredictable results will occur.

After a cluster node that is part of a cluster administrative domain is ended, monitored resources can still be changed on the inactive node. When the node is started again, the changes will be resynchronized with the rest of the cluster administrative domain. During the resynchronization process, the cluster administrative domain applies any changes from the node that was inactive to the rest of the active nodes in the domain, unless changes had also been made in the active domain while the node was inactive. If changes were made to a monitored resource both in the active domain and on the inactive node, the changes made in the active domain are applied to the joining node. In other words, no changes made to any monitored resource are lost, regardless of the status of the node. You can specify the synchronization option to control synchronization behavior.

If you want to end a cluster node that is part of a cluster administrative domain, and not allow changes made on the inactive node to be propagated back to the active domain when the node is started (for example, when ending the cluster node to do testing on it), you must remove the node from the administrative domain peer CRG before you end the cluster node.

Related concepts

[Remove Admin Domain Node Entry \(RMVCADNODE\) command](#)

Related tasks

[Starting a cluster administrative domain](#)

Cluster administrative domains provide environment resiliency for resources within an IBM i high-availability solution.

Related information

[Remove CRG Node Entry \(RMVCRGNODE\) command](#)

Adding monitored resource entries

You can add a monitored resource entry (MRE) to a cluster administrative domain. Monitored resource entries define critical resources so that changes made to these resources are kept consistent across a high-availability environment.

To add a monitored resource entry using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click on **PowerHA** in the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Administrative Domains**.
5. On the **Cluster Administrative Domains** page, select **Monitored Resource ...** from the context menu of the cluster administrative domain to which you want to add the monitored resource.
6. On the **Monitored Resource** page, select **Add Monitored Resource(s) ...** from the **Select Action** menu.
7. Specify the monitored resource information and click **OK**.

Related tasks

[Selecting attributes to monitor](#)

After you have added monitored resource entries, you can select attributes associated with that resource to be monitored by the cluster administrative domain.

Related information

[Add Admin Domain MRE \(ADDCADMRE\) command](#)

Configuring independent disk pools

In a high-availability environment, it is necessary that the application and its data remain consistent among the nodes that participate in high availability. An independent disk pool is a disk pool that contains objects, the directories, or libraries that contain the objects, and other object attributes such as authorization and ownership attributes. An independent disk pool can be used to help in this situation.

Related information

[Disk pools](#)

[IBM eServer iSeries Independent ASPs: A Guide to Moving Applications to IASPs](#)

Creating an independent disk pool

To create an independent disk pool, you can use the PowerHA graphical interface, the Configuration and Service graphical interface, or the Configure Device ASP (CFGDEVASP) command.

To use the PowerHA graphical interface to create a new independent disk pool, follow these steps:

1. In a web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click on **PowerHA** in the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Independent ASPs**.
5. On the **Independent ASPs** page, click on **Show All Others ...** to show any independent disk pools that are not currently highly available.
6. Select **Create Independent ASP ...** from the **Select Action** menu in the All Other table.
7. Follow the instructions in the **Create Independent ASP** wizard to create an independent disk pool.

To use the Configuration and Service graphical interface to create a new independent disk pool, follow these steps:

Note:

To work with disk within the IBM Navigator for i, you must have the appropriate password configuration for Dedicated Service Tools. With the New Disk Pool wizard you can include unconfigured disk units in a parity set, and you can start device parity protection and disk compression. As you add disk units, do not spread disk units that are in same parity set across multiple disk pools, because failure to one parity set would affect multiple disk pools.

Configuration and Service graphical interface (IBM Navigator for i)

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Select **Configuration and Service** from your IBM Navigator for i window.
4. Select **Disk Units**.
5. From the **Select Actions** menu, select **New Disk Pool**.
6. Follow the wizard's instructions to add disk units to a new disk pool.
7. Print your disk configuration to have it available in a recovery situation.
8. Record the relationship between the independent disk pool name and number.

The Configure Device ASP (CFGDEVASP) command can be used to create a new independent disk pool. For more information on the command, see [CFGDEVASP \(Configure Device ASP\) command](#).

Note: Add independent disk pools when your system is fully restarted. If you must use the New Disk Pool wizard in the dedicated service tools (DST) mode, you need to create an associated device description for the independent disk pool when the system is fully restarted. Use the Create Device Description (ASP) (CRTDEVASP) command to create the device description. Name the device description and resource name the same as you name the independent disk pool. You can use the Work with Device Descriptions (WRKDEVVD) command to verify that the device description and independent disk pool name match.

Starting mirrored protection

The Add Disk Unit and New Disk Pool wizards guide you through the process of adding pairs of similar capacity disk units to a protected disk pool. After your disk units are configured correctly, you are ready to start mirrored protection.

Mirrored protection is local to a single system and is distinct from cross-site mirroring. If you want to start mirroring on an independent disk pool that is unavailable, you can do so when your system is fully restarted. For all other disk pools, you must restart your system to dedicated service tools (DST) mode before you start mirror-protection.

There are restrictions to be aware of when you start mirror-protection on the disk units of a disk pool.

- Disk units that are device parity protected cannot be mirrored.
- Disk units must be mirrored to disk units that have the identical sector size.
- Disk units must be mirrored to disk units that have approximately the same capacity.

There are restrictions to follow when you start mirror-protection on the load source disk unit.

- The smaller capacity disk must start out as the load source device when two disks of unequal capacity are matched as a mirrored pair. The load source can then be matched with the larger capacity disk unit. For example, if the load source disk unit is a 35 G-Byte disk, it can be matched with a 36 GB disk. If the load source is a 36 G-Byte disk, it cannot be matched with the 35 G-Byte disk.
- The system must be directed to match the load source disk unit with a disk unit that is in a physical location the service processor cannot use to IPL the partition. From SST, select **Work with disk units->Work with disk configuration->Enable remote load source mirroring**. The **Enable remote load source mirroring** function allows a disk unit to be matched with the load source disk unit even though the disk unit is in a physical location the service processor is not able to use to IPL the partition.

To start mirroring using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Pools**.
3. Select the disk pool that you want to mirror.
4. Right-click the disk pool that you want to mirror, and select **Start Mirroring**.

Stopping mirrored protection

When you stop mirrored protection, one disk unit from each mirrored pair is unconfigured. Before you can stop mirrored protection for a disk pool, at least one disk unit in each mirrored pair in that disk pool must be present and active.

To control which mirrored disk unit of each pair is unconfigured, you suspend the disk units that you want to become unconfigured. For disk units that are not suspended, the selection is automatic.

If you want to stop mirroring on an independent disk pool that is unavailable, you can do so when your system is fully restarted. For all other disk pools, you need to restart your system to the dedicated service tools (DST) mode before stopping mirrored protection.

Mirrored protection is dedicated to a single system, and is distinct from cross-site mirroring.

To stop mirrored protection using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Pools**.
3. Right-click the disk pool that you want to stop.
4. Select **Stop Mirroring**.

Adding a disk unit or disk pool

The Add Disk Unit wizard can be used to add new or non-configured disk units to an existing disk pool.

The Add Disk Unit and Disk Pool wizards save you time by bundling several time-consuming configuration functions into one efficient process. They also take the guesswork out of disk unit configuration because

they understand the capabilities of your system and only offer valid choices. For instance, the wizard does not list the option to start compression unless your system has that capability.

When you choose to add disk units to a protected disk pool, the wizard forces you to include the disk units in device parity protection or to add enough disk units of the same capacity to start mirrored protection. The wizard also gives you the option of balancing data across the disk pool or starting disk compression if these options are permissible actions for your system configuration. You decide which options to choose so that the operation is tailored to your system.

To add a disk unit or disk pool using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Units**.
3. From the **Actions** menu, select **Add Disk Unit**.
4. Follow the wizard's instructions to add disk units to your disk pool.

You can also add a disk unit to a disk pool using your IBM Navigator for i by following these steps:

1. Select **Configuration and Service** from your IBM Navigator for i
2. Select **Disk Pools**.
3. Right-click the disk pool that you want to add a disk unit to and select **Add Disk Unit**.
4. Follow the wizard's instructions to add disk units to your disk pool.

Note: You must perform an initial program load (IPL) of the system before the new disk pool is available through the Integrated File System (IFS).

Evaluating the current configuration

Before you change the disk configuration of your system, it is important to know exactly where the existing disk units are located in relation to disk pools, IOAs, and frames.

The graphical view of IBM Navigator for i eliminates the process of compiling all this information by providing a graphical representation of how your system is configured. You can use the graphical view to perform any function that is possible through the Disk Units list view of IBM Navigator for i, with the added benefit of being able to see a visual representation. If you right-click any object in the table, such as a specific disk unit, disk pool, parity set, or frame, you see the same options as in the main IBM Navigator for i table.

You can choose how to view the hardware in the Disk Unit Graphical View window. For example, you can select to view by disk pools, and then select a disk pool in the list to display only those frames that contain the disk units that make up the selected disk pool. You can select Show all frames to see whether or not they contain disk units in the selected disk pool. You can also select Show device positions to associate disk unit names with the device position where they are inserted.

You can right-click any highlighted blue disk unit in the graphical view and select an action to perform on the disk unit. For example, you can select to start or stop compression on a disk unit, include or exclude the disk unit in a parity set, or rename the disk unit. If the disk unit has mirrored protection, you can suspend or resume mirroring on the disk unit. If you right-click an empty disk unit slot, you can start the Install Disk Unit wizard.

To activate the graphical view from your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Units** or **Disk Pools**.
3. From the **Actions** menu, select **Graphical View**.

Making an independent disk pool highly available

To make an independent disk pool highly available, it becomes managed by clustering using a device cluster resource group.

To use the PowerHA graphical interface to make an independent disk pool highly available, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** in the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs Details** page, click **Make Highly Available...** from the context menu of the Independent ASP that you want to make highly available.
6. Follow the instructions in the **Make Highly Available** wizard to make the independent disk pool highly available.

Configuring geographic mirroring

Geographic mirroring is a subfunction of cross-site mirroring. To configure a high-availability solution by using geographic mirroring, you need to configure a mirroring session between the production system and the backup system.

Before configuring geographic mirroring, you must have an active cluster, nodes, and CRG. The independent disk pools which you plan to use for geographic mirroring must also be varied off (unavailable) to complete configuration. The topic, [Scenario: Geographic mirroring](#), provides step-by-step instructions for setting up a high-availability solution that is based on geographic mirroring.

To use the PowerHA graphical interface to configure geographic mirroring follow these steps:

1. In a web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** in the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, select **Details ...** from the context menu of the independent disk pool for which you want to configure geographic mirroring.
6. On the **Independent ASP Details** page, select **Configuring Mirroring ...** from the context menu of the production copy of the independent disk pool.
7. Follow the instructions in the **Create Independent ASP** wizard to create an independent disk pool.

To use the Configuration and Service graphical interface to configure geographic mirroring, follow the steps below. The Configuration and Service graphical interface does not show up if IBM PowerHA SystemMirror for i licensed program is installed.

Configuration and Service graphical interface (IBM Navigator for i)

1. Log on to the system with your user profile and password.
2. Log on to the system with your user profile and password.
3. Select **Configuration and Service** from your IBM Systems Director Navigator for i window.
4. Select **Disk Pools**.
5. Select the disk pool that you want to use as the production (source) copy.
6. On the **Independent ASP Details** page, select **Configure Mirroring...** from the context menu of the production copy of the independent disk pool.
7. Follow the instructions in the **Configure Mirroring** wizard to configure geographic mirroring.


Related information

[Configure Geographic Mirror \(CFGGEOMIR\) command](#)

Configuring Metro Mirror

For IBM i high availability solutions that use IBM System Storage Metro Mirror technology, you need to configure a session between the IBM i machine and IBM System Storage external storage units that have Metro Mirror configured. In IBM i, Metro Mirror sessions do not set up the mirroring on the external storage units, but rather sets up a relationship between the IBM i systems and the existing Metro Mirror configuration on external storage units.

Before creating a Metro Mirror session on IBM i, you should have configured Metro Mirror on the IBM System Storage external storage units.

- For information about using Metro Mirror on IBM System Storage DS8000, see [Copy Services in the IBM System Storage DS8000 Information Center](#) .
- For information about using Metro Mirror with SAN Volume Controller (SVC) or Storwize storage devices, see [IBM SAN Volume Controller Information Center](#).

To use the PowerHA graphical interface to configure Metro Mirror, follow these steps:

1. In a web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** in the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, select **Details ...** from the context menu of the independent disk pool for which you want to configure Metro Mirror.
6. On the **Independent ASP Details** page, select **Configure Mirroring ...** from the context menu of the production copy of the independent disk pool.
7. Follow the instructions in the **Configure Mirroring** wizard to configure Metro Mirror.

To use the Configuration and Service graphical interface to configure Metro Mirror session, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Select **Configuration and Service** from your IBM Navigator for i window.
4. Select **Disk Pools**.
5. Select the disk pool that you want to use as the production (source) copy.
6. From the **Select Actions** menu, select **New Session**.
7. Follow the wizard's instructions to complete the task.

Related concepts

[PowerHA supported storage servers](#)

IBM System Storage provides enhanced storage capabilities.

Related information


[Add ASP Copy Description \(ADDASPCPYD\) command](#)

[Start ASP Session \(STRASPSSN\) command](#)

Configuring Global Mirror

For IBM i high availability solutions that use IBM System Storage Global Mirror technology, you need to configure a session between the IBM i machine and IBM System Storage external storage units that have Global Mirror configured. In IBM i, Global Mirror sessions do not set up the mirroring on the external storage units, but rather sets up a relationship between the IBM i systems and the existing Global Mirror configuration on external storage units.

Before creating a Global Mirror session on IBM i, you should have configured Global Mirror on the IBM System Storage external storage units.

- For information about using Global Mirror on IBM System Storage DS8000, see [Copy Services in the IBM System Storage DS8000 Information Center](#) .
- For information about using Global Mirror with SAN Volume Controller (SVC) or Storwize storage devices, see [IBM SAN Volume Controller Information Center](#).

To use the PowerHA graphical interface to create a new independent disk pool, follow these steps:

1. In a web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.

2. Log on to the system with your user profile and password.
3. Click **PowerHA** in the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, select **Details ...** from the context menu of the independent disk pool for which you want to configure Global Mirror.
6. On the **Independent ASPs Details** page, select **Configure Mirroring ...** from the context menu of the production copy of the independent disk pool.
7. Follow the instructions in the **Create Independent ASP** wizard to create an independent disk pool.

To configure global mirroring, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Select **Configuration and Service** from your IBM Navigator for i window.
4. Select **Disk Pools**.
5. Select the disk pool that you want to use as the production (source) copy.
6. From the **Select Actions** menu, select **New Session**.
7. Follow the wizard's instructions to complete the task.

Related concepts

[PowerHA supported storage servers](#)

IBM System Storage provides enhanced storage capabilities.

Related information

[Add ASP Copy Description \(ADDASPCPYD\) command](#)

[Start ASP Session \(STRASPSSN\) command](#)

Configuring switched logical units (LUNs)

Switched logical units is an independent disk pool that is controlled by a device cluster resource group and can be switched between nodes within a cluster. For IBM i high availability solutions that use switchable logical units with IBM® System Storage™, you need to configure an ASP Copy Description that defines the host identifiers and volume groups for each cluster node in which the ASP device can be switched.

In IBM i, the creation of an ASP Copy Description for switchable logical units does not set up the host identifiers and volume groups on the external storage units, but rather sets up a relationship between the IBM i systems and existing host identifiers and volume groups on the external storage units.

Before creating a copy description for switchable logical units on IBM i, you should have configured the logical units on the IBM System Storage external storage unit as well as created the independent auxiliary storage pool and device cluster resource group that will control the switching. For information about using the IBM System Storage DS8000®, see IBM System Storage DS8000 Information Center.

For details on using other storage technologies, see [“PowerHA supported storage servers” on page 37](#).

To configure switchable logical units, follow these steps:

1. Enter the Add ASP Copy Description (**ADDASPCPYD**) command at an IBM i command prompt.
2. Enter the appropriate names for ASP device, cluster resource group, cluster resource group site, IBM System Storage host, location, and logical unit name.
3. Enter the appropriate cluster node names, host identifiers, and volume groups in the recovery domain field.
4. Create an ASP Copy Description for every ASP device that is to use switchable logical units. All ASP devices in the cluster resource group must be defined by a separate ASP copy description.

Related concepts

[PowerHA supported storage servers](#)

IBM System Storage provides enhanced storage capabilities.

Configuring a FlashCopy session

For IBM i high-availability environments that use IBM System Storage technology, you can configure a FlashCopy session to create a point-in-time copy of data.

For information on using the FlashCopy feature on IBM System Storage, see [“PowerHA supported storage servers” on page 37](#).

Configuring a FlashCopy when IBM PowerHA SystemMirror for i licensed program is installed

To configure FlashCopy by using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** in the PowerHA window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, select **Details...** from the context menu of the independent disk pool for which you want to create a FlashCopy.
6. On the **Independent ASP Details** page, select **Create FlashCopy...** from the context menu of which you want to create the FlashCopy.
7. Follow the instructions in the Create FlashCopy wizard to create a FlashCopy.

Related concepts

[PowerHA supported storage servers](#)

IBM System Storage provides enhanced storage capabilities.

Configuring DS8000 Full System HyperSwap

For System i high availability solutions using HyperSwap, detection of the HyperSwap configuration is automatic on the System i machine after appropriate configuration has been performed on the IBM System Storage external storage units.

The steps assume knowledge of configuring and managing an IBM System Storage DS8000. For information about using an IBM System Storage DS8000 unit, see [Copy Services in the IBM System Storage DS8000 Information Center](#).

To configure HyperSwap, follow these steps:

1. Create LUNs on the source DS8000.
2. Assign LUNs on source DS8000 to the System i machine.
3. Create LUNs on the target DS8000. Do not assign them to the System i until after Metro Mirror is configured.
4. Configure Metro Mirror between the source and target DS8000.
5. Assign the LUNs on the target DS8000 to the same System i machine.
6. Verify the HyperSwap configuration with the Display HyperSwap Status (DSPHYSSTS) command.

To understand how HyperSwap interfaces with affinity, see [“Defining HyperSwap Affinity” on page 73](#)

Related tasks

[Defining HyperSwap Affinity](#)

By using DS8000 HyperSwap combined with logical unit (LUN) level switching and or PowerVM Live Partition Mobility, it is possible to define which IBM i systems should use which IBM System Storage DS8000 as their primary. Upon accomplishing a live partition mobility switch or a LUN switching CRG switch, if the partition is no longer running on the preferred IBM System Storage unit, a HyperSwap switchover to the secondary IBM System Storage unit is done as part of the switching process.

Related information

[Change HyperSwap Status \(CHGHYSSTS\) command](#)

[Display HyperSwap Status \(DSPHYSSTS\) command](#)

[Live Partition Mobility](#)

Defining HyperSwap Affinity

By using DS8000 HyperSwap combined with logical unit (LUN) level switching and or PowerVM Live Partition Mobility, it is possible to define which IBM i systems should use which IBM System Storage DS8000 as their primary. Upon accomplishing a live partition mobility switch or a LUN switching CRG switch, if the partition is no longer running on the preferred IBM System Storage unit, a HyperSwap switchover to the secondary IBM System Storage unit is done as part of the switching process.

To define which DS8000 unit can be used with which IBM i systems, HA configuration descriptions are used. The Add HA Configuration Description (ADDHACFGD) command can be used to define an affinity relationship between the serial number of the server that is hosting the IBM i partition and the IBM System Storage unit identifier of the DS8000 storage server.

Related tasks

[Configuring DS8000 Full System HyperSwap](#)

For System i high availability solutions using HyperSwap, detection of the HyperSwap configuration is automatic on the System i machine after appropriate configuration has been performed on the IBM System Storage external storage units.

Related information

[Add HyperSwap Storage Description \(ADDHYSSTGD\) command](#)

[Change HyperSwap Storage Description \(CHGHYSSTGD\) command](#)

[Remove HyperSwap Storage Description \(RMVHYSSTGD\) command](#)

[Live Partition Mobility](#)

Configuring DS8000 HyperSwap with independent auxiliary storage pools (IASPs)

IASP-based HyperSwap environments require both HyperSwap and logical unit number (LUN) switching configuration.

These steps assume the user has experience configuring and managing an IBM System Storage DS8000. For information about using an IBM System Storage DS8000 unit, see [Copy Services in the IBM System Storage DS8000 Information Center](#).

If necessary, use HyperSwap independently as a disk protection solution in a single-system environment without any LUN configuration. This has limited uses.

Inside a high availability environment with IASPs, both HyperSwap and LUN switching must be configured. The IASPs in the HyperSwap configuration require LUN volume groups separate from any of the SYSBAS LUNs. Any secondary IASPs in the same HyperSwap relationship associate with the primary IASP only. Primary and secondary IASPs cannot split into separate relationships. Other primary IASPs or user defined file system (UDFS IASP) outside of the HyperSwap environment have separate defined relationships.

IBM i high availability solutions using DS8000 HyperSwap automatically detect the HyperSwap IASPs after the DS8000 is configured appropriately.

To configure HyperSwap for use with IASPs, follow these steps:

1. Create LUNs on the source DS8000 system.
2. Assign those LUNs on the source DS8000 to the System i partition.

Note: The IASP LUNs must be in separate volume groups from any SYSBAS LUNs, on both the source and target DS8000.

3. Create LUNs on the target DS8000 but do not assign them to the IBM i until after the Metro Mirror has been configured.

Note: The IBM i device driver ensures data consistency for HyperSwap switchover and failover operations. Do not use Metro Mirror consistency groups here.

4. Configure Metro Mirror between the source and target DS8000 with the **mkpprcpath** command on the DS8000 from the DSCLI.
5. Start Peer-to-Peer Remote Copy (PPRC) with the **mkpprc** command on the DS8000 from the DSCLI.
6. Assign the LUNs on the target DS8000 to the same IBM i partition.
7. Verify the HyperSwap configuration with the Display HyperSwap Status (**DSPHYSSTS**) command.
The status of the HyperSwap drive pairs configured should display in the DSCLI
8. **If your established HyperSwap environment requires additional disk units for more capacity, complete the following procedure to add units to an existing IASP in a HyperSwap environment:**
 - a) For new disk units to be added to the IASP, pause PPRC with the **pausepprc** command on the DS8000 from the DSCLI
 - b) Stop replication with the Change HyperSwap Status command, **CHGHYSSTS OPTION(*STOP) ASPDEV(iasp-name)**.
 - c) Add the new units to the IASP with the System Service Tools (**STRSST**) command.
 - d) Resume replication with the Change HyperSwap Status command, **CHGHYSSTS OPTION(*START) ASPDEV(iasp-name)**.

The expanded storage in your HyperSwap environment on your IBM i partition is configured for switching.

Configuring LUN level switching in the HyperSwap environment on IBM i

Perform the following steps to set up the LUN level switching in an IASP-based HyperSwap environment on the IBM i partition:

1. Pause PPRC with the **pausepprc** command on the DS8000 from the DSCLI.
2. Create the IASP with the Configure Device ASP (**CFGDEVASP**) command.
To see the units for the IASP, use the Display HyperSwap Status (**DSPHYSSTS**) command.
3. Resume replication with the Change HyperSwap Status command, **CHGHYSSTS OPTION(*START) ASPDEV(iasp-name)**.
4. Create the cluster, the device domain, and device CRG to control the switching.
All nodes in the recovery domain of the device CRG must be in the same device domain.
5. Create the ASP device description on the backup node.
6. Use the Add ASP Copy Description (**ADDASPCPYD**) command to configure both LUN level switching and HyperSwap.
The nodes, host connections, and volume groups, which can access the IASP are entered in the recovery domain parameter. The storage host information, LUN ranges, and recovery domain information for the second DS8000 in the HyperSwap relationship are entered by using the **STGHOST2**, **LUN2**, and **RCYDMN2** parameters.
7. Create an ASP copy description for every ASP device that is to use LUN level switching with HyperSwap.
All ASP devices in the cluster resource group must be defined by a separate ASP copy description.
8. Verify the HyperSwap configuration using the Display HyperSwap Status (**DSPHYSSTS**) command and the correct ASP copy description configuration using the Display ASP Copy Description (**DSPASPCPYD**) command.
To see the new units for the IASP, use the Display HyperSwap Status (**DSPHYSSTS**) command.

