IBM

# IBM i 6.1 Independent ASPs:
## A Guide to Quick Implementation of Independent ASPs

- Take advantage of the IBM i 6.1 support for independent ASPs
- Improve your uptime by using IASPs
- Easily install and configure IASPs for your environment

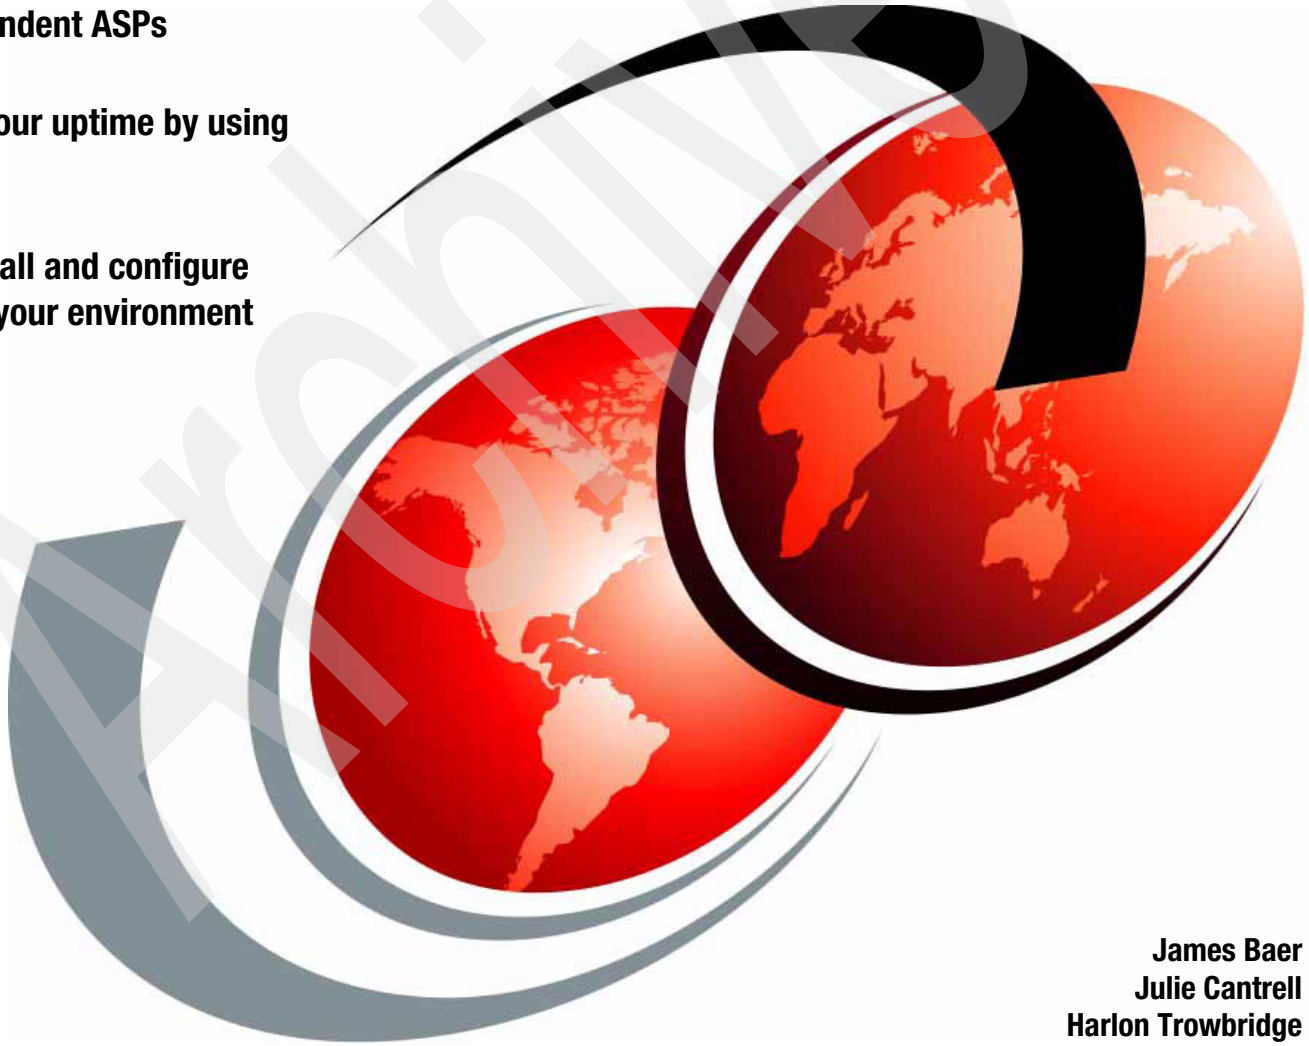



James Baer
Julie Cantrell
Harlon Trowbridge

Redbooks

ibm.com/redbooks

International Technical Support Organization

# IBM i 6.1 Independent ASPs: A Guide to Quick Implementation of Independent ASPs

December 2009



SG24-7811-00

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



**First Edition (December 2009)**

This edition applies to Version 6, Release 1, Modification 0 of IBM i (5761-SS1) and related licensed program products.

# Contents



 

# Figures

# Tables





 

# 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 give you any license to these patents. You can send license inquiries, in writing, to:
*IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.*

**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.

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 measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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

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

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

# Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX®
AnyNet®
AS/400®
DB2®
Distributed Relational Database Architecture™
DRDA®
i5/OS®
IBM®
iSeries®
OS/400®
Power Systems™
PowerHA™
PowerVM™
POWER®
Redbooks®
Redbooks (logo) ®
System i®
xSeries®

The following terms are trademarks of other companies:

Java, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

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

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

# Preface

This IBM® Redbooks® publication explains how to configure and manage independent disk pool (IASP) functionality of IBM i 6.1. It is designed to help IBM technical professionals, business partners, and customers understand and implement independent disk pools in the IBM i 6.1.

In addition, this publication provides the background information that is necessary to plan, implement, and customize this functionality to your particular environment. It provides guidance on running user applications with either application data or most application objects residing in an independent disk pool. Considering that you can also use independent disk pools in a cluster environment, this publication shows you the basic steps to make your independent disk pool switchable between two Power Systems™ servers or a single server with multiple LPARs.

Independent auxiliary storage pools have many business and technical advantages for Power Systems using IBM i. Not only are independent auxiliary storage pools (IASPs) easy to create and maintain, most applications can use them by simple work management changes. IASPs can provide immediate benefits to your enterprise.

## The team who wrote this book

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

**James Baer** is a retired IBM employee now working as an independent consultant in Missouri. He worked for over 31 years for IBM as a Systems Engineer and a Client Representative. After retiring from IBM he worked for Huber and Associates, an IBM Business Partner, in Springfield, MO. Since 2006 he has worked as an independent consultant, helping IBM i customers install, migrate, and upgrade Power Systems servers.

**Julie Cantrell** is an HA and Power Systems Specialist with Mainline Information Systems. For 26 years Julie has specialized in IBM SMB Systems and IBM i architecture. Prior to joining Mainline, Julie was an IS Director and Consultant for several Chicago-based corporations. After seven years with Lakeview Technology, she joined the Mainline team with a specialty in HA for Power Systems.

**Harlon Trowbridge** is a Senior Software Engineer at the IBM Innovation Center in Dallas, Texas. He has over 35 years of experience in IBM systems management, IBM Business Partner support, and application design and development. He has specialized in S/38, AS/400®, and IBM i architecture for the last 28 years. Harlon implemented the IBM Linux® on iSeries® Sandbox for Business Partners environment and is currently the Systems Management Team Lead for the IBM Virtual Loaner Program, which provides cloud-like access to IBM Power systems for IBM Business Partners.

The project that produced this publication was managed by:
**Scott Vetter**, PMP

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

Ann Lund, Don Brennan, Rich Conway, and Dave Bennin
International Technical Support Organization, Poughkeepsie Center

Niel Brieger
IBM ISV and Developer Relations, IBM Innovation Center - Dallas

Mark Roy
IBM i Specialist, Sysarb Pty Ltd, Melbourne, Australia

Chuck Stupca
IBM Rochester

## Become a published author

Join us for a two- to six-week residency program! Help write a book dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You will have the opportunity to team with IBM technical professionals, Business Partners, and Clients.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you will develop a network of contacts in IBM development labs, and increase your productivity and marketability.

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

**ibm.com**/redbooks/residencies.html

## Comments welcome

Your comments are important to us!

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

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

  **ibm.com**/redbooks

- Send your comments in an e-mail to:

  redbooks@us.ibm.com

- Mail your comments to:

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

1

# Introduction to independent disk pools (IASPs)

Independent auxiliary storage pools have many business and technical advantages for Power Systems using IBM i. Not only are independent auxiliary storage pool (IASPs) easy to create and maintain, most applications can use them with simple work management changes. IASPs can provide immediate benefits to your enterprise.

An independent disk pool, or independent auxiliary storage pool, is a collection of disk units that can be brought online or taken offline, independent of the rest of the storage on a system, including the system ASP, user ASPs, and other independent disk pools. The act of making the contents of an independent disk pool visible is known as *varying on* the independent disk pool. This vary-on step is similar to the IPL processing steps, except that it is limited to the objects residing within the independent disk pool. An independent disk pool can be either:

**Non-switchable** An independent disk pool that is used with a single system, also known as an independent auxiliary storage pool. The contents of the independent disk pool can dynamically be made available or unavailable for use on that system only using *vary on/off*.

**Switchable** An independent disk pool used across multiple partitions is referred to as a switchable independent disk pool. When you switch an independent disk pool to another partition, the entire contents of the pool can be accessed from that system without having to restart (IPL) the entire system. However, a *vary on* (mini-IPL) does take place. It is *not* instantaneous.

Enhancements made to IBM i make using independent disk pools an attractive option for many clients who are looking for higher levels of availability and server consolidation.

This chapter looks at independent disk pools and introduces enhancements included in IBM i 6.1. It also explains how independent disk pools work, the benefits of implementing independent disk pools, and some restrictions and considerations.

Throughout this book, the terms *independent disk pool, independent auxiliary storage pool,* and *IASP* are synonymous. See 1.3, “Terminology” on page 4, for more information about the terminology used in this book.

**Note:** Although clustering support is required to enable switchable disk pools, this publication does not go into detail about creating and managing clusters. Instead, refer to *Implementing PowerHA for IBM i*, SG24-7405, for more information about clustering.

## 1.1 Business benefits

A business need or requirement generally drives the upgrade of every system. When carrying out a plan for independent disk pools, keep in mind the business needs. In each step of the planning process, you must satisfy the business need before you approve and adopt the plan. This constant verification adds to support and justification for your environment. It may also assist with determining and moving to higher levels of availability.

The advantages of implementing independent disk pools can include:

- Server consolidation
- Workload balancing
- Isolation of historical and archived data
- Isolation of application data
- Independent save/restore
- Application maintenance by IASP
- Multiple application versions
- Support test and development environments
- Availability, including the following attributes:
  - Disk isolation
  - Application availability
  - Alternative to data replication

## 1.2 Technical advantages

An additional method of Power Systems server availability is enabled through the use of independent disk pools. This section lists some of the advantages of using independent disk pools:

- For disk drives in the independent disk pool, device parity protection can be stopped or started from within IBM i 6.1.

  For regular ASPs, stopping and starting device parity protection is a dedicated service tools (DST) function.

- For disk drives in the independent disk pool, mirroring can be turned on and off from within IBM i 6.1.

  For regular ASPs, stopping and starting mirroring is a DST function.

- Independent disk pools enable a higher level of availability without needing to buy a duplicate set of disks for the backup system.

- The contents of a switchable independent disk pool can be made available to the backup system without any kind of replication or extra replication software, provided that some objects are already available on the target system.

- It is not necessary to maintain multiple copies of data, programs, and other objects.

  Multiple copies of objects is a function of replication.

- There is minimal additional system overhead with independent disk pools.

  Replication requires more CPU cycles when replicating to a backup system.

- Reclaim storage by independent disk pool.

- No network traffic is associated with independent disk pools.

  Replication across a LAN or WAN involves network traffic.

- There is less work for system functions such as IPL, reclaim storage, and some save operations.

  In a single system environment, an independent disk pool can be used to store certain data offline except for periods when it is actually needed. The isolation provided by storing data offline means that there is less work necessary for system functions.
- Objects are not *in flight* in the event of a failure.

  With replication, journal entries can become *trapped* on the source system at the time of failure and do not arrive at the target machine.
- When used in conjunction with journaling, independent disk pools can be used as archives, or save/restore repositories, without needing high-availability business partner (HABP) software.

## 1.3 Terminology

Throughout this book, the terms *independent disk pool*, *independent auxiliary storage pool*, and *IASP are synonymous*. While this may seem confusing to the reader, perhaps a brief discussion of the evolution of disk pools and IBM i will assist you.

The concept of single-level storage uses disk storage as the permanent addressed location of all objects on the system. When executing programs or accessing data files, only the necessary pieces are brought into memory on a *as needed* basis. This technique is referred to as *paging*. The term *main storage* is used to refer to temporary memory space, while *auxiliary storage* is used to refer to the permanent disk space.

As the need to segregate groups of programs and data on the same system emerged, the concept of pools developed and was included as part of the operating system. The pools were referred to as auxiliary storage pools (ASPs) because they pertained to areas of auxiliary storage (disk space). The new command structures within the operating system used the letters ASP when referring to the auxiliary storage pools.

Enhancements to the concept of pools has led to independent auxiliary storage pools. These are pools that can be brought online, taken offline, and accessed, independently of the other pools on the system. They can even be logically switched between another system or logical partition.

Because these unique pools now refer to one or more disk drives, the term *independent disk pool* has emerged to reference them. The newer tools, such as IBM Systems Director Navigator for i5/OS®, use the term *disk pools* instead of the older term, *independent auxiliary storage pool*.

Throughout this manual we try to use the term that matches the topic that we are discussing. In our general discussion and the sections discussing use of IBM Systems Director Navigator for i5/OS, we use the newer term *disk pools*. While discussing items within the IBM i operating system we use auxiliary storage pools (ASP) to be consistent with the *green screen* command entry screens.

## 1.4 What is new in IBM i 6.1

Several enhancements have been made to independent disk pool support in IBM i 6.1. These enhancements include:

- Support of job queues in independent disk pools

  This allows applications to run with fewer changes. Job queue entries are lost (not persistent) when varying off/on independent disk pools.

- Support for associating an independent disk pool with a subsystem description

  Independent disk pools must be available before the subsystem is activated.

- Quiesce function

  Transactions and operations are suspended to ensure that as much in-flight data as possible is written to disk.

- Encrypted independent disk pools

  This protects data transmission to and from the disk drive and in the cross-site mirroring environment (only when the data being mirrored is on an encrypted independent disk pool). This protects disk data in case of theft, return, or resale. This must be specified at independent disk pool creation time and cannot be applied to existing independent disk pools.

## 1.5 Types of disk pools

**Note:** The terms *disk pool* and *auxiliary storage pool* are synonymous.

The following definitions describe the different types of disk pools:

- System disk pool (disk pool 1)

  This contains IBM i 6.1 and licensed program products (LPPs) and any user objects.

- Basic disk pool (disk pools 2 to 32)

  These isolate objects from the objects stored in the system disk pool. Basic disk pools are always assessable whenever the system is up and running.

- Primary disk pool

  This is an independent disk pool that defines a collection of directories and libraries and may have other secondary disk pools associated with it. Primary disk pools and any associated secondary pools can be taken offline or brought online independent of system activity on other disk pools.

- Secondary disk pool

  This is an independent disk pool that defines a collection of directories and libraries and must be associated with a primary disk pool.

- User-defined file system

  This is an independent disk pool that contains only user-defined file systems. It cannot be a member of a disk pool group unless it is converted to a primary or a secondary disk pool.

- Disk pool groups

  Disk pool groups consist of a primary disk pool and zero or more secondary disk pools. Each disk pool is independent in regard to data storage, but in the disk pool group they

combine to act as one entity. Making disk pools available to the users is accomplished by using the disk pool group name.

Figure 1-1 shows these types.

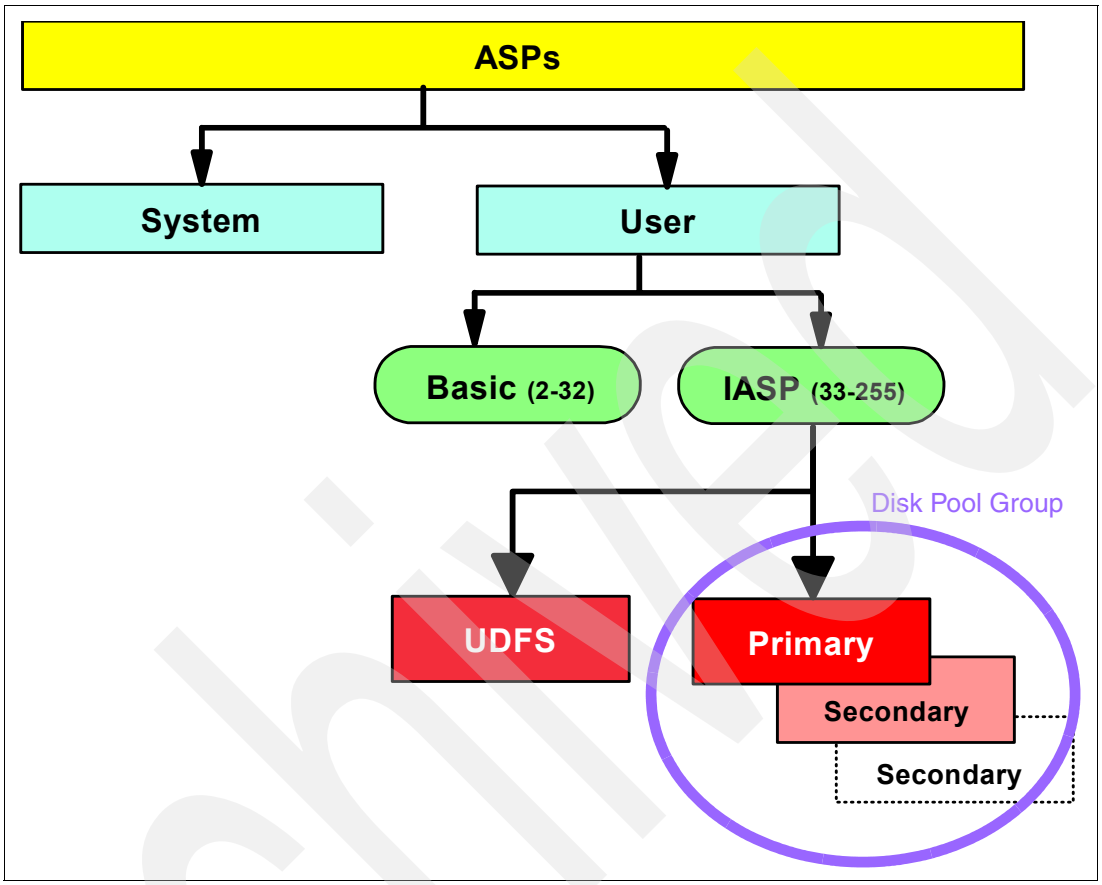


*Figure 1-1 Disk pool explanation*

## 1.6 Independent disk pool object support for IBM i 6.1

As a general rule, IBM i 6.1 objects and LPPs may not exist in an independent disk pool. See Appendix D, “Supported and unsupported objects” on page 197, for a complete list.

## 1.7 Positioning independent disk pools

There are several possible configurations, but only two basic environments in which independent disk pools can be used:

- A single-system environment with a single Power Systems server (not partitioned)
- A multisystem or multi-partition environment managed by an Power Systems server cluster

## Single-partition environment

In a single-partition environment, you can take independent disk pools offline or make them unavailable, independent of other disk pools. You can also bring the independent disk pool online or make it available while the system is active without performing an initial program load (IPL).

Figure 1-2 shows an example of multiple databases that reside in independent disk pools. This example has independent disk pools for the payroll data, order entry data, and data for companies 1, 2, and 3. The actual application code could reside in the system disk pool (ASP) or another disk pool (either a user ASP or another IASP).

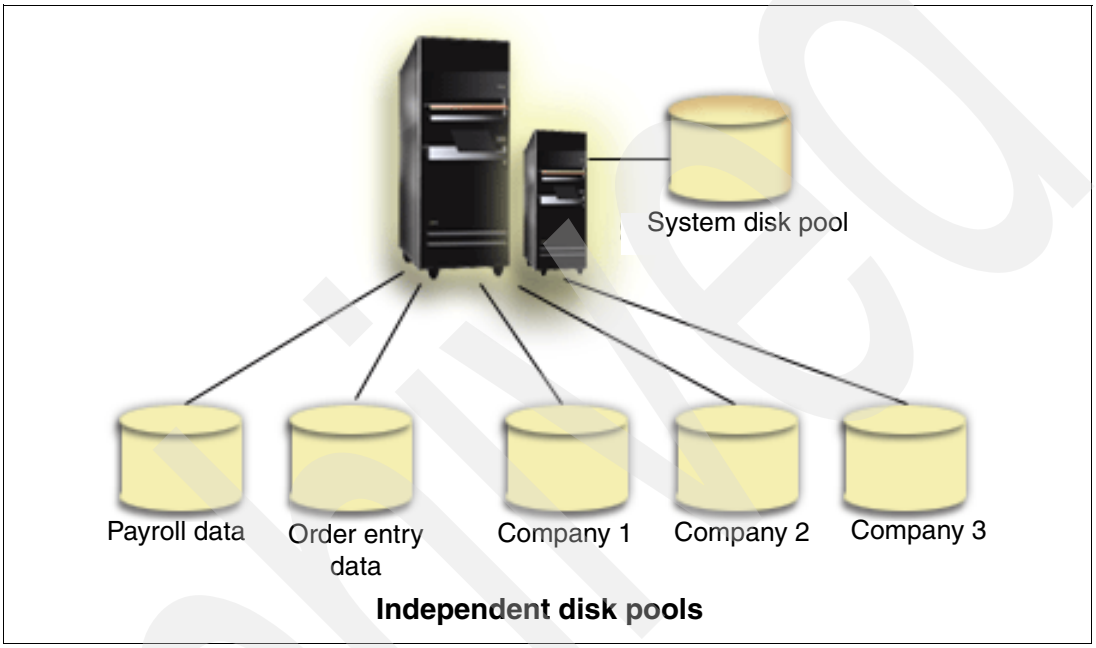


*Figure 1-2 Independent disk pools*

A typical use of independent disk pools as shown in this example is for server consolidation of multiple branch office or store systems. Corporate data can reside in the other independent disk pools. Segmenting your databases in this manner allows for greater control and flexibility.

2

# Creating an IASP

This chapter explains how to build an independent auxiliary storage pool (IASP) using the IBM Systems Director Navigator for i5/OS. When creating an IASP the terminology used is *disk pools*.

## 2.1 Disk pool (IASP) creation prerequisites

Prior to creating an IASP, you must complete the following prerequisites:

- IBM Systems Director Navigator for i5/OS must have a connection to the systems that are using the independent disks to create the IASP.
- The user must have *IOSYSCFG and *ALLOBJ authority.
- A minimum of one unconfigured disk unit (physical or virtual).

**Important:** Be sure to follow the steps in Appendix A, “Prerequisite steps” on page 187, to enable IBM Systems Director Navigator for i5/OS disk pool capabilities.

## 2.2 Creating a primary disk pool (IASP)

Creating an IASP begins with defining the disk pool and the disk units to be included. You can create a disk pool and add disk units to it by using the IBM Systems Director Navigator for i5/OS.

To create a new disk pool group, follow these steps:

1. In IBM Systems Director Navigator for i5/OS, expand **Configuration and Service** → **New Disk Pool** (Figure 2-1).

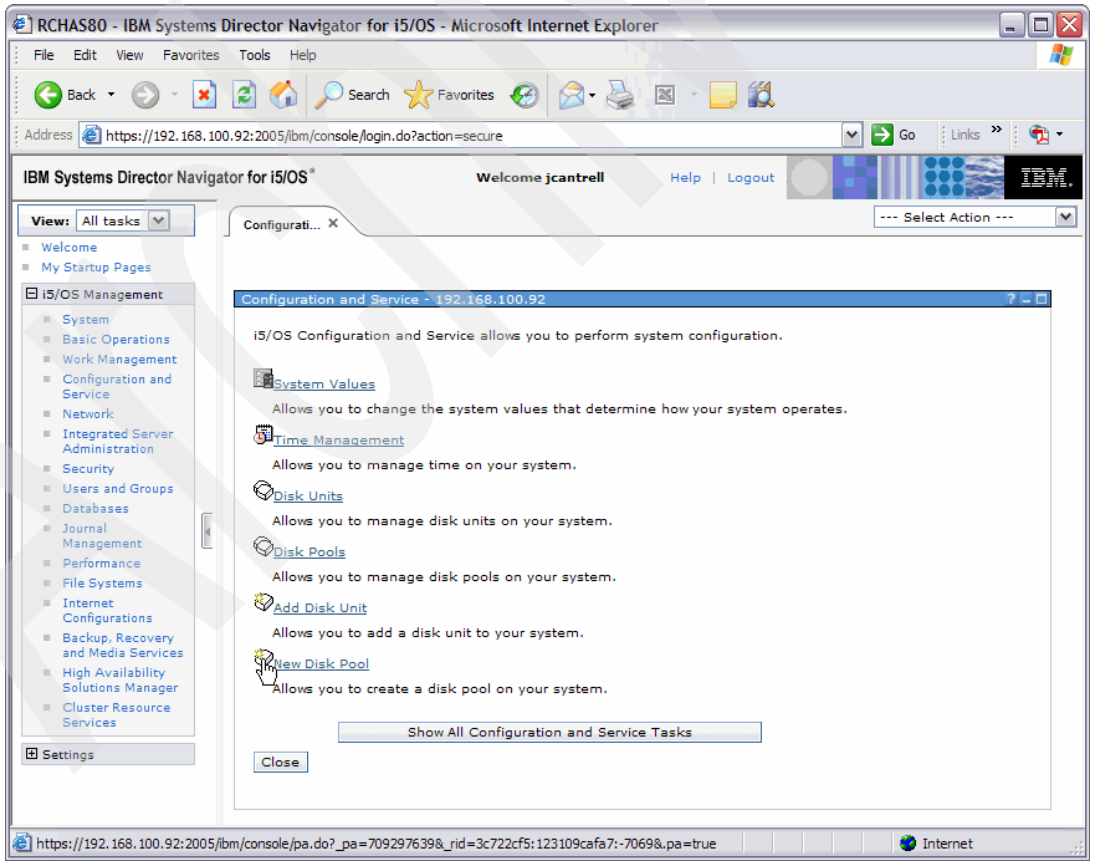


*Figure 2-1 Director Navigator → expand i5/OS Management → Configuration and Service*

2. At the New Disk Pool - Welcome display (Figure 2-2), select **Next**.

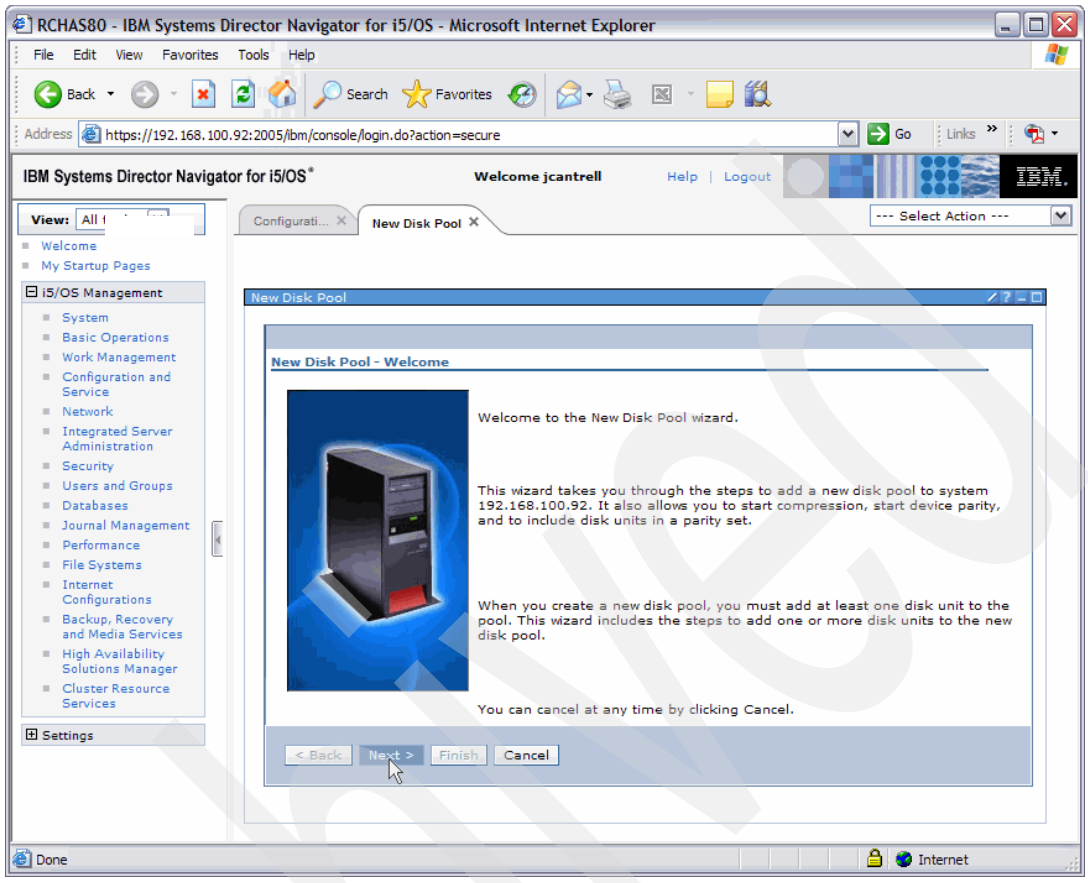


*Figure 2-2 New Disk Pool - Welcome → Next*

3. Use the Type of disk pool drop-down to select **Primary** (Figure 2-3). Select **Go**.

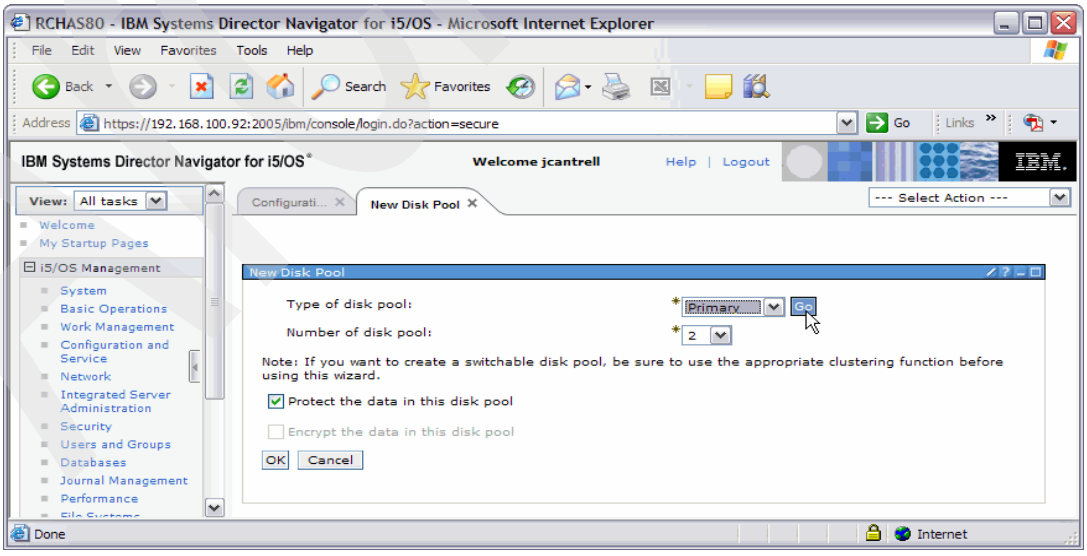


*Figure 2-3 New Disk Pool → Type of disk pool → Go*

4. Enter the name of the disk pool and select **OK** (Figure 2-4).

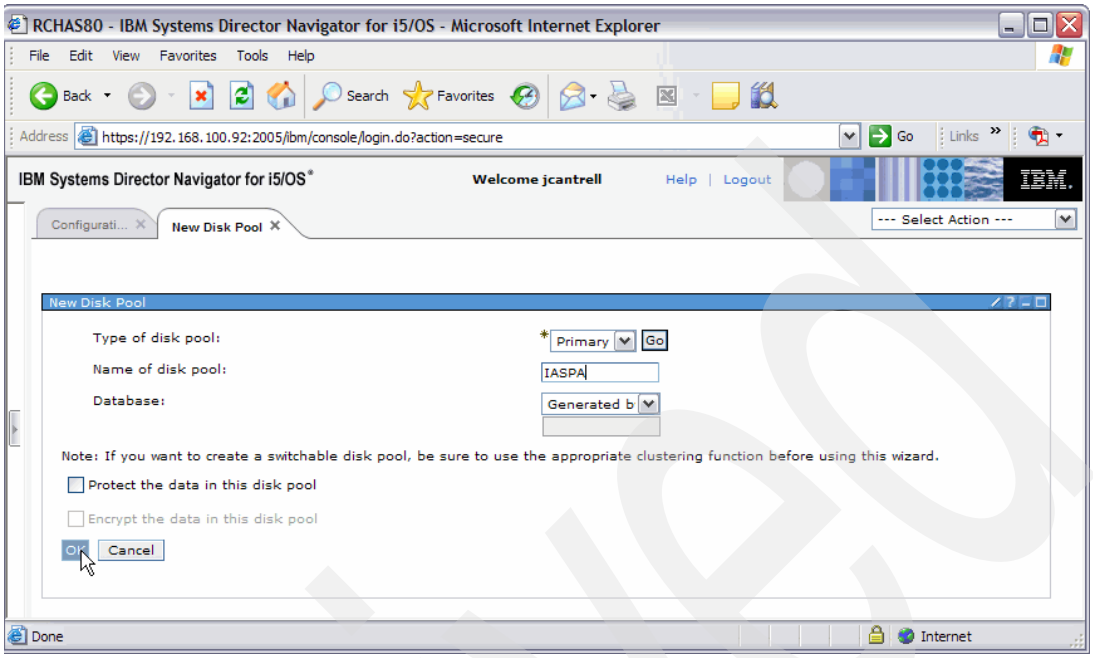


*Figure 2-4 New Disk Pool → Name of disk pool → OK*

5. Disk units are added to the disk pool using the New Disk Pool - Select Disk Pool window (Figure 2-5). On the right side of the window, under Add disk units to these pools, select the disk pool that you created. Select **Next**.

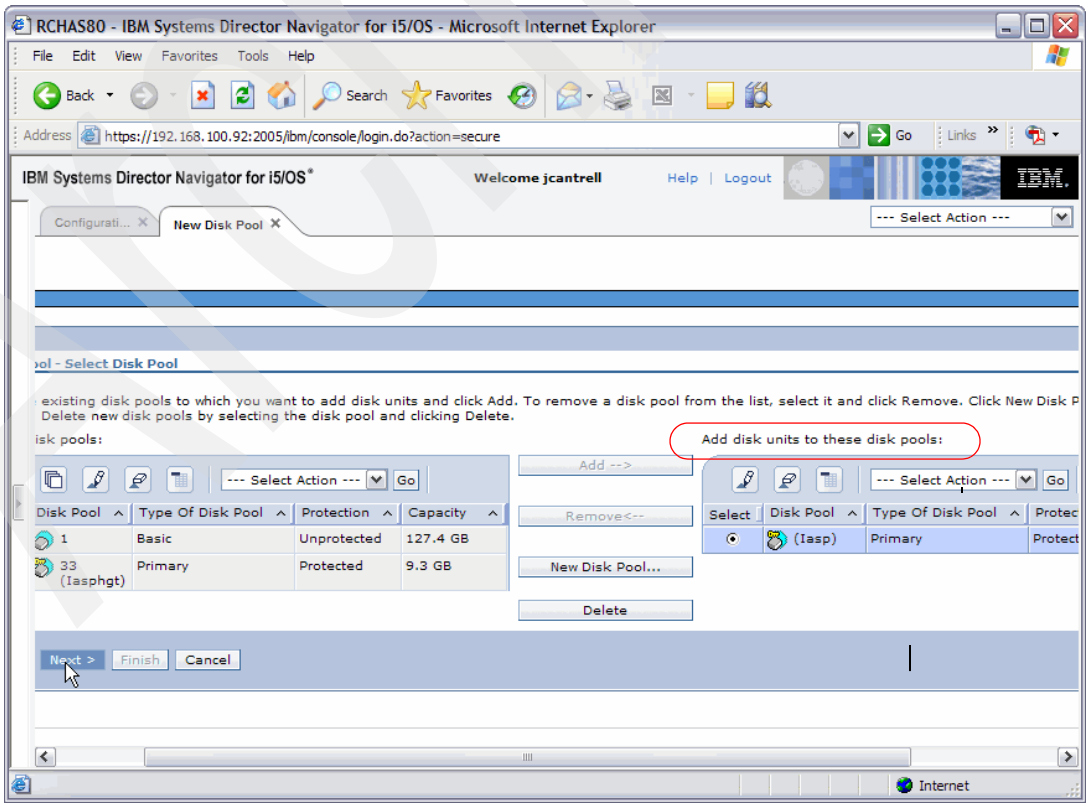


*Figure 2-5 New Disk Pool - Select Disk Pool → Add disk units to these pools → Next*

6. Select **Add Disks to be Mirrored** (Figure 2-6).

> **Note:** Disks to be mirrored requires pairs of disk. Parity protected disk requires physical drives. We are using virtual drives. If you are creating a disk pool that is not protected select Add Disks.

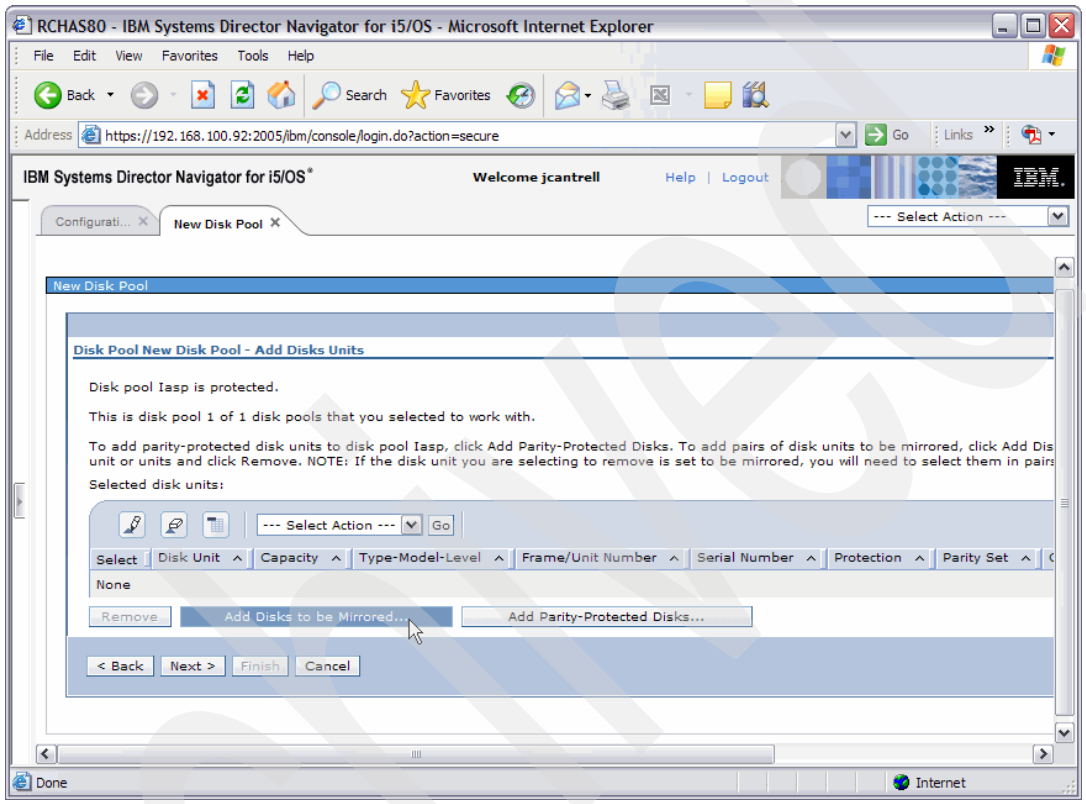


*Figure 2-6 Disk Pool: New Disk Pool - Add Disk Units → Add Disks to be MIrrored*

7. When the disk list appears on the Disk Pool IASP - Add Disks to be Mirrored window (Figure 2-7), select the appropriate disk units, then select **Add**.

