

System i

Availability

Implementing High Availability with the solution-based approach—Operator's guide

Version 6 Release 1





System i

Availability

Implementing High Availability with the solution-based approach—Operator's guide

Version 6 Release 1

Note

Before using this information and the product it supports, read the information in "Notices," on page 105.

This edition applies to version 6, release 1, modification 0 of IBM i5/OS (product number 5761-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.

© **Copyright IBM Corporation 1998, 2008.**

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

Contents

Implementing high availability with the solution-based approach 1

Selecting a high availability solution	6
Overview: Switched disk between logical partitions	6
Overview: Switched disk between systems	7
Overview: Switched disk with geographic mirroring	8
Overview: Cross-site mirroring with geographic mirroring	9
Verifying requirements for your high availability solution	10
Setting up a high availability solution	11
Placing high availability systems in dedicated state	13
Changing backup procedures for a high-availability solution	13
Verifying applications after each setup step	14
Managing a high availability solution	14
Working with high availability status messages	15
Working with high availability solution tasks	17
Shutting down a high availability solution with data available	17
Shutting down a high availability solution with data unavailable	18
Resuming a high availability solution	19
Performing switchovers	19
Switchover for a two-node solution	19
Switchover for a three-node solution	20
Recovering from partition state	20
Collecting service information	21
Managing high availability resources	21
Managing nodes	21
Starting nodes	22
Stopping nodes	23
Working with all nodes	23
Displaying node properties	24
Managing cluster resource groups (CRG)	24
Starting a CRG	26
Stopping a CRG	26
Deleting a CRG	27
Working with all CRGs	27
Displaying CRG properties	28
Managing monitored resources	28
Working with all monitored resources	29
Displaying monitored resource properties	29
Managing independent disk pools	30
Starting mirroring	31
Stopping mirroring	31
Working with all independent disk pools	32
Displaying independent disk pool properties	32

Managing TCP/IP interfaces	33
Starting TCP/IP interfaces	33
Stopping TCP/IP interfaces	34
Working with all TCP/IP interfaces	34
Displaying properties for TCP/IP interfaces	35
Managing policies	35
Working with high availability event log	37
Working with messages	37
Reference: Panel help for the solution-based interface	41
High Availability Solutions Manager welcome	41
Select your high availability solution	43
Customized shopping list	43
Executive summary	48
Save or print high availability solution information	48
Verify requirements list	49
Set up your high availability solution	54
Display high availability log	57
Set up high availability policies	57
Set up high availability environment	58
Verify administrative switchover	61
Migrate user profiles	63
Migrate libraries	66
Migrate directories	70
Finish set up and clean up files	75
Manage your high availability solution - Nodes	76
Manage high availability solution - Cluster Resource Groups	80
Manage high-availability solution - Monitored resources	83
Manage high-availability solution - Independent disk pools	86
Manage high-availability solution - TCP/IP interfaces	89
Manage high availability solution - Policies	93
Collect service information	95
Verify administrative switchover	96
Recover from partition state	98
Shut down high availability solution - keep data available	100
Shut down high availability solution - make data not available	101
Resume high availability solution	102
No help is available	103

Appendix. Notices 105

Programming interface information	107
Trademarks	107
Terms and conditions	107

Implementing high availability with the solution-based approach

The High Availability Solutions Manager graphical interface provides a solution-based approach to selecting, configuring, and managing high availability in your business.

With High Availability Solutions Manager, which is part of the IBM® System i® High Availability Solutions Manager (iHASM) licensed program (5761-HAS) you can select from several predefined high-availability solutions. For each of these solutions, the dependent technologies are configured based on your selection. The High Availability Solutions Manager graphical interface provides easy-to-use tools to manage your high-availability solution.

Implementing a high-availability solution can be a complex task, and each business environment is unique. Working with a business partner skilled in high availability is recommended, both when selecting and configuring a high-availability solution.

View the following Flash demo for an animated overview of the high availability solutions that High Availability Solutions Manager provides. The following Flash demo requires the Flash plug-in .

(“Transcript”).

Transcript

This topic provides a description of the Flash demo, which provides an overview of High Availability Solutions Manager.

IBM High Availability Solutions Manager

Audio

It's 10 p.m. Do you know where your business is? Is it processing bank transactions for a customer? Or is it selling products to a shopper in another state?

Whatever the case, whatever the need, are you prepared—for this?

Most businesses aren't equipped for unexpected outages or server failures. A high-availability solution provides your business with an answer to planned or unplanned outages, keeping your business running.

IBM has developed a solution to ensure your business is ready before the problem occurs. It's called IBM High Availability Solutions Manager, and it's more than a disaster recovery solution. It's disaster prevention. In a business environment, failure to access data is only one element to a server or system outage.

A true high-availability solution also includes redirected access to applications, environment, and policies. IBM has designed the High Availability Solutions Manager to cover access to everything, ensuring that, along with your data, your business applications, user information, passwords, and identification are also available.

It provides you with information to make the right high-availability choice to meet your business needs. IBM's High Availability Solutions Manager removes the complexity of managing your high-availability solution with a built-in and easy-to-use monitoring tool.

To learn more about IBM's high-availability solutions select one of the options.

Switched disk between logical partitions

Slide 1 of 5

Description: Image of a system with two logical partitions labeled LPAR1 and LPAR2 external to system

Audio: IBM's switched disk between logical partitions solution includes a single system with two logical partitions.

Slide 2 of 5

Description: Image of a system and two logical partitions animates.

Audio: Within each partition is SYSBAS, which is storage that contains the operating system for each partition.

Slide 3 of 5

Description: Image of system with two logical partitions and independent disk pool animates.

Audio: Also included is a single independent auxiliary storage pool, otherwise known as an independent ASP. The independent ASP disks can reside either internally or externally.

Slide 4 of 5

Description: Image of system with two logical partition and an independent disk pool animates showing data transfer.

Audio: The independent ASP provides the key mechanisms for each partition sharing access to the data, applications, environment, and policies.

Slide 5 of 5

Description: Image of system with two logical partition and an independent disk pool animates showing switch over between logical partitions.

Audio: If you have a planned outage or failure on one partition, the ownership of the independent ASP switches over to the other logical partition.

Switched disk between systems

Slide 1 of 8

Description: Animated images of two system

Audio: IBM's switched disk between systems solution consists of two single-partition independent systems.

Slide 2 of 8

Description: Animated images of two system with SYSBAS

Audio: Within each system is SYSBAS, which is storage that contains the operating system for each system.

Slide 3 of 8

Description: Animated images of two system with SYSBAS with an external tower containing independent disk pools

Audio: They are connected to an external tower that has a single independent auxiliary storage pool, otherwise known as an independent disk pool.

Slide 4 of 8

Description: Animated images of two system with SYSBAS with an external tower containing independent disk pools. Shows data stored on independent disk pools.

Audio: The independent ASP provides the key mechanisms for each system sharing access to the data, applications, environment, and policies.

Slide 5 of 8

Description: Animated images of two system with SYSBAS with an external tower containing independent disk pools. Shows connection between independent disk pools and systems.

Audio: The connection between the independent disk pool and the two systems is done through a high-speed link that provides optimal transfer of data.

Slide 6 of 8

Description: Animated images of two system with SYSBAS with an external tower containing independent disk pools. Shows relationship between independent disk pools and System 1.

Audio: In this example, the independent disk pool is owned by system 1 and the information flow is continuous between the independent disk pool and system 1.

Slide 7 of 8

Description: Animated images of two system with an external tower containing independent disk pools. Shows data switchover to System 2 during a planned outage.

Audio: If you have a planned outage or failure on system 1, the ownership of the independent disk pool switches over to System 2.

Slide 8 of 8

Description: Animated images of two system with an external tower containing independent disk pools. Shows a user connection during an outage.

Audio: This solution also works if a user is connected into system 1 with an IP address. If system 1 fails, the IP address is dynamically reassigned to system 2, and system 2 begins to communicate with the independent disk pool.

Cross-site mirroring with geographic mirroring

Slide 1 of 8

Description: Animated images of two system one located uptown and the other located downtown.

Audio: IBM's cross-site mirroring with geographic mirroring solution consists of two systems, each located in a different geographical place. The distance between the systems is limited to a maximum of 300 kilometers. In this example, System 1 is located uptown and System 2 is downtown.

Slide 2 of 8

Description: Animated images of two system with SYSBAS

Audio: Within each system is SYSBAS, which is storage that contains the operating system for the system.

Slide 3 of 8

Description: Animated images of two system with SYSBAS with an external tower containing independent disk pools

Audio: Both systems are connected to external towers that have single independent auxiliary storage pools otherwise known as an independent disk pools.

Slide 4 of 8

Description: Animated images of two system with SYSBAS with an external tower containing independent disk pools. Shows data stored on independent disk pools.

Audio: The independent disk pools provides the key mechanisms for each system sharing access to the data, applications, environment, and policies.

Slide 5 of 8

Description: Animated images of data being mirrored from independent disk pools on System 1 to independent disk pools on System 2.

Audio: With IBM's cross-site mirroring with geographic mirroring, the independent disk pool data is mirrored from one set of disks to another at a page level, so the data from the production independent disk pool is copied over to the mirrored independent disk pool.

Slide 6 of 8

Description: Animated images of two system with SYSBAS with an external tower containing independent disk pools. Shows data being updated and mirrored synchronously.

Audio: The transfer of data from the production independent disk pool to the mirrored independent disk pool is done in a synchronous way, so a user making a change on the production independent disk pool has to wait until the transmission happens to the mirrored independent disk pool.

Slide 7 of 8

Description: Animated images of two system with an external tower containing independent disk pools. Shows data stored in the mirrored independent disk pool is not accessible to System 2.

Audio: The mirrored independent disk pool is only used for high-availability purposes and is not currently accessible by system 2.

Slide 8 of 8

Description: Animated images of two system with an external tower containing independent disk pools. Shows data being switched to the mirrored copy during an outage.

Audio: Only one copy of the data is accessible at any given time. If you have a planned outage or failure on System 1, the ownership of the Independent ASP switches over to System 2 and the mirrored Independent ASP becomes the production Independent ASP.

Switched disk with geographic mirroring

Slide 1 of 9

Description: Animated images of two system, where one contains logical partitions. One system is located uptown and the other is located downtown.

Audio: IBM's switched disk with geographic mirroring solution includes one system with two logical partitions and another physical system located in another geographical region. The distance between the systems is limited to a maximum of 300 kilometers. In this example system 1 is located uptown and system 2 is downtown.

Slide 2 of 9

Description: Animated images of system 1 containing two logical partitions.

Audio: Within each logical partition of System 1 is SYSBAS, which is storage that contains the operating system for that partition.

Slide 3 of 9

Description: Animated images of system 1 containing independent disk pools.

Audio: Also included is a single independent auxiliary storage pool, otherwise known as an independent disk pool.

Slide 4 of 9

Description: Animated images of independent disk pool and the type of data it contains.

Audio: The independent disk pool provides the key mechanisms for each partition sharing access to the data, applications, environment, and policies.

Slide 5 of 9

Description: Animated image of data flow between partitions and independent disk pool.

Audio: If you have a planned outage or failure on one partition, the ownership of the independent disk pool switches over to the other logical partition.

Slide 6 of 9

Description: Animated image of data flow from independent disk pool to remote independent disk pool.

Audio: IBM's switched disk with geographic mirroring provides disaster recovery because the independent disk pool data is mirrored from one set of disks to another at a page level, so the data from the production independent disk pool is copied over to the mirrored independent disk pool .

Slide 7 of 9

Description: Animated image of data flow from user to mirrored independent disk pool.

Audio: The transfer of data from the production independent disk pool to the mirrored independent disk pool is done in a synchronous way so a user making a change on the production independent disk pool has to wait until the transmission happens to the mirrored independent disk pool.

Slide 8 of 9

Description: Animated image show no concurrent access to the second copy of data.

Audio: The mirrored independent disk pool is only used for disaster recovery purposes and is not concurrently accessible by system 2. Only one copy of the data is accessible at any given time.

Slide 9 of 9

Description:

Audio: If you have an outage to both partitions in system 1, the ownership of the independent disk pool switches over to system 2, and the mirrored independent disk pool becomes the production independent disk pool.

Selecting a high availability solution

The High Availability Solutions Manager graphical interface provides several predefined solutions, which use high-availability technologies supported by the i5/OS operating system. Each solution provides different levels of high-availability coverage and has specific advantages, restrictions and requirements.

About this task

To select one of the high-availability solutions, complete the following steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS[®] navigation area.
4. On the **Welcome** page, select **Select your high-availability solution**.
5. From the list of supported solutions, select the desired solution and click **Select**. You can select from the following solutions:
 - Switched disk between logical partitions
 - Switched disk between systems
 - Switched disk with geographic mirroring
 - Cross-site mirroring with geographic mirroring

For each solution, you can review overviews and executive summaries, and generate a dynamic list of requirements for that solution.

Overview: Switched disk between logical partitions

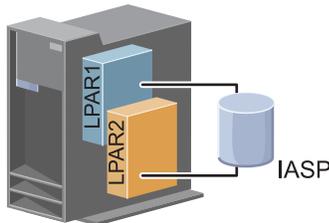
The High Availability Solutions Manager graphical interface defines several solutions for an i5/OS high-availability environment. This high-availability solution uses disk pools that are switched between two logical partitions that reside on a single system.

Logical partitioning is the ability to make a single i5/OS system function as if it were two or more independent systems. This solution is a good choice for businesses that already have logical partitions configured in their environment.

The High Availability Solutions Manager graphical interface does not perform any logical partition configuration. It validates that logical partitions exist in your environment and makes recommendations based on that discovery.

It also configures all the other necessary technologies, such as clusters and cluster administrative domains, and tests the setup by performing a series of switchovers that validates the configuration.

The following graphic illustrates the environment that uses this solution:



This solution has the following advantages:

- This is low-cost solution that uses available system resources.
- It provides availability for your business resources during planned outages.
- It provides availability for business resources during some unplanned outages, such as a single logical partition failure.
- Because this solution uses a single copy of data, it minimizes the number of disk units that are required.
- This solution contains current data that does not need to be synchronized.

This solution has the following restrictions:

- There is no disaster recovery available for a site-wide outage.
- There is a requirement that you configure a logical partition.
- There is a possible requirement for redundant hardware between partitions.
- There is only one logical copy of the data that resides in the independent disk pool. This can be a single point of failure, although the data may be protected with RAID protection.
- There is not any concurrent access to the disk pool from both logical partitions.

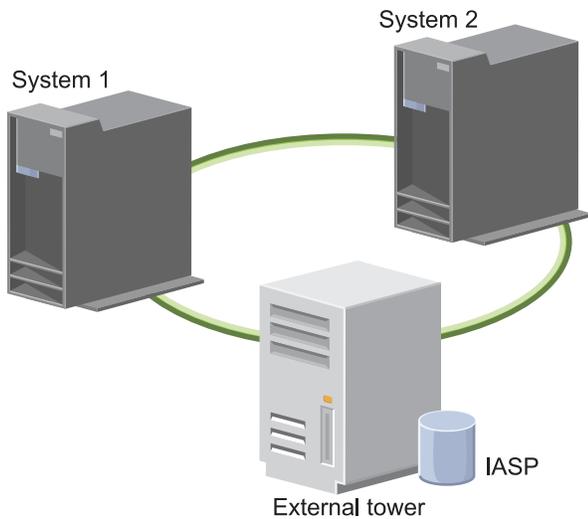
Overview: Switched disk between systems

The High Availability Solutions Manager graphical interface defines several solutions for an i5/OS high-availability environment. This solution uses switched disks between two systems and provides high availability for data, applications, or devices during planned and unplanned outages.

By using switched disk technology, this solution provides a simple high-availability solution. With this solution, a single copy of data that is stored in the switched disk always remains current, which eliminates the need to synchronize data among systems and eliminates the risk of losing data during transmission.

With this solution, the High Availability Solutions Manager graphical interface configures disks and moves data and applications to the independent disk pool. It also configures all the other necessary technologies, such as clusters and cluster administrative domains, and tests the setup by performing a series of switchovers that validates the configuration.

The following graphic illustrates this solution:



This solution has the following advantages:

- Provides availability for your business resources during planned outages
- Provides availability for business resources during some unplanned outages
- Enables a single copy of data, which minimizes the number of disk units that are required
- Provides minimal performance overhead
- Enables data to remain current and does not need to be synchronized

This solution has the following restrictions:

- There is not any disaster recovery for a site-wide outage
- There is only one logical copy of the data resides in the independent disk pool. This can be a single point of failure, although the data may be protected with RAID protection.
- There is not any concurrent access to the disk pool from both systems

Overview: Switched disk with geographic mirroring

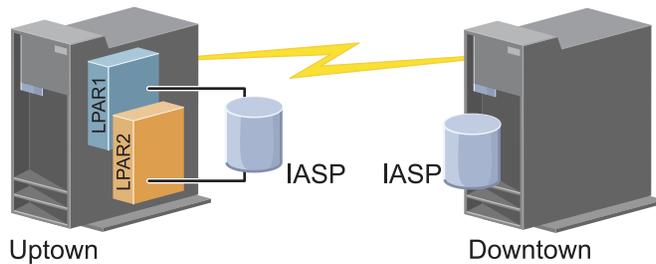
The High Availability Solutions Manager graphical interface provides several predefined i5/OS high-availability solutions. This solution uses a three-node cross site mirroring environment to provide both disaster recovery and high availability.

At the production site (Uptown), switched disks are used to move independent disk pools between two nodes. The solution also uses geographic mirroring to generate a copy of the independent disk at a second site (Downtown). Thus, this solution provides both disaster recovery and high availability. The benefits of this solution are essentially the same as the basic switched disk solution with the added advantage of providing disaster recovery to application data by duplicating that data at another location. The production site (Uptown) has an independent disk pool that can be switched between logical partitions to provide high availability with fast switchover times for planned outages, such as applying fixes. This solution also provides disaster recovery with cross-site and geographic mirroring.

Geographic mirroring is a subfunction of cross-site mirroring where data is mirrored to a copy of the independent disk pool at the remote location. Data from the independent disk pool at the production site (Uptown) is mirrored to an independent disk pool on the backup site (Downtown). 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 storages solutions provide.

The High Availability Solutions Manager graphical interface verifies the requirements and configures the associated high availability technologies for this solution. However, it does not configure logical partitions for this solution. This solution can be implemented with existing logical partitions, or logical partitions can be created prior to the high availability configuration. It is recommended that you perform the necessary hardware and software planning for logical partitions prior to any configuration of your high-availability solution.

The following graphic illustrates this solution:



This solution has the following advantages:

- Provides availability for your business resources during planned outages
- Provides availability for business resources during some unplanned outages
- Provides availability for business resources during site-wide disasters
- Enables each site to have a single copy of data which minimizes the number of disk units that are required
- Enables data to remain current and does not need to be synchronized

This solution has the following restrictions:

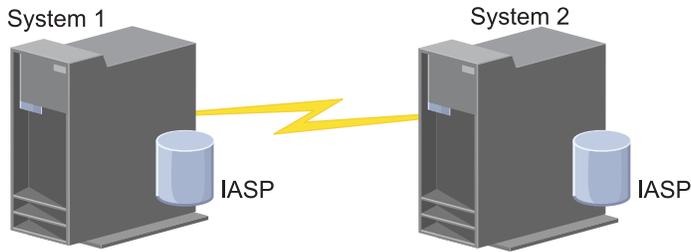
- | • There is no concurrent access to the disk pool. However, you can detach the mirror copy for offline processing of a second copy of the data.
- | • There is potential performance effects with increased central processing unit (CPU) that is required to support geographic mirroring
- | • Consider using redundant communication paths and adequate bandwidth

Overview: Cross-site mirroring with geographic mirroring

The High Availability Solutions Manager graphical interface defines several solutions for a high-availability environment. This solution provides high availability and disaster recovery by maintaining identical copies of disk pools at two sites that are geographically separated from one another.

Geographic mirroring is a subfunction of cross-site mirroring where data is mirrored to a copy of the independent disk pool at the remote location. In 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.

The following graphic illustrates this solution:



This solution has the following advantages:

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

This solution has the following restrictions:

- There is no concurrent access to the disk pool. 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

Verifying requirements for your high availability solution

The High Availability Solutions Manager graphical interface displays a list of requirements for your selected high-availability solution. Before configuring a solution, you need to provide configuration information that will enable setup of your selected solution.

About this task

After you have selected a solution for high availability, you should verify that the necessary software and hardware requirements are met by completing these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. On the **Welcome** page, select **Verify requirements before setting up your high-availability solution**.
5. On the **Verify Requirements List** page, check the list of requirement. The interface provides the following statuses for each requirement:

Example

Table 1. Verify requirement status

Status	Description
	<p>The requirement must be met to set up your high availability solution. You can meet a requirement by supplying missing identification information for resources which will be used later to set up your high-availability solution or by installing missing hardware and software.</p> <p>After entering or selecting missing information click OK. After installing missing hardware or software, click Refresh so the requirements list can be updated to reflect the changes.</p>

Table 1. Verify requirement status (continued)

Status	Description
	<p>The requirement is optional but might be recommended based on your specific business requirements. For example, a minimum of one disk drive is required to configure an independent disk pool but one disk might not be sufficient to store all of your data.</p> <p>You can meet a requirement by supplying missing identification information for resources which will be used later to set up your high-availability solution or by installing missing hardware and software.</p> <p>After entering or selecting missing information click OK. After installing missing hardware or software click Refresh so the requirements list can be updated to reflect the changes.</p>
	The requirement has been met.

What to do next

After all required hardware, software and information has been identified and successfully verified, you can set up your solution.

Setting up a high availability solution

The High Availability Solutions Manager interface configures your selected high-availability solution.

Before setting up your solution, 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. Select **High Availability Solutions Manager** from your IBM Systems Director Navigator for i5/OS window.
4. On the **Welcome** page, select **Set up your high availability solution**.

To set up your selected high-availability solution, follow these steps. When each step is completed, the Status column indicates whether the step completed successfully, failed, or was incomplete. If the step completed successfully, the  arrow moves to the next step.

1. Click **Go** to perform the **Set up high availability policies** step.
 - a. On the Set Up High Availability Policies page, select the policies that you want to use and click **OK**.
 - b. When this step is completed, the Status column indicates whether the step completed successfully, failed, or was incomplete. If the step completed successfully, the  arrow moves to the next step.
2. Click **Go** to perform the **Set up high availability environment** step.
 - a. On the Set Up High Availability Environment page, click **Run Now**.
 - b. When this step is completed, click **Close** to return to this page.
3. Click **Go** to perform the **Verify administrative switchover from {1} to {2}** step, where {1} represents the primary node for your solution and {2} represents the backup node for your solution. (If you are setting up a three-node solution, this switchover is from the primary node to the local backup node.)
 - a. On the Verify Administrative Switchover page, click **Run Now**.
 - b. When this step is completed, click **Close** to return to this page.
 - c. Manually verify that the administrative switchover was successful.

4. Click **Go** to perform the **Verify administrative switchover from {2} to {3}** step, where {2} represents the backup node for your solution and {3} represents the primary node for your solution. (If you are setting up a three-node solution, this switchover is from the local backup node to the remote backup node.)
 - a. On the Verify Administrative Switchover page, click **Run Now**.
 - b. When this step is completed, click **Close** to return to this page.
 - c. Manually verify that the administrative switchover was successful.
5. Click **Go** to perform the **Verify administrative switchover from {3} to {1}** step. This step displays only if you are setting up a three-node solution. This switchover is from the remote backup node to the primary node.
6. Click **Go** to perform the **Migrate user profiles** step.
 - a. On the Migrate User Profiles page, select the profiles that you want to migrate to the high-availability solution.
 - b. When this step is completed, click **Close** to return to this page.
7. Click **Go** to perform the **Migrate libraries** step.
 - a. On the Migrate Libraries page, select the libraries that you want to migrate to the high-availability solution.
 - b. When this step is completed, click **Close** to return to this page.
8. Click **Go** to perform the **Migrate directories** step.
 - a. On the Migrate Directories page, select the directories that you want to migrate to the high-availability solution.
 - b. When this step is completed, click **Close** to return to this page.
9. Click **Go** to perform the **Verify administrative switchover from {1} to {2}** step, where {1} represents the primary node for your solution and {2} represents the backup node for your solution. (If you are setting up a three-node solution, this switchover is from the primary node to the local backup node.)
 - a. On the Verify Administrative Switchover page, click **Run Now**.
 - b. When this step is completed, click **Close** to return to this page.
 - c. Manually verify that the administrative switchover was successful.