Managing PowerHA

After you have configured a PowerHA solution, you can manage that solution by using several interfaces that are related to high availability.

Scenarios: Managing high availability solutions

As a system operator or administrator of your high-availability solution, you need to perform common tasks like backup and system maintenance in your high-availability environment.

The following scenarios provide instructions on performing common system tasks, such as backups and upgrades, as well as examples of managing high-availability events, such as cluster partitions and failover. For each scenario, a model environment has been chosen. The instructions for each scenario correspond to that particular high-availability solution and are meant for example purposes only.

Scenarios: Performing backups in a high-availability environment

Depending on your high-availability solution and your backup strategy, the method for backing up data can be different. However, there is a common set of tasks when you perform backup operations for systems in a high-availability environment.

In several high-availability solutions, you have the capability of performing remote backups from the second copy of data that is stored on the backup system. Remote backups allow you to keep your production system operational, while the second system is backed up. Each of these scenarios provides examples of two high-availability solutions where backups are performed remotely on the backup system.

In the first scenario, remote backups are performed in a high-availability solution that uses geographic mirroring technology. The second scenario shows how the FlashCopy feature can be used in a high-availability environment that uses IBM System Storage solutions, such as metro or global mirror.

Scenario: Performing backups in geographic mirroring environment

This scenario provides an overview of tasks that are necessary when performing a remote backup in a PowerHA solution that uses geographic mirroring.


Overview

In this example, a system administrator needs to perform a backup of data stored on independent disk pools that are used in a high-availability solution based on geographic mirroring technology. The administrator does not want to affect the production system by taking it offline to perform the backup. Instead, the administrator plans to temporarily detach the mirrored copy, and then perform a backup from the second copy of data located on independent disk pools at a remote location.

Note: Detaching the mirrored copy essentially ends geographic mirroring until the copy is reattached to the production. During the time it is detached, high availability and disaster recovery are not operational. If an outage to the production system occurs during this process, some data will be lost.

Details

The following graphic illustrates this environment:



Configuration steps

Complete the following steps when performing a backup of a mirror copy using the PowerHA graphical interface:

1. [“Detaching mirror copy” on page 121](#)
2. [Making disk pool available](#)
3. [Backing up independent disk pool](#)
4. [Make disk pool unavailable](#)
5. [“Reattaching mirror copy” on page 122](#)

Complete the following steps when performing a backup of a mirror copy using other interfaces:

1. [“Quiescing an independent disk pool” on page 119](#)
2. [“Detaching mirror copy” on page 121](#)
3. [“Resuming an independent disk pool” on page 119](#)
4. [Making disk pool available](#)
5. [Backing up independent disk pool](#)
6. [Make disk pool unavailable](#)
7. [“Reattaching mirror copy” on page 122](#)

Scenario: Performing a FlashCopy function

In this example, an administrator wants to perform a backup from the remote copy of data that is stored in an external storage unit at the backup site. Using the FlashCopy function available with IBM Storage Solutions, the administrator reduces his backup time considerably.

Overview


In this example, a system administrator needs to perform a backup of data that is stored on IBM System Storage external storage units. The administrator does not want to affect the production system by taking it offline to perform the backup. Instead, the administrator plans to perform a FlashCopy operation, which takes a point-in-time capture of the data. From this data, the administrator backs up the data to external media. The FlashCopy operation only takes a few seconds to complete, thus reducing the time for the entire backup process.

Although in this example the FlashCopy feature is being used to for backup operations, it should be noted that the FlashCopy feature has multiple uses. For example, FlashCopy can be used for data warehousing to reduce query workload on production systems, or for duplicating production data to create a test environment.

A similar scenario can also be accomplished with other system storage technologies. For more information on storage technologies that support FlashCopy, see [“PowerHA supported storage servers” on page 37](#).

Configuration steps

Complete the following steps when performing a FlashCopy using the PowerHA graphical interface:

1. Configure the IBMSystem Storage external storage units forFlashCopy. For information about using the FlashCopy function on IBM System StorageDS8000, see [Copy Services in the IBM System Storage DS8000 Information Center](#) .
2. [“Configuring a FlashCopy session” on page 72](#)
3. [Make disk pool available](#)
4. [Backing up independent disk pool](#)

Complete the following steps when performing a FlashCopy using other interfaces:

1. Configure the IBM System Storage external storage units for FlashCopy. For information about using the FlashCopy function on IBM System Storage DS8000, see [Copy Services in the IBM System Storage DS8000 Information Center](#).
2. [“Quiescing an independent disk pool” on page 119](#)
3. [“Configuring a FlashCopy session” on page 72](#)
4. [“Resuming an independent disk pool” on page 119](#)
5. [Make disk pool available](#)
6. [Backing up independent disk pool](#)

Related concepts

PowerHA supported storage servers
 IBM System Storage provides enhanced storage capabilities.

Scenario: Upgrading operating system in a high-availability environment


In this example, a system administrator is upgrading the operating system for two IBM i systems in a PowerHA solution that is based on geographic mirroring.

Overview

The system administrator needs to upgrade the operating system for two systems in the high-availability environment. In this example, there are two nodes: System 1 and System 2. System 1 is the production copy and System 2 is the mirror copy. Both systems are at IBM i V7R1. The independent disk pool is online, geographic mirroring is active, and the systems are synchronized. The system administrator wants to upgrade both of these systems to IBM i 7.2.

Details

The following graphic illustrates the environment:



Configuration steps

1. [Detach the mirror copy \(System 2\)](#).
2. [End the CRG \(System 2\)](#).
3. [Stop the node \(System 2\)](#).
4. Upgrading System 2 to the new release. See [Upgrading or replacing IBM i and related software](#) for details.
5. [Install IBM PowerHA SystemMirror for i licensed program](#).
6. [Make disk pool available](#) and test applications on System 2. Testing the applications ensures that they operate as expected within the new release. After application tests are completed, you can finish the upgrade by completing the rest of these steps.
7. [Make disk pool unavailable](#) on the detached mirrored copy (System 2).
8. [Reattach mirrored copy](#). This initiates a resynchronization of the mirrored data. After the resynchronization is completed, you can continue the upgrade process.
9. [Performing switchovers](#). This makes the mirrored copy (System 2) the new production copy and the production copy (System 1) becomes the new mirrored copy.

Note: Geographic mirroring is suspended because you cannot perform geographic mirroring from n-1 to n. You can perform geographic mirroring from n to n-1 without problems. In this scenario, geographic mirroring is suspended after a switchover is completed. Data is now no longer mirrored during the remainder of the upgrade process because there is no longer a valid backup system.

10. [End the CRG \(System 1\)](#).
11. [Stop the node \(System 1\)](#).
12. Upgrade System 1 to the new release. See [Upgrading or replacing IBM i and related software](#) for details.
13. [Install IBM PowerHA SystemMirror for i licensed program](#).
14. [Start nodes \(System 1\)](#).
15. [Start CRGs \(System 1\)](#).
16. [Resume mirroring](#)
17. [Perform switchover](#). This switches the current mirrored copy (System 1) back to the production copy and the production copy (System 2) to become the mirrored copy. This is the original configuration before the upgrade.
18. [Adjusting the cluster version of a cluster](#)
19. [Adjusting the high availability version of PowerHA LP](#)

Example: Upgrading operating system

In a high-availability environments, you must perform specific actions prior to performing operating system upgrades.

The following examples can help you determine what you need to do to perform an upgrade in your cluster environment. Before performing the upgrade or any actions you should first determine the current cluster version for your cluster.

Notes:

1. 7.1 represents the current release of the operating system.
2. 7.2 represents the new release of the operating system.
3. V6R1 represents the prior release of the operating system.

Example 1: The node to be upgraded is at IBM i 7.1. All other nodes in the cluster are at IBM i 7.1 or higher. The current cluster version is 7.

Action:

1. [Upgrade](#) node to IBM i 7.2.
2. [Start](#) the upgraded node.

Example 2: The node to be upgraded is at IBM i 7.1. All other nodes in the cluster are at IBM i 7.1 or higher. The current cluster version is 6.

Action:

1. [Upgrade](#) the node to IBM i 7.2.
2. [Start](#) the upgraded node.

Example 3: The node to be upgraded is IBM i V5R4. All other nodes in the cluster are at IBM i V5R4 or higher. The current cluster version is 5.

Action:

1. [Upgrade](#) the node to IBM i 7.1.
2. [Start](#) the upgraded node.

Example 4: The node to be upgraded is at IBM i V6R1. All other nodes in the cluster are at IBM i V5R4 or higher. The current cluster version is 5.

Actions:

1. [Upgrade](#) all nodes to V6R1.

2. Start all of the upgraded nodes.
3. Change the cluster version to 6.
4. Upgrade the node to 7.2.
5. Start the upgraded node.

Example 5: The node to be upgraded is at IBM i V5R4 or lower. All other nodes in the cluster are at IBM i V5R4 or lower. The current cluster version is less than or equal to 5.

Actions:

1. Upgrade all nodes to V6R1.
2. Start all of the upgraded nodes.
3. Change the cluster version to 6.
4. Upgrade the node to 7.2.
5. Start the upgraded node.

The following table provides actions you need to take when performing an upgrade in a cluster environment.

Current® release of node you are upgrading	Current cluster version	Actions
V7R1	6 or 7	<ol style="list-style-type: none"> 1. <u>Upgrade</u> the node to IBM i 7.2. 2. <u>Start</u> the upgraded node.
V6R1	5 or 6	<ol style="list-style-type: none"> 1. <u>Upgrade</u> the node to IBM i 7.2. 2. <u>Start</u> the upgraded node.
V5R4 or lower	less than or equal to 5	<ol style="list-style-type: none"> 1. <u>Upgrade</u> all nodes to V6R1. 2. <u>Start</u> all of the upgraded node. 3. <u>Change</u> the cluster version to 6. 4. <u>Upgrade</u> the node to IBM i 7.2. 5. <u>Start</u> the upgraded node.

Managing clusters

Using the PowerHA graphical interface, you can perform many tasks that are associated with the cluster technology that is the basis of your IBM i high availability solution. These tasks help you manage and maintain your cluster.

Note: The PowerHA graphical interface allows you to perform these tasks from any node in the cluster.

Some of the changes that you can make to the cluster after you configure it include the following:

Cluster tasks

- [“Creating a cluster” on page 49](#)
- [“Deleting a cluster” on page 82](#)
- [“Specifying message queues” on page 53](#)
- [Add a node to a cluster](#)

- [Remove nodes from a cluster](#)
- [Start a cluster node](#)
- [End a cluster node](#)
- [“Displaying node properties” on page 85](#)
- [“Adding a node to a device domain” on page 56](#)
- [“Removing a node from a device domain” on page 87](#)
- [“Add a cluster monitor to a node” on page 87](#)
- [“Removing a cluster monitor ” on page 88](#)
- [“Monitoring cluster status” on page 84](#)
- [Adjust the cluster version of a cluster to the latest level](#)
- [“Adjusting the PowerHA version” on page 81](#)

Cluster resource group tasks

- [“Creating cluster resource groups \(CRGs\)” on page 61](#)
- [“Deleting a CRG” on page 90](#)
- [“Starting a CRG” on page 53](#)
- [“Stopping a CRG” on page 89](#)
- Add a node to a cluster resource group
- Remove a node from a cluster resource group
- [“Changing the recovery domain for a CRG” on page 90](#)
- [“Performing switchovers” on page 54](#)
- [“Displaying CRG status” on page 88](#)
- Add a node to a device domain
- Remove a node from a device domain

Cluster administrative domain tasks

- [“Creating a cluster administrative domain ” on page 63](#)
- [“Deleting a cluster administrative domain” on page 96](#)
- [“Changing the properties of a cluster administrative domain” on page 97](#)
- [“Starting a cluster administrative domain” on page 64](#)
- [“Stopping a cluster administrative domain” on page 96](#)
- [“Adding a node to the cluster administrative domain” on page 63](#)
- [“Adding monitored resource entries” on page 65](#)
- [“Removing monitored resource entries” on page 98](#)
- [“Listing monitored resource entries” on page 98](#)
- [“Displaying monitored resource entry messages” on page 117](#)

Miscellaneous tasks

- [“Displaying cluster configuration” on page 83](#)
- [“Saving and restoring cluster configuration” on page 83](#)
- [“Cluster deconfiguration checklist” on page 85](#)
- [“Managing failover outage events” on page 91](#)

Adjusting the PowerHA version

The PowerHA version is the version at which the nodes in the cluster that is managed by the PowerHA product are actively communicating with each other.

The PowerHA version values determine which functions can be used by the PowerHA product. The PowerHA version may require a certain cluster version in order to operate. For example, PowerHA version 2.0 requires a current cluster version of 7.

The current PowerHA version is set when a cluster is created. If a cluster exists, the current PowerHA version is set to the lowest supported version.

Like cluster version, PowerHA has a current and potential version level. The current PowerHA version is the version at which the nodes in the cluster, which are known by the PowerHA product are actively communicating with each other. The potential PowerHA version is the highest PowerHA version the node can support. The PowerHA version cannot be changed until all PowerHA nodes are installed with a common potential PowerHA version. The potential PowerHA version can be between n and $n+3$. For example, the current PowerHA version is 2.0, NODE1 has a potential PowerHA version of 3.0, NODE2 has a potential PowerHA version of 4.0, and NODE3 has a potential PowerHA version of 5.0. All three nodes can support version 3.0, so the current PowerHA version can be adjusted to 3.0.

Beginning with PowerHA version 2.0, if a node with an incompatible potential PowerHA version is added to the cluster, the node will successfully be added, but the node will be considered "unknown" to PowerHA. If a node is unknown to PowerHA, certain product functions cannot be performed on the node. A node is known to PowerHA if the node has the PowerHA product installed, and the potential PowerHA version is compatible with the current PowerHA version.

For users of Cluster Administrative Domain, when the current PowerHA version is adjusted to 3.0 or higher, monitored resource entries in the cluster administrative domain may be affected. If all attributes supported at the previous cluster and PowerHA versions are being monitored for a monitored resource entry, then PowerHA will automatically update that monitored resource entry to also monitor any new attributes supported at the current cluster version and current PowerHA version.

The current PowerHA version can be changed with the **Change Cluster Version (CHGCLUVER)** command.

The **Change Cluster Version (CHGCLUVER)** command can only be used to adjust to a higher cluster or PowerHA version. If you want to adjust the PowerHA version by two, the **CHGCLUVER** command must be run twice.

The current cluster version cannot be set higher than the lowest potential node version in the cluster. Likewise, the current PowerHA version cannot be set higher than the lowest potential PowerHA version of any node in the cluster. PowerHA potential version of any node in the cluster. To view the potential node and PowerHA versions, use the **Display Cluster Information (DSPCLUINF)** command.

To adjust the PowerHA version using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Select **PowerHA** from the IBM Navigator for i window.
4. On the PowerHA page, select **Properties ...** from the Select Action menu next to the cluster name.
5. Click **Edit** in the General section.
6. Select a new PowerHA version and click **Save**.

Related concepts

[Cluster version](#)

Related information

[Change Cluster Version \(CHGCLUVER\) command](#)

Adjusting the cluster version of a cluster

The cluster version defines the level at which all the nodes in the cluster are actively communicating with each other.

Cluster versioning is a technique that allows the cluster to contain systems at multiple release levels and fully interoperate by determining the communications protocol level to be used.

To change the cluster version, all nodes in the cluster must be at the same potential version. The cluster version can then be changed to match the potential version. This allows the new function to be used. The version can only be increased by one. It cannot be decremented without deleting the cluster and re-creating it at a lower version. The current cluster version is initially set by the first node that is defined in the cluster. Subsequent nodes added to the cluster must be equal to the current cluster version or the next level version; otherwise, they cannot be added to the cluster.

If you are upgrading a node to a new release, you must ensure that the node has the appropriate cluster version. Cluster only supports a one version difference. If all the nodes in the cluster are at the same release, you should upgrade to the new release, before changing the cluster version. This ensures that all functions associated with the new release are available. See the topic, [“Scenario: Upgrading operating system in a high-availability environment”](#) on page 77 for detailed actions for upgrading to a new release.

To adjust the cluster version using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, select **Properties ...** from the **Select Action** menu next to the cluster name.
5. Click **Edit** in the General section.
6. Select a new cluster version and click **Save**.

Related concepts

[Cluster version](#)

Related information

[Change Cluster Version \(CHGCLUVER\) command](#)

[Adjust Cluster Version \(QcstAdjustClusterVersion\) API](#)

Deleting a cluster

When you delete a cluster, cluster resource services ends on all active cluster nodes and they are removed from the cluster.

You must have at least one active node before you can delete a cluster. If you have switched disks or other switchable devices in your cluster, you must first remove each node from the device domain before you delete your cluster. Otherwise, you might not be able to add the disks back into another cluster.

To delete a cluster using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, select **Delete Cluster** from the Select Action menu next to the cluster name.
5. Click **Yes** on the confirmation panel.

Related tasks

[Removing a node from a device domain](#)

A *device domain* is a subset of nodes in a cluster that share device resources.

Related information

[Delete Cluster \(DLTCLU\) command](#)

[Delete Cluster \(QcstDeleteCluster\) API](#)

Displaying cluster configuration

You can display a detailed report that provides information on the cluster configuration. The cluster configuration report provides detailed information about the cluster, node membership list, configuration and tuning parameters, and each cluster resource group in the cluster.

To display cluster configuration, complete the following steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **Cluster Resource Services** page, select the **Display Configuration Information** task. This displays the Cluster Configuration and Properties page. You can save this page as a file or print it.

To display cluster configuration from the command-line, use either of these commands:

- [Display Cluster Information \(DSPCLUINF\) command](#)
- [Work with Cluster \(WRKCLU\) command](#)

Saving and restoring cluster configuration

If you use clustering on your systems, it is still important that you create a backup and recovery strategy to protect your data.

If you are planning on using clustering as your backup strategy so that you have one system up and running while the other system is down when its being backed up, it is recommended that you have a minimum of three systems in the cluster. By having three systems in the cluster, you will always have one system to switch over to should a failure occur.

Saving and restoring cluster resource groups

You can save a cluster resource group whether the cluster is active or inactive. The following restrictions apply for restoring a cluster resource group:

- If the cluster is up and the cluster resource group is not known to that cluster, you cannot restore the cluster resource group.
- If the node is not configured for a cluster, you cannot restore a cluster resource group.

You can restore a cluster resource group if the cluster is active, the cluster resource group is not known to that cluster, the node is in the recovery domain of that cluster resource group, and the cluster name matches that in the cluster resource group. You can restore a cluster resource group if the cluster is configured but is not active on that node and if that node is in the recovery domain of that cluster resource group.

Preparing for a disaster

In the event of a disaster, you might need to reconfigure your cluster. In order to prepare for such a scenario, it is recommended that you save your cluster configuration information and keep a hardcopy printout of that information.

1. Use the **Save Configuration (SAVCFG)** command or the **Save System (SAVSYS)** command after making cluster configuration changes so that the restored internal cluster information is current and consistent with other nodes in the cluster. See Saving configuration information for details on performing a SAVCFG or SAVSYS operation.
2. Print a copy of the cluster configuration information every time you change it. You can use the **Display Cluster Information (DSPCLUINF)** command to print the cluster configuration. Keep a copy with your backup tapes. In the event of a disaster, you might need to reconfigure your entire cluster.

Related information

[Saving configuration information](#)

[Save Configuration \(SAVCFG\) command](#)

[Save System \(SAVSYS\) command](#)

[Display Cluster Information \(DSPCLUINF\) command](#)

Monitoring cluster status

The PowerHA graphical interface presents the overall status of the cluster via status icons.

To obtain consistent information, you can either access the cluster information from an active node in the cluster, or start this node and retry the request.

To monitor cluster status using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, status icons show the overall status of the cluster.

To see the overall status of the cluster from the command-line, use either of these commands:

- [Display Cluster Information \(DSPCLUINF\) command](#)
- [Display Cluster Resource Group Information \(DSPCRGINF\) command](#)
- [Work with Cluster \(WRKCLU\) command](#)

Related tasks

Starting nodes

Starting a cluster node activates clustering and cluster resource services on a node in an IBM i high availability environment.

Related information

[List Cluster Information \(QcstListClusterInfo\) API](#)

[List Device Domain Info \(QcstListDeviceDomainInfo\) API](#)

[Retrieve Cluster Resource Services Information \(QcstRetrieveCRSInfo\) API](#)

[Retrieve Cluster Information \(QcstRetrieveClusterInfo\) API](#)

[List Cluster Resource Groups \(QcstListClusterResourceGroups\) API](#)

[List Cluster Resource Group Information \(QcstListClusterResourceGroupInf\) API](#)

Specifying message queues

You can either specify a cluster message queue or a failover message queue. These message queues help you determine causes of failures in your PowerHA environment.

A cluster message queue is used for cluster-level messages and provides one message, which controls all cluster resource groups (CRGs) failing over to a specific node. A failover message queue is used for CRG-level messages and provides one message for each CRG that is failing over.

Specifying a cluster message queue

Note: You can also configure a cluster to use a cluster message queue by specifying the message queue while running the Create Cluster wizard.

To specify a cluster message queue using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, select **Properties ...** from the **Select Action** menu.
5. Click **Edit** in the Advanced section of the **Properties** page.
6. Specify the cluster message queue information in the **Cluster Message Queue** field and click **Save**.

Specifying a failover message queue

To specify a failover message queue using the PowerHA graphical interface, complete these steps:

1. In a web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** in the IBM Navigator for i window.

4. On the **PowerHA** page, click **Cluster Resource Groups**.
5. On the **Cluster Resource Group** page, select **Properties ...** from the context menu of the CRG for which you want a failover message queue.
6. Click **Edit** in the Advanced section of the **Properties** page.
7. Specify the failover message queue information in the **Failover Message Queue** field and click **Save**.

Cluster deconfiguration checklist

To ensure complete deconfiguration of a cluster, you must systematically remove different cluster components.

<i>Table 8. Independent disk pool deconfiguration checklist for clusters</i>	
Independent disk pool requirements	
---	If you are using switched disk pools, the tower should be switched to the node which is the SPCN owner before deconfiguring the cluster resource group. You can use Initiate Switchover (QcstInitiateSwitchOver) API or the Change Cluster Resource Group Primary (CHGCRGPRI) command to move the CRG back to the SPCN owner. If this step is not performed, you will not be able to mark the tower private for that system.
---	If you plan to remove a subset of an independent disk pool group or remove the last independent disk pool in the switchable devices, you must end the CRG first. Use the End Cluster Resource Group (ENDCRG) command.
---	If you want delete an independent disk pool that is participating in a cluster, it is strongly recommended that you first delete the device cluster resource group (CRG). See “Deleting a CRG” on page 90 for details. You can also use the Remove CRG Device Entry (RMVCRGDEVE) command to remove the configuration object of the independent disk pool from the CRG.
---	After you have removed the configuration object of the independent disk pool from the cluster switchable device, you can delete an independent disk pool .
---	Delete the device description for an independent disk pool by completing these tasks: <ol style="list-style-type: none"> 1. On a command-line interface, type WRKDEVD DEVD(*ASP) and press Enter. 2. Page down until you see the device description for the independent disk pool that you want to delete. 3. Select Option 4 (Delete) by the name of the device description and press Enter.