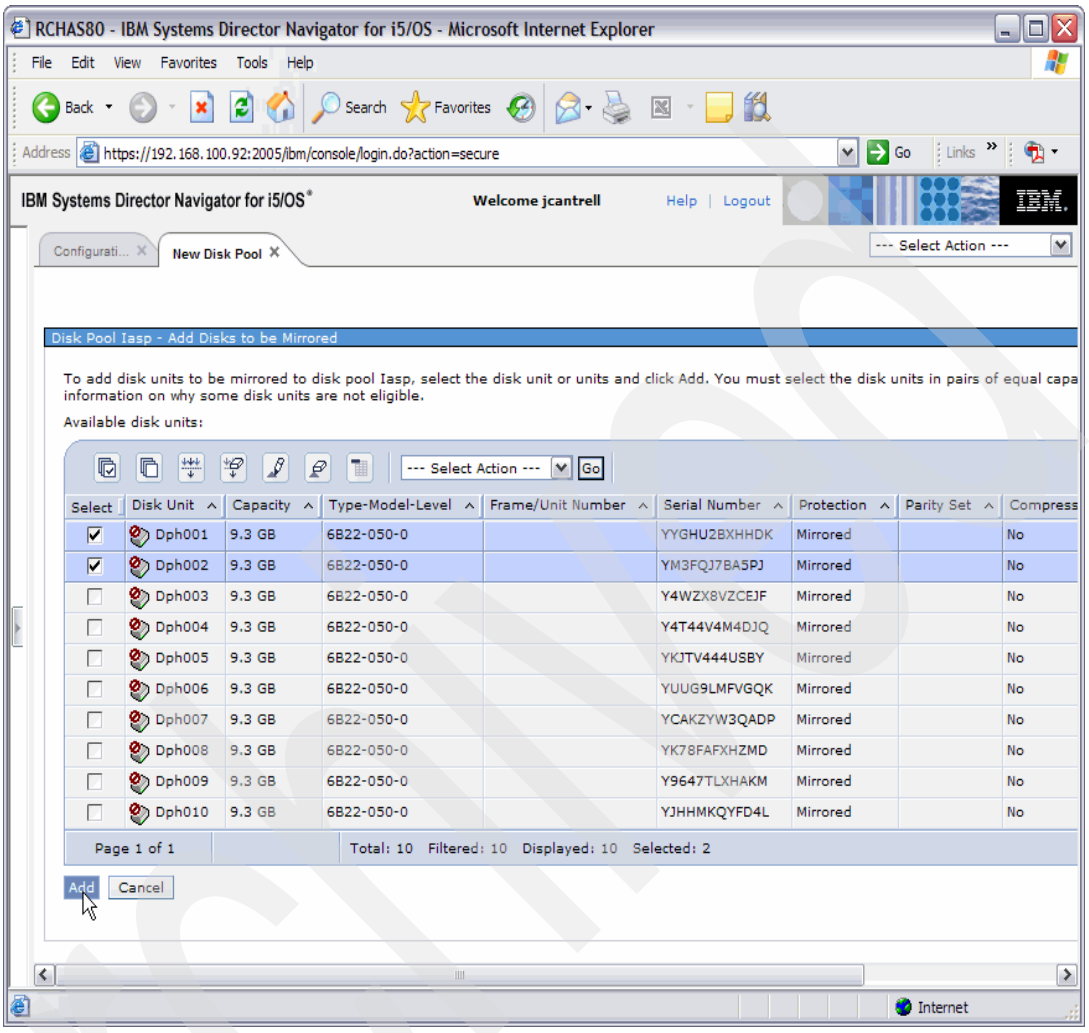


*Figure 2-7 Disk Pool IASP - Add Disks to be Mirrored → OK*

8. When the disk that you selected appears in the Disk Pool IASP - Add Disks Units (Figure 2-8) select **Next**.

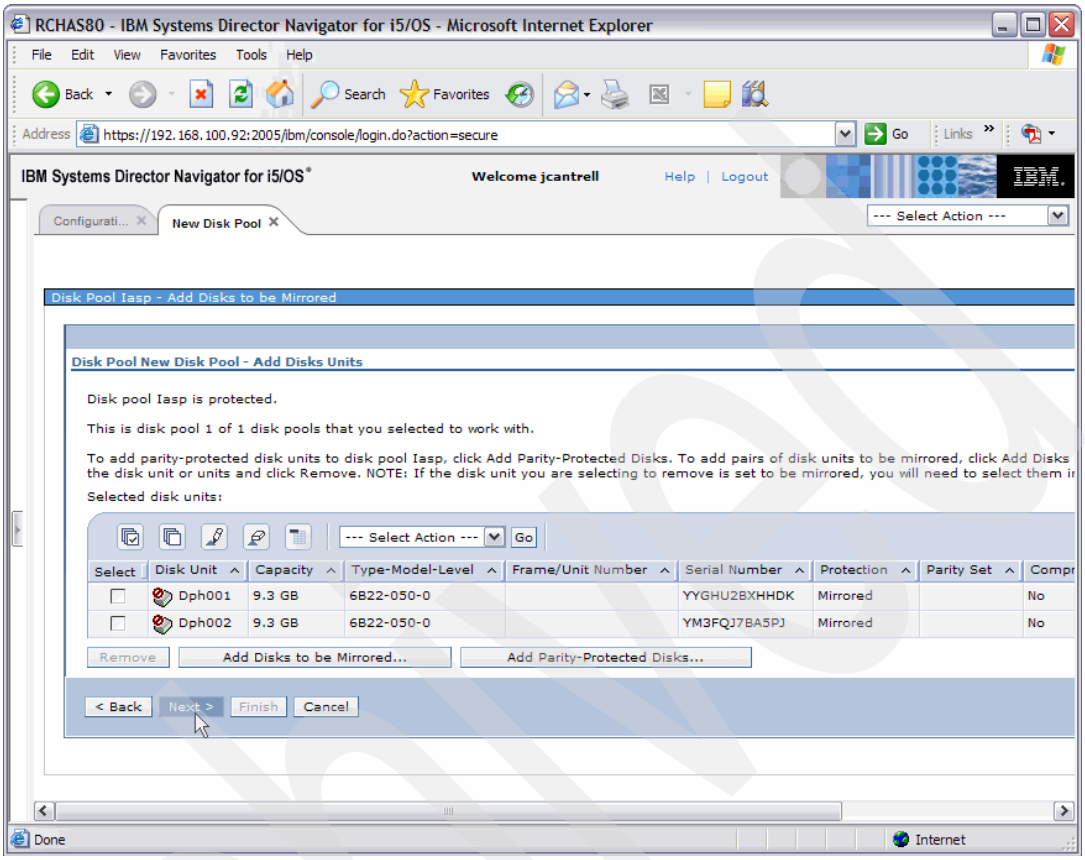


*Figure 2-8 Disk Pool IASP - Add Disks Units → Next*

9. On the New Disk Pool - Summary window (Figure 2-9), select **Finish** to confirm the addition of the disks.

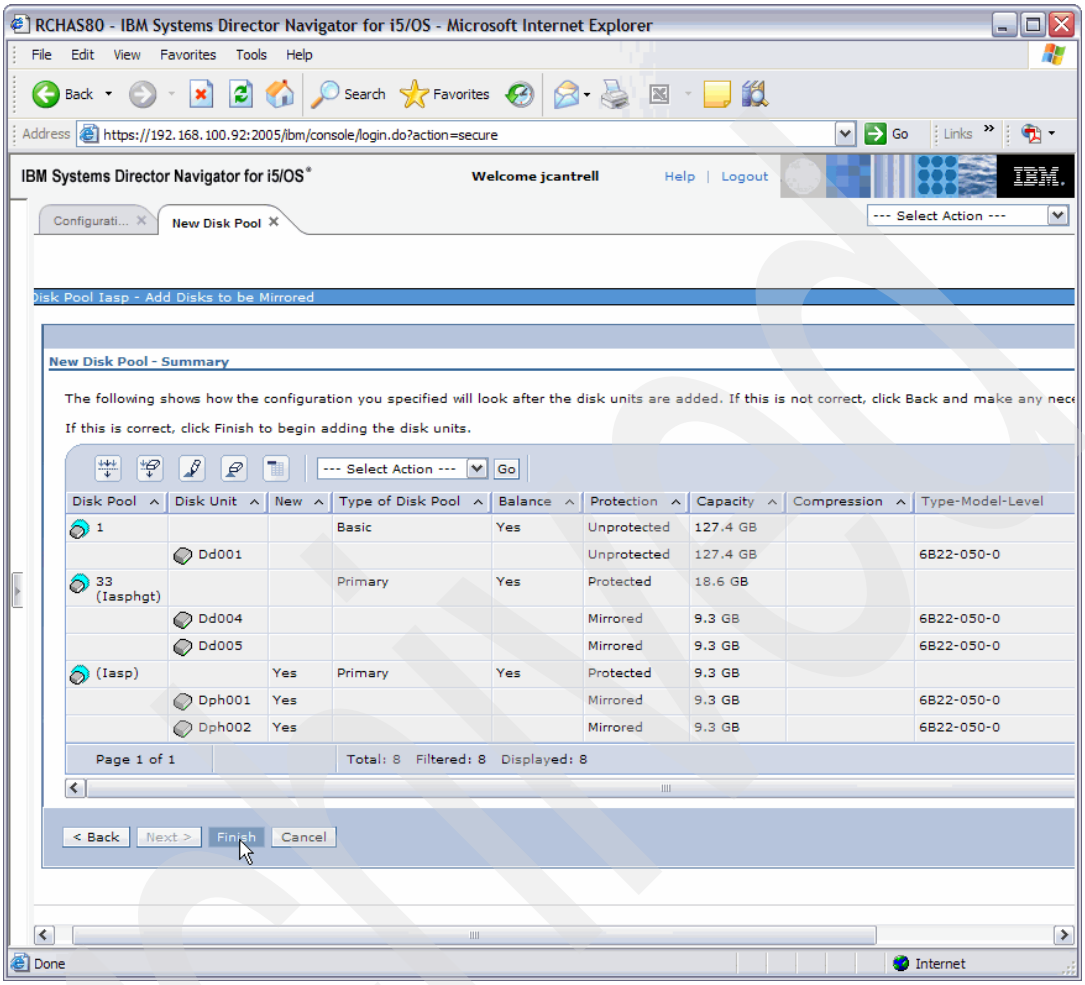


*Figure 2-9 New Disk Pool – Summary → Finish*

10. The New Disk Pool Status window displays indicating the current action. Select **Refresh** to update this window (Figure 2-10).

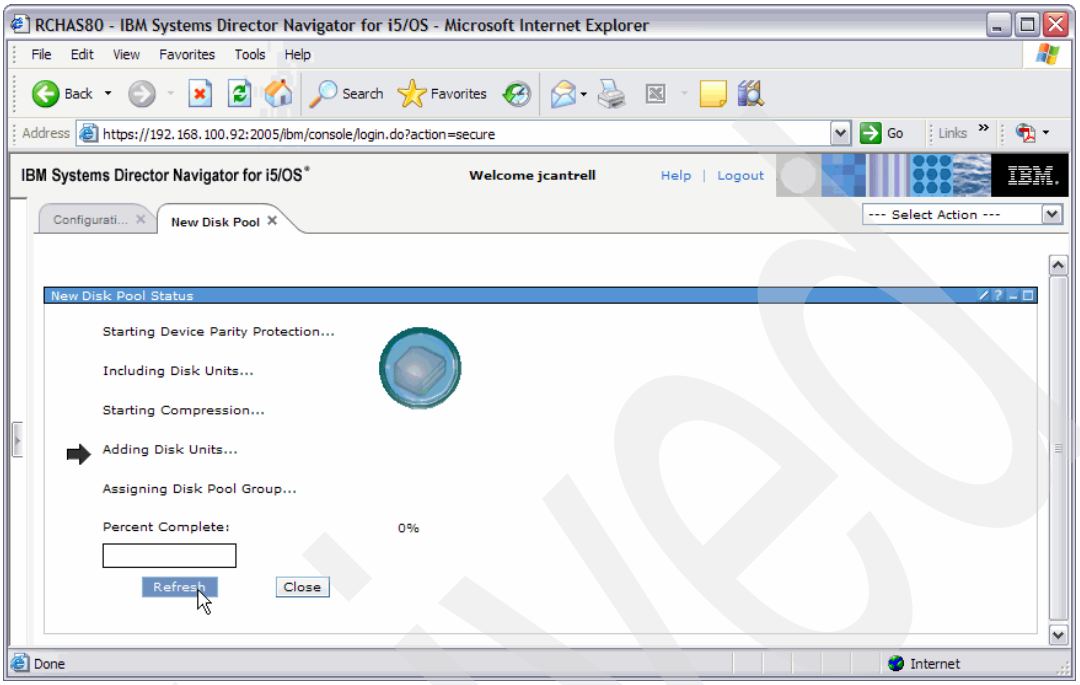


*Figure 2-10 New Disk Pool Summary → Refresh*

11. The warning message shown in Figure 2-11 may appear. Select **Details** for an explanation or select **Continue**.

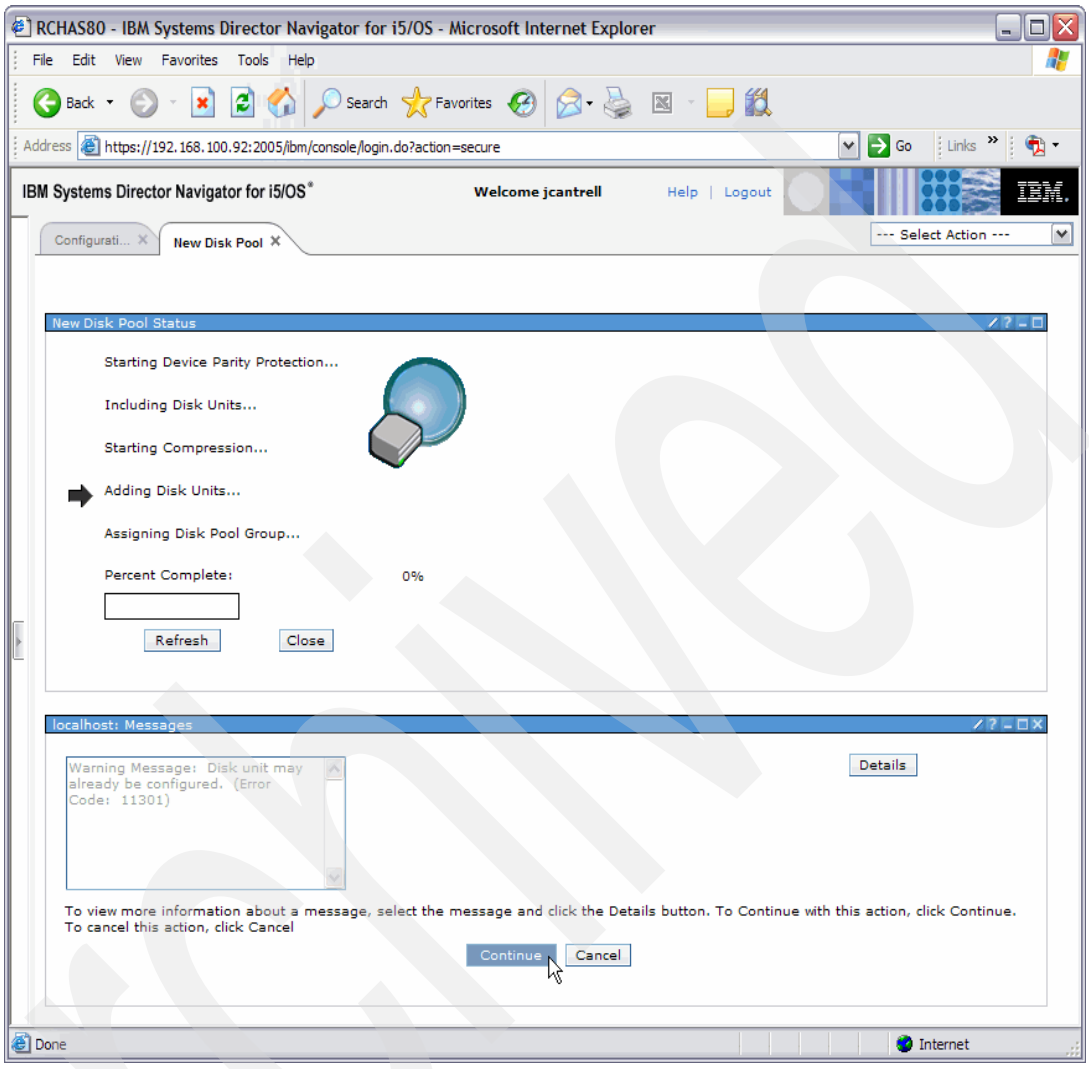


*Figure 2-11 New Disk Pool Status → localhost: messages → Continue*

12. Periodically click **Refresh** to update the status (Figure 2-12).

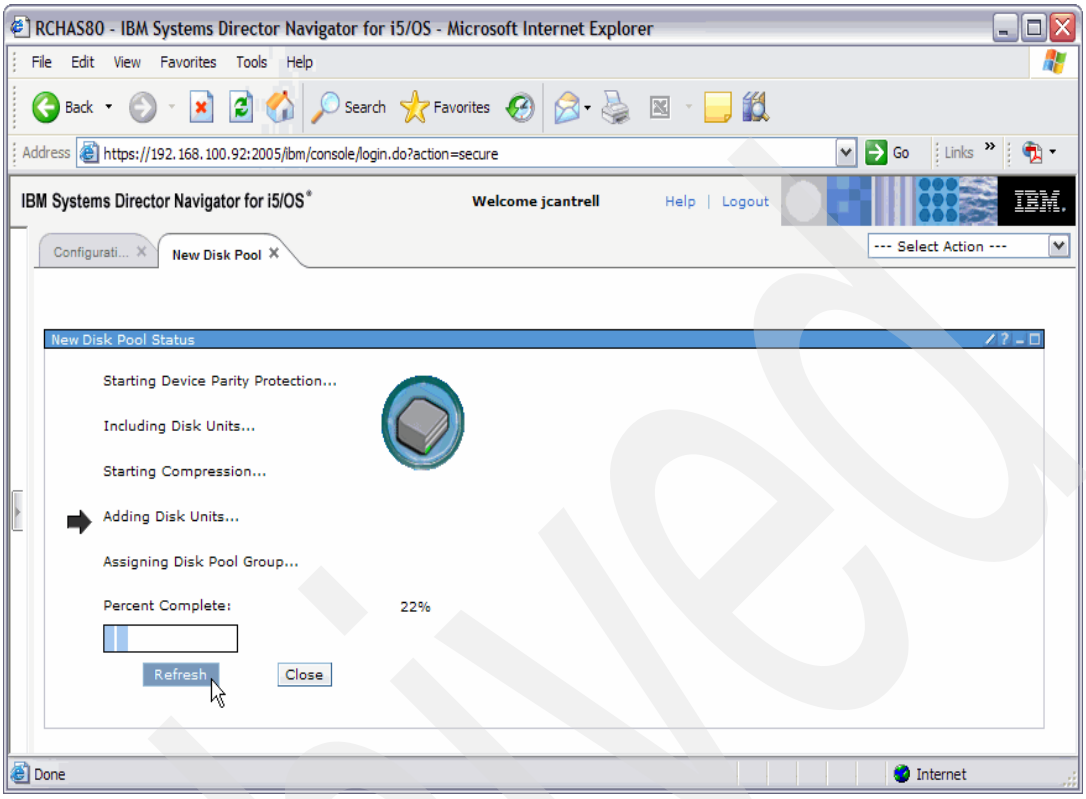


*Figure 2-12 New Disk Pool Status → Refresh*

13. Upon completion of adding the disk units, you will receive a window stating that the action that you requested completed successfully (Figure 2-13). Select **OK**.

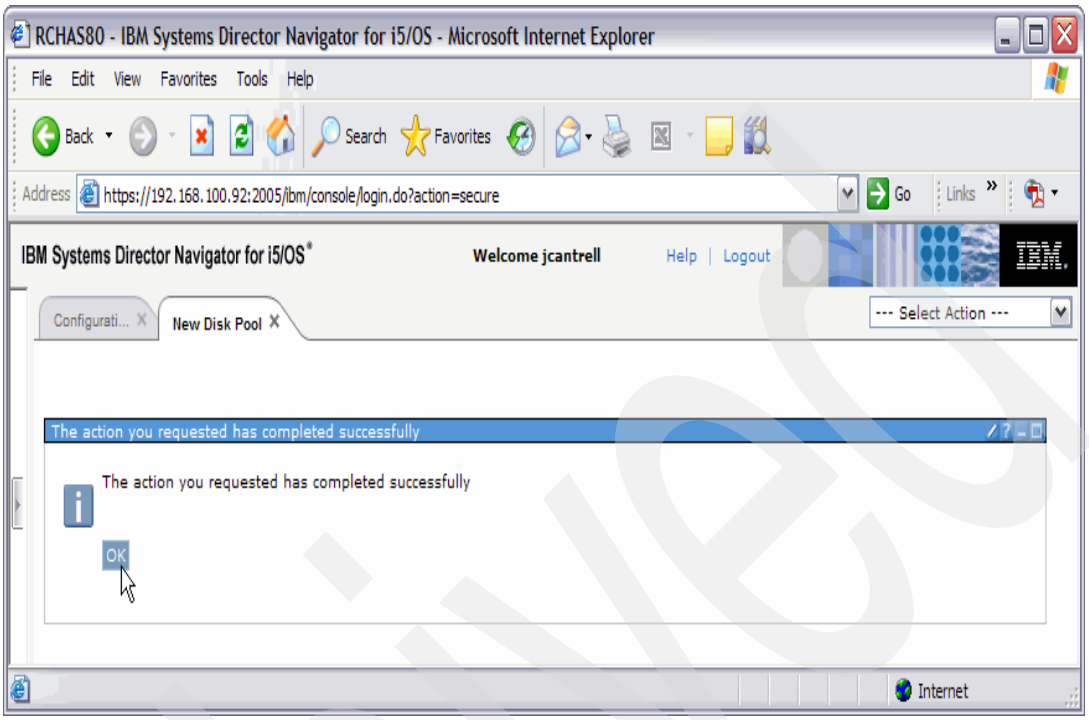


*Figure 2-13 New Disk Pool → The action you requested has completed successfully → OK*

14. If you have selected the disk units to be mirrored the window shown in Figure 2-14 appears. If you have not selected mirroring, go to 2.4, “Creating a secondary disk pool” on page 29.

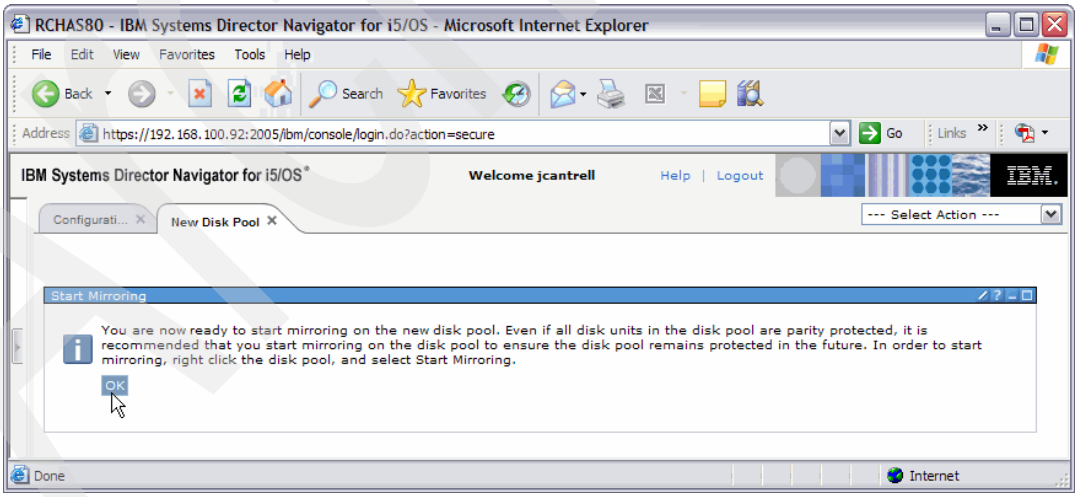


*Figure 2-14 New Disk Pool → Start Mirroring → OK*

## 2.3 Start mirroring

Starting mirroring can take some time. To start mirroring of the disk units selected for the pool:

1. On the Configuration and Service Task window (Figure 2-15) select **Disk Pools**.

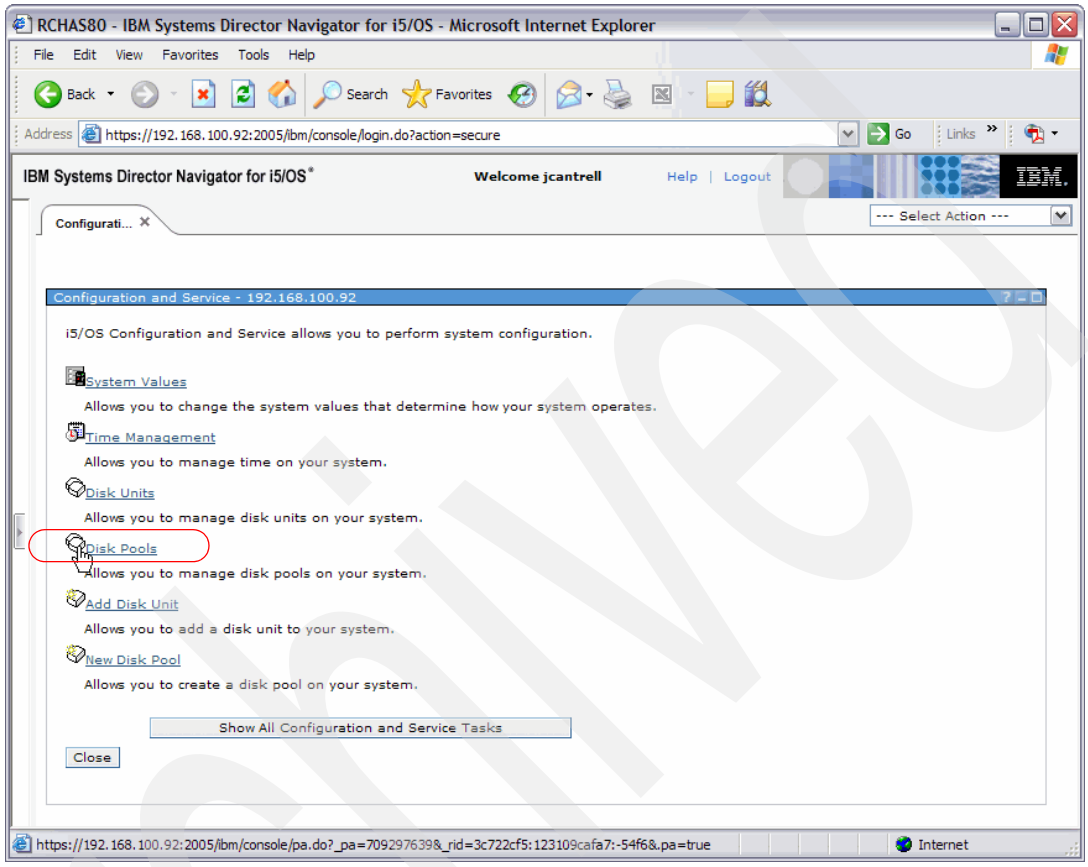


*Figure 2-15 Director Navigator → i5/OS Management → Configuration and Service → Disk Pools*

2. View the pop-up menu for the disk pool that you want to mirror by clicking the arrow next to the disk pool name (Figure 2-16).

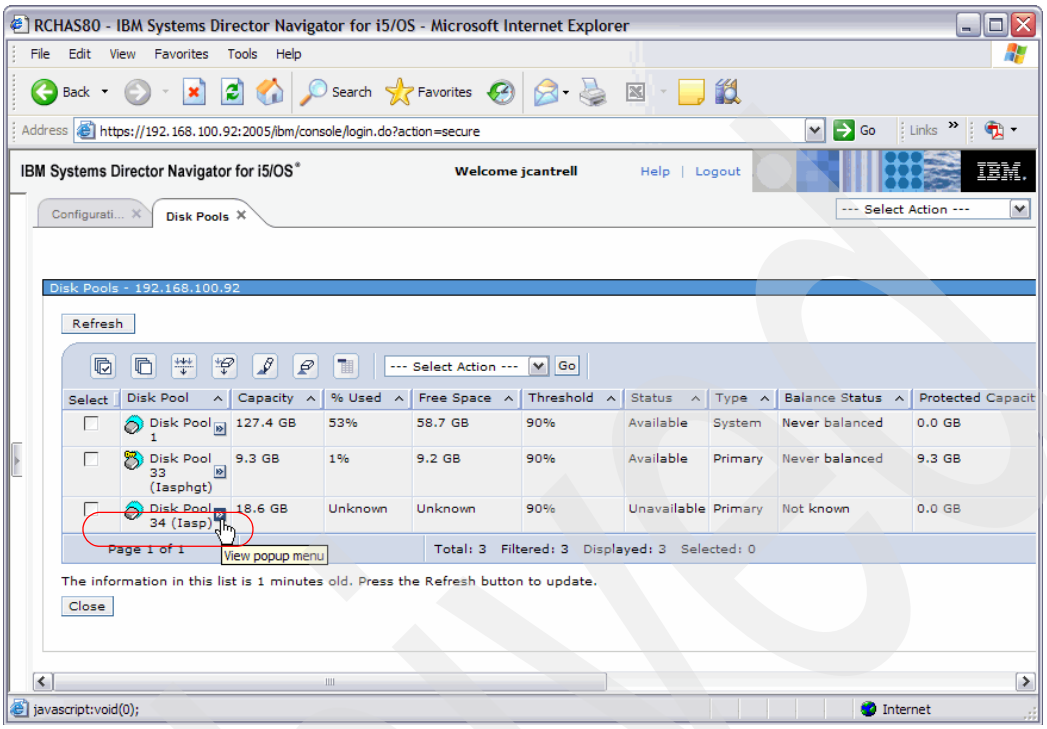


*Figure 2-16 Disk Pools → view pull-down menu*

3. Select **Start Mirroring** from the pull-down menu (Figure 2-17).

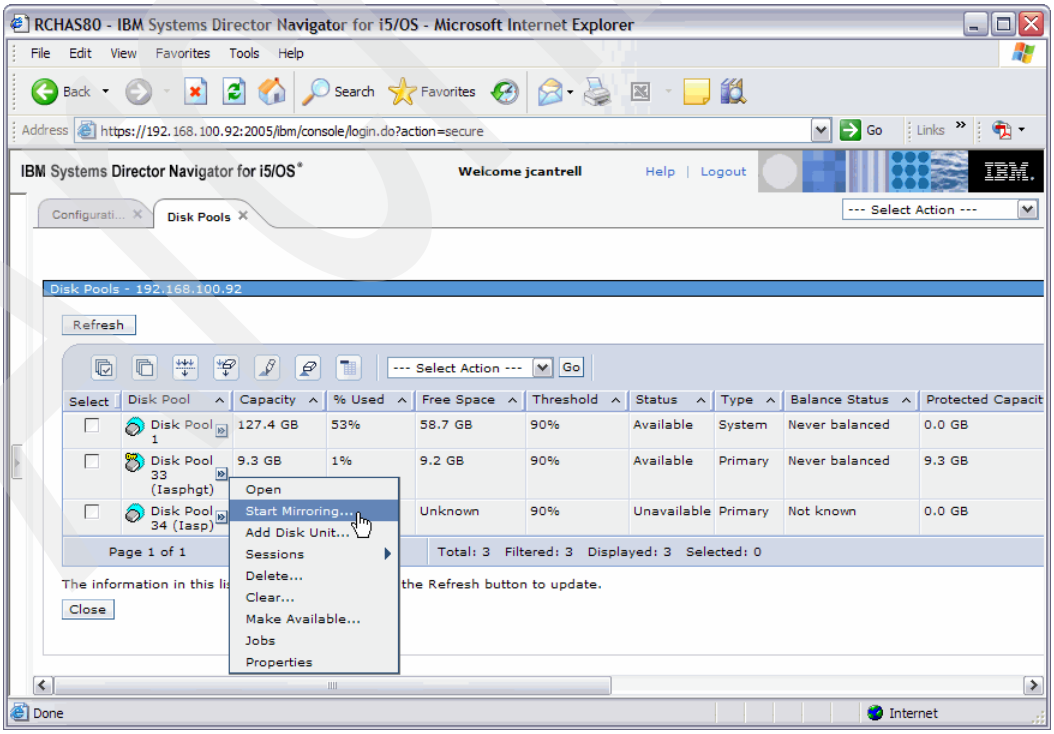


*Figure 2-17 Disk Pools → pull-down menu → Start Mirroring*

4. Confirm your start mirroring request by clicking **Start Mirroring** (Figure 2-18).

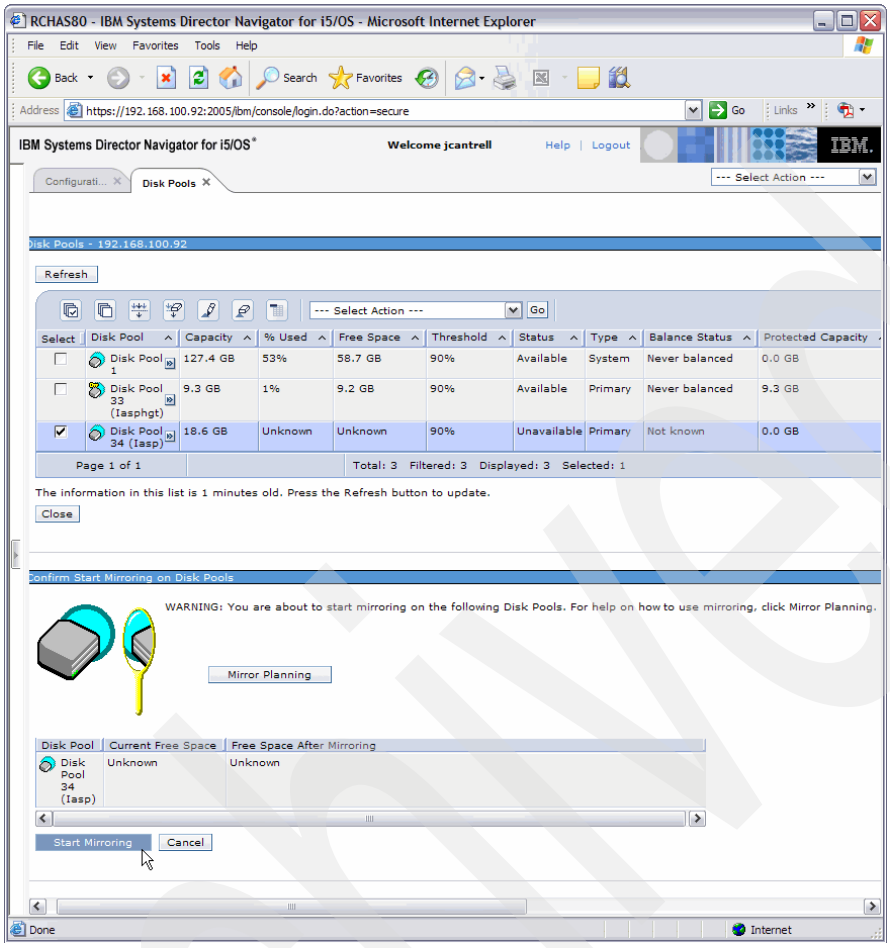


*Figure 2-18 Disk Pools → Confirm Start Mirroring on Disk Pools → Start Mirroring*