Note: Before completing the setup of your high-availability solution, you must make sure that your business applications will work appropriately on all nodes in your solution. This step involves switching the nodes on which your applications will be running. When all the nodes in the solution have been verified, you can proceed to complete setting up the high-availability solution.

10. Click **Go** to perform the **Verify administrative switchover from {2} to {3}** step, where {2} represents the backup node for your solution and {3} represents the primary node for your solution. (If you are setting up a three-node solution, this switchover is from the local backup node to the remote backup node.)
 - a. On the Verify Administrative Switchover page, click **Run Now**.
 - b. When this step is completed, click **Close** to return to this page.
 - c. Manually verify that the administrative switchover was successful.
11. Click **Go** to perform the **Verify administrative switchover from {3} to {1}** step. This step displays only if you are setting up a three-node solution. This switchover is from the remote backup node to the primary node.
12. Click **Go** to perform the **Finish set up and clean up work files** step.
 - a. On the Finish Set Up and Clean Up Work Files page, click **Run Now**.
 - b. When this step is completed, click **Close** to return to this page.
13. Click **Close** on the Set Up Your High Availability Solution page. You have finished setting up your high-availability solution. Your high-availability solution is now active and ready to be managed.

Placing high availability systems in dedicated state

Before setting up your high-availability solution, each system that is participating in the high-availability solution must be in dedicated state. The setup process for high availability is disruptive to the system to the point where normal business operations are not possible. It is recommended that setup be completed during off-hours to eliminate this effect.

Before you begin

Placing your systems in dedicated state is specific to your environment. However, several common actions should be taken to place your systems in dedicated state. The High Availability Solutions Manager graphical interface provides a message area that displays subsystems and jobs that may need to be ended in order to reach dedicated state.

Procedure

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 QSECOFR user profile and password.
3. End all user jobs. See Ending jobs for details. The message area displays jobs that should ended.
4. Ensure that TCP/IP is active. See Tools for verifying your network structure for information on different tools that can validate TCP/IP connections.
5. End all jobs and subsystems associated with all licensed program (all LPs).
6. Ensure that all subsystem jobs are ended except QCTL, QBATCH, QSYSWRK, and QUSRWRK.

Changing backup procedures for a high-availability solution

As part of configuring a high-availability solution with the High Availability Solution Manager interface, data is migrated to independent disk pools. To ensure that these objects continue to be backed up appropriately, you might need to change existing backup procedures.

Before you begin

Before you setup your high-availability solution, you should back up all of your data and system information. Use Customizing your GO SAVE backup to produce a set of instructions that is tailored to your save environment.

About this task

After you set up the high-availability solution and before migrating user profiles or user data into the high-availability environment, you need to ensure that the objects migrated to the independent disk pools can continue to be backed up. The following are some backup procedures that you might need to change if you are using GO SAVE options:

Procedure

1. You can save independent disk pools separately, you can save them as part of a full system save (GO SAVE: Option 21), or you can save all user data (GO SAVE: Option 23). You must make the independent disk pools available before you perform the save operation.
2. If your system includes independent disk pools that use cross-site mirroring with geographical mirroring, metro mirror, or global mirror, eliminate them from the GO SAVE option by varying off the independent disk pools. You should save these independent disk pools separately from this GO SAVE operation. If these independent disk pools remain available during the GO SAVE operation, cross-site mirroring is suspended when the system becomes restricted. When you resume mirroring after the save operation, a partial synchronization is required.
3. After your high-availability solution is set up, if you want to save independent disk pools without causing a failover or if you want to save the cluster environment for a node, you must shut down the

high availability solution first. Select the task, Shut down high availability solution - keep data available, from dropdown list in the Manage Your High Availability Solution section.

When the save operation completes, you can resume the high-availability solution.

Related tasks

“Shutting down a high availability solution with data available” on page 17

The High Availability Solutions Manager interface allows you to quickly and easily perform the necessary steps for shutting down your high-availability solution while keeping data available to your customers and users.

“Resuming a high availability solution” on page 19

The High Availability Solutions Manager graphical interface enables you to resume a high-availability solution that has been shut down.

Related information

Performing a complete save with the GO SAVE checklist

Saving independent ASPs

Verifying applications after each setup step

After the High Availability Solutions Manager graphical interface completes each setup step for your high-availability solution, you should verify that your applications still work properly.

Before you begin

Setting up your high-availability solution must be done when the systems in your high-availability solution are not being accessed by any internal or external users. It is strongly recommended that you verify your business applications work appropriately after each step in the setup process completes. Verifying after each setup step will enable you to identify migration problems before they could impact your end users.

Determining that your applications work in an independent disk pool environment is potentially a long process and it is a step that should be done on a non-production system prior to any migration to your production system. It is common to find that some changes are required to your applications before they will work in an independent disk pool environment. It is strongly recommended that you complete thorough testing of your applications in an independent disk pool environment before setting up your high-availability solution with the High Availability Solutions Manager graphical interface.

As an example, if you migrate a library into the high-availability environment that has hard coded library references, the applications using that library may no longer work after the migration. By verifying your applications after the migration of LIBA, you will discover the problem and be able to migrate LIBA out of the high-availability solution so that the applications will continue to work as before.

Related information

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

Managing a high availability solution

The High Availability Solutions Manager graphical interface enables you to manage your high-availability solution through solution-level tasks that are generated dynamically based on the current status of your solution. In addition to these tasks, you can also manage high availability resources that comprise your solution and view event messages.

Before you begin

You can monitor and perform high-availability solution-level actions, monitor and manage high-availability solution resources, and monitor solution events in the event log. The dynamically changing status icons makes this page ideal to use as a *dashboard*.

There are three sections on the Manage Your High Availability Solution page:

- The **Manage Your High Availability Solution** section provides an at-glance view of the status of the high-availability solution and quick access to solution-level actions.
- The **High Availability Solution Resources** section provides a tabbed list of all of the high-availability solution resources. Each tab gives a detailed view of each resource along with possible actions to perform on the resource.
- The **Event Log** section presents the list of events that have occurred in the high-availability solution.

After you have set up your high-availability solution, you can manage it. You must have the IBM System i High Availability Solutions Manager (iHASM) licensed product installed to be able to manage the high-availability solution.

About this task

To manage your high-availability solution, follow these steps:

Procedure

1. In a web browser, enter `http://mysystem:2001`, where `mysystem` is the host name of the system with the high-availability solution.
2. Log on to the system with your user profile and password.
3. Under **i5/OS Management**, select **High Availability Solution Manager**.
4. On the High Availability Solutions Manager welcome page, select **Manage your high-availability solution**.

Working with high availability status messages

The High Availability Solutions Manager graphical interface provides overall status messages for your high-availability solution and for your high availability resources. These messages indicate whether the high-availability solution is working correctly, needs attention, or is not working correctly.

The following table describes the status graphics that are displayed for all elements in the high-availability solution.

Status	Description
	Consistent. Indicates that the solution or resource is working correctly. You can initiate an administrative switchover, shut down the high-availability solution, or collect service information.
	Error. Indicates that one or more resources in the solution may have errors which require attention. If the overall status is in error, look at the high availability resources section to see which resource is causing the problem, and take actions to resolve that problem. For example, if the solution is indicating error status because one of the nodes is stopped, you can restart the node to correct the problem.
	Pending. Indicates that one or more resources in the solution are pending.
	Service. Indicates that one or more resources in the solution need requires service.
	Stopped. Indicates that the solution or resource is stopped. You need to restart the solution or resource.
	Unknown. Indicates that the resource is unknown to the system.

Status	Description
	Warning. Indicates a possible problem with the solution or resource, or that a lengthy operation is in progress. You can shut down and restart the high-availability solution, or collect service information.

To work with high availability status messages, follow these steps:

- Review the event log to find any new messages about high-availability solution events.
- Look at the status of the high-availability solution in the Manage Your High Availability Solution section. This information indicates whether the solution is available for switchover, and the status of all nodes, cluster resource groups, monitored resources, independent disk pools, and TCP/IP interfaces. The actions that you can take depend on the status of the high-availability solution.
- Look at the status of each resource in the High Availability Solution Resources section. This information indicates whether each resource is consistent or inconsistent. The actions that you can take depends on the status of each resource:
 - Nodes
 - Cluster resource groups
 - Monitored resources
 - Independent disk pools
 - TCP/IP interfaces
 - Policies

These are the actions that you can take at the high-availability solution level.

Condition	Status	Action you can take	Result
Overall status is consistent		Select Switchover to initiate an administrative switchover of the high-availability solution. If you are using cross-site mirroring with switched disks, you can select Switchover at same site or Switchover to different site .	The high-availability solution switches from the primary node to the backup node.
Overall status is consistent		Select Shut down - Make data not available to shut down the high-availability solution for IPL.	The high-availability solution ends all cluster resource groups and nodes, varies off the independent disk pool, and prepares the environment for a power off.
Overall status is consistent		Select Shut down - Keep data available to shut down the high-availability solution for a dedicated system backup.	The high-availability solution ends all cluster resource groups and nodes, but does not vary off the independent disk pool, and prepares the environment for a backup. The data on the independent disk pool is still available.
All overall statuses have this option		Select Service info to collect or print the service information.	High Availability Solutions Manager collects all of the service-related information for the high-availability solution. You have the option to save or print this information.

Condition	Status	Action you can take	Result
Overall status is in error. The independent disk pool is manually suspended with or without tracking, or is suspended because of a communications failure.		Select Resume high availability solution to resume geographic mirroring.	High Availability Solutions Manager resumes geographic mirroring of the independent disk pool.
Overall status is in error. The independent disk pool is detached with tracking.		Select Resume high availability solution .	High Availability Solutions Manager reattaches and resumes geographic mirroring, and starts a partial synchronization of the independent disk pool. While the solution is being synchronized, the overall status is warning. You will not be able to work with your solution until the synchronization process completes.
Overall status is in error. The independent disk pool is detached without tracking.		Select Resume high availability solution .	High Availability Solutions Manager reattaches and resumes geographic mirroring, and starts a full synchronization of the independent disk pool. While the solution is being synchronized, the overall status is warning. You will not be able to work with your solution until the synchronization process completes.
Overall status is in error. One of the cluster resource groups is stopped.		Select Resume high availability solution .	High Availability Solutions Manager starts the cluster resource groups that are stopped.
Overall status is in error. One of the nodes is stopped.		Select Resume high availability solution .	High Availability Solutions Manager starts the nodes that are stopped.
Overall status is warning.			

Working with high availability solution tasks

The High Availability Solutions Manager graphical interface simplifies various administrative tasks, such as initiating an administrative switchover, shutting down and restarting the high-availability solution, and collecting service information. The High Availability Solutions Manager graphical interface also recovers from specific problems automatically, such as cluster partition state, stopped or locked resources, or a suspended independent disk pool.

About this task

Shutting down a high availability solution with data available

The High Availability Solutions Manager interface allows you to quickly and easily perform the necessary steps for shutting down your high-availability solution while keeping data available to your customers and users.

About this task

To shut down a high-availability solution before performing a dedicated backup, perform these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. On the High Availability Solutions Manager welcome page, select **Manage your high-availability solution**.
5. Select **Shut Down - Keep Data Available** from the dropdown list in the Manage Your High Availability Solution section.
6. From the Shut Down - Keep Data Available page, click **Run Now**. The High Availability Solutions Manager graphical interface performs the steps for shutting down the high-availability solution. It stops all high availability resources, but keeps the independent disk pool varied on.
7. When the high-availability solution is shut down, click **Close** to return to the Manage Your High Availability Solution page. The overall status of the high-availability solution is refreshed. You can now place your system in restricted state without switching the independent disk pool to the backup system and perform your backup.

What to do next

Related tasks

“Changing backup procedures for a high-availability solution” on page 13

As part of configuring a high-availability solution with the High Availability Solution Manager interface, data is migrated to independent disk pools. To ensure that these objects continue to be backed up appropriately, you might need to change existing backup procedures.

Shutting down a high availability solution with data unavailable

The High Availability Solutions Manager graphical interface allows you to end the high-availability solution in situations where you need to IPL a system. The system within the high-availability environment needs to be shut down for operating system or hardware upgrades. When you perform this action, applications and data are not failed over to a backup system.

About this task

To shut down a high-availability solution before performing a system IPL, perform these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. In High Availability Solutions Manager, select **Manage your high-availability solution**.
5. Select **Shut Down - Make Data Not Available** from the dropdown list in the Manage Your High Availability Solution section.
6. From the Shut Down - Make Data Not Available page, click **Run Now**. The High Availability Solutions Manager graphical interface performs the steps for shutting down the high-availability solution. It stops all high availability resources and varies off the independent disk pool. Data becomes unavailable to end users.
7. When the solution is shut down, click **Close** to return to the Manage Your High Availability Solution page. The overall status of the solution is refreshed. The system is now ready for IPL.

| **Resuming a high availability solution**

| The High Availability Solutions Manager graphical interface enables you to resume a high-availability solution that has been shut down.

| **Before you begin**

| You can resume a high-availability solution in the following situations:

- | • After the high-availability solution has shut down automatically because of a problem with resources or geographic mirroring.
- | • After you have shut down the high-availability solution to perform a dedicated backup or system IPL.

| **About this task**

| To resume a high-availability solution, perform these steps:

| **Procedure**

- | 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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
- | 4. From the High Availability Solutions Manager welcome page, select **Manage your high availability solution**.
- | 5. Select **Resume** from the options list in the Manage Your High Availability Solution section.
- | 6. From the Resume High Availability Solution page, click **Run Now**. The interface performs the steps for resuming the high-availability solution. It restarts the nodes, cluster resource groups, and monitored resources, and if necessary, varies on the independent disk pool.
- | 7. When the high-availability solution is resumed, click **Close** to return to the Manage Your High Availability Solution page. The overall status of the high-availability solution is refreshed.

| **Results**

| **Related tasks**

| “Changing backup procedures for a high-availability solution” on page 13
| As part of configuring a high-availability solution with the High Availability Solution Manager interface, data is migrated to independent disk pools. To ensure that these objects continue to be backed up appropriately, you might need to change existing backup procedures.

Performing switchovers

The High Availability Solutions Manager graphical interface allows you to quickly and easily to perform administrative switchovers. A switchover allows you to change ownership from a production system to a backup system. You can perform switchovers for a variety of reasons, such as performing system maintenance.

About this task

Three of the four high-availability solutions use a two-node system, where you can switch over from the production system to the backup system, and back again. However, the Switched disk with cross-site mirroring solution is a three-node solution, which allows you to perform a three-way administrative switchover. You can switch over from the production system to a backup system at the same site, or to another backup system at a remote site, and then back again to the original production system.

To perform an administrative switchover of the high-availability solution, perform these steps:

Switchover for a two-node solution:

About this task

To perform an administrative switchover of the high availability solution, perform these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. In High Availability Solutions Manager, select **Manage your high availability solution**.
5. If you are using a two-node, high-availability solution, select **Switchover** from the dropdown list in the Manage Your High Availability Solution section.
6. From the Perform Administrative Switchover page, click **Run Now**. The High Availability Solutions Manager graphical interface performs the switchover to the backup system.
7. When the switchover completes, click **Close** to return to the Manage Your High Availability Solution page. The overall status of the high-availability solution is refreshed.

Switchover for a three-node solution:

About this task

If you are using the Switched disk with cross-site mirroring solution, you either switch over to another node at the same site or a different site:

Procedure

1. Select the type of switchover you want to perform:
 - a. Select **Switchover at same site** from the dropdown list. The High Availability Solutions Manager graphical interface performs the switchover to the backup system at the same site, or back again.
 - b. Select **Switchover to different site** from the dropdown list. The High Availability Solutions Manager graphical interface performs the switchover to the backup system at the remote site, or back again.
2. From the Perform Administrative Switchover page, click **Run Now**. The High Availability Solutions Manager graphical interface performs the switchover to the backup system.
3. When the switchover completes, click **Close** to return to the Manage Your High Availability Solution page. The overall status of the high-availability solution is refreshed.

Recovering from partition state

A cluster partition occurs in a cluster whenever communication is lost between one or more nodes in the cluster, and a failure of the lost nodes cannot be confirmed.

About this task

The High Availability Solutions Manager graphical interface allows you to quickly and easily change the partitioned node to failed and switch the data and applications to the first backup node. To recover from a cluster partition state, perform these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. In High Availability Solutions Manager, select **Manage your high availability solution**.

5. Select **Recover from partition state** from the dropdown list in the “Manage Your High Availability Solution” section.
6. From the Recover from a Partition State page, click **Run Now**. The High Availability Solutions Manager graphical interface performs the steps for recovering the node from the partition state.
7. When the cluster node is restarted, click **Close** to return to the Manage Your High Availability Solution page. The overall status of the high-availability solution is refreshed.

Collecting service information

If you are having a problem with your high-availability solution or simply want detailed information for your records, you can use High Availability Solution Manager graphical interface to collect or print service information.

About this task

To collect service information, follow these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. In High Availability Solutions Manager, select **Manage your high-availability solution**.
5. Select **Collect info** from the dropdown list in the Manage Your High Availability Solution section.
6. From the Collect Service Information page, click **Run Now**. The High Availability Solutions Manager graphical interface collects the service information for the high availability solution.
7. You can choose whether to print or save the service information.

Managing high availability resources

Use the High Availability Solutions Manager graphical interface to monitor and work with high availability resources. These resources include nodes, cluster resource groups, monitored resources, independent disk pools, TCP/IP interfaces, and high availability policies.

Managing nodes

Nodes are systems or logical partitions which have been defined within your high-availability solution. You can manage the nodes in your high-availability solution by using the High Availability Solutions Manager graphical interface.

About this task

To work with all nodes in your high-availability solution, follow these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. From the High Availability Solutions Manager welcome page, select **Manage Your High Availability Solution**.
5. Select the **Nodes** tab. You can perform any of the following functions on the nodes in the high-availability solution:
 - Monitor the status of nodes
 - Display or edit node properties

- Start a node
- Stop a node
- Work with all nodes

Example

The following are possible status values for nodes:

Table 2. Node status indicators

Icon	Status	Description
	Active	The node has been started by using a Create Cluster operation, an Add Cluster Node Entry operation, which started the cluster node; or with a Start Cluster Node operation. Cluster Resource Services is active on the node.
	Active Pending	The node is in the process of being started as the result of a Create Cluster operation, an Add Cluster Node Entry operation, which started the cluster node, or a Start Cluster Node operation. In addition, the node could have previously had a status of Partition and will change to the Active Pending status, as a result of the partitions being merged.
	Inactive Pending	Cluster Resource Services is in the process of ending on this node as the result of an end cluster node operation. The node is still in the cluster membership list.
	Remove Pending	The node is in the process of being removed from the cluster membership list as the result of a Remove Cluster Node Entry operation.
	New	A node has been added to the cluster membership list, but Cluster Resource Services has never been started on that node. The Cluster Resource Service data structures have not been created on the node. The Cluster Resource Service data structures are created only on the node running the Create Cluster operation.
	Inactive	Cluster Resource Services has been ended on the node as the result of an end cluster node operation. The node is still in the cluster membership list, but is no longer communicating with other nodes in the cluster.
	Failed	A previously active node has failed. A failure is either a system or clustering failure detected by Cluster Resource Services.
	Partition	The node is only communicating with a subset of the cluster because of a network failure detected by Cluster Resource Services, which has resulted in the loss of communications to one or more nodes in the cluster. When the partitioned nodes are merged back into a whole cluster, the node will change to Active status without operator intervention. Any node that had a status of Failed in any partition will still have a status of Failed after the merge.
	Unknown	This node is not an active member of the cluster, so the status of other nodes cannot be determined.

Related information

Cluster node

Starting nodes:

Use the High Availability Solutions Manager graphical interface to start nodes that are part of your high-availability solution.

About this task

To start a node within the high-availability solution, complete these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. From the High Availability Solutions Manager welcome page, select **Manage Your High Availability Solution**.
5. Select the **Nodes** tab.
6. Click the context menu next to the node and select **Start**.

Results

When the node starts, you are returned to the Manage Your High Availability Solution page. The node status information is refreshed.

Stopping nodes:

You can use the High Availability Solutions Manager graphical interface to work with high availability resources, such as nodes.

About this task

To stop a node within the high-availability solution, complete these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. From the High Availability Solutions Manager welcome page, select **Manage Your High Availability Solution**.
5. Select the **Nodes** tab.
6. Click the context menu next to the node and select **Stop**.

Results

When the node stops, you are returned to the Manage Your High Availability Solution page. The node status information is refreshed.

Working with all nodes:

You can work with all of the nodes in your high-availability solution by using the High Availability Solutions Manager graphical interface.

About this task

To work with all the nodes within the high-availability solution, complete these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. From the High Availability Solutions Manager welcome page, select **Manage Your High Availability Solution**.
5. Select the **Nodes** tab.
6. Click the context menu next to the node and select **Work with All Nodes...**

Results

This displays the Cluster Resource Services graphical interface from which you can work with all the nodes within the high-availability solution.

Displaying node properties:

Display node properties to view or edit information about the cluster node.

About this task

To display node properties, follow these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. From the High Availability Solutions Manager welcome page, select **Manage Your High Availability Solution**.
5. Select the **Nodes** tab.
6. Click the context menu next to the node and select **Properties**. The Node Properties page is shown.

Results

When you close the Nodes Properties page, you are returned to the Manage Your High Availability Solution page. The node status information is refreshed.

Managing cluster resource groups (CRG)

You can manage the CRGs in your high-availability solution by using the High Availability Solutions Manager graphical interface.

Before you begin

Cluster resource groups manage high-availability resources and define relationships within the high-availability solution.

About this task

To manage CRGs in your high-availability solution, follow these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.

4. From the High Availability Solutions Manager welcome page, select **Manage Your High Availability Solution**.
5. Select the **Cluster Resource Groups** tab. You can perform any of the following functions on the CRGs in the high availability solution:
 - Monitor the status of CRGs
 - Start a CRG
 - Stop a CRG
 - Delete a CRG
 - Work with all CRGs
 - Display or edit CRG properties

Example

The following are possible status values for CRGs:

Table 3. Cluster resource group status indicators

Icon	Status	Description
	Active	The resources managed by the cluster resource group are currently resilient.
	Inactive	The resources managed by the cluster resource group are currently not resilient.
	Indoubt	The information contained within the cluster resource group object might not be accurate. This status occurs when an exit program is called with an action of Undo and fails to complete successfully.
	Restored	The cluster resource group object was restored on this node and was not copied to the other nodes in the recovery domain. When Cluster Resource Services is started on this node, the cluster resource group will be synchronized with the other nodes in the recovery domain and its status set to Inactive .
	Add Pending	A new node is in the process of being added to the recovery domain of a cluster resource group. If the exit program is successful, the status is reset to its value at the time the API was called. If the exit program fails and the original state cannot be recovered, the status is set to Indoubt .
	Delete Pending	Cluster resource group is in the process of being deleted. When the exit program completes, the cluster resource group is deleted from all nodes in the recovery domain.
	Change Pending	The cluster resource group is in the process of being changed. If the exit program is successful, the status is reset to the value at the time the API was called. If the exit program fails and the original state cannot be recovered, the status is set to Indoubt .
	End Pending	Resiliency for the cluster resource group is in the process of ending. If the exit program is successful, the status is set to Inactive . If the exit program fails and the original state cannot be recovered, the status is set to Indoubt .
	Initialize Pending	A cluster resource group is being created and is in the process of being initialized. If the exit program is successful, the status is set to Inactive . If the exit program fails, the cluster resource group will be deleted from all nodes.
	Remove Pending	A node is in the process of being removed from the recovery domain of the cluster resource group. If the exit program is successful, the status is reset to the value at the time the API was called. If the exit program fails and the original state cannot be recovered, the status is set to Indoubt .