<i>Table 9. Cluster resource group deconfiguration checklist for clusters</i>	
Cluster resource group requirement	
---	Delete cluster resource group by completing the either of the following steps: <ol style="list-style-type: none"> 1. If clustering is not active on the node, then type DLTCRG CRG(CRGNAME) on a command-line interface. CRGNAME is the name of the CRG that you want to delete. Press Enter. 2. If clustering is active on the node, then type DLTCRGCLU CLUSTER(CLUSTERNAME) CRG(CRGNAME) on a command-line interface. CLUSTERNAME is the name of the cluster. CRGNAME is the name of the CRG that you want to delete. Press Enter.

Managing nodes

System and logical partitions that are a part of an IBM i high availability environment are called nodes. You can perform several managing tasks that pertain to nodes.

Displaying node properties

You can display and manage properties that are associated with nodes that are configured as part of your high-availability environment by using the PowerHA graphical interface.

To display node properties using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Nodes**.
5. On the **Cluster Nodes** page, select **Properties ...** from the context menu of a node view with its properties.
 - The General page displays the name of the node and the system IP address for that node.
 - The Clustering page displays the following information:
 - The cluster interface IP addresses that are used by clustering to communicate with other nodes in the cluster.
 - The potential version of the node specifies the version and modification level at which the nodes in the cluster are actively communicating with each other.
 - The device domains that are configured in the selected cluster. If you select a device domain in the list, the nodes that belong to the selected device domain are also displayed.

Related information

[Work with cluster \(WRKCLU\) comand](#)

Stopping nodes

Stopping or ending a node ends clustering and cluster resource services on that node.

To stop clustering on a node using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click **Cluster Nodes**.
5. On the **Cluster Nodes** page, select **Stop** from the context menu of the nodes you want to stop.
6. Click **Yes** on the confirmation panel.

Related information

[End Cluster Node \(ENDCLUNOD\) command](#)

[End Cluster Node \(QcstEndClusterNode\) API](#)

Removing nodes

You might need to remove a node from a cluster if you are performing an upgrade of that node or if the node no longer needs to participate in the IBM i high-availability environment.

To remove a node from an existing cluster using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Nodes**.
5. On the **Cluster Nodes** page, select **Remove** from the context menu of the node you want to remove.
6. Click **Yes** on the confirmation window.

Related tasks

[Deconfiguring geographic mirroring](#)

If you no longer want the capability to use geographic mirroring for a specific disk pool or disk pool group, you can select to **Deconfigure Geographic Mirroring**. If you deconfigure geographic mirroring, the system stops geographic mirroring and deletes the mirror copy of the disk pools on the nodes in the mirror copy site.

Related information

[Remove Cluster Node Entry \(RMVCLUNODE\) command](#)

Remove Cluster Node Entry (QcstRemoveClusterNodeEntry) API

Removing a node from a device domain

A *device domain* is a subset of nodes in a cluster that share device resources.

Important:

Be cautious when removing a node from a device domain. If you remove a node from a device domain, and that node is the current primary point of access for any independent disk pools, those independent disk pools remain with the node being removed. This means that those independent disk pools are no longer accessible from the remaining nodes in the device domain.

After a node is removed from a device domain, it cannot be added back to the same device domain if one or more of the existing cluster nodes still belong to that same device domain. To add the node back to the device domain, you must:

1. Delete the independent disk pools currently owned by the node being added to the device domain.
2. Restart the system by performing an IPL on the node.
3. Add the node to the device domain.
4. Re-create the independent disk pools that were deleted in step 1.

To remove a node from a device domain using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Nodes**.
5. On the **Cluster Nodes** page, select **Properties ...** from the context menu of the node you want to add to a device domain.
6. Click **Edit** in the General section of the **Properties** page.
7. Clear the **Device Domain** field and click **Save**.

Related tasks

Deleting a cluster

When you delete a cluster, cluster resource services ends on all active cluster nodes and they are removed from the cluster.

Related information

Remove Device Domain Entry (RMVDEVDMNE) command

Remove Device Domain Entry (QcstRemoveDeviceDomainEntry) API

Add a cluster monitor to a node

The PowerHA® graphical interface can now use Hardware Management Console (HMC) or a Virtual I/O Server (VIOS) partition on an Integrated Virtualization Manager (IVM) managed server to detect when a cluster node fails. This new capability allows more failure scenarios to be positively identified and avoids cluster partition situations.

The PowerHA® graphical interface allows you to use the HMC or VIOS on an IVM managed server to monitor and manage the state of each system. Once a monitor is set up, HMC or IVM provide notification of node failures. A cluster monitor can be used to reduce the number of failure scenarios which result in cluster partitions.

PowerHA GUI only supports the Common Informational Model (CIM) server for the cluster monitor. The Add Cluster Monitor (**ADDCLUMON**) command must be used if you want to use the representational state transfer (REST) server.

To add a cluster monitor to an existing cluster using the PowerHA® graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.

3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Nodes** .
5. On the **Cluster Nodes** page, select **Properties ...** from the context menu of the node to which you want to add a cluster monitor.
6. In the Cluster Monitors section of the **Properties** page, select **Add Cluster Monitor ...** from the **Select Action** menu.
7. Specify the cluster monitor information and click **OK**.

Related information

[Add Cluster Monitor \(ADDCLUMON\) command](#)

Removing a cluster monitor

A *cluster monitor* provides another source of information to allow cluster resource services to determine when a cluster node has failed.

Important:

Be cautious when removing a cluster monitor. If you remove a node from a cluster monitor, and that node is the current primary point of access for any CRG, that node could partition when in fact, the node really failed. This means the user must now do manual steps to become highly available again.

PowerHA GUI only supports the Common Informational Model (CIM) server for the cluster monitor. The Remove Cluster Monitor command must be used if you want to use the representational state transfer (REST) server.

To remove a cluster monitor using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Nodes** .
5. On the **Cluster Nodes** page, select **Properties ...** from the context menu of the node to which you want to remove a cluster monitor.
6. In the **Cluster Monitors** section of the **Properties** page, select **Remove ...** from the context menu of the cluster monitors you want to remove.
7. Click **Yes** on the confirmation panel.

Related information

[Remove Cluster Monitor \(RMVCLUMON\) command](#)

Managing cluster resource groups (CRGs)

Cluster resource groups (CRGs) manage resilient resources within an IBM i high availability environment. They are a cluster technology that defines and controls switching resources to backup systems in the event of an outage.

Displaying CRG status

You can monitor the status of cluster resource groups (CRG) in your high-availability environment. You can use these status messages to validate changes in the CRG or to determine problems with the CRG.

To display the CRG status using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Resource Groups**.

The following are possible status values for a CRG:

Table 10. Status values for CRGs

Possible values	Description
Started	The CRG is currently started.
Stopped	The CRG is currently stopped.
Indoubt	The information about this CRG within the high availability solution might not be accurate. This status occurs when the CRG exit program is called with an action of undo and fails to complete successfully.
Restored	The CRG was restored on its node and has not been copied to other nodes in the cluster. When clustering is started on the node, the CRG will be synchronized with the other nodes and its status is set to inactive.
Inactive	Cluster resource services for the CRG is not active on the node. The node might have failed, the node might have been ended, or the CRG job on that node might not be running.
Deleting	The CRG is in the process of being deleted from the cluster.
Changing	The CRG is in the process of being changed. The CRG is reset to its previous status when the change has been successfully completed.
Stopping	The CRG is in the process of being stopped.
Adding	The CRG is in the process of being added to the cluster.
Starting	The CRG is in the process of being started.
Switching	The CRG is in the process of switching over to another node.
Adding node	A new node is in the process of being added to the cluster. The CRG is reset to its previous status when the node has been successfully added.
Removing node	A node is in the process of being removed from the CRG. The CRG is reset to its previous status when the node has been successfully removed.
Changing node status	The status of a node in the recovery domain for a CRG is currently being changed.

Related information

[Display CRG Information \(DSPCRGINF\) command](#)

[Work with Cluster \(WRKCLU\) command, option 9 \(Work with cluster resource groups\)](#)

Stopping a CRG

Cluster resource groups (CRGs) manage resilient resources within an IBM i high availability environment. They are a cluster technology that defines and controls switching resilient resources to backup systems in the event of an outage.

You might want to stop the CRG to end automatic failover capability in your high-availability environment. For example, you might be performing an IPL on one of the systems that is defined in the CRG.

To stop a CRG using PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Resource Groups**.
5. On the **Cluster Resource Group** page, select **Stop** from the context menu of the CRG that you want to stop.
6. Click **Yes** on the confirmation panel.

Related information

[End Cluster Resource Group \(ENDCRG\) command](#)

[End Cluster Resource Group \(QcstEndClusterResourceGroup\) API](#)

Deleting a CRG

You can delete a cluster resource group by using the PowerHA graphical interface or PowerHA commands.

You can delete a cluster resource group by using the PowerHA graphical interface. In order to delete a CRG, it must be inactive.

To delete a CRG using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Resource Groups**.
5. On the **Cluster Resource Group** page, select **Delete** from the context menu of the CRG that you want to delete.
6. Click **Yes** on the confirmation panel.

Related information

[Delete CRG Cluster \(DLTCRGCLU\) command](#)

[Delete Cluster Resource Group \(DLTCRG\) command](#)

[Delete Cluster Resource Group \(QcstDeleteClusterResourceGroup\) API](#)

Changing the recovery domain for a CRG

The recovery domain controls recovery actions for a subset of nodes that are defined in a cluster resource group (CRG).

To use the PowerHA graphical interface to make an independent disk pool available, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Resource Groups**.
5. On the **Cluster Resource Group** page, select **Recovery Domain ...** from the context menu of the CRG of which you want to change the recovery domain.
6. Once on the **Recovery Domain** page, you can add and remove nodes, change node roles, change site names, and change data port IP addresses.

Related information

[Add Cluster Resource Group Node Entry \(ADDCRGNODE\) command](#)

[Change Cluster Resource Group \(CHGCRG\) command](#)

[Remove Cluster Resource Group Node Entry \(RMVCRGNODE\) command](#)

[Add a Node to Recovery Domain \(QcstAddNodeToRcvyDomain\) API](#)

[Change Cluster Resource Group \(QcstChangeClusterResourceGroup\) API](#)

Remove Node from Recovery Domain (QcstRemoveNodeFromRcvyDomain) API

Creating site names and data port IP addresses

If you are using geographic mirroring, the nodes defined in the recovery domain node of the device cluster resource group must have a data port IP address and site name.

The site name is associated with a node in the recovery domain for a device cluster resource group, applicable only to geographic mirroring. When you are configuring a geographic mirroring environment for high availability, each node at different sites must be assigned to a different site name.

To create the data port IP address and site names for nodes in the recovery domain, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the Cluster Resource Services page, click the **Work with Cluster Resource Groups** task to show a list of cluster resource groups in the cluster.
5. On the Cluster Resource Group tab, click the context icon next to the device cluster resource group, and then select **Properties**.
6. On the Recovery Domain page, select **Edit**.
7. To use an existing data port IP address, select it from the list and click **OK**. To add a new data port IP address, click **Add**. In the Add Data Port IP Address window, enter the IP address.
8. In the Edit window, you can specify the Site name.

Managing failover outage events

Typically, a failover results from a node outage, but there are other reasons that can also generate a failover. Different system or user actions can potentially cause failover situations.

It is possible for a problem to affect only a single cluster resource group (CRG) that can cause a failover for that CRG but not for any other CRG.

Four categories of outages can occur within a cluster. Some of these events are true failover situations where the node is experiencing an outage, while others require investigation to determine the cause and the appropriate response. The following tables describe each of these categories of outages, the types of outage events that fall into that category and the appropriate recovery action you should take to recover.

Category 1 outages: Node outage causing failover

Node-level failover occurs, causing the following to happen:

- For each CRG, the primary node is marked *inactive* and made the last backup node.
- The node that was the first backup becomes the new primary node.

Failovers happen in this order:

1. All device CRGs
2. All data CRGs
3. All application CRGs

Notes:

1. If a failover for any CRG detects that none of the backup nodes are active, the status of the CRG is set to *indoubt* and the CRG recovery domain does not change.
2. If all of cluster resource services fails, then the resources (CRGs) that are managed by cluster resource services go through the failover process.

Failover outage event
ENDTCP(*IMMED or *CNTRLD with a time limit) is issued.

Table 11. Category 1 outages: Node outage causing failover (continued)

Failover outage event
ENDSYS (*IMMED or *CNTRLD) is issued.
PWRDWN SYS(*IMMED or *CNTRLD) is issued.
Initial program load (IPL) button is pressed while cluster resource services is active on the system.
End Cluster Node (API or command) is called on the primary node in the CRG recovery domain.
Remove Cluster Node (API or command) is called on the primary node in the CRG recovery domain.
HMC delayed power down of the partition or panel option 7 is issued.
ENDSBS QSYSWRK(*IMMED or *CNTRLD) is issued.

Category 2 outages: Node outage causing partition or failover

These outages will cause either a partition or a failover depending on whether advanced node failure detection is configured. Refer to the columns in the table. If advanced node failure detection is configured, failover occurs in most cases and Category 1 outage information applies. If advanced node failure detection is not configured, partition occurs and the following applies:

- The status of the nodes not communicating by cluster messaging is set to a Partition status. See [Cluster partition](#) for information about partitions.
- All nodes in the cluster partition that do not have the primary node as a member of the partition will end the active cluster resource group.

Notes:

1. If a node really failed but is detected only as a partition problem and the failed node was the primary node, you lose all the data and application services on that node and no automatic failover is started.
2. You must either declare the node as failed or bring the node back up and start clustering on that node again. See [Change partitioned nodes to failed](#) for more information.

Table 12. Category 2 outages: Node outage causing partition

Failover outage event	No advanced node failure detection	HMC	VIOS on IVM
CEC hardware outage (CPU, for example) occurs.	partition	failover	partition or failover
Operating system software machine check occurs.	partition	failover	failover
HMC immediate power off or panel option 8 is issued.	partition	failover	failover
HMC partition restart or panel option 3 is issued.	partition	failover	failover
Power loss to the CEC occurs.	partition	partition	partition

Category 3 outages: CRG fault causing failover

For a system using a Virtual I/O Server (VIOS) partition on an Integrated Virtualization Manager (IVM) managed server, a CEC hardware failure could result in either failover or partition. Which occurs depends

upon the type of system and the hardware failure. For example, in a blade system, a CEC failure that prevents VIOS on an IVM managed server from running results in a partition since VIOS is unable to report any failure. In the same system in which a single blade fails but VIOS on an IVM managed server continues to run, failover results since VIOS is able to report the failure.

When a CRG fault causes a failover, the following happens:

- If only a single CRG is affected, failover occurs on an individual CRG basis. This is because CRGs are independent of each other.
- If someone cancels several cluster resource jobs, so that several CRGs are affected at the same time, no coordinated failover between CRGs is performed.
- The primary node is marked as Inactive in each CRG and made the last backup node.
- The node that was the first backup node becomes the new primary node.
- If there is no active backup node, the status of the CRG is set to Indoubt and the recovery domain remains unchanged.

<i>Table 13. Category 3 outages: CRG fault causing failover</i>
Failover outage event
The CRG job has a software error that causes it to end abnormally.
Application exit program failure for an application CRG.

Category 4 outages: Communication outage causing partition

This category is similar to category 2. These events occur:

- The status of the nodes not communicating by cluster messaging are set to Partition status. See [Cluster partition](#) for information about partitions.
- All nodes and cluster resource services on the nodes are still operational, but not all nodes can communicate with each other.
- The cluster is partitioned, but each CRG's primary node is still providing services.

The normal recovery for this partition state should be to repair the communication problem that caused the cluster partition. The cluster will resolve the partition state without any additional intervention.

Note: If you want the CRGs to fail over to a new primary node, ensure that the old primary node is not using the resources before the node is marked as failed. See [Change partitioned nodes to failed](#) for more information.

<i>Table 14. Category 4 outages: Communication outage causing partition</i>
Failover outage event
Communications adapter, line, or router failure on cluster heartbeat IP address lines occurs.
ENDTCPIFC is affecting all cluster heartbeat IP addresses on a cluster node.

Outages with active CRGs

- If the CRG is Active and the failing node is *not* the primary node, the following results:
 - The failover updates the status of the failed recovery domain member in the CRG's recovery domain.
 - If the failing node is a backup node, the list of backup nodes is reordered so that active nodes are at the beginning of the list.
- If the CRG is Active and the recovery domain member is the primary node, the actions the system performs depend on which type of outage has occurred.
 - Category 1 outages: Node outage causing failover
 - Category 2 outages: Node outage causing partition

- Category 3 outages: CRG fault causing failover
- Category 4 outages: Communication outage causing partition

Outages with inactive CRGs

When there is an outage with CRGs, the following occur:

- The membership status of the failed node in the cluster resource group's recovery domain is changed to either Inactive or Partition status.
- The node roles are not changed, and the backup nodes are not reordered automatically.
- The backup nodes are reordered in an Inactive CRG when the **Start Cluster Resource Group (STRCRG)** command or the `Start Cluster Resource Group (QcstStartClusterResourceGroup)` API is called.

Note: The Start Cluster Resource Group API will fail if the primary node is not active. You must issue the **Change Cluster Resource Group (CHGCRG)** command or the `Change Cluster Resource Group (QcstChangeClusterResourceGroup)` API to designate an active node as the primary node, and then call the Start Cluster Resource Group API again.

Managing cluster administrative domains

After a cluster administrative domain is created and the appropriate monitored resource entries (MREs) are added, the cluster administrator should monitor the activity within the administrative domain to ensure that the monitored resources remain consistent. Using PowerHA graphical interface or the PowerHA command-line interfaces, you can manage and monitor a cluster administrative domain.

Related tasks

Displaying cluster administrative domains

You can manage and monitor the status of cluster administrative domains in your high-availability environment.

Stopping a cluster administrative domain

Cluster administrative domains provide environment resiliency for resources within an IBM i high-availability solution. You might need to stop a cluster administrative domain to temporarily end synchronization of monitored resources.

Deleting a cluster administrative domain

Using the PowerHA graphical interface, you can delete a cluster administrative domain. Deleting a cluster administrative domain ends synchronization of monitored resources that are defined in the cluster administrative domain. In order to delete a cluster administrative domain, it must be inactive.

Changing the properties of a cluster administrative domain

Using the PowerHA graphical interface, you can change properties to an existing cluster administrative domain. These properties control synchronizations of monitored resource entries that are defined in the cluster administrative domain.

Managing monitored resource entries

The PowerHA graphical interface allows you to manage monitored resource entries in your cluster administrative domain. A cluster administrative domain ensures that changes made to these monitored resources remain consistent on each node within the high-availability environment.

Displaying cluster administrative domains

You can manage and monitor the status of cluster administrative domains in your high-availability environment.

This graphical interface provides the ability to list the MREs along with the global status for each resource. Detailed information can be displayed by selecting an MRE. This information includes the global value for each attribute that is associated with the MRE, along with an indication whether the attribute is consistent or inconsistent with the domain. If the global status of a monitored resource is inconsistent, the administrator should take the necessary steps to determine why the resource is inconsistent, correct the problem, and resynchronize the resource.

If the resource is inconsistent because an update failed on one or more nodes, information is kept for the MRE that can help you determine the cause of the failure. On the node where the failure occurred, a message is logged with the MRE as to the cause of the failed update. On the other nodes, there is an informational message logged internally, which tells you there was a failure, along with the list of nodes where the update failed. These messages are available through the PowerHA graphical interface or by calling the [Retrieve Monitored Resource Information \(QfpadRtvMonitoredResourceInfo\)](#) API. Failure messages are also logged in the job log of the administrative domain job.

After the cause of the inconsistency is determined, the resource can be resynchronized, either as a result of an update operation on the node where the failure occurred, or by ending and restarting the administrative domain. For example, an MRE for a user profile could be inconsistent because you changed the UID for the user profile on one node in the administrative domain, but the UID you specified was already in use by another user profile on one of the nodes. If you change the value of the UID again to something that is not used by another user profile within the administrative domain, the change will be made by the cluster administrative domain on all nodes and the global status for the user profile MRE is set to consistent. You do not need to take any further action to resynchronize the user profile MRE.

In some cases, you need to end and restart the cluster administrative domain CRG in order for the inconsistent resources to be resynchronized. For example, if you change the UID for a user profile that has an MRE associated with it, but the user profile is active in a job on one of the other cluster nodes in the administrative domain, the global value for the MRE associated with the user profile will be set to inconsistent because the change operation failed on the node where the user profile was active in a job. In order to correct this situation, you need to wait until the job has ended and then end the cluster administrative domain. When the administrative domain is started again, the global value for each attribute that is inconsistent will be used to change the resource to a consistent state.

The global status for a monitored resource is always set to failed if the resource is deleted, renamed, or moved on any node in the domain. If this is the case, the MRE should be removed because the resource is no longer synchronized by the cluster administrative domain.

When you restore a monitored resource on any system that is part of a cluster administrative domain, the resource is resynchronized to the global value currently known in the cluster administrative domain when the cluster administrative domain is active.

The following restore commands result in a resynchronization of system objects: RSTLIB, RSTOBJ, RSTUSRPRF, and RSTCFG. In addition, RSTSYSINF and UPDSYSINF result in a resynchronization of system values and network attributes. To resynchronize system environment variables after running the RSTSYSINF or UPDSYSINF commands, the peer CRG that represents the cluster administrative domain must be ended and started again.

If you want to restore your monitored resources to a previous state, remove the MRE that represents the resource that you want to restore. Then, after restoring the resource, add an MRE for the resource from the system where the restore operation was done. The cluster administrative domain will synchronize the monitored resource across the domain by using the values from the restored resource.

To display a cluster administrative domain using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click **Cluster Administrative Domain**.

Related tasks

[Managing cluster administrative domains](#)

After a cluster administrative domain is created and the appropriate monitored resource entries (MREs) are added, the cluster administrator should monitor the activity within the administrative domain to ensure that the monitored resources remain consistent. Using PowerHA graphical interface or the PowerHA command-line interfaces, you can manage and monitor a cluster administrative domain.

Related information

[Work with Monitored Resources \(WRKCADMRE\) command](#)

Stopping a cluster administrative domain

Cluster administrative domains provide environment resiliency for resources within an IBM i high-availability solution. You might need to stop a cluster administrative domain to temporarily end synchronization of monitored resources.

A cluster administrative domain becomes inactive when it is stopped. While the cluster administrative domain is inactive, all of the monitored resources are considered to be inconsistent because changes to them are not being synchronized. Although changes to monitored resources continue to be tracked, the global value is not changed and changes are not propagated to the rest of the administrative domain. Any changes that are made to any monitored resource while the cluster administrative domain is inactive are synchronized across all active nodes when the cluster administrative domain is restarted.

Note: The cluster administrative domain and its associated exit program are IBM-supplied objects. They should not be changed using the QcstChangeClusterResourceGroup API or the **CHGCRG** command. Making these changes will cause unpredictable results.

To stop a cluster administrative domain using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** in the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Administrative Domains**.
5. On the **Cluster Administrative Domains** page, select **Stop** from the context menu of the cluster administrative domain that you want to stop.
6. Click **Yes** on the confirmation panel.

Related tasks

Managing cluster administrative domains

After a cluster administrative domain is created and the appropriate monitored resource entries (MREs) are added, the cluster administrator should monitor the activity within the administrative domain to ensure that the monitored resources remain consistent. Using PowerHA graphical interface or the PowerHA command-line interfaces, you can manage and monitor a cluster administrative domain.

Related information

End Cluster Administrative Domain (ENDCAD) command

Deleting a cluster administrative domain

Using the PowerHA graphical interface, you can delete a cluster administrative domain. Deleting a cluster administrative domain ends synchronization of monitored resources that are defined in the cluster administrative domain. In order to delete a cluster administrative domain, it must be inactive.