5. Periodically refresh the window to show any warning messages that may exist (Figure 2-19).

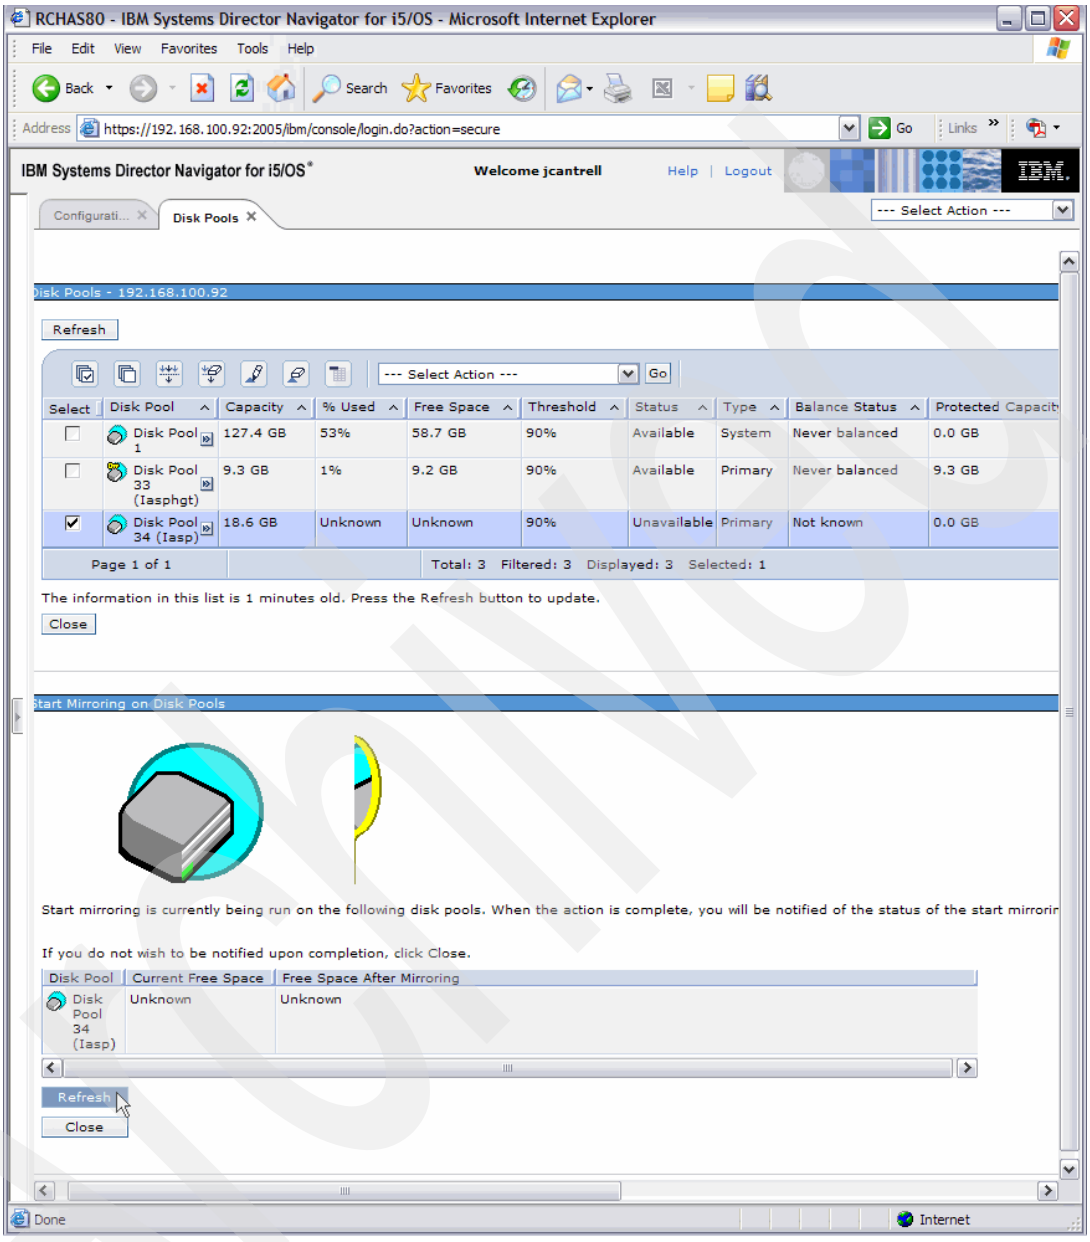


*Figure 2-19 Disk Pools → Start mirroring on Disk Pools → Refresh*

6. The warning message shown in Figure 2-20 may appear. Click **Details** or **Continue**.

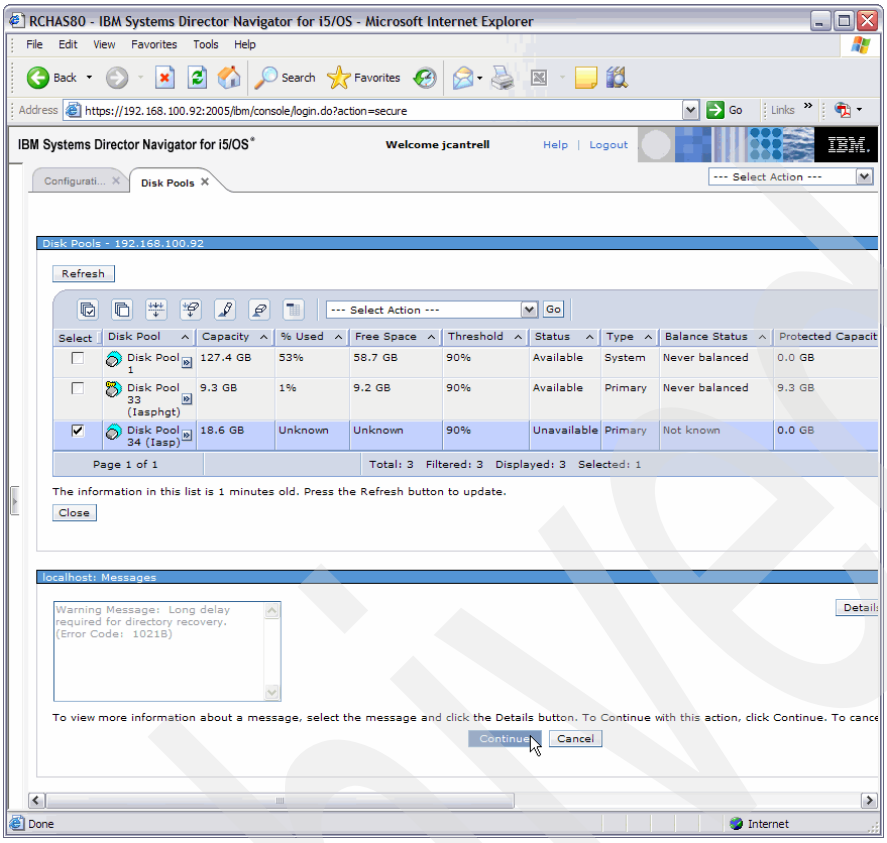


*Figure 2-20 Disk Pools → localhost: messages → Continue*

7. Virtual disk users may see the message shown in Figure 2-21. Select **Details** or **Continue**.

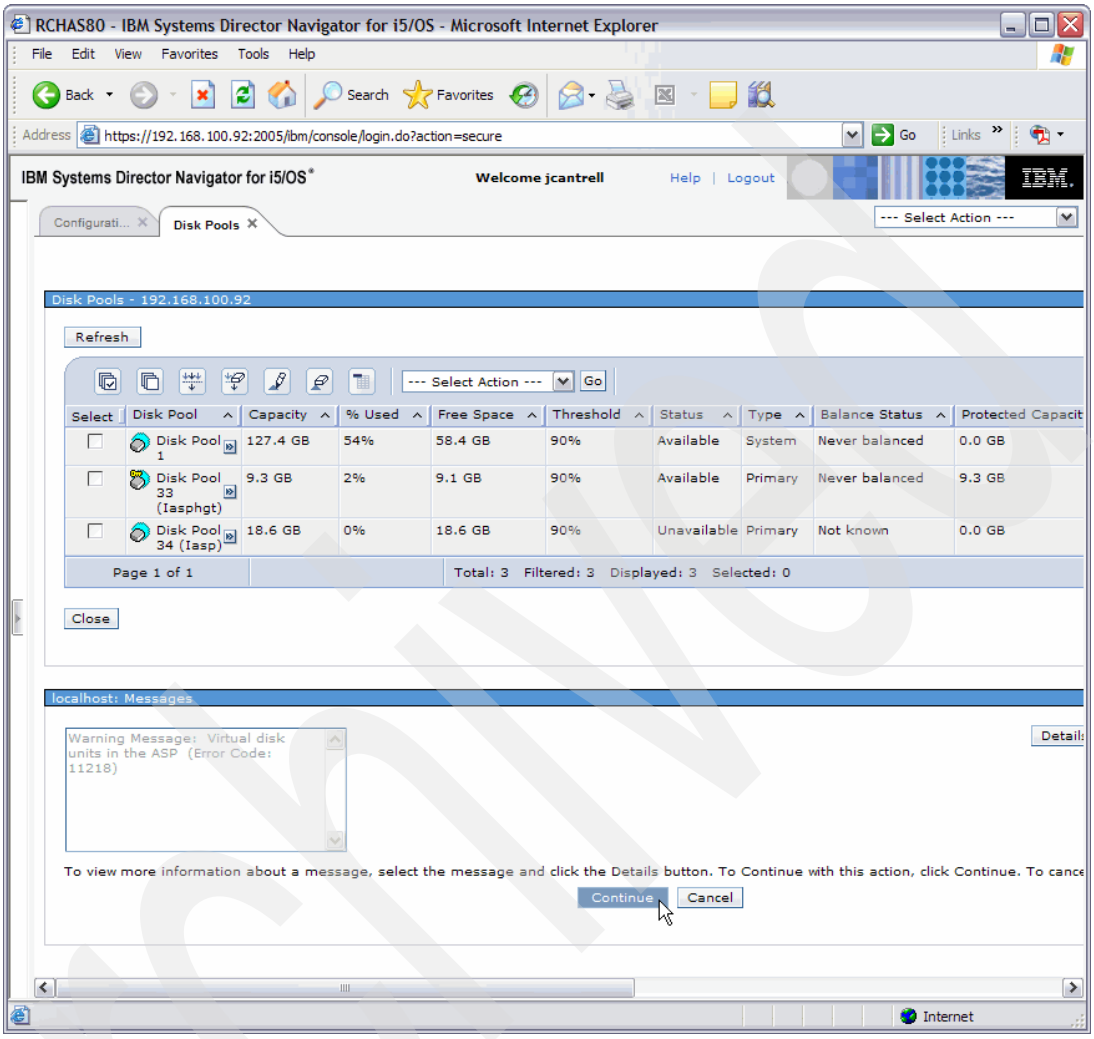


*Figure 2-21 Disk Pools → localhost:messages → Continue*

8. Continue clicking **Refresh** (Figure 2-22) periodically to view the updated status.

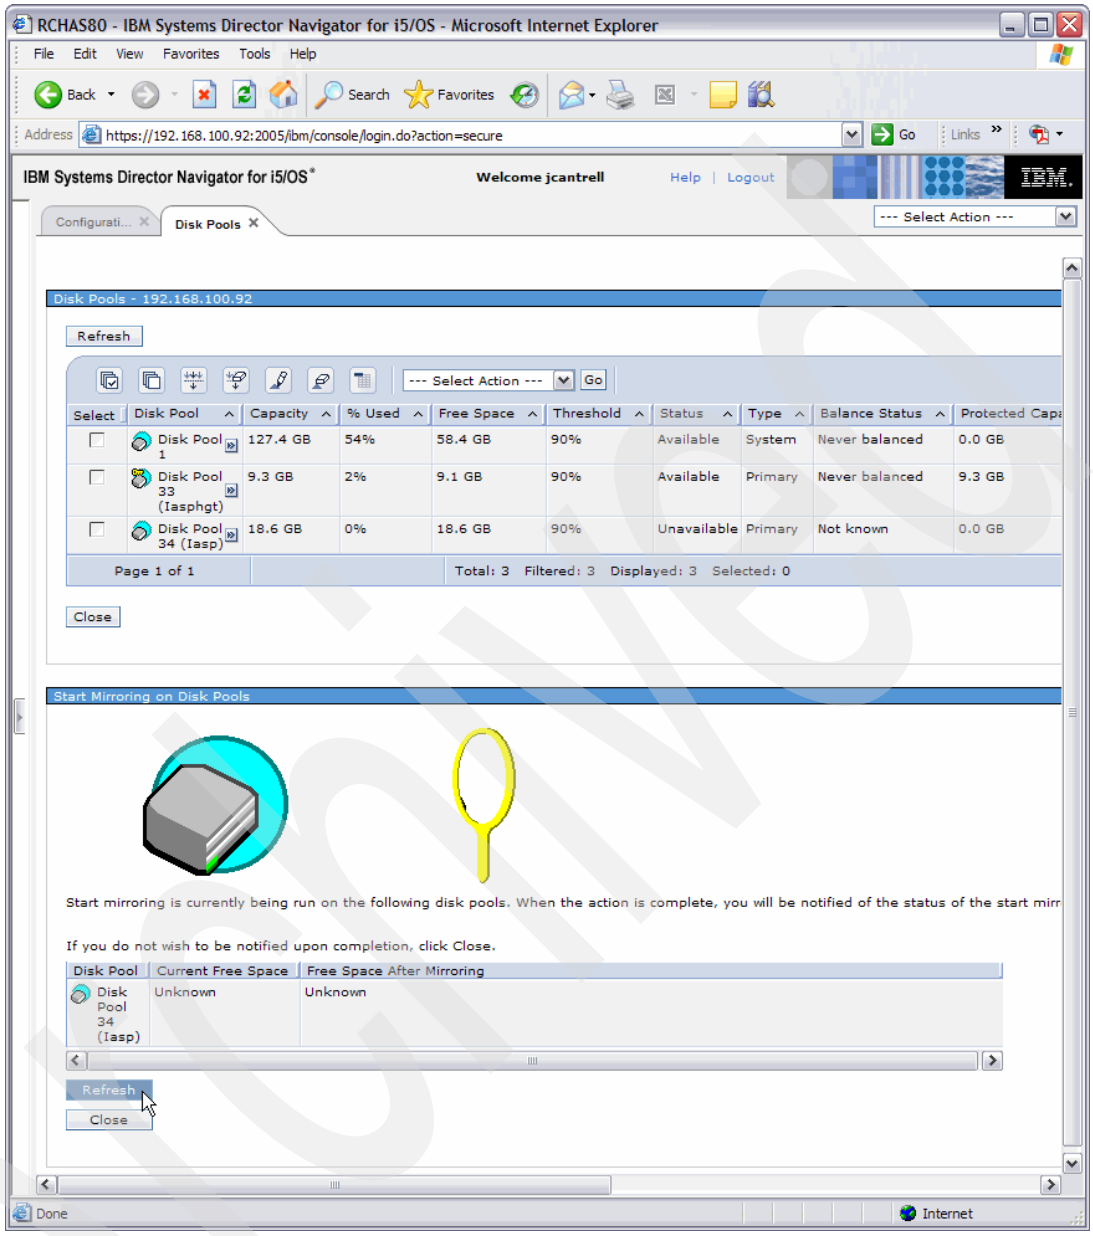


*Figure 2-22 Disk Pools → Start Mirroring on Disk Pools → Refresh*

9. Upon successful completion the window shown in Figure 2-23 is displayed. Select **OK**.

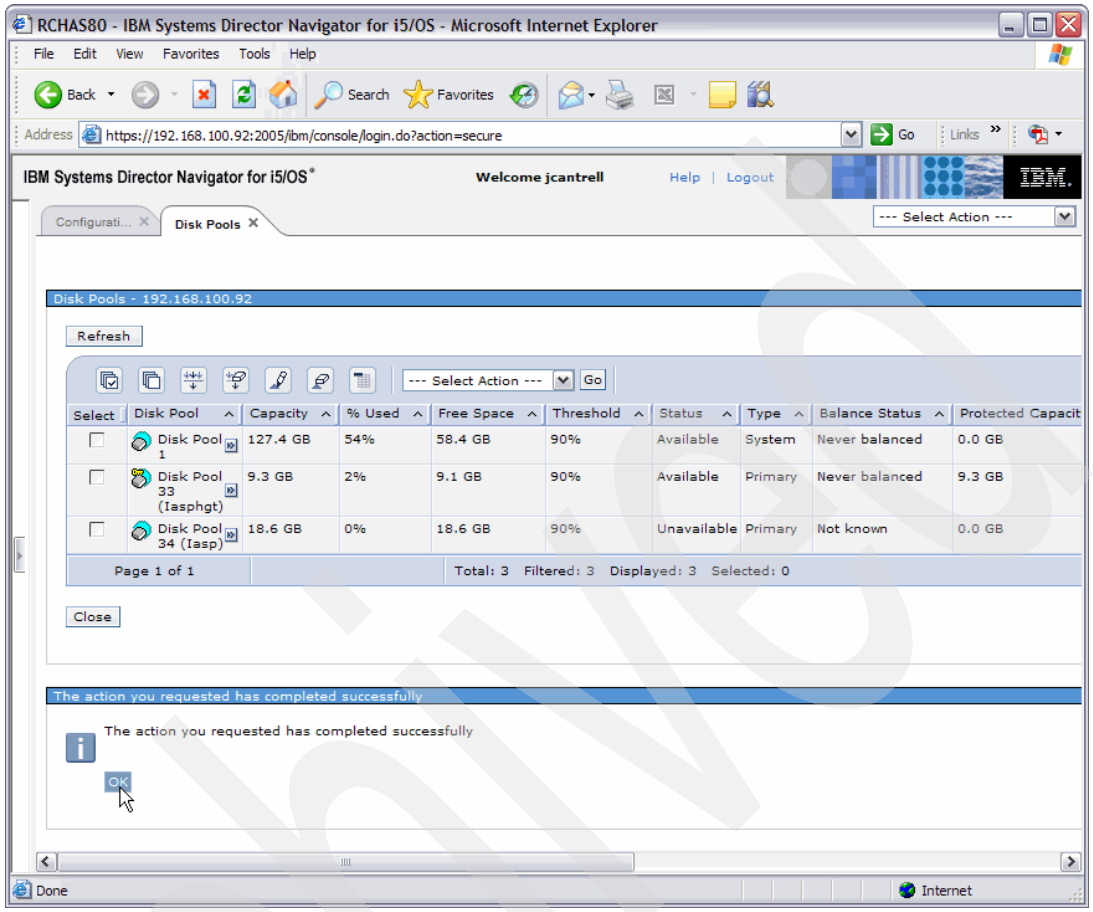


*Figure 2-23 Disk Pools → The job you requested has completed successfully*

10. You must make your disk pool available before you can use it. Refer to 2.5, “Making a disk pool available” on page 33.

## 2.4 Creating a secondary disk pool

A secondary disk pool is simply another IASP. It is attached to a primary IASP and then made unavailable or available with the primary. To create the secondary disk pool, follow the same steps for creating a primary disk pool as in 2.2, “Creating a primary disk pool (IASP)” on page 10. However, on the New Disk Pool window, for Type of disk pool, select **Secondary**, Then assign it to an existing primary IASP.

1. From Director Navigator → Configuration and Service → Disk Pools use the pull-down to select **New Disk Pool** (Figure 2-24) or use Director Navigator → Configuration and Service → New Disk Pool.

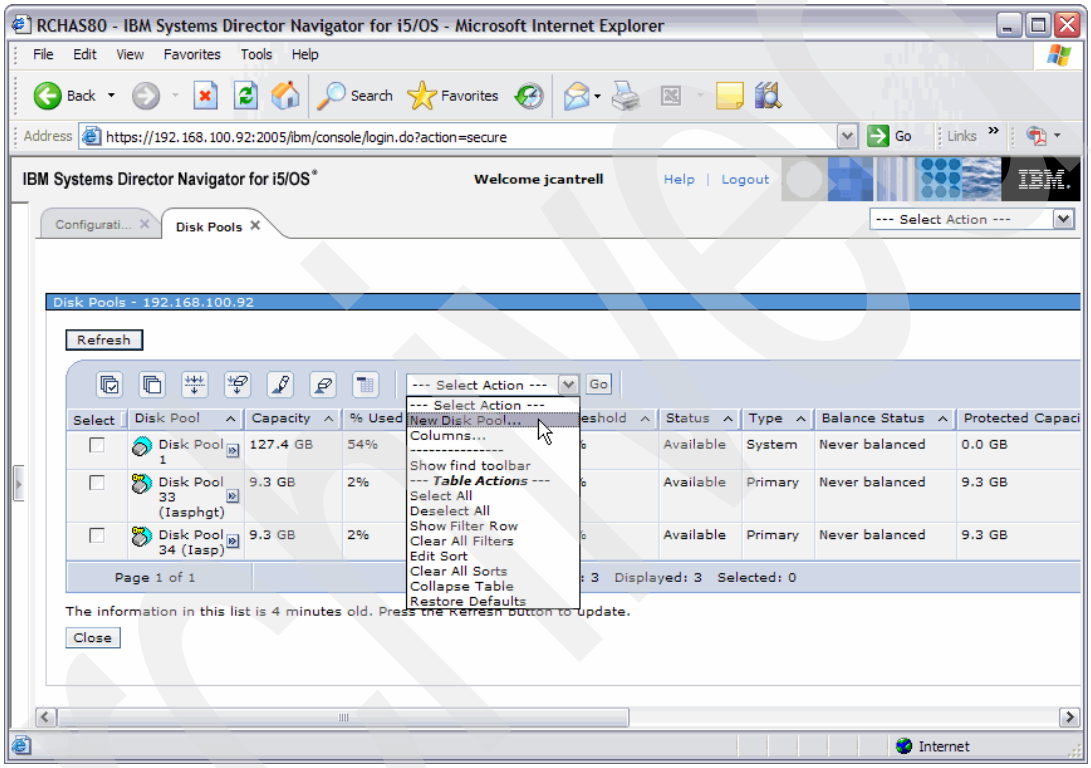


*Figure 2-24 Configuration and Service → Disk Pools → New Disk Pool*

2. Once you have selected New Disk Pool, select **Go** (Figure 2-25).

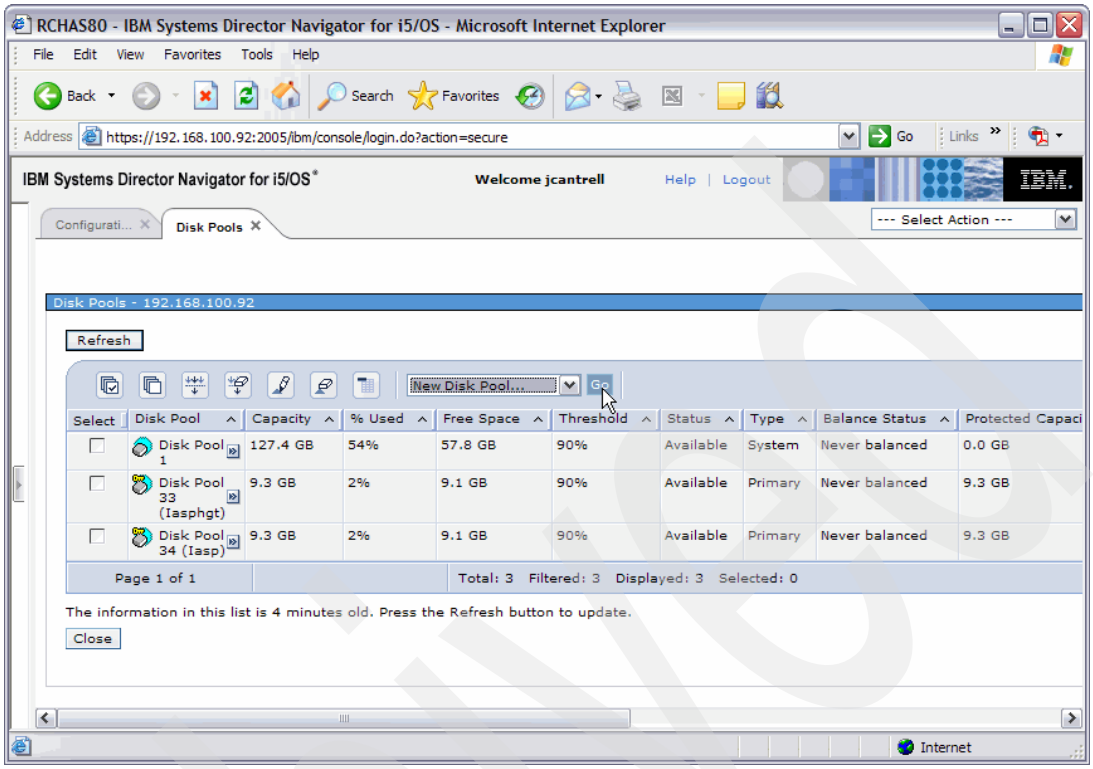


*Figure 2-25 Configuration and Service → Disk Pools → New Disk Pool → Go*

3. The New Disk Pool - Welcome Display is shown. Select **Next** (Figure 2-26).

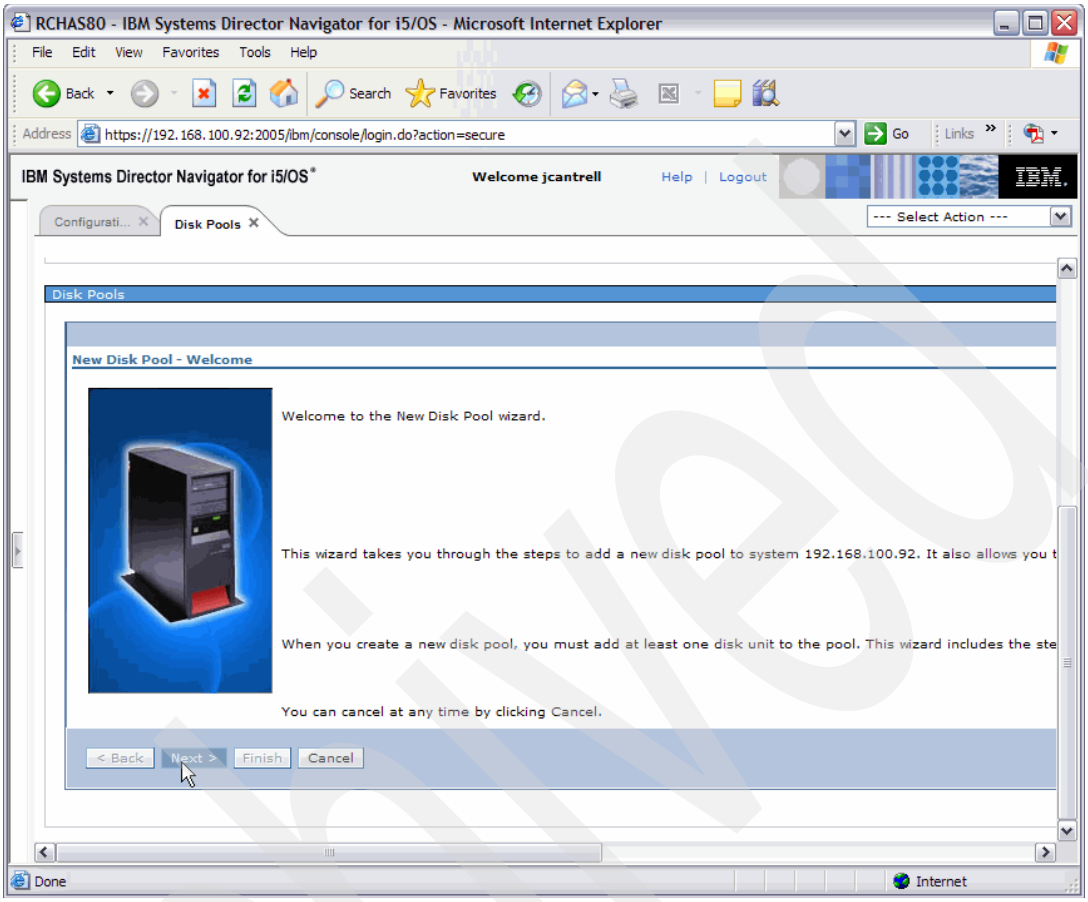


*Figure 2-26 Disk Pools → New DIsk Pool - Welcome → Next*

4. Select **Secondary** and then **Go** (Figure 2-27).

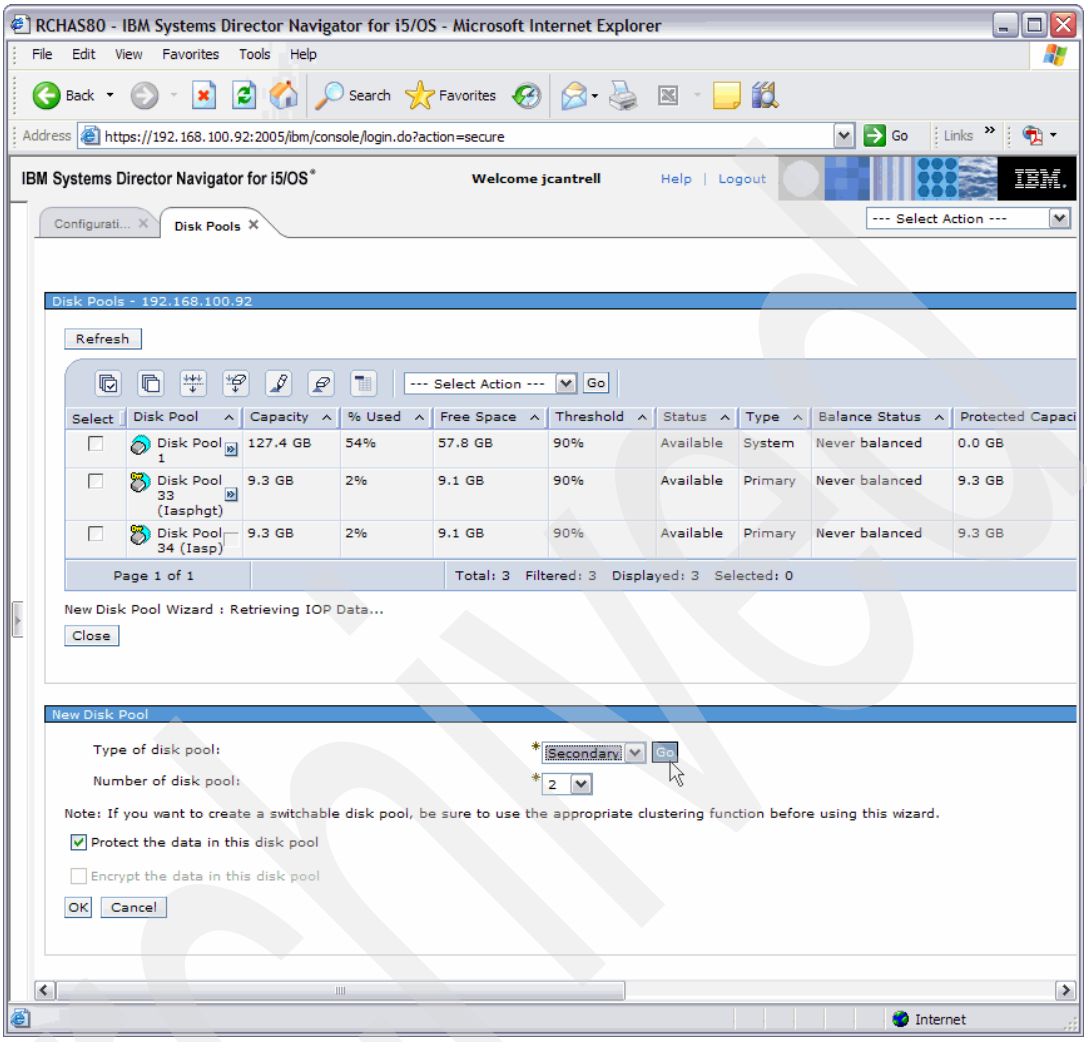


*Figure 2-27 Disk Pools → New Disk Pools → Type of Disk Pool → Go*

5. Enter the name of your secondary disk pool, then select the corresponding primary pool (Figure 2-28).

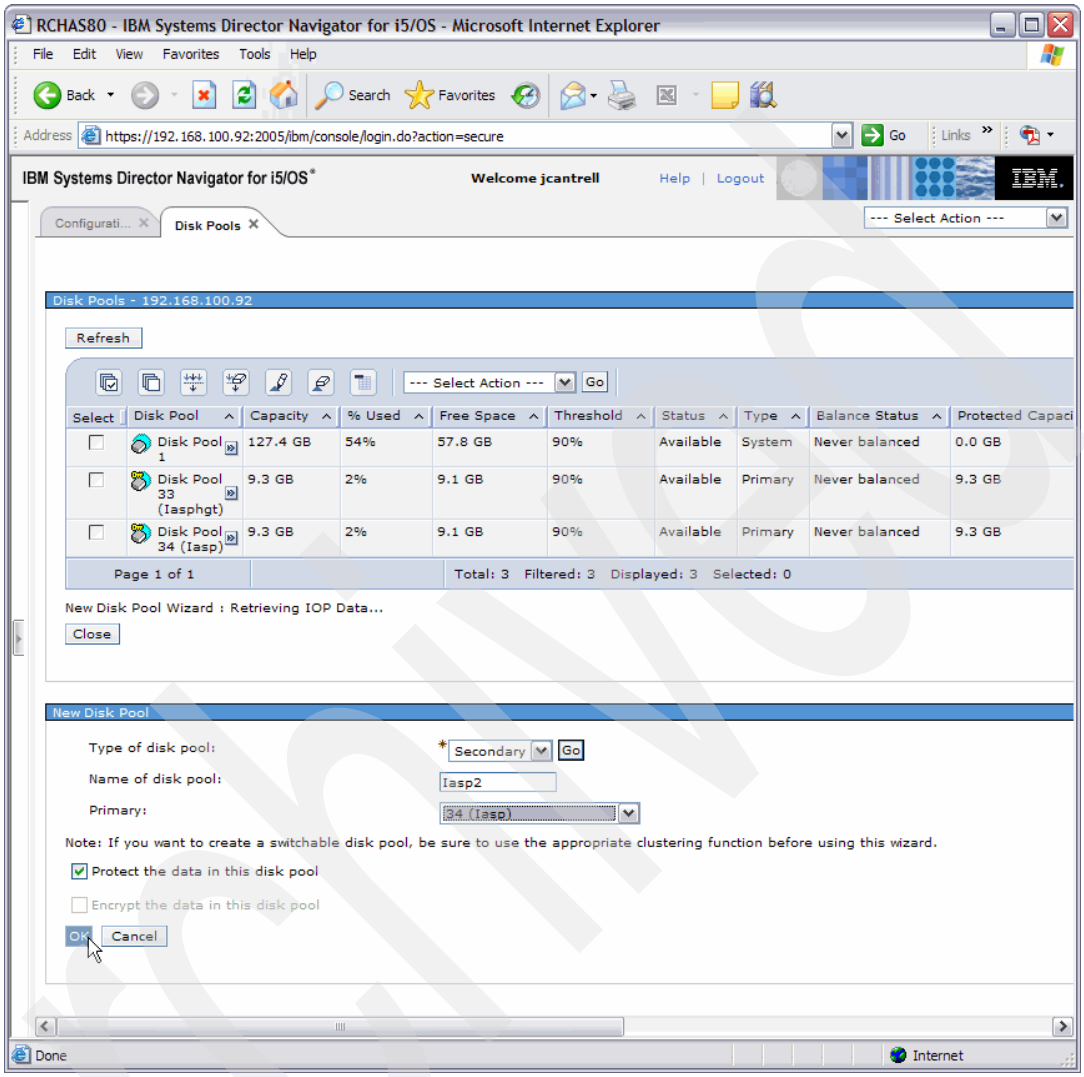


*Figure 2-28 Disk Pool → New Disk Pool → Ok*

## 2.5 Making a disk pool available

A disk pool must be made available to access and work with its data.

**Note:** If you make a primary or secondary disk pool available, all of the disk pools in the disk pool group are also made available at the same time.

The associated disk pool group name (primary disk pool name) for a secondary disk pool is available by viewing the Properties option of the pull-down menu for the selected disk pool.

## 2.5.1 Activating a disk pool using the System Director Navigator

To make a disk pool available:

1. From the System Director Task List, expand **i5/OS Management**.
2. Select **Configuration and Service**.
3. Select the **Disk Pools** link.
4. Left-click the double arrow icon  next to the disk pool name and number on the entry in the disk pools list. This opens a pull-down menu.
5. Select the **Make Available** option from the pull-down menu, as shown in Figure 2-29.

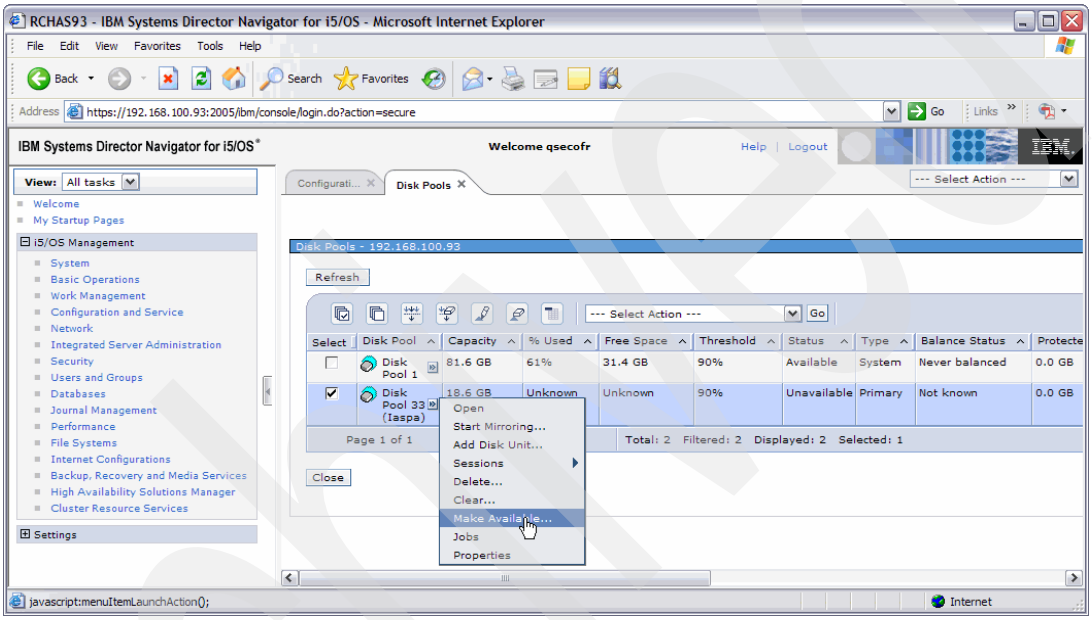


*Figure 2-29 Choosing the Make Available option*

6. Confirm the option by clicking the **Make Available** button on the confirmation panel, as shown in Figure 2-30.

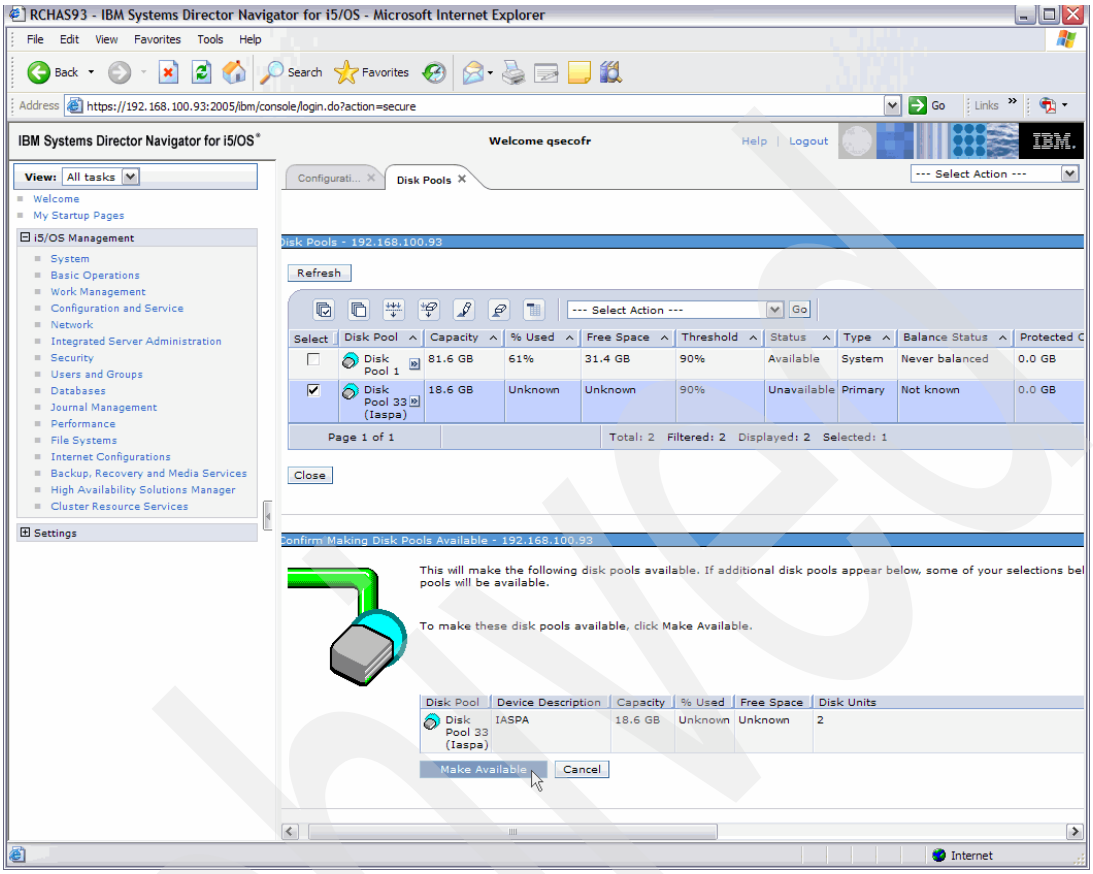


*Figure 2-30 Confirming the Make Available option*

7. Select the **Refresh** button on the Making Disk Pools Available panel to check the progress of the Make Available function, as indicated in Figure 2-31. This panel does not refresh automatically, so refresh it until the Make Available is complete. Each time that you refresh, the status and elapsed time will be updated. If an error occurs, you will not see the error condition until you refresh. If you receive an error message that is in its own panel that has a Continue button, you must click **Continue** for the process to continue.

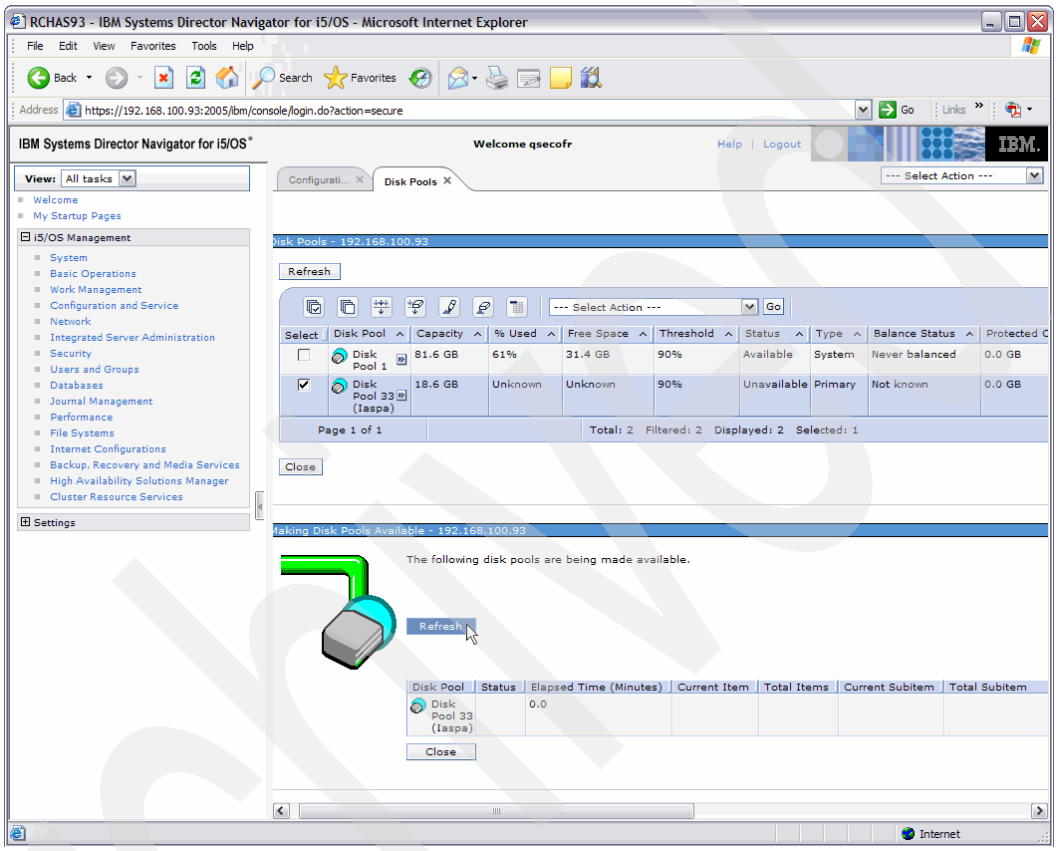


*Figure 2-31 Refreshing the Make Available panel*

8. When the Make Available function has successfully completed, a refresh will show a status of Complete, as indicated in Figure 2-32.
9. Select **Close** on the Making Disk Pools Available panel.