Table 3. Cluster resource group status indicators (continued)

Icon	Status	Description
	Start Pending	Resiliency is in the process of starting for the cluster resource group. If the exit program is successful, the status is set to Active . If the exit program fails and the original state cannot be recovered, the status is set to Indoubt . For peer cluster resource groups, all nodes defined with a role of peer are active access points for the cluster resources.
	Switchover Pending	The Initiate Switchover API was called, a failure of a cluster resource group occurred, or a node failed, causing a switchover or failover to begin. The first backup node is in the process of becoming the primary node. If the exit program is successful, the status is set to Active . If the exit program fails and the original state cannot be recovered, the status is set to Indoubt . While the switchover function is not valid for a peer cluster resource group, users may see the status Switchover pending during a node failure.

Related information

Cluster resource group

Starting a CRG:

You can start a cluster resource group (CRG) by using the High Availability Solutions Manager interface.

About this task

To start a CRG, complete these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. From the High Availability Solutions Manager welcome page, select **Manage Your High Availability Solution**.
5. Select the **Cluster Resource Groups** tab.
6. Click the context menu next to the CRG and select **Start**.

Results

When the CRG starts, you are returned to the Manage Your High Availability Solution page. The status information is refreshed.

Stopping a CRG:

You can stop a cluster resource group (CRG) by using the High Availability Solutions Manager interface.

About this task

To stop a CRG, complete these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. From the High Availability Solutions Manager welcome page, select **Manage Your High Availability Solution**.
5. Select the **Cluster Resource Groups** tab.
6. Click the context menu next to the CRG and select **Stop**.

Results

When the CRG stops, you are returned to the Manage Your High Availability Solution page. The status information is refreshed.

Deleting a CRG:

You can delete a cluster resource group (CRG) by using the High Availability Solutions Manager interface.

About this task

To delete a CRG, complete these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. From the High Availability Solutions Manager welcome page, select **Manage Your High Availability Solution**.
5. Select the **Cluster Resource Groups** tab.
6. Click the context menu next to the CRG and select **Delete**.

Results

When the CRG is deleted, you are returned to the Manage Your High Availability Solution page. The status information is refreshed.

Working with all CRGs:

You can work with all the CRGs that are defined in the high availability solution by using the High Availability Solutions Manager graphical interface.

About this task

To work with all the CRGs, complete these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. From the High Availability Solutions Manager welcome page, select **Manage Your High Availability Solution**.
5. Select the **Cluster Resource Groups** tab.

6. Click the context menu next to the CRG and select **Work with All CRGs**.

Results

This will display the Cluster Resource Services graphical interface from which you can work with all the CRGs within the high-availability solution.

Displaying CRG properties:

You can display CRG properties that are defined in the high availability solution by using the High Availability Solutions Manager interface.

About this task

To display the properties of a CRG, complete these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. From the High Availability Solutions Manager welcome page, select **Manage Your High Availability Solution**.
5. Select the **Cluster Resource Groups** tab.
6. Click the context menu next to the CRG and select **Properties** to display the properties for the selected CRG.

Managing monitored resources

You can manage monitored resources in your high-availability solution by using the High Availability Solutions Manager graphical interface.

Before you begin

Monitored resources are values associated with the operational environment for the high-availability solution. These resources are monitored throughout the high-availability solution. When they are changed on one node, those changes are propagated to other nodes in the high-availability solution. A cluster administrative domain manages the monitoring and synchronization of these resources, which eliminates the need to manually synchronize these resources throughout the environment.

About this task

To manage monitored resources in your high-availability solution, follow these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. From the High Availability Solutions Manager welcome page, select **Manage Your High Availability Solution**.
5. Select the **Monitored Resources** tab. You can perform any of the following functions on the monitored resources the high-availability solution:
 - Monitor the status of monitored resources

- Work with all monitored resources
- Display or edit properties

Example

The following table shows the possible status values for monitored resources:

Table 4. Global status indicators

Icon	Status	Description
	Consistent	The values for all the resource's attributes monitored by the system are the same on all active nodes within the cluster administrative domain.
	Inconsistent	The values for all the resource's attributes monitored by the system are not the same on all active nodes within the cluster administrative domain.
	Pending	The values of the monitored attributes are in the process of being synchronized across the cluster administrative domain.
	Added	The monitored resource entry has been added to the monitored resource directory in the cluster administrative domain but has not yet been synchronized.
	Ended	The monitored resource is in an unknown state because the cluster administrative domain has been ended, and changes to the resource are no longer being processed.
	Failed	The resource is no longer being monitored by the cluster administrative domain and the MRE should be removed. Certain resource actions are not recommended when a resource is being synchronized by a cluster administrative domain. If the resource represented by an MRE is a system object, it should not be deleted, renamed, or moved to a different library without removing the MRE first. If a resource is deleted, renamed or moved to a different library, the global status for the MRE is Failed and any changes made to the resource on any node after that are not propagated to any node in the cluster administrative domain.

Related information

Cluster administrative domain

Working with all monitored resources:

You can work with all monitored resources in your high availability solution by using the High Availability Solutions Manager graphical interface.

About this task

To work with all monitored resources, complete these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. From the High Availability Solutions Manager welcome page, select **Manage Your High Availability Solution**.
5. Select the **Monitored Resources** tab.
6. Click the context menu next to the monitored resource and select **Work with All Monitored Resources**. This will display the Cluster Resource Service interface from which you can work with all the monitored resource within the high-availability solution.

Displaying monitored resource properties:

You can display properties of monitored resources in your high availability solution by using the High Availability Solutions Manager graphical interface.

About this task

To display the properties of a monitored resource, complete these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. From the High Availability Solutions Manager welcome page, select **Manage Your High Availability Solution**.
5. Select the **Monitored Resources** tab.
6. Click the context menu next to the monitored resource and select **Properties**.

Managing independent disk pools

You can manage independent disk pools in your high availability solution by using the High Availability Solutions Manager graphical interface.

About this task

To manage independent disk pools in your high-availability solution, follow these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. From the High Availability Solutions Manager welcome page, select **Manage Your High Availability Solution**.
5. Select the **Independent disk pools** tab. You can perform any of the following functions on the independent disk pools in the high-availability solution:
 - Monitor the status of independent disk pools
 - Start mirroring
 - Stop mirroring
 - Work with all independent disk pools
 - Display or edit properties

Example

The Status column of the table contains an icon representing the status of the independent disk pool. See the following table for a list of icons and the status each represents. The following are possible status values for independent disk pools:

Table 5. Independent disk pool status indicators

Icon	Status
	Available.

Table 5. Independent disk pool status indicators (continued)

Icon	Status
	Available and mirroring is suspended with tracking.
	Available and mirroring is suspended without tracking.
	Available and mirroring is detached.
	Synchronization is in progress.
	Unavailable.
	Unavailable and mirroring is suspended with tracking.
	Unavailable and mirroring is suspended without tracking.
	Unavailable and mirroring is detached.
	Failed.

Related information

Independent disk pools

Starting mirroring:

You can start mirroring an independent disk pool in your high availability solution by using the High Availability Solutions Manager graphical interface.

About this task

To start mirroring, complete these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. From the High Availability Solutions Manager welcome page, select **Manage Your High Availability Solution**.
5. Select the **Independent disk pools** tab.
6. Click the context menu next to an independent disk pool and select **Start Mirroring**.

Stopping mirroring:

You can stop mirroring an independent disk pool in your high availability solution by using the High Availability Solutions Manager graphical interface.

About this task

To stop mirroring, complete these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. From the High Availability Solutions Manager welcome page, select **Manage Your High Availability Solution**.
5. Select the **Independent disk pools** tab.
6. Click the context menu next to an independent disk pool and select **Stop Mirroring**.

Working with all independent disk pools:

You can work with all independent disk pools in your high-availability solution by using the High Availability Solutions Manager graphical interface.

About this task

To work with all independent disk pools, complete these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. From the High Availability Solutions Manager welcome page, select **Manage Your High Availability Solution**.
5. Select the **Independent disk pools** tab.
6. Click the context menu next to an independent disk pool and select **Work with All Independent Disk Pools**. This will display the Disk Management graphical interface from which you can work with all the independent disk pools within the high-availability solution

Displaying independent disk pool properties:

You can work with all the independent disk pools in your high availability solution by using the High Availability Solutions Manager graphical interface.

About this task

To display independent disk pool properties, complete these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. From the High Availability Solutions Manager welcome page, select **Manage Your High Availability Solution**.
5. Select the **Independent disk pools** tab.
6. Click the context menu next to an independent disk pool and select **Properties**.

Managing TCP/IP interfaces

You can manage TCP/IP interfaces in your high-availability solution by using the High Availability Solutions Manager graphical interface.

About this task

To manage TCP/IP interfaces in your high-availability solution, follow these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. From the High Availability Solutions Manager welcome page, select **Manage Your High Availability Solution**.
5. Select the **TCP/IP Interfaces** tab. You can perform any of the following functions on the TCP/IP interfaces in the high availability solution:
 - Monitor the status of TCP/IP interfaces
 - Start TCP/IP interfaces
 - Stop TCP/IP interfaces
 - Work with all TCP/IP interfaces
 - Display or edit properties

Example

The following are possible status values for TCP/IP interfaces:

Table 6. TCP/IP interface status indicators

Icon	Status	Description
	Active	The interface has been started and is running.
	Inactive	The interface has not been started. The interface is not active.
	Starting	The system is processing the request to start this interface.
	Recovery Pending	An error with the physical line associated with this interface was detected by the system. The line description associated with this interface is in the recovery pending state.
	Recovery Canceled	A hardware failure has occurred. The line description associated with this interface is in the recovery canceled state.
	Failed	The line description associated with this interface has entered the failed state.
	Failed (TCP)	An error was detected in the IBM TCP/IP Vertical Licensed Internal Code.

Related information

TCP/IP Setup

Starting TCP/IP interfaces:

You can start TCP/IP interfaces in your high-availability solution by using the High Availability Solutions Manager graphical interface.

About this task

To start a TCP/IP interface, complete these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. From the High Availability Solutions Manager welcome page, select **Manage Your High Availability Solution**.
5. Select the **TCP/IP Interface** tab.
6. Click the context menu next to a TCP/IP Interface and select **Start TCP/IP Interface**.

Stopping TCP/IP interfaces:

When you select to stop TCP/IP interface, the High Availability Solutions Manager graphical interface responds based on the policy actions that have been defined for TCP/IP interfaces for this solution.

About this task

To stop a TCP/IP interface, complete these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. From the High Availability Solutions Manager welcome page, select **Manage Your High Availability Solution**.
5. Select the **TCP/IP Interface** tab.
6. Click the context menu next to a TCP/IP Interface and select **Stop TCP/IP Interface**.

Working with all TCP/IP interfaces:

You can work with all TCP/IP interfaces in your high availability solution by using the High Availability Solution Manager interface.

About this task

To work with all TCP/IP interfaces, complete these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. From the High Availability Solutions Manager welcome page, select **Manage Your High Availability Solution**.

5. Select the **TCP/IP Interface** tab.
6. Click the context menu next to a TCP/IP Interface and select **Work with All TCP/IP Interfaces**.

| **Displaying properties for TCP/IP interfaces:**

| You can display TCP/IP properties in your high-availability solution by using the High Availability Solutions Manager graphical interface.

| **About this task**

| To display or edit TCP/IP interface properties, complete these steps:

| **Procedure**

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. From the High Availability Solutions Manager welcome page, select **Manage Your High Availability Solution**.
5. Select the **TCP/IP Interface** tab.
6. Click the context menu next to a TCP/IP Interface and select **Properties**.

| **Managing policies**

You can manage policies in your high-availability solution by using the High Availability Solutions Manager graphical interface.

About this task

Policies define automatic responses to common actions that might occur within your high-availability environment. These policies control how the High Availability Solutions Manager graphical interface responds within your solution. These policies were initially created when you set up your high-availability solution. To change and manage policies in your high-availability solution, follow these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. From the High Availability Solutions Manager welcome page, select **Manage Your High Availability Solution**.
5. Select the **Policy** tab, and click **Edit** to change policies associated with your high-availability solution. You can work with the following policies for your high-availability solution:

Example

Table 7. Policies and options for high-availability solutions

Policy	Policy options
Action when a user profile is created	<ul style="list-style-type: none"> Automatically create the user profile on all other nodes in the high availability solution and add the monitored resource entry (MRE) to the administrative domain to ensure the user profile is synchronized on all the nodes. This is the default option. Take no action when a user profile is created.
Action when a user profile is deleted	<ul style="list-style-type: none"> Automatically remove the MRE from the administrative domain for the user profile. Do not delete the user profile on other nodes in the high availability solution. This is the default option. Automatically remove the MRE from the administrative domain for the user profile. Delete the user profile on all other nodes in the high availability solution. All objects owned by that user profile on all nodes is deleted. Automatically remove the MRE from the administrative domain for the user profile. Delete the user profile on all other nodes in the high availability solution. All objects owned by that user profile on all nodes is owned by the QDFTOWN user profile.
Action before the primary node enters restricted state	<ul style="list-style-type: none"> Shut down the high-availability solution without performing an administrative switchover. Vary off the independent disk pool so that all the data is made unavailable before entering restricted state. This is the default option. Shut down the high-availability solution without performing an administrative switchover. The independent disk pool and all the data it contains is available while in restricted state. Perform an administrative switchover of the high-availability solution from the primary node to an available backup node before entering restricted state on the primary node.
Action before the primary node performs a power down	<ul style="list-style-type: none"> Shut down the high-availability solution without performing an administrative switchover. The independent disk pool is varied off and all the data it contains is made unavailable before powering down the system. This is the default option. Perform an administrative switchover of the high-availability solution from the primary node to an available backup node before powering down the primary node.
Action when performing a failover to a backup node	<ul style="list-style-type: none"> If possible, perform a failover option to a backup node at the same site as the primary node. This is the default option. Perform a failover from the primary node to the next available node in the recovery domain of the device cluster resource group (CRG) associated with the high-availability solution.

Working with high availability event log

Use the event log in High Availability Solutions Manager graphical interface to view informational, warning, and error messages for your high-availability solution.

About this task

Each message in the event log on the Manage Your High Availability Solution page has a date and time stamp, severity level, and description. This table describes the severity level of messages that display in the event log.

Table 8. Event log status indicators

Severity	Description
	The log contains informational messages.
	The log contains warning messages and should be examined.

To work with the high availability event log, follow these steps:

Procedure

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 **High Availability Solutions Manager** from the IBM Systems Director Navigator for i5/OS navigation area.
4. From the High Availability Solutions Manager welcome page, select **Manage Your High Availability Solution**.
5. Review the event log to see if there are any new messages about events that are related to your high-availability solution.
6. If a warning or error message is displayed, attempt to correct the problem with the high-availability solution or resource.

Results

After you have corrected the problem with your solution or resource, an informational message indicates that your actions were successful.

Working with messages

The High Availability Solutions Manager graphical interface provides informational and error messages for your selected high-availability solution.

had000b

Dedicated Service Tools (DST) must be set up as follows:

- A service tools user ID and password for {0} must be configured.
- Service tools user IDs are configured using Dedicated Service Tools (DST).
- Service tools user ID and password must match the i5/OS user profile and password of the user that is calling this API.
- The i5/OS user profile password and the service tools user ID password must be all uppercase.
- Password level of the Dedicated Service Tools (DST) must be set. This is done through the service tools security data option, followed by the password level option. When this option has been activated, service tools user ID passwords are case sensitive.

hae0027

Communications could not be established with node {0}. The following are required on node {0} for communications to work.

- TCP server *INETD must be started
- IBM System i High Availability Solutions Manager (iHASM) licensed program number (5761-HAS) must be installed.
- IP address {1} must be active and reachable from the local system.

hae0028

The Share Memory Control (QSHRMEMCTL) system value must be set to a value of 1. In order to use the High Availability Solutions Manager graphical interface, the system value must be changed. Use the WRKSYSVAL command to change the value of QSHRMEMCTL to 1.

hai0006

Dedicated Service Tools (DST) must be set up for the QYHCHCOP API for the {0} user profile.

hat001b

Primary Node

hat001c

Backup Node

hat001d

Logical Partition

hat001e

System Serial Number

hat001f

Tower {0}

hat002c

V6R1 i5/OS

hat002d

Cluster IP Address {0}

hat002e

Line Description

hat002f

OptiConnect Enabled

hat0030

Tower contains required local communications hardware that will be switched with the tower.

hat0031

Present.

hat0032

Needed.

hat0033

Omit from solution.

hat0034

All unconfigured DASD under the IOP will be selected if any are selected.

hat0035

DASD.

hat0036

5761HAS - IBM System i High Availability Solutions Manager.

- hat0037**
Switchable IO Pool.
- hat0038**
Include in solution.
- hat0039**
Dependent Library {0} not allowed in independent disk pool.
- hat003a**
Tower contains configured DASD.
- hat003b**
Tower cannot be made switchable.
- hat003c**
Tower is private, may or may not be made switchable (unknown).
- hat003d**
IP address {0} not reachable.
- hat003e**
Cannot split up a parity set.
- hat003f**
Insufficient capacity on source node.
- hat0040**
Insufficient capacity on backup node.
- hat0041**
IOP
- hat0042**
All unconfigured DASD under the tower will be selected if any are selected.
- hat0044**
{0} contains object not supported in independent disk pool.
- hat0045**
{0} contains journal {1} to IFS objects.
- hat0046**
{0} contains job descriptions.
- hat0047**
{0} in use by System Value {1}.
- hat0048**
{0} in use by Network Attribute {1}.
- hat0049**
{0} contains journaled IFS objects.
- hat004a**
{0} contains class objects.
- hat004b**
{0} contains subsystem descriptions.
- hat004c**
{0} in use by subsystem description {1}.
- hat004d**
{0} contains database file {1} with data link control.

- hat004e**
{0} contains multisystem database file {1} (node group).
- hat0053**
DataPortIP {0}.
- hat0057**
Backup node at remote site.
- hat0058**
System Name.
- hat005a**
Configure independent disk pool.
- hat005b**
Configure geographic mirroring.
- hat005c**
Backup node at primary site.
- hat005d**
System value {0} must be set to {1}.
- hat005e**
Server takeover IP address.
- hat005f**
Tower contains non-switchable hardware.
- hat0060**
IOP {0} has configured DASD.
- hat0061**
IOP {0} contains communications hardware which will be switched with the IOP.
- hat0062**
Could not collect inventory for node {0}.
- hat0063**
Could not retrieve inventory for node {0}.
- hat0064**
Inventory collection processing step {0} of {1}.
- hat0065**
Tower{0} not owned by primary node.
- hat0066**
Tower {0} has configured DASD.
- hat0067**
Tower {0} has communications hardware or other non-switchable devices.
- hat0068**
Tower {0} is private.
- hat0069**
Tower {0} not visible from node {1}.
- hat006a**
IOP {0} has communications hardware or other non-switchable devices.
- hat006b**
The backup system for the primary site is not valid. It must be another partition of the same system.

hat006c

The remote backup system is not valid. It must be another physical system.

hat006d

IP address {1} not found.

hat006e

Current selection {1} not valid.

hat006f

The setup of your high-availability solution is now complete. If your normal security policy is to have the QSECOFR profile disabled and you enabled the QSECOFR profile to set up your high-availability solution, then you should return to your normal security policy and disable the QSECOFR profile.

hat0070

IOP {0} not owned by primary node.

hat0071

Licensed program product {1}

Related tasks

Enabling and accessing disk units

Related information

Share Memory Control (QSHRMEMCTL)

Reference: Panel help for the solution-based interface

This topic collection includes the panel help for the High Availability Solutions Manager graphical interface. You also can display the panel help by clicking the question mark ? icon on the page.

High Availability Solutions Manager welcome

High Availability Solutions Manager is an easy-to-use interface that helps you select, set up, and manage your high-availability solution. High Availability Solutions Manager is included in the IBM System i High Availability Solutions Manager licensed program, 5761-HAS.

Setting up your high-availability solution involves migrating your business applications to run in a high-availability environment. This environment includes an independent disk pool and two or more instances of i5/OS running V6R1. Setting up a high-availability solution results in the data, applications, and environment from one system replacing the contents of each backup system in the high-availability solution.

You need to perform the five tasks on the High Availability Solutions Manager welcome page in sequence. The only task that you can skip is the first one, See how IBM high availability solutions work. When you are finished using this panel, click **Close**.

See how IBM high availability solutions work

Displays a flash demo that introduces high availability concepts and explains how your enterprise can benefit from using the High Availability Solutions Manager interface to easily set up and manage a highly-available solution. The flash demo is always available.

Select your high-availability solution

Allows you to select from one of four high-availability solutions:

- Switched disk between logical partitions
- Switched disk between systems
- Switched disk with geographic mirroring
- Cross-site mirroring with geographic mirroring

The Select your high availability solution task is available until you start setting up the solution.

Verify requirements before setting up your high-availability solution

Verifies that your systems contain the required hardware and software resources for setting up the selected high-availability solution. This task is available only if you have selected a high-availability solution, and have not yet started setting it up. If your systems do not meet the requirements for a particular high-availability solution, you can either add the required resources or select a different high-availability solution.

Set up your high-availability solution (license required)

Sets up the selected high-availability solution automatically, step-by-step. This task is available only if you have verified the high availability requirements.

During the set up of the high-availability environment, the node from where you are running the setup becomes the primary node of the high-availability environment. The other node or nodes in your solution become the backup nodes. Because the setup of the high-availability environment deletes data from the backup node, save important data that exists on the backup node before beginning the setup step.

You must have a license for the IBM System i High Availability Solutions Manager, 5761-HAS licensed program to be able to set up the high-availability solution. You have a 70-day trial period in which you can use the High Availability Solutions Manager. When the trial period ends, you will not be able to use the High Availability Solutions Manager unless you purchase a license key through your standard software ordering process. Any high-availability solution that you have designed and set up during the trial period will be reactivated when the license key has been purchased.

Manage your high-availability solution (license required)

Allows you to manage your high-availability solution. The Manage your high-availability solution task is available after you have successfully set up the high-availability solution. You can monitor the status of your high-availability solution, perform an administrative switchover, shut down and resume the high-availability solution, and perform many other tasks.

You must have a license for the IBM System i High Availability Solutions Manager, 5761-HAS, to be able to manage the high-availability solution.

Terms and Concepts

Displays the *High availability technologies* PDF, which provides the infrastructure and tools to help achieve your availability goals. The key technologies in i5/OS high-availability solutions are cluster resource services (clusters) and independent disk pools.

High Availability Operator's Guide

Displays the Implementing high availability with the High Availability Solutions Manager – Operator's guide PDF, which includes all of the tasks for working with your high-availability solution.

An icon next to the task number indicates the status of the high availability tasks.

Icon	Description
	Indicates that the task has been completed. If this icon appears next to the Manage your high-availability solution task, it indicates that the high-availability solution is working properly.
	Indicates that the task is incomplete.
	Indicates that there is a problem with one or more resources of the high-availability solution resources. This icon is used only for the Manage your high-availability solution task.

Select your high availability solution

The Select Your High Availability Solution page describes the predefined high-availability solutions available in High Availability Solutions Manager. Select a high-availability solution from the list on the left to see a brief description of the solution.

For each high-availability solution, you can view an overview of the solution, a customized shopping list, and an executive summary.