To delete a cluster administrative domain using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Administrative Domains**.
5. On the **Cluster Administrative Domains** page, select **Delete** from the context menu of the cluster administrative domain that you want to stop.
6. Click **Yes** on the confirmation panel.

Related tasks

Managing cluster administrative domains

After a cluster administrative domain is created and the appropriate monitored resource entries (MREs) are added, the cluster administrator should monitor the activity within the administrative domain to ensure that the monitored resources remain consistent. Using PowerHA graphical interface or the PowerHA command-line interfaces, you can manage and monitor a cluster administrative domain.

Related information

Delete Cluster Admin Domain (DLTCAD) command

Changing the properties of a cluster administrative domain

Using the PowerHA graphical interface, you can change properties to an existing cluster administrative domain. These properties control synchronizations of monitored resource entries that are defined in the cluster administrative domain.

To change the properties of a cluster administrative domain using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Administrative Domains**.
5. On the **Cluster Administrative Domains** page, select **Properties...** from the context menu of the cluster administrative domain of which you want to change the properties.
6. On the **Properties** page, click **Edit**.
7. Make changes to the properties and click **Save**.

Related tasks

Managing cluster administrative domains

After a cluster administrative domain is created and the appropriate monitored resource entries (MREs) are added, the cluster administrator should monitor the activity within the administrative domain to ensure that the monitored resources remain consistent. Using PowerHA graphical interface or the PowerHA command-line interfaces, you can manage and monitor a cluster administrative domain.

Managing monitored resource entries

The PowerHA graphical interface allows you to manage monitored resource entries in your cluster administrative domain. A cluster administrative domain ensures that changes made to these monitored resources remain consistent on each node within the high-availability environment.

Related tasks

Managing cluster administrative domains

After a cluster administrative domain is created and the appropriate monitored resource entries (MREs) are added, the cluster administrator should monitor the activity within the administrative domain to ensure that the monitored resources remain consistent. Using PowerHA graphical interface or the PowerHA command-line interfaces, you can manage and monitor a cluster administrative domain.

Working with monitored resource entry status

The PowerHA graphical interface provides status messages for monitored resource entries within a cluster administrative domain.

After an MRE is added to the cluster administrative domain, the resource is monitored for changes on all administrative domain nodes so that the values of the resource attributes can be synchronized across the nodes in the cluster administrative domain. The synchronization behavior is dependent on a number of factors:

- Status of the cluster
- Status of the cluster administrative domain
- Status of the node
- Particular actions on the resource

To work with monitored resource entry status using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Administrative Domains**.
5. On the **Cluster Administrative Domains** page, select **Monitored Resource ...** from the context menu of the cluster administrative domain to which you want to add the monitored resource.

The following status information is available from the **Monitored Resources** page.

- The status of each monitored resource is shown in the **Global Status** column.
- The status of each individual attribute of a monitored resource is shown by selecting **Attributes...** from the context menu of the monitored resource of interest.
- The reason a monitored resource is inconsistent is shown by selecting **Node Details...** from the context menu of the monitored resource of interest.

Related information

[Work with Monitored Resources \(WRKCADMRE\) command](#)

Removing monitored resource entries

Monitored resource entries (MREs) are resources that are currently used within the high-availability environment and are monitored for changes through a cluster administrative domain. You might want to remove MREs when you no longer need them to be monitored. You can remove monitored resource entries (MREs) by using the PowerHA graphical interface.

To remove a monitored resource entry using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Administrative Domains**.
5. On the **Cluster Administrative Domains** page, select **Monitored Resource ...** from the context menu of the cluster administrative domain to which you want to remove the monitored resource.
6. On the **Monitored Resource Entries** page, select **Remove** from the context menu of the monitored resource that you want to remove.
7. Click **Yes** on the confirmation panel.

Related information

[Remove Admin Domain MRE \(RMVCADMRE\) command](#)

[Remove Monitored Resource Entry \(QfpadRmvMonitoredResourceEntry\) API](#)

Listing monitored resource entries

Monitored resource entries are resources, such as user profiles and environment variables, that have been defined in a cluster administrative domain. You can use the PowerHA graphical interface to list monitored resource entries that are currently defined in a cluster administrative domain.

To list monitored resource entries using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Administrative Domains**.
5. On the **Cluster Administrative Domains** page, select **Monitored Resource ...** from the context menu of the cluster administrative domain to which you want to add the monitored resource.

Selecting attributes to monitor

After you have added monitored resource entries, you can select attributes associated with that resource to be monitored by the cluster administrative domain.

To select attributes to monitor for a monitored resource entry (MRE), follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Administrative Domains**.

5. On the **Cluster Administrative Domains** page, select **Monitored Resource ...** from the context menu of the cluster administrative domain to which you want to add the monitored resource.
6. In the list of monitored resource types, click the context icon next to the monitored resource type, and select **Monitored Resource Entries....** The MRE object list is shown.
7. Click the context icon next to the MRE object, such as user profile or system value, and select **Work with Attributes**. The MRE Attributes List is shown.
8. In the MRE Attribute List window, select the attributes that you want to monitor, and then click **Close**. For example, if you want to monitor Ethernet line description for changes to its resource name attribute, you would select resource name as the attribute.

Related tasks

Adding monitored resource entries

You can add a monitored resource entry (MRE) to a cluster administrative domain. Monitored resource entries define critical resources so that changes made to these resources are kept consistent across a high-availability environment.

Attributes that can be monitored

A monitored resource entry can be added to the cluster administrative domain for various types of resources. This topic lists the attributes that each resource type can be monitored.

Resource types

- Authorization lists (*AUTL)
- Classes (*CLS)
- Ethernet line descriptions (*ETHLIN)
- Independent disk pools device descriptions (*ASPDEV)
- Job descriptions (*JOBDD)
- Network attributes (*NETA)
- Network server configuration for connection security (*NWSCFG)
- Network server configuration for remote systems (*NWSCFG)
- Network server configurations for service processors (*NWSCFG)
- Network server descriptions for iSCSI connections (*NWSD)
- Network server descriptions for integrated network servers (*NWSD)
- Network server descriptions for integrated network servers (*NWSD) (server operating system *AUX)
- Network server storage spaces (*NWSSTG)
- Network server host adapter device descriptions (*NWSHDEV)
- Optical device descriptions (*OPTDEV)
- Printer device descriptions for LAN connections (*PRTDEV)
- Printer device descriptions for virtual connections (*PRTDEV)
- Subsystem descriptions (*SBSD)
- System environment variables (*ENVVAR)
- System values (*SYSVAL)
- Tape device descriptions (*TAPDEV)
- Token-ring line descriptions (*TRNLIN)
- TCP/IP attributes (*TCPA)
- User profiles (*USRPRF)

*Table 15. Attributes that can be monitored for authorization lists (*AUTL)*

Attribute name	Description
AUT	Authority
OBJAUTE	Authority entry
OBJPGP	Primary group
OBJOWNER	Object owner
TEXT	Text description

*Table 16. Attributes that can be monitored for classes (*CLS)*

Attribute name	Description
CPUTIME	Maximum CPU time
DFTWAIT	Default wait time
MAXTHD	Maximum threads
MAXTMPSTG	Maximum temporary storage
OBJAUTE	Authority entry
OBJAUTL	Authorization list
OBJOWNER	Object owner
OBJPGP	Primary group
RUNPTY	Run priority
TEXT	Text description
TIMESLICE	Time slice

*Table 17. Attributes that can be monitored for Ethernet line descriptions (*ETHLIN)*

Attribute name	Description
ASSOCPORT	Associated port resource name
AUTOCRTCTL	Autocreate controller
AUTODLTCTL	Autodelete controller
CMNRCYLMT	Recovery limits
COSTBYTE	Relative cost per byte for sending and receiving data on the line
COSTCNN	Relative cost of being connected on the line
DUPLEX	Duplex
GENTSTFRM	Generate test frames
GRPADR	Group address
LINESPEED	Line speed
LINKSPEED	Link speed
MAXFRAME	Maximum frame size

*Table 17. Attributes that can be monitored for Ethernet line descriptions (*ETHLIN) (continued)*

Attribute name	Description
MAXCTL	Maximum controllers
MSGQ	Message queue
OBJAUTE	Authority entry
OBJAUTL	Authorization list
OBJOWNER	Object owner
OBJPGP	Primary group
ONLINE	Online at IPL
PRPDLY	Propagation delay
RSRCNAME	Resource name
SECURITY	Security level of the physical line
SSAP	Source service access point (SSAP) information list
TEXT	Text description
USRDFN1	First user-defined
USRDFN2	Second user-defined
USRDFN3	Third user-defined
VRYWAIT	Vary on wait

*Table 18. Attributes that can be monitored for independent disk pools device descriptions (*ASPDEV)*

Attribute name	Description
MSGQ	Message queue
OBJAUTE	Authority entry
OBJAUTL	Authorization list
OBJOWNER	Object owner
OBJPGP	Primary group
RDB	Relational database
RSRCNAME	Resource name
TEXT	Text description

*Table 19. Attributes that can be monitored for job descriptions (*JOBDD)*

Attribute name	Description
ACGCDE	Accounting code
ALWMLTTHD	Allow multiple threads
DDMCNV	DDM conversation
DEVRCYACN	Device recovery action
ENDSEV	End severity

Table 19. Attributes that can be monitored for job descriptions (*JOBQ) (continued)

Attribute name	Description
HOLD	Hold on job queue
INLASPGRP	Initial ASP group
INLLIBL	Initial library list
INQMSGRPY	Inquiry message reply
JOBMSGQFL	Job message queue full action
JOBMSGQMX	Job message queue maximum size
JOBPTY	Job priority (on JOBQ)
JOBQ	Job queue
LOG	Message logging
LOGCLPGM	Log CL program commands
OBJAUTE	Authority entry
OBJAUTL	Authorization list
OBJOWNER	Object owner
OBJPGP	Primary group
OUTPTY	Output priority (on OUTQ)
OUTQ	Output queue
PRTDEV	Print device
PRTTXT	Print text
RQSDTA	Request data or command
RTGDTA	Routing data
SPLFACN	Spoiled file action
SWS	Job switches
SYNTAX	CL syntax check
TEXT	Text description
TSEPOOL	Time slice end pool
USER	User

Table 20. Attributes that can be monitored for network attributes (*NETA)

Attribute name	Description
ALWADDCLU	Allow add to cluster
DDMACC	DDM/DRDA request access
NWSDOMAIN	Network server domain
PCSACC	Client request access
<p>Note: Each network attribute is treated as its own monitored resource entry. For these, the resource type and attribute names are identical.</p>	

*Table 21. Attributes that can be monitored for network server configurations for service processors (*NWSCFG)*

Attribute name	Description
EID	Enclosure identifier
INZSP	Initialize service processor
OBJAUTE	Authority entry
OBJAUTL	Authorization list
OBJOWNER	Object owner
OBJPGP	Primary group
SPAUT	Service processor authority
SPCERTID	Service processor certificate identifier
SPINTNETA	Service processor Internet address
SPNAME	Service processor name
TEXT	Text description

*Table 22. Attributes that can be monitored for network server configuration for remote systems (*NWSCFG)*

Attribute name	Description
BOOTDEVID	Boot device identifier
CHAPAUT	Target CHAP authentication
DELIVERY	Delivery method
DYNBOOTOPT	Dynamic boot options
INRCHAPAUT	Initiator CHAP authentication
OBJAUTE	Authority entry
OBJAUTL	Authorization list
OBJOWNER	Object owner
OBJPGP	Primary group
RMTIFC	Remote interfaces
RMTSYSID	Remote system identifier
SPNWSCFG	Service processor network server configuration that is used to manage the remote server
TEXT	Text description

*Table 23. Attributes that can be monitored for network server configuration for connection security (*NWSCFG)*

Attribute name	Description
IPSECRULE	IP security rules
OBJAUTE	Authority entry
OBJAUTL	Authorization list

Table 23. Attributes that can be monitored for network server configuration for connection security (*NWSCFG) (continued)

Attribute name	Description
OBJOWNER	Object owner
OBJPGP	Primary group
TEXT	Text description

Table 24. Attributes that can be monitored for Network server descriptions for integrated network servers (*NWSD)

Attribute name	Description
ALWDEVRSC	Allowed device resources
CFGFILE	Configuration file
CODEPAGE	ASCII code page representing the character set to be used by this network server
EVTLOG	Event log
MSGQ	Message queue
NWSSTGL	Storage space links
OBJAUTE	Authority entry
OBJAUTL	Authorization list
OBJOWNER	Object owner
OBJPGP	Primary group
PRPDMNUSR	Propagate domain user
RSRCNAME	Resource name
RSTDDEVRSC	Restricted device resources
SHUTDTIMO	Shut down time out
SYNCTIME	Synchronize date and time
TCPDMNNAME	TCP/IP local domain name
TCPHOSTNAM	TCP/IP host name
TCPPORTCFG	TCP/IP port configuration
TCPNAMSVR	TCP/IP name server system
TEXT	Text description
VRYWAIT	Vary on wait
WINDOWSNT	Windows network server description

Table 25. Attributes that can be monitored for Network server descriptions for integrated network servers (server operating system *AUX) (*NWSD)

Attribute name	Description
CODEPAGE	ASCII code page representing the character set to be used by this network server

*Table 25. Attributes that can be monitored for Network server descriptions for integrated network servers (server operating system *AUX) (*NWSD) (continued)*

Attribute name	Description
DSBUSRPRF	Disable user profiles
EVTLOG	Event log
MSGQ	Message queue
OBJAUTE	Authority entry
OBJAUTL	Authorization list
OBJOWNER	Object owner
OBJPGP	Primary group
PRPDMNUSR	Propagate domain user
RSRCNAME	Resource name
SHUTDTIMO	Shut down time out
SRVOPT	Serviceability options
SYNCTIME	Synchronize date and time
TCPDMNNAME	TCP/IP local domain name
TCPHOSTNAM	TCP/IP host name
TCPPORTCFG	TCP/IP port configuration
TEXT	Text description
VRYWAIT	Vary on wait

*Table 26. Attributes that can be monitored for network server descriptions for iSCSI connections (*NWSD)*

Attribute name	Description
ACTTMR	Activation timer
ALWDEVRSC	Allowed device resources
CFGFILE	Configuration file
CMNMSGQ	Communications message queue
CODEPAGE	ASCII code page representing the character set to be used by this network server
DFTSECRULE	Default IP security rule
DFTSTGPTH	Default storage path
EVTLOG	Event log
MLTPTHGRP	Multi-path group
MSGQ	Message queue
NWSCFG	Network server configuration
NWSSTGL	Storage space links
OBJAUTE	Authority entry

Table 26. Attributes that can be monitored for network server descriptions for iSCSI connections (*NWSD) (continued)

Attribute name	Description
OBJAUTL	Authorization list
OBJOWNER	Object owner
OBJPGP	Primary group
PRPDMNUSR	Propagate domain user
RMVMEDPTH	Removable media path
RSRCNAME	Resource name
RSTDDEVRSC	Restricted device resources
SHUTDTIMO	Shut down time out
STGPTH	iSCSI storage paths of the network server
SVROPT	Serviceability options
SYNCTIME	Synchronize date and time
TCPDMNNAME	TCP/IP local domain name
TCPHOSTNAM	TCP/IP host name
TCPNAMSVR	TCP/IP name server system
TCPPORTCFG	TCP/IP port configuration
TEXT	Text description
VRTETHCTLP	Virtual Ethernet control port
VRTETHPTH	Virtual Ethernet path
VRYWAIT	Vary on wait

Table 27. Attributes that can be monitored for network server storage spaces (*NWSSTG)

Attribute name	Description
OBJAUTE	Authority entry
OBJAUTL	Authorization list
OBJOWNER	Object owner
OBJPGP	Primary group
RSCALCPTY	Resource allocate priority
SIZE	Size
TEXT	Text description
TOTALFILES	Total files

Table 28. Attributes that can be monitored for network server host adapter device descriptions (*NWSHDEV)

Attribute name	Description
CMNRCYLMT	Recovery limits

*Table 28. Attributes that can be monitored for network server host adapter device descriptions (*NWSHDEV) (continued)*

Attribute name	Description
LCLIFC	Associated local interface
MSGQ	Message queue
OBJAUTE	Authority entry
OBJAUTL	Authorization list
OBJOWNER	Object owner
OBJPGP	Primary group
ONLINE	Online at IPL
RSRCNAME	Resource name
TEXT	Text description

*Table 29. Attributes that can be monitored for optical device descriptions *OPTDEV)*

Attribute name	Description
MSGQ	Message queue
OBJAUTE	Authority entry
OBJAUTL	Authorization list
OBJOWNER	Object owner
OBJPGP	Primary group
ONLINE	Online at IPL
RSRCNAME	Resource name
TEXT	Text description

*Table 30. Attributes that can be monitored for printer device descriptions for *LAN connections (*PRTDEV)*

Attribute name	Description
ACTTMR	Activation timer
ADPTADR	LAN remote adapter address
ADPTTYPE	Adapter type
ADPTCNNTYP	Adapter connection type
AFP	Advanced function printing
CHRID	Character identifier
FONT	Font
FORMFEED	Formfeed
IMGCFG	Image configuration
INACTTMR	Inactivity timer
LNGTYPE	Language type

Table 30. Attributes that can be monitored for printer device descriptions for *LAN connections (*PRTDEV) (continued)

Attribute name	Description
LOCADR	Location location address
MAXPNDRQS	Maximum pending request
MFRTYPMDL	Manufacturer type and model
MSGQ	Message queue
OBJAUTE	Authority entry
OBJAUTL	Authorization list
OBJOWNER	Object owner
OBJPGP	Primary group
ONLINE	Online at IPL
PORT	Port number
PRTERMSG	Print error message
PUBLISHINF	Publishing information
RMTLOCNAME	Remote location
SEPDRAWER	Separator drawer
SEPPGM	Separator program
SWTLINLST	Switched line list
SYSDRVPGM	System driver program
TEXT	Text description
TRANSFORM	Host printer transform
USRDFNOBJ	User-defined object
USRDFNOPT	User-defined options
USRDRVPGM	User-defined driver program
USRDTATFM	Data transform program
WSCST	Workstation customizing object

Table 31. Attributes that can be monitored for printer device descriptions for *VRT connections (*PRTDEV)

Attribute name	Description
CHRID	Character identifier
FORMFEED	Form feed
IGCFEAT	DBCS FEATURE
IMGCFG	Image configuration
MAXLENRU	Maximum length of request unit
MFRTYPMDL	Manufacturer type and model
MSGQ	Message queue

*Table 31. Attributes that can be monitored for printer device descriptions for *VRT connections (*PRTDEV) (continued)*

Attribute name	Description
OBJAUTE	Authority entry
OBJAUTL	Authorization list
OBJOWNER	Object owner
OBJPGP	Primary group
ONLINE	Online at IPL
PRTERMSG	Print error message
PUBLISHINF	Publishing information
SEPDRAWER	Separator drawer
SEPPGM	Separator program
TEXT	Text description
TRANSFORM	Host print transform
USRDFNOBJ	User-defined object
USRDFNOPT	User-defined options
USRDRVPGM	User-defined driver program
USRDTAFM	Data transform program
WSCST	Workstation customizing object
SEPPGM	Separator program
SWTLINLST	Switched line list
SYSDRVPGM	System driver program
TEXT	Text description
TRANSFORM	Host printer transform
USRDFNOBJ	User-defined object
USRDFNOPT	User-defined options
USRDRVPGM	User-defined driver program
USRDTATFM	Data transform program
WSCST	Workstation customizing object

*Table 32. Attributes that can be monitored for subsystem descriptions (*SBSD)*

Attribute name	Description
AJE	Autostart job entry
CMNE	Online at IPL
JOBQE	Job queue
MAXJOBS	Maximum number of jobs
OBJAUTE	Authority entry

Table 32. Attributes that can be monitored for subsystem descriptions (*SBSD) (continued)

Attribute name	Description
OBJAUTL	Authorization list
OBJOWNER	Object owner
OBJPGP	Primary group
PJE	Prestart job entry
RMTLOCNAME	Remote location name
RTGE	Routing entry
SGNDSPF	Sign on display
SYSLIBLE	Subsystem library
TEXT	Text description
WSNE	Workstation name entry
WSTE	Workstation type entry

Table 33. Attributes that can be monitored for system environment variables (*ENVVAR)

Attribute name	Description
	Any *SYS level environment variable can be monitored. The attribute and resource name are both the same as the environment variable's name.
	Note: Each environment variable is treated as its own monitored resource entry. For these, the resource type and attribute names are identical.

Table 34. Attributes that can be monitored for system values (*SYSVAL)

Attribute name	Description
QACGLVL	Accounting level
QACTJOBITP	Allow jobs to be interrupted
QALWOBJRST	Prevents anyone from restoring a system-state object or an object that adopts authority
QALWUSRDMN	Allows user domain objects
QASTLVL	Assistance level
QATNPGM	Attention program
QAUDCTL	Audit control
QAUDENDACN	Audit journal error action
QAUDFRCLVL	Auditing force level
QAUDLVL	Auditing level
QAUDLVL2	Auditing level extension
QAUTOCFG	Automatic device configuration

Table 34. Attributes that can be monitored for system values (*SYSVAL) (continued)

Attribute name	Description
QAUTORMT	Remote controllers and devices
QAUTOVRT	Automatic virtual device configuration
QCCSID	Coded character set identifier
QCFGMSGQ	Message queue for lines, controllers, and devices
QCHRID	Default graphic character set and code page used for displaying or printing data
QCHRIDCTL	Character identifier control for the job
QCMNRCYLMT	Automatic communications error recovery
QCNTYID	Country or region identifier
QCRTAUT	Authority for new objects
QCRTOBJAUD	Auditing new objects
QCTLSBSD	Controlling subsystem or library
QCURSYM	Currency symbol
QDATFMT	Date format
QDATSEP	Date separator
QDBRCVYWT	Wait for database recovery before completing restart
QDECfmt	Decimal format
QDEVNAMING	Device naming convention
QDEVRcyACN	Device recovery action
QDSCJOBITV	Time out interval for disconnected jobs
QDSPSGNINF	Controls the display of sign-on information
QENDJOB LMT	Maximum time for immediate end
QFRCCVNRST	Force conversion on restore
QHSTLOGSIZ	History log file size
QIGCCDEFNT	Coded font name
QIGCFNTSIZ	Coded font point size
QINACTIV	Inactive job time-out interval
QINACTMSGQ	Timeout interval action
QIPLTYPE	Type of restart
QJOBMSGQFL	Job message queue full action
QJOBMSGQMX	Job message queue maximum size
QJOBMSGQSZ	Initial size of job message queue in kilobytes (KB)
QJOBMSGQTL	Maximum size of job message queue (in KB)

Table 34. Attributes that can be monitored for system values (*SYSVAL) (continued)

Attribute name	Description
QJOBSPLA	Initial size of spooling control block for a job (in bytes)
QKBDBUF	Keyboard buffer
QKBDTYPE	Keyboard language character set
QLANGID	Default language identifier
QLIBLCKLVL	Lock libraries in a user job's library search list
QLMTDEVSSN	Limit device sessions
QLMTSECOFR	Limit security officer device access
QLOCALE	Locale
QLOGOUTPUT	Produce printer output for job log
QMAXACTLVL	Maximum activity level of the system
QMAXJOB	Maximum number of jobs that are allowed on the system
QMAXSGNACN	The system's response when the limit imposed by QMAXSIGN system value is reached
QMAXSIGN	Maximum number of not valid sign-on attempts allowed
QMAXSPLF	Maximum printer output files
QMLTTHDACN	When a function in a multithreaded job is not threadsafe
QPASTHRSVR	Available display station pass-through server jobs
QPRBFTR	Problem log filter
QPRBHDITV	Minimum retention
QPRTDEV	Default printer
QPRTKEYFMT	Print key format
QPRTTXT	Up to 30 characters of text that can be printed at the bottom of listings and separator pages
QPWDCHGBLK	Minimum time between password changes
QPWDEXPITV	Number of days for which a password is valid
QPWDEXPWRN	Password expiration warning interval system
QPWDLMTACJ	Limits the use of adjacent numbers in a password
QPWDLMTCHR	Limits the use of certain characters in a password
QPWDLMTREP	Limits the use of repeating characters in a password
QPWDLVL	Password level
QPWDMAXLEN	Maximum number of characters in a password
QPWDMINLEN	Minimum number of characters in a password

Table 34. Attributes that can be monitored for system values (*SYSVAL) (continued)

Attribute name	Description
QPWDPOSDIF	Controls the position of characters in a new password
QPWDRQDDGT	Require a number in a new password
QPWDRQDDIF	Controls whether the password must be different from the previous passwords
QPWDRULES	Password rules
QPWDVLDPGM	Password approval program
QPWRDWNLMT	Maximum time for immediate shutdown
QRCLSPLSTG	Automatically clean up unused printer output storage
QRETSVRSEC	Retain server security data indicator
QRMTSIGN	Remote sign-on
QRMTSRVATR	Remote service attribute
QSCANFS	Scan file systems
QSCANFCTL	Scan control
QSCPFCONS	Console problem occurs
QSECURITY	System security level
QSETJOBATR	Set job attributes
QSFWERRLOG	Software error log
QSHRMEMCTL	Allow use of shared or mapped memory with write capability
QSPCENV	Default user environment
QSPLFACN	Spooled file action
QSRTSEQ	Sort sequence
QSRVDMP	Service log for unmonitored escape messages
QSSLCSL	Secure Sockets Layer cipher specification list
QSSLCSLCTL	Secure Sockets Layer cipher control
QSSLPCL	Secure Sockets Layer protocols
QSTRUPPGM	Set startup program
QSTMSG	Display status messages
QSYSLIBL	System library list
QTIMSEP	Time separator
QTSEPOOL	Indicates whether interactive jobs should be moved to another main storage pool when they reach time slice end
QUPSMGQ	Uninterruptible power supply message queue

Table 34. Attributes that can be monitored for system values (*SYSVAL) (continued)

Attribute name	Description
QUSEADPAUT	Use adopted authority
QUSRLIBL	User part of the library list
QVFOBJRST	Verify object on restore
Note: Each system value is treated as its own monitored resource entry. For these, the resource type and attribute names are identical.	