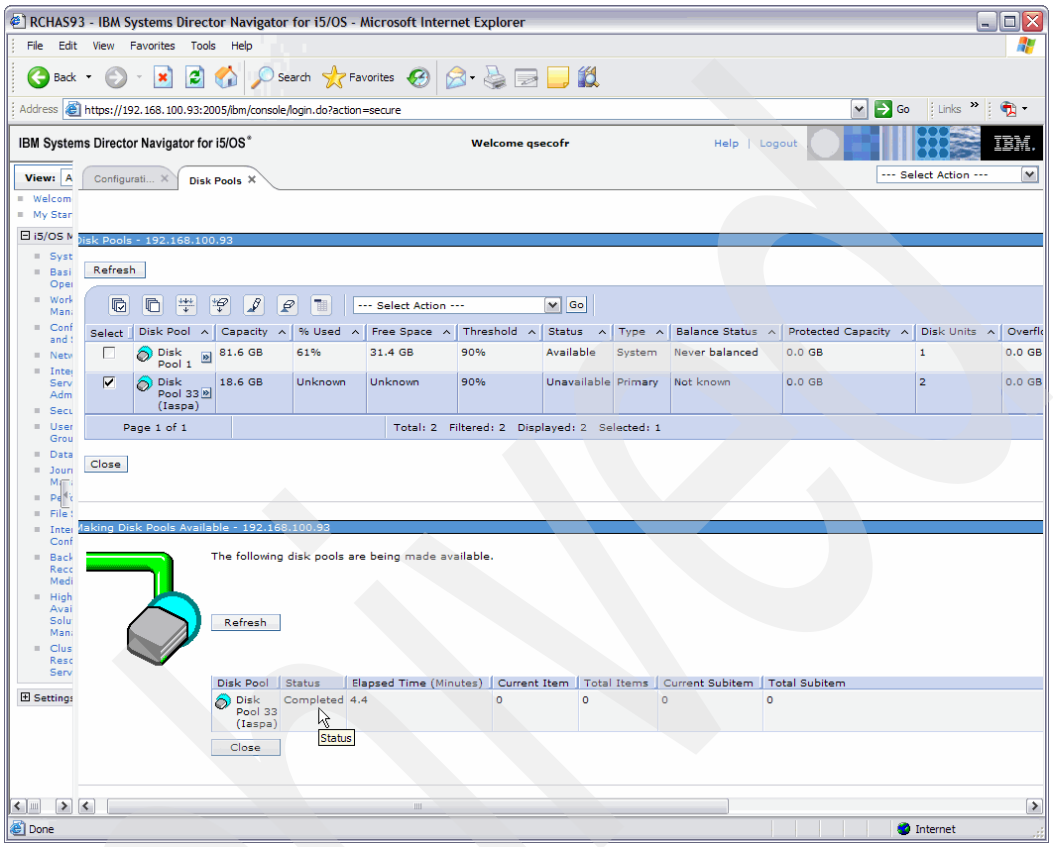


*Figure 2-32 Making Disk Pools Available complete status*

10. Click **Refresh** on the Disk Pools panel to show the current status of the disk pools. The disk pool specified should now show a status of Available, as shown in Figure 2-33.

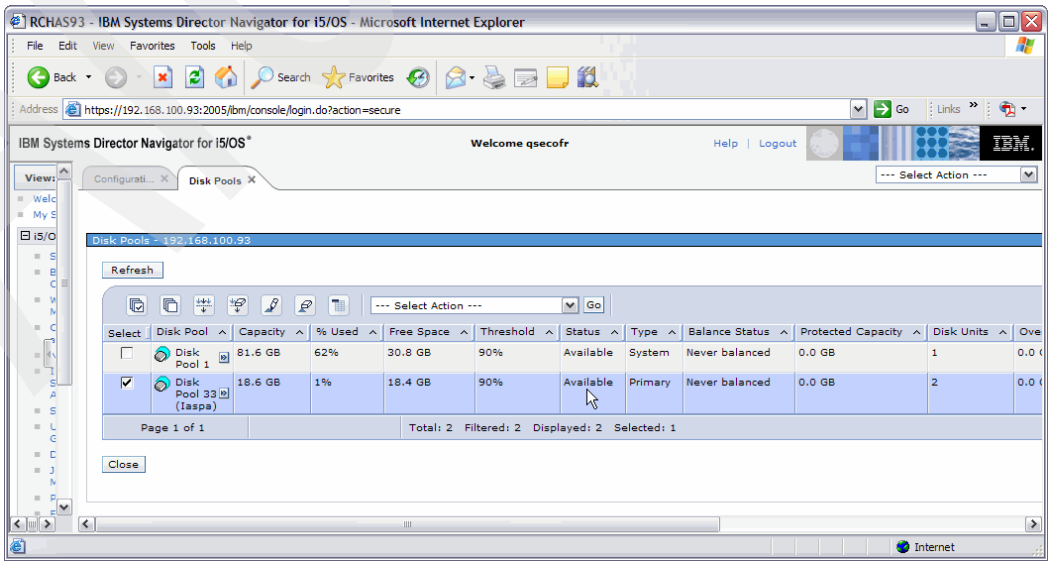


*Figure 2-33 Disk Pool Available*

11. Double-click the disk pool name and the number entry will open a list of the disks allocated to this disk pool, as shown in Figure 2-34. Observe the disk units beginning with the numeral 4 in the thousands position. These are the newly available IASP disk units.

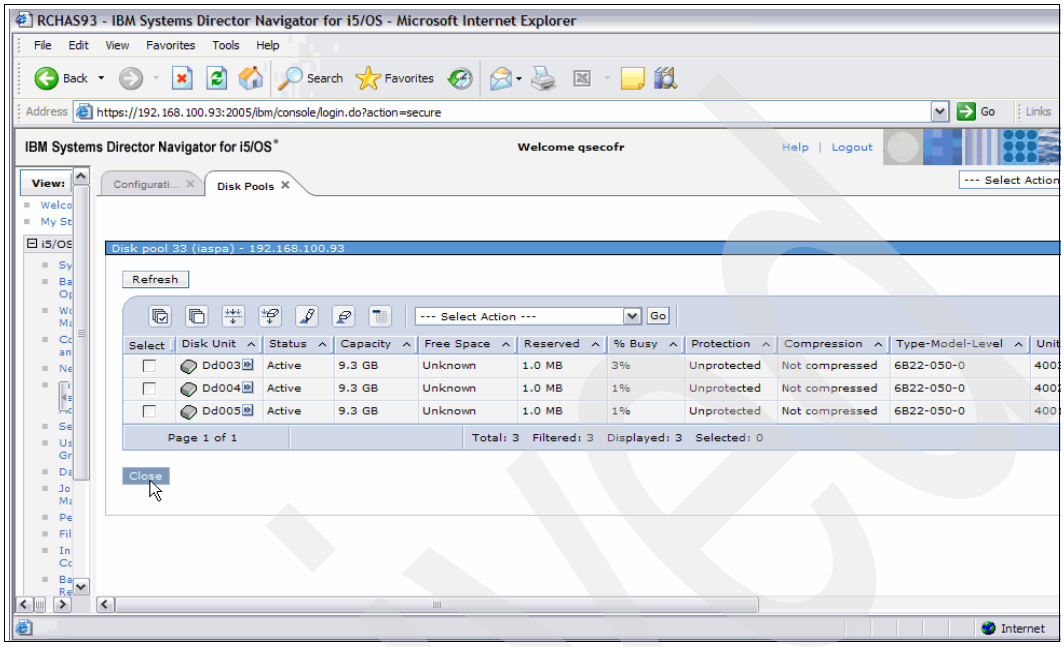


*Figure 2-34 Disk pool disk listing*

12. The IASP disk units are now available for use. You have successfully created them.

    After the IASP is available you can visit other sections in this book for additional topics.

## 2.5.2 Activating a disk pool using the VRYCFG command.

The command VRYCFG can also be used to enable an IASP or disk pool. The VRYCFG command can be included in the system startup program to enable an IASP on system IPL or in an application process. The VRYCFG command appears as in Figure 2-35. Note that CFGOBJ is the name of the disk pool as specified when it was created and CFGTYPE is *DEV.

```
Type command, press Enter.
===> VRYCFG CFGOBJ(IASPA) CFGTYPE(*DEV) STATUS(*ON)

     Vary on completed for device IASPA.
```

*Figure 2-35 Disk pool VRYCFG *ON command*

## 2.5.3 Duration of the Make Available option

The time that it takes for the Make Available (VRYCFG *ON) option to execute is sensitive to the number of objects in *SYSBAS. If the Make Available time is more that two minutes, many application functions may start to time out and fail.

**Note:** The first time that you make a new disk pool available can take considerably longer due to disk formatting and file creation activities that are only executed on the first Make Available.

On systems where the number of objects in *SYSBAS is relatively small (IBM i and licensed programs only installed), the Make Available time is relatively quick—under a minute in test environments. On systems where the number of objects in *SYSBAS is extremely large, the Make Available time has taken quite a bit longer.

3

# Planning for independent disk pools

This chapter discusses the various elements of planning for independent auxiliary storage pools (independent disk pools). The planning steps illustrated here represent the process involved. The full extent to which planning for independent disk pools can be taken is a subjective process.

## 3.1 Independent disk pool characteristics

Each independent disk poll group is represented as a separate DB2® database instance. Duplicate library names are allowed within different independent disk pool group on the same system.

Each job or thread always has visibility to objects in disk pools 1–32 (system disk pool and all basic user disk pools), but can only be attached to one disk pool group at a time.

Independent disk pools can be replicated or switched between multiple systems, however, they work with only one system at any one time. Since independent disk pools are self-contained, they have unique characteristics. All of the necessary system information associated with the objects contained in an independent disk pool is contained in that pool. Because of this characteristic, there are certain restrictions when using independent disk pools.

### 3.1.1 Migrating independent disk pools between release levels

After an independent disk pool is made available on the IBM i 6.1 system, it cannot be made available on earlier release levels. It is possible to move an earlier version independent disk pool to an IBM i6.1 system and make it available. However, after it is made available on the IBM i 6.1 system, its internal contents are changed, and it cannot be made available to the earlier version system.

### 3.1.2 Spool file considerations

Spool files can exist in an independent disk pool. If spooled output is in an independent disk pool, when the pool is switched to another system, the spool files become available to that system. The ability to move SPOOL files requires that the spool files be detached from the job/thread that produces the output.The ability to detach SPOOL files is controlled by setting the SPLFACN parameter on the job description to *DETACH or setting the system value QSPLFACN to *DETACH. The default for the system value is *KEEP, and the default for the job description is *SYSVAL (refer to the system value). If the default values are used, SPOOL files will remain attached to the job/thread, and the job/thread will remain in the system until the SPOOL files are deleted. Often times this will clutter the system with thousands of jobs that are no longer active. Performance of several systems functions and IPL time are adversely affected by too many jobs in the system.

### 3.1.3 Job queue considerations

Job queue objects can be created in independent disk pools, allowing applications to run with fewer changes when using independent disk pools. Operationally, they are identical to job queues in SYSBAS. Users can manipulate the jobs (submit, hold, release, and so on) or the job queues themselves (clear, hold, release, delete, and so on). It is important to note that jobs on JOBQs in independent disk pools will not persist across vary on/off.

## 3.2 Independent disk pool performance considerations

Because of the structure of independent disk pools and their relationship with the system ASP, there is a performance impact to the system for each independent disk pool in use. The more independent disk pools are involved, the bigger the impact. During normal operations, this overhead is in communications between the independent disk pool tables and system tables, keeping them synchronized. There is a much greater impact when these independent disk pools first come online to the system. A few reasons for this are:

- If this is the first vary on since an IPL, the objects are verified for consistency.
- After each vary on of the independent disk pool, the user IDs and AUTLs are created as required, if they do not exist.
- Making more than one independent disk pool available at a time to the system means that this synchronization between the system and the ASP must take place concurrently (serially can yield better system performance).

### 3.2.1 Performance overview

Performance is usually one metric of business needs. Even if it is not a primary consideration, maximizing throughput of the configuration is a small effort with a potentially high return. Consider the following key performance areas:

- Processor capability

  In an independent disk pool environment, the processing resources are not separated or divided among the various applications. They are available to everyone. Control is managed using IBM i 6.1 work management techniques.

- Number of disk arms

  This is important when moving from a single auxiliary storage pool (ASP) environment to an independent disk pool environment. If an application requires 20 arms to achieve good performance in the system ASP, it may take 20 arms in an independent disk pool. Arms are also required in *SYSBAS, so do not just remove them all and leave the load source unit.

- Size of system ASP

  The size of the system ASP is important for two reasons:

  – Performance of the operating system and any applications that remain in *SYSBAS.

  – Performance of SQL functions related to the application in the independent disk pool

    The temporary space used for running queries is still found from the system ASP. There must be enough space to run the queries and sufficient arms to allow the application to perform to requirements.

- Size of the system ASP related to independent disk pool IPL time

  The relationship between the size and number of objects in the system ASP and the independent disk pool affects the vary-on time of the independent disk pool. Keep the system ASP size and number of objects at a minimum, or the independent disk pool IPL may be protracted.

- On demand processors

  With the new IBM @server On/Off Capacity on Demand for Power Systems servers, you have the potential for considerable changes in processing environments. These changes can severely impact performance. Changes to the number of processors cause the access plan to be re-validated and may require a different access plan. For a few not-too-complex

queries, this should go unnoticed. But for thousands of small queries, the addition of a second processor to re-validate the plan can cause a significant performance degradation. Similarly, changes in the number of processors increases or decreases the amount of parallelism. Very large queries may change their access plan, resulting in variable performance.

- SQL packages

  These packages may also be affected by moving their data to an independent disk pool that is switchable. If the independent disk pools are different sizes and use the same SQL package, the access plan changes each time that a different independent disk pool is selected by a user.

## 3.2.2 Disk drives: arms versus capacity

The number of disk arms on any system is vital for good application performance. This section discusses the impact of disk arms and why you need to achieve some balance between the number of arms and disk capacity.

Mixing drive capacities within a pool is the most difficult for tracking performance characteristics, at the pool level. If application performance is tracked, the job is easier if all disk drives within a disk pool are the same speed and size.

Mixing drives distorts the performance characteristics available though performance tools. Mixing drive sizes may significantly reduce performance. The amount and type of data to be placed within the independent disk pool determines the type and placement of the disk.

The type of data to be placed in an independent disk pool may also determine the performance requirements. If the data is historical and infrequently accessed, slow access or reduced performance may be perfectly acceptable.

### Arms

Systems can run (slowly) with as few as one disk arm in the system ASP (you would *not* actually do this). In testing, as few as five disk arms in the system ASP showed degradation when supporting an independent disk pool of 30 arms. When 10 arms supported the same independent disk pool of 30 arms, performance was acceptable.

Moving applications to independent disk pools and sharing *SYSBAS is not the same as moving applications from a single partition to a multi-partition server with a managing partition (also known as a *thin primary*). You should not remove all resources from *SYSBAS.

There is not a rule of thumb for a ratio of arms between the system ASP and an independent disk pool. However, arms in *SYSBAS are very important for performance.

### Capacity

The size relationship between *SYSBAS and an independent disk pool is important. The ratio of one to two (1:2) or one to three (1:3) for disk arms can be a good start for sizing disk capacity. If 30 GB of disk is required in the independent disk pool, then 10 GB of disk is required for the system ASP. Keep in mind that all temporary objects for applications residing on an independent disk pool are created in the system ASP.

When the application runs in the system ASP, the amount of storage used is a combination of data storage and temporary storage plus other stored objects. The real data storage requirement for an independent disk pool is somewhat less than the total original storage capacity.

As for disk arms, a balance between *SYSBAS and an independent disk pool must be gained. If not, there is the risk of serious performance problems. Disks and arms should be spread evenly along among the available input/output adapters (IOAs) in the I/O enclosure

**Important:** When you run a query over data in an independent disk pool, temporary objects are automatically created in the system ASP. If you do not allow sufficient capacity in the system ASP for the temporary objects, the system ASP will fill. If the query is allowed to run and fill the system ASP, the system will crash.

**Note:** Every situation is going to be different. The rule of thumb above is presented as a starting point. Your own requirements may vary.

### 3.2.3 Disk protection and failures

In the same fashion as disks in *SYSBAS, independent disk pools represent a single point of failure in the system. If the disks in the independent disk pool are permanently damaged and the data is unrecoverable, data must be restored from the latest backup copy.

IBM Systems Director Navigator for i5/OS provides the capability to protect the disks when creating the independent disk pool. Mirroring, Device Parity, or No Protection can be chosen. This is useful when new unprotected drives have been added to the system for use in an independent disk pool. The independent disk pool can be created and protected without impacting the entire system.

**Note:** Care should be taken in selecting protection options within IBM Systems Director Navigator for i5/OS, as it does not consider the physical drives themselves and the protection results may not be as expected.

## 3.3 Independent disk pool system settings

Before you implement independent disk pools, you should examine how the following system settings will be affected and how the settings will affect your application. As a general rule, all objects or libraries that are specified in system values that are resolved during the IPL process must be in *SYSBAS. In addition, all objects and libraries that are specified in network attributes must be in *SYSBAS. In the following section we discuss settings to consider.

### 3.3.1 System values

The following system values are *not* resolved during the IPL process. They are resolved when they are used and may reference items that are in an independent disk pool.

**QALWUSRDMN** Allow user domain objects in libraries.

This value specifies which libraries may contain user domain user (*USRxxx) objects. You can specify up to 50 individual libraries or all libraries on the system.

**QBOOKPATH** Book and bookshelf search path.

Regardless of how you set this value, it does not affect the operating system. The operating system no longer uses this system value. However, your application programs may use it.

**QLOCALE** Locale path name.

The locale (QLOCALE) system value specifies a locale object that can determine how data is processed, printed, and displayed. Locales can define the language used by the system, cultural data of that language, and the type of characters displayed or printed.

All objects and libraries that are specified in the following system values *are* resolved during IPL and must be in *SYSBAS:

**QATNPGM** Attention program

This value specifies the name and library of the attention program. This program must exist in the system ASP or in a basic user ASP.

**QCFGMSGQ** Configuration message queue

This system value allows you to specify the default message queue that the system uses when sending messages for lines, controllers, and devices. The message queue must exist in the system ASP or in a basic user ASP.

**QCTLSBSD** Controlling subsystem

The controlling subsystem is the first subsystem to start after an IPL. One subsystem must be active while the system is running. This is the controlling subsystem. Other subsystems can be started and stopped. If this subsystem description cannot be used (for example, it is damaged), the backup subsystem description QSYSSBSD in the library QSYS can be used. A subsystem description specified as the controlling subsystem cannot be deleted or renamed after the system is fully operational.

**QIGCCDEFNT** Double-byte code font

This value is used when transforming an SNA character string (SCS) into an Advanced Function Printing Data Stream (AFPDS). It is also used when creating an AFPDS spooled file with shift in/shift out (SI/SO) characters present in the data. The IGC coded font must exist in the system ASP or in a basic user ASP. The shipped value is different for different countries or regions.

**QINACTMSGQ** Inactive job message queue

This value specifies the action that the system takes when an interactive job has been inactive for an interval of time (the time interval is specified by the system value QINACTITV). The interactive job can be ended, disconnected, or message CPI1126 may be sent to the message queue that you specify. The message queue must exist in the system ASP or in a basic user ASP. If the specified message queue does not exist or is damaged when the inactive time-out interval is reached, the messages are sent to the QSYSOPR message queue. All of the messages in the specified message queue are cleared during an IPL. If you assign a user's message queue as QINACTMSGQ, the user loses all messages that are in the user's message queue during each IPL.

**QPRBFTR** Problem log filter

This value specifies the name of the filter object used by the Service Activity Manager when processing problems. The filter must exist in the system ASP or in a basic user ASP.

**QPWDVLDPGM** Password validation program

This value provides the ability for a user-written program to perform additional validation on passwords. The program must exist in the system ASP or in a basic user ASP.

**QRMTSIGN** Remote sign-on control

This system value specifies how the system handles remote sign-on requests.

**Remote session program**

All values are supported as described for display station pass-through, IBM System i® Access for Windows® Work Station Function (WSF), and other 5250 emulation programs on programmable workstations. For information about the level of support provided for TELNET sessions, review the documentation on this system value in the IBM i Information Center on the Web at: http://publib.boulder.ibm.com/iseries/

**Program** This option allows you to specify the name of a program and library to decide which remote sessions to allow and which user profiles to automatically sign on from which locations. The program must exist in the system ASP or in a basic user ASP.

**QSRTSEQ** Sort sequence

This system value specifies the default sort sequence algorithm to be used by the system.

*Sort sequence table name* specifies the name and library of the sort sequence table to be used. The sort sequence table must exist in the system ASP or in a basic user ASP.

**QSTRUPPGM** Startup program

This value specifies the name of the program called from an autostart job when the controlling subsystem is started. This program performs setup functions, such as starting subsystems and printers. The program must exist in the system ASP or in a basic user ASP.

**QSYSLIBL** System part of the library list

When searching for an object in the library list, the libraries in the system part are searched before any libraries in the user part are searched. The list can contain as many as 15 library names. The libraries must exist in the system ASP or in a basic user ASP.

**QUPSMSGQ** Uninterrupted power supply (UPS) message queue

This value specifies the name and library of the message queue that will receive UPS messages. It allows you to monitor the message queue and control the power down. If the message queue is not the system operator message queue (QSYS/QSYSOPR), all UPS messages are also sent to the system operator message queue.

**QUSRLIBL** User part of the library list

When searching for an object in the library list, the libraries in this part are searched after the libraries in the system part and after the product library and current library entries. The list may contain as many as 25 library names. The libraries must exist in the system ASP or in a basic user ASP.

## 3.3.2 Network attribute settings

When you set up independent disk pools for the first time or move applications to independent disk pools, consider some of the keywords and parameters for the system network attributes. If the keywords and parameters highlighted in the following sections are in use, review them for the impact that independent disk pools may have on their use. These parameters are on the Change Network Attributes (CHGNETA) command. Some of them are on the Retrieve Network Attributes (RVTNETA) command. For more information about these commands see the IBM i Information Center:

`http://publib.boulder.ibm.com/iseries/`

To access this function, type `CL Command Finder` in the Search field.

### Alert Filters (ALRFTR)

This parameter specifies the qualified name of the alert filter used by the alert manager when processing alerts. The alert filter must exist in the system ASP or in a basic user ASP. The possible values are:

**SAME** The value does not change.

**NONE** An alert filter is not used.

The name of the alert filter can be qualified by one of the following library values:

***LIBL** All libraries in the job's library list are searched.

***CURLIB** The current library for the job is searched. If no library is specified as the current library for the job, the QGPL library is used.

**library-name** Specify the name of the library to be searched.

**filter-name** Specify the name of the alert filter.

### Message queue (MSGQ)

This parameter specifies the qualified name of the message queue where messages received through the SNADS network are sent for users with no message queue specified in their user profile or whose message queue is not available. The message queue must exist in the system ASP or in a basic user ASP. The possible value is *SAME. This value does not change.

The possible library values are:

***LIBL** The library list is used to locate the message queue.

When *LIBL is used as the library name, the library list of the job calling this command is searched to find a message queue with the specified object name. If the message queue is found, the name of the library in which it is found is used in the fully qualified name and it is stored. If the message queue is not found, an exception is signaled, and no network attributes are changed.

***CURLIB** The current library for the job is used to locate the message queue. If no library is specified as the current library for the job, the QGPL library is used.

**library-name** This specifies the name of the library where the message queue is located. When the library name or *CURLIB is specified, this command attempts to find the message queue. If the message queue cannot be found in the specified library, a diagnostic message is sent. If all other parameters on the command are specified correctly, as is whether this command can find the message queue in the library specified, the MSGQ network attribute is changed to the qualified message queue name.

**message-queue-name**
The specifies the name of the message queue.

## Distributed data management access (DDMACC)

This parameter specifies how the system processes distributed data management (DDM) and DRDA® requests from remote systems for access to the data resources of the system. The DDM and DRDA connections refer to APPC conversations or active TCP/IP or OptiConnect connections. Changes to this parameter are immediate and apply to DRDA, DDM, or DB2 Multisystem applications. However, jobs that are currently running on the system do not use the new value. The DDMACC value is accessed only when a job is first started. You must specify a special value or program name that dictates how the requests are to be handled.

If a program name is specified, the program must exist in the system ASP or in a basic user ASP. The possible values are:

***SAME** The value does not change.

***REJECT** This system does not allow DDM and DRDA requests from remote systems. However, this system can still use DDM or DRDA to access files or SQL tables on remote systems. Source (client) systems cannot access files or SQL tables on a destination system when *REJECT is used.

***OBJAUT** If the user profile associated with the DDM or DRDA job is authorized to the files, all file or remote SQL requests are accepted. Object authorities, such as read, write, or update, must also exist for the files.

The possible library values are:

***LIBL** The library list is used to locate the validation program.

If *LIBL is used as the library name, the library list of the job calling this command is searched to find the program name with the specified object name. If the program name is found, the name of the library in which it is found is used in the fully qualified name and it is stored. If the program name is not found, an error message is sent, and no network attributes are changed.

***CURLIB** The current library for the job is used to locate the validation program. If no library is specified as the current library for the job, the QGPL library is used.

**library-name** This specifies the name of the library where the validation program is located. When the library name or *CURLIB is specified, this command attempts to find the program name. If the program name cannot be found in the specified library, a diagnostic message is sent to the user. If this command can find the program name in the

specified library, the DDMACC network attribute is changed to the qualified program name if all other parameters on the command are specified correctly.

**program-name** This specifies the name of the validation program.

The program name is the name of the client validation program that can supplement system object-level security. This user-exit program can restrict user access to *PUBLIC and privately authorized files. The target DDM support calls the user program each time that a file is read. The user exit program indicates to DDM whether the request must proceed or end. The IBM i object-level security still applies.

## PC support access (PCSACC)

This parameter specifies how IBM System i Access for Windows requests are handled. You must specify a special value or program name that dictates how the requests must be handled. This permits greater control over IBM System i Access for Windows applications. Changes to this parameter are immediate. However, jobs currently running on the system do not use the new value. The PCSACC value is used only when a job is first started.

If a program name is specified, the program must exist in the system ASP or in a basic user ASP. The following values are possible:

***SAME** The value does not change.

***REJECT** Normal object authorizations are checked for this client request (for example, authorization to retrieve data from a database file for a transfer facility request).

***REGFAC** The system uses the system's registration facility to determine which exit program (if any) to run. If no exit program is defined for an exit point and this value is specified, *OBJAUT is used.

The possible library values are:

***LIBL** The library list is used to locate the program.

When *LIBL is used as the library name, the library list of the job calling this command is searched to find the program name with the specified object name. If the program name is found, the name of the library in which it is found is used in the fully qualified name, and it is stored. If the program name is not found, an error message is sent and no network attributes are changed.

***CURLIB** The current library for the job is used to locate the program. If no library is specified as the current library for the job, the QGPL library is used.

**library-name** This specifies the name of the library where the program is located. When the library name or *CURLIB is specified, this command attempts to find the program name. If the program name cannot be found in the specified library, a diagnostic message is sent to the user. If all other parameters on the command were specified correctly, as is whether this command can find the program name in the specified library, the PCSACC network attribute is changed to the qualified program name.

**program-name** The specifies the name of the program. The program name is the name of the client-supplied IBM System i Access for Windows host system application exit program that can supplement system object-level security. This user-exit program can restrict requests

handled from the client. Each personal computer support application calls the exit program for requests from the client. Two parameters are passed to the user-exit program. The first describes the client request (which application and what kind of request). The second is used by the exit program to indicate to the client support application whether this client request must be handled.

### Allow add to cluster (ALWADDCLU)

Before you can add a node to a cluster, network attribute ALWADDCLU must be set to *RQSAUT or *ANY to allow the add. If the attribute is set to *RQSAUT, it will require validation using X.509 digital certificates, and the node running the command and the node being added must have the following features installed:

- Operating system option 34 (Digital Certificate Manager)
- Cryptographic Access Provider Product (AC2 or AC3)

**Tip:** You must have *IOSYSCFG authority to change the network attribute ALWASSCLU

The following values specify the authorizations that control how a node is added to a cluster:

| | |
|---|---|
| ***SAME** | The value does not change. They system is shipped with a value of *NONE. |
| ***NONE** | No other system can add this system as a node in a cluster. |
| ***ANY** | Any other system can add this system as a node in a cluster. |
| ***RQSAUT** | Any other system can add this system as a node in a cluster only after the cluster add request has been authenticated. |

## 3.4 Software requirements

Independent disk pool support is included as part of the operating system. You may need to consider other software components that may be required for your particular configuration.

### 3.4.1 Required software

The Systems Director Navigator for i5/OS is included as part of the IBM i 6.1 operating system. Either Systems Director Navigator for i5/OS or System i Navigator is required to perform some of the disk management tasks necessary to implement independent disk pools. Some functions are not available using *green screen* service tools. In this manual we use Systems Director for i5/OS for all examples. See Chapter 2, “Creating an IASP” on page 9, for information about creating disk pools.

### 3.4.2 Optional software

The following products are optional:

- ObjectConnect, Product Option 22 (5761-SS1)

  This product provides support to simply and efficiently move individual objects, entire libraries, or entire integrated file system (IFS) directories from one Power Systems server to another over a standard communications connection or over a high-speed fiber optic bus. Systems can be connected using:

  - Standard advanced program-to-program communication (APPC) using Advanced Peer-to-Peer Networking (APPN)
  - TCP/IP communications lines using AnyNet®
  - A fiber optic bus using OptiConnect for Power Systems servers

  The economy of not requiring intermediate-save file procedures and copies to distribution queues saves DASD. It improves performance in a manner that is nondisruptive to system operations. This product is not required for normal independent disk pool functionality. This is a no-charge feature.

- OptiConnect, Product Option 23 (5761-SS1)

  This product provides high-speed transparent access to data through fiber optic bus connections and performance enhancements to IBM i Distributed Data Management (DDM). It also provides virtual OptiConnect for communications when switching independent disk pools between LPARs. This product is not required for normal independent disk pool functionality. This is a charged feature.

- HA Switchable Resources, Product Option 41 (5761-SS1)

  This is required when setting up simple clustering for switching between two systems.

  **Note:** Keep in mind that switchable independent disk pools require IBM i clustering support, which is available through separately priced IBM License Product 5761-SS1 Option 41.

- HA Journal Performance, Product Option 42 (5761-SS1)

  This is a journal caching feature. You should review it if there is a plan to put journals in secondary independent disk pools.

## 3.5 Application integration

You must review the applications to be loaded into independent disk pools in terms of their current level of integration between themselves and the operating system. Use the checklist in Appendix B, “Application object planning checklist” on page 189, to assist you in placement of various objects types. Traditionally, data objects pertaining to an application area are stored in a data library. Program objects pertaining to an application area are stored in a program library. Other objects common to the application area are stored in system libraries or libraries designated as common to that application. Keep in mind that library lists are maintained differently when using independent disk pools. When an independent disk pool is made unavailable to a user, library names that reside in the independent disk pool are removed from the user’s library list. Library names are *not* automatically added to a user’s library list when the disk pools are made available. The list must be managed by the application.

See Chapter 4, “Accessing an independent auxiliary storage pool (IASP)” on page 57, for detailed steps for loading applications into independent disk pools. Make a careful note of the considerations by object type for supported objects.

## 3.6 Authority considerations

Considerations relating to authority are a key part of moving applications to independent disk pools. This section discusses the characteristics of user profiles and authorization lists as they relate to the security of an independent disk pool.

### 3.6.1 User profiles and independent disk pools

User profile information is stored in the system disk pool. Each user profile object is an object type of *USRPRF. Copies of *USRPRF objects are not in any independent disk pool. However, some user profile information is maintained in the independent disk pool itself.

Each object in an independent disk pool requires this user profile information:

- The owner of the object
- The primary group of the object
- The private authority entries for the object

Figure 3-1 illustrates this concept.

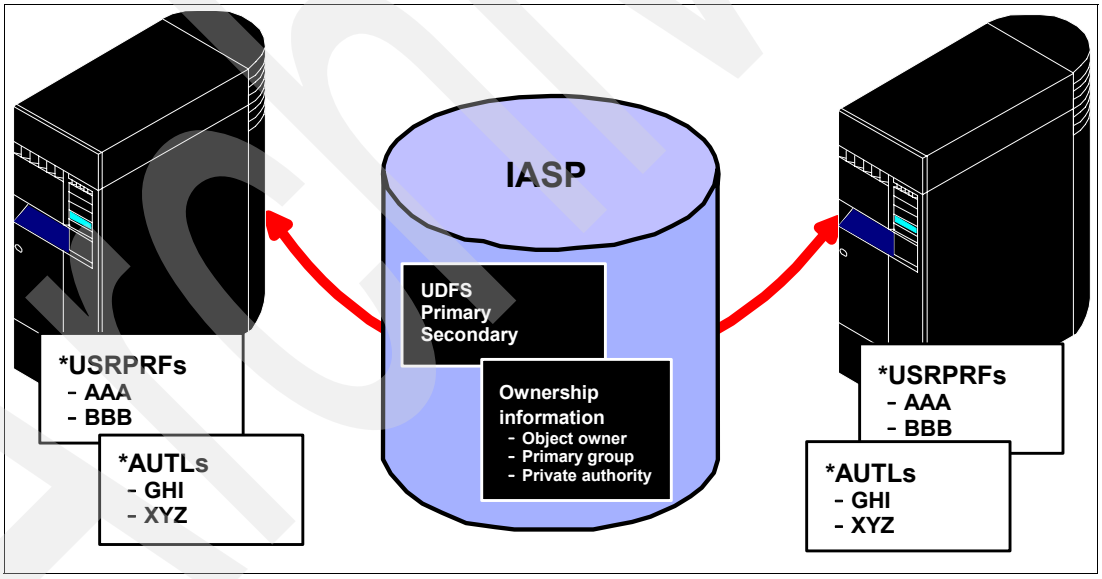


*Figure 3-1 Independent disk pools and user profiles*

Additional storage (above that consumed by objects) is required for these system security structures. This is necessary to make the independent ASP self contained. These structures consume disk space within the independent disk pool. The percentage varies and depends on the size and number of objects referenced by these structures.

For each user profile that owns or has private authority to an object in an independent disk pool, the system stores information about the user profile in an internal structure called a *user profile extension*. The user profile extensions, while not visible nor accessible to users, are stored in the independent disk pool.

An independent disk pool may be varied on by the vary configuration (VRYCFG) command, an independent disk pool failover, or may be switched from another node. Anytime that this happens, the system matches the user profile extension information in the independent disk pool with a corresponding user profile object in the system ASP. This match is based solely on the name of the user profile.

If a user profile object does not exist with the name saved in a user profile extension, then the system creates a user profile object with the saved name. In this case, the user profile is created with the attributes listed in Table 3-1.

*Table 3-1 User profile attribute when its object does not exist with the name of the extension*

| Keyword | Parameter description |
|---|---|
| USRPRF | Saved name |
| PASSWORD | *NONE |
| STATUS | *DISABLED |
| UID | Saved value or *GEN |
| GID | Saved value or *GEN |
| TEXT | Created by auto-configuration |
| All others | CRTUSRPRF command default values |

In addition, during vary on of an independent disk pool, the system verifies that the user ID number (UID) and group ID number (GID) values saved in the user profile extension are the same as the values in the matching user profile object. If these values do not match, the system performs recovery functions to ensure that the system is using a consistent set of UIDs and GIDs for all user profiles and for all objects on the system that use the UID or GID values.

The result of these system operations during the vary on of an independent disk pool means that:

- If a user profile is deleted while an independent disk pool is varied off or switched to another node, the user profile may reappear when the independent disk pool is varied on.
- If the UID or GID values of a user profile are changed (including changing a user profile to be a group profile) while an independent disk pool is varied off or switched to another node, then the UID or GID values associated with objects on the independent disk pool may change during vary on processing.
- The time to vary on an independent disk pool may be longer if user profiles are deleted or their UID or GID values changed.

## 3.6.2 Authorization lists (AUTL) and independent disk pools

Authorization lists are handled in a manner similar to user profiles. Authorization lists are only stored in the system ASP. The system maintains an *extension* in the independent disk pool. *AUTL objects are matched to extensions while varying on. An *AUTL is created if one does not exist.

Authorization lists are used to secure access to objects in a system, irrespective of ASP boundaries. An authorization list (similar to a user profile) is implemented as a distributed object. That is, it can secure objects within the system ASP and any independent disk pools.

There is only one copy of each authorization list (*AUTL) object in the system ASP. Copies of *AUTL objects are not kept in independent disk pools. However, when an object in an independent disk pool is secured by an authorization list, the name of that list is kept in the independent disk pool.

Figure 3-2 illustrates this concept. Authorization list names AUTL111 and AUTL222 are stored in the independent disk pool, because the objects HJK and RST are secured by these lists. There is no reference in the independent disk pool itself to authorization list name AUTL333 because the object LHR is not stored in the independent disk pool.

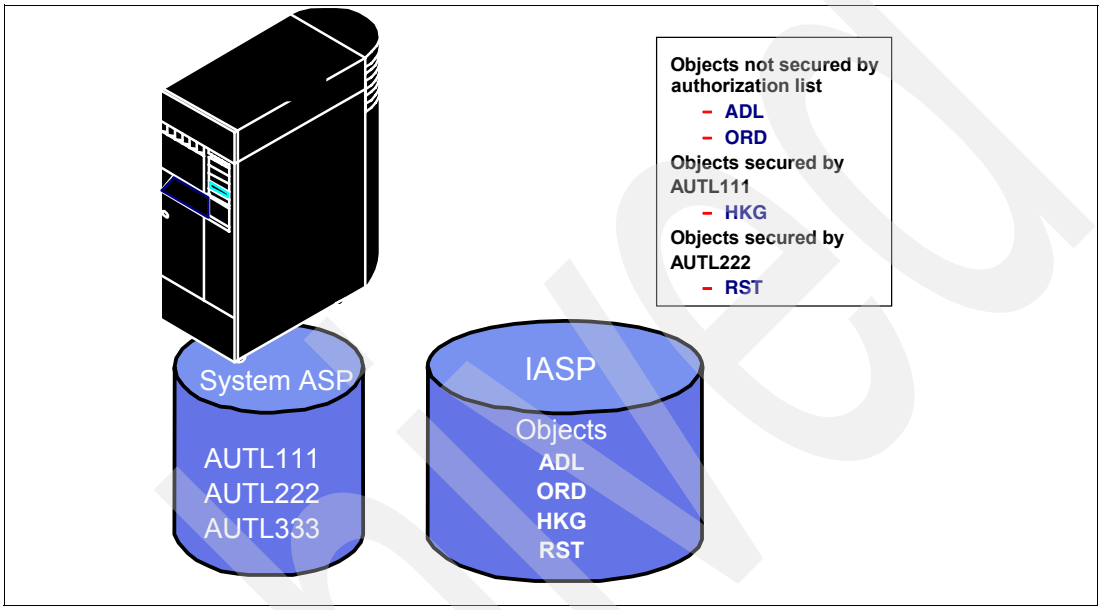


*Figure 3-2 independent disk pools and authorization lists*

As with user profiles, the System Licensed Internal Code (SLIC) is responsible for handling this function. The user interface does not change for storing, retrieving, changing, or deleting the authorization list information in an independent disk pool.