You can perform the following actions from this page:

View solution overview

| Click to display a flash demo of the selected high-availability solution. You can save this
| information.

View customized shopping list

| Click to see a customized list of hardware and software products that are required to set up the
| high-availability solution. You can save or print this information.

View executive summary

| Click to see a list of the benefits and limitations of the selected high-availability solution. Use this
| information to help you decide which solution to select. You can save or print this information.

Select Select the desired high-availability solution from the list and click **Select** to identify the solution to be configured and deployed. (*Deploy* means to place files or install software into an operational environment.)

Cancel

Click to return to the High Availability Solutions Manager welcome page.

If these high-availability solutions do not meet your needs, contact your IBM representative for a customized solution, or use the Cluster Resource Services and Disk Management graphical interfaces, or use the IBM System i High Availability Solutions Manager (iHASM) commands and APIs to design your own high-availability solution.

For more information, see these information center topics:

- “Selecting a high availability solution” on page 6
- Switched disk between logical partitions
- Switched disk between systems
- Switched disk with geographic mirroring
- Cross-site mirroring with geographic mirroring

Customized shopping list

The Customized Shopping List page in High Availability Solutions Manager displays a list of the minimum requirements for your selected high-availability solution. Before configuring a solution, you need to ensure that all the minimum requirements have been met.

| Each high-availability solution listed on the Select Your High Availability Solution page has its own
| customized shopping list. If your system does not meet the requirements for a particular high-availability
| solution, look at the shopping list for one of the other solutions. You also can save or print the list of
| requirements.

The system looks for the hardware devices and software licenses that are required to configure the selected high-availability solution. While retrieving this information, the inventory collection animation is displayed. An image of the selected high-availability solution is displayed in the upper right corner of the page. You cannot add extra resources to the minimum required list.

The **Requirements** column describes the hardware or software components that are required for the selected high-availability solution. This column also includes a link to a specific article or Web site with the technical details about the selected requirement.

The **Status** column indicates the following:

Table 9. Verify requirement status

Status	Description
	The requirement must be met to set up your high availability solution. You can meet a requirement by supplying missing identification information for resources which will be used later to set up your high-availability solution or by installing missing hardware and software. After entering or selecting missing information click OK . After installing missing hardware or software, click Refresh so the requirements list can be updated to reflect the changes.
	The requirement is optional but might be recommended based on your specific business requirements. For example, a minimum of one disk drive is required to configure an independent disk pool but one disk might not be sufficient to store all of your data. You can meet a requirement by supplying missing identification information for resources which will be used later to set up your high-availability solution or by installing missing hardware and software. After entering or selecting missing information click OK . After installing missing hardware or software click Refresh so the requirements list can be updated to reflect the changes.
	The requirement has been met.

The **Information** column lists the valid choices that meet the requirement. You can select the value that meets the requirement from the choices listed in the combo box. If no valid choices are listed, or if an existing resource is listed as not valid, check the message area for reasons why existing resources cannot be selected.

After all required hardware, software and information has been identified and successfully verified, you can set up your solution.

You also can perform the following actions on the requirements list:

Action	Description
Print	Prints the list of requirements. This action brings up the Print High Availability Solution Information page.
Refresh	Refreshes the list of requirements with the latest data.
Save	Saves this list of requirements to a file. This action brings up the Save High Availability Solution Information page.
OK	After you have finished viewing the list of requirements, click OK to save the current list and return to the previous page.
Cancel	Discards any changes that you have made, closes this page, and returns to the previous page.

Messages

You might see the following informational and error messages for your selected high-availability solution:

had000b

Dedicated Service Tools (DST) must be set up as follows:

- A service tools user ID and password for {0} must be configured.
- Service tools user IDs are configured using Dedicated Service Tools (DST).
- Service tools user ID and password must match the i5/OS user profile and password of the user that is calling this API.
- The i5/OS user profile password and the service tools user ID password must be all uppercase.
- Password level of the Dedicated Service Tools (DST) must be set. This is done through the service tools security data option, followed by the password level option. When this option has been activated, service tools user ID passwords are case sensitive.

hae0027

Communications could not be established with node {0}. The following are required on node {0} for communications to work.

- TCP server *INETD must be started
- IBM System i High Availability Solutions Manager (iHASM) licensed program number (5761-HAS) must be installed.
- IP address {1} must be active and reachable from the local system.

hai0006

Dedicated Service Tools (DST) must be set up for the QYHCHCOP API for the {0} user profile.

hat001b

Primary Node

hat001c

Backup Node

hat001d

Logical Partition

hat001e

System Serial Number

hat001f

Tower {0}

hat002c

V6R1 i5/OS

hat002d

Cluster IP Address {0}

hat002e

Line Description

hat002f

OptiConnect Enabled

hat0030

Tower contains required local communications hardware that will be switched with the tower.

hat0031

Present.

hat0032

Needed.

hat0033

Omit from solution.

hat0034

All unconfigured DASD under the IOP will be selected if any are selected.

- hat0035**
DASD.
- hat0036**
5761HAS - IBM System i High Availability Solutions Manager.
- hat0037**
Switchable IO Pool.
- hat0038**
Include in solution.
- hat0039**
Dependent Library {0} not allowed in independent disk pool.
- hat003a**
Tower contains configured DASD.
- hat003b**
Tower cannot be made switchable.
- hat003c**
Tower is private, may or may not be made switchable (unknown).
- hat003d**
IP address {0} not reachable.
- hat003e**
Cannot split up a parity set.
- hat003f**
Insufficient capacity on source node.
- hat0040**
Insufficient capacity on backup node.
- hat0041**
IOP
- hat0042**
All unconfigured DASD under the tower will be selected if any are selected.
- hat0044**
{0} contains object not supported in independent disk pool.
- hat0045**
{0} contains journal {1} to IFS objects.
- hat0046**
{0} contains job descriptions.
- hat0047**
{0} in use by System Value {1}.
- hat0048**
{0} in use by Network Attribute {1}.
- hat0049**
{0} contains journaled IFS objects.
- hat004a**
{0} contains class objects.
- hat004b**
{0} contains subsystem descriptions.

- hat004c**
{0} in use by subsystem description {1}.
- hat004d**
{0} contains database file {1} with data link control.
- hat004e**
{0} contains multisystem database file {1} (node group).
- hat0053**
DataPortIP {0}.
- hat0057**
Backup node at remote site.
- hat0058**
System Name.
- hat005c**
Backup node at primary site.
- hat005d**
System value {0} must be set to {1}.
- hat005e**
Server takeover IP address.
- hat005f**
Tower contains non-switchable hardware.
- hat0060**
IOP {0} has configured DASD.
- hat0061**
IOP {0} contains communications hardware which will be switched with the IOP.
- hat0062**
Could not collect inventory for node {0}.
- hat0063**
Could not retrieve inventory for node {0}.
- hat0064**
Inventory collection processing step {0} of {1}.
- hat0065**
Tower{0} not owned by primary node.
- hat0066**
Tower {0} has configured DASD.
- hat0067**
Tower {0} has communications hardware or other non-switchable devices.
- hat0068**
Tower {0} is private.
- hat0069**
Tower {0} not visible from node {1}.
- hat006a**
IOP {0} has communications hardware or other non-switchable devices.
- hat006b**
The backup system for the primary site is not valid. It must be another partition of the same system.

hat006c

The remote backup system is not valid. It must be another physical system.

hat006d

IP address {1} not found.

hat006e

Current selection {1} not valid.

hat0071

Licensed program product {1}

For the complete list of messages, see the “Working with messages” on page 37 information center topic.

Executive summary

The Executive Summary page in High Availability Solutions Manager describes the benefits and restrictions for the selected high availability solution. Use this information to help your company decide which solution to set up.

The following table describes the actions that you can perform on this page.

Action	Result
More details about this high availability solution	Displays additional information about the selected high availability solution.
Print	Prints the executive summary. This action brings up the Print High Availability Solution Information page.
Save	Saves the executive summary to a file. This action brings up the Save High Availability Solution Information page.
Close	Closes the Executive Summary panel and returns to the Select Your High Availability Solution page.

Save or print high availability solution information

You can use the Save or Print High Availability Solutions Information page in High Availability Solutions Manager to save or print information regarding a high-availability solution.

The **Solution Overview (flash demo)** information source is a visual demonstration of how the High Availability Solutions Manager graphical interface works. The **Solution Overview** can only be saved. There is no printable format.

The **Customized Shopping List** information source is a list of hardware and software items needed to configure the selected high-availability solution.

The **Executive Summary** information source is a list of benefits and restrictions that the selected high-availability solution provides.

The **Display Log** information source lists the results of performing a step during the setup of the selected high-availability solution.

The information option that is selected by default on this page depends on which function requested the action. For example, if you clicked **Save** on the Executive Summary page, the **Executive Summary** information source is selected on this page.

- To save or print additional information sources, check the box to the left of the information source.
- Clicking on a checked box deselects the information source.

To select all of the options available, click the **Select All** button above the information source table, or select the **Select All** menu option in the **Select Action** dropdown menu and click **Go**.

To uncheck all checked boxes, click the **Deselect All** button (located next to the **Select All** button), or select the **Deselect All** menu option from the **Select Action** dropdown menu and click **Go**.

- Click **OK** to save or print the selected information. Click **Cancel** to return to the previous page.

The following table describes what happens after clicking **OK**:

Action	Result
Save	<p>The following dialog boxes might vary between browsers. This example is based on an Internet Explorer browser.</p> <ul style="list-style-type: none"> The File Download dialog box appears. Options are Open, Save, or Cancel. Choosing Save in the File Download dialog box brings up a Save As dialog box. Select the location to save the file and click Save. The Download Complete dialog appears when the save is complete. Options are Open, Open Folder, or Close.
Print	Launches the standard Print dialog box.

Verify requirements list

- The Verify Requirements List page in High Availability Solutions Manager displays a list of the minimum requirements for your selected high-availability solution. Before configuring a solution, you need to ensure that all the minimum requirements have been met. You can save or print the requirements list.

The system looks for the hardware devices and software licenses that are required to configure the selected high-availability solution. While retrieving this information, the inventory collection animation is displayed. An image of the selected high-availability solution is displayed in the upper right corner of the page. You cannot add extra resources to the minimum required list.

The **Requirements** column describes the hardware or software components that are required for the selected high-availability solution. This column also includes a link to a specific article or Web site with the technical details about the selected requirement.

The **Status** column indicates the following:

Table 10. Verify requirement status

Status	Description
	<p>The requirement must be met to set up your high availability solution. You can meet a requirement by supplying missing identification information for resources which will be used later to set up your high-availability solution or by installing missing hardware and software.</p> <p>After entering or selecting missing information click OK. After installing missing hardware or software, click Refresh so the requirements list can be updated to reflect the changes.</p>

Table 10. Verify requirement status (continued)

Status	Description
	<p>The requirement is optional but might be recommended based on your specific business requirements. For example, a minimum of one disk drive is required to configure an independent disk pool but one disk might not be sufficient to store all of your data.</p> <p>You can meet a requirement by supplying missing identification information for resources which will be used later to set up your high-availability solution or by installing missing hardware and software.</p> <p>After entering or selecting missing information click OK. After installing missing hardware or software click Refresh so the requirements list can be updated to reflect the changes.</p>
	The requirement has been met.

The **Information** column lists the valid choices that meet the requirement. You can select the value that meets the requirement from the choices listed in the combo box. If no valid choices are listed, or if an existing resource is listed as not valid, check the message area for reasons why existing resources cannot be selected.

After all required hardware, software and information has been identified and successfully verified, you can set up your solution.

You also can perform the following actions on the requirements list:

Action	Description
Print	Prints the list of requirements. This action brings up the Print High Availability Solution Information page.
Refresh	Refreshes the list of requirements with the latest data.
Save	Saves the list of requirements to a file. This action brings up the Save High Availability Solution Information page.
OK	After you have finished viewing the list of requirements, click OK to save the current list and return to the previous page. OK is enabled only when all mandatory requirements have been met. You are now ready to set up your high-availability solution.
Cancel	Discards any changes that you have made, closes this page, and returns to the previous page.

For more information, see the “Verifying requirements for your high availability solution” on page 10 information center topic.

Messages

You might see the following informational and error messages for your selected high-availability solution:

had000b

Dedicated Service Tools (DST) must be set up as follows:

- A service tools user ID and password for {0} must be configured.
- Service tools user IDs are configured using Dedicated Service Tools (DST).
- Service tools user ID and password must match the i5/OS user profile and password of the user that is calling this API.
- The i5/OS user profile password and the service tools user ID password must be all uppercase.

- Password level of the Dedicated Service Tools (DST) must be set. This is done through the service tools security data option, followed by the password level option. When this option has been activated, service tools user ID passwords are case sensitive.

hae0027

Communications could not be established with node {0}. The following are required on node {0} for communications to work.

- TCP server *INETD must be started
- IBM System i High Availability Solutions Manager (iHASM) licensed program number (5761-HAS) must be installed.
- IP address {1} must be active and reachable from the local system.

hai0006

Dedicated Service Tools (DST) must be set up for the QYHCHCOP API for the {0} user profile.

hat001b

Primary Node

hat001c

Backup Node

hat001d

Logical Partition

hat001e

System Serial Number

hat001f

Tower {0}

hat002c

V6R1 i5/OS

hat002d

Cluster IP Address {0}

hat002e

Line Description

hat002f

OptiConnect Enabled

hat0030

Tower contains required local communications hardware that will be switched with the tower.

hat0031

Present.

hat0032

Needed.

hat0033

Omit from solution.

hat0034

All unconfigured DASD under the IOP will be selected if any are selected.

hat0035

DASD.

hat0036

5761HAS - IBM System i High Availability Solutions Manager.

hat0037

Switchable IO Pool.

- hat0038**
Include in solution.
- hat0039**
Dependent Library {0} not allowed in independent disk pool.
- hat003a**
Tower contains configured DASD.
- hat003b**
Tower cannot be made switchable.
- hat003c**
Tower is private, may or may not be made switchable (unknown).
- hat003d**
IP address {0} not reachable.
- hat003e**
Cannot split up a parity set.
- hat003f**
Insufficient capacity on source node.
- hat0040**
Insufficient capacity on backup node.
- hat0041**
IOP
- hat0042**
All unconfigured DASD under the tower will be selected if any are selected.
- hat0044**
{0} contains object not supported in independent disk pool.
- hat0045**
{0} contains journal {1} to IFS objects.
- hat0046**
{0} contains job descriptions.
- hat0047**
{0} in use by System Value {1}.
- hat0048**
{0} in use by Network Attribute {1}.
- hat0049**
{0} contains journaled IFS objects.
- hat004a**
{0} contains class objects.
- hat004b**
{0} contains subsystem descriptions.
- hat004c**
{0} in use by subsystem description {1}.
- hat004d**
{0} contains database file {1} with data link control.
- hat004e**
{0} contains multisystem database file {1} (node group).

hat0053
DataPortIP {0}.

hat0057
Backup node at remote site.

hat0058
System Name.

hat005c
Backup node at primary site.

hat005d
System value {0} must be set to {1}.

hat005e
Server takeover IP address.

hat005f
Tower contains non-switchable hardware.

hat0060
IOP {0} has configured DASD.

hat0061
IOP {0} contains communications hardware which will be switched with the IOP.

hat0062
Could not collect inventory for node {0}.

hat0063
Could not retrieve inventory for node {0}.

hat0064
Inventory collection processing step {0} of {1}.

hat0065
Tower{0} not owned by primary node.

hat0066
Tower {0} has configured DASD.

hat0067
Tower {0} has communications hardware or other non-switchable devices.

hat0068
Tower {0} is private.

hat0069
Tower {0} not visible from node {1}.

hat006a
IOP {0} has communications hardware or other non-switchable devices.

hat006b
The backup system for the primary site is not valid. It must be another partition of the same system.

hat006c
The remote backup system is not valid. It must be another physical system.

hat006d
IP address {1} not found.

hat006e
Current selection {1} not valid.

hat0071

Licensed program product {1}

For the complete list of messages, see the “Working with messages” on page 37 information center topic.

Set up your high availability solution

Use the Set Up Your High Availability Solution page in High Availability Solutions Manager to set up the selected high-availability solution. After you have verified that your data has been backed up, you are signed in as QSECOFR, and that you are running the setup from the node that you want to be the primary node in the high-availability solution, click **Go** to start setting up your solution.

Before you can set up the high-availability solution, your system must meet the following requirements:

- Your system must have the IBM System i High Availability Solutions Manager (5761-HAS) license key installed or be within the 70-day trial period.

If the trial period is running, the High Availability Solutions Manager Trial Period message displays. If the trial period has expired, the High Availability Solutions Manager Trial Period Ended message displays, and you will not be able to continue setting up the solution.

- Your systems must be in dedicated state. This means that you are signed on as QSECOFR, no user jobs are running, and no users are signed on to the system.

If you are not signed in as QSECOFR, the user authority error message (HAI007E) displays. Click **All the systems involved in the solution must be in dedicated state** during solution setup for more details.

Set up your high-availability solution table

This table lists the steps for configuring your selected high-availability solution. Each configuration step contains a substep page. You must complete all configuration steps to set up the solution. However, you can stop in the middle of setup and return to setup later on. When you click **Go** to start the current step, the substep page for that step displays. After you complete the current step, you are returned to this page. The **Status** column indicates whether the step was completed successfully. Continue with each setup step until you have completed all steps.

Column	Description
Step	The configuration step name. The  arrow points to the current step. When you complete a step, the arrow moves to the next step.
Estimated Time	The estimated time, in minutes, to complete the current configuration step.
Actual Time	The length of time, in minutes, that the configuration step took to complete.
Status	The status of the current step. This column is blank for steps that have not yet been started. When a configuration step is completed, the status column is set as follows: Complete - the step has completed successfully. Failed - the step failed. Running - the step is running. Undoing - the step is being undone. You can select Display log , Retry , or Undo step from the context menu  for each completed, failed, or incomplete step.

To set up your selected high-availability solution, follow these steps. When each step is completed, the Status column indicates whether the step completed successfully, failed, or was incomplete. If the step completed successfully, the  arrow moves to the next step.

1. Click **Go** to perform the **Set up high availability policies** step.
 - a. On the Set Up High Availability Policies page, select the policies that you want to use and click **OK**.
 - b. When this step is completed, the Status column indicates whether the step completed successfully, failed, or was incomplete. If the step completed successfully, the  arrow moves to the next step.
2. Click **Go** to perform the **Set up high availability environment** step.
 - a. On the Set Up High Availability Environment page, click **Run Now**.
 - b. When this step is completed, click **Close** to return to this page.
3. Click **Go** to perform the **Verify administrative switchover from {1} to {2}** step, where {1} represents the primary node for your solution and {2} represents the backup node for your solution. (If you are setting up a three-node solution, this switchover is from the primary node to the local backup node.)
 - a. On the Verify Administrative Switchover page, click **Run Now**.
 - b. When this step is completed, click **Close** to return to this page.
 - c. Manually verify that the administrative switchover was successful.
4. Click **Go** to perform the **Verify administrative switchover from {2} to {3}** step, where {2} represents the backup node for your solution and {3} represents the primary node for your solution. (If you are setting up a three-node solution, this switchover is from the local backup node to the remote backup node.)
 - a. On the Verify Administrative Switchover page, click **Run Now**.
 - b. When this step is completed, click **Close** to return to this page.
 - c. Manually verify that the administrative switchover was successful.
5. Click **Go** to perform the **Verify administrative switchover from {3} to {1}** step. This step displays only if you are setting up a three-node solution. This switchover is from the remote backup node to the primary node.
6. Click **Go** to perform the **Migrate user profiles** step.
 - a. On the Migrate User Profiles page, select the profiles that you want to migrate to the high-availability solution.
 - b. When this step is completed, click **Close** to return to this page.
7. Click **Go** to perform the **Migrate libraries** step.
 - a. On the Migrate Libraries page, select the libraries that you want to migrate to the high-availability solution.
 - b. When this step is completed, click **Close** to return to this page.
8. Click **Go** to perform the **Migrate directories** step.
 - a. On the Migrate Directories page, select the directories that you want to migrate to the high-availability solution.
 - b. When this step is completed, click **Close** to return to this page.
9. Click **Go** to perform the **Verify administrative switchover from {1} to {2}** step, where {1} represents the primary node for your solution and {2} represents the backup node for your solution. (If you are setting up a three-node solution, this switchover is from the primary node to the local backup node.)
 - a. On the Verify Administrative Switchover page, click **Run Now**.
 - b. When this step is completed, click **Close** to return to this page.
 - c. Manually verify that the administrative switchover was successful.

Note: Before completing the setup of your high-availability solution, you must make sure that your business applications will work appropriately on all nodes in your solution. This step involves switching the nodes on which your applications will be running. When all the nodes in the solution have been verified, you can proceed to complete setting up the high-availability solution.

10. Click **Go** to perform the **Verify administrative switchover from {2} to {3}** step, where {2} represents the backup node for your solution and {3} represents the primary node for your solution. (If you are setting up a three-node solution, this switchover is from the local backup node to the remote backup node.)
 - a. On the Verify Administrative Switchover page, click **Run Now**.
 - b. When this step is completed, click **Close** to return to this page.
 - c. Manually verify that the administrative switchover was successful.
11. Click **Go** to perform the **Verify administrative switchover from {3} to {1}** step. This step displays only if you are setting up a three-node solution. This switchover is from the remote backup node to the primary node.
12. Click **Go** to perform the **Finish set up and clean up work files** step.
 - a. On the Finish Set Up and Clean Up Work Files page, click **Run Now**.
 - b. When this step is completed, click **Close** to return to this page.
13. Click **Close** on the Set Up Your High Availability Solution page. You have finished setting up your high-availability solution. Your high-availability solution is now active and ready to be managed.

Actions

You can perform the following actions while setting up the high-availability solution:

Action	Result
Close	To return to the High Availability Solutions Manager page, click Close . The Set up high availability solution step is marked as completed only if all configuration steps are completed. Otherwise, it is marked as incomplete.
Display log	To display the log that was created while performing the configuration step, select Display Log from the context menu  for the step. This log contains the commands that were run, the parameters, and the results. For example, if a setup step fails, display the log to read the error messages and troubleshoot the problem.
Go	To display the substep page for the current configuration step, click Go . This action is available until all configuration steps are completed, unless the current step has failed.
Retry	To perform the current configuration step again, click Retry . You also can select Retry from the context menu  for a step. This action is available only when the current step has failed, or when the step can be run any time. For example, you can set the high availability policies or verify administrative switchover at any time.
Undo previous step	To perform a rollback to all changes made in the previous configuration step, click Undo previous step . You also can select Undo step from the context menu  for a step. The previous step is highlighted and ready to be configured. This action is available after completing the first configuration step.

Messages

During the setup of the high-availability solution, the following messages might display:

Backup procedure warning

Displays before every step which migrates objects to the high-availability solution.

Dedicated state

Displays before each major setup step if the system is not in dedicated state.

High Availability Solutions Manager trial period

Displays if you do not have an iHASM license.

High Availability Solutions Manager trial period ended

Displays if the 70-day trial period has expired.

Migration error

Displays if you try to migrate more data than the independent disk pool (IASP) can hold. You can either add more disk units to the independent disk pool or migrate less data to the independent disk pool.

Step completed

Displays after every major setup step. You should verify that your applications still work properly after each set up step.

User authority error

Displays if you are not signed in as QSECOFR.

For more information, see the following information center topics:

- “Changing backup procedures for a high-availability solution” on page 13
- “Placing high availability systems in dedicated state” on page 13
- “Verifying applications after each setup step” on page 14

Display high availability log

The Display Log page in High Availability Solutions Manager lists the commands, APIs, and other system functions that were used to perform a step during the setup of your high-availability solution. The log also includes completion and error messages.

After you finish running a setup step, select **Display log** from the context menu  for that step. For example, if a setup step fails, display the log to read the error messages and troubleshoot the problem.