Table 35. Attributes that can be monitored for tape device descriptions (*TAPDEV)

Attribute name	Description
ASSIGN	Assign device at vary on
MSGQ	Message queue
OBJAUTE	Authority entry
OBJAUTL	Authorization list
OBJOWNER	Object owner
OBJPGP	Primary object
ONLINE	Online at IPL
RSRCNAME	Resource name
TEXT	Text description
UNLOAD	Unload device at vary off

Table 36. Attributes that can be monitored for token-ring descriptions (*TRNLIN)

Attribute name	Description
ACTLANMGR	Activate LAN manager
ADPTADR	Local adapter address
AUOCRTCTL	Autocreate controller
AUTODLTCTL	Autodelete controller
CMNRCYLMT	Recovery limits
COSTBYTE	Relative cost per byte for sending and receiving data on the line
COSTCNN	Relative cost of being connected on the line
DUPLEX	Duplex
ELYTKNRLS	Early token release
FCNADR	Functional address
LINESPEED	Line speed
LINKSPEED	Link speed
LOGCFGCHG	Log configuration changes
MAXCTL	Maximum controllers

Table 36. Attributes that can be monitored for token-ring descriptions (*TRNLIN) (continued)

Attribute name	Description
MAXFRAME	Maximum frame size
MSGQ	Message queue
OBJAUTE	Authority entry
OBJAUTL	Authorization list
OBJOWNER	Object owner
OBJPGP	Primary group
ONLINE	Online at IPL
PRPDLY	Propagation delay
RSRCNAME	Resource name
SECURITY	Security for line
SSAP	Source service access point (SSAP) information list
TRNINFBDN	Token-ring inform of beacon
TRNLOGLVL	TRLAN manager logging level
TRNMGRMODE	TRLAN manager mode
TEXT	Text description of the token-ring line
USRDFN1	First user-defined
USRDFN2	Second user-defined
USRDFN3	Third user-defined
VRYWAIT	Vary on wait

Table 37. Attributes that can be monitored for TCP/IP attributes (*TCPA)

Attribute name	Description
ARPTIMO	Address resolution protocol (ARP) cache timeout
ECN	Enable explicit congestion notification (ECN)
IP6TMPAXP	IPv6 temporary address excluded prefix
IPDEADGATE	IP dead gateway detection
IPDTGFWD	IP datagram forwarding
IPPATHMTU	Path maximum transmission unit (MTU) discovery
IPQOSBCH	IP QoS datagram batching
IPQOSEN	IP QoS enablement
IPQOSTMR	IP QoS timer resolution
IPRSBTIMO	IP reassembly timeout
IPSRCRTG	IP source routing
IPTTL	IP time to live (hop limit)
LOGPCLERR	Log protocol errors

Table 37. Attributes that can be monitored for TCP/IP attributes (*TCPA) (continued)

Attribute name	Description
NFC	Network file cache
TCPCLOTIMO	TCP time-wait timeout
TCPCNNMSG	TCP close connection message
TCPKEEPALV	TCP keep alive
TCPMINRTM	TCP minimum retransmit time
TCPR1CNT	TCP R1 retransmission count
TCPR2CNT	TCP R2 retransmission count
TCPRCVBUF	TCP receive buffer size
TCPSNDBUF	TCP send buffer size
TCPURGPTR	TCP urgent pointer
UDPCKS	UDP checksum
Note: Each TCP/IP attribute is treated as its own monitored resource entry. For these, the resource type and attribute names are identical.	

Table 38. Attributes that can be monitored for user profiles (*USRPRF)

Attribute name	Description
ACGCDE	Accounting code
ASTLVL	Assistance level
ATNPGM	Attention program
CCSID	Coded character set ID
CHRIDCTL	Character identifier control
CNTRYID	Country or region ID
CURLIB	Current library
DLVRY	Delivery
DSPSGNINF	Display sign-on information
GID	Group ID number
GRPAUT	Group authority
GRPAUTYP	Group authority type
GRPPRF	Group profile
HOMEDIR	Home directory
INLMNU	Initial menu
INLPGM	Initial program to call
JOB	Job description
KBDBUF	Keyboard buffering
LANGID	Language ID

Table 38. Attributes that can be monitored for user profiles (*USRPRF) (continued)

Attribute name	Description
LCLPWDMGT	Local password management
LMTCPB	Limit capabilities
LMTDEVSSN	Limit device sessions
LOCALE	Locale
MAXSTG	Maximum allowed storage
MAXSTGLRG	Maximum allowed storage large
MSGQ	Message queue
OBJAUTE	Authority entry
OBJOWNER	Object owner
OBJPGP	Primary group
OUTQ	Output queue
OWNER	Owner
PASSWORD	User password
PRTDEV	Print device
PTYLMT	Highest schedule priority
PWDEXP	Set password to expired
PWDEXPITV	Password expiration interval
SETJOBATR	Locale job attributes
SEV	Severity code filter
SPCAUT	Special authority
SPCENV	Special environment
SRTSEQ	Sort sequence
STATUS	Status
SUPGRPPRF	Supplemental groups
TEXT	Text description
UID	User ID number
USRCLS	User class
USREXPDATE	User expiration date
USREXPITV	User expiration interval
USROPT	User options

Displaying monitored resource entry messages

Using the PowerHA graphical interface, you can display messages that are associated with monitored resource entries. These messages can help explain why a particular monitored resource entry is inconsistent.

To view monitored resource entry messages using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Administrative Domains**.
5. On the **Cluster Administrative Domains** page, select **Monitored Resource ...** from the context menu of the cluster administrative domain to which you want to add the monitored resource.
6. On the **Monitored Resources** page, select **Node Details ...** from the context menu of the monitored resource entry of which you want to see messages.

Managing independent disk pools

In a high-availability environment, it is necessary that the application and its data remain consistent among the nodes that participate in high availability. An independent disk pool, also called an independent auxiliary storage pool (IASP), is a disk pool that contains objects, the directories, or libraries that contain the objects, and other object attributes such as authorization and ownership attributes. An independent disk pool can be used to help in this situation. Ownership of data and applications that are stored in an independent disk pool can be switched to other systems that are defined in the device CRG. PowerHA technology provides high availability during planned and some unplanned outages.

Related information

[IBM eServer iSeries Independent ASPs: A Guide to Moving Applications to IASPs](#)

Making a disk pool unavailable

You can select an independent disk pool to make it unavailable (vary it off). You cannot access any of the disk units or objects in the independent disk pool or its corresponding database until it is made available (varied on) again. The pool can be made available again on the same system or another system in the recovery domain of the cluster resource group.

Important: Before an independent disk pool can be made unavailable, no jobs can hold reservations on the disk pool. See [Release job reservations on an independent disk](#) for details on determining whether jobs are using an independent disk pool and how to release the job reservations.

To make an independent disk pool unavailable using the PowerHA graphical interface, follow these steps:

1. In a web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system using your user profile and password.
3. Click **PowerHA** in the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, click **Show All Others...** to show any independent disk pools that are not currently highly available.
6. Select **Vary Off** from the context menu of the independent disk pool that you want to make unavailable.
7. A progress panel is displayed that shows the current step in the process.

When making a UDFS disk pool unavailable using IBM Navigator for i, messages might be generated that require a response in the character-based interface. IBM Navigator for i will not provide any indication that a message is waiting.

To make an independent disk pool unavailable using your IBM Navigator for i, follow these steps::

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Pools**.
3. Right-click the disk pool that you want to make unavailable and select **Make Unavailable**.
4. Click **Make Unavailable** on the dialog box that is displayed to make the disk pool unavailable.

You can use the Vary Configuration (VRYCFG) command in the character-based interface to make the disk pool unavailable.

Use the Display ASP Status (DSPASPSTS) command to identify where a step is in the process.

Use the Control ASP Access (QYASPCTLAA) API to restrict the processes that have access to the ASP.

Use the Start DASD Management Operation (QYASSDMO) API to reduce the amount of time it takes to make a disk pool unavailable.

Quiescing an independent disk pool

In an IBM i high-availability solution, independent disk pools are used to store resilient data and applications. Some system functions, such as performing backups, require that you temporarily suspend changes to that data while the operation occurs.

To decrease the amount of time it takes to quiesce an independent disk pool, you might want to hold batch job queues, end some subsystems, or send a break message to interactive users, advising them to postpone new work.

To quiesce an independent disk pool, complete these steps.

In a command line interface, enter the following command: **CHGASPACT ASPDEV(name) OPTION(*SUSPEND) SSPTIMO(30) SSPTIMOACN(*CONT)**, where *name* is the name of the independent disk pool that you want to suspend. In this command you are specifying to suspend the independent disk pool with a 30-second timeout, and to continue with the next step even if the timeout limit has been exceeded.

Resuming an independent disk pool

After you have quiesced an independent disk pool in an IBM i high availability environment for backup operations, you will need to resume the independent disk pool to ensure changes that are made to the data during the quiesce are updated.

Complete these steps to resume an independent disk pool:

In a command-line interface, enter the following command: **CHGASPACT ASPDEV(name) OPTION(*RESUME)**, where *name* is the name of the independent disk pool that you want to resume.

Managing geographic mirroring

Use the following information to help you manage geographic mirroring. Geographic mirroring is a sub-function of cross-site mirroring, where data is mirrored to independent disk pools in an IBM i environment.

Suspending geographic mirroring

If you need to end TCP communications for any reason, such as placing your system in restricted state, you should suspend geographic mirroring first. This action temporarily stops mirroring between systems in a high-availability solution.

When you suspend mirroring, any changes that are made on the production copy of the independent disk pool are not being transmitted to the mirror copy.

Note: When you resume geographic mirroring, synchronization is required between the production and mirror copies. If geographic mirroring was suspended without tracking, then full synchronization occurs. This can be a lengthy process.

Suspending geographic mirroring when IBM PowerHA for i is installed

To use the PowerHA graphical interface to suspend geographic mirroring, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where *mysystem* is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, select **Details...** from the context menu of the independent disk pool for which you want to suspend geographic mirroring.
6. On the **Independent ASP Details** page, select **Suspend Mirroring** from the **Select Action** menu below the mirroring status.
7. Click **Yes** on the confirmation panel.

Suspending geographic mirroring with command line interface when IBM PowerHA SystemMirror for i licensed program is installed

To suspend a geographic mirroring by using the command line interface, issue the following command:

- CHGASPSSN SESSION(session-name) OPTION(*SUSPEND) where session-name is the name of your geographic mirroring session.

Suspending geographic mirroring when IBM PowerHA for i is not installed

To suspend geographic mirroring with IBM Navigator for i, follow these steps:

1. In IBM Navigator for i, expand **My Connections** (or your active environment).
2. Expand the system that owns the production copy of the geographically mirrored disk pool that you want to suspend.
3. Expand **Configuration and Service > Hardware > Disk Units > Disk Pools**.
4. Right-click the production copy of the **Disk Pool** you want to suspend and select **Geographic Mirroring > Suspend Geographic Mirroring**.

If you suspend with tracking, the system attempts to track changes that are made to those disk pools. This might shorten the length of the synchronization process by performing partial synchronization when you resume geographic mirroring. If tracking space is exhausted, then when you resume geographic mirroring, complete synchronization is required.

Note: If you suspend geographic mirroring without tracking changes, then when you resume geographic mirroring, a complete synchronization is required between the production and mirror copies. If you suspend geographic mirroring and you do track changes, then only a partial synchronization is required. Complete synchronization can be a lengthy process, anywhere from one to several hours, or longer. The length of time it takes to synchronize is dependent on the amount of data being synchronized, the speed of TCP/IP connections, and the number of communication lines that are used for geographic mirroring.

Related information

[Change ASP Session \(CHGASPSSN\) command](#)

[Change SVC Session \(CHGSVCSSN\) command](#)

Resuming geographic mirroring

If you suspend geographic mirroring, you must resume it in order to reactivate mirroring between the production and mirrored copies again.

Note: When you resume geographic mirroring, the production and mirror copies are synchronized concurrent with performing geographic mirroring. Synchronization can be a lengthy process. If a disk pool becoming unavailable interrupts synchronization, then synchronization continues from where it was interrupted when the disk pool becomes available again. When an interrupted synchronization is continued, the first message (CPI0985D) states that the synchronization is 0% complete.

Resuming geographic mirroring when PowerHA is installed

To use the PowerHA graphical interface to resume geographic mirroring, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, select **Details...** from the context menu of the independent disk pool for which you want to resume geographic mirroring.
6. On the **Independent ASP Details** page, select **Resume Mirroring** from the **Select Action** menu below the mirroring status.
7. Click **Yes** on the confirmation panel.

Resuming geographic mirroring with command line interface when IBM PowerHA SystemMirror for i licensed program is installed

To resume a geographic mirroring session by using the command line interface, issue the following command:

- CHGASPSSN SESSION(session-name) OPTION(*RESUME) where session-name is the name of your geographic mirroring session.

Resuming geographic mirroring when PowerHA is not installed

To resume geographic mirroring using PowerHA, follow these steps:

1. In IBM Navigator for i, expand **My Connections** (or your active environment).
2. Expand the system that owns the production copy of the disk pool for which you want to resume geographic mirroring.
3. Expand **Configuration and Service > Hardware > Disk Units > Disk Pools**.
4. Right-click the **Disk Pool** you want to resume and select **Geographic Mirroring > Resume Geographic Mirroring**.

Related information

[Change ASP Session \(CHGASPSSN\) command](#)

[Change SVC Session \(CHGSVCSSN\) command](#)

Detaching mirror copy

If you are using geographic mirroring and want to access the mirror copy to perform save operations or data mining, or to create reports, you must detach the mirror copy from the production copy.

Detaching the mirror copy when IBM PowerHA for i is installed

To use the PowerHA graphical interface to detach the mirror copy, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, select **Details...** from the context menu of the independent disk pool for which you want to detach the mirror copy.
6. On the **Independent ASP Details** page, select **Detach** from the context menu of the mirror copy.
7. Click **Yes** on the confirmation panel.

Detaching the geographic mirroring with command line interface when IBM PowerHA SystemMirror for i licensed program is installed

To detach geographic mirroring using the command line interface, issue the following command:

- CHGASPSSN SESSION(session-name) OPTION(*DETACH) where session-name is the name of your geographic mirroring session.

Detaching the mirror copy when IBM PowerHA for i is not installed

It is recommended, but not required, that you make the independent disk pool unavailable to ensure that the production copy is not altered while the detachment is being performed.

To detach the mirror copy by using IBM Navigator for i, follow these steps:

1. In IBM Navigator for i, expand **My Connections** (or your active environment).
2. Expand the system that owns the production copy of the disk pool from which you want to detach the mirror copy.
3. Expand **Configuration and Service > Hardware > Disk Units > Disk Pools**.
4. Right-click the production copy of the **Disk Pool** you want to detach and select **Geographic Mirroring > Detach Mirror Copy**.

If **Geographic Mirroring > Detach Mirror Copy** cannot be clicked because it is disabled, the mirror copy is not in sync with the production copy, Geographic mirroring must be resumed, the disk pool varied on, and production and mirror copies that are synchronized before the mirror copy can be detached.

Related information

[Change ASP Session \(CHGASPSSN\) command](#)

[Change SVC Session \(CHGSVCSSN\) command](#)

Reattaching mirror copy

If you detached the mirror copy and have completed your work with the detached mirror copy, you must reattach the detached mirror copy to resume using geographic mirroring.

The detached mirror copy must be unavailable when you reattach it to the production copy.

Note: When you reattach the detached mirror copy, as of V6R1, there is the option to detach with tracking, which only requires a partial synchronization on reattach. Changes that are made on the Production Copy will be applied to the Mirror Copy, and all changes that are made on the Detached Mirror copy will be overwritten with the contents from the Production Copy.

Reattaching the mirror copy when IBM PowerHA SystemMirror for i licensed program is installed

To use the PowerHA graphical interface to reattach the mirror copy, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, select **Details...** from the context menu of the independent disk pool for which you want to reattach the mirror copy.
6. On the **Independent ASP Details** page, select **Reattach** from the context menu of the mirror copy.
7. Click **Yes** on the confirmation panel.

Reattaching geographic mirroring with command line interface when IBM PowerHA SystemMirror for i licensed program is installed

To reattach geographic mirroring using the command line interface, issue the following command:

- `CHGASPSSN SESSION(session-name) OPTION(*REATTACH)` where `session-name` is the name of your geographic mirroring session.

Reattaching the mirror copy when IBM PowerHA SystemMirror for i is not installed

To reattach the mirror copy using IBM Navigator for i, follow these steps:

1. In IBM Navigator for i, expand **My Connections** (or your active environment).
2. Expand the system that owns the production copy of the disk pool to which you want to reattach the detached mirror copy.
3. Expand **Configuration and Service > Hardware > Disk Units > Disk Pools**.
4. Right-click the production copy of the **Disk Pool** you want to reattach and select **Geographic Mirroring > Reattach Mirror Copy**.

Related information

[Change ASP Session \(CHGASPSSN\) command](#)

[Change SVC Session \(CHGSVCSSN\) command](#)

Deconfiguring geographic mirroring

If you no longer want the capability to use geographic mirroring for a specific disk pool or disk pool group, you can select to **Deconfigure Geographic Mirroring**. If you deconfigure geographic mirroring, the system

stops geographic mirroring and deletes the mirror copy of the disk pools on the nodes in the mirror copy site.

The disk pool must be offline to deconfigure geographic mirroring and the cluster resource group (CRG) must have a status of INACTIVE.

Deconfigure geographic mirroring when IBM PowerHA SystemMirror for i licensed program is installed

To use the IBM Navigator for i deconfigure geographic mirroring with, follow these steps:

1. In a web browser, enter **http://mysystem:2001**, where **mysystem** is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** in the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, select **Details...** from the context menu of the independent disk pool for which you want to deconfigure geographic mirroring.
6. On the **Independent ASP Details** page, select **Deconfigure Geographic Mirroring...** from the **Select Action** menu below the mirroring status.
7. Click **OK** on the confirmation panel.

Deconfiguring geographic mirroring with command line interface when IBM PowerHA SystemMirror for i licensed program is installed

To deconfigure geographic mirroring using the command line interface, issue the following command:

- CHGASPSN SESSION(session-name) OPTION(*DELETE) where session-name is the name of the independent ASP to deconfigure geographic mirroring on.

Deconfigure geographic mirroring when IBM PowerHA SystemMirror for i licensed program is not installed

1. In IBM Navigator for i, expand **My Connections** (or your active environment).
2. Expand the system that you want to examine, **Configuration and Service > Hardware > Disk Units > Disk Pools**.
3. Right-click the production copy of the **Disk Pool** you want to deconfigure and select **Geographic Mirroring>Deconfigure Geographic Mirroring**.
4. Update your cluster configuration, as follows:
 - a) Remove the nodes that are associated with the mirror copy from the device cluster resource group (CRG) recovery domain.
 - b) Remove the site name and data port IP addresses from the remaining nodes in the cluster.

Related tasks

Removing nodes

You might need to remove a node from a cluster if you are performing an upgrade of that node or if the node no longer needs to participate in the IBM i high-availability environment.

Related information

[Change ASP Session \(CHGASPSN\) command](#)

Changing geographic mirroring properties

All geographic mirroring properties can be changed when the production copy of the independent disk pool is varied off. The Synchronization Priority property can be also changed with the independent disk pool is Available and geographic mirroring is Suspended.

Changing geographic mirroring properties with IBM PowerHA SystemMirror for i licensed program

To change the geographic mirroring properties by using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter **http://mysystem:2001**, where **mysystem** is the host name of the system.

2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, select **Details...** from the context menu of the independent disk pool for which you want to change the geographic mirroring properties.
6. On the **Independent ASP Details** page, select **Mirroring Properties...** from the **Select Action** menu below the mirroring status.
7. On the **Properties** page, click **Edit**.
8. Make changes to the properties and click **Save**.

Changing geographic mirroring with command line interface when IBM PowerHA SystemMirror for i licensed program is installed

To change the geographic mirroring session using the command line interface, issue the following command:

- CHGASPSN SESSION(session-name) OPTION(*CHGATTR) where session-name is the name of your geographic mirroring Mirror session. You need to list the attributes that should be changed.

Changing geographic mirroring properties with IBM Navigator for i

To change the geographic mirroring properties by using IBM Navigator for i, follow these steps:

1. In IBM Navigator for i, expand **My Connections** (or your active environment).
2. Expand the system that owns the production copy of the geographically mirrored disk pool associated with the geographic mirror session for which you want to edit the attributes, **Configuration and Service > Hardware > Disk Units > Disk Pools**.
3. Right-click the production copy of the **Disk Pool** for which you want to edit the attributes and select **Sessions > Open**.
4. Right-click the production copy of the **Session** for which you want to edit the attributes and select **Properties**. To change an associated copy description, select the copy description and click **Edit**.

Managing Metro Mirror

In PowerHA environment that use IBM System Storage Metro Mirror technology, you must configure a metro mirroring session between the IBM i systems and the external disk units with Metro Mirror configured. From the system, you can manage these sessions.

Suspending Metro Mirror

You might need to suspend Metro Mirror sessions to perform maintenance on the system.

Suspending metro mirroring with PowerHA graphical interface

To suspend a Metro Mirror session by using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, select **Details...** from the context menu of the independent disk pool for which you want to suspend Metro Mirror.
6. On the **Independent ASP Details** page, select **Suspend Mirroring** from the **Select Action** menu below the mirroring status.
7. Click **Yes** on the confirmation panel.