## 3.7 Hardware requirements

To create a non-switchable independent disk pool, you only need one or more non-configured disk units. The disks can be physical or virtual disk drives. Non-configured disks can be obtained by removing existing disk units from existing ASPs or adding new non-configured disks to the system.The disk units will be formatted during the independent disk pool definition process using Systems Director Navigator for i5/OS.

Independent disk pools can also be switched between a multi-partitioned system. There are additional hardware and software requirements when using switchable independent disk pools. When a tower is referred to as being switchable, it is meant in the logical sense of the word. The software in the operating system performs the switching of the resources from use by one partition to use by another. There is no physical switching of cables. Keep in mind that you can configure multiple independent disk pools on an IOP or in a tower, and everything in the tower switches when the switch-over is performed.

4

# Accessing an independent auxiliary storage pool (IASP)

This chapter discusses accessing the IASP and the application considerations. Namespace is a key concept in understanding how to access the ASP and a relational database. The implications of addressing multiple relational databases (RDBs) in IBM i and Structured Query Language (SQL) environments is also addressed.

## 4.1 Namespace and relational database

Namespace and relational database concepts are explained and explored here. Definitions are provided as a level set.

**Namespace** The namespace is the group of libraries accessible to a thread.

**ASP group** An ASP group consists of a primary IASP and zero or more secondary IASPS linked to it, reference by the name of the primary. Only one ASP group can be associated with a thread at one time.

**Relational database** A relational database is a database that can be perceived as a set of tables and can be manipulated in accordance with the relational model of data. The relational database contains a set of objects used to store, access, and manage data. The set of objects includes tables, views, indexes, aliases, distinct types, functions, procedures, sequences, and packages. IBM i has three types: system, user (includes IASP), and remote relational database.

When a user signs on, all objects in *SYSBAS (default namespace) become available to them (Figure 4-1).

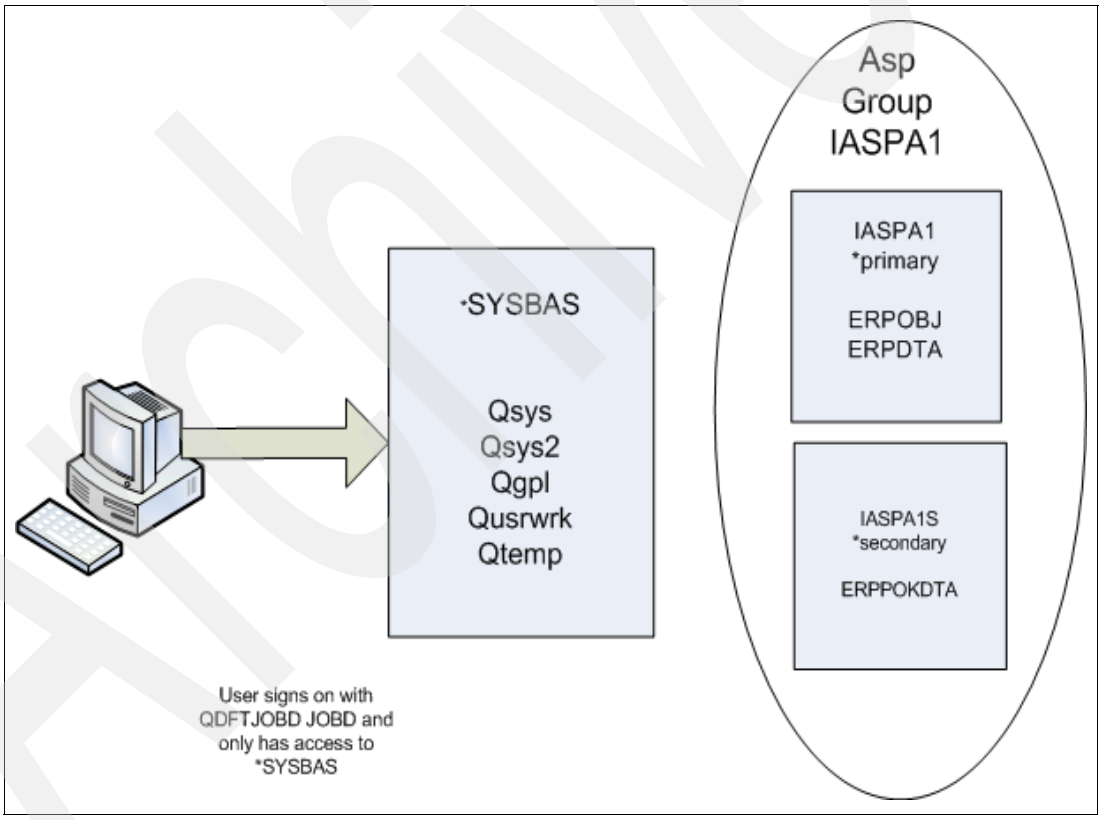


*Figure 4-1 User signs on: access to *SYSBAS*

An IASP that contains libraries must be varied on in order for the libraries to be accessible. In addition, the IASP must also be explicitly associated with a thread to allow for the thread to access its libraries (this is referred to as adding the IASP into the threads namespace).

There are several methods used to add an IASP to a user's namespace:

- Use the SETASPGRP command to direct a job to explicitly use an IASP.
- Use the INLASPGRP parameter on the *JOBD associated with the job to set the IASP when the job starts. (JOBDs beginning with Q should *not* be modified to include an INLASGRP.)

> **Tip:** If you use QDFTJOBD, make a copy of it and modify the new description using the INLASPGRP and INLLIBL parameters. Then modify the user profiles to specify the new job description.

- Use the INLASPGRP parameter in the SBMJOB command (the default is *CURRENT).
- Using the Java™ toolbox, use the connect class to connect to the IASP.

In Figure 4-2, a user signs on with the appropriate JOBD using the initial ASP parameter (INLASPGRP) and now has access to all the libraries in ASP group IASPA1.

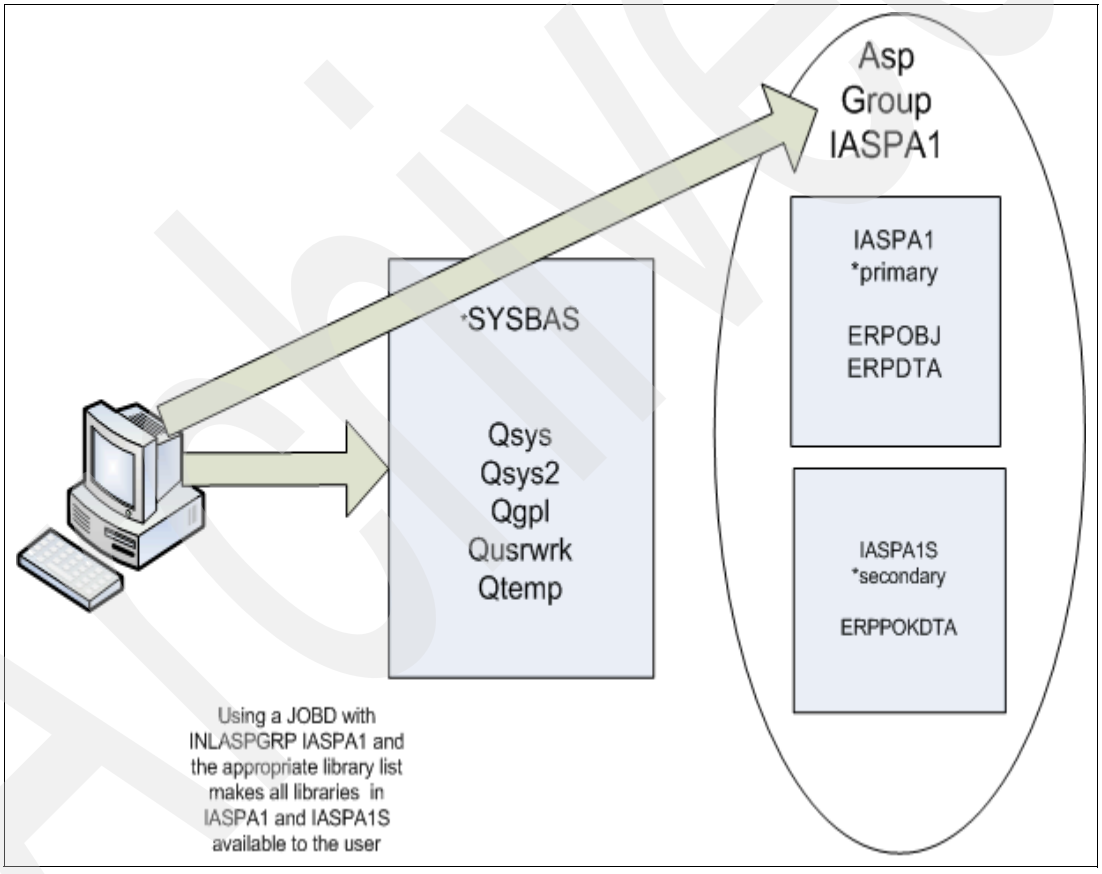


*Figure 4-2 User signs on with appropriate initial ASP group parameter set to IASPA1*

A thread can reference, by name, all of the libraries in the IASPs of one ASP group. A thread that does not have an ASP group component in its namespace has its library references limited to the *SYSBAS component. A thread with an ASP group component in its library namespace can reference libraries in both the *SYSBAS and the ASP group.

When attaching a job/thread to an IASP and establishing a library list for the job/thread, it is important to understand the work management initiation procedure. Subsystem or interactive job initiation will only search the system ASP for subsystem and job descriptions, classes, and routing programs. An attempt to place an IASP library in either the QSYSLIBL or QUSRLIBL will fail. The libraries must be added to the job/thread list using the INLASPGRP/INLIBLL job description parameter, the SETASPGRP command, or a combination of the SETASPGRP command and the ADDLIBLE command.

Library names do not need to be unique on a system. Library names must be unique in every possible namespace. Since *SYSBAS is a component of every namespace, the presence of a library name in *SYSBAS precludes its use within any IASP. All libraries in all IASPs of an ASP group are part of a namespace. Existence of a library name within one IASP of an ASP group precludes its use within any other IASP of the same ASP group. A namespace can have only one ASP group component. A library name that is not used in *SYSBAS can be used in any or all ASP groups.

If an IASP is not available, a library with the same name as a library in the IASP may be created in *SYSBAS. Attempting to vary on the IASP will fail. The IASP status will be *Active* but you will not be able to access the IASP. You must delete either of the conflicting libraries. Then the IASP must be varied off and back on to be accessed.

IBM i has a file interface and an SQL interface to its databases. The file interface uses the namespace to locate database objects. Each namespace is treated as a separate relational database by SQL. For compatibility, SQL maintains a catalog for each ASP group. This catalog resides in the primary IASP of the ASP group. The catalog is built from the objects that are in a namespace that has the ASP group and *SYSBAS as its two components. The RDB and the namespace are somewhat interchangeable because they refer to the same set of database objects. It is required that all relational databases accessed by SQL are defined in the RDB directory on the system.

The default relational database name is the name of the primary IASP. You can assign unique database names to each IASP for a more meaningful RDB schema. Figure 4-3 shows the RDB entries on a system.

```
Work with Relational Database Directory Entries

Position to  . . . . . .

Type options, press Enter.
  1=Add   2=Change   4=Remove   5=Display details   6=Print details

                            Remote
Option  Entry               Location                  Text

        IASPA               LOOPBACK                  Entry added by system
        IASPA1              LOOPBACK                  Entry added by system
        IASP93ADB           LOOPBACK                  Entry added by system
        RCHAS93             *LOCAL

                                                                        Bottom
F3=Exit   F5=Refresh   F6=Print list   F12=Cancel   F22=Display entire field
(C) COPYRIGHT IBM CORP. 1980, 2007.
```

*Figure 4-3 WRKRDBDIRE on system RCHAS93*

The system (RCHAS93) has four RDBs:

- The system ASP (RCHAS93)
- Three ASP groups:
  - ASPA
  - IASPA1
  - IASP93ADB

On this system, RDB IASPA1 includes disk pools IASPA1 and IASPA1S.

## 4.2 Relational database directory

IBM i uses the RDB directory to define the relational database names that can be accessed by system applications and to associate these relational database names with their corresponding network parameters. Each Power Server in the distributed relational database network must have a relational database directory configured. There is only one relational database directory on a system. The system also uses the directory to specify whether the connection uses SNA or IP.

The RDB name assigned to the local RDB must be different from any other RDB in the network. It must be unique. Names assigned to other RDBs in the directory identify remote RDBs or local user databases. The names of remote RDBs must match the name that an AS uses to identify its local system database or one of its user databases, if configured. If the local system RDB name entry for an application server does not exist when it is needed, one is created automatically in the directory. The name used is the current system name displayed by the display network attributes (DSPNETA) command.

The RDB directory allows entry of an application requester driver (ARD) program to specify communication information. An application requester driver program is an exit program that enables SQL applications to access data managed by a database management system other than DB2 for IBM i. The ARD program must reside in *SYSBAS.

The ARD program allows the RDB name from the application to translate into the appropriate IP address or host name and port. Or it can translate it into the appropriate SNA network identifier and logical unit (LU) name values for communications processing. The RDB directory entry can specify the user's preferred outbound connection security mechanism. Each ARD in the distributed relational database network must have an entry in its relational database directory for its local RDB and one for each remote and local user RDB that the AR accesses. Any system in the distributed RDB network that acts only as an application server does not need to include the RDB names of other remote RDBs in its directory.

## 4.3 Integrated File System (IFS)

IFS objects are stored in a directory structure. Access to the objects is by a path that navigates the directory structure to reach the object. An available IASP has a directory in the root directory. The directory has the same name as the IASP. When the IASP is available, the contents of the IASP are *mounted* to the IASP directory. In the example in Figure 4-4, the system has an IASP named IASPA1, using the WRKLNK command to show the IASPA1 directory and its contents.

```
Work with Object Links

Directory  . . . . :   /IASPA1

Type options, press Enter.
  2=Edit   3=Copy   4=Remove   5=Display   7=Rename   8=Display attributes
  11=Change current directory ...

Opt   Object link            Type     Attribute    Text
      QIBM                   DIR
      QSYS.LIB               DIR

                                                                    Bottom
Parameters or command
===>
F3=Exit   F4=Prompt   F5=Refresh   F9=Retrieve   F12=Cancel   F17=Position to
F22=Display entire field          F23=More options
```

*Figure 4-4 WRKLNK command to view directory of IASP named IASPA1*

When an IASP is not available (or before the IASP is created), it is possible to create a directory with the name of the IASP. If there is a directory with the same name as an IASP, when the IASP is varied on:

- The MOUNT operation will be successful if the existing directory is empty.
- The MOUNT operation will fail if there are any objects in the existing directory. The vary-on will not fail. The first indication of failure is likely to occur when users try to access objects in the directory. The objects will be missing or incorrect. The only indication that the MOUNT operation failed is a message in the joblog of the thread that performed the vary-on operation. If the IASP environment uses IFS, each vary-on operation should be checked to ensure that the IFS mounted properly.

Applications that have a hard-coded path can access IFS objects in IASPs by using a symbolic link from the original location to the IASP location.

**Tip:** Check to see whether you have any IFS directories that exist using the same name as the primary IASP that you plan to create.

## 4.4 IASP and RDB distinctions

The distinctions between the IASP and the RDB entry created when setting up an IASP are subtle, yet important.

The difference between CRTLIB and CRT COLLECTION is that the CRTLIB command defaults to *SYSBAS. CRT COLLECTION defaults to the current library namespace.

Using the RESTORE or RSTLIB command runs the CRTLIB command under the covers (unless the library already exists). For the IBM i command specify the *ASPDEV or *RSTASPDEV parameters to put an object in the appropriate IASP.

## 4.5 IASP as a separate database

Each ASP group is given a database name. The name of the primary IASP in an ASP group is used if the RDB name is not assigned. It is a namespace that is treated as a separate RDB by SQL. A namespace, which the RDB may be part of, consists of the system ASP (*sysbas), user ASPs, and an ASP group.

The user may create additional RDBs on a Power Server by configuring IASPs on the system. Each *primary* IASP is an RDB. It consists of all the database objects that exist on the IASP disks. Additionally, all database objects in the system RDB of the Power Server, to which the IASP is connected, are logically included in a user RDB. Therefore, the name of any schema created in a user RDB must not already exist in that user RDB or in the associated system RDB.

Although the objects in the system RDB are logically included in a user RDB, certain dependencies between the objects in the system RDB and the user RDB are *not* allowed. These include:

- Creating a view into a schema that exists in the same RDB as its referenced tables, views, or functions
- Creating an index into a schema that exists in the same RDB as its referenced table
- Creating a trigger or constraint into a schema that exists in the same RDB as its base table
- Ensuring that the parent table and dependent table in a referential constraint both exist in the same RDB
- Creating a table into a schema that exists in the same RDB as any referenced distinct types
- Ensuring that the parent table and the dependent table in a referential constraint both exist in the same RDB

Other dependencies between the objects in the system RDB and the user RDB are allowed. For example, a procedure in a schema in a user RDB may reference objects in the system RDB. However, operations on such an object may fail if the other RDB is not available, such as when a user RDB is varied off and then varied on to another system. A user RDB is local to the IBM i server while the IASP is varied on. IASPs can be varied off on one Power Server and then varied on to another Power Server. Therefore, a user RDB may be local to a given Power Server at one point in time and remote at a different point in time.

### 4.5.1 SQL connections

In an SQL environment, SQL CONNECT is used to specify the correct database. To achieve best performance, make sure that the database being connected to corresponds with your current library namespace. You can use the SETASPGRP command to achieve this. If the SQL CONNECT function is not operating within the same library namespace, the application uses Distributed Relational Database Architecture™ (DRDA) support, which can affect performance.

There are two types of connections:

- Application connections
- System connections

#### Application connections

Application connections are established based on the following rules:

- These connections are similar to the types of connections supported in previous releases. Under this type of connection, only one local connection per activation group is allowed at a time. Any other connections use DRDA.
- Distributed connection rules for both distributed unit of work (DUW) and remote unit of work (RUW) are applied equally to all RDBs and to implicit and explicit connections. The rules include:
  - If RUW is used, a connection request fails with the SQL0752 `Connection cannot be changed` error message if there are pending transactions.
  - If RUW is used, only one active connection is allowed. All the resources associated with the current connection are released before a new connection is started. For example, all open cursors are closed, all prepared statements are destroyed, and all normal SQL locks are released.
- The CONNECT statement to a local RDB is processed as follows:
  - For RUW, the local connection is started if the target RDB corresponds to the current namespace. Otherwise, DRDA is started.
  - For DUW, the local connection is started if the target RDB corresponds to the current namespace and there are no existing local connections. Otherwise, DRDA is started.
- If the first SQL statement in the activation group is not CONNECT, an implicit connection occurs. For a local program, SQL starts the local connection to the current namespace (RDB). For a distributed program, SQL starts the DRDA connection to the RDB specified in the RDB parameter at compile or SQL package creation time.
- If the activation group is already connected to an IASP, the first SQL statement following a namespace change to an IASP that is different from the current connection must be CONNECT or SET CONNECTION. Otherwise the SQL statement causes the SQL0752 `Connection cannot be changed` error message with reason code 9.

### 4.5.2 Changing RDBs

The RDB (namespace) can be changed using the SETASPGRP command. A RDB cannot be changed by using the CONNECT statement with a new database name. You must use one of the three methods prior to using the CONNECT statement. RDB changes are not allowed in stored procedures, user-defined functions, or triggers.

### 4.5.3 Object creation

Objects that are tightly coupled to other objects must be created in the same IASP. Notice that this is not in the same namespace, but the same IASP.

SQL objects do not allow spanning across IASP boundaries. Spanning between primary and secondary IASPs in an ASP group *is* allowed. Spanning between *SYSBAS and an IASP is *not* allowed. For example, applications cannot create a view across libraries in both *SYSBAS and an IASP.

Applications that use commitment control may not update objects in *SYSBAS and the IASP within the same scope of work or within the same connection. Often the application programmer may be unaware that there are objects in *SYSBAS. These may be control tables or fields in some product library or they may be views or indexes where the name was not qualified.

While it is possible to create files, tables, and so on, in QSYS2, the corresponding library in the independent disk pool prevents this from occurring. Most applications that create data in QSYS2 do not realize it and fail when running in an independent disk pool. Consider the following example:

```
CHGCURLIB DEMO10
create view ICTABLES(Owner, tabname, type) as select table_schema, TABLE_NAME, TABLE_TYPE
from SYSTABLES where table_name like'IC%'
```

In this example, the view ICTABLES is not built in the current library (DEMO10) as you would expect. It is built in the library of the first table that is mentioned, which is QSYS2. It fails when accessing the independent disk pool because creation of objects in QSYS2XXXXX is prevented.

## 4.6 ODBC and JDBC considerations

JDBC and ODBC operations access information in a database. If the data is moved from *SYSBAS to an IASP, the applications must be capable of accessing the data in the IASP. In order to allow a job/thread that uses JDBC or ODBC to access data in an IASP without application modification, the job/thread should be initiated with a job description that attached the job/thread to the IASP. In addition, both ODBC and JDBC statements have new parameters that will retrieve data from an IASP database.

The parameter added to a JDBC operation is ‘ds.setDatabaseName’. Use this parameter to specify the name of the IASP database.

in ODBC, the DATABASE parameter is used to specify the name of the IASP database.

## 4.7 System-managed access-path protection considerations

System-managed access-path protection (SMAPP) can reduce the amount of time that it takes to restart your system or vary on an independent disk pool after an abnormal end. When you create a new disk pool, the access recovery time for that disk pool is set to *NONE. You can use the EDTRCYAP command to set a target recovery time for the disk pool if desired (Figure 4-5).

```
Edit Recovery for Access Paths                RCHAS93
                                                        08/28/09  13:00:55
Estimated system access path recovery time  . . . :           3   Minutes
Total not eligible recovery time  . . . . . . . . :           0   Minutes
Total disk storage used . . . . . . . . . . . . . :        .458   MB
% of disk storage used  . . . . . . . . . . . . . :        .000

Type changes, press Enter.
  System access path recovery time  . . .   50         *SYSDFT, *NONE, *MIN,
                                                         *OFF, Recovery time
  Include access paths  . . . . . . . . .   *ALL       *ALL, *ELIGIBLE

            ------Access Path Recovery Time------  -----Disk Storage Used-----
ASP               Target          Estimated          Megabytes     ASP %
1                 *NONE                1                .090      .000
IASPA1S           *MIN                 0                .090      .000
IASP93A           *NONE                0                .090      .000
IASP93AS1         *MIN                 0                .090      .000
IASPA1            *NONE                1                .098      .000

                                                                     Bottom
F3=Exit   F4=Prompt   F5=Refresh        F12=Cancel
F13=Display not eligible access paths   F24=More keys
```

*Figure 4-5 EDTRCYAP screen*

Using SMAPP helps limit the vary-on duration as well as the quantity of background job activity, which must make each access path whole when you vary on your independent disk pools after an abnormal vary off.

The recovery time that you specify becomes an attribute of the independent disk pool and moves with the pool. The only occasion when the specified recovery time is not moved is when the system that you are moving the independent disk pool to has its system-wide recover time specified as *OFF.

## 4.8 IASPs with SQL

When dealing with applications in an SQL environment, the main considerations are:

- Determine the local database and how it is different from the *local relational database entry.
- Understand the method used to access data in an IASP not in the job/thread namespace.

- Understand the restrictions regarding creation of tables in a job/thread namespace.
- Understand the performance implications for the placement of status SQL procedures.

A system with one or more IASPs will have a relational database entry for each IASP and a database entry that is *local. Typically, *local is the data located in *SYSBAS.

A user now running a system with an IASP has the ASP group in the job/threads namespace. The local database is now *sysbas (*local) plus the IASP. As long as SQL operations us the default for the local database, there are no application issues. If the application retrieves the name of the *local database and uses that name as the local database, only data in *sysbas will be local and the application will not function properly when the data is in an IASP. Specifying the name of the IASP database will work properly.

A CONNECT statement can be used to change the namespace. The CONNECT statement will remove the current disk pool database from local and add a new database to local. Optionally, the job could use the SETASPGRP command to attach to a new database.

In general, when SQL is active, commitment control is activated for the database tables. For a job/thread with a local database that includes an IASP, the scope of the commit function is the IASP database only. From the viewpoint of commitment control, the *local database and the IASP database are two independent databases and commitment control will not span databases unless two-phase commit is implemented.

Some tips and recommendations:

- Many characteristics of IASPs are the same as separate systems. Ensure that you understand what is meant by local.
- Compile applications in a job/thread with an IASP in its namespace and specify *LOCAL for the RDB in order to avoid unnecessary connect/overhead. This also reduces the necessity of managing SQL packages.
- Applications that will require data across multiple IASPs reside in either *SYSBAS or an IASP. Use DUW with the set connect function of SQL to change the connections to the required data.
- When moving data into IASPs, current DRDA connections must be handled. The easiest way to handle this is to set the name of the *LOCAL database to something else and make the IASP primary the name of the old *LOCAL database.
- A static SQL application should be placed in IASPs if they will be used with multiple namespaces for performance reasons.

## 4.9 STRQMQRY and STRQMPRC RDB support

You can resolve the SQL objects (tables, functions, views, types) that are referenced in a query management query (QMQRY) object. To do this use the RDB specified on the RDB parameter or the RDB specified on the CONNECT/SET CONNECTION commands. This RDB may be an IASP. The query management objects referenced must be in the current RDB (namespace).

When output from a STRQQRY command is directed to an output file, query management ensures that the output file is created on the RDB (namespace) that was current at the time that the SRQMQRY is executed.

## 4.10 Web query

Web query only references objects in the current RDB (namespace). A *QRYDFN object created in *SYSBAS may reference files in an IASP and vice versa. If a *QRYDFN object created to reference objects in an IASP runs when a different IASP is set as the current RDB (namespace), the *QRYDFN runs successfully if the new IASP contains objects with the same name and the file formats are compatible.

## 4.11 Journaling considerations

This section explains some of the rules for journaling where IASPs are concerned.

The objects being journaled, the journal (*JRN) itself, and the associated receivers (*JRNRCV) must be in an ASP group. The function of journaling requires access to all three sets of objects at all times.

Figure 4-6 represents the concept of keeping your journaling objects in an ASP group.

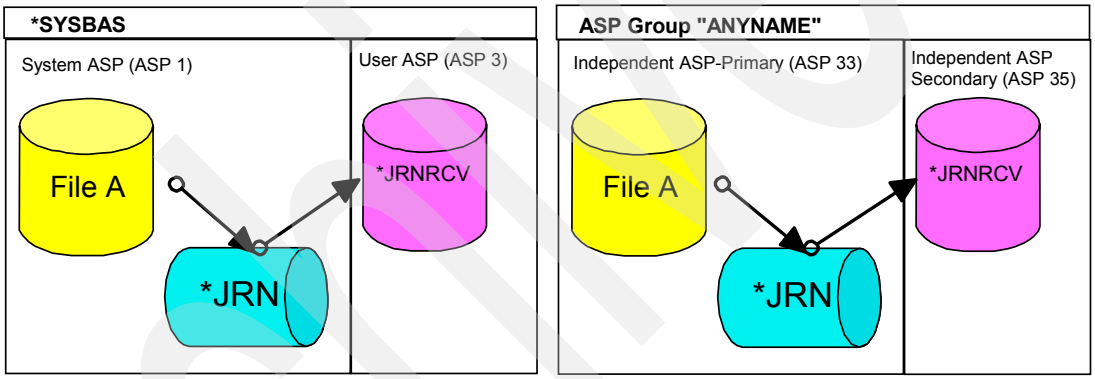


*Figure 4-6 Journaling boundaries*

You can group independent disk pools into *disk pool groups*. Disk pool groups consist of one primary disk pool and one or more secondary disk pools. If you are going to journal an object in a disk pool group, the object and the journal must be in the same disk pool. The journal receiver can be in a different disk pool, but must be in the same disk pool group as the journal and journaled object. Use the following rules when journaling objects on independent disk pools:

- The disk pool must be available on the system on which you are working.
- In a disk pool group, the journaled object and the journal must be in the same disk pool.
- In a disk pool group, the journal receiver can be in a different disk pool, but must be in the same disk pool group.

## 4.12 Subsystem considerations

Subsystem descriptions must be in *SYSBAS. The CTRSBSD and CHGSBSD commands have a new parameter, ASPGRP. You can specify the name of the IASP that the subsystem will use to look for its objects (AJs, PJs, routing entries).

An autostart job entry may be added to a subsystem. When the subsystem starts, the autostart jobs will be initiated by the subsystem. Autostart jobs are merely links that the

subsystem uses to locate a job description. The JOBD specifies information used by a subsystem to locate a routing entry for the job. If the program used by an autostart job is in an IASP, the job description must also specify the RQSDTA (request data) parameter to call the program. Since the program is in an IASP, it cannot be specified as the program on the routing entry.

## 4.13 Job queues

Job queue objects can be created in an IASP. This allows applications to run in the IASP with fewer changes. The JOBQ operates the same as a JOBQ in *SYSBAS. Jobs in the JOBQ can be managed and the JOBQ itself can be managed in the IASP. Behavioral differences do exist. Jobs on a JOBQ in an IASP will not persist across vary off/ vary on, nor will jobs on the JOBQ be available after a switch or failover.

## 4.14 DRDA considerations

There are certain DRDA-related objects that cannot be contained in user databases. DDM user exit programs must reside in libraries in the system database, as must any ARD programs.

You should be aware that the process of varying on a user database causes the RDB directory to be unavailable for a period of time. This can cause attempts by a DRDA application requester or application server to use the directory that is to be delayed or to time out. The exposure to having directory operations time out due to unavailability caused by varying on a database is much greater if multiple databases are varied on at the same time. As noted below, the first time that a user database is varied on, an attempt is made by the server to add a directory entry for that database. If the directory is unavailable due to a concurrent vary-on operation, the addition will fail, in which case you must manually add the entry.

Another consideration for using user databases is in regard to configuration of entries in the RDB directory. One of the rules for naming user databases is that user RDB names cannot match the system name specified in the network attributes (as displayed by the display network attributes (DSPNETA) command).

Local user database entries in the RDB directory are added automatically the first time that the associated databases are varied on. They are created using the *IP protocol type and with the remote location designated as LOOPBACK. LOOPBACK indicates that the database is on the same server as the directory.

## 4.15 Commitment control considerations

Independent disk pools and independent disk pool groups can each have a separate OS/400® SQL database. Commitment control can be used with these databases. However, since each independent disk pool or independent disk pool group has a separate SQL database, we make the recommendations offered in the following sections.

### 4.15.1 Commitment definitions

When commitment control is started, the commitment definition is created in the QRECOVERY library. Each independent disk pool or independent disk pool group has its own version of a QRECOVERY library. On an independent disk pool, the name of the QRECOVERY library is QRCY*xxxxx*, where *xxxxx* is the number of the independent disk pool. For example, the name of the QRECOVERY library for independent disk pool 39 is QRCY00039. Furthermore, if the independent disk pool is part of a disk pool group, only the primary disk pool has a QRCYxxxxx library. When you start commitment control, the commitment definition is created in the QRECOVERY library of the independent disk pool that is associated with that job. This makes commitment control active on the independent disk pool.

Using the Set ASP Group (SETASPGRP) command while commitment control is active on an independent disk pool has the following effects:

- If you change from one disk pool to another and resources are registered with commitment control on the disk pool, the SETASPGRP command fails with message CPDB8EC, reason code 2: `The thread has an uncommitted transaction.` This message if followed by message CPFB8E9. If you change disk pools and no resources are registered with commitment control, the commitment definitions are moved to the independent disk pool to which you are switching. If you change from the system disk pool (ASP group *NONE), commitment control is not affected. The commitment definitions stay on the system disk pool.
- If you use a notify object, the notify object must reside on the same independent disk pool or independent disk pool group as the commitment definition. If you move the commitment definition to another independent disk pool or independent disk pool group, the notify object must also reside on that other independent disk pool or independent disk pool group. The notify object on the other independent disk pool or independent disk pool group is updated if the commitment definition ends abnormally. If the notify object is not found on the other independent disk pool or independent disk pool group, the update fails with message CPF8358.

  **Note:** One exception to the notify object rule is with the QSYSOPR message queue. It is possible to use this queue in *SYSBAS even when the commit definition is in the independent ASP.

Recovery of commitment definitions residing on an independent disk pool is performed during independent disk pool vary-on processing. This is similar to IPL recovery. Commitment definitions in an independent disk pool are not recovered during the system IPL. The vary-off of an independent disk pool has the following effects on commitment definitions:

- Jobs associated with the independent disk pool end.
- No new commitment definitions are allowed to be created on the independent disk pool.
- Commitment definitions residing on the independent disk pool become unusable.
- If the same definitions are not attached to a job, the release transaction that is scoped locks.

You cannot use an LU6.2 SNA connection (protected conversations or DUW) to connect to a remote database from an independent disk pool database. You can use unprotected SNA conversations to connect from an independent disk pool database to a remote database.

When commitment control is active for a job or thread, access to data outside the independent disk pool or disk pool group to which the commitment definition belongs is only possible remotely, as though it were data that resides on another system. When you issue an SQL CONNECT statement to connect to the RDB on the independent disk pool, the system makes the connection a remote connection. The system disk pool and basic disk pools do not require a remote connection for read-only access to data that resides on an independent disk pool. Likewise, an independent disk pool does not require a remote connection for read-only access to data that resides on the system disk pool or a basic disk pool.

### 4.15.2 Considerations for XA transactions

In the XA environment, each database is considered a separate resource manager. When a transaction manager wants to access two databases under the same transaction, it must use the XA protocols to perform a two-phase commit with the two resource managers. Since each independent disk pool is a separate SQL database, in the XA environment each independent disk pool is also considered a separate resource manager. For an application server to perform a transaction that targets two different independent disk pools, the transaction manager must also use a two-phase commit protocol.

## 4.16 Exit programs

Exit programs are a special consideration when dealing with independent disk pools. If your independent disk pools contain primarily data, the concern is not as great. However, you may have multiple independent disk pools on your system, containing application programs. In this case, it may be worthwhile to consider an application library in *SYSBAS where commonly used application programs and exit programs can reside. This makes those programs accessible even when one or more of the independent disk pools are unavailable.

## 4.17 System libraries

Avoid or eliminate the use of system libraries, such as QGPL, QSYS, QSYS2, and essentially any library beginning with the letter Q, whenever possible. Applications that are being written or moved to independent disk pools should *not* use system libraries. Those libraries remain in the system ASP.

Consider a user-created library in the IASP first when eliminating your use of system libraries. If this does not work for you, then consider using a separate user library in the system ASP.

## 4.18 System ASP and all basic user ASPs (*SYSBAS)

The ASP group of applications that are being written or moved to independent disk pools is designated only as permanent storage. The exception is if the application is written or designed to currently work across multiple systems or partitions.

Although permanent storage for the system ASPs is available to the applications running from an independent disk pool or ASP group, avoid using this storage whenever possible. You must view *SYSBAS (system ASP and all basic user ASPs) as the domain and storage for operating system code, LPP code, objects that are not supported in independent disk pools, and temporary working storage for applications.

## 4.19 Other system considerations

Consider the additional items in the following sections when planning application migration to independent ASPs.

### 4.19.1 System-wide is no longer system-wide

Designers and developers must be aware that jobs that used to accomplish work on a system-wide basis now only accomplish such work on a name-space basis or database basis. Or the jobs may accomplish this work across the set of available (varied on) independent disk pools and *SYSBAS. This is expressed in the parameters used for several keywords. The namespace-wide basis is referred to as *ALL for all available independent disk pools. *SYSBAS is referred to as *ALLAVL.

### 4.19.2 May need ASP group, library, or object to identify an object

The object name and library name may not be enough to uniquely identify an object any longer. Suppose that you have a job that provides services for other jobs. Also consider that your server takes work requests for several different independent disk pools. Someone requests a task with LIBNAME/OBJNAME in one independent disk pool and the next user of this job's services asks for a task with LIBNAME/OBJNAME in a different independent disk pool. These are different objects if the library is in the independent disk pools. However, they are the same object if the library is in *SYSBAS. To the application, this means that:

- You need the independent disk pool identification in addition to the library name and object name.
- Various places where data may be cached must be redone to include the independent disk pool as part of the object designation. Keep in mind that the independent disk pool may go offline at any time.

### 4.19.3 Locking

Locking no longer prevents an object from changing or being deleted in some cases. When you lock an object that is in an independent disk pool group and then change the namespace for the job or thread to work with *SYSBAS or a different independent disk pool group, you still have a lock on the object.

Note that the deallocate object (DLCOBJ) running in the new namespace cannot deallocate something that it cannot find. It cannot find anything outside of the current namespace, so you cannot get rid of a lock unless you are running in the correct namespace.

Consider this problem: While you are not running in the namespace for the locked object, no jobs may be running in that namespace. Therefore, there is no reason for the system to prevent the independent disk pool from being taken offline.

### 4.19.4 Unable to change a namespace

There are times when you cannot change a namespace. SETASPGRP rejects the request if there is an uncommitted transaction or if an operation in progress blocked the namespace change. You cannot change the namespace when:

- Running within PDM
- Running in a program invoked as a result of pressing the attention key (SETATNPGM)
- Running in a pre-attention exit program
- Running in a pre-system-request exit program
- Running in a message queue's break handling program
- Any user exits from certain system jobs

The other side of this is true. That is, it is a really bad idea to produce an open, uncommitted transaction from a user exit program. You can prevent whatever is running (in the environment from which this user exit was invoked) from a namespace change.

### 4.19.5 Workflow design and control for use of independent disk pools

Consider doing workflow design and controlling which jobs use the independent disk pool at certain times. Having an independent disk pool in the namespace of a job or thread prevents the independent disk pool from being made unavailable. Therefore, some servers should avoid waiting for new work while the independent disk pool is in the namespace. This is particularly true for servers that span independent disk pool boundaries. You do not want to optimize out namespace changes by waiting to see if a switch is needed. In doing so, you interfere with make unavailable, administrative switchback, and other scenarios. Moreover, if your server works for different users, you must swap users and then try the SETASPGRP command so that you get proper security checking. The new user is required to have *USE authority to the device descriptions for each ASP in the group.

Prestart jobs may need to avoid waiting for work in a state that holds the independent disk pool active. On the other hand, you can avoid this requirement by shutting down the subsystem before you vary off the independent disk pool.

When a subsystem takes work from a JOBQ and the jobs have an associated ASP group, you have two considerations:

- For one-at-a-time batch processing, no jobs may be using the independent disk pool for a short time between jobs. This allows the independent disk pools to be made unavailable before you really wanted that done. The *jobs using an IASP* interfaces do not tell you about jobs that want to use the independent disk pool in the immediate future.
- When the independent disk pool is taken down, job initiation fails for jobs that were submitted to run with that independent disk pool. An amazing number of jobs can be removed from a JOBQ and killed before you realize that anything is wrong.

### 4.19.6 System values

There are three system values that are not resolved during the IPL process. These system values are resolved when they are used and may reference items that are in an IASP:

- QALWUSRDMN can reference libraries in an IASP.
- QBOOKPATH can reference directories in an IASP.
- QLOCALE can reference directories in an IASP.

### 4.19.7 Restoring IBM Licensed Programs Products

You are not able to restore license programs if you have an IASP in your namespace.