The most recent completion or error messages are appended to the bottom of the log.

Messages sent to the log during a setup step are never removed, even if you perform an undo or retry of the step.

You can perform the following actions on this page:

Action	Result
Print	Click to print the log.
Save	Click to save the log to a file.
Close	Click to return to the Set Up Your High Availability Solution page.

Set up high availability policies

| The Set Up High Availability Policies page in High Availability Solutions Manager asks you to choose
| default actions for various events that take place on the system after the high-availability solution is set
| up. The default action is automatically taken when the specified event occurs.

You can set up policies for the following default actions. Select one option from each section. To change the selected option in a section, select a different radio button in that section:

| Default action when a user profile is created

| Having user profiles automatically created on the nodes in the high availability solution can
| simplify system maintenance because you create only one profile for all the nodes.

| **Default action when a user profile is deleted**

| Having user profiles automatically deleted on the nodes in the high availability solution can simplify system maintenance because you delete only one profile for all the nodes.

| **Default action before the primary node enters the restricted state**

| With this action, you can keep the independent disk pool data available, either on the current primary node or on a backup node, using an administrative switchover.

| **Default action before the primary node performs a power-off operation**

| With this action, you can keep the independent disk pool data available on a backup node, using an administrative switchover.

| **Default action when performing a failover to a backup node**

| This policy is available only if the selected solution is Switched Disk with Geographic Mirroring.

After you have made the selections, click **OK** to save the selected options and return to the Set Up High Availability Solution page.

Click **Cancel** to discard the changes and return to the Set Up High Availability Solution page.

For more information, see the “Managing policies” on page 35 information center topic.

Set up high availability environment

The Set Up High Availability Environment page in High Availability Solutions Manager lists the substeps to be performed to complete the setup of the high-availability environment. This step uses the data that you provided earlier to configure each system to be part of the high-availability environment.

Important: You must be signed on using the QSECOFR user profile. The password for QSECOFR must be the same across all the nodes to be added to the high-availability solution. The password for the dedicated service tools (DST) QSECOFR must be the same as for the QSECOFR user profile, but in upper case.

The system on which the High Availability Solutions Manager web interface is running becomes the primary node of the high-availability solution. If this web interface is not running on the system that is designated as the primary node, you must close this browser window and open a new browser window on the system that will become the primary node before continuing with the setup of your high-availability solution.

Note: Verify that the TCP/IP *INETD server is started on all systems that will be involved in the deployment of the high-availability solution before beginning the deployment.

Field information

The substep table has the following columns.

Column	Description
Substep	Short description of the command or API that will be run.
Estimated time	Estimation of how much time the substep will take to complete. The estimated time changes dynamically to show the best estimate of the remaining time.
Actual time	The actual amount of time that has elapsed since the substep began running.
Status	The status of the substep. Possible values include: Blank - The substep has not yet been run or was undone. Complete - The substep has completed without error. Failed - The substep encountered errors while running. Running - The substep is currently running.
Command/API	The command or API that is currently being processed to complete the substep.

Completion and error messages that are received as the substeps are run are shown in the message area above the **Run Now**, **Cancel**, **Undo**, and **Close** buttons. You can attempt the task again after taking corrective action based on the error messages provided.

Click **Run Now** to begin the substep that the  arrow is pointing to.

Click **Cancel** to return to the Set Up High Availability Solution page without setting up the high-availability environment.

While the substeps are running, click **Cancel** to stop after the current substep has finished running, but before the next substep starts running. After the cancel processing is complete, the following options are available:

- Click **Undo** to undo the previously completed substeps.
- Click **Run Now** to continue running with the next substep.
- Click **Cancel** to return to the previous page without continuing or undoing any substeps. When setting up the high-availability solution, future set up processing will remember which substep is the next substep to be run.

Click **Close** to return to the Set Up High Availability Solution page after all substeps have finished running.

Substeps performed to set up the high-availability environment

High Availability Solutions Manager performs the following substeps to set up the high-availability environment:

- The Allow Add to Cluster Network attribute is changed to *ANY on all nodes in the high-availability solution.
- The cluster is created on all systems (nodes) that were previously specified.
- Device domain entries are added to the cluster.
- The independent disk pool device description is created.
- The Allow Server Security Information to be Retained (QRETSVRSEC) system value is changed to 1 for user profiles and other object types that need to be treated as secure.
- A cluster administrative domain is created and started.
- Monitored resources, such as system values, network attributes, TCP/IP attributes, environment variables, job descriptions, classes, and subsystem descriptions, are enrolled in the cluster administrative domain.
- The device cluster resource group is created to manage the independent disk pool.
- The disk units in the independent disk pool are configured.
- Geographical mirroring is configured, if you selected a high-availability solution that uses geographical mirroring.
- The independent disk pool is varied on.
- The device cluster resource group is started.
- Depending on which policies were selected, exit programs are added to five system exit points. The exit programs that are registered during this step are run after any existing exit programs are run. The VRYEXIT, CRT_PROFILE, and DLT_PROFILE exit programs are added.

See the table below for the exit points used, and a description of what the registered exit programs will do. For more information about exit points and exit programs, see Using Registration Facility APIs and Registration Facility Preprocessing Exit Programs in the i5/OS Information Center.

Exit point	Exit program name Exit program format	Exit program function
QIBM_QDC_VRYEXIT	QSBVRYEXT PROF0100	Calls the program specified in the QUSRHASM/QSHUTDOWN data area. This data area is updated with the name of the program which will shut down the application.
QIBM_QSY_CRT_PROFILE	QSBCTRTEXT CRTP0100	Performs the policy selected for Default action when a user profile is created on the Set Up High Availability Policies page.
QIBM_QSY_DLT_PROFILE	QSBDLTEXT DLTP0100	Performs the policy selected for Default action when a user profile is deleted on the Set Up High Availability Policies page.
QIBM_QWC_PRERESTRICT	QSBPREEXT PRSE0200	Performs the policy selected for Default action before the primary node enters restricted state on the Set Up High Availability Policies page.
QIBM_QWC_PWRDWN SYS	QSBPDSEXT PDPF0200	Performs the policy selected for Default action before the primary node performs a power down on the Set Up High Availability Policies page.

- Data areas are created for startup and shutdown processing. The data areas listed are used so the high-availability solution environment is started and ended in a specific way. If this processing is not done, unintended results such as unwanted failovers or unreachable data, could occur. See the table below for the names of the data areas and a description of what the data areas contain.

Note: End users can update the values in the data areas at any time.

Qualified data area name	Data area contents
QUSRHASM/QSTARTAPP	The 10-character name and 10-character library that is referenced during the switchover processing to start the user applications. This data area references the user program that does whatever processing is needed to start jobs, such as user applications and subsystems. This program might be the same program that is referenced from the QUSRHASM/QSTRUPPGM data area if the user program is called when the system is already running.
QUSRHASM/QSTRUPPGM	The 10-character name and 10-character library that is referenced in the QSTRUPPGM system value. The setup step puts the program for the QSTRUPPGM system value into the QUSRHASM/QSTRUPPGM data area. The QSTRUPPGM system value is replaced with QHASM/QSTRUPPGM because the high-availability solution environment needs to be started in a special order. At each IPL, the high availability environment is started, then the program referenced in the QUSRHASM/QSTRUPPGM data area is called to finish startup processing. The contents of the data area might be MYPGM MYLIB.

Qualified data area name	Data area contents
QUSRHASM/QSHUTDOWN	The 10-character name and 10-character library of the program that shuts down the highly available user applications. The contents of the data area might be MYPGM MYLIB. You need to update the QUSRHASM/QSHUTDOWN data area with the name of the program that shuts down the highly-available user applications. If a program name is not provided in the QUSRHASM/QSHUTDOWN data area, when the independent disk pool is varied off, the system ends all jobs (ENDJOB *IMMED) with references to the independent disk pool. By providing a program here, you can shut down your applications in a more orderly way. This program is called from the VRYEXIT program.

Messages

You might see the following informational and error messages for your selected high-availability solution:

hat005a

Configure independent disk pool.

hat005b

Configure geographic mirroring.

For the complete list of messages, see the “Working with messages” on page 37 information center topic.

Verify administrative switchover

Use the High Availability Solutions Manager interface to run the substeps to perform an administrative switchover. The title of this page might be Verify Administrative Switchover, Verify Administrative Switchover at Same Site, or Verify Administrative Switchover at Different Site.

Important: Ensure that the TCP/IP *INETD server is started on the nodes involved in the switchover.

Substeps performed to complete the administrative switchover

High Availability Solutions Manager performs the following substeps to complete the administrative switchover:

- Verifies that the nodes involved in the administrative switchover are active.
- Verifies that the device CRG is active.
- Verifies that the independent disk pool is active.
- Verifies that no duplicate libraries exist on the backup node.
- Changes the device CRG primary node by using the Change CRG Primary (CHGCRGPRI) command. See the Change CRG Primary information center topic for a detailed description of the functions that the CHGCRGPRI command performs. The CHGCRGPRI command does the following for a device CRG:
 - Varies off the independent disk pool. The exit program registered on the QIBM_QDC_VRYEXIT exit point is called. The exit program was registered on this exit point during the Set Up High Availability Environment step of the deployment.
 - Ends the server takeover IP address that is associated with the independent disk pool.
 - Varies on the independent disk pool on the node that is becoming the new primary node.
 - Starts the TCP/IP interface that is associated with the independent disk pool on the new primary node.

- The device CRG exit program calls the program referenced by the QUSRHASM/QSTARTAPP data area that was set up in the Set Up High Availability Environment step of the deployment. This program starts all user applications that reference the independent disk pool.
- Sets the device CRG status to Active if the switchover completes successfully.
- If the solution is Switched Disk with Geographic Mirroring, the recovery domain of the device CRG is modified, as described in the high-availability policies. For example, suppose the policy is to fail over to a node at the same site. After successful switchover, the recovery domain is changed so that the first available backup is the other node at the same site as the new primary node.

The end result of a successful administrative switchover depends on the high-availability solution that was deployed.

If the solution that has been set up is Switched Disk with Logical Partitions or Switched Disk with Two Systems, the independent disk pool moves to the other node. For example, NODE1 is the primary node and has the independent disk pool, and NODE2 is the backup node. During the switchover, the independent disk pool moves to NODE2, NODE2 becomes the primary node, and NODE1 becomes the backup node.

If the solution that has been set up is Cross-Site Mirroring with Geographic Mirroring, the independent disk pools change roles. For example, there are two nodes: NODE1 and NODE2. NODE1 has the production copy of the independent disk pool and NODE2 has the mirror copy. During the switchover, the mirror copy of the independent disk pool becomes the production copy, and the production copy becomes the mirror copy. Thus, NODE2 now has the production copy and NODE1 has the mirror copy.

If the solution that has been set up is Switched Disk with Geographic Mirroring, and the switchover is going to be to the same site, the hardware switches the same way as for the Switched Disk with Logical Partitions or Switched Disk with Two Systems solutions. If the switchover is to a different site, then the mirror copy of the independent disk pool becomes the production copy, and no hardware switches. For example, NODE1 and NODE2 are part of site NEWYORK. NODE1 is the primary node and has the production copy of the independent disk pool. NODE3 is part of site LONDON and has the mirror copy of the independent disk pool. During the switchover, NODE3 becomes the primary node, the production copy attached to NODE1 becomes the mirror copy, and the mirror copy attached to NODE3 becomes the production copy.

If the administrative switchover fails, the independent disk pool is switched back to the original primary system, and an error message displays in the message area.

Field information

The substep table has the following columns.

Column	Description
Substep	Short description of the command or API that will be run.
Estimated time	Estimation of how much time the substep will take to complete. The estimated time changes dynamically to show the best estimate of the remaining time.
Actual time	The actual amount of time that has elapsed since the substep began running.
Status	The status of the substep. Possible values include: Blank - The substep has not yet been run or was undone. Complete - The substep has completed without error. Failed - The substep encountered errors while running. Running - The substep is currently running.
Command/API	The command or API that is currently being processed to complete the substep.

Completion and error messages that are received as the substeps are run are shown in the message area above the **Run Now**, **Cancel**, **Undo**, and **Close** buttons. You can attempt the switchover again after taking corrective action based on the error messages provided.

Click **Run Now** to begin the substep that the  arrow is pointing to.

Click **Cancel** to exit without completing the administrative switchover.

While the substeps are running, click **Cancel** to stop after a substep has finished running, but before the next substep starts running. After the cancel processing is complete, the following options are available:

- Click **Undo** to undo the previously completed substeps.
- Click **Run Now** to continue running with the first substep that has not yet completed.
- Click **Cancel** to return to the previous page without continuing or undoing any substeps. When setting up the high-availability solution, future set up processing will remember which substep is the next substep to be run.

Click **Close** to exit after all substeps have finished running.

For more information, see the “Performing switchovers” on page 19 information center topic.

Migrate user profiles

The Migrate User Profiles page in High Availability Solutions Manager provides a list of user profiles to migrate to and from the high-availability solution.

The high-availability environment appears as a single system to end users. The number of backup systems or which system is currently active does not matter. The end user sees only one system, which is called the *active primary node*. You should migrate all end user profiles to the high-availability solution. Because the end user's data is only available on the active primary node, their profile is only able to sign on to that node.

Other users perform administrative tasks such as applying fixes or performing system backups. Those user profiles are called *administrative profiles*, and need to access all systems in the high-availability environment. Therefore, you should not migrate administrative profiles to the high-availability solution.

Each user profile references a job description. Job description objects must be in a library in the system auxiliary storage pool (ASP) to be used. If a library containing a job description is migrated to an independent disk pool, the job description becomes unusable. To avoid this potential problem, the High Availability Solutions Manager copies the job description that is associated with each user profile to be migrated to the QUSRHASM library.

Additionally, the newly copied job description in the QUSRHASM library has the initial ASP group (INLASPGRP) parameter updated to the name of the independent disk pool of the high-availability solution. This action is done so that the user is able to access the independent disk pool immediately after signing onto the system. The user profile being migrated is then updated to reference the copied job description in the QUSRHASM library.

To migrate the user profiles, High Availability Solutions Manager performs the following steps:

- The job description that is associated with the profile is copied to the QUSRHASM library and the INLASPGRP parameter is updated, as described above.
- The job description is enrolled in the cluster administrative domain by using the Add Cluster Administrative Domain Monitored Resource Entry (ADDCADMRE) command.
- The user profile is updated to use the job description in the QUSRHASM library.

- The user profile is enrolled in the cluster administrative domain by using the ADDCADMRE command. If the profile does not exist on the other nodes in the solution, the profile is created. If the user profile already exists, it is changed to match the profile on the primary system.

l The ADDCADMRE command might fail if the user ID or group ID numbers that are associated with the user profile are already taken on the other system. If this happens, the High Availability Solutions Manager determines which user profile on the other system is currently taken by using the user or group ID number. That profile's user or group ID number is regenerated. After the user or group ID numbers are freed, the migrated user profile is changed to have the same user or group ID number as the primary system.

Attention: Before you migrate user profiles, libraries, and directories to the high-availability environment, perform a full-system backup. The goal of setting up a high-availability solution is to create a single image of your operating system environment regardless of the number of nodes in your high-availability solution. Libraries, directories, and user profiles are deleted on the backup nodes. For example, when library LIBA is migrated from the primary node to the independent disk pool, all versions of LIBA that exist on all other nodes in the solution are deleted. For more information, see "Changing backup procedures for a high-availability solution" on page 13.

Migrating user profiles to the high-availability environment

To migrate user profiles to the high-availability environment, follow these steps:

1. To select the profiles to be migrated from the Objects Not in the High Availability Solution table, click the box in the Select column.
2. Click **Migrate >**.
3. While migration is taking place, each profile that has been migrated moves from the Objects Not in the High Availability Solution table to the High Availability Solution Objects table. The totals for **Migrated objects** and **Elapsed time** are updated after each profile is migrated. Completion and error messages are displayed in the message area.
4. To cancel the migration, click **Cancel**. The migration stops after the current user profile has completed migration.
5. To undo a user profile migration or to migrate profiles back to the Objects Not in the High-Availability Solution table, follow the Migrating user profiles from the high-availability environment instructions below.
6. Once migration is complete, click **Close** to return to the Set Up High Availability Solution page.

Migrating user profiles from the high availability environment

To migrate user profiles from the high-availability environment, follow these steps:

1. To select the profiles to be migrated from the High-Availability Solution Objects table, click the box in the Select column.
2. Click **< Migrate**.
3. While migration is taking place, each profile that has been migrated moves from the High-Availability Solution Objects table to the Objects Not in the High Availability Solution table. The totals for **Migrated objects** and **Elapsed time** are updated after each profile is migrated. Completion and error messages are displayed in the message area.
4. To cancel the migration, click **Cancel**. The migration stops after the current user profile has completed migration.
5. To undo a user profile migration or to migrate profiles back to the High-Availability Solution Objects table, follow the Migrating user profiles to the high-availability environment instructions above.
6. Once migration is complete, click **Close** to return to the Set Up High Availability Solution page.

Note: Profiles that are migrated out of the high-availability solution are not changed back to their original state before migration.

Note: Migration cannot happen both ways at the same time. Clicking **Migrate >** moves user profiles into the high-availability solution. Clicking **< Migrate** moves user profiles out of the high-availability solution.

Migration tables

Each Migration table contains four columns: Select, Profile, Status, Description.

Click the boxes in the Select column to select the user profiles to migrate.

The Status column contains an icon representing whether or not the user profile can be migrated. The following table shows the icons that might appear and what each icon means.

Status	Description
	The user profile is ready for migration.
	The user profile can be migrated, but you might need to perform additional steps. Click the context menu  next to the user profile name and select View restrictions .
	The user profile cannot be migrated. Click the context menu  next to the user profile name and select View restrictions .

Migration totals for selected objects section:

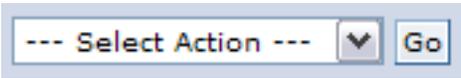
- The **Selected objects** field is updated as user profiles are selected for migration.
- The **Migrated objects** field is updated after each user profile has been migrated.
- The **Estimated time** field is an estimate of the time remaining to migrate all selected user profiles. This field is updated while migration is taking place to reflect the estimated time remaining for user profiles that have not yet been migrated.
- The **Elapsed time** field is updated during migration to show how long the migration has been running.
- Completion and error messages received during migration are shown in the message area.

Click **Cancel** to stop migrating after the current user profile has been migrated.

Click **Close** after user profile migration has completed successfully.

The toolbar for each migration table contains the following toolbar buttons and menu:

Icon	Name	Description
	Select all	Selects all user profiles in the table. You also can select Select all in the Select Action drop-down menu and click Go .
	Deselect all	Deselects all user profiles in the table. You also can select Deselect all in the Select Action drop-down menu and click Go .
	Edit sort	Allows you to specify up to 3 columns from the list to sort the list by and specify ascending or descending sorting for each of the column. You also can select Edit Sort in the Select Action drop-down menu and click Go .

Icon	Name	Description
	Clear all sorts	Removes all current sorted views from the table. You also can select Clear Sort in the Select Action drop-down menu and click Go .
	Select Action drop-down menu	Contains actions for the resources in the table view. Actions are listed in the Name column of this table.

Migrate libraries

Use the Migrate Libraries page in High Availability Solutions Manager to migrate libraries to and from the high-availability solution.

Deploying a high-availability solution implies that data is to be migrated from the primary system into the switchable independent disk pool. Therefore, only one copy of the data will exist. The data residing on the system where the deployment is originating from is moved to the independent disk pool. All other copies of the data are deleted from the backup systems.

High Availability Solutions Manager performs the following steps to migrate a library to the high-availability solution:

1. The library is saved with the SAVLIB command to a user-specified media, such as a tape device, save file, or optical device. (You would specify this device in the **Device used by migration** menu.)
2. The library is deleted from the current system.
3. The library is restored with the RSTLIB command into the independent disk pool.
4. Libraries with the same name as the previously restored library are deleted from the other nodes in the high-availability solution. This step is done to avoid any namespace collisions during switchovers. Private authorities are saved and restored with the library.

Note: The library on the backup node is not saved before deletion.

After the library has completed migration, any job description objects that specify INLLIBL(*SYSVAL) are updated. If the migrated library is referenced in the QUSRLIBL or QSYSLIBL system values, it is removed from those system values. Job descriptions which specify INLLIBL(*SYSVAL) are duplicated into the QUSRHASM library. The duplicated job description are modified to specify all libraries in the QUSRLIBL system value plus the migrated library. This step is done because system values cannot reference any libraries that exist in independent disk pools.

Some object types are allowed to reside in an independent disk pool, but cannot be used in an independent disk pool. Examples of these object types include *JOB, *SBS, and *CLS. If an object of that type exists in a library that needs to be migrated, consider some of the following options so that applications are not disrupted by a migration:

- Move the object to a library that will remain in the system disk pool.
- Do not migrate the library containing the object.
- Create a system disk pool library called **SYSBASE**, and move all objects that should not be migrated into the new SYSBASE library.
- Change all references to the moved objects at their new location in the SYSBASE library.

Job schedule entries might be affected if the independent disk pool must be varied on or varied off when the job is scheduled to run. For example, if the job that is scheduled to run must access the independent disk pool and the independent disk pool is not available at the scheduled time, then the job will not complete successfully. Modify the job schedule entry to run at a time when the independent disk pool is most likely to be available, or modify the job to not require the independent disk pool.

During the migration, a message might appear stating: The total size of the objects selected for migration and the work space required to do the migration exceeds the space available in the high-availability environment. To continue with the migration, you need to add additional unconfigured disk units to the independent disk pool. For details on adding unconfigured disk units to the independent disk pool, see Adding a disk unit or disk pool in the i5/OS Information Center.

Unconfigured disk units come from two sources: a) adding new disk units to the system, or b) removing disk units from the system disk pool. For details about removing disk units from the system disk pool, see Moving or removing a disk unit from a disk pool.

Attention: Before you migrate user profiles, libraries, and directories to the high-availability environment, perform a full-system backup. The goal of setting up a high-availability solution is to create a single image of your operating system environment regardless of the number of nodes in your high-availability solution. Libraries, directories, and user profiles are deleted on the backup nodes. For example, when library LIBA is migrated from the primary node to the independent disk pool, all versions of LIBA that exist on all other nodes in the solution are deleted. For more information, see “Changing backup procedures for a high-availability solution” on page 13.

Migrating libraries to the high availability environment

To migrate libraries to the high-availability environment, follow these steps:

1. Select the device to be used for the migration from the **Device used by migration** drop-down menu.
2. To select the libraries to be migrated from the Objects Not in the High Availability Solution table, click the box in the Select column.
3. Click **Migrate >**.

Note: Migration cannot happen both ways at the same time. Clicking **Migrate >** moves libraries into the high-availability solution. Clicking **< Migrate** moves libraries out of the high-availability solution.

4. While migration is taking place, each library that has been migrated moves from the Objects Not in the High-Availability Solution table to the High-Availability Solution Objects table. The totals for Migrated objects and Elapsed time are updated after each library is migrated. Completion and error messages are displayed in the message area.
5. To cancel the migration, click **Cancel**. The migration stops after the current library has completed migration.
6. To undo a library migration or to migrate libraries back to the Objects Not in the High Availability Solution table, follow the Migrating libraries from the high-availability environment instructions below.
7. Once migration is complete, click **Close** to return to the Set Up High Availability Solution page.

Migrating libraries from the high availability environment

To migrate libraries from the high-availability environment, follow these steps:

1. Select the device to be used for the migration from the **Device used by migration** drop-down menu.
2. To select the libraries to be migrated from the High-Availability Solution Objects table, click the box in the Select column.
3. Click **< Migrate**.

Note: Migration cannot happen both ways at the same time. Clicking **Migrate >** moves libraries into the high-availability solution. Clicking **< Migrate** moves libraries out of the high-availability solution.