Suspending metro mirroring with command-line interface

To suspend a Metro Mirror session by using the command-line interface, use the following command:

- CHGASPSSN SESSION(session-name) OPTION(*SUSPEND) where session-name is the name of your DS8000 Metro Mirror session.
- CHGSVCSSN SESSION(session-name) OPTION(*SUSPEND) where session-name is the name of your SAN Volume Controller(SVC) Metro Mirror session.

Related information

[Change ASP Session \(CHGASPSSN\) command](#)

[Change SVC Session \(CHGSVCSSN\) command](#)

Resuming Metro Mirror sessions

After you have completed routine operations, such as performing maintenance on your system, you need to resume a suspended Metro Mirror session to re-enable high availability.

Resuming metro mirroring with PowerHA graphical interface

To resume a Metro Mirror session by using the PowerHA graphical interface, follow these steps:

1. In a web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** in the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, select **Details...** from the context menu of the independent disk pool for which you want to resume geographic mirroring.
6. On the **Independent ASP Details** page, select **Resume Mirroring** from the **Select Action** menu below the mirroring status.
7. Click **Yes** on the confirmation panel.

Resuming metro mirroring with command-line interface

To resume a Metro Mirror session by using the command-line interface, issue the following command:

- CHGASPSSN SESSION(session-name) OPTION(*RESUME) where session-name is the name of your DS8000 Metro Mirror session.
- CHGSVCSSN SESSION(session-name) OPTION(*RESUME) where session-name is the name of your SAN Volume Controller (SVC) Metro Mirror session.

Related information

[Change ASP Session \(CHGASPSSN\) command](#)

[Change SVC Session \(CHGSVCSSN\) command](#)

Detaching Metro Mirror copy

If you are using Metro Mirror and want to access the target copy to perform save operations or data mining, you must detach the target copy from the source copy.

Detaching the mirror copy with PowerHA graphical interface

To use the PowerHA graphical interface to detach the mirror copy, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, select **Details...** from the context menu of the independent disk pool for which you want to detach the mirror copy.
6. On the **Independent ASP Details** page, select **Detach** from the context menu of the mirror copy.
7. Click **Yes** on the confirmation panel.

Detaching the mirror copy with command-line interface

To detach the mirror copy using the command-line interface, issue the following command:

- CHGASPSSN SESSION(session-name) OPTION(*DETACH) where session-name is the name of your DS8000 Metro Mirror session.
- CHGSVCSSN SESSION(session-name) OPTION(*DETACH) where session-name is the name of your SAN Volume Controller (SVC) Metro Mirror session.

Reattaching Metro Mirror

If you detached the target copy and have completed your work with the detached target copy, you must reattach the detached target copy to resume using Metro Mirroring.

The detached target copy must be unavailable when you reattach it to the source copy.

Note: When you reattach the detached mirror copy, as of V6R1, there is the option to detach with tracking, which only requires a partial synchronization on reattach. Changes that are made on the source copy will be applied to the target copy, and all changes that are made on the detached target copy will be overwritten with the contents from the source copy.

Reattaching the target copy with PowerHA graphical interface

To use the PowerHA graphical interface to reattach the mirror copy, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, select **Details...** from the context menu of the independent disk pool for which you want to reattach the mirror copy.
6. On the **Independent ASP Details** page, select **Reattach** from the context menu of the mirror copy.
7. Click **Yes** on the confirmation panel.

Reattaching the target copy with command-line interface

To reattach metro mirroring using the command-line interface, use the following command:

- CHGASPSSN SESSION(session-name) OPTION(*REATTACH) where session-name is the name of your DS8000 Metro Mirror session.
- CHGSVCSSN SESSION(session-name) OPTION(*REATTACH) where session-name is the name of your SAN Volume Controller (SVC) Metro Mirror session.

Ending Metro Mirror

You can end the Metro Mirror session to no longer use the session for high availability and disaster recovery.

Ending the Metro Mirror session does not deconfigure metro mirroring on the IBM System Storage, it just removes our connection to the IBM System Storage.

Ending metro mirroring with PowerHA graphical interface

To end a Metro Mirror session by using the PowerHA graphical interface, follow these steps:

1. In a web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** in the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, select **Details...** from the context menu of the independent disk pool for which you want to delete the Metro Mirror session.
6. On the **Independent ASP Details** page, select **Deconfigure Mirroring...** from the **Select Action** menu below the mirroring status.
7. Click **OK** on the confirmation panel.


Ending metro mirroring with command-line interface

To end a Metro Mirror session by using the command-line interface, use the following command:

- `ENDASPSSN SESSION(session-name)` where `session-name` is the name of your DS8000 Metro Mirror session.
- `ENDSVCSSN SESSION(session-name)` where `session-name` is the name of your SAN Volume Controller (SVC) Metro Mirror session.

Deleting metro mirroring configuration

Ending the Metro Mirror session does not deconfigure metro mirroring on the IBM System Storage, it just removes the connection to the IBM System Storage.

To deconfigure Metro Mirror on the IBM i, you should deconfigure the IBM System Storage external storage units. For information about Metro Mirror on IBM System Storage DS8000, see [Copy Services in the IBM System Storage DS8000 Information Center](#) .

Related information

[End ASP Session \(ENDASPSSN\) command](#)

[End SVC Session \(ENDSVCSSN\) command](#)

Displaying or changing Metro Mirror properties

Display information about a metro mirroring session to change the associated copy descriptions.

Displaying or changing metro mirroring properties with PowerHA licensed program is installed

To change metro mirroring properties by using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** in the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, select **Details...** from the context menu of the independent disk pool for which you want to change the metro mirroring properties.
6. On the **Independent ASP Details** page, select **Properties...** from the context menu of either the production or mirror copy.
7. On the **Properties** page, click **Edit** in the Advanced section.
8. Make changes to the properties and click **Save**.

Managing Global Mirror

In IBM i high availability environment that use IBM System Storage Global Mirror technology, you must configure a global mirroring session between the IBM i systems and the external disk units with Global Mirror configured. From the system, you can manage these sessions.

Suspending Global Mirror

You might need to suspend Global Mirror sessions to perform maintenance on the system.

Suspending Global Mirror with PowerHA graphical interface

To suspend a Global Mirror session by using the PowerHA graphical interface, follow these steps:

1. In a web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** in the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, select **Details...** from the context menu of the independent disk pool for which you want to suspend Global Mirror.
6. On the **Independent ASP Details** page, select **Suspend Mirroring** from the **Select Action** menu below the mirroring status.

7. Click **Yes** on the confirmation panel.

Suspending Global Mirror with command-line interface

To suspend a Global Mirror session by using the command-line interface, use the following command:

- CHGASPSSN SESSION(session-name) OPTION(*SUSPEND) where session-name is the name of your DS8000 Global Mirror session.
- CHGSVCSSN SESSION(session-name) OPTION(*SUSPEND) where session-name is the name of your SAN Volume Controller (SVC) Global Mirror session.

Resuming Global Mirror

After you have completed routine operations, such as performing maintenance on your system, you need to resume a suspended Global Mirror session to re-enable high availability.

Resuming Global Mirror with PowerHA graphical interface

To resume a suspended global mirroring session by using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, select **Details...** from the context menu of the independent disk pool for which you want to resume Global Mirror.
6. On the **Independent ASP Details** page, select **Resume Mirroring** from the **Select Action** menu below the mirroring status.
7. Click **Yes** on the confirmation panel.

Resuming Global Mirror with command-line interface

To resume a suspended global mirroring session using the command-line interface, use this command:

- CHGASPSSN SESSION(session-name) OPTION(*RESUME) where session-name is the name of your DS8000 Global Mirror session.
- CHGSVCSSN SESSION(session-name) OPTION(*RESUME) where session-name is the name of your SAN Volume Controller(SVC) Global Mirror session.

Detaching Global Mirror copy

If you are using Global Mirror and want to access the target copy to perform save operations or data mining, you must detach the target copy from the source copy.

Detaching the mirror copy with PowerHA graphical interface

To use the PowerHA graphical interface to detach the mirror copy, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, select **Details...** from the context menu of the independent disk pool for which you want to detach the mirror copy.
6. On the **Independent ASP Details** page, select **Detach** from the context menu of the mirror copy.
7. Click **Yes** on the confirmation panel.

Detaching the mirror copy with command-line interface

To detach the mirror copy using the command-line interface, use the following command:

- CHGASPSSN SESSION(session-name) OPTION(*DETACH) where session-name is the name of your DS8000 Global Mirror session.

- CHGSVCSSN SESSION(session-name) OPTION(*DETACH) where session-name is the name of your SAN Volume Controller (SVC) Global Mirror session.

Reattaching Global Mirror

If you detached the target copy and have completed your work with the detached target copy, you must reattach the detached target copy to resume using Global Mirror.

The detached target copy must be unavailable when you reattach it to the source copy.

Note: When you reattach the detached mirror copy, as of V6R1, there is the option to detach with tracking, which only requires a partial synchronization on reattach. Changes that are made on the source copy will be applied to the target copy, and all changes that are made on the detached target copy will be overwritten with the contents from the source copy.

Reattaching the target copy when PowerHA graphical interface

To use the PowerHA graphical interface to reattach the mirror copy, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, select **Details...** from the context menu of the independent disk pool for which you want to reattach the mirror copy.
6. On the **Independent ASP Details** page, select **Reattach** from the context menu of the mirror copy.
7. Click **Yes** on the confirmation panel.

Reattaching the target copy with command-line interface

To reattach mirror copy using the command-line interface, use the following command:

- CHGASPSSN SESSION(session-name) OPTION(*REATTACH) where session-name is the name of your DS8000 Global Mirror session.
- CHGSVCSSN SESSION(session-name) OPTION(*REATTACH) where session-name is the name of your SAN Volume Controller (SVC) Global Mirror session.

Ending Global Mirror

You can end the Global Mirror session to no longer use the session for high availability and disaster recovery.

Ending global mirroring with PowerHA graphical interface

To delete a global mirroring session by using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, select **Details...** from the context menu of the independent disk pool for which you want to delete the Global Mirror session.
6. On the **Independent ASP Details** page, select **Deconfigure Mirroring...** from the **Select Action** menu below the mirroring status.
7. Click **OK** on the confirmation panel.

Ending global mirroring with the command-line interface


To end a Global Mirror session using the command-line interface, use this command:

- ENDASPSSN SESSION(session-name) where session-name is the name of your DS8000 Metro Mirror session.

- ENDSVCSSN SESSION(session-name) where session-name is the name of your SAN Volume Controller (SVC) Metro Mirror session.

Deleting Global Mirror configuration

Ending the Global Mirror session does not deconfigure global mirroring on the IBM System Storage, it just removes the connection to the IBM System Storage.

To deconfigure Global Mirror on the IBM i, you should deconfigure the IBM System Storage external storage units. For information about Global Mirror on IBM System Storage DS8000, see [Copy Services in the IBM System Storage DS8000 Information Center](#) .

Changing Global Mirror session properties

Display information about a Global Mirror session to change the associated copy descriptions.

Displaying or changing Global Mirror properties with PowerHA graphical interface

To change Global Mirror properties by using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** in the IBM Navigator for i window
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, select **Details...** from the context menu of the independent disk pool for which you want to change Global Mirror properties.
6. On the **Independent ASP Details** page, select **Properties...** from the context menu of either the production or mirror copy.
7. On the **Properties** page, click **Edit** in the Advanced section.
8. Make changes to the properties and click **Save**.

Managing switched logical units (LUNs)

Switched logical units are independent disk pools created from logical units created in an IBM System Storage that have been configured as part of a device cluster resource group (CRG).

Ownership of data and applications that stored in a switched logical unit can be switched to other systems that have been defined in the device CRG. Switched disk technology provides high availability during planned and some unplanned outages.

Related concepts

[PowerHA supported storage servers](#)

IBM System Storage provides enhanced storage capabilities.

Making switched logical units (LUNs) available and unavailable

You can select an independent disk pool to make it unavailable or available. You cannot access any of the disk units or objects in the independent disk pool or its corresponding database until it is made available again. The pool can be made available again on the same system or another system in the recovery domain of the cluster resource group.

An independent disk pool can be made unavailable by varying it off. Access to any of the disk units or objects in the independent disk pool or its corresponding database are not available, until they are varied on. The pool can be made available on the same system or another system in the recovery domain of the cluster resource group.

Managing the FlashCopy technology

FlashCopy is an IBM System Storage technology that allows you to take a point-in-time copy of external disk units. In PowerHA solutions that use Metro Mirror or Global Mirror, The FlashCopy technology can be used for backup window reduction by taking a copy of data that then can be backed up to media. To use the FlashCopy technology, a session must be created between the system and the external storage units.

Updating a FlashCopy session

You can update a FlashCopy session when you are performing resynchronization of FlashCopy volumes on your IBM System Storage external storage units. Resynchronization allows you to make a copy without recopying the entire volume. This process is only possible with a persistent relationship, whereby the storage unit continually tracks updates to the source and target volumes. With persistent relationships, the relationship between the source and target volumes is maintained after the background copy has completed. The FlashCopy session that is created on the IBM i provides a means to manage and monitor activity that is related to the FlashCopy session on the IBM System Storage units.

Updating a FlashCopy when IBM PowerHA SystemMirror for i licensed program is installed

To update a FlashCopy session by using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** in the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, select **Details...** from the context menu of the independent disk pool that has a FlashCopy you want to update.
6. On the **Independent ASP Details** page, select **Update** from the context menu of the FlashCopy you want to update.
7. Click **Yes** on the confirmation panel.

Reattaching a FlashCopy session

Reattach a FlashCopy session.

Reattaching a FlashCopy when IBM PowerHA SystemMirror for i licensed program is installed

To reattach (enable) a FlashCopy session by using the IBM PowerHA for i graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** in the IBM Navigator for i window
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, select **Details...** from the context menu of the independent disk pool for that has a FlashCopy you want to reattach (enable).
6. On the **Independent ASP Details** page, select **Enable** from the context menu of the FlashCopy you want to reattach (enable).

Reattaching a FlashCopy when IBM PowerHA SystemMirror for i licensed program is not installed

To reattach a FlashCopy session by using the Configuration and Service graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Select **Configuration and Service** from your IBM Navigator for i window.
4. Select **Disk Pools**.
5. Select the disk pool that is associated with the session that you want to reattach.
6. From the **Select Actions** menu, select **Sessions**.
7. Select the session that you want to reattach.
8. From the **Select Actions** menu, select **Reattach FlashCopy**.

Detaching a FlashCopy session

You can detach the target volumes from the source for a selected FlashCopy session.

Detaching FlashCopy when IBM PowerHA for i is installed

To detach (disable) a FlashCopy session by using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** in the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, select **Details...** from the context menu of the independent disk pool that has a FlashCopy you want to detach (disable).
6. On the **Independent ASP Details** page, select **Disable** from the context menu of the FlashCopy you want to detach (disable).
7. Click **Yes** on the confirmation panel.

Deleting a FlashCopy session

Delete a FlashCopy session.

Deleting a FlashCopy when IBM PowerHA SystemMirror for i licensed program is installed

To delete a FlashCopy session by using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** in the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, select **Details...** from the context menu of the independent disk pool that has a FlashCopy you want to delete.
6. On the **Independent ASP Details** page, select **Delete** from the context menu of the FlashCopy you want to delete.
7. Click **Yes** on the confirmation panel.

Restoring data from a FlashCopy session

After a FlashCopy session has been completed on the IBM System Storage units, you can restore that data from target volume to the source volume in the event of an outage at the source copy of data. To do this, you need to reverse the FlashCopy session that is created on IBM i. However, reversing the session copies data from the target back to the source and returns the source to it an earlier version.

If the FlashCopy source is a member of a cluster resource group, then reverse FlashCopy are not allowed unless the original FlashCopy was performed as a *NOCOPY flash. Full-copy reverse flashes are not supported when the original FlashCopy source is a member of a cluster resource group.



Attention: Reversing a FlashCopy session backs out the changes that are made on the source copy by copying the target's data back to the source. This returns the source to that earlier point in time.

To reverse a FlashCopy session, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Select **Configuration and Service** from your IBM Navigator for i window.
4. Select **Disk Pools**.
5. Select the disk pool of the source copy.
6. From the **Select Actions** menu, select **Open Sessions**.

7. Select the session.
8. From the **Select Actions** menu, select **Reverse FlashCopy**.

Changing FlashCopy properties

Display information about a FlashCopy session to change the associated copy descriptions.

Changing FlashCopy properties when IBM PowerHA SystemMirror for i licensed program is installed

To change information about a FlashCopy session by using the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** in the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, select **Details...** from the context menu of the independent disk pool for which you want to change FlashCopy properties.
6. On the **Independent ASP Details** page, select **Properties...** from the context menu FlashCopy.
7. On the **Properties** page, click **Edit**.
8. Make changes to the properties and click **Save**.

Managing DS8000 Full System HyperSwap

DS8000 Full System HyperSwap is a full system solution that allows logical units that are mirrored between two IBM System Storage DS8000 units to be switched between without an outage, providing a minimal impact high availability solution.

To use DS8000 Full System HyperSwap requires IBM PowerHA for i Express Edition be installed, with a valid license key.

Displaying HyperSwap Status

To display the current HyperSwap Status for the system, use the DSPHYSSTS command. The output from the command indicates the current direction and status of HyperSwap replication.

Related information

[Display HyperSwap Status \(DSPHYSSTS\) command](#)

Perform a Planned HyperSwap

To perform a planned HyperSwap from the primary IBM System Storage unit to the secondary, and start replication in the reverse direction, use the Change HyperSwap Status (CHGHYSSTS) command with the *SWAP option.

Related information

[Change HyperSwap Status \(CHGHYSSTS\) command](#)

Suspend HyperSwap Replication

To suspend HyperSwap replication for the system, use the Change HyperSwap Status (CHGHYSSTS) command with the *STOP option.

Note: When replication is suspended, a HyperSwap cannot be performed. Hourly messages are posted to the QSYSOPR message queue.

Related information

[Change HyperSwap Status \(CHGHYSSTS\) command](#)

Resume HyperSwap Replication

To resume HyperSwap replication that was suspended, use the Change HyperSwap Status (CHGHYSSTS) command with the *START option.

Note: The system is able to perform a HyperSwap as soon as the resynchronization has completed.

Related information

[Change HyperSwap Status \(CHGHYSSTS\) command](#)

Recovering from an unplanned HyperSwap Failover

In the event of an unplanned HyperSwap failover, the system will automatically switch to use the secondary IBM System Storage device as the primary. When the issue that caused the HyperSwap failover has been resolved, the system indicates that there are two primary IBM System Storage devices. The replication direction is indicated on the Display HyperSwap Status (DSPHYSSTS) command output.

To recover perform the following steps:

1. Issue a Change HyperSwap Status (CHGHYSSTS) command with the *START option to resume replication.
2. Display the HyperSwap Status (DSPHYSSTS) command waiting until all units are fully synchronized.
3. To switch back to the original primary, issue a Change HyperSwap Status (CHGHYSSTS) command with the *SWAP option.

Related information

[Change HyperSwap Status \(CHGHYSSTS\) command](#)

[Display HyperSwap Status \(DSPHYSSTS\) command](#)

Managing DS8000 HyperSwap with independent auxiliary storage pools (IASPs)

DS8000 HyperSwap with IASPs is typically used with PowerHA LUN switching technology to provide coverage for planned and unplanned storage and server outages, and can also be used with live partition mobility to further minimize downtime for planned server outages.

The following table contains the recommended technology to use for several outages and scenarios, both planned and unplanned.

Outage type	Recovery type
Planned IBM i OS upgrade or maintenance	CRG switchover. For more information, see “CRG Planned switchover with HyperSwap and LUN switching” on page 136.
Planned server firmware upgrade or outage	LPM switch or CRG switchover. For more information, see “CRG Planned switchover with HyperSwap and LUN switching” on page 136 or “Live partition mobility switch with HyperSwap and affinity” on page 139.
Planned storage upgrade or outage	HyperSwap switchover. For more information, see “HyperSwap planned switchover of SYSBAS” on page 135, “HyperSwap planned switchover of independent auxiliary storage pool (IASP)” on page 135, or “HyperSwap planned switchover with HyperSwap and LUN switching” on page 135.
Planned server hardware upgrade or outage	CRG switchover. For more information, see “CRG Planned switchover with HyperSwap and LUN switching” on page 136.
Unplanned IBM i OS outage	CRG failover. For more information, see “CRG Failover HyperSwap and LUN switching” on page 138.
Unplanned server firmware outage	CRG failover. For more information, see “CRG Failover HyperSwap and LUN switching” on page 138.

Table 39. Outage scenarios (continued)

Outage type	Recovery type
Unplanned storage outage	HyperSwap failover. For more information, see “HyperSwap failover with HyperSwap and LUN switching” on page 137.
Unplanned server hardware outage	CRG failover. For more information, see “CRG Failover HyperSwap and LUN switching” on page 138.

Displaying HyperSwap Status

To display the current HyperSwap Status for the system, use the DSPHYSSTS command. The output from the command indicates the current direction and status of HyperSwap replication.

Related information

[Display HyperSwap Status \(DSPHYSSTS\) command](#)

HyperSwap planned switchover of SYSBAS

To do a planned switchover of SYSBAS, enter the following command:

```
CHGHYSSTS OPTION(*SWAP) NODE(*) ASPDEV(*SYSBAS)
```

This command switches primary access of the SYSBAS from one DS8000 storage server to the other DS8000 storage server in the HyperSwap relationship. If an IASP exists on the system, its HyperSwap relationship stays the same.

HyperSwap planned switchover of independent auxiliary storage pool (IASP)

To do a planned switchover of an IASP, enter the following command:

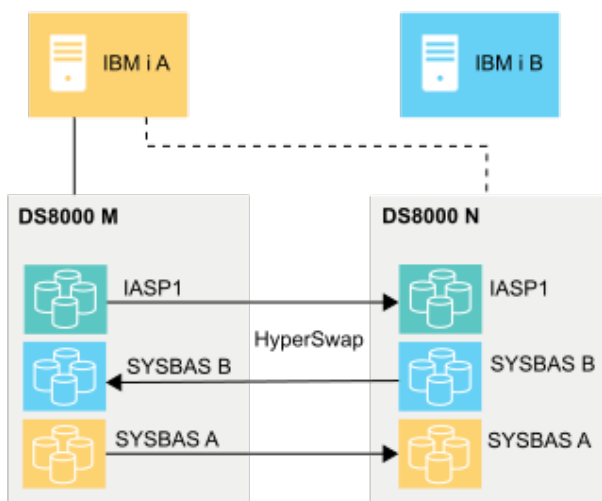
```
CHGHYSSTS OPTION(*SWAP) NODE(*) ASPDEV(<iasp-name>)
```

This command switches primary access of the asp-name IASP from one DS8000 storage server to the other DS8000 storage server in the HyperSwap relationship. It does not change the relationship for SYSBAS or any other IASPs, which exist on the system.

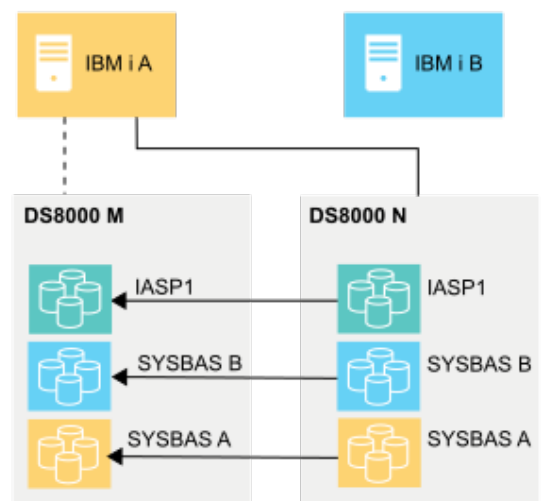
HyperSwap planned switchover with HyperSwap and LUN switching

In the case of a planned outage of a storage server, a HyperSwap switchover can be initiated, resulting in near-zero downtime from a user perspective. In the following figure on the left, all disk units in IASP1 are metro mirrored to the second DS8000 to create a HyperSwap relationship. In addition to being configured for HyperSwap, IASP1 is also configured for LUN level switching to a second IBM i partition, named IBM i B in the figure.