## 4.20 Creating an image catalog to be used with an IASP

The advantage to using virtual tape (image catalog) is that the security is saved with the objects, just like a save to tape. When moving objects from *SYSBAS to an IASP there is no move capability. Save/restore (SAVRST) to physical tape media or virtual tape is the only mechanism available. Save files do not save object authorities, creating authority issues when restored on the target IASP.

Some items to consider prior to using an image catalog are:

- The image size (IMGSIZ) parameter must be set to the maximum size to which you think the save will grow. If you leave it with the default of *IMGCLGTYP, it will default to 1 MB, then automatically create a second volume. You could end up with multiple volumes with which to work.
- Appendix H, “Using virtual tape on IBM i” on page 221, describes how you can overallocate, but not exceed, the system available disk. When actually performing a write to the virtual tape the system will not exceed 95% of available storage.
- There is no compression using an image catalog. Disk must be considered.

See also Appendix C in *i5/OS V5R4 Virtual Tape: A Guide to Planning and Implementation*, SG24-7164, for more information.

5

# The IASP project: installing/converting your application

This chapter discusses the majority of items that you must consider when implementing an independent auxiliary storage pool (IASP) or converting your application. Most changes to support IASPs are typically work management related. In general, IASP migration can be transparent to most users. Our intent is to provide you with a framework and guide to your project, pointing out as many of the considerations of which we are currently aware.

## 5.1 Installing or converting ISV applications

Check with your application provider. Many ISV applications are already IASP compatible.

## 5.2 Typical IASP migration project outline

A typical IASP migration project should contain the following steps. Consider creating a second IASP that can be used for initial proof of concept and migration testing. Supporting multiple versions of a product is another strategy for IASPs. To test a new version of a product, install the product, programs, and data on an independent disk pool:

1. Determine how IASPs fit into your business strategy.
2. Perform initial proof of concept on a subset of applications.
3. Educate administrative and programming staff on IASP enablement considerations, such as:
   - Connectivity to databases
   - Namespace
   - IASP boundaries for databases
   - DB2 triggers and commitment control cycles.
   - DDM (Use *RDB.)
   - Journaling: ASP group boundaries
4. Set up a test environment in which application and work management changes for IASP enablement will be performed and tested.
5. Devise a plan for synchronizing *sysbas with primary and secondary systems, if applicable.
6. Test all process, application, and work management changes.
7. Determine production migration strategy based on available hardware options.
8. Test the migration process in a sandbox environment, if possible. Provide any necessary education for support personnel and users on the process change, and so on.
9. Execute the production migration and implement the revised HA strategy based on IASPs, if applicable.

## 5.3 IASP enablement considerations

When enabling IASPs, several factors should be carefully considered. At the application level you must understand where the objects are, who the users are, and how they access programs and data.

1. The location of application objects is important. An object cannot reside in the same library name in *SYSBAS and the ASP group. Understand which object types must exist in the same ASP group.
2. Consider how you will set up and load the IASPs. If you are converting from a user ASP or *SYSBAS consider the impact that your save/restore strategy will have on authorities. Document your naming strategy for disk pools and relational databases.
   a. Before creating your disk pool, make sure that the relational database name does not exist in your current environment.

b. Make sure that a directory in the /root does not exist with the name of any disk pools that you will create. A directory is created for every disk pool, *primary and *secondary.

3. Work management decisions will affect the transparency to the user. Decide the best method for managing proper workflow and consider the namespace, library list, and job initiation.
4. Take into consideration how the application connects to the database. If you have DDM files consider using the RMTLOCNAM parameter *RDB and the relational database name. JDBC and ODBC connections should also use the appropriate RDB name.
5. If the application uses commitment control or join logical files, or both, make note. There are special considerations for these.
6. Additional considerations for multiple IASPs on a system include changing the namespace, work management, and database connectivity.
7. Consider editing SMAPP access recovery path time values for the IASPs.

## 5.4 IASP setup considerations

Consider the following:

- Decide whether program objects will be in *SYSBAS or in the IASP. Consider the following guidelines for object placement. Refer to Appendix B, “Application object planning checklist” on page 189.
  - *SYSBAS
    - Objects not supported in IASP (See Appendix D, “Supported and unsupported objects” on page 197.)
    - Exit programs, and so on, that must be found in *SYSBAS
    - Operational version of job control objects (SBSD, JOBQ, JOBD, and so on)
    - Temporary objects
    - System libraries
    - Application program objects
  - ASP Group
    - Permanent application data objects
    - Spool files
    - Journals and journal receivers
    - Definitional version of job control objects (SBSD, JOBQ, JOBD, and so on)
    - Application program objects
- Evaluate IFS directories and files to move to the IASP. Make sure that an IFS directory with the name of your disk pool does not exist before you start. Making a disk pool available creates an IFS directory called `/diskpoolname`. You can migrate your IFS to a disk pool. See Appendix F, “Migrating Integrated File Systems (IFS)” on page 209. If necessary, you can create a symbolic link in the old *SYSBAS IFS location pointing to the new IASP location.
- Populate IASP with data.
  - Consider using virtual tape to ensure proper object authority retention. See Appendix H, “Using virtual tape on IBM i” on page 221.
  - Use the ASPDEV parameter since the RSTLIB default will put objects in *SYSBAS.

  - Permanent SQL objects cannot span IASP boundaries.
  - You cannot create files, tables, and so on, in QSYS2nnnnn.
  - Journal objects must be in the same ASP group as objects being journaled.
  - You may want to create new companion libraries for library content to be split between *SYSBAS and IASP.
- Modify your startup program to vary on the IASP during IPL. Use the VRYCFG command and vary on the *DEVD for your IASP.
- Understand your SMAPP setting. It can greatly reduce the amount of time that it takes to vary on an IASP after an abnormal shut down. However, a low SMAPP setting can have an effect on system performance. An IASP used for archive probably does not require as low a target access recovery path time as a critical production IASP.
- The user ID or group ID (UID/GID) should be the same across multiple systems if using synchronization. Objects owned by IBM profiles, such as QPGMR, can cause challenges for LPAR-to-LPAR synchronization. APIs may be required to change the UID/GID on target LPARS. Refer to the IBM i InfoCenter for more information.

## 5.5 Work management considerations

Consider the following:

- Determine how applications will access an IASP (namespace). The most user-transparent method will be using the INLASPGRP parameter on the job description used by the user profile. Changing the QDFTJOBD is *not* recommended. Instead, make a copy of the QDFTJOBD job description and change the new JOBD to include the INLASPGRP and the appropriate library list. Change the user profile to use the new JOBD.
- Understand where the SETASPGRP command can and cannot be used. See Appendix E, “SETASPGRP command” on page 203.
- When using statements, ensure that you are using the appropriate relational database name.
- Verify your Library List processing and double check your system values in QUSRLIBL and QSYSLIBL. If you move a library from *SYSBAS to an IASP that is in QUSRLIBL or QSYSLIBL you must adjust the application JOBD accordingly. If you are using 36 environment (QS36F) this presents additional library list challenges in an IASP environment.
- Evaluate SBSDs for changes or attachment to IASP. Analyze all components and know where they all reside:
  - Autostart job entries (AJE)
  - Prestart job entries (PJE)
  - Routing entries *CLS object type
  - JOBQ entries
  - Communication entries (JOBD)
  - Remote location name entries (JOBD)

## 5.6 Database considerations

Consider the following:

- Views and tables cannot span IASP boundaries. If you use join logicals over physicals they must reside in the same IASP.
- Commit block cannot span IASP boundaries. If connected to an IASP RDB, you cannot commit changes against both the IASP and the *SYBAS (except QTEMP).
- Decide on the RDB name for the IASP.
- DDM files should be configured to use *RBD for connectivity.
- When creating new permanent libraries for an application, use the *ASPDEV parameter on the CRTLIB command. If you create a collection, the default for INLASP is the current namespace.
- When connecting to a database using JDBC, ODBC, or FTP connectivity, use the *JOBD in the *USRPRF where possible.

## 5.7 Moving applications from *SYSBAS to an IASP

We recommend deleting all objects used and created by the application (especially the unqualified objects that may go in surprising places) before you run an application in *SYSBAS and later in an independent disk pool (or visa versa). Failure to do this may cause the independent disk pool to be ineligible to be varied on due to duplicate libraries and objects.

6

# Managing an IASP

This chapter illustrates additional management functions for independent disk pools. It looks at the various functions that are available from the IBM Systems Director Navigator for i5/OS and some command functions. We discuss creating a new disk pool in Chapter 2, “Creating an IASP” on page 9.

# 6.1 Independent disk pool management

Independent disk pools must be managed, as do any other objects on the system. Systems with only a system ASP do not have to worry about such tasks. Remember that with independent auxiliary storage pools (IASPs) you are dealing with multiple databases.

## 6.1.1 Disk pool (IASP) operation

A system can bring an IASP online and make it active at any time during IPL after the QSYSWRK subsystem is active. The system can bring an IASP online during IPL by using the vary configuration (VRYCFG) command in the startup program. After the IASP is online and active, the objects within the IASP are accessible and usable.

**Important:** An IASP can go offline and yet the rest of the system remains functional. However, after an IASP is offline, the objects that it contains are no longer visible to, accessible from, or usable by the system where it was previously online and active.

A system in which an IASP is online can vary that IASP offline (make unavailable). The system can do this explicitly with a user request or it can do this implicitly through system termination.

## 6.1.2 Creating an independent disk pool

Before you create a disk pool, be sure to complete the steps in Appendix A, “Prerequisite steps” on page 187. Then refer to Chapter 2, “Creating an IASP” on page 9, to learn how to create a disk pool.

## 6.1.3 Disk pool and disk pool group

To access disk pool and disk pool group information, complete the following steps using IBM System Director Navigator for i5/OS:

1. From the System Director Task List, expand **i5/OS Management**.
2. Select **Configuration and Service**.

3. Select the **Disk Pools** link, as shown in Figure 6-1.

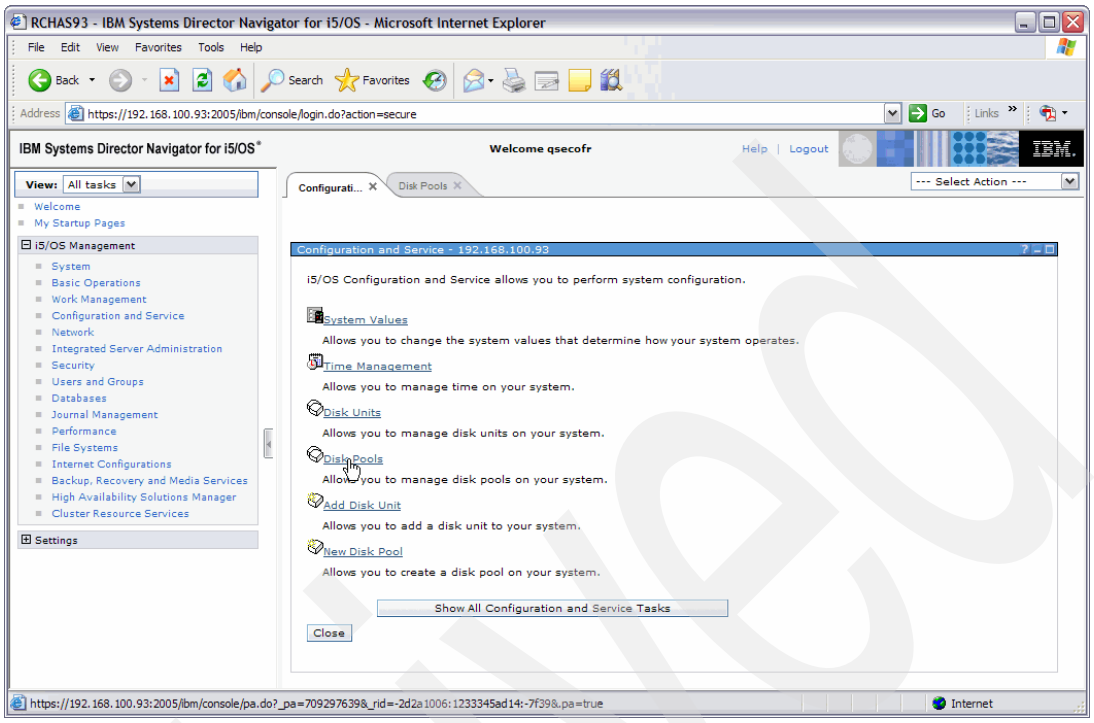


*Figure 6-1 Configuration and services*

## 6.1.4 Making an independent disk pool unavailable

You can select an independent disk pool to make it unavailable (varied off) in the same way that it was made available (varied on). You cannot access any of the disk units in the independent disk pool until it is made available (varied on) again. The pool can be made available again on the same system or another system in the recovery domain of the cluster resource group (CRG) after a switch occurs.

**Important:** Make sure that no active jobs are using the disk pool prior to making the pool unavailable.

To make an independent disk pool unavailable, follow these steps:

1. Open the disk pools list as described in 6.1.3, “Disk pool and disk pool group” on page 84.
2. Click the double arrow icon next to the disk pool name and number on the entry in the disk pools list. This opens a pop-up menu.

3. Select the **Make Unavailable** option, as shown in Figure 6-2.

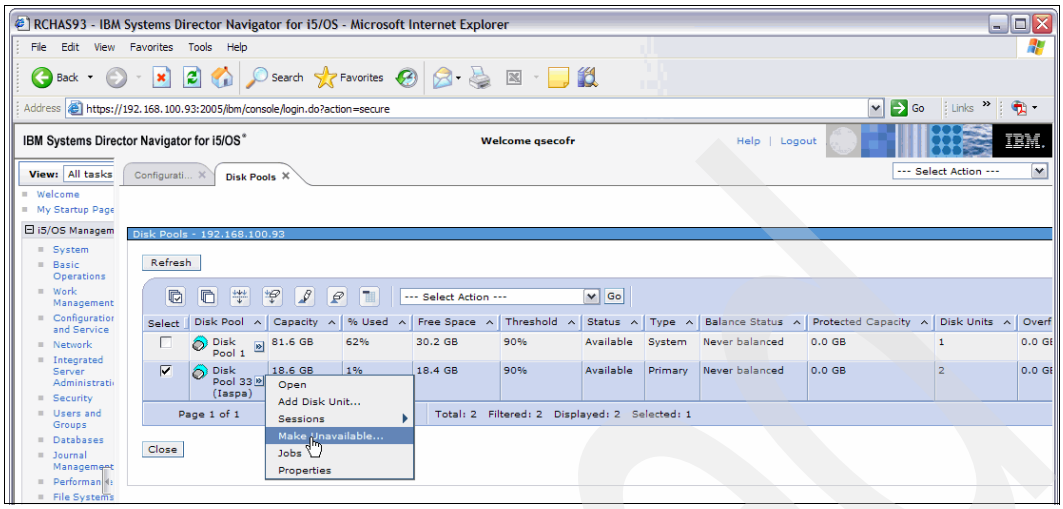


*Figure 6-2 Make Unavailable option*

4. On the panel that follows, confirm the selection to make the disk pool unavailable, as shown in Figure 6-3.

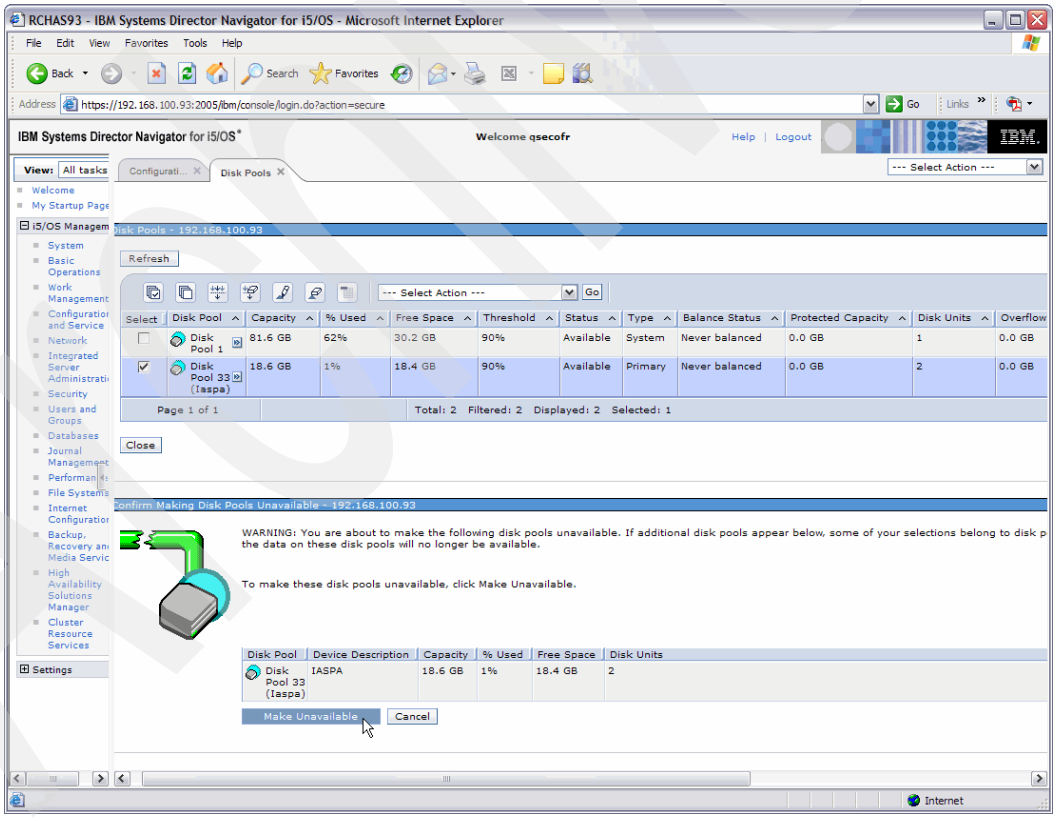


*Figure 6-3 Make Unavailable confirmation*

5. You must select **Refresh** on the Making Disk Pools Unavailable panel to check the progress of the make unavailable function, as indicated in Figure 6-4. This panel does not refresh automatically, so refresh it until the make unavailable is complete. Each time that you refresh, the status and elapsed time will be updated. If an error occurs, you will not see the error condition until you refresh. If you receive an error message that is in its own panel with a Continue button, the **Continue** button must be clicked for the process to continue.

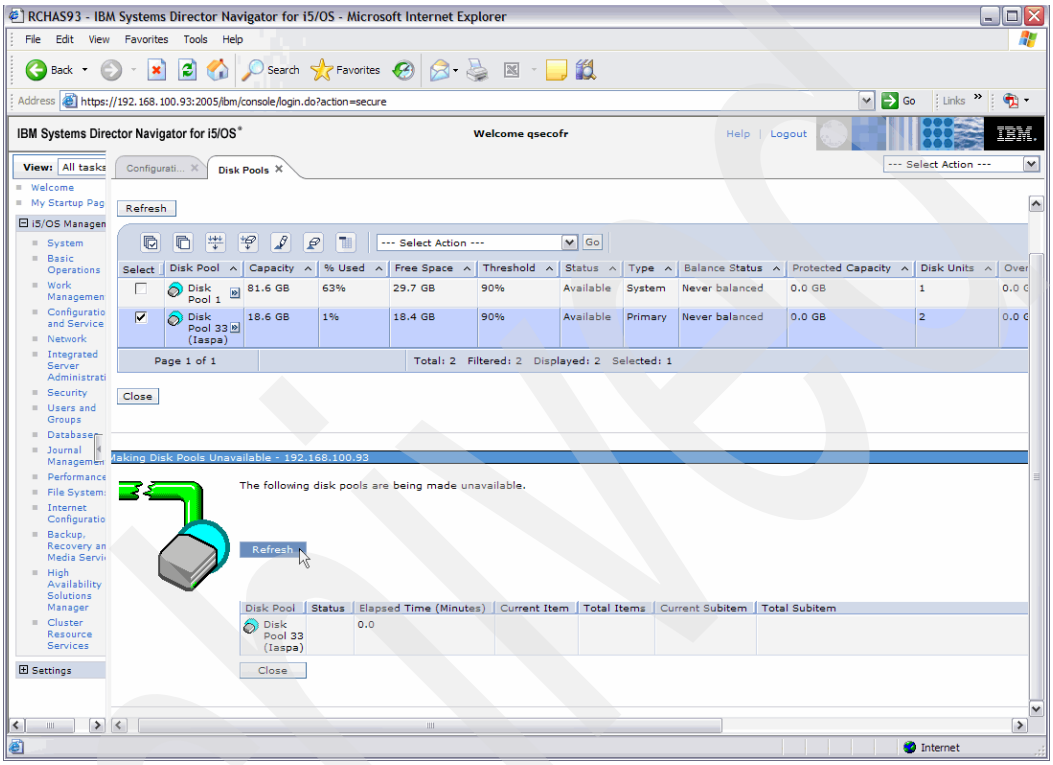


*Figure 6-4 Make Unavailable Refresh*

6. When the make unavailable function has successfully completed, the refresh will show a status of complete, as shown in Figure 6-5.
7. Select **Close** on the Making Disk Pools Unavailable panel.

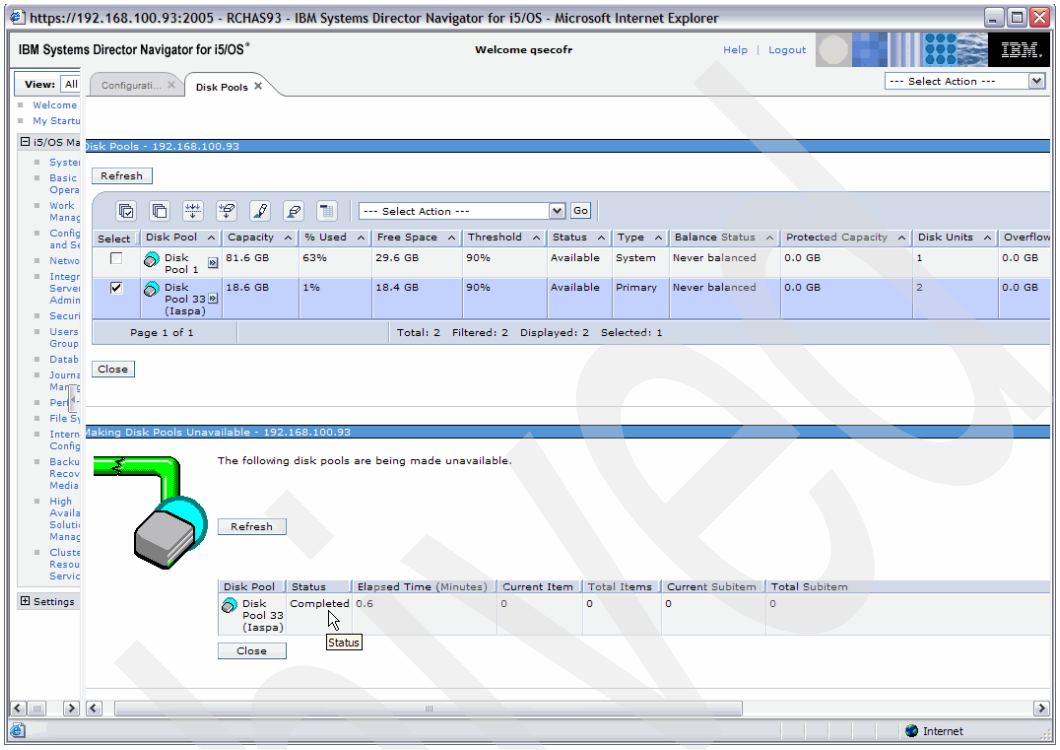


*Figure 6-5 Making Disk Pools Unavailable complete status*

8. Click **Refresh** on the Disk Pools panel to show the current status of the disk pools. The disk pool specified should now show a status of unavailable, as shown in Figure 6-6.

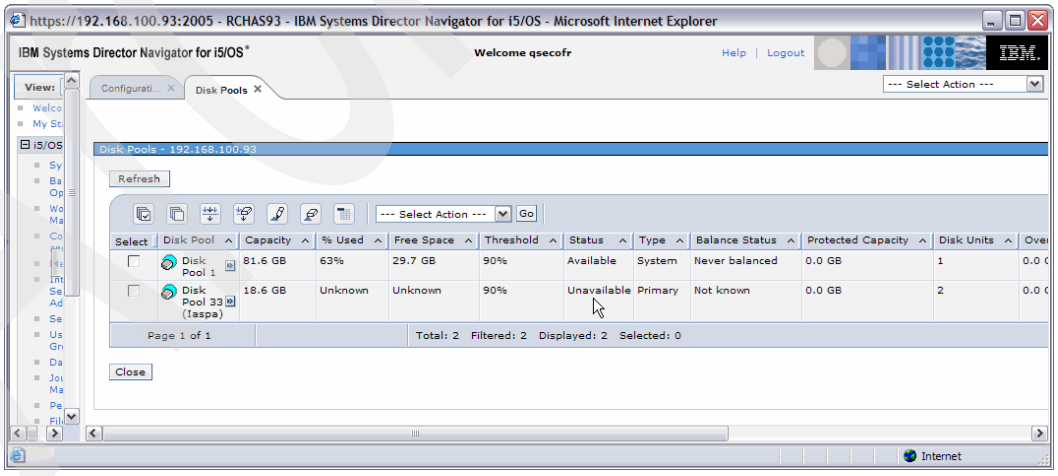


*Figure 6-6 Disk Pool Unavailable*

## 6.1.5 Deactivating a disk pool using the VRYCFG command

The command VRYCFG can also be used to disable an IASP or disk pool. The VRYCFG command can be included in an application process if needed. The VRYCFG command

would appear as in Figure 6-7. Note that CFGOBJ is the name of the disk pool as specified when it was created and CFGTYPE is *DEV.

```
Type command, press Enter.
===> VRYCFG CFGOBJ(IASPA) CFGTYPE(*DEV) STATUS(*OFF)

     Vary off completed for device IASPA.
```

*Figure 6-7 Disk pool VRYCFG *OFF command*

### 6.1.6 IASP save/restore

Refer to Chapter 7, “Backup and recovery” on page 125, for a complete discussion about saving and restoring IASPs.

### 6.1.7 Clearing the data from an independent disk pool

You can clear the data from a disk pool from your system. When you clear a disk pool, you destroy all data on the disk units in the pool. An independent disk pool must be unavailable before it can be cleared. To clear a disk pool, follow these steps:

1. Open the disk pools list as described in 6.1.3, “Disk pool and disk pool group” on page 84.
2. Ensure that the disk pool is unavailable. If it is available, see 6.1.4, “Making an independent disk pool unavailable” on page 85, to make it unavailable.
3. Click the double arrow icon next to the disk pool name and number on the entry in the disk pools list. This opens a pop-up menu.

4. Select the **Clear** option, as shown in Figure 6-8.

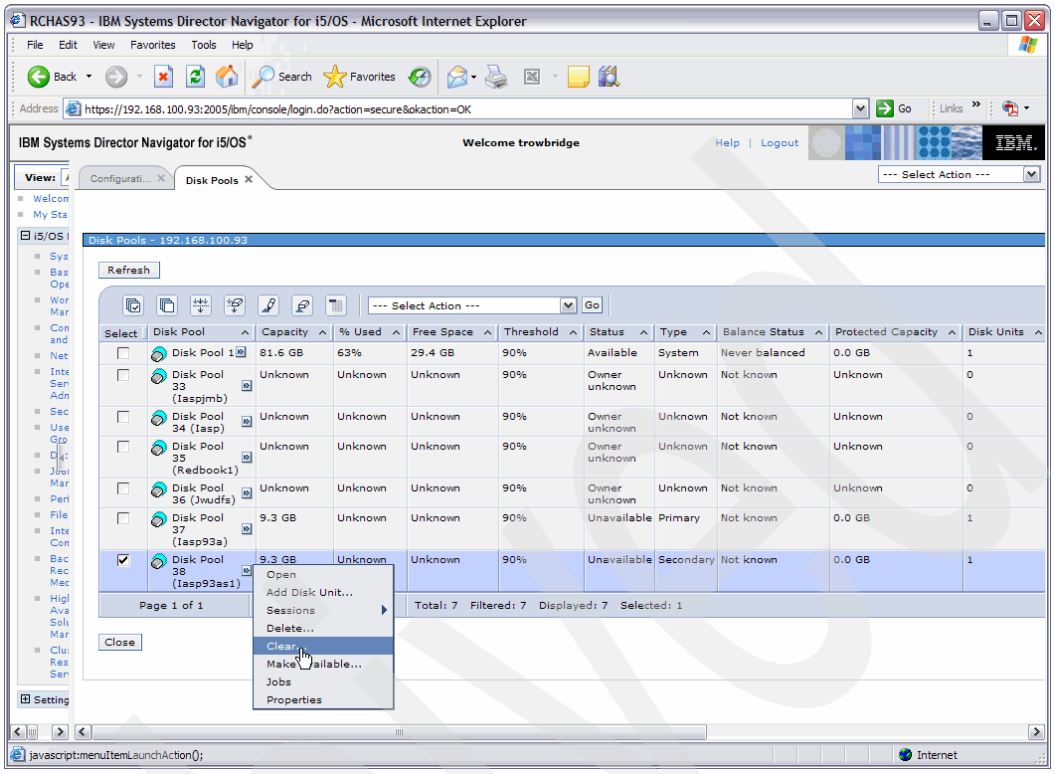


*Figure 6-8 Disk Pool list: Clear Disk Pool option*

**Note**: Multiple disk pools can be cleared on a single request by selecting the check box in the disk pool list for each disk pool to be cleared, then using the Select Action drop-down list to select the **Clear** action.

5. The Confirm Clearing Disk Pools panel is displayed listing the disk pools to be cleared. Select **Clear Disk Pools**, as shown in Figure 6-9.

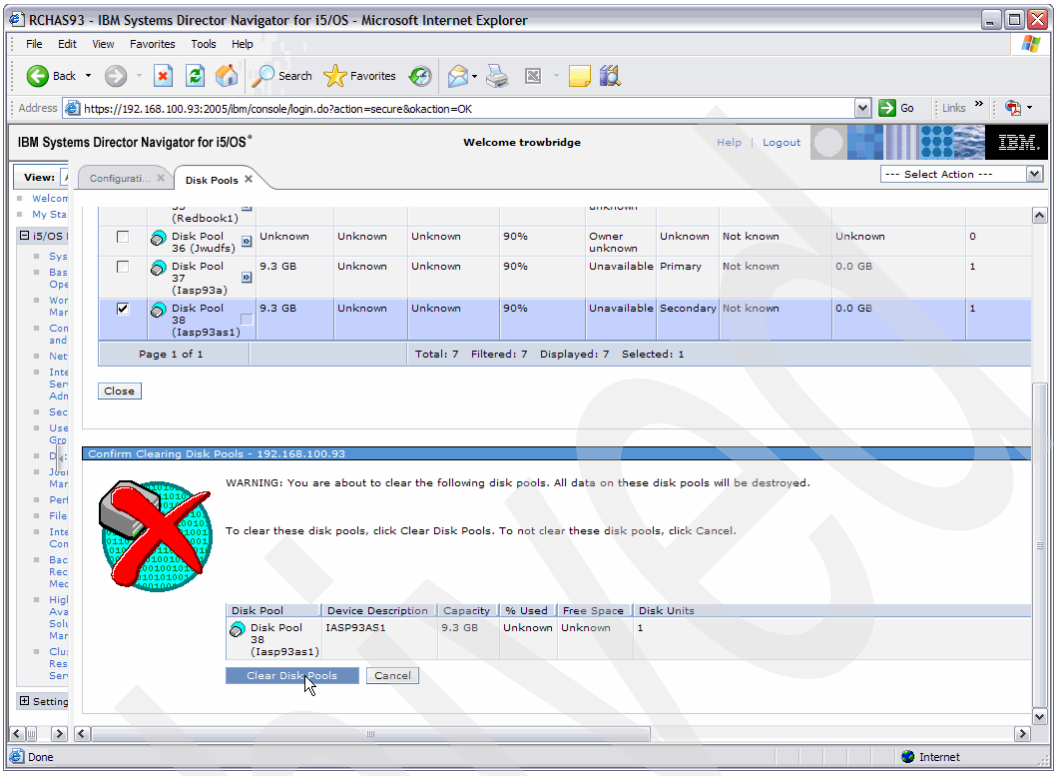


*Figure 6-9 Confirm clearing disks in disk pools*

6. A completion message is returned when the clear is completed, as shown in Figure 6-10.

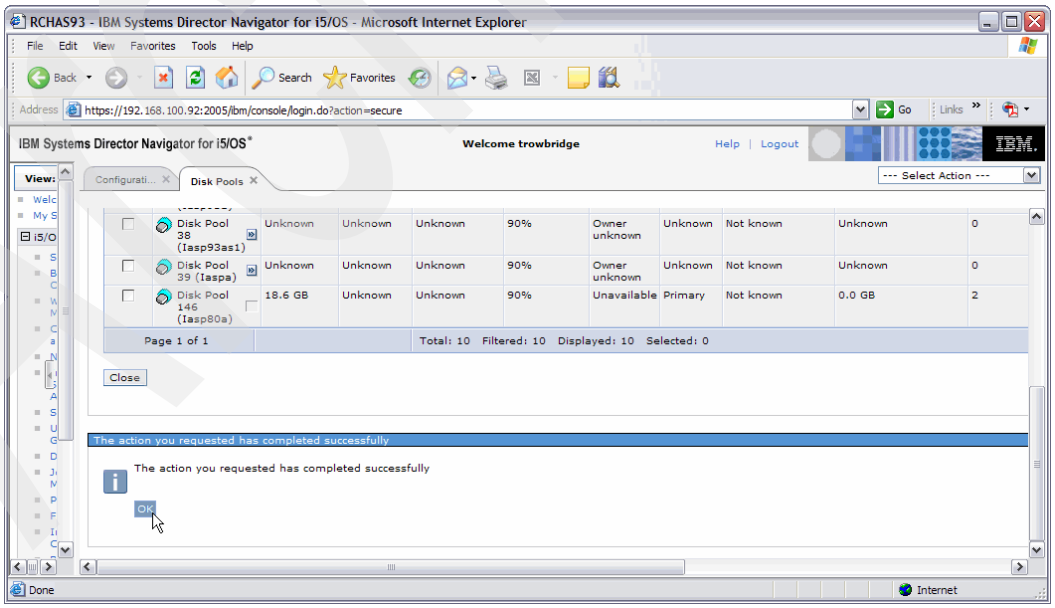


*Figure 6-10 Clear disk pool completion*

## 6.1.8 Recovering the disk pool group

If the primary disk pool for a secondary disk pool is deleted, or if the primary disk pool is not aware of the secondary disk pool, the secondary disk pool must be re-associated with a primary disk pool. You can recover the disk pool group using System Director Navigator.

To recover a disk pool group, follow these steps:

1. Open the disk pools list as described in 6.1.3, “Disk pool and disk pool group” on page 84.
2. Click the double arrow icon next to the disk pool name and number on the entry in the disk pools list. This opens a pop-up menu.
3. Select the **Recover disk pool group** option, as shown in Figure 6-11. This option will only be available on secondary disk pools that have an inconsistent link status.

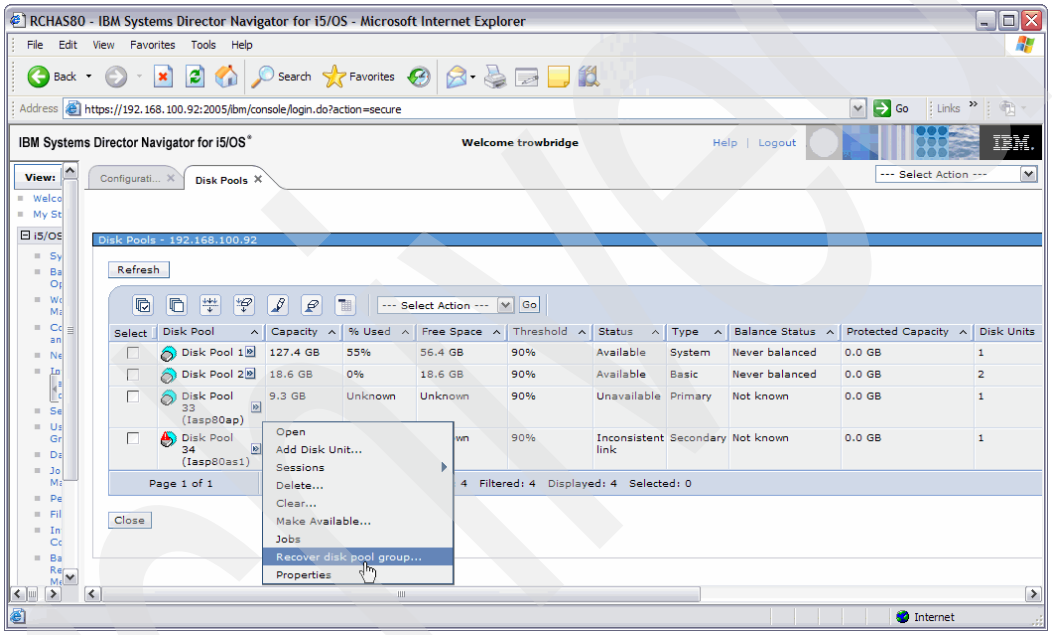


*Figure 6-11 Select to recover a secondary disk pool association*

4. On the Recover Disk Pool Group window, select the primary disk pool that you want to associate with the secondary disk pools and select the **GO** button, as shown in Figure 6-12 on page 93. Only the primary disk pools that are currently owned by the system are available for selection.

5. Select **Recover Group**, as shown in Figure 6-12. You cannot change the primary disk pool after you perform this action.

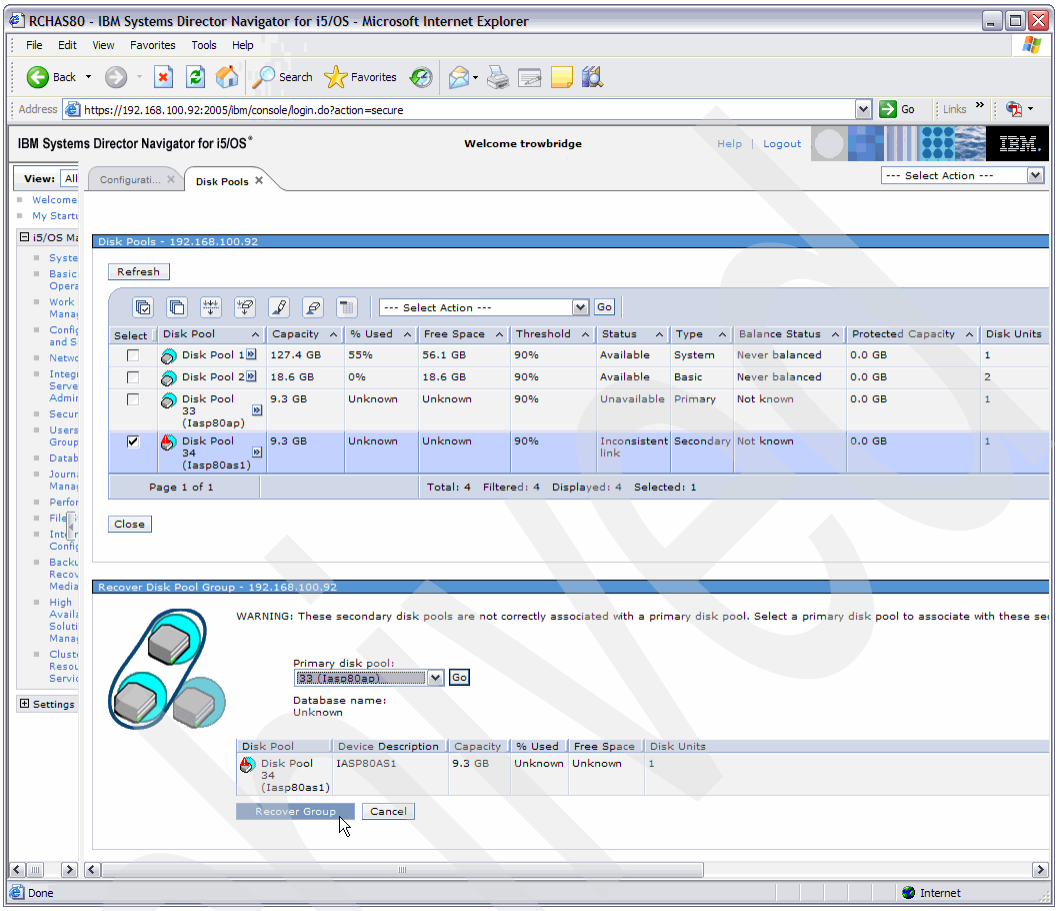


*Figure 6-12 Confirm recover disk pool group*

6. A completion message is returned when the recovery action has finished.
7. Refreshing the disk pool list shows that the status is now normal.