4. While migration is taking place, each library that has been migrated moves from the High-Availability Solution Objects table to the Objects Not in the High-Availability Solution table. The totals for Migrated objects and Elapsed time are updated after each library is migrated. Completion and error messages are displayed in the message area.
5. To cancel the migration, click **Cancel**. The migration stops after the current library has completed migration.
6. To undo a library migration or to migrate libraries back to the High Availability Solution Objects table, follow the Migrating libraries to the high availability environment instructions above.
7. Once migration is complete, click **Close** to return to the Set Up High Availability Solution page.

Note: Libraries that are migrated from the high-availability solution are not migrated to all of the nodes in the high-availability solution. They are migrated back to only the primary system.

Migration tables

The tables of libraries have the following columns: Select, Library, Status, Size, Estimated Time, and Dependent Library.

Click the boxes in the **Select** column to select the library in the **Library** column.

The **Status** column contains an icon representing whether or not the library can be migrated. The following table shows the icons that might appear and what each icon means.

Status	Description
	The library is ready for migration.
	The library can be migrated, but dependent libraries will be migrated with this library. Click the context menu  next to the path name and select View dependencies .
	The library cannot be migrated. Click the context menu  next to the path name and select View restrictions .

The **Device used by migration** menu lists the available tape device descriptions on the systems. Possible values include:

- | *VRT,*TAP,63B0,001 (default)
- | *SAVF
- | *tape1*
- | *tape2*

The **Size** column indicates the size (in megabytes) of the library.

The **Estimated Time** column shows an approximation of how long it will take to migrate the library.

The **Dependent Library** column identifies other libraries that need to be migrated with the selected library. For a list of the reasons why libraries are dependent, see “Library dependencies” on page 69.

Migration totals for selected objects section:

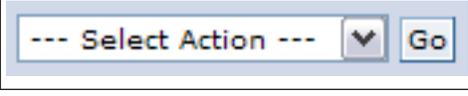
- The **Selected objects** field is updated as libraries are selected for migration.
- The **Migrated objects** field is updated after each library has been migrated.
- The **Estimated time** field is an estimate of the time remaining to migrate all selected libraries. This field is updated while migration is taking place to reflect the estimated time remaining for libraries that have not yet been migrated.
- The **Elapsed time** field is updated during migration to show how long the migration has been running.

- Completion and error messages received during migration are shown in the message area.

Click **Cancel** to stop migrating after the current library has been migrated.

Click **Close** after library migration has completed successfully.

The toolbar for each migration table contains the following toolbar buttons and menu:

Icon	Name	Description
	Select all	Selects all items in the table. You also can select Select all in the Select Action drop-down menu and click Go .
	Deselect all	Deselects all items in the table. You also can select Deselect all in the Select Action drop-down menu and click Go .
	Edit sort	Allows you to specify up to 3 columns from the list to sort the list by and specify ascending or descending sorting for each of the column. You also can select Edit Sort in the Select Action drop-down menu and click Go .
	Clear all sorts	Removes all current sorted views from the table. You also can select Clear Sort in the Select Action drop-down menu and click Go .
	Select Action drop-down menu	Contains actions for the resources in the table view. Actions are listed in the Name column of this table.

Library dependencies

The following is a list of possible reasons why a library might be dependent upon another library:

- The library contains a journal that has journaled objects in a different library.
- The library contains a journal that has journal receivers in a different library.
- The library contains a physical file that has a logical file associated with it in a different library.
- The library contains a logical file that is associated with a physical file in a different library.

Library restrictions

The following is a list of possible reasons why a library cannot be migrated to the high-availability solution:

- The library is part of i5/OS and can only reside in the system disk pool.
- The library contains a journal that has one or more integrated file system objects journaled to it.
- The library is in the library list of the current thread.
- The library is in the library list of any primary thread that is active on the system when the QLIBLCKLVL system value is set to lock libraries in the library list.
- The library contains an allocated job queue or output queue.
- The independent disk pool does not have enough space for the library and its objects.
- The library contains object types that cannot reside in an independent disk pool. For a complete list of objects that cannot reside in an independent disk pool, see Supported and unsupported object types.
- The library is referenced by any of the system values shown below.

QACGLVL	QATNPGM	QAUDCTL	QCFGMSGQ
---------	---------	---------	----------

QCONSOLE	QCTLSBSD	QIGCCDEFNT	QINACTMSGQ
QPRBFTR	QPRTDEV	QPWDVLDPGM	QRMTSIGN
QSRTSEQ	QSTRUPPGM	QUPSMMSGQ	QUSEADPAUT

- The library is referenced by any of the network attributes shown below.

ALRFTR	DDMACC	DFTMODE
MSGQ	OUTQ	PCSACC

- The library is referenced by a user profile that has not been migrated to the independent disk pool.

Messages

You might see the following informational and error messages for your selected high-availability solution:

hat0042

All unconfigured DASD under the tower will be selected if any are selected.

hat0044

{0} contains object not supported in independent disk pool.

hat0045

{0} contains journal {1} to IFS objects.

hat0046

{0} contains job descriptions.

hat0047

{0} in use by System Value {1}.

hat0048

{0} in use by Network Attribute {1}.

hat0049

{0} contains journaled IFS objects.

hat004a

{0} contains class objects.

hat004b

{0} contains subsystem descriptions.

hat004c

{0} in use by subsystem description {1}.

hat004d

{0} contains database file {1} with data link control.

hat004e

{0} contains multisystem database file {1} (node group).

For the complete list of messages, see the “Working with messages” on page 37 information center topic.

Migrate directories

Use the Migrate Directories page in High Availability Solutions Manager to migrate directories to and from the high-availability solution.

Deploying a high-availability solution implies that data is to be migrated from the primary system into the switchable independent disk pool. Therefore, only one copy of the data exists. The data residing on

the system where the deployment is originating from is moved to the independent disk pool. All other copies of the data are deleted from the backup systems.

High Availability Solutions Manager performs the following steps to migrate a directory to the high-availability solution:

1. The directories and their subdirectories are saved with the SAV command to a user-specified media, such as a tape device, save file, or optical device. (You would specify this device in the **Device used by migration** menu.)
2. The directory is deleted from the system.
3. The directory is restored with the RST command into the independent disk pool. For example, if directory **/mydir** is to be migrated, the system would restore the directory to **/IndependentDiskPoolName/mydir**.
4. A symbolic link is added to the restored directory so that existing applications continue to work without modification. For example, the system would add a symbolic link named **/mydir**, whose contents would link to **/IndependentDiskPoolName/mydir**.
5. Directories with the same name as the migrated directory are deleted from the other nodes in the high-availability solution.

Note: The directory on the backup node is not saved before deletion.

6. A symbolic link is created on every backup node so that when the independent disk pool is switched to the backup node, the applications running on the backup node can find the migrated data.

During the migration, a message might appear stating The total size of the objects selected for migration and the work space required to do the migration exceeds the space available in the high-availability environment. To continue with the migration, you need to add additional unconfigured disk units to the independent disk pool. For details on adding unconfigured disk units to the independent disk pool, see Adding a disk unit or disk pool in the i5/OS Information Center.

Unconfigured disk units come from two sources: a) adding new disk units to the system, or b) removing disk units from the system disk pool. For details about removing disk units from the system disk pool, see Moving or removing a disk unit from a disk pool in the i5/OS Information Center.

Attention: Before you migrate user profiles, libraries, and directories to the high-availability environment, perform a full-system backup. The goal of setting up a high-availability solution is to create a single image of your operating system environment regardless of the number of nodes in your high-availability solution. Libraries, directories, and user profiles are deleted on the backup nodes. For example, when library LIBA is migrated from the primary node to the independent disk pool, all versions of LIBA that exist on all other nodes in the solution are deleted. For more information, see “Changing backup procedures for a high-availability solution” on page 13.

Migrating directories to the high-availability environment

To migrate directories to the high-availability environment, follow these steps:

1. Select the device to be used for the migration from the **Device used by migration** drop-down menu.
2. To select the directories to be migrated from the Objects Not in the High Availability Solution table, click the box in the Select column.
3. Click **Migrate >**.

Note: Migration cannot happen both ways at the same time. Clicking on **Migrate >** moves directories into the high-availability solution. Clicking on **< Migrate** moves directories out of the high-availability solution.

4. While migration is taking place, each directory that has been migrated moves from the Objects Not in the High-Availability Solution table to the High-Availability Solution Objects table. The totals for Migrated objects and Elapsed time are updated after each directory is migrated. Completion and error messages are displayed in the message area.
5. To cancel the migration, click **Cancel**. The migration stops after the current directory has completed migration.
6. To undo a directory migration or to migrate directories back to the Objects Not in the High-Availability Solution table, follow the Migrating directories from the high-availability environment instructions below.
7. Once migration is complete, click **Close** to return to the Set Up High Availability Solution page.

Migrating directories from the high-availability environment

To migrate directories from the high-availability environment, follow these steps:

1. Select the device to be used for the migration from the **Device used by migration** drop-down menu.
2. To select the directories to be migrated from the High Availability Solution Objects table, click the box in the Select column.
3. Click **< Migrate**.

Note: Migration cannot happen both ways at the same time. Clicking on **Migrate >** moves directories into the high-availability solution. Clicking on **< Migrate** moves directories out of the high-availability solution.

4. While migration is taking place, each directory that has been migrated moves from the High-Availability Solution Objects table to the Objects Not in the High-Availability Solution table. The totals for Migrated objects and Elapsed time are updated after each directory is migrated. Completion and error messages are displayed in the message area.
5. To cancel the migration, click **Cancel**. The migration stops after the current directory has completed migration.
6. To undo a directory migration or to migrate directories back to the High-Availability Solution Objects table, follow the Migrating directories to the high-availability environment instructions above.
7. Once migration is complete, click **Close** to return to the Set Up High Availability Solution page.

Note: Directories that are migrated from the high-availability solution are not migrated to all of the nodes in the high-availability solution. They are migrated back to the primary system.

Migration tables

The tables of directories have the following columns: Select, Path, Status, Size, and Estimated Time.

Click the boxes in the **Select** column to select the directory in the **Path** column.

The **Status** column contains an icon representing whether or not the directory can be migrated. The following table shows the icons that might appear and what each icon means.

Status	Description
	The directory is ready for migration.
	The directory can be migrated, but there are warnings to consider before migrating the directory. Click the context menu  next to the path name and select View warnings .
	The directory cannot be migrated. Click the context menu  next to the path name and select View restrictions .

The **Device used by migration** menu lists the available tape device descriptions on the systems. Possible values include:

- | *VRT,*TAP,63B0,001 (default)
- | *SAVF
- | *tape1*
- | *tape2*

The **Size** column indicates the size of the directory.

The **Estimated Time** column shows an approximation of how long it takes to migrate the directory.

Migration totals for selected objects section:

- The **Selected objects** field is updated as directories are selected for migration.
- The **Migrated objects** field is updated after each directory has been migrated.
- | • The **Size selected** field is the size of all selected directories.
- | • The **Size available** field is the size available for migrating the selected directories. The **Size available** value must be larger than the **Size selected** value for a successful migration.
- | • The **Estimated time** field is an estimate of the time remaining to migrate all selected directories. This field is updated while migration is taking place to reflect the estimated time remaining for directories that have not yet been migrated.
- The **Elapsed time** field is updated during migration to show how long the migration has been running.
- Completion and error messages received during migration are shown in the message area.

Click **Cancel** to stop migrating after the current directory has been migrated.

Click **Close** after directory migration has completed successfully.

The toolbar for each migration table contains the following toolbar buttons and menu:

Icon	Name	Description
	Select all	Selects all items in the table. You also can select Select all in the Select Action drop-down menu and click Go .
	Deselect all	Deselects all items in the table. You also can select Deselect all in the Select Action drop-down menu and click Go .
	Edit sort	Allows you to specify up to 3 columns from the list to sort the list by and specify ascending or descending sorting for each of the column. You also can select Edit Sort in the Select Action drop-down menu and click Go .
	Clear all sorts	Removes all current sorted views from the table. You also can select Clear Sort in the Select Action drop-down menu and click Go .
	Select Action drop-down menu	Contains actions for the resources in the table view. Actions are listed in the Name column of this table.

Migration warnings

The following is a list of possible warnings to be considered before migrating the directory to the high-availability solution:

- The path resolves to a symbolic link. Only the link is migrated. The directory or file pointed to by the link is not migrated.
- The path resolves to an object that is contained within a mounted local file system. Migrating objects in a mounted local file system is permitted, but the object is restored into the default file system of the independent disk pool. These objects no longer exist in the mounted local file system.

Migration restrictions

The following is a list of possible reasons why a directory cannot be migrated to the high-availability solution:

- The directory subtree contains an object that is currently journaled. Migrating integrated file system objects that are journaled requires you to take the following actions before these objects can be migrated:
 - Run `ENDJRN OBJ('/mydir') SUBTREE(*ALL) JRN(*OBJ) LOGLVL(*ALL)`
Make note of all of the journaled objects, journals, and journal libraries that are associated with the objects.
 - Migrate each library, which contained journals associated with the objects, into the high-availability solution.
 - Migrate the `/mydir` directory subtree.
 - Run the `STRJRN` command to start journaling for each of the objects which had previously been journaled
- The directory is provided by the system. There is no recovery or work around for this restriction. It is expected that the directories provided by the system are already in the system disk pool on every system in the high-availability solution.
- The directory is a mount point. If a file system mounted over the directory, there is no recovery or workaround. Mounting and unmounting file systems provides a dynamic way to change the integrated file system namespace. File systems can only be mounted over directories. The result of migrating a directory leaves a symbolic link in place of the original directory. The file system can no longer be mounted at the same location.
- The path resolves to an object within a remote file system. There is no recovery or work around for this restriction. The remote file system is mounted on every system and, therefore, is not needed in the switchable independent disk pool.

Messages

You might see the following informational and error messages for your selected high-availability solution:

hat0042

All unconfigured DASD under the tower will be selected if any are selected.

hat0044

{0} contains object not supported in independent disk pool.

hat0045

{0} contains journal {1} to IFS objects.

hat0046

{0} contains job descriptions.

hat0047

{0} in use by System Value {1}.

hat0048

{0} in use by Network Attribute {1}.

hat0049

{0} contains journaled IFS objects.

- hat004a**
{0} contains class objects.
- hat004b**
{0} contains subsystem descriptions.
- hat004c**
{0} in use by subsystem description {1}.
- hat004d**
{0} contains database file {1} with data link control.
- hat004e**
{0} contains multisystem database file {1} (node group).

For the complete list of messages, see the “Working with messages” on page 37 information center topic.

Finish set up and clean up files

The Finish Set Up and Clean Up Work Files page in High Availability Solutions Manager processes the substeps that are required to complete the setup of the high-availability solution, and clean up the work files that were created during the setup. This is the final step before your high-availability solution becomes active.

Important: This step cannot be undone. Once you have completed the substeps, the high-availability solution will be deployed. The policies chosen during the deployment are activated to ensure that the solution is highly available.

High Availability Solutions Manager performs the following substeps:

1. The QSTRUPPGM system value is updated with QHASM/QSBSTRUPPM. This program does everything that is required, including starting any cluster resource groups, to make the system highly available at startup.
2. The exit programs registered on the QIBM_QDC_VRYEXIT, QIBM_QWC_PRERESTRICT, and QIBM_QWC_PWRDWN SYS exit points are updated to include ending the cluster resource groups when varying off the independent disk pool, entering restricted state, or powering down the system.
3. Internal work files created during the setup of the high availability infrastructure are deleted.
4. The cluster nodes are started.
5. The cluster resource groups are started.
6. The independent disk pool is varied on.

Field information

The substep table has the following columns.

Column	Description
Substep	Short description of the command or API that will be run.
Estimated time	Estimation of how much time the substep will take to complete. The estimated time changes dynamically to show the best estimate of the remaining time.
Actual time	The actual amount of time that has elapsed since the substep began running.
Status	The status of the substep. Possible values include: Blank - The substep has not yet been run or was undone. Complete - The substep has completed without error. Failed - The substep encountered errors while running. Running - The substep is currently running.
Command/API	The command or API that is currently being processed to complete the substep.

The substep table has the following columns:

Column	Description
Substep	Short description of the command or API that will be run.
Estimated time	Estimation of how much time the substep will take to complete. The estimated time changes dynamically to show the best estimate of the remaining time.
Actual time	The actual amount of time that has elapsed since the substep began running.
Status	The status of the substep. Possible values include: Blank - The substep has not yet been run or was undone. Complete - The substep has completed without error. Failed - The substep encountered errors while running. Running - The substep is currently running.
Command/API	The command or API that is currently being processed to complete the substep. This field is blank if the substep has not been run or has been undone.

Click **Run Now** to begin the substep that the  arrow is pointing to.

Click **Cancel** to return to the Set Up High Availability Solution page without completing the deployment.

While the substeps are running, click **Cancel** to stop after the current substep has finished running, but before the next substep starts running. After the cancel processing is complete, the following options are available:

- Click **Run Now** to continue running with the first substep that has not completed.
- Click **Cancel** to return to the Set Up High Availability Solution page. Future set up processing will remember which substep is the next substep to be run.

Click **Close** to return to the Set Up High Availability Solution page after all substeps have finished running.

The message area below the substep table shows the completion and error messages that are received as the substeps are run.

Messages

You might see the following informational and error messages for your selected high-availability solution:

hat006f

The setup of your high-availability solution is now complete. If your normal security policy is to have the QSECOFR profile disabled and you enabled the QSECOFR profile to set up your high-availability solution, then you should return to your normal security policy and disable the QSECOFR profile.

For the complete list of messages, see the “Working with messages” on page 37 information center topic.

Manage your high availability solution - Nodes

The Manage Your High Availability Solution page is used to monitor and manage your high-availability solution. From this page, you also can manage the cluster nodes that your high-availability solution uses.

You can monitor and perform high-availability solution-level actions, monitor and manage high-availability solution resources, and monitor solution events in the event log. The dynamically changing status icons makes this page ideal to use as a *dashboard*.

There are three sections on the Manage Your High Availability Solution page:

- The **Manage Your High Availability Solution** section provides an at-glance view of the status of the high-availability solution and quick access to solution-level actions.
- The **High Availability Solution Resources** section provides a tabbed list of all of the high-availability solution resources. Each tab gives a detailed view of each resource along with possible actions to perform on the resource.
- The **Event Log** section presents the list of events that have occurred in the high-availability solution.

The **Manage Your High Availability Solution** section gives a graphical summary of the status of the high-availability solution. See the table below for a description of the possible statuses.

Table 11. Status of the high-availability solution

Status	Description
	High-availability solution is available for switchover.
	High-availability solution is not available for switchover, but no action is required.
	High-availability solution is not available for switchover. User action is required.

A dropdown menu with the following solution-level actions is included. To perform one of the actions, select it from the menu and click **Go**.

Switchover

Select this action if you want to move the high-availability solution to the backup node. This process makes the backup node the new primary node. Perform this action before a planned outage of the primary node if the application is to remain available during the outage.

Shut Down - Keep Data Available

Select this action if you want to perform system maintenance or a system backup. When the action is performed, your system still is usable, but is not highly available should an unplanned outage occur. Although the application is still available, it will be unavailable in either a planned or unplanned outage.

Shut Down - Make Data Not Available

This action is typically run for system maintenance or if the system is going to be powered down. When this action is performed, your system is accessible only by administrative users. The application is not available on either the primary or backup nodes.

Resume

Select this action if the high-availability solution was shut down and needs to be resumed.

Service Info

Select this action if you want to prepare data to send to a service representative. The system creates a save file named QUSRHASM/QSBSERVICE on the node from which the collection is initiated. The save file should be sent when service information is requested.

Recover Partition

Select this action to recover from a partition state. Cluster Resource Services cannot detect certain failure conditions as a node failure. Recovery from partition state requires operator intervention to determine the cause of the partition state. Selecting this action is only done in certain rare cases. The help for this task describes the process for deciding if selecting this action is necessary.

In addition to the previously listed actions, the Switched disk with geographic mirroring solution has the following possible actions:

Switchover to Different Site

Select this action if you want to run the high-availability solution on a backup node located at a geographically-different site. This action is typically done when all of the nodes at one site need to be taken down for maintenance.

Switchover to Same Site

Select this action if you want to run the high-availability solution on a backup node located at the same geographic location as the current primary node. This process makes the backup node at that site the new primary node.

The **High Availability Solution Resources** section provides a tabbed list of resources including Nodes, Cluster Resource Groups, Monitored Resources, Independent Disk Pools, TCP/IP Interfaces, and Policies. To navigate through this list, click the tab on the left side of the section.

Nodes

The high-availability solution resources **Nodes** tab of the Manage Your High Availability Solution page lists all of the nodes in the high-availability solution, as well as the status and role of each node.

The Name column of the table contains the name of each node in the solution. Beside the name is the  button. Click this button to see what actions are available for the node. The status of the node determines which actions are available. For example, if the node is already started, the **Start** action is not available. The possible actions are:

Start

Start the node.

Stop

Stop the node.

Work with All Nodes

View and work with all nodes, including nodes which might not be included in the high-availability solution.

The Status column of the table contains an icon representing the status of the node. See the following table for a list of icons and the status each represents. Another option is to hold the mouse over the icon to see a banner pop up with the description.

Table 12. Node status indicators

Icon	Status	Description
	Active	The node has been started by using a Create Cluster operation, an Add Cluster Node Entry operation, which started the cluster node; or with a Start Cluster Node operation. Cluster Resource Services is active on the node.
	Active Pending	The node is in the process of being started as the result of a Create Cluster operation, an Add Cluster Node Entry operation, which started the cluster node, or a Start Cluster Node operation. In addition, the node could have previously had a status of Partition and will change to the Active Pending status, as a result of the partitions being merged.
	Inactive Pending	Cluster Resource Services is in the process of ending on this node as the result of an end cluster node operation. The node is still in the cluster membership list.
	Remove Pending	The node is in the process of being removed from the cluster membership list as the result of a Remove Cluster Node Entry operation.
	New	A node has been added to the cluster membership list, but Cluster Resource Services has never been started on that node. The Cluster Resource Service data structures have not been created on the node. The Cluster Resource Service data structures are created only on the node running the Create Cluster operation.

Table 12. Node status indicators (continued)

Icon	Status	Description
	Inactive	Cluster Resource Services has been ended on the node as the result of an end cluster node operation. The node is still in the cluster membership list, but is no longer communicating with other nodes in the cluster.
	Failed	A previously active node has failed. A failure is either a system or clustering failure detected by Cluster Resource Services.
	Partition	The node is only communicating with a subset of the cluster because of a network failure detected by Cluster Resource Services, which has resulted in the loss of communications to one or more nodes in the cluster. When the partitioned nodes are merged back into a whole cluster, the node will change to Active status without operator intervention. Any node that had a status of Failed in any partition will still have a status of Failed after the merge.
	Unknown	This node is not an active member of the cluster, so the status of other nodes cannot be determined.

The Solution Role column of the table shows the current role of the node in the solution. See the table below for a list of the possible solution roles and a description of each.

Table 13. Current role of each node in the solution

Solution Role	Description
Primary	The node that is currently running the high-availability solution. The independent disk pool is owned by the primary node.
Backup	The node that will take over the role of primary if the current primary node fails or a manual switchover is initiated. The backup node does not own the independent disk pool until a switchover occurs.
Replicate	A node that has a copy of the cluster resources, but cannot assume the role of primary or backup, unless you assign the node a new role of primary or backup.
Peer	A node that has no order. This node can be an access point for cluster resources. All peer nodes should be equivalent.

Event log information

The **Event Log** section enables you to enter how many days event information should be kept on the system, and see the list of events that have occurred within the number of days specified. The columns shown in the table of logs are Time, Severity, and Information.

Time Shows the date and time the event was logged.

Severity

Shows an icon that represents the urgency of the event. See the table below for a description of each icon.

Information

Gives a brief description of the event.