1 Normal Operation



2 HyperSwap Switchover of IASP1 to secondary DS8000 Near-zero downtime



For the HyperSwap switchover to be initiated, the following conditions must be met.

1. All disk units in the ASP group must be configured in a HyperSwap relationship.
2. All of the HyperSwap relationships in the ASP group must be fully synchronized.
3. All of the HyperSwap relationships in the ASP group must be going the same direction (some disk units in the ASP group cannot have one DS8000 as the primary while other disk units have the other DS8000 as primary).

When the user wants to initiate a planned HyperSwap switchover of IASP1, the following command must be entered:

```
CHGHYSSTS OPTION(*SWAP) NODE(*) ASPDEV(*ALL)
```

After you enter the command, the primary access for IASP1 and SYSBAS will be the second DS8000, as shown in the figure on the right side.

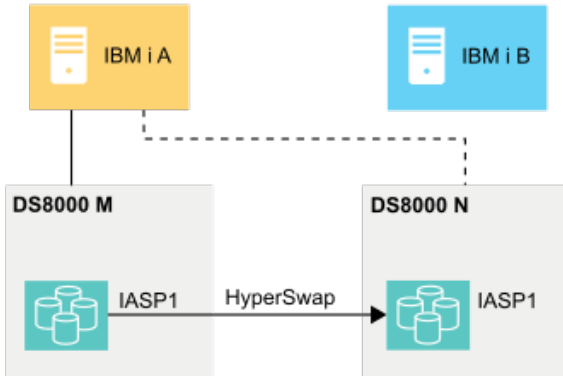
CRG Planned switchover with HyperSwap and LUN switching

In the case of a planned outage of the system server, a CRG switchover can be initiated.

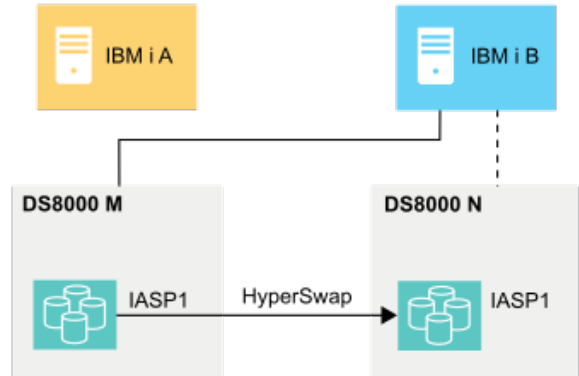
The following conditions must be met before the switchover can be initiated.

1. Both DS8000s must be accessible by PowerHA through DSCLI.
2. None of the HyperSwap relationships in the IASP can have two primary volumes. This is shown by XXXX for the copy status when a Display HyperSwap Status or Work with HyperSwap Status command is run for the IASP.

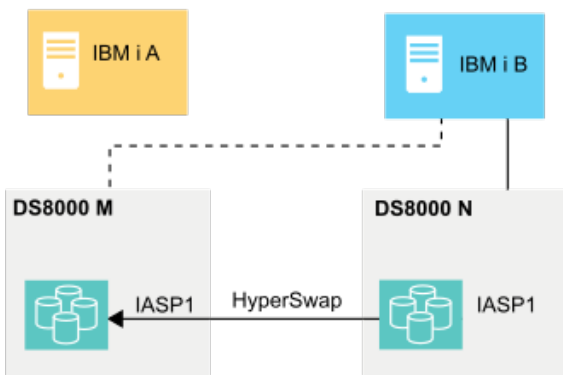
1 Normal Operation



2 CRG Switchover (LUN Switch) from IBM i A to IBM i B



3 PowerHA automatically performs a HyperSwap switchover to correct affinity



The following command must be entered to perform the CRG switchover.

```
CHGCRGPRI CLUSTER(<cluster-name>) CRG(<CRG-name>)
```

The **2** in the figure shows the configuration after IASP1 is varied on for IBM i B. If affinity was defined by the ADDHACFGD command and the affinity is not correct after the CRG switchover, then PowerHA will also initiate a HyperSwap switchover of IASP1, shown by **3** in the figure.

When a HyperSwap switchover cannot be completed, the CRG switchover completes and a diagnostic message is sent indicating that the HyperSwap switchover failed. Examples of why the HyperSwap switchover could not complete are HyperSwap is being suspended, or is not fully synchronized.

HyperSwap failover with HyperSwap and LUN switching

In the case of an unplanned outage of a storage server, a HyperSwap failover is initiated, resulting in near-zero downtime from a user perspective. The figure on the left shows a HyperSwap relationship between all of the disk units in SYSBAS A and SYSBAS B, and in IASP1. The primary access for SYSBAS A and IASP1 is DS8000 M, and the primary access for SYSBAS B is DS8000 N. In addition to being configured for HyperSwap, IASP1 is also configured for LUN level switching to a second IBM i partition, named IBM i B in the figure.

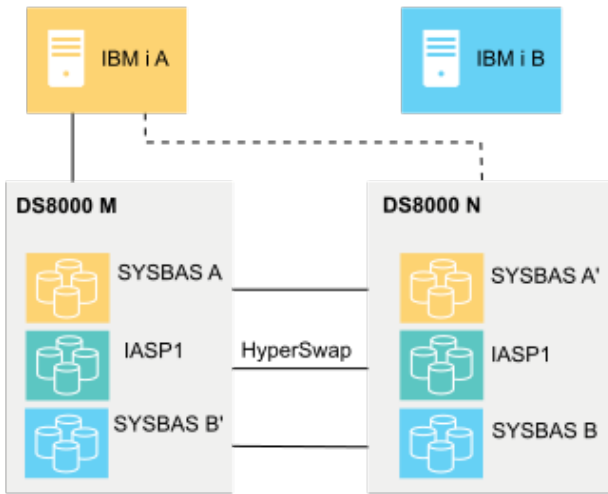
In order for the HyperSwap failover to be initiated, the following conditions must be met.

1. All disk units in the ASP group must be configured in a HyperSwap relationship.
2. All of the HyperSwap relationships in the ASP group must be fully synchronized.


rzaij0011-1

- All of the HyperSwap relationships in the ASP group must be going the same direction (some disk units in the ASP group can't have one DS8000 as the primary while other disk units have the other DS8000 as primary).

1 Normal Operation



2 Unplanned HyperSwap failover to secondary DS8000 Minimal downtime



After the HyperSwap failover completes, the primary access for IASP1 and SYSBAS will be the second DS8000, as shown in the figure on the right side. If the HyperSwap failover cannot be completed, the IASP goes into DASD attention.

After the failover occurs and the failed DS8000 comes back online, the HyperSwap relationship must be restarted by entering the following command.

```
CHGHYSSTS OPTION(*START) NODE(*) ASPDEV(*ALL)
```

This will restart replication for all ASP devices. The Display HyperSwap Status and Work with HyperSwap Status commands can then be used to monitor the resynchronization process. After resynchronization is completed, a HyperSwap switchover can be initiated if desired.

Two primary case

After a HyperSwap failover and until the HyperSwap relationship is restarted, each DS8000s reports that it is the primary. If the mirroring relationship is not restarted before the IBM i partition is IPLed, the IPL will fail, requiring the user to manually correct the condition from the DS8000.

The two primary case can be determined in one of the following ways.

- By entering the **DSPHYSSTS** command, the HyperSwap copy status is shown as XXXX.
- Hourly messages are posted to QSYSOPR indicating there is a problem with the HyperSwap relationship.

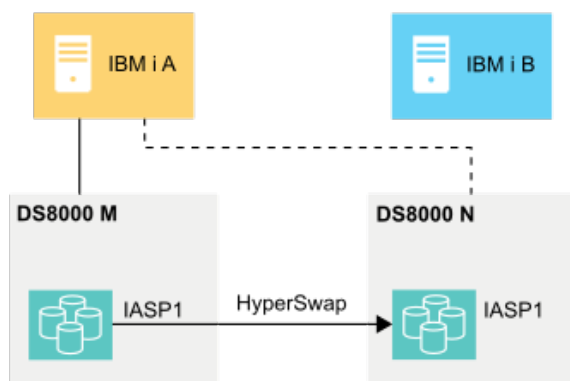
CRG Failover HyperSwap and LUN switching

In the case of a planned outage of the system server, a CRG failover can be initiated.

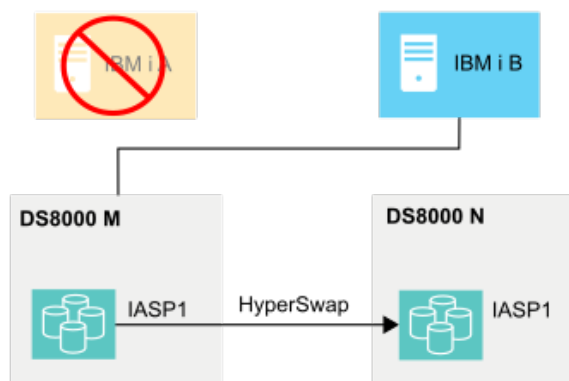
The following conditions must be met before the failover is initiated.

- Both DS8000s must be accessible by PowerHA through DSCLI.
- None of the HyperSwap relationships in the IASP can have two primary volumes. This is shown by **XXXX** for the copy status when a Display HyperSwap Status or Work with HyperSwap Status command is run for the IASP.

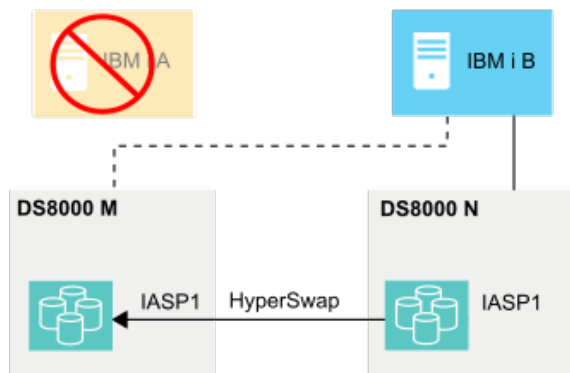
1 Normal Operation



2 IBM i A fails, CRG Failover (LUN Switch) from IBM i A to IBM i B



3 PowerHA automatically performs a HyperSwap switchover to correct affinity



rzaig0012-1

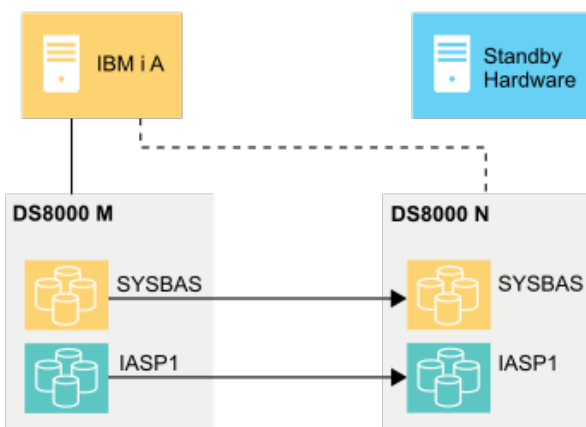
In the figure, first the CRG failover takes place, resulting in **2**, where IASP1 is now varied on to IBM i B. If affinity was defined by the **ADDHACFGD** command and the affinity is not correct after the CRG failover, then PowerHA will also initiate a HyperSwap switchover of IASP1, shown in **3**.

If a HyperSwap switchover cannot be completed due to some reason such as HyperSwap being suspended or not fully synchronized, the CRG failover completes and a diagnostic message is sent indicating that the HyperSwap switchover failed.

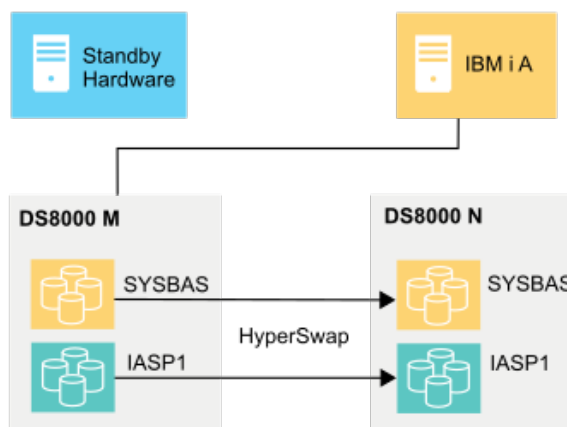
Live partition mobility switch with HyperSwap and affinity

Live partition mobility (LPM) allows a running IBM i partition to migrate from one physical server to another server with near-zero downtime. HyperSwap can be combined with live partition mobility to move the data access to the DS8000 storage server with the most affinity to the physical server hosting the IBM i partition. LPM switch with HyperSwap and affinity is illustrated in the figure.

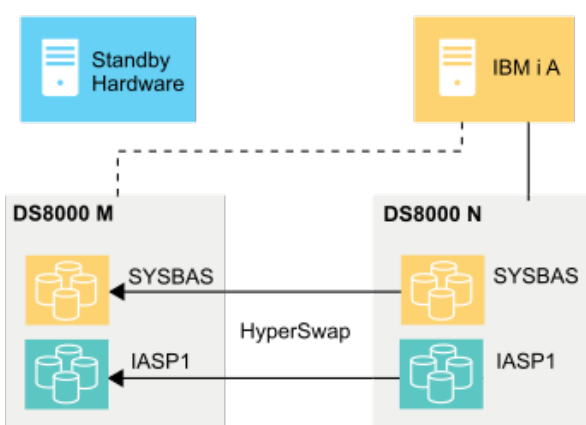
1 Normal Operation



2 After LPM, there is an I/O performance penalty



3 Automatic HyperSwap switchover to eliminate the penalty



When HyperSwap affinity is configured by using the Add HA Configuration Description (**ADDHACFGD**) command, the HyperSwap switchover can automatically occur when a live partition mobility switch occurs and the current primary DS8000 does not have affinity to the new server.

DS8000 HyperSwap IASP with FlashCopy

In addition to integrating HyperSwap with PowerHA logical unit (LUN) level switching, HyperSwap can also be integrated with the PowerHA FlashCopy technology.

The following rules apply to a configuration that has both HyperSwap and FlashCopy

1. A FlashCopy source can be a HyperSwap IASP (as identified by two storage hosts and LUN ranges that are defined in the source copy description).
2. A FlashCopy target cannot be a HyperSwap IASP. It might have only one storage host and LUN range defined in the target copy description.

Depending on which storage server is currently the primary DS8000 in the HyperSwap relationship, the FlashCopy target might be in the secondary DS8000. In this case, the FlashCopy target gets the data, which is in the target DS8000. If the HyperSwap relationship is not active, then the FlashCopy is stale data as well.

rzsig0008-1

Troubleshooting your high availability solution

After you have configured your IBM i high-availability solution, you may encounter problems with different technologies, including clusters and cross-site mirroring.

Troubleshooting clusters

Find error recovery solutions for problems that are specific to clusters.

At times, it may appear that the cluster is not working properly. This topic covers information about problems that you may encounter with clusters.

Determine if a cluster problem exists

Start here to diagnose your cluster problems.

At times, it may seem that your cluster is not operating correctly. When you think a problem exists, you can use the following to help determine if a problem exists and the nature of the problem.

- **Use the PowerHA graphical interface to examine the cluster.**

To use the PowerHA graphical interface, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** in the IBM Navigator for i window
4. Verify if the status icons indicate a problem and investigate further if needed.

- **Determine if clustering is active on your system.**

To determine if cluster resource services are active, look for the two jobs - QCSTCTL and QCSTCRGM - in the list of system jobs. If these jobs are active, then cluster resource services are active. You can use the Work Management function in IBM Navigator for i or in IBM Navigator for i to View jobs or use the **WRKACTJOB (Work with Active Jobs)** command to do this. You can also use the **DSPCLUINF (Display Cluster Information)** command to view status information for the cluster.

– Additional jobs for cluster resource services may also be active. [Cluster jobs](#) provides information about how cluster resource services jobs are formatted.

- **Determine the cause of a CPFBB26 message.**

```
Message . . . . : Cluster Resource Services not active or not responding.  
Cause . . . . . : Cluster Resource Services is either not active or cannot  
respond to this request because a resource is unavailable or damaged.
```

This error can mean that either the CRG job is not active or the cluster is not active. Use the **DSPCLUINF (Display Cluster Information)** command to determine if the node is active. If the node is not active, start the cluster node. If it is active, you should also check the CRG to determine whether the CRG has problems.

Look for the CRG job in the list of system jobs. You can use the Work Management function in IBM Navigator for i or in IBM Navigator for i to View jobs or use the **WRKACTJOB (Work with Active Jobs)** command to do this. You can also use the **DSPCRGINF (Display CRG Information)** command to view status information for the specific CRG, by specifying the CRG name in the command. If the CRG job is not active, look for the CRG job log to determine the cause of why it was ended. Once the problem is fixed, you could restart the CRG job with **CHGCLURCY (Change Cluster Recovery) command** or by ending and restarting cluster on that node.

- **Look for messages indicating a problem.**

– Ensure that you can review all messages associated with a cluster command, by selecting F10, which toggles between "Include detailed messages" and "Exclude detailed messages". Select to include all detailed messages and review them to determine if other actions are necessary.

- Look for inquiry messages in QSYSOPR that are waiting for a response.
- Look for error messages in QSYSOPR that indicate a cluster problem. Generally, these are in the CPFBB00 to CPFBBFF range.
- Display the history log (**DSPLOG** CL command) for messages that indicate a cluster problem. Generally, these are in the CPFBB00 to CPFBBFF range.
- **Look at job logs for the cluster jobs for severe errors.**
 These jobs are initially set with a logging level at (4 0 *SECLVL) so that you can see the necessary error messages. You should ensure that these jobs and the exit program jobs have the logging level set appropriately. If clustering is not active, you can still look for spool files for the cluster jobs and exit program jobs.
- **If you suspect some kind of hang condition, look at call stacks of cluster jobs.**
 Determine if there is any program in some kind of DEQW (dequeue wait). If so, check the call stack of each thread and see if any of them have getSpecialMsg in the call stack.
- **Check for cluster vertical Licensed Internal Code (VLIC) logs entries.**
 These log entries have a 4800 major code.
- **Use NETSTAT command to determine if there are any abnormalities in your communications environment.**
NETSTAT returns information about the status of Internet Protocol network routes, interfaces, TCP connections, and UDP ports on your system.
 - Use Netstat Option 1 (Work with TCP/IP interface status) to ensure that the IP addresses chosen to be used for clustering show an 'Active' status. Also ensure that the LOOPBACK address (127.0.0.1) is also active.
 - Use **NETSTAT** Option 3 (Work with TCP/IP Connection Status) to display the port numbers (F14). Local port 5550 should be in a 'Listen' state. This port must be opened using the **STRTCPSVR *INETD** command evidenced by the existence of a QTOGINTD (User QTCP) job in the Active Jobs list. If clustering is started on a node, local port 5551 must be opened and be in a '*UDP' state. If clustering is not started, port 5551 must not be opened or it will, in fact, prevent the successful start of clustering on the subject node.
- Use **PING** to verify if there is a communications problem. If you try to start a cluster node and there is a communications problem, you may receive an internal clustering error (CPFBB46). However, **PING** does not work between IPv4 and IPv6 addresses, or if a firewall is blocking it.

Gathering recovery information for a cluster

You can use the **Work with Cluster (WRKCLU)** command to collect information for a complete picture of your cluster. This information can be used to aid in error resolution.

The **Work with Cluster (WRKCLU)** command is used to display and to work with cluster nodes and objects. When you run this command, the Work with Cluster display is shown. In addition to displaying nodes in a cluster and cluster information, you can use this command to view cluster information and to gather data about your cluster

To gather error recovery information, complete these steps:

1. On a character-based interface, type **WRKCLU OPTION(OPTION)**. You can specify the following options to indicate with which cluster status information you want to work with.
 - *SELECT**
 Display the Work with Cluster menu.
 - *CLUINF**
 Display cluster information.
 - *CFG**
 Display the performance and configuration parameters for the cluster.
 - *NODE**
 Display the Work with Cluster Nodes panel which is a list of nodes in the cluster.

***DEVDMN**

Display the Work with Device Domains panel which is a list of device domains in the cluster.

***CRG**

Display the Work with Cluster Resource Groups panel which is a list of cluster resource groups in the cluster.

***ADMDMN**

Display the Work with Administrative Domains panel which is a list of administrative domains in the cluster.

***SERVICE**

Gathers related trace and debug information for all cluster resource service jobs in the cluster. This information is written to a file with a member for each cluster resource service job. Use this option only when directed by your service provider. It will display a prompt panel for the **Dump Cluster Trace (DMPCLUTRC)**.

Common cluster problems

Lists some of the most common problems that can occur in a cluster, as well as ways to avoid and recover from them.

The following common problems are easily avoidable or easily correctable.

Use the PowerHA graphical interface to identify current or potential problems.

To use the PowerHA graphical interface to identify a current problem, follow these steps:

1. In a Web browser, enter *http://mysystem:2001*, where *mysystem* is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click on **PowerHA** in the IBM Navigator for iwindow.
4. On the **PowerHA** page, verify if any status icons indicate a problem and investigate further if needed.
- 5.

To use the PowerHA graphical interface to identify a potential problem, follow these steps:

1. In a Web browser, enter *http://mysystem:2001*, where *mysystem* is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click on **PowerHA** in the IBM Navigator for iwindow.
4. On the PowerHA, select **Check Requirements...** from the **Select Action** menu next to the cluster name.
5. Fix any problems identified on the **Check Requirements** page. Some problems can be fixed immediately by selecting Fix from the context menu of the identified problem.

You cannot start or restart a cluster node

This situation is typically due to some problem with your communications environment. To avoid this situation, ensure that your network attributes are set correctly, including the loopback address, INETD settings, ALWADDCLU attribute, and the IP addresses for cluster communications.

- The ALWADDCLU network attribute must be appropriately set on the target node if trying to start a remote node. This should be set to either *ANY or *RQSAUT depending on your environment.
- The IP addresses chosen to be used for clustering locally and on the target node must show an *Active* status.
- The LOOPBACK address (127.0.0.1) locally and on the target node must also be active.
- Verify that network routing is active by attempting to PING using the IP addresses used for clustering on the local and remote nodes; however, PING does not work between IPv4 and IPv6 addresses, or if a firewall is blocking it. If any cluster node uses an IPv4 address, than every node in the cluster needs to have an active IPv4 address (not necessarily configured as a Cluster IP address) that can route to and send TCP packets to that address. Also, if any cluster node uses an IPv6 address, than every node in the

cluster needs to have an active IPv6 address (not necessarily configured as a Cluster IP address) that can route to and send TCP packets to that address.

- INETD must be active on the target node. When INETD is active, port 5550 on the target node should be in a *Listen* state. See INETD server for information about starting the INETD server.
- Prior to attempting to start a node, port 5551 on the node to be started must not be opened or it will, in fact, prevent the successful start of clustering on the subject node.

You end up with several, disjointed one-node clusters

This can occur when the node being started cannot communicate with the rest of the cluster nodes. Check the communications paths.

The response from exit programs is slow.

A common cause for this situation is incorrect setting for the job description used by the exit program. The MAXACT parameter may be set too low so that, for example, only one instance of the exit program can be active at any point in time. It is recommended that this be set to *NOMAX.