## 6.1.9 Balancing a disk pool

You can balance the data on a disk pool in your system. Balancing a disk pool improves system performance by balancing disk capacity across all the disk units in a disk pool.

There are two ways to balance a disk pool using System Director Navigator:

- Use the Add Disk Unit wizard when you add disk units to a pool.
- Use the Add Disk Unit wizard when you create a new disk pool.

See 6.1.16, “Adding a disk unit to an existing IASP” on page 106, for instructions for the Add the Disk Unit Wizard.

The balancing occurs immediately when the pool is created or when the disk are added, as shown in Figure 6-13 through Figure 6-15.

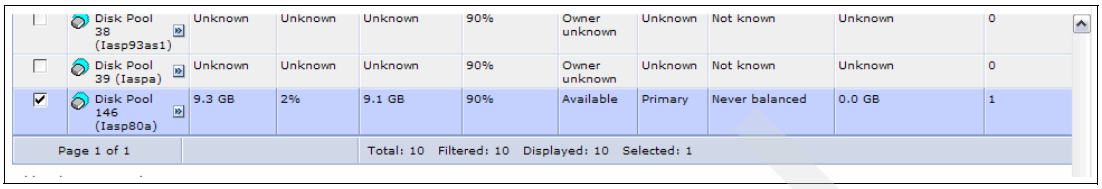


*Figure 6-13 Disk pool not balanced*

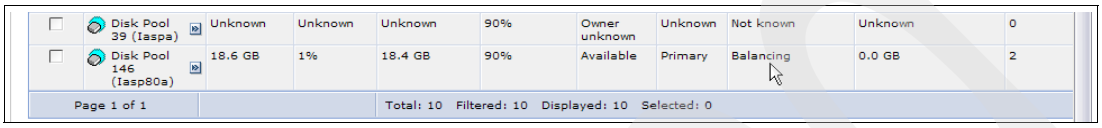


*Figure 6-14 Disk pool balancing*

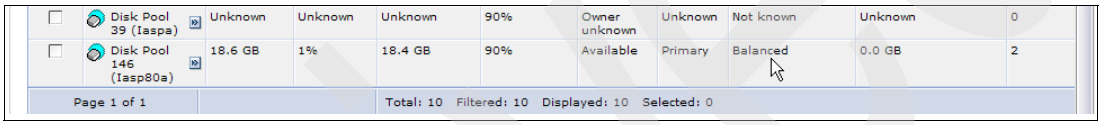


*Figure 6-15 Disk pool balanced*

## 6.1.10 Deleting a disk pool

You can delete a disk pool from your system. When you delete a disk pool, you remove all disk units from the pool, and the disk units are designated as unconfigured. All data on the disk units in a deleted disk pool is destroyed.

The disk pool must be in an unavailable status in order to be deleted.

To delete a disk pool, follow these steps:

1. Open the disk pools list as described in 6.1.3, “Disk pool and disk pool group” on page 84.
2. Click the double arrow icon next to the disk pool name and number on the entry in the disk pools list. This opens a pop-up menu.

3. Select the **Delete** option, as shown in Figure 6-16.

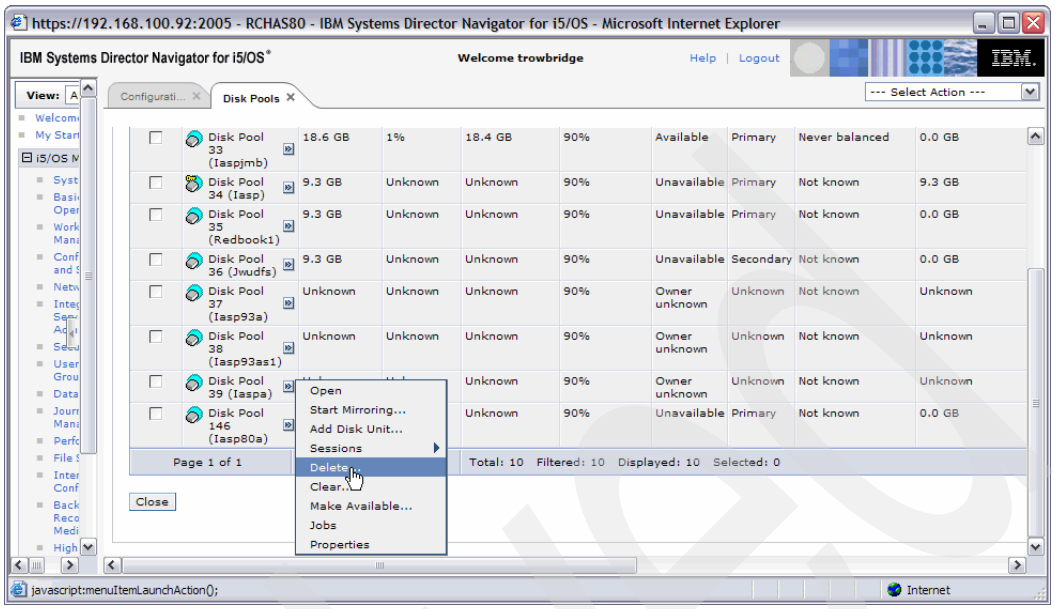


*Figure 6-16 Delete a disk pool pop-up menu selection*

4. The Confirm Deleting Disk Pools panel is displayed listing the disk pools to be deleted. Select **Delete Disk Pools**, as shown in Figure 6-17.

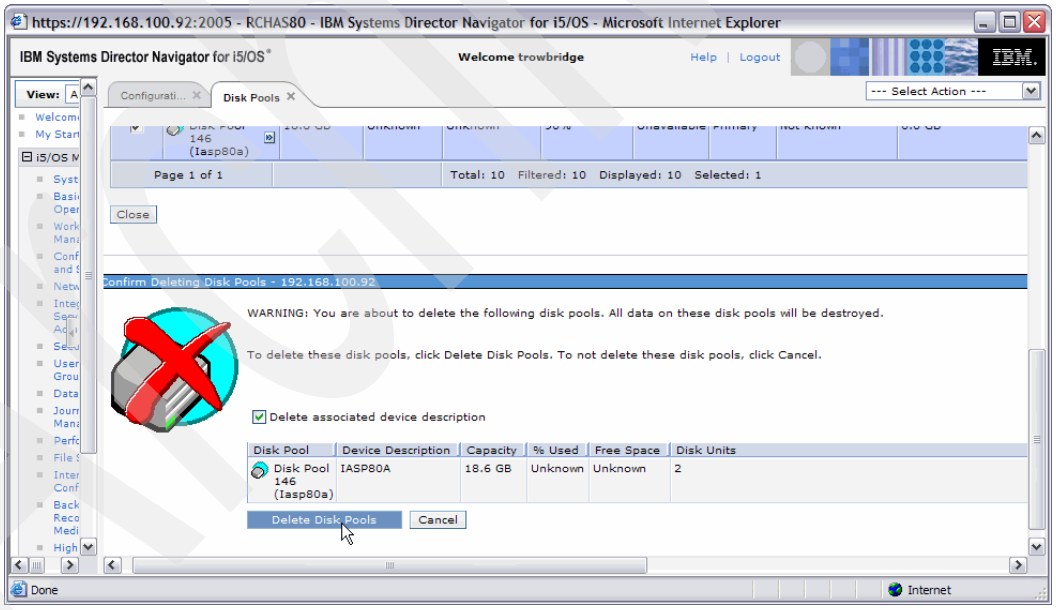


*Figure 6-17 Delete disk pools confirmation*

5. A completion message is returned when the delete is completed, as shown in Figure 6-18.

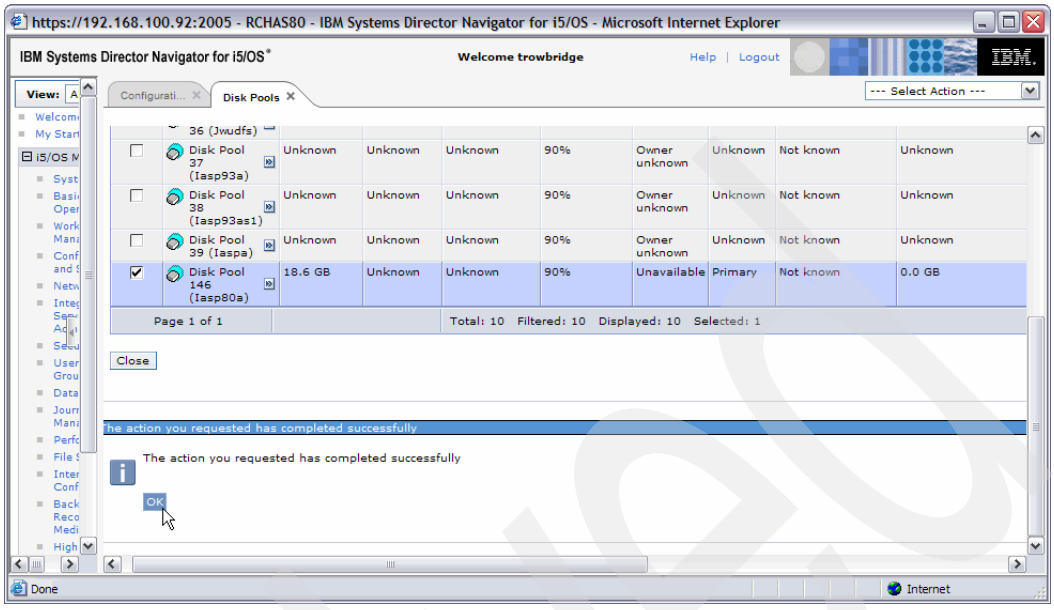


*Figure 6-18 Delete disk pool completion message*

**Note:** Multiple disk pools can be deleted on a single request by selecting the check box in the disk pool list for each disk pool to be deleted, then using the Select Action drop-down list to select the **Delete** action.

## 6.1.11 Converting a UDFS disk pool to a primary or secondary disk pool

You can convert UDFS disk pools to library-capable primary or secondary disk pools. Library-capable disk pools support library-based objects. You must convert UDFS disk pools if you want them to participate in a disk pool group. Before you create a secondary disk pool, you must create its primary disk pool.

**Important:** You cannot reverse this action.

To convert an UDFS disk pools to a library-capable primary or secondary disk pool, follow these steps:

1. Open the disk pools list as described in 6.1.3, “Disk pool and disk pool group” on page 84.
2. Check the check box for the entry in the disk pools list for the UDFS type disk pool that is to be converted.

3. Open the **Select Action** drop-down list. There are two convert options, as shown in Figure 6-19.

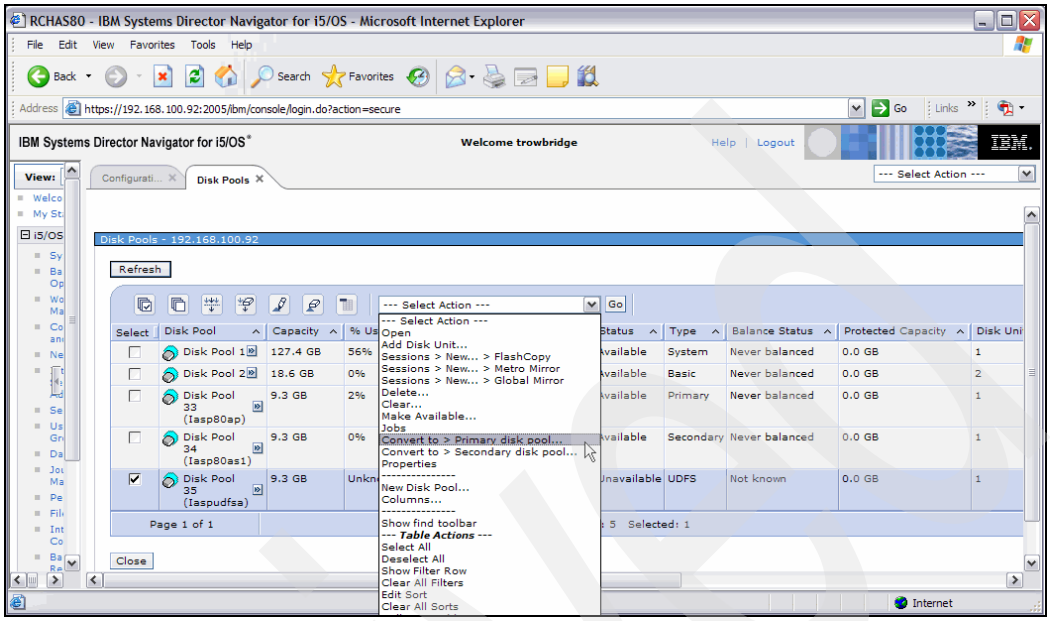


*Figure 6-19 Select Action drop-down: Convert to selection*

4. In the drop-down action list, if Convert to → Primary disk pool is selected, the Confirm Convert to Primary Disk Pool panel is displayed. This also allows the database name for the new primary disk pool to be specified, as shown in Figure 6-20. Select **Convert Disk Pool** to start the conversion.

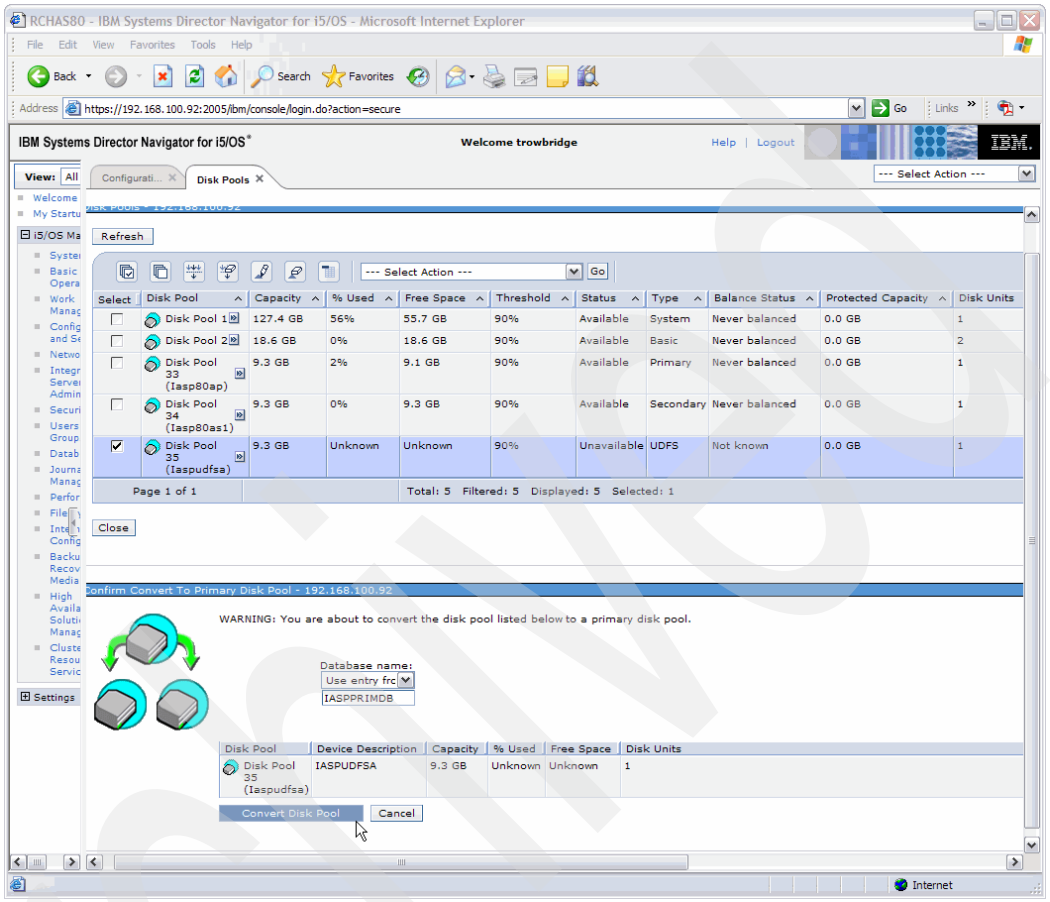


*Figure 6-20 Confirm conversion to primary disk pool*

5. In the drop-down action list, if Convert to → Secondary disk pool is selected, the Confirm Convert to Secondary Disk Pool panel is displayed. This allows the primary disk pool that this new secondary disk pool will be associated with to be selected. Only the primary disk pools that are currently owned by the system are available for selection. Select the primary disk pool from the drop-down list of Primary Disk Pools and select **GO** beside the list. Then select **Convert Disk Pool** to start the conversion, as shown in Figure 6-21. You cannot change the primary disk pool association after you perform this action.

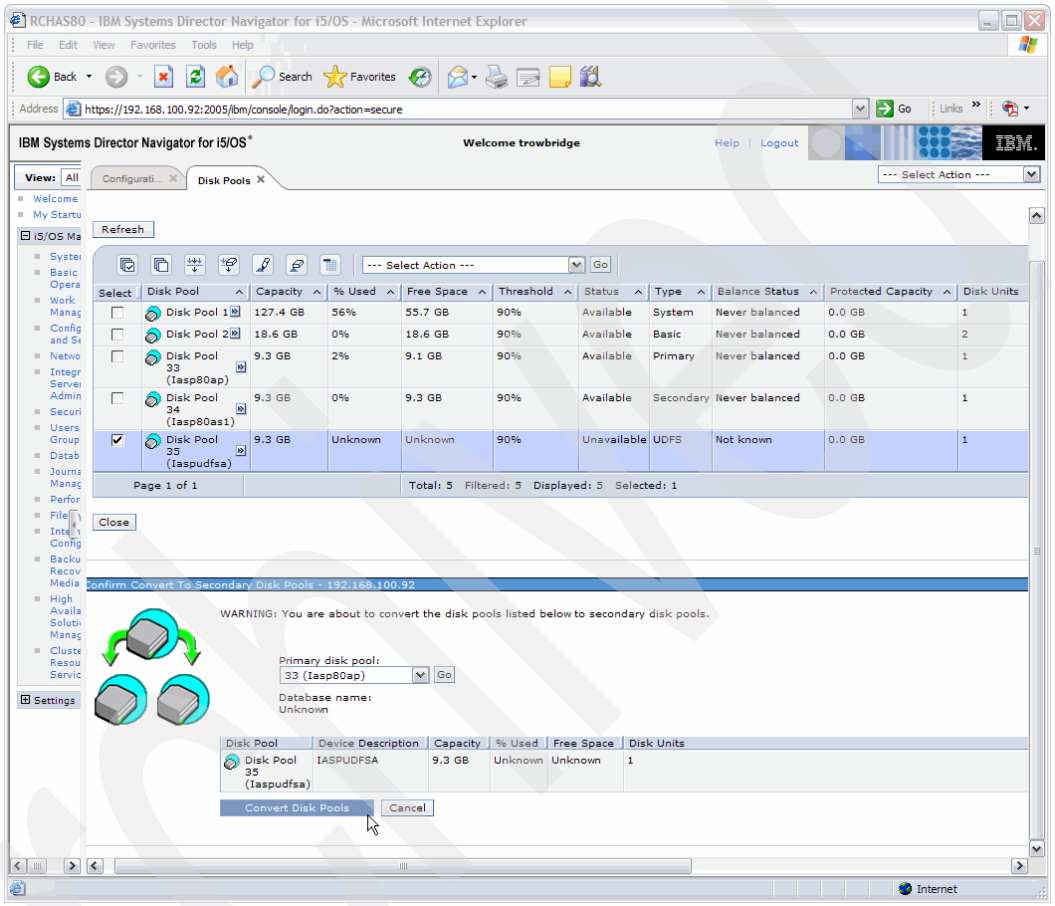


*Figure 6-21 Confirm conversion to secondary disk pool*

6. A completion message is returned when the conversion completes. Selecting **Refresh** on the Disk Pools list panel shows that the selected pool is now a secondary type pool, but it is still unavailable, as shown in Figure 6-22.

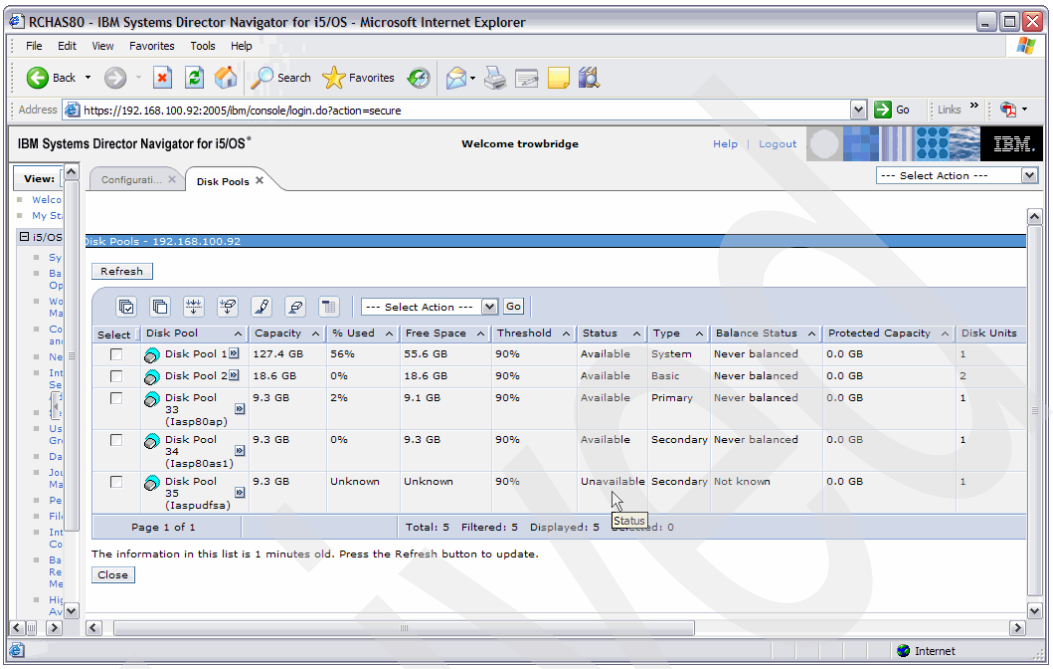


*Figure 6-22 Converted disk pool unavailable*

7. The disk pool must be made available. See 6.1.4, “Making an independent disk pool unavailable” on page 85.

8. Displaying the properties of the disk pool shows that it is now a secondary pool associated with the selected primary disk pool and database and is available, as illustrated in Figure 6-23.

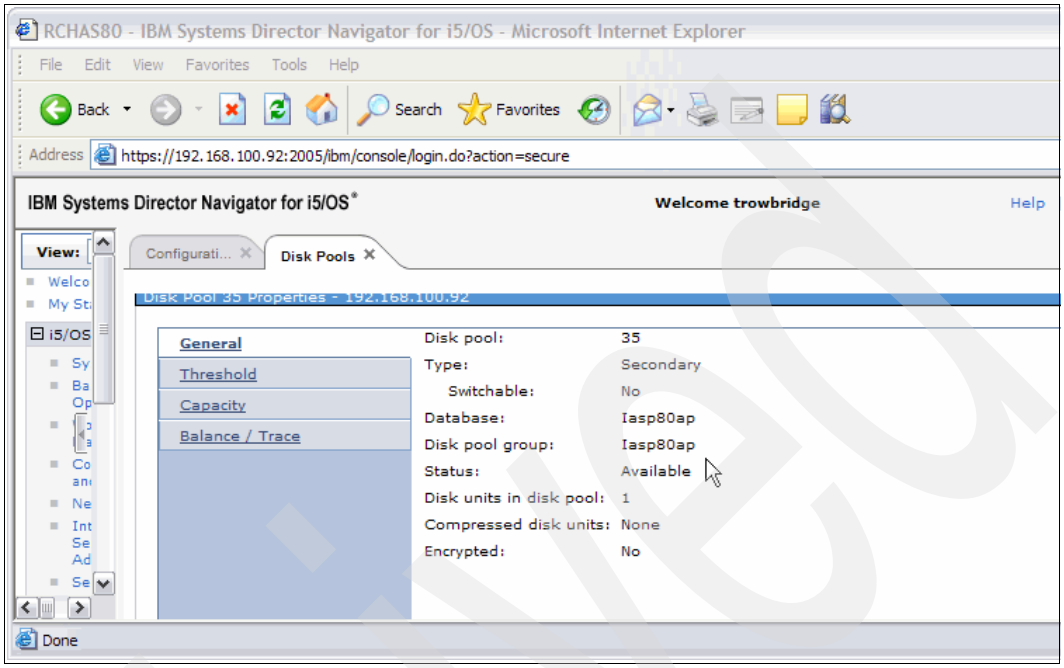


*Figure 6-23 Disk pool properties*

## 6.1.12 Setting the threshold of a disk pool

You can set the threshold of a disk pool by following these steps:

1. Open the Disk Pools list as described in 6.1.3, “Disk pool and disk pool group” on page 84.
2. Click the double arrow icon next to the disk pool name and number on the entry in the disk pools list. This opens a pop-up menu.

3. Select the **Properties** option, as shown in Figure 6-24.

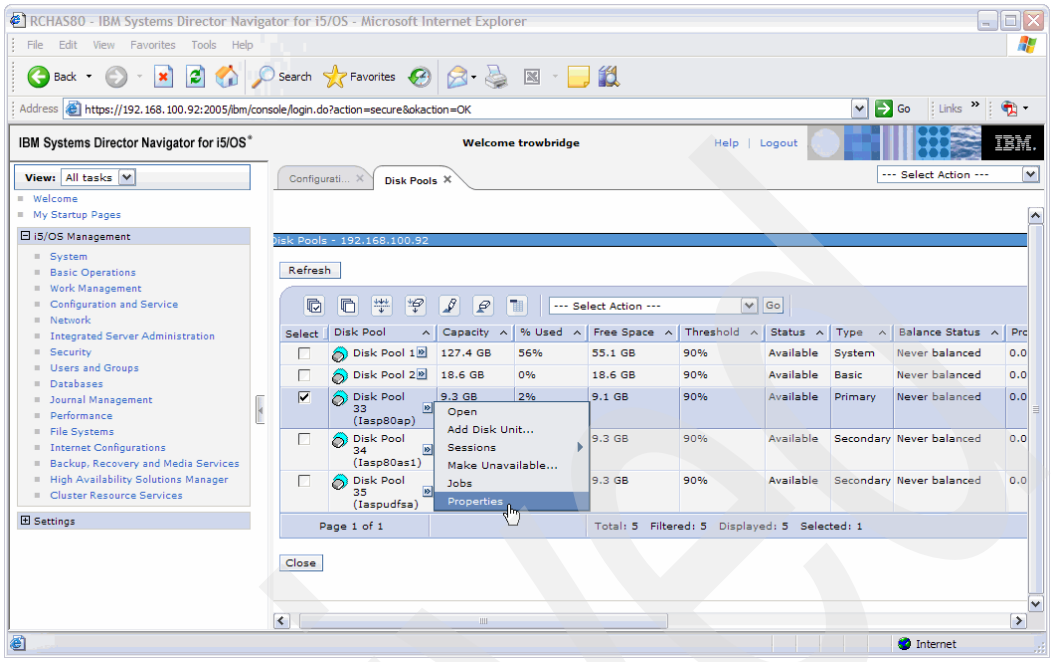


*Figure 6-24 Properties selection for a disk pool entry*

4. Select the **Threshold** tab. On this page, specify the new threshold value for the disk pool, as shown on Figure 6-25, and click the **OK** button.

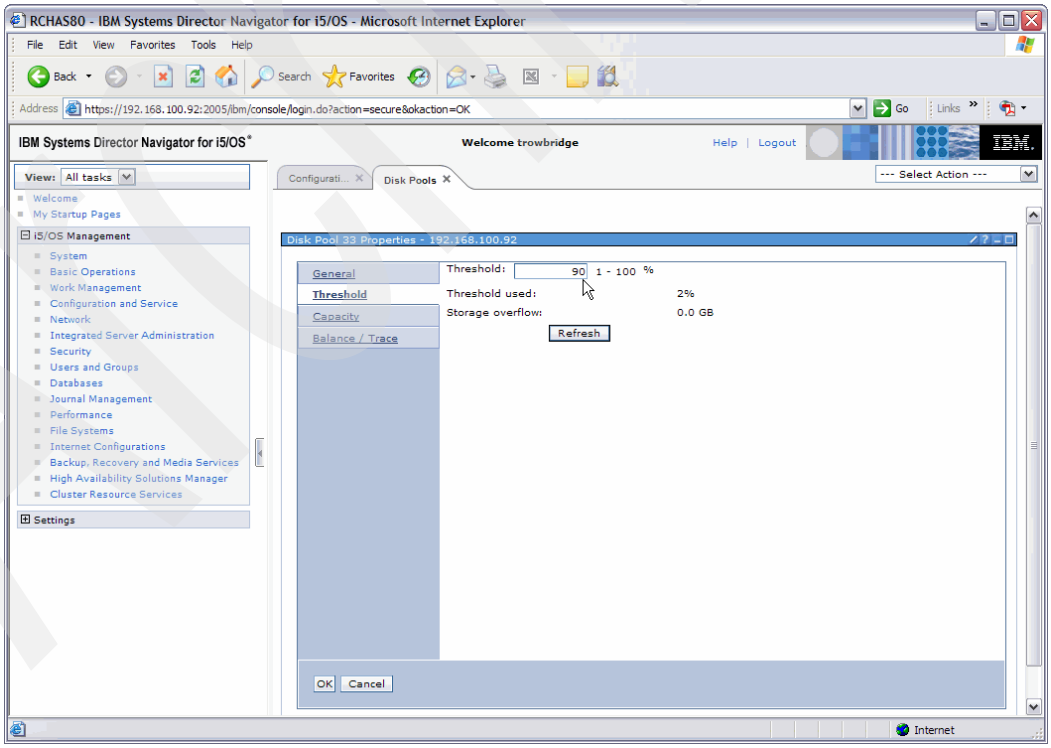


*Figure 6-25 Changing the threshold value for a disk pool*

5. A completion message is returned when the change is made. Select **OK** on the message to return to the disk pool list. Refresh the list and it should show the new threshold value.

During system operation, if the threshold value for the disk pool is exceeded, the message CPI0953, as shown in Figure 6-26, is sent to the QSYSOPR message queue.

```
                        Additional Message Information

 Message ID . . . . . :   CPI0953       Severity . . . . . :   90
 Message type . . . . . :   Information
 Date sent  . . . . . :   08/26/09      Time sent  . . . . . :   14:06:20

 Message . . . . :   ASP 33 storage threshold reached.
 Cause . . . . . :   The amount of storage used in auxiliary storage pool (ASP)
   33 has reached the threshold value of 50.000000 percent.  This is a serious
   system condition.  The auxiliary storage capacity for ASP 33 is 9544368128
   bytes. This message will be repeated until the amount of storage used is
   reduced to less than 50.000000 percent.
 Recovery  . . . :   The amount of storage used by ASP 33 must be reduced below
   the threshold value of 50.000000 percent.  The amount of storage used in ASP
   33 can be monitored by using the Work with Disk Status (WRKDSKSTS) command
   or by using the System Service Tools (SST) function.  To reduce the amount
   of storage used, do the following:
     -- Delete objects from ASP 33 that are not needed.
     -- Save objects from ASP 33 that are not needed online by specifying
   STG(*FREE) on the Save Object (SAVOBJ) command.

     For additional information on ASP management, see the Recovering your
   system book, SC41-5304.
                                                                      Bottom
Press Enter to continue.

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

*Figure 6-26 CPI0953 ASP storage threshold reached message*

## 6.1.13 Disk pool overflow

There is a difference between user ASPs and independent ASPs when it comes to data overflow. That is, user ASPs overflow and independent ASPs do not. An overflow of a basic user ASP occurs when the ASP fills. The excess data spills into the system ASP.

IASPs are designed so that they cannot overflow. Otherwise, they would not be considered independent or switchable. An IASP is allowed to fill up and the application that is responsible for filling it up simply halts. There is no automatic cancellation of the responsible job. If this job is running from a single-threaded JOBQ, in a single-threaded subsystem all further processing is stopped until user action is initiated.

## 6.1.14 What to do when a disk pool fills up

When a disk pool fills up, the job that generates the data that filled up the disk pool may not be complete. The system generates an MCH2814 message indicating this condition, as shown in Figure 6-27.

```
                        Additional Message Information

Message ID . . . . . :   MCH2814       Severity . . . . . :   30
Message type . . . . . :   Escape
Date sent  . . . . . :   03/22/02      Time sent  . . . . . :   17:35:06