Click the **High Availability Operator's Guide** link to view or download the *Implementing high availability with the High Availability Solutions Manager – Operator's guide* PDF.

Click **Close** to return to the High Availability Solutions Manager welcome page.

Table 14. Event log status indicators

Severity	Description
	The log contains informational messages.
	The log contains warning messages and should be examined.

For more information, see the following information center topics:

- “Managing a high availability solution” on page 14
- “Working with high availability status messages” on page 15
- “Managing nodes” on page 21

Manage high availability solution - Cluster Resource Groups

The Manage Your High Availability Solution page is used to monitor and manage your high-availability solution. From this page, you also can manage the cluster resource groups (CRGs) that your high-availability solution uses.

You can monitor and perform high-availability solution-level actions, monitor and manage high-availability solution resources, and monitor solution events in the event log. The dynamically changing status icons makes this page ideal to use as a *dashboard*.

There are three sections on the Manage Your High Availability Solution page:

- The **Manage Your High Availability Solution** section provides an at-glance view of the status of the high-availability solution and quick access to solution-level actions.
- The **High Availability Solution Resources** section provides a tabbed list of all of the high-availability solution resources. Each tab gives a detailed view of each resource along with possible actions to perform on the resource.
- The **Event Log** section presents the list of events that have occurred in the high-availability solution.

- 1 The **Manage Your High Availability Solution** section gives a graphical summary of the status of the high-availability solution. See the table below for a description of the possible statuses.

Table 15. Status of the high-availability solution

Status	Description
	High-availability solution is available for switchover.
	High-availability solution is not available for switchover, but no action is required.
	High-availability solution is not available for switchover. User action is required.

A dropdown menu with the following solution-level actions is included. To perform one of the actions, select it from the menu and click **Go**.

Switchover

Select this action if you want to move the high-availability solution to the backup node. This process makes the backup node the new primary node. Perform this action before a planned outage of the primary node if the application is to remain available during the outage.

Shut Down - Keep Data Available

Select this action if you want to perform system maintenance or a system backup. When the

action is performed, your system still is usable, but is not highly available should an unplanned outage occur. Although the application is still available, it will be unavailable in either a planned or unplanned outage.

Shut Down - Make Data Not Available

This action is typically run for system maintenance or if the system is going to be powered down. When this action is performed, your system is accessible only by administrative users. The application is not available on either the primary or backup nodes.

Resume

Select this action if the high-availability solution was shut down and needs to be resumed.

Service Info

Select this action if you want to prepare data to send to a service representative. The system creates a save file named QUSRHASM/QSBSERVICE on the node from which the collection is initiated. The save file should be sent when service information is requested.

Recover Partition

Select this action to recover from a partition state. Cluster Resource Services cannot detect certain failure conditions as a node failure. Recovery from partition state requires operator intervention to determine the cause of the partition state. Selecting this action is only done in certain rare cases. The help for this task describes the process for deciding if selecting this action is necessary.

In addition to the previously listed actions, the Switched disk with geographic mirroring solution has the following possible actions:

Switchover to Different Site

Select this action if you want to run the high-availability solution on a backup node located at a geographically-different site. This action is typically done when all of the nodes at one site need to be taken down for maintenance.

Switchover to Same Site

Select this action if you want to run the high-availability solution on a backup node located at the same geographic location as the current primary node. This process makes the backup node at that site the new primary node.

The **High Availability Solution Resources** section provides a tabbed list of resources including Nodes, Cluster Resource Groups, Monitored Resources, Independent Disk Pools, TCP/IP Interfaces, and Policies. To navigate through this list, click the tab on the left side of the section.

Cluster Resource Groups

The high-availability solution resources **Cluster Resource Groups** tab of the Manage Your High Availability Solution page shows a list of all of the cluster resource groups (CRGs) in the high-availability solution, as well as the status and role of each CRG.

The Name column of the table contains the name of each CRG in the solution. Beside the name is the  button. Click this button to see what actions are available for the CRG. The status of the CRG determines which actions are available. For example, if the CRG is already started, the **Start** action is not available. The possible actions are:

Start Start the CRG.

Stop Stop the CRG.

Work with All CRGs

View and work with all CRGs, including CRGs which might not be included in the high-availability solution.

The Status column of the table contains an icon representing the status of the CRG. See the following table for a list of icons and the status each represents. Another option is to hold the mouse over the icon to see a banner pop up with the description.

Table 16. Cluster resource group status indicators

Icon	Status	Description
	Active	The resources managed by the cluster resource group are currently resilient.
	Inactive	The resources managed by the cluster resource group are currently not resilient.
	Indoubt	The information contained within the cluster resource group object might not be accurate. This status occurs when an exit program is called with an action of Undo and fails to complete successfully.
	Restored	The cluster resource group object was restored on this node and was not copied to the other nodes in the recovery domain. When Cluster Resource Services is started on this node, the cluster resource group will be synchronized with the other nodes in the recovery domain and its status set to Inactive .
	Add Pending	A new node is in the process of being added to the recovery domain of a cluster resource group. If the exit program is successful, the status is reset to its value at the time the API was called. If the exit program fails and the original state cannot be recovered, the status is set to Indoubt .
	Delete Pending	Cluster resource group is in the process of being deleted. When the exit program completes, the cluster resource group is deleted from all nodes in the recovery domain.
	Change Pending	The cluster resource group is in the process of being changed. If the exit program is successful, the status is reset to the value at the time the API was called. If the exit program fails and the original state cannot be recovered, the status is set to Indoubt .
	End Pending	Resiliency for the cluster resource group is in the process of ending. If the exit program is successful, the status is set to Inactive . If the exit program fails and the original state cannot be recovered, the status is set to Indoubt .
	Initialize Pending	A cluster resource group is being created and is in the process of being initialized. If the exit program is successful, the status is set to Inactive . If the exit program fails, the cluster resource group will be deleted from all nodes.
	Remove Pending	A node is in the process of being removed from the recovery domain of the cluster resource group. If the exit program is successful, the status is reset to the value at the time the API was called. If the exit program fails and the original state cannot be recovered, the status is set to Indoubt .
	Start Pending	Resiliency is in the process of starting for the cluster resource group. If the exit program is successful, the status is set to Active . If the exit program fails and the original state cannot be recovered, the status is set to Indoubt . For peer cluster resource groups, all nodes defined with a role of peer are active access points for the cluster resources.
	Switchover Pending	The Initiate Switchover API was called, a failure of a cluster resource group occurred, or a node failed, causing a switchover or failover to begin. The first backup node is in the process of becoming the primary node. If the exit program is successful, the status is set to Active . If the exit program fails and the original state cannot be recovered, the status is set to Indoubt . While the switchover function is not valid for a peer cluster resource group, users may see the status Switchover pending during a node failure.

The Type column of the table shows the type of CRG. See the table below for a list of the possible CRG types.

Table 17. Cluster resource group types

Type	Description
Administrative domain	An administrative domain is a special case of a peer CRG. An administrative domain provides a way for objects, such as user profiles and system values, to be synchronized across all of the nodes in the recovery domain.
Application	Application resiliency enables an application program to be restarted on either the same node or a different node in the cluster.
Data	Data resiliency enables multiple copies of data to be maintained on more than one node in a cluster, and enables the point of access to be changed to a backup node.
Device	Device resiliency enables a device resource to be moved (switched) to a backup node.
Peer	Peer resiliency enables all nodes within the CRG to have the same data. When a failure occurs, the point of failure is communicated to the other nodes in the CRG and those nodes continue the operation from the failure point.

Event log information

The **Event Log** section enables you to enter how many days event information should be kept on the system, and see the list of events that have occurred within the number of days specified. The columns shown in the table of logs are Time, Severity, and Information.

Time Shows the date and time the event was logged.

Severity

Shows an icon that represents the urgency of the event. See the table below for a description of each icon.

Information

Gives a brief description of the event.

Click the **High Availability Operator's Guide** link to view or download the *Implementing high availability with the High Availability Solutions Manager – Operator's guide* PDF.

Click **Close** to return to the High Availability Solutions Manager welcome page.

Table 18. Event log status indicators

Severity	Description
	The log contains informational messages.
	The log contains warning messages and should be examined.

For more information, see the following information center topics:

- "Managing a high availability solution" on page 14
- "Working with high availability status messages" on page 15
- "Managing cluster resource groups (CRG)" on page 24

Manage high-availability solution - Monitored resources

The Manage Your High Availability Solution page is used to monitor and manage your high-availability solution. From this page, you also can manage the monitored resources that your high-availability solution uses.

You can monitor and perform high-availability solution-level actions, monitor and manage high-availability solution resources, and monitor solution events in the event log. The dynamically changing status icons makes this page ideal to use as a *dashboard*.

There are three sections on the Manage Your High Availability Solution page:

- The **Manage Your High Availability Solution** section provides an at-glance view of the status of the high-availability solution and quick access to solution-level actions.
- The **High Availability Solution Resources** section provides a tabbed list of all of the high-availability solution resources. Each tab gives a detailed view of each resource along with possible actions to perform on the resource.
- The **Event Log** section presents the list of events that have occurred in the high-availability solution.

The **Manage Your High Availability Solution** section gives a graphical summary of the status of the high-availability solution. See the table below for a description of the possible statuses.

Table 19. Status of the high-availability solution

Status	Description
	High-availability solution is available for switchover.
	High-availability solution is not available for switchover, but no action is required.
	High-availability solution is not available for switchover. User action is required.

A dropdown menu with the following solution-level actions is included. To perform one of the actions, select it from the menu and click **Go**.

Switchover

Select this action if you want to move the high-availability solution to the backup node. This process makes the backup node the new primary node. Perform this action before a planned outage of the primary node if the application is to remain available during the outage.

Shut Down - Keep Data Available

Select this action if you want to perform system maintenance or a system backup. When the action is performed, your system still is usable, but is not highly available should an unplanned outage occur. Although the application is still available, it will be unavailable in either a planned or unplanned outage.

Shut Down - Make Data Not Available

This action is typically run for system maintenance or if the system is going to be powered down. When this action is performed, your system is accessible only by administrative users. The application is not available on either the primary or backup nodes.

Resume

Select this action if the high-availability solution was shut down and needs to be resumed.

Service Info

Select this action if you want to prepare data to send to a service representative. The system creates a save file named QUSRHASM/QSBSERVICE on the node from which the collection is initiated. The save file should be sent when service information is requested.

Recover Partition

Select this action to recover from a partition state. Cluster Resource Services cannot detect certain failure conditions as a node failure. Recovery from partition state requires operator intervention to determine the cause of the partition state. Selecting this action is only done in certain rare cases. The help for this task describes the process for deciding if selecting this action is necessary.

In addition to the previously listed actions, the Switched disk with geographic mirroring solution has the following possible actions:

Switchover to Different Site

Select this action if you want to run the high-availability solution on a backup node located at a geographically-different site. This action is typically done when all of the nodes at one site need to be taken down for maintenance.

Switchover to Same Site

Select this action if you want to run the high-availability solution on a backup node located at the same geographic location as the current primary node. This process makes the backup node at that site the new primary node.

The **High Availability Solution Resources** section provides a tabbed list of resources including Nodes, Cluster Resource Groups, Monitored Resources, Independent Disk Pools, TCP/IP Interfaces, and Policies. To navigate through this list, click the tab on the left side of the section.

Monitored Resources

The high-availability solution resources **Monitored resources** tab of the Manage Your High Availability Solution page lists all of the monitored resources in the high-availability solution, as well as the status and type of each monitored resource.

The Global Status column of the table contains an icon representing the global status of the monitored resource. See the following table for a list of icons and the status each represents. Another option is to hold the mouse over the icon to see a banner pop up with the global status.

Table 20. Global status indicators

Icon	Status	Description
	Consistent	The values for all the resource's attributes monitored by the system are the same on all active nodes within the cluster administrative domain.
	Inconsistent	The values for all the resource's attributes monitored by the system are not the same on all active nodes within the cluster administrative domain.
	Pending	The values of the monitored attributes are in the process of being synchronized across the cluster administrative domain.
	Added	The monitored resource entry has been added to the monitored resource directory in the cluster administrative domain but has not yet been synchronized.
	Ended	The monitored resource is in an unknown state because the cluster administrative domain has been ended, and changes to the resource are no longer being processed.
	Failed	The resource is no longer being monitored by the cluster administrative domain and the MRE should be removed. Certain resource actions are not recommended when a resource is being synchronized by a cluster administrative domain. If the resource represented by an MRE is a system object, it should not be deleted, renamed, or moved to a different library without removing the MRE first. If a resource is deleted, renamed or moved to a different library, the global status for the MRE is Failed and any changes made to the resource on any node after that are not propagated to any node in the cluster administrative domain.

The Type column of the table shows the type of the monitored resource in the solution. See the table below for a list of the possible types of monitored resources.

Table 21. Types of monitored resources

ASP devices	Optical devices
Classes	Subsystem descriptions

Table 21. Types of monitored resources (continued)

Ethernet lines	System environment variables
Job descriptions	System values
Network attributes	User profiles
Network host adapters	Tape devices
Network servers	TCP/IP attributes
NWS configurations	Token ring lines
NWS storage spaces	

Event log information

The **Event Log** section enables you to enter how many days event information should be kept on the system, and see the list of events that have occurred within the number of days specified. The columns shown in the table of logs are Time, Severity, and Information.

Time Shows the date and time the event was logged.

Severity

Shows an icon that represents the urgency of the event. See the table below for a description of each icon.

Information

Gives a brief description of the event.

Click the **High Availability Operator's Guide** link to view or download the *Implementing high availability with the High Availability Solutions Manager – Operator's guide* PDF.

Click **Close** to return to the High Availability Solutions Manager welcome page.

Table 22. Event log status indicators

Severity	Description
	The log contains informational messages.
	The log contains warning messages and should be examined.

For more information, see the following information center topics:

- "Managing a high availability solution" on page 14
- "Working with high availability status messages" on page 15
- "Managing monitored resources" on page 28

Manage high-availability solution - Independent disk pools

The Manage Your High Availability Solution page is used to monitor and manage your high-availability solution. From this page, you also can manage the independent disk pools that your high-availability solution uses.

You can monitor and perform high-availability solution-level actions, monitor and manage high-availability solution resources, and monitor solution events in the event log. The dynamically changing status icons makes this page ideal to use as a *dashboard*.

There are three sections on the Manage Your High Availability Solution page:

- The **Manage Your High Availability Solution** section provides an at-glance view of the status of the high-availability solution and quick access to solution-level actions.
- The **High Availability Solution Resources** section provides a tabbed list of all of the high-availability solution resources. Each tab gives a detailed view of each resource along with possible actions to perform on the resource.
- The **Event Log** section presents the list of events that have occurred in the high-availability solution.

The **Manage Your High Availability Solution** section gives a graphical summary of the status of the high-availability solution. See the table below for a description of the possible statuses.

Table 23. Status of the high-availability solution

Status	Description
	High-availability solution is available for switchover.
	High-availability solution is not available for switchover, but no action is required.
	High-availability solution is not available for switchover. User action is required.

A dropdown menu with the following solution-level actions is included. To perform one of the actions, select it from the menu and click **Go**.

Switchover

Select this action if you want to move the high-availability solution to the backup node. This process makes the backup node the new primary node. Perform this action before a planned outage of the primary node if the application is to remain available during the outage.

Shut Down - Keep Data Available

Select this action if you want to perform system maintenance or a system backup. When the action is performed, your system still is usable, but is not highly available should an unplanned outage occur. Although the application is still available, it will be unavailable in either a planned or unplanned outage.

Shut Down - Make Data Not Available

This action is typically run for system maintenance or if the system is going to be powered down. When this action is performed, your system is accessible only by administrative users. The application is not available on either the primary or backup nodes.

Resume

Select this action if the high-availability solution was shut down and needs to be resumed.

Service Info

Select this action if you want to prepare data to send to a service representative. The system creates a save file named QUSRHASM/QSBSERVICE on the node from which the collection is initiated. The save file should be sent when service information is requested.

Recover Partition

Select this action to recover from a partition state. Cluster Resource Services cannot detect certain failure conditions as a node failure. Recovery from partition state requires operator intervention to determine the cause of the partition state. Selecting this action is only done in certain rare cases. The help for this task describes the process for deciding if selecting this action is necessary.

In addition to the previously listed actions, the Switched disk with geographic mirroring solution has the following possible actions:

Switchover to Different Site

Select this action if you want to run the high-availability solution on a backup node located at a geographically-different site. This action is typically done when all of the nodes at one site need to be taken down for maintenance.

Switchover to Same Site

Select this action if you want to run the high-availability solution on a backup node located at the same geographic location as the current primary node. This process makes the backup node at that site the new primary node.

The **High Availability Solution Resources** section provides a tabbed list of resources including Nodes, Cluster Resource Groups, Monitored Resources, Independent Disk Pools, TCP/IP Interfaces, and Policies. To navigate through this list, click the tab on the left side of the section.

Independent Disk Pools

The high-availability solution resources **Independent disk pools** tab of the Manage Your High Availability Solution page lists all of the independent disk pools in the high-availability solution, as well as the status and capacity of each independent disk pools.

The Name column of the table contains the name of each independent disk pool in the solution. Beside the name is the  button. Click this button to see what actions are available for the independent disk pool. The status of the independent disk pool determines which actions are available. The possible action is:

Work with All Independent Disk Pools

View and work with all independent disk pools, including those which might not be included in the high-availability solution.

In addition to the action listed above, the following actions are available if a solution involving geographic mirroring has been deployed.

Start mirroring

Select this option if mirroring was stopped. This option is not available when the independent disk pool is already being mirrored.

Stop mirroring

Choose this option to stop mirroring. This option is not available if the independent disk pool is not currently being mirrored.

The Status column of the table contains an icon representing the status of the independent disk pool. See the following table for a list of icons and the status each represents. Another option is to hold the mouse over the icon to see a banner pop up with the description.

Table 24. Independent disk pool status indicators

Icon	Status
	Available.
	Available and mirroring is suspended with tracking.
	Available and mirroring is suspended without tracking.
	Available and mirroring is detached.
	Synchronization is in progress.

Table 24. Independent disk pool status indicators (continued)

Icon	Status
	Unavailable.
	Unavailable and mirroring is suspended with tracking.
	Unavailable and mirroring is suspended without tracking.
	Unavailable and mirroring is detached.
	Failed.

The Capacity column of the table shows the current capacity, in gigabytes, of the independent disk pool.

Event log information

The **Event Log** section enables you to enter how many days event information should be kept on the system, and see the list of events that have occurred within the number of days specified. The columns shown in the table of logs are Time, Severity, and Information.

Time Shows the date and time the event was logged.

Severity

Shows an icon that represents the urgency of the event. See the table below for a description of each icon.

Information

Gives a brief description of the event.

Click the **High Availability Operator's Guide** link to view or download the *Implementing high availability with the High Availability Solutions Manager – Operator's guide* PDF.

Click **Close** to return to the High Availability Solutions Manager welcome page.

Table 25. Event log status indicators

Severity	Description
	The log contains informational messages.
	The log contains warning messages and should be examined.

For more information, see the following information center topics:

- “Managing a high availability solution” on page 14
- “Working with high availability status messages” on page 15
- “Managing independent disk pools” on page 30

Manage high-availability solution - TCP/IP interfaces

The Manage Your High Availability Solution page is used to monitor and manage your high-availability solution. From this page, you also can manage the TCP/IP interfaces that your high-availability solution uses.

You can monitor and perform high-availability solution-level actions, monitor and manage high-availability solution resources, and monitor solution events in the event log. The dynamically changing status icons makes this page ideal to use as a *dashboard*.

There are three sections on the Manage Your High Availability Solution page:

- The **Manage Your High Availability Solution** section provides an at-glance view of the status of the high-availability solution and quick access to solution-level actions.
- The **High Availability Solution Resources** section provides a tabbed list of all of the high-availability solution resources. Each tab gives a detailed view of each resource along with possible actions to perform on the resource.
- The **Event Log** section presents the list of events that have occurred in the high-availability solution.

The **Manage Your High Availability Solution** section gives a graphical summary of the status of the high-availability solution. See the table below for a description of the possible statuses.

Table 26. Status of the high-availability solution

Status	Description
	High-availability solution is available for switchover.
	High-availability solution is not available for switchover, but no action is required.
	High-availability solution is not available for switchover. User action is required.

A dropdown menu with the following solution-level actions is included. To perform one of the actions, select it from the menu and click **Go**.

Switchover

Select this action if you want to move the high-availability solution to the backup node. This process makes the backup node the new primary node. Perform this action before a planned outage of the primary node if the application is to remain available during the outage.

Shut Down - Keep Data Available

Select this action if you want to perform system maintenance or a system backup. When the action is performed, your system still is usable, but is not highly available should an unplanned outage occur. Although the application is still available, it will be unavailable in either a planned or unplanned outage.

Shut Down - Make Data Not Available

This action is typically run for system maintenance or if the system is going to be powered down. When this action is performed, your system is accessible only by administrative users. The application is not available on either the primary or backup nodes.

Resume

Select this action if the high-availability solution was shut down and needs to be resumed.

Service Info

Select this action if you want to prepare data to send to a service representative. The system creates a save file named QUSRHASM/QSBSERVICE on the node from which the collection is initiated. The save file should be sent when service information is requested.

Recover Partition

Select this action to recover from a partition state. Cluster Resource Services cannot detect certain failure conditions as a node failure. Recovery from partition state requires operator intervention to determine the cause of the partition state. Selecting this action is only done in certain rare cases. The help for this task describes the process for deciding if selecting this action is necessary.

In addition to the previously listed actions, the Switched disk with geographic mirroring solution has the following possible actions:

Switchover to Different Site

Select this action if you want to run the high-availability solution on a backup node located at a geographically-different site. This action is typically done when all of the nodes at one site need to be taken down for maintenance.

Switchover to Same Site

Select this action if you want to run the high-availability solution on a backup node located at the same geographic location as the current primary node. This process makes the backup node at that site the new primary node.

The **High Availability Solution Resources** section provides a tabbed list of resources including Nodes, Cluster Resource Groups, Monitored Resources, Independent Disk Pools, TCP/IP Interfaces, and Policies. To navigate through this list, click the tab on the left side of the section.

TCP/IP Interfaces

The high-availability solution resources **TCP/IP interfaces** tab of the Manage Your High Availability Solution page lists all of the TCP/IP interfaces in the high-availability solution, as well as the status and role of each TCP/IP interfaces.

The Internet Address column of the table contains the IP address of each interface in the solution. Beside the internet address is the  button. Click this button to see what actions are available for the TCP/IP interface. The status of the TCP/IP interface determines which actions are available. For example, if the interface is already started, the **Start** action is not available. The possible actions are:

Start TCP/IP Interface

Start the TCP/IP interface.

Stop TCP/IP Interface

Stop the TCP/IP interface.

Work with All TCP/IP Interfaces

View and work with all TCP/IP interfaces, including those which might not be included in the high-availability solution.

The Status column of the table contains an icon representing the status of the TCP/IP interface. See the following table for a list of icons and the status each represents. Another option is to hold the mouse over the icon to see a banner pop up with the description.

Table 27. TCP/IP interface status indicators

Icon	Status	Description
	Active	The interface has been started and is running.
	Inactive	The interface has not been started. The interface is not active.
	Starting	The system is processing the request to start this interface.
	Recovery Pending	An error with the physical line associated with this interface was detected by the system. The line description associated with this interface is in the recovery pending state.
	Recovery Canceled	A hardware failure has occurred. The line description associated with this interface is in the recovery canceled state.

Table 27. TCP/IP interface status indicators (continued)

Icon	Status	Description
	Failed	The line description associated with this interface has entered the failed state.
	Failed (TCP)	An error was detected in the IBM TCP/IP Vertical Licensed Internal Code.

The Solution Role column of the table shows the current role of the TCP/IP interface in the solution. See the table below for a list of the possible solution roles and a description of each.