Performance in general seems to be slow.

There are several common causes for this symptom.

- The most likely cause is heavy communications traffic over a shared communications line.
- Another likely cause is an inconsistency between the communications environment and the cluster message tuning parameters. You can use the [Retrieve Cluster Resource Services Information \(QcstRetrieveCRSInfo\) API](#) to view the current settings of the tuning parameters and the [Change Cluster Resource Services \(QcstChgClusterResourceServices\) API](#) to change the settings. Cluster performance may be degraded under default cluster tuning parameter settings if using old adapter hardware. The adapter hardware types included in the definition of *old* are 2617, 2618, 2619, 2626, and 2665. In this case, setting of the *Performance class* tuning parameter to *Normal* is desired.
- If all the nodes of a cluster are on a local LAN or have routing capabilities which can handle Maximum Transmission Unit (MTU) packet sizes of greater than 1,464 bytes throughout the network routes, large cluster message transfers (greater than 1,536K bytes) can be greatly speeded up by increasing the cluster tuning parameter value for *Message fragment size* to better match the route MTUs.

You cannot use any of the function of the new release.

If you attempt to use new release function and you see error message CPFBB70, then your current cluster version is still set at the prior version level. You must upgrade all cluster nodes to the new release level and then use the adjust cluster version interface to set the current cluster version to the new level. See [Adjust the cluster version of a cluster](#) for more information.

You cannot add a node to a device domain or access the System i Navigator cluster management interface.

To access the System i Navigator cluster management interface, or to use switchable devices, you must have IBM i Option 41, HA Switchable Resources installed on your system. You must also have a valid license key for this option.

You applied a cluster PTF and it does not seem to be working.

You should ensure that you have completed the following tasks after applying the PTF:

1. [End the cluster](#)
2. Signoff then signon

The old program is still active in the activation group until the activation group is destroyed. All of the cluster code (even the cluster APIs) run in the default activation group.

3. Start the cluster

Most cluster PTFs require clustering to be ended and restarted on the node to activate the PTF.

CEE0200 appears in the exit program joblog.

On this error message, the from module is QLEPM and the from procedure is Q_LE_leBdyPeilog. Any program that the exit program invokes must run in either *CALLER or a named activation group. You must change your exit program or the program in error to correct this condition.

CPD000D followed by CPF0001 appears in the cluster resource services joblog.

When you receive this error message, make sure the QMLTTHDACN system value is set to either 1 or 2.

Cluster appears hung.

Make sure cluster resource group exit programs are outstanding. To check the exit program, use the **WRKACTJOB (Work with Active Jobs)** command, then look in the Function column for the presence of PGM-QCSTCRGEXT.

Partition errors

Certain cluster conditions are easily corrected. If a cluster partition has occurred, you can learn how to recover. This topic also tells you how to avoid a cluster partition and gives you an example of how to merge partitions back together.

A cluster partition occurs in a cluster whenever contact is lost between one or more nodes in the cluster and a failure of the lost nodes cannot be confirmed. This is not to be confused with a partition in a logical partition (LPAR) environment.

If you receive error message CPFBB20 in either the history log (QHST) or the QCSTCTL joblog, a cluster partition has occurred and you need to know how to recover. The following example shows a cluster partition that involves a cluster made up of four nodes: A, B, C, and D. The example shows a loss of communication between cluster nodes B and C has occurred, which results in the cluster dividing into two cluster partitions. Before the cluster partition occurred, there were four cluster resource groups, which can be of any type, called CRG A, CRG B, CRG C, and CRG D. The example shows the recovery domain of each cluster resource group.

Table 40. Example of a recovery domain during a cluster partition

Node A	Node B	x	Node C	Node D
CRG A (backup1)	CRG A (primary)			
	CRG B (primary)		CRG B (backup1)	
	CRG C (primary)		CRG C (backup1)	CRG C (backup2)
CRG D (backup2)	CRG D (primary)		CRG D (backup1)	
Partition 1			Partition 2	

A cluster may partition if the maximum transmission unit (MTU) at any point in the communication path is less than the cluster communications tuneable parameter, message fragment size. MTU for a cluster IP address can be verified by using the **Work with TCP/IP Network Status (WRKTCPSTS)** command on the subject node. The MTU must also be verified at each step along the entire communication path. If the MTU is less than the message fragment size, either raise the MTU of the path or lower the message fragment size. You can use the Retrieve Cluster Resource Services Information (QcstRetrieveCRSInfo) API to view the current settings of the tuning parameters and the Change Cluster Resource Services (QcstChgClusterResourceServices) API to change the settings.

Once the cause of the cluster partition condition has been corrected, the cluster will detect the re-established communication link and issue the message CPFBB21 in either the history log (QHST) or the QCSTCTL joblog. This informs the operator that the cluster has recovered from the cluster partition. Be

aware that once the cluster partition condition has been corrected, it may be a few minutes before the cluster merges back together.

Determining primary and secondary cluster partitions

In order to determine the types of cluster resource group actions that you can take within a cluster partition, you need to know whether the partition is a primary or a secondary cluster partition. When a partition is detected, each partition is designated as a primary or secondary partition for each cluster resource group defined in the cluster.

For primary-backup model, the primary partition contains the node that has the current node role of primary. All other partitions are secondary. The primary partition may not be the same for all cluster resource groups.

A peer model has the following partition rules:

- If the recovery domain nodes are fully contained within one partition, it will be the primary partition.
- If the recovery domain nodes span a partition, there will be no primary partition. Both partitions will be secondary partitions.
- If the cluster resource group is active and there are no peer nodes in the given partition, the cluster resource group will be ended in that partition.
- Operational changes are allowed in a secondary partition as long as the restrictions for the operational changes are met.
- No configuration changes are allowed in a secondary partition.

The restrictions for each Cluster Resource Group API are:

<i>Table 41. Cluster Resource Group API Partition Restrictions</i>		
Cluster Resource Group API	Allowed in primary partition	Allowed in secondary partitions
Add Node to Recovery Domain	X	
Add CRG Device Entry		
Change Cluster Resource Group	X	
Change CRG Device Entry	X	X
Create Cluster Resource Group		
Delete Cluster Resource Group	X	X
Distribute Information	X	X
End Cluster Resource Group ¹	X	
Initiate Switchover	X	
List Cluster Resource Groups	X	X
List Cluster Resource Group Information	X	X
Remove Node from Recovery Domain	X	
Remove CRG Device Entry	X	
Start Cluster Resource Group ¹	X	
Note:		
1. Allowed in all partitions for peer cluster resource groups, but only affects the partition running the API.		

By applying these restrictions, cluster resource groups can be synchronized when the cluster is no longer partitioned. As nodes rejoin the cluster from a partitioned status, the version of the cluster resource group in the primary partition is copied to nodes from a secondary partition.

When merging two secondary partitions for peer model, the partition which has cluster resource group with status of Active will be declared the winner. If both partitions have the same status for cluster resource group, the partition which contains the first node listed in the cluster resource group recovery domain will be declared the winner. The version of the cluster resource group in the winning partition will be copied to nodes in another partition.

When a partition is detected, the Add Cluster Node Entry, Adjust Cluster Version, and the Create Cluster API cannot be run in any of the partitions. The Add Device Domain Entry API can only be run if none of the nodes in the device domain are partitioned. All of the other Cluster Control APIs may be run in any partition. However, the action performed by the API takes affect only in the partition running the API.

Changing partitioned nodes to failed

Sometimes, a partitioned condition is reported when there really was a node outage. This can occur when cluster resource service loses communications with one or more nodes, but cannot detect if the nodes are still operational. When this condition occurs, a simple mechanism exists for you to indicate that the node has failed.



Attention: When you tell cluster resource services that a node has failed, it makes recovery from the partition state simpler. However, changing the node status to failed when, in fact, the node is still active and a true partition has occurred should not be done. Doing so can cause a node in more than one partition to assume the primary role for a cluster resource group. When two nodes think they are the primary node, data such as files or databases can become disjoint or corrupted if multiple nodes are each independently making changes to their copies of files. In addition, the two partitions cannot be merged back together when a node in each partition has been assigned the primary role.

When the status of a node is changed to Failed, the role of nodes in the recovery domain for each cluster resource group in the partition may be reordered. The node being set to Failed will be assigned as the last backup. If multiple nodes have failed and their status needs to be changed, the order in which the nodes are changed will affect the final order of the recovery domain's backup nodes. If the failed node was the primary node for a CRG, the first active backup will be reassigned as the new primary node.

When cluster resource services has lost communications with a node but cannot detect if the node is still operational, a cluster node will have a status of **Not communicating**. You may need to change the status of the node from **Not communicating** to **Failed**. You will then be able to restart the node.

To change the status of a node from **Not communicating** to **Failed**, follow these steps:

1. In a Web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system.
2. Log on to the system with your user profile and password.
3. Click **PowerHA** from the IBM Navigator for i window.
4. On the **PowerHA** page, click on **Cluster Nodes**.
5. On the **Cluster Nodes** page, select **Change Status** from the context menu of the node for which you want to change its status.
6. Click **Yes** on the confirmation panel.

Related information

[Change Cluster Node \(CHGCLUNODE\) command](#)

[Change Cluster Node Entry \(QcstChangeClusterNodeEntry\) API](#)

Partitioned cluster administrative domains

Consider the following information when working with partitioned cluster administrative domains.

If a cluster administrative domain is partitioned, changes continue to be synchronized among the active nodes in each partition. When the nodes are merged back together again, the cluster administrative domain propagates all changes made in every partition so that the resources are consistent within the

active domain. There are several considerations regarding the merge processing for a cluster administrative domain:

- If all partitions were active and changes were made to the same resource in different partitions, the most recent change is applied to resource on all nodes during the merge. The most recent change is determined by using Coordinated Universal Time (UTC) from each node where a change initiated.
- If all partitions were inactive, the global values for each resource are resolved based on the last change made while any partition was active. The actual application of these changes to the monitored resources does not happen until the peer CRG that represents the cluster administrative domain is started.
- If some partitions were active and some were inactive prior to the merge, the global values representing changes made in the active partitions are propagated to the inactive partitions. The inactive partitions are then started, causing any pending changes made on the nodes in the inactive partitions to propagate to the merged domain.

Tips: Cluster partitions

Use these tips for cluster partitions.

1. The rules for restricting operations within a partition are designed to make merging the partitions feasible. Without these restrictions, reconstructing the cluster requires extensive work.
2. If the nodes in the primary partition have been destroyed, special processing may be necessary in a secondary partition. The most common scenario that causes this condition is the loss of the site that made up the primary partition. Use the example in recovering from partition errors and assume that Partition 1 was destroyed. In this case, the primary node for Cluster Resource Groups B, C, and D must be located in Partition 2. The simplest recovery is to use Change Cluster Node Entry to set both Node A and Node B to failed. See changing partitioned nodes to failed for more information about how to do this. Recovery can also be achieved manually. In order to do this, perform these operations:
 - a) Remove Nodes A and B from the cluster in Partition 2. Partition 2 is now the cluster.
 - b) Establish any logical replication environments needed in the new cluster. IE. Start Cluster Resource Group API/CL command, and so on.

Since nodes have been removed from the cluster definition in Partition 2, an attempt to merge Partition 1 and Partition 2 will fail. In order to correct the mismatch in cluster definitions, run the Delete Cluster (QcstDeleteCluster) API on each node in Partition 1. Then add the nodes from Partition 1 to the cluster, and reestablish all the cluster resource group definitions, recovery domains, and logical replication. This requires a great deal of work and is also prone to errors. It is very important that you do this procedure only in a site loss situation.

3. Processing a start node operation is dependent on the status of the node that is being started:

The node either failed or an End Node operation ended the node:

- a) Cluster resource services is started on the node that is being added
- b) Cluster definition is copied from an active node in the cluster to the node that is being started.
- c) Any cluster resource group that has the node being started in the recovery domain is copied from an active node in the cluster to the node being started. No cluster resource groups are copied from the node that is being started to an active node in the cluster.

The node is a partitioned node:

- a) The cluster definition of an active node is compared to the cluster definition of the node that is being started. If the definitions are the same, the start will continue as a merge operation. If the definitions do not match, the merge will stop, and the user will need to intervene.
- b) If the merge continues, the node that is being started is set to an active status.
- c) Any cluster resource group that has the node being started in the recovery domain is copied from the primary partition of the cluster resource group to the secondary partition of the cluster resource group. Cluster resource groups may be copied from the node that is being started to nodes that are already active in the cluster.

Cluster recovery

Read about how to recover from other cluster failures that may occur.

Recovering from cluster job failures

Failure of a cluster resource services job is usually indicative of some other problem.

You should look for the job log associated with the failed job and look for messages that describe why it failed. Correct any error situations.

You can use the **Change Cluster Recovery (CHGCLURCY) command** to restart a cluster resource group job that was ended without having to end and restart clustering on a node.

1. CHGCLURCY CLUSTER (EXAMPLE) CRG (CRG1) NODE (NODE1) ACTION (*STRCRGJOB) This command will cause cluster resource group job, CRG1, on node NODE1 to be submitted. To start the cluster resource group job on NODE1 requires clustering to be active on NODE1.
2. Restart clustering on the node.

If you are using an IBM Business Partner cluster management product, refer to the documentation that came with the product.

Related information

[Change Cluster Recovery \(CHGCLURCY\) command](#)

Recovering a damaged cluster object

While it is unlikely you will ever experience a damaged object, it may be possible for cluster resource services objects to become damaged.

The system, if it is an active node, will attempt to recover from another active node in the cluster. The system will perform the following recovery steps:

For a damaged internal object

1. The node that has the damage ends.
2. If there is at least one other active node within the cluster, the damaged node will automatically restart itself and rejoin the cluster. The process of rejoining will correct the damaged situation.

For a damaged cluster resource group

1. The node that has a damaged CRG will fail any operation currently in process that is associated with that CRG. The system will then attempt to automatically recover the CRG from another active node.
2. If there is at least one active member in the recovery domain, the CRG recovery will work. Otherwise, the CRG job ends.

If the system cannot identify or reach any other active node, you will need to perform these recovery steps.

For a damaged internal object

You receive an internal clustering error (CPFBB46, CPFBB47, or CPFBB48).

1. End clustering for the node that contains the damage.
2. Restart clustering for the node that contains the damage. Do this from another active node in the cluster.
3. If Steps 1 and 2 do not solve the problem, remove the damaged node from the cluster.
4. Add the system back into the cluster and into the recovery domain for the appropriate cluster resource groups.

For a damaged cluster resource group

You receive an error stating that an object is damaged (CPF9804).

1. End clustering on the node that contains the damaged cluster resource group.
2. Delete the CRG by using the **DLTCRG** command.
3. If there is no other node active in the cluster that contains the CRG object, restore from media.
4. Start clustering on the node that contains the damaged cluster resource group. This can be done from any active node.
5. When you start clustering, the system resynchronizes all of the cluster resource groups. You may need to recreate the CRG if no other node in the cluster contains the CRG.

Recovering a cluster after a complete system loss

Use this information with the appropriate checklist in the *Recovering your system* topic for recovering your entire system after a complete system loss when your system loses power unexpectedly.

Scenario 1: Restoring to the same system

1. In order to prevent inconsistencies in the device domain information between the Licensed Internal Code and IBM i, it is recommended that you install the Licensed Internal Code by using option 3 (Install Licensed Internal Code and Recover Configuration).

Note: For the Install Licensed Internal Code and Recover Configuration operation to succeed, you must have the same disk units -- with exception of the load source disk unit if it has failed. You must also be recovering the same release.

2. After you have installed the Licensed Internal Code, follow the [Recovering Your Disk Configuration](#) procedure in the *Recovering your system* topic. These steps will help you avoid having to reconfigure the disk pools.
3. After you have recovered your system information and are ready to start clustering on the node you just recovered, you must start clustering from the active node. This will propagate the most current configuration information to the recovered node.

Scenario 2: Restoring to a different system

After you have recovered your system information and checked the job log to make sure that all objects have restored, you must perform the following steps to obtain the correct cluster device domain configuration.

1. From the node you just restored, delete the cluster.
2. From the active node, perform these steps:
 - a. Remove the recovered node from the cluster.
 - b. Add the recovered node back into the cluster.
 - c. Add the recovered node to the device domain.
 - d. Create the cluster resource group or add the node to the recovery domain.

Recovering a cluster after a disaster

In the case of a disaster where all your nodes are lost, you will need to reconfigure your cluster.

In order to prepare for such a scenario, it is recommended that you save your cluster configuration information and keep a hardcopy printout of that information.

Restoring a cluster from backup tapes

During normal operations, you should never restore from a backup tape.

This is only necessary when a disaster occurs and all nodes were lost in your cluster. If a disaster should occur, you recover by following your normal recovery procedures that you have put in place after you created your backup and recovery strategy.

Troubleshooting geographic mirroring

This information can help you solve problems related to geographic mirroring that you might encounter.

Geographic mirroring messages

Review the geographic mirroring message descriptions and recoveries to resolve your geographic mirroring problems.

0x00010259

Description: Operation failed because the system did not find the mirror copy.

Recovery: Not all the nodes in the device domain responded. Make sure that clustering is active. If necessary, start clusters on the node. See “Starting nodes” on page 55 for details. Try the request again. If the problem persists, contact your technical support provider.

0x0001025A

Description: Not all of the disk pools in the disk pool group are geographically mirrored.

Recovery: If one disk pool in a disk pool group is geographically mirrored, all of the disk pools in the disk pool group must be geographically mirrored. Take one of the following actions:

1. Configure geographic mirroring for the disk pools which are not geographically mirrored.
2. Deconfigure geographic mirroring for the disk pools that are geographically mirrored.

0x00010265

Description: The detached mirrored copy is available.

Recovery: Make the detached mirrored copy unavailable and then try the reattach operation again.

0x00010380

Description: A disk unit is missing from the configuration of the mirror copy.

Recovery: Find or fix the missing disk unit in the mirror copy. Check the Product Activity Log on destination node. Reclaim IOP cache storage.

0x00011210

Description: The proposed secondary disk pool for the disk pool group is not geographically mirrored.

Recovery: If one disk pool in a disk pool group is geographically mirrored, all of the disk pools in the disk pool group must be geographically mirrored. You must configure geographic mirroring for the proposed secondary disk pool which is not geographically mirrored, either now or after completing this operation.

0x00011211

Description: Duplicate mirror copies exist.

Recovery: Check for locally mirrored disk units that may exist on two systems, Enterprise Storage Server® FlashCopy, or back level independent disk pool copies. See the Product Activity Log on the mirror copy node for more information. Eliminate duplication and try the request again. If the problem persists, contact your technical support provider, or see IBM i Technical Support for information about IBM support and services.

Troubleshooting Metro Mirror, Global Mirror, and FlashCopy

This information can help you solve problems that are related to Metro Mirror, Global Mirror, and FlashCopy mirroring that you might encounter.

For more troubleshooting ideas for Metro Mirror, Global Mirror and FlashCopy, check out the [IBM PowerHA SystemMirror for i wiki](#).

Troubleshooting HyperSwap

Find error recovery solutions for problems that are specific to clusters.

The following are HyperSwap messages and recommended recovery.








- **System fails to IPL with SRC A6005090**

- Power down the System i and verify that the Metro Mirror replication on the IBM System Storage device is started and running in the correct direction for all logical units. Contact your next level of support if the problem persists.







Related information for Implementing high availability

Product manuals, IBM Redbooks publications, websites, and other information center topic collections contain information that relates to the High availability topic collection. You also can find related information about implementing independent disk pools, PowerHA technologies, and disaster recovery. You can view or print any of the PDF files.



IBM Redbooks

- [IBM i and IBM Storwize Family: A Practical Guide to Usage Scenarios](#) 
- [IBM i and IBM System Storage: A Guide to Implementing External Disks on IBM i](#) 
- [Implementing PowerHA for IBM i](#) 
- [Introduction to Storage Area Networks](#) 
- [iSeries in Storage Area Networks: A Guide to Implementing FC Disk and Tape with iSeries](#) 
- [PowerHA SystemMirror for IBM i Cookbook](#) 
- [Simple Configuration Example for Storwize V7000 FlashCopy and PowerHA SystemMirror for i](#) 

Websites

- [High availability with IBM PowerHA](#)  This is the IBM site for High Availability and Clusters for i, UNIX, and Linux.
- [IBM i](#) 
- [IBM PowerHA SystemMirror for i](#) 
- [IBM PowerHA SystemMirror for i wiki](#) 
- [IBM Storage](#) 
- [Power® Services](#) 



This IBM site contains the Systems lab services and training that is offered for IBM i.

- [IBM System Storage Interoperation Center \(SSIC\)](#) 
- [IBM Techdocs Library](#) 

This site gives you access to the most current installation, planning, and technical support information available from IBM pre-sales support, and is constantly updated. Look for the latest and greatest technical documents on high availability, independent disk pools, SAP, JD Edwards, etc.

- [Learning Services US](#) 

This is the IBM site for IT product training, custom solutions, and e-Learning. You can search for courses that are offered on clustering and independent disk pools.

- [Performance Management on IBM i](#) 
- [Recommended fixes](#) 

This site provides links to available PTFs for several IBM i products. For PTFs related to high availability, select the topic High Availability: Cluster, IASP, XSM, and Journal.

Information Center topic collections

- [Availability roadmap](#)
- [High availability overview](#)
- [High availability technologies](#)
- [IBM SAN Volume Controller Knowledge Center](#)
- [IBM Storwize V7000 Knowledge Center](#)
- [IBM Storwize V3700 Knowledge Center](#)
- [IBM DS8000 Knowledge Center](#)
- [Implementing high availability](#)

Other information

- [Disk management](#)
- [Resource Monitoring and Control \(RMC\)](#)

Related reference

[PDF file for Implementing high availability](#)

You can view and print a PDF file of this information about implementing high availability.

Code license and disclaimer information

IBM grants you a nonexclusive copyright license to use all programming code examples from which you can generate similar function tailored to your own specific needs.

SUBJECT TO ANY STATUTORY WARRANTIES WHICH CANNOT BE EXCLUDED, IBM, ITS PROGRAM DEVELOPERS AND SUPPLIERS MAKE NO WARRANTIES OR CONDITIONS EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OR CONDITIONS OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT, REGARDING THE PROGRAM OR TECHNICAL SUPPORT, IF ANY.

UNDER NO CIRCUMSTANCES IS IBM, ITS PROGRAM DEVELOPERS OR SUPPLIERS LIABLE FOR ANY OF THE FOLLOWING, EVEN IF INFORMED OF THEIR POSSIBILITY:

1. LOSS OF, OR DAMAGE TO, DATA;
2. DIRECT, SPECIAL, INCIDENTAL, OR INDIRECT DAMAGES, OR FOR ANY ECONOMIC CONSEQUENTIAL DAMAGES; OR
3. LOST PROFITS, BUSINESS, REVENUE, GOODWILL, OR ANTICIPATED SAVINGS.

SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF DIRECT, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, SO SOME OR ALL OF THE ABOVE LIMITATIONS OR EXCLUSIONS MAY NOT APPLY TO YOU.

Notices

This information was developed for products and services offered in the U.S.A.

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

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS 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 states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

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

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

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Software Interoperability Coordinator, Department YBWA
3605 Highway 52 N
Rochester, MN 55901
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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

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

This information is for planning purposes only. The information herein is subject to change before the products described become available.

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

COPYRIGHT LICENSE:

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

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. _enter the year or years_.

Programming interface information

This Implementing high availability publication documents intended Programming Interfaces that allow the customer to write programs to obtain the services of IBM i.

Trademarks

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

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

Other product and service names might be trademarks of IBM or other companies.

Terms and conditions

Permissions for the use of these publications is granted subject to the following terms and conditions.

Personal Use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of IBM.

Commercial Use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.



Product Number: 5770-SS1