Table 28. Current role of each TCP/IP interface

Solution Role	Description
Device takeover IP	The TCP/IP interface is used for an application cluster resource group.
Data port IP	The TCP/IP interface is used for a device cluster resource group for geographic mirroring.
Node communication IP	The TCP/IP interface is used for cluster node communication.
Server takeover IP	The TCP/IP interface is used for systems associated with the relational database name in the device description for an independent disk pool. The specified address must exist on all nodes in the recovery domain if the cluster resource group is active.

Event log information

The **Event Log** section enables you to enter how many days event information should be kept on the system, and see the list of events that have occurred within the number of days specified. The columns shown in the table of logs are Time, Severity, and Information.

Time Shows the date and time the event was logged.

Severity

Shows an icon that represents the urgency of the event. See the table below for a description of each icon.

Information

Gives a brief description of the event.

Click the **High Availability Operator's Guide** link to view or download the *Implementing high availability with the High Availability Solutions Manager – Operator's guide* PDF.

Click **Close** to return to the High Availability Solutions Manager welcome page.

Table 29. Event log status indicators

Severity	Description
	The log contains informational messages.
	The log contains warning messages and should be examined.

For more information, see the following information center topics:

- “Managing a high availability solution” on page 14
- “Working with high availability status messages” on page 15

- “Managing TCP/IP interfaces” on page 33

Manage high availability solution - Policies

The Manage Your High Availability Solution page is used to monitor and manage your high-availability solution. From this page, you also can manage the policies that your high-availability solution uses.

You can monitor and perform high-availability solution-level actions, monitor and manage high-availability solution resources, and monitor solution events in the event log. The dynamically changing status icons makes this page ideal to use as a *dashboard*.

There are three sections on the Manage Your High Availability Solution page:

- The **Manage Your High Availability Solution** section provides an at-glance view of the status of the high-availability solution and quick access to solution-level actions.
- The **High Availability Solution Resources** section provides a tabbed list of all of the high-availability solution resources. Each tab gives a detailed view of each resource along with possible actions to perform on the resource.
- The **Event Log** section presents the list of events that have occurred in the high-availability solution.

| The **Manage Your High Availability Solution** section gives a graphical summary of the status of the high-availability solution. See the table below for a description of the possible statuses.

Table 30. Status of the high-availability solution

Status	Description
	High-availability solution is available for switchover.
	High-availability solution is not available for switchover, but no action is required.
	High-availability solution is not available for switchover. User action is required.

A dropdown menu with the following solution-level actions is included. To perform one of the actions, select it from the menu and click **Go**.

Switchover

Select this action if you want to move the high-availability solution to the backup node. This process makes the backup node the new primary node. Perform this action before a planned outage of the primary node if the application is to remain available during the outage.

Shut Down - Keep Data Available

Select this action if you want to perform system maintenance or a system backup. When the action is performed, your system still is usable, but is not highly available should an unplanned outage occur. Although the application is still available, it will be unavailable in either a planned or unplanned outage.

Shut Down - Make Data Not Available

This action is typically run for system maintenance or if the system is going to be powered down. When this action is performed, your system is accessible only by administrative users. The application is not available on either the primary or backup nodes.

Resume

Select this action if the high-availability solution was shut down and needs to be resumed.

Service Info

| Select this action if you want to prepare data to send to a service representative. The system
| creates a save file named QUSRHASM/QSBSERVICE on the node from which the collection is
| initiated. The save file should be sent when service information is requested.

Recover Partition

Select this action to recover from a partition state. Cluster Resource Services cannot detect certain failure conditions as a node failure. Recovery from partition state requires operator intervention to determine the cause of the partition state. Selecting this action is only done in certain rare cases. The help for this task describes the process for deciding if selecting this action is necessary.

In addition to the previously listed actions, the Switched disk with geographic mirroring solution has the following possible actions:

Switchover to Different Site

Select this action if you want to run the high-availability solution on a backup node located at a geographically-different site. This action is typically done when all of the nodes at one site need to be taken down for maintenance.

Switchover to Same Site

Select this action if you want to run the high-availability solution on a backup node located at the same geographic location as the current primary node. This process makes the backup node at that site the new primary node.

The **High Availability Solution Resources** section provides a tabbed list of resources including Nodes, Cluster Resource Groups, Monitored Resources, Independent Disk Pools, TCP/IP Interfaces, and Policies. To navigate through this list, click the tab on the left side of the section.

Policies

The high-availability solution resources **Policies** tab of the Manage Your High Availability Solution panel enables you to change the high availability policies at any time.

To change the currently selected option, select the radio button next to the desired option.

Click **OK** to save the new selections.

For more information about each high-availability policy, see “Managing policies” on page 35 in the information center.

Event log information

The **Event Log** section enables you to enter how many days event information should be kept on the system, and see the list of events that have occurred within the number of days specified. The columns shown in the table of logs are Time, Severity, and Information.

Time Shows the date and time the event was logged.

Severity

Shows an icon that represents the urgency of the event. See the table below for a description of each icon.

Information

Gives a brief description of the event.

Click the **High Availability Operator’s Guide** link to view or download the *Implementing high availability with the High Availability Solutions Manager – Operator’s guide* PDF.

Click **Close** to return to the High Availability Solutions Manager welcome page.

Table 31. Event log status indicators

Severity	Description
	The log contains informational messages.
	The log contains warning messages and should be examined.

For more information, see the following information center topics:

- “Managing a high availability solution” on page 14
- “Working with high availability status messages” on page 15

Collect service information

The Collect Service Information page in High Availability Solutions Manager contains a list of substeps to be performed to collect service information, including traces, logs, and flight recorders, from every node in the high-availability solution.

The information is collected into a save file which you can then send to an IBM service representative. Established service procedures should be used when sending the data to IBM. Do this task only when an IBM service representative has asked you to collect this information.

Substeps performed to collect service information

High Availability Solutions Manager performs the following substeps to collect service information from every node in the high-availability solution:

- Creates library QHASMTEMP.
- If the QHASMTEMP library already exists, it is cleared and processing continues.
- Locks library QHASMTEMP. Locking this library ensures only one person at a time can collect service information.
- Uses the Dump Cluster Trace (DMPCLUTRC) command to dump the cluster trace queues to a database file in QHASMTEMP.
- Saves the cluster trace dump file in a save file in QHASMTEMP.
- Saves the output from the level-1 Integrated Solutions Console to a save file.
- Saves the flight recorder from the High Availability Solutions Manager graphical interface to a save file.
- Saves all other internal infrastructure data for the high-availability solution to a save file.
- Saves library QHASMTEMP to a save file called QHASMSDATA in library QGPL. If the save file already exists, it is cleared before the new data is saved.
- Deletes library QHASMTEMP.

At this point, you can send all of the QHASMSDATA save files that were created on each node to IBM.

Field information

The substep table has the following columns.

Column	Description
Substep	Short description of the command or API that will be run.
Estimated time	Estimation of how much time the substep will take to complete. The estimated time changes dynamically to show the best estimate of the remaining time.
Actual time	The actual amount of time that has elapsed since the substep began running.

Column	Description
Status	The status of the substep. Possible values include: Blank - The substep has not yet been run or was undone. Complete - The substep has completed without error. Failed - The substep encountered errors while running. Running - The substep is currently running.
Command/API	The command or API that is currently being processed to complete the substep.

Completion and error messages that are received as the substeps are run are shown in the message area above the **Run Now**, **Cancel**, **Undo**, and **Close** buttons.

Click **Run Now** to begin the substep that the  arrow is pointing to.

Click **Cancel** to return to the Manage Your High Availability Solution page without collecting the service information.

While the substeps are running, click **Cancel** to stop after the current substep has finished running, but before the next substep starts running. After the cancel processing is complete, the following options are available:

- Click **Undo** to undo the previously completed substeps.
- Click **Run Now** to continue running with the next substep.
- Click **Cancel** to return to the Manage Your High Availability Solution page without continuing or undoing any substeps.

Click **Close** to return to the Manage Your High Availability Solution page after all substeps have finished running.

For more information, see the “Collecting service information” on page 21 information center topic.

Verify administrative switchover

Use the High Availability Solutions Manager interface to run the substeps to perform an administrative switchover. The title of this page might be Verify Administrative Switchover, Verify Administrative Switchover at Same Site, or Verify Administrative Switchover at Different Site.

Important: Ensure that the TCP/IP *INETD server is started on the nodes involved in the switchover.

Substeps performed to complete the administrative switchover

High Availability Solutions Manager performs the following substeps to complete the administrative switchover:

- Verifies that the nodes involved in the administrative switchover are active.
- Verifies that the device CRG is active.
- Verifies that the independent disk pool is active.
- Verifies that no duplicate libraries exist on the backup node.
- Changes the device CRG primary node by using the Change CRG Primary (CHGCRGPRI) command. See the Change CRG Primary information center topic for a detailed description of the functions that the CHGCRGPRI command performs. The CHGCRGPRI command does the following for a device CRG:
 - Varies off the independent disk pool. The exit program registered on the QIBM_QDC_VRYEXIT exit point is called. The exit program was registered on this exit point during the Set Up High Availability Environment step of the deployment.

- Ends the server takeover IP address that is associated with the independent disk pool.
- Varies on the independent disk pool on the node that is becoming the new primary node.
- Starts the TCP/IP interface that is associated with the independent disk pool on the new primary node.
- The device CRG exit program calls the program referenced by the QUSRHASM/QSTARTAPP data area that was set up in the Set Up High Availability Environment step of the deployment. This program starts all user applications that reference the independent disk pool.
- Sets the device CRG status to Active if the switchover completes successfully.
- If the solution is Switched Disk with Geographic Mirroring, the recovery domain of the device CRG is modified, as described in the high-availability policies. For example, suppose the policy is to fail over to a node at the same site. After successful switchover, the recovery domain is changed so that the first available backup is the other node at the same site as the new primary node.

The end result of a successful administrative switchover depends on the high-availability solution that was deployed.

If the solution that has been set up is Switched Disk with Logical Partitions or Switched Disk with Two Systems, the independent disk pool moves to the other node. For example, NODE1 is the primary node and has the independent disk pool, and NODE2 is the backup node. During the switchover, the independent disk pool moves to NODE2, NODE2 becomes the primary node, and NODE1 becomes the backup node.

If the solution that has been set up is Cross-Site Mirroring with Geographic Mirroring, the independent disk pools change roles. For example, there are two nodes: NODE1 and NODE2. NODE1 has the production copy of the independent disk pool and NODE2 has the mirror copy. During the switchover, the mirror copy of the independent disk pool becomes the production copy, and the production copy becomes the mirror copy. Thus, NODE2 now has the production copy and NODE1 has the mirror copy.

If the solution that has been set up is Switched Disk with Geographic Mirroring, and the switchover is going to be to the same site, the hardware switches the same way as for the Switched Disk with Logical Partitions or Switched Disk with Two Systems solutions. If the switchover is to a different site, then the mirror copy of the independent disk pool becomes the production copy, and no hardware switches. For example, NODE1 and NODE2 are part of site NEWYORK. NODE1 is the primary node and has the production copy of the independent disk pool. NODE3 is part of site LONDON and has the mirror copy of the independent disk pool. During the switchover, NODE3 becomes the primary node, the production copy attached to NODE1 becomes the mirror copy, and the mirror copy attached to NODE3 becomes the production copy.

If the administrative switchover fails, the independent disk pool is switched back to the original primary system, and an error message displays in the message area.

Field information

The substep table has the following columns.

Column	Description
Substep	Short description of the command or API that will be run.
Estimated time	Estimation of how much time the substep will take to complete. The estimated time changes dynamically to show the best estimate of the remaining time.
Actual time	The actual amount of time that has elapsed since the substep began running.

Column	Description
Status	The status of the substep. Possible values include: Blank - The substep has not yet been run or was undone. Complete - The substep has completed without error. Failed - The substep encountered errors while running. Running - The substep is currently running.
Command/API	The command or API that is currently being processed to complete the substep.

Completion and error messages that are received as the substeps are run are shown in the message area above the **Run Now**, **Cancel**, **Undo**, and **Close** buttons. You can attempt the switchover again after taking corrective action based on the error messages provided.

Click **Run Now** to begin the substep that the  arrow is pointing to.

Click **Cancel** to exit without completing the administrative switchover.

While the substeps are running, click **Cancel** to stop after a substep has finished running, but before the next substep starts running. After the cancel processing is complete, the following options are available:

- Click **Undo** to undo the previously completed substeps.
- Click **Run Now** to continue running with the first substep that has not yet completed.
- Click **Cancel** to return to the previous page without continuing or undoing any substeps. When setting up the high-availability solution, future set up processing will remember which substep is the next substep to be run.

Click **Close** to exit after all substeps have finished running.

For more information, see the “Performing switchovers” on page 19 information center topic.

Recover from partition state

You can run the Recover from Partition State task in High Availability Solutions Manager when a node has failed, but the failure has been detected as a partition condition. The Recover from Partition State page contains a list of substeps that you can run when a partitioned node has failed. If a node is still active, but is partitioned, you should not run this task.

Important: Take extreme caution when running this task. You need to run the substeps to recover from partition state only in rare cases where a node has failed, but the failure has been detected as a partition condition. Unpredictable results can occur if you run this task when a node has not failed. You can check the following conditions to determine if you need to run this task:

- A partition state occurs when the system cannot determine whether another system is down, or is unreachable. It might be a cable problem, which you can resolve by plugging the cable back in. The system self-heals on the partition state when the communications problem is fixed. This type of recovery can take 1 – 15 minutes, based on your parameter settings. If the problem can be resolved in this manner, you should not run the Recover from Partition task.
- If you are on the primary node and the backup node has a status of **Partition**, check the status of the backup system, and check the communication status between the primary node and the backup node. Take necessary corrective actions to return the backup system to a usable state. When the backup system is returned to a usable state and the communications link between the systems is active, then the partition state might recover without operator intervention. If not, you can run the Recover from Partition task from the primary node. You also need to restart clustering on the backup node.

- If you are on the backup node and the primary node has a status of **Partition**, check the status of the primary node. If the primary node is still operational but not communicating to the backup node, take appropriate recovery actions on the communications interface. When the communications interface is active, automatic recovery might take place. If the primary node is no longer operational, then you must decide whether to force a switchover to the backup node by running the Recover from Partition task.

Note: Be careful about forcing a switchover. If the independent disk on the primary node is still varied on, vary off the independent disk on the primary node before running this task. When you have verified that the independent disk is varied off, or if the primary node is no longer operational, and you decide to force a switchover to the backup node, you can run the Recover from Partition task.

For more information about how a cluster is partitioned and how to avoid and recover from partition errors, see Partition errors and Recovering from partition state.

Substeps performed to recover from partition state

If a backup node has a status of **Partition** as seen from another node, the following steps are run when you select **Run Now**:

- The partitioned backup node is set to the **Failed** status. For more details, see the Changing partitioned nodes to failed topic.
- If cross-site mirroring is configured for a two-node solution, tracking will be suspended by using the Start ASP Session (CHGASPSSN) command.

If the primary node has a status of **Partition** as seen from one of the backup nodes, the following steps are run when you select **Run Now**:

- The node is set to the **Failed** status. For more details, see the Changing partitioned nodes to failed topic.
- If cross-site mirroring is configured for a two-node solution, or the primary node site is different from the backup node site, tracking will be suspended by using the CHGASPSSN command.
- The cluster resource group is changed with the Change Cluster Resource Group (CHGCRG) command to make the first backup node the primary node. The node roles are changed, and the backup node owns the hardware.
- The cluster resource group is started with the Start Cluster Resource Group (STRCRG) command.
- The independent disk pool is varied on.

Important: These substeps do not recover the failed node. If the failed system is recovered and usable, you need to perform additional recovery steps to return the system to an active node in the cluster. If the node failed because of a complete system loss, for example, the system disk pool had to be reloaded, see the Recovering a cluster after a complete system loss topic. If the system does not need to be reloaded, manually issue the Start Cluster Node (STRCLUNOD) command from the system whose node was set to **Failed** by the Recover from Partition task.

Field information

The substep table has the following columns.

Column	Description
Substep	Short description of the command or API that will be run.
Estimated time	Estimation of how much time the substep will take to complete. The estimated time changes dynamically to show the best estimate of the remaining time.
Actual time	The actual amount of time that has elapsed since the substep began running.

Column	Description
Status	The status of the substep. Possible values include: Blank - The substep has not yet been run or was undone. Complete - The substep has completed without error. Failed - The substep encountered errors while running. Running - The substep is currently running.
Command/API	The command or API that is currently being processed to complete the substep.

Completion and error messages that are received when the substeps are run are shown in the message area above the **Run Now**, **Cancel**, and **Close** buttons.

Click **Run Now** to begin the substep that the  arrow is pointing to.

Click **Cancel** to return to the Manage Your High Availability Solution page without recovering from the partition state.

While the substeps are running, click **Cancel** to stop after the current substep has finished running, but before the next substep starts running. After the cancel processing is complete, the following options are available:

- Click **Run Now** to continue running with the next substep.
- Click **Cancel** to return to the Manage Your High Availability Solution page.

Click **Close** to return to the Manage Your High Availability Solution page after all substeps have finished running.

Related tasks

Changing partitioned nodes to failed

Sometimes, a partitioned condition is reported when there really was a node outage. This can occur when cluster resource services 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.

Shut down high availability solution - keep data available

The Shut Down High Availability Solution - Keep Data Available page contains a list of substeps to be performed to shut down the high-availability environment but keep the independent disk pool available on the current node. This is typically done when a full-system save or data backup needs to be performed. No switchover will occur as a result of this action.

Substeps performed to shut down the high-availability environment while keeping data available

High Availability Solutions Manager performs the following substeps to shut down the high-availability environment and keep the data available:

- If you have chosen the Switched Disk with Geographic Mirroring or Cross Site Mirroring with Geographic Mirroring solution, geographic mirroring is suspended.
- The device cluster resource group is ended.
- All of the cluster nodes in the high-availability solution are ended.

When the substeps have completed successfully, the high-availability solution is shut down, and the independent disk pool is still available.

Field information

The substep table has the following columns.

Column	Description
Substep	Short description of the command or API that will be run.
Estimated time	Estimation of how much time the substep will take to complete. The estimated time changes dynamically to show the best estimate of the remaining time.
Actual time	The actual amount of time that has elapsed since the substep began running.
Status	The status of the substep. Possible values include: Blank - The substep has not yet been run or was undone. Complete - The substep has completed without error. Failed - The substep encountered errors while running. Running - The substep is currently running.
Command/API	The command or API that is currently being processed to complete the substep.

- | Completion and error messages that are received as the substeps are run are shown in the message area
- | above the **Run Now**, **Cancel**, and **Close** buttons.

Click **Run Now** to begin the substep that the  arrow is pointing to.

Click **Cancel** to return to the Manage Your High Availability Solution page without shutting down the high-availability environment.

While the substeps are running, click **Cancel** to stop after the current substep has finished running, but before the next substep starts running. After the cancel processing is complete, the following options are available:

- Click **Run Now** to continue running with the next substep.
- | • Click **Cancel** to return to the Manage Your High Availability Solution page.

Click **Close** to return to the Manage Your High Availability Solution page after all substeps have finished running.

For more information, see the “Shutting down a high availability solution with data available” on page 17 information center topic.

Shut down high availability solution - make data not available

The Shutdown High Availability Solution – Make Data Not Available page contains a list of substeps to be performed to shut down the high-availability environment and also make the independent disk pool unavailable. Typically, this is done in preparation of an IPL of the primary node, where you do not need to perform an administrative switchover.

Substeps performed to shut down the high-availability environment

High Availability Solutions Manager performs the following substeps to shut down the high-availability environment and make the independent disk pool unavailable:

Note: The high-availability solution policies are not consulted during this task. The independent disk pool will not be switched over to a backup node. If you want to switch over the high-availability solution, this is not the task to use.

- The independent disk pool is varied off.
 - This substep automatically ends any user jobs that are associated with the independent disk pool.

- During this substep, the exit program registered on the QIBM_QDC_VRYEXIT exit point is called. This exit program invokes any program that is defined in the QUSRHASM/QSHUTDOWN data area.
- The device cluster resource group is ended.
- All of the cluster nodes in the high-availability solution are ended.

When all of the substeps have completed successfully, the high-availability solution is shut down and the independent disk pool is not available on any nodes.

Field information

The substep table has the following columns.

Column	Description
Substep	Short description of the command or API that will be run.
Estimated time	Estimation of how much time the substep will take to complete. The estimated time changes dynamically to show the best estimate of the remaining time.
Actual time	The actual amount of time that has elapsed since the substep began running.
Status	The status of the substep. Possible values include: Blank - The substep has not yet been run or was undone. Complete - The substep has completed without error. Failed - The substep encountered errors while running. Running - The substep is currently running.
Command/API	The command or API that is currently being processed to complete the substep.

- | Completion and error messages that are received as the substeps are run are shown in the message area above the **Run Now**, **Cancel**, and **Close** buttons.

Click **Run Now** to begin the substep that the  arrow is pointing to.

Click **Cancel** to return to the Manage Your High Availability Solution page without shutting down the high-availability environment.

While the substeps are running, click **Cancel** to stop after the current substep has finished running, but before the next substep starts running. After the cancel processing is complete, the following options are available:

- Click **Run Now** to continue running with the next substep.
- | • Click **Cancel** to return to the Manage Your High Availability Solution page.

Click **Close** to return to the Manage Your High Availability Solution page after all substeps have finished running.

For more information, see the “Shutting down a high availability solution with data unavailable” on page 18 information center topic.

Resume high availability solution

The Resume High Availability Solution page in High Availability Solutions Manager contains a list of substeps to be performed to resume the high-availability solution and make the data available. The high availability solution is set up to automatically resume when the system restarts after an IPL. This task is typically used only after you have shut down the high availability solution and need to resume it manually.

Substeps performed to resume the high-availability solution

High Availability Solutions Manager performs the following substeps to resume the high-availability solution:

- The cluster nodes are started.
- The cluster resource groups are started.
- If the independent disk pool is varied off, it will be varied on.
- If the high-availability solution is Switched Disk with Geographical Mirroring or Cross-Site Mirroring with Geographical Mirroring, mirroring of the independent disk pool is resumed.

Field information

The substep table has the following columns.

Column	Description
Substep	Short description of the command or API that will be run.
Estimated time	Estimation of how much time the substep will take to complete. The estimated time changes dynamically to show the best estimate of the remaining time.
Actual time	The actual amount of time that has elapsed since the substep began running.
Status	The status of the substep. Possible values include: Blank - The substep has not yet been run or was undone. Complete - The substep has completed without error. Failed - The substep encountered errors while running. Running - The substep is currently running.
Command/API	The command or API that is currently being processed to complete the substep.

- | Completion and error messages that are received as the substeps are run are shown in the message area above the **Run Now**, **Cancel**, and **Close** buttons.

Click **Run Now** to begin the substep that the  arrow is pointing to.

Click **Cancel** to return to the Manage Your High Availability Solution page without resuming the high-availability environment.

While the substeps are running, click **Cancel** to stop after the current substep has finished running, but before the next substep starts running. After the cancel processing is complete, the following options are available:

- Click **Run Now** to continue running with the next substep.
- | • Click **Cancel** to return to the Manage Your High Availability Solution page.

Click **Close** to return to the Manage Your High Availability Solution page after all substeps have finished running.

For more information, see the “Resuming a high availability solution” on page 19 information center topic.

| No help is available

- | No help is available for this page.

|

Appendix. 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:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, 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,
- | IBM License Agreement for Machine Code, 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.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

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

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_. All rights reserved.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Programming interface information

This "Implementing high availability with the solution-based approach" publication documents intended Programming Interfaces that allow the customer to write programs to obtain the services of IBM i5/OS.

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

i5/OS
IBM
IBM (logo)
System i
System i5
System Storage
TotalStorage
FlashCopy

- | 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.
- | Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
- | Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

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.



Printed in USA