Message . . . . :   ASP resources exceeded.
Cause . . . . . :   Resources of ASP  &1, ASP number 34, were exceeded.
    The reason code is 2.  The reason codes are:
    0 - Unspecified reason code.
    1 - ASP addresses are unavailable.
    2 - The storage space of the ASP has been exceeded.
Recovery  . . . :   For reason code 2, either destroy some objects that reside
  in the ASP or add an additional disk unit to the ASP.
```

*Figure 6-27 ASP resources exceeded*

This may have serious ramifications. Cancelling the offending job relieves the problem in most cases. The system does not automatically cancel the offending job. If the job is from a single-threaded JOBQ or a single-threaded subsystem, other jobs behind it are held up until the offending job is handled. Possible scheduling impacts may occur.

## 6.1.15 Removing a disk unit from an IASP

You can remove a disk unit from an IASP when it is unavailable by following these steps:

1. Access the disk pools listing as described in 6.1.3, “Disk pool and disk pool group” on page 84.
2. Click the disk pool that will have the disk removed or select the double arrow icon on that list entry, opening the pop-up menu. Then select the **Open** option, which is the default, as shown in Figure 6-28.

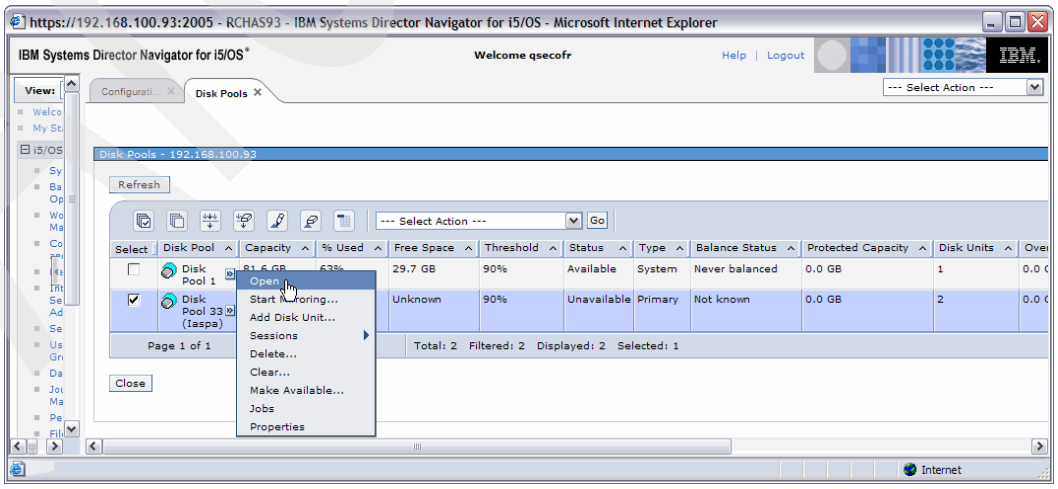


*Figure 6-28 Open disk pool to show disk units in pool*

3. Select the disk unit to be removed, select the double arrow icon, and in the pop-up menu select **Remove**, as shown in Figure 6-29.

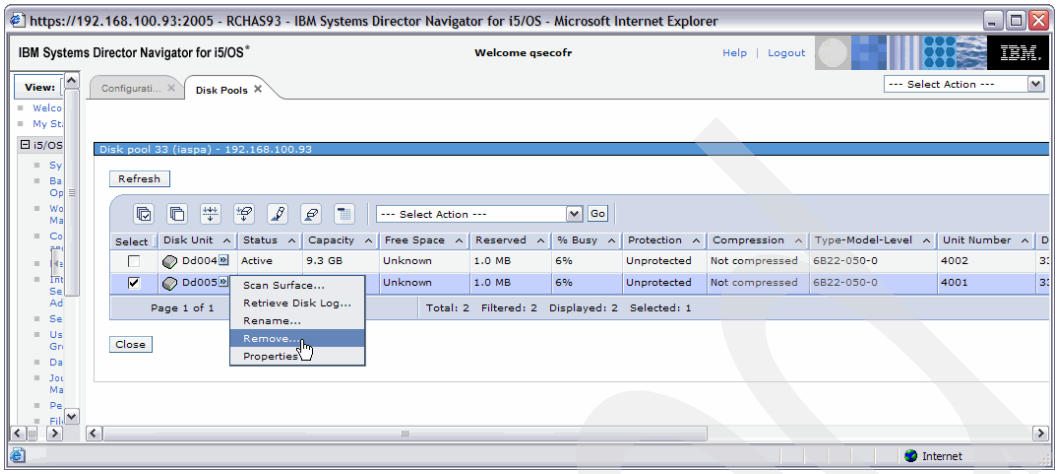


*Figure 6-29 Remove disk unit request*

4. On the confirmation panel displayed, select **Remove Disk Units**, as shown in Figure 6-30.

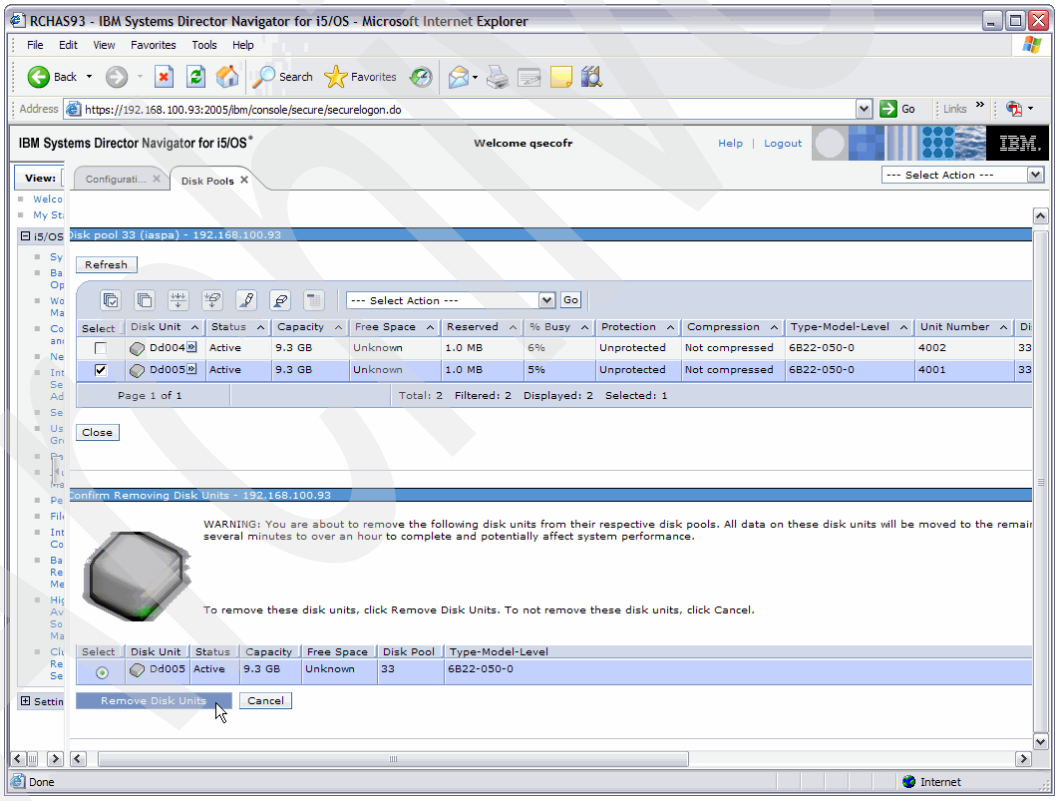


*Figure 6-30 Remove disk unit confirmation*

5. To determine the results of the request, select **Refresh**, as shown in Figure 6-31, until the remove is completed, then close the panel.

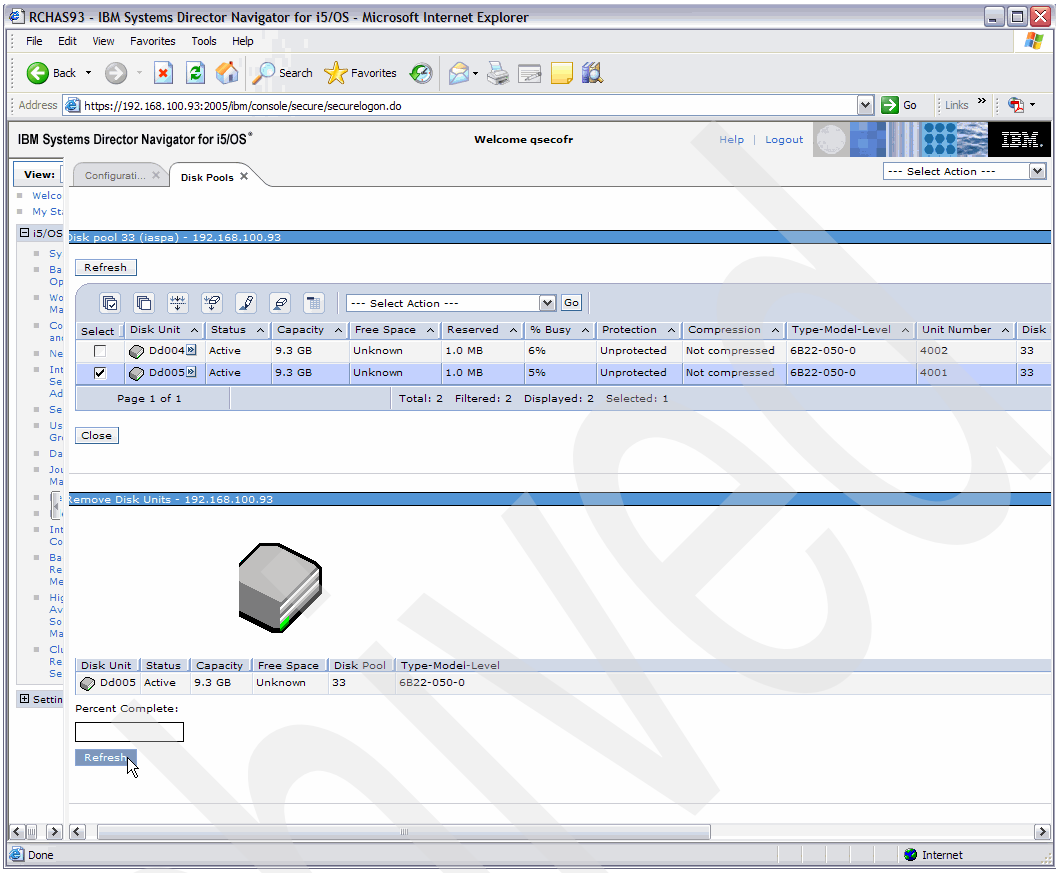


*Figure 6-31 Remove disk unit refresh until complete*

## 6.1.16 Adding a disk unit to an existing IASP

You can add a non-configured disk unit to an IASP by following these steps:

1. Access the disk pools listing as described in 6.1.3, “Disk pool and disk pool group” on page 84.

2. Select the double arrow icon on the disk pools list entry, opening the pop-up menu, then select the **Add Disk Unit** option, as shown in Figure 6-32.

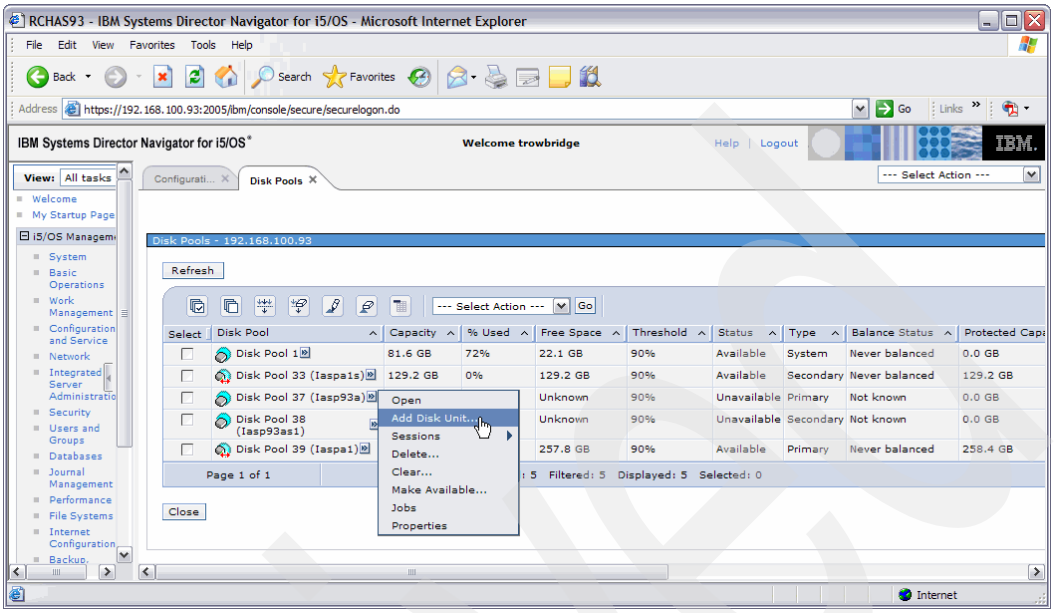


*Figure 6-32 Add Disk Units option on the Disk Pool panel*

3. The Add Disk Unit Wizard is opened, as shown in Figure 6-33. Select **Next**.

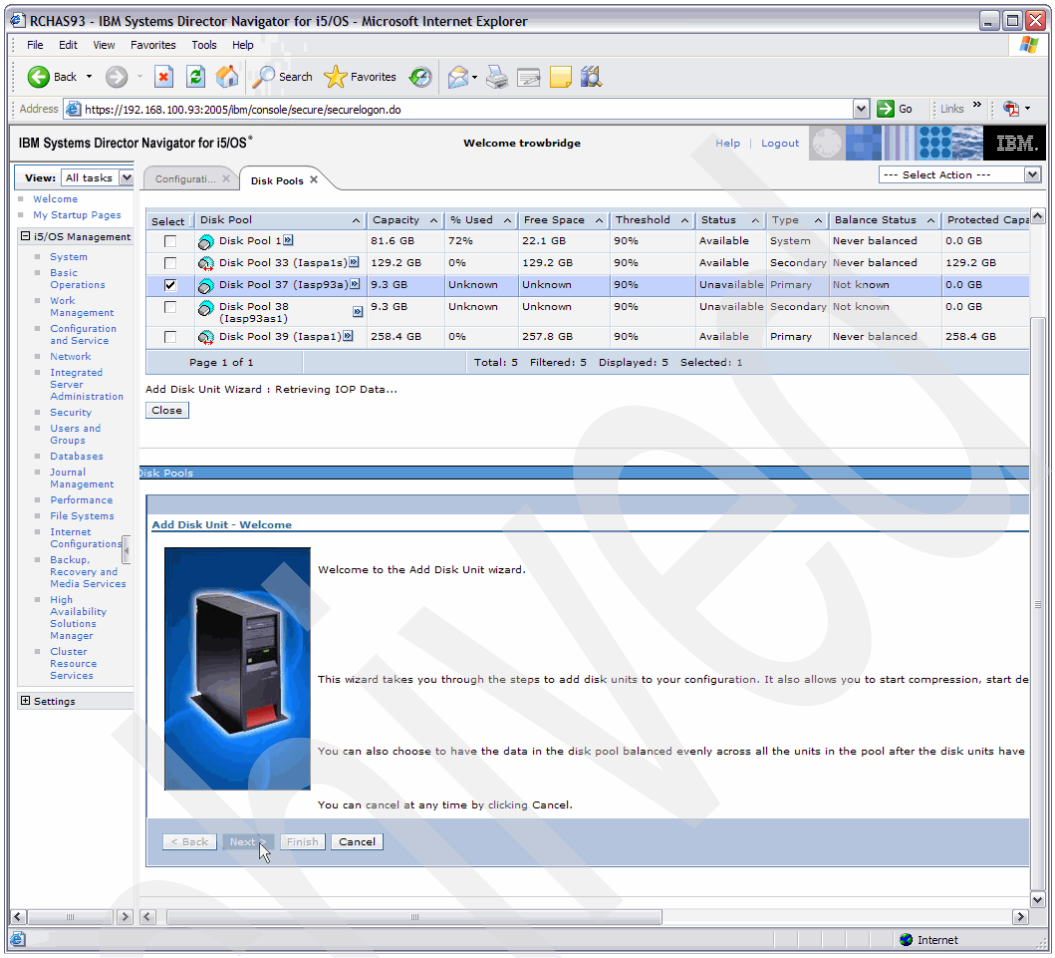


*Figure 6-33 The Add Disk Wizard*

4. Select **Add Disks**, as shown in Figure 6-34.

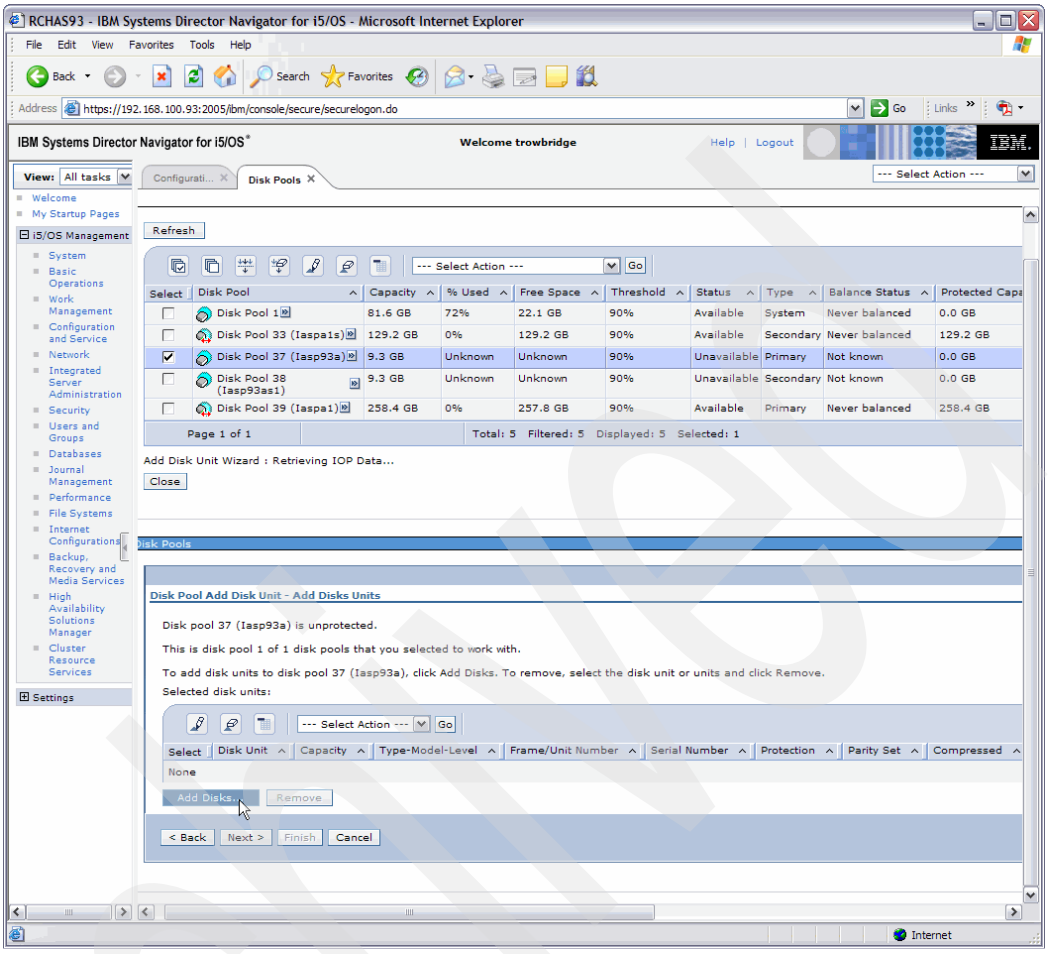


*Figure 6-34 Add Disk button in the Add Disk Wizard*

5. A list of unconfigured disks available to be added to the disk pool is displayed. Select the disks that you want to add to the disk pool, then select **Add**, as shown in Figure 6-35.

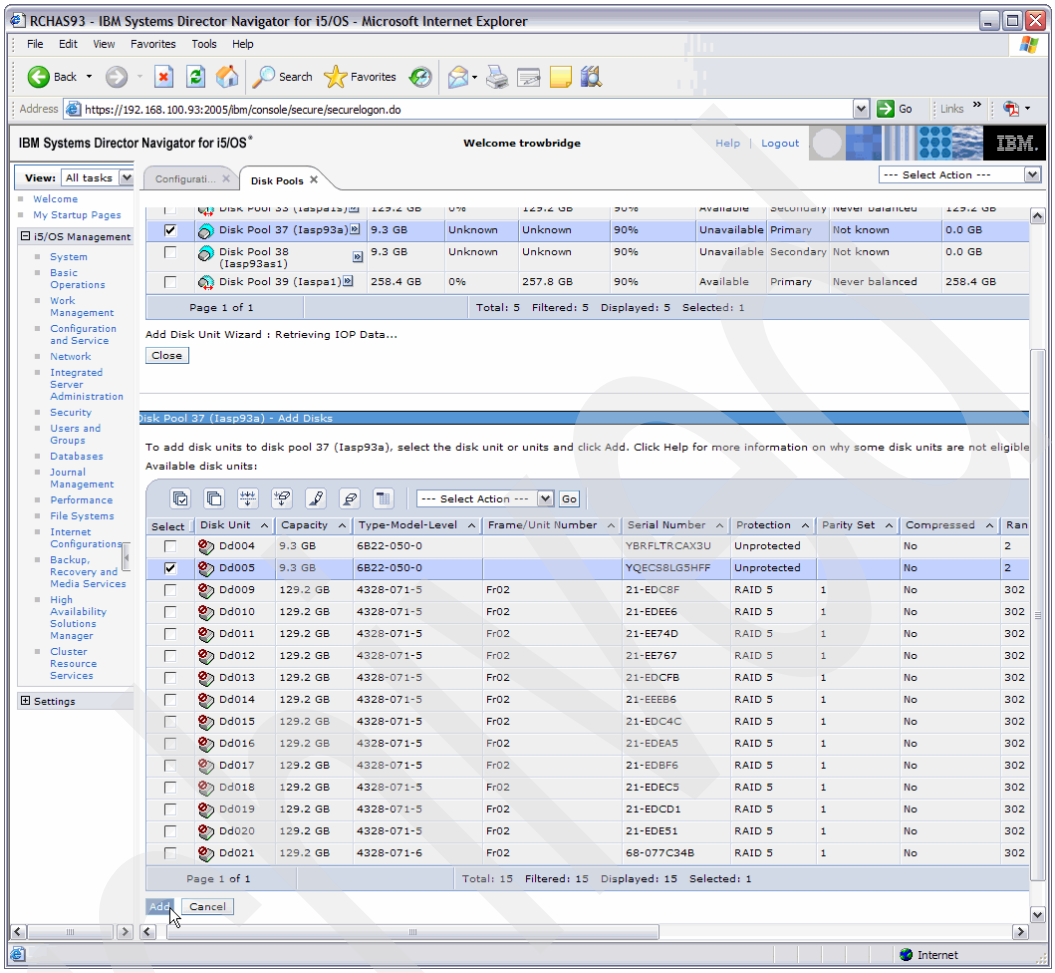


*Figure 6-35 Add button on the Disk Select panel*

6. A confirmation panel is returned listing the disk units that have been selected to be added to the disk pool. If there are no more disks to be added to this pool, select **Next**, as shown in Figure 6-36. Otherwise, select **Add Disks** again to go back to step 5 on page 110 to select additional disks to be added.

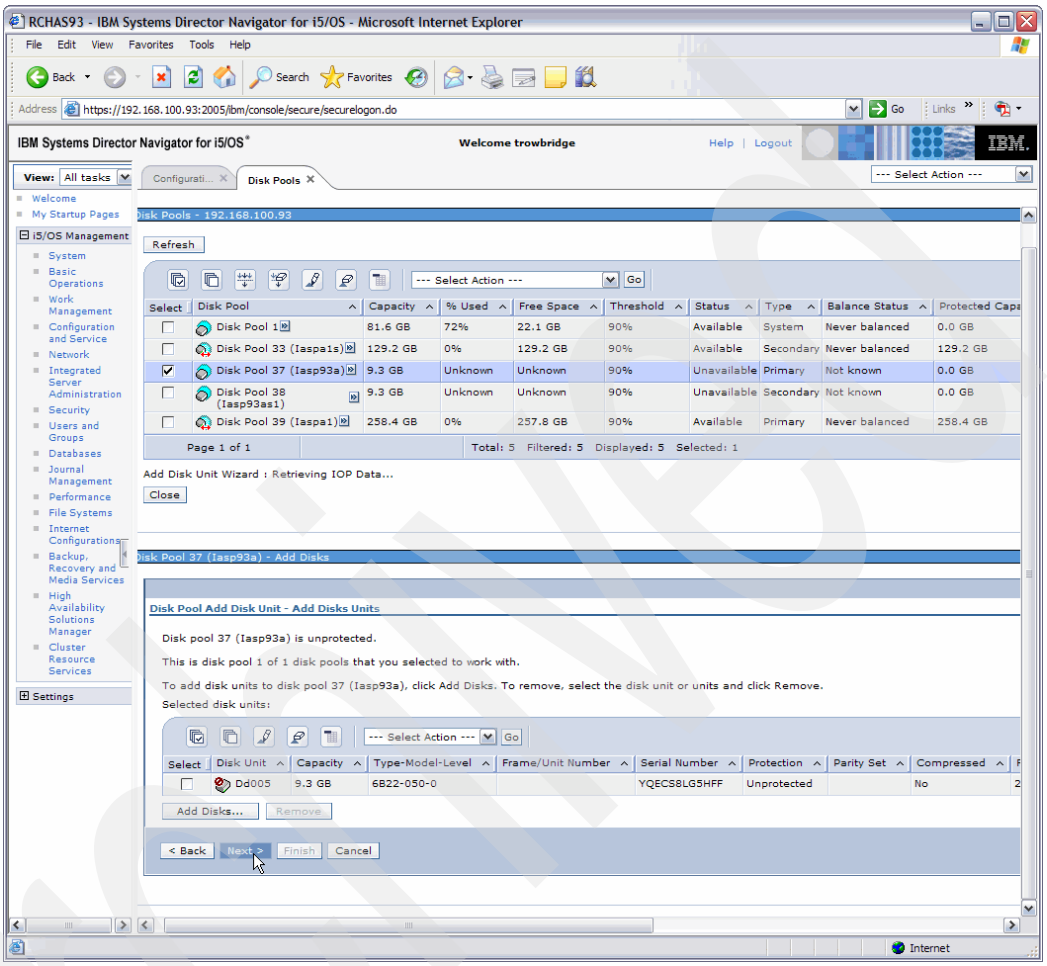


*Figure 6-36 Add disk selection confirmation and continuation*

7. On the next panel, indicate whether data balancing is to be performed when the disk are added to the pool, as shown in Figure 6-37. Balancing occurs during the process of adding the disk to the pool if requested.

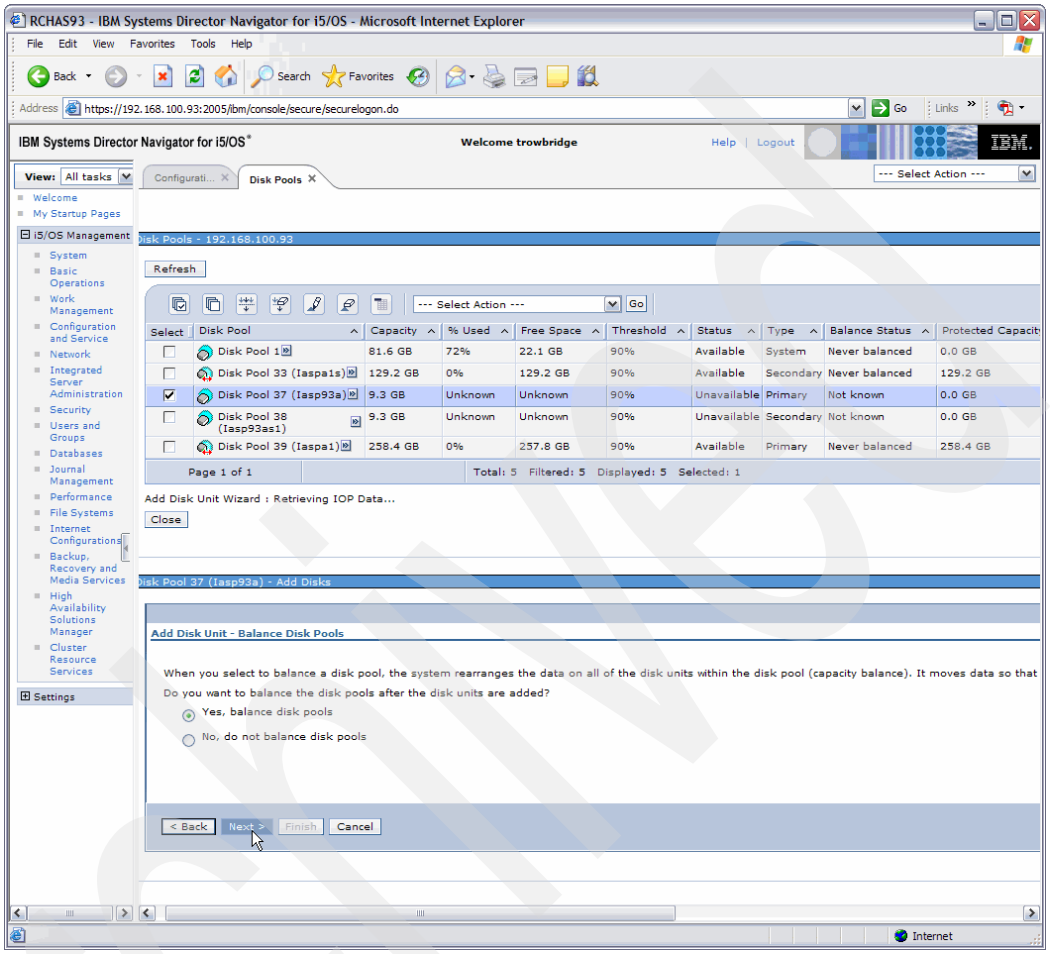


*Figure 6-37 Add Disk Wizard, balancing request*

8. The Add Disk Unit - Summary panel is then displayed, as shown in Figure 6-38, with the disk pool and disk configuration that will be implemented when **Finish** is selected on this panel. If this configuration is correct, then select **Finish** to start the update of the disk pool configuration. If it is not, select **Back** to make additional changes or **Cancel** to exit the wizard without performing any updates.

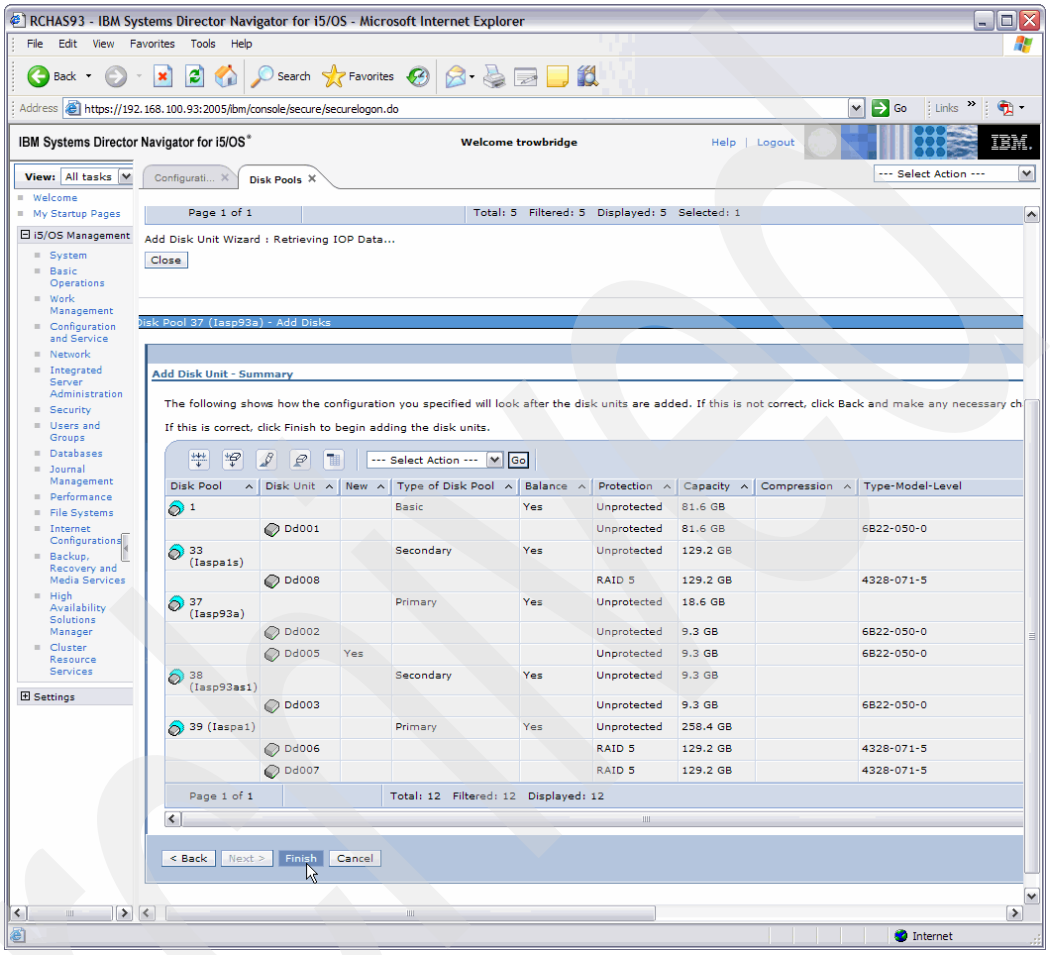


*Figure 6-38 Add disk units summary*

9. The Add Disk Unit Status panel is then presented. Select **Refresh**, as shown in Figure 6-39, to update the status. This panel does not automatically update. Continue to select **Refresh** to see the current status.

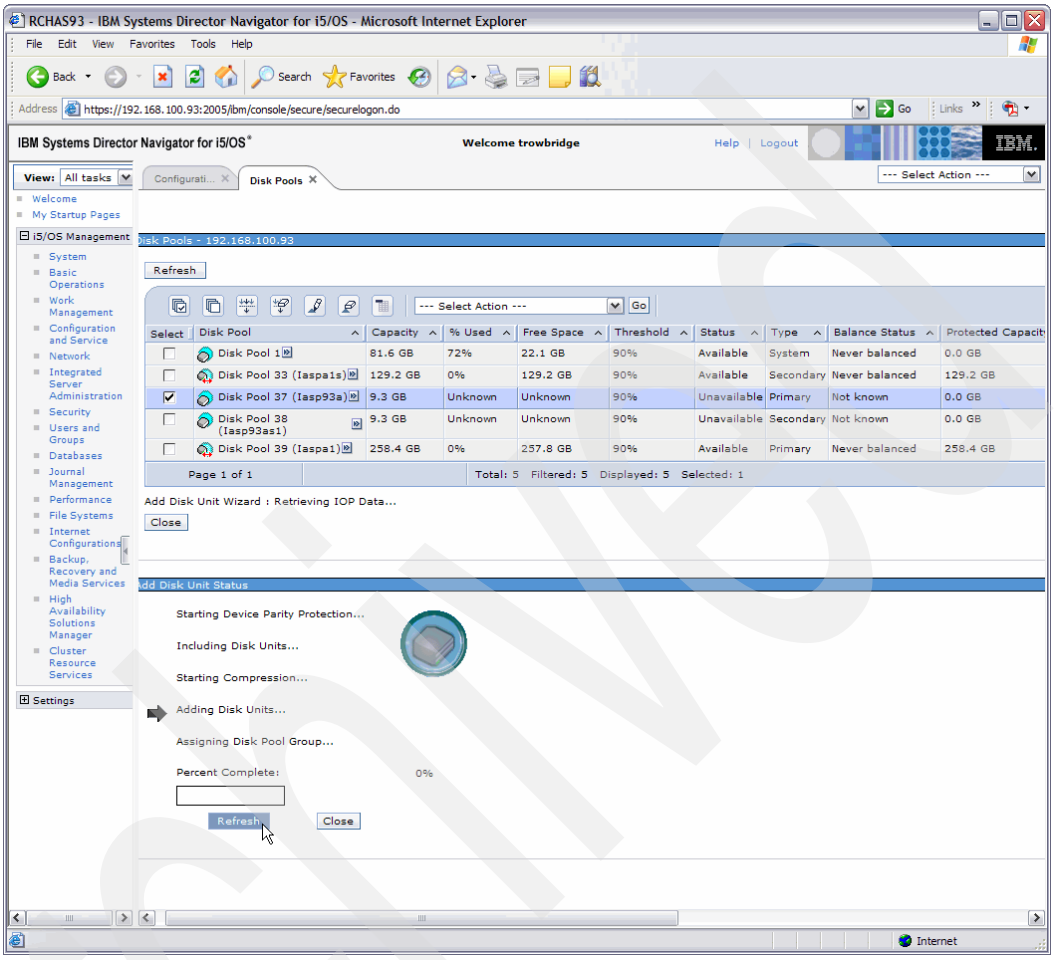


*Figure 6-39 Add Disk Unit Status*

10. Once the add disk unit function has completed and you select **Refresh**, the action complete message will be returned. Select **OK** on this panel, as shown in Figure 6-40.

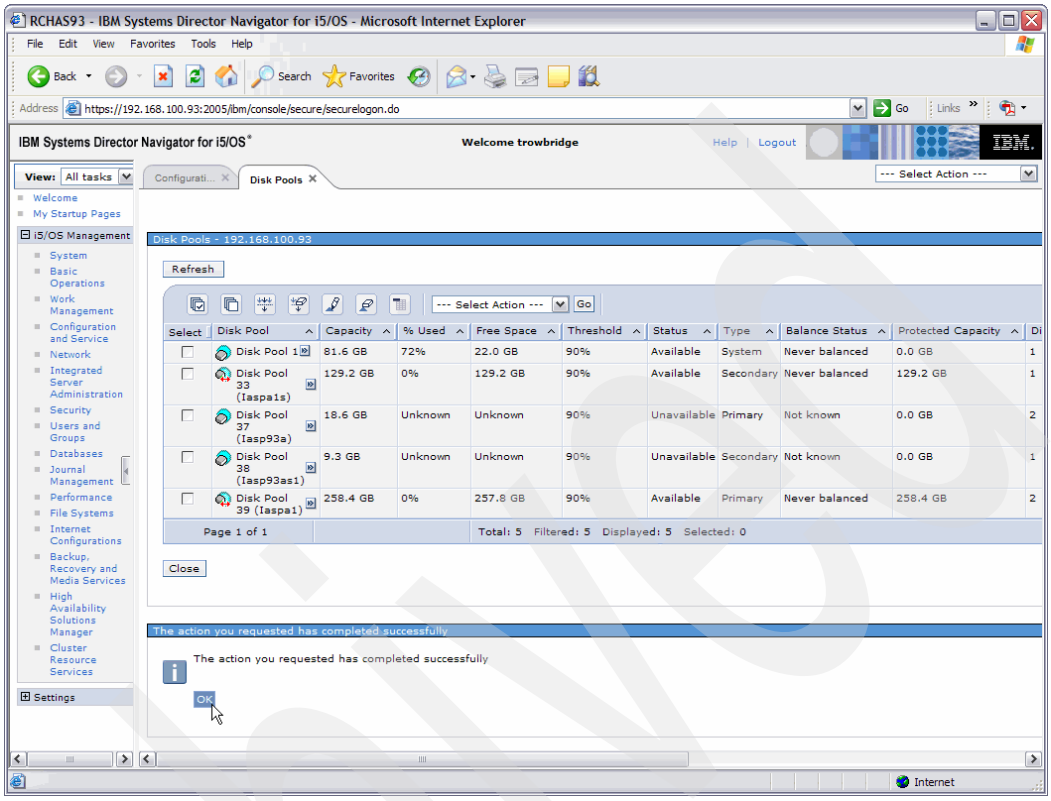


*Figure 6-40 Add Disk Unit Completion message*

11. To see the new configuration of the disk pool after the addition of new disk units, open the disk pool entry from the disk pool listing, as shown in Figure 6-41 (either double-click the disk pool name or select the double arrow and click **Open**.

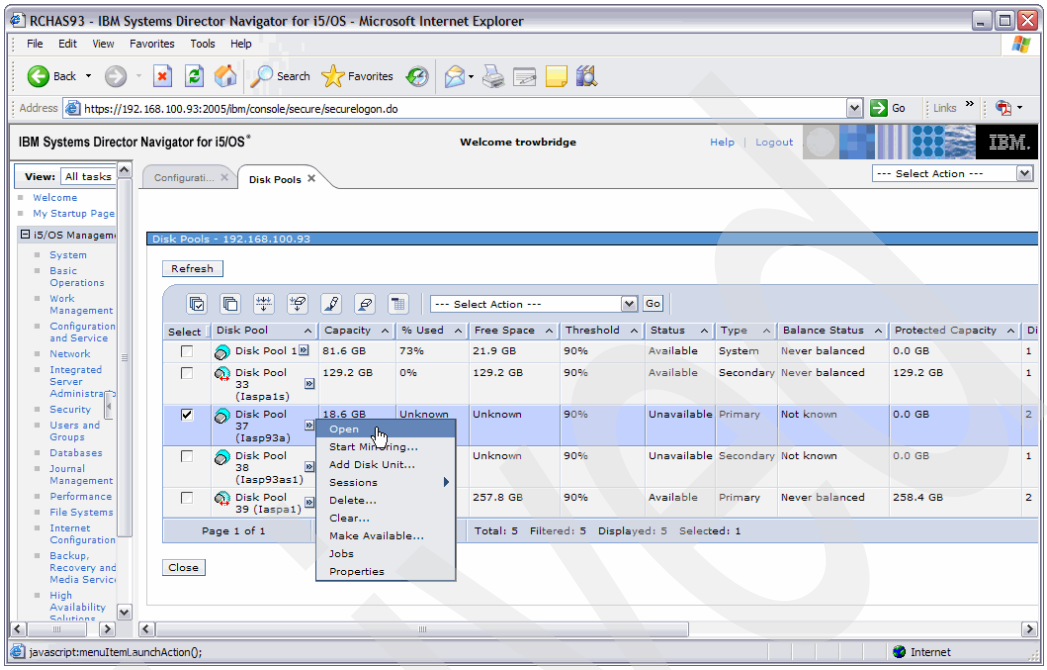


*Figure 6-41 Reviewing disk pool disk configuration (1)*

12. The disk pool disk unit listing is displayed, as shown in Figure 6-42.

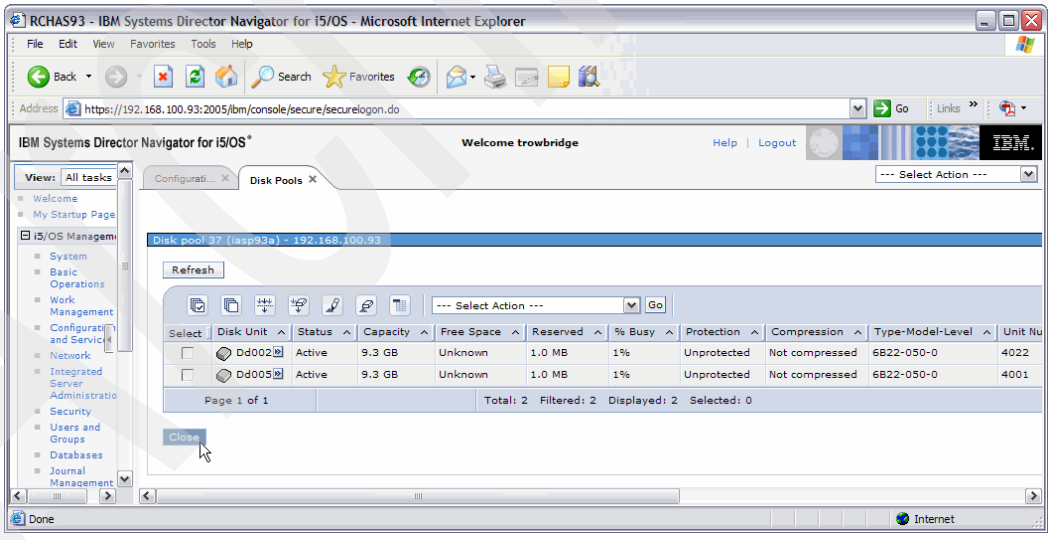


*Figure 6-42 Reviewing disk pool disk configuration (2)*

## 6.1.17 Reclaim storage and IASPs

With the introduction of IASPs you can run reclaim storage (RCLSTG) on an IASP while the rest of the system keeps running. This means that multiple IASP RCLSTG processes can run concurrently—one for each IASP on the system.

**Note:** Reclaiming an IASP device requires that there can be no active users of the IASP device that is the subject of the reclaim.

The RCLSTG command supports two value sets for the ASPDEV parameter:

- *SYSBAS

  If the *SYSBAS value is specified for the ASP device, the reclaim storage operation is performed on the system and on all traditional user-defined ASPs (2–32). The system must be in a restricted state to run this. In addition, all IASP must be in an unavailable (varied off) state.

- IASP device or group name

  Reclaim storage for an IASP device can run without the system being in restricted state. In fact, because the IASP must be available or varied on to be seen, the system *cannot* be in a restricted state. The RCLSTG command is only valid in interactive mode, so it must be run from a terminal session, but multiple jobs can be executing, each performing RCLSTG on a different ASP device. Multiple ASP devices can be reclaimed in parallel.

Figure 6-43 shows an example of a reclaim storage command execution specifying an IASP.

```
Command Entry                          RCHAS80
                                                        Request level:   4
Previous commands and messages:
   > RCLSTG ASPDEV(IASP80AP)
     Database cross reference files processed successfully.
     RCLSTG processing complete. 574 objects processed. 9 deleted.

                                                                      Bottom
Type command, press Enter.
===>



F3=Exit   F4=Prompt   F9=Retrieve   F10=Include detailed messages
F11=Display full      F12=Cancel    F13=Information Assistant   F24=More keys
```

*Figure 6-43 RCLSTG command*

**Note:** Additional restrictions and considerations exist for running the RCLSTG command. Review the interactive help for this command before using it.

## 6.2 ASP and disk unit numbering

The system assigns an ASP number to the system ASP and IASPs. The user assigns ASP numbers to user ASPs. Disks are also assigned unique numbers by the system as they are assigned to disk pools, depending on their usage.

### 6.2.1 The numbering scheme follows these guidelines

ASPs are numbered as follows:

- The system ASP is always number 1.
- User ASPs are assigned by the user in the range of 2 through 32.
- IASP numbers range of 33 through 255.

Disk are assigned numbers in the following ranges:

- Disk drive numbers in the system ASP and user ASPs (1–33) range from 1 through 2047.
- Disk drive numbers in IASPs are assigned numbers in the range 4001 through 6047.
- Disk drive numbers in the range 2048 through 4000, and greater than 6047, are reserved for future use.

### 6.2.2 Device domains and disk numbering

The construct known as a *device domain* is a subset of cluster nodes that share device resources. A device resource may be a disk pool.

A function of a device domain is to prevent conflicts that can cause the failure of an attempt to switch a resilient device between systems. The resources involved in a device domain include the structures used to identify and manage the content of the structures across the multiple systems that are involved in the domain.

Across an IASP device domain, numbers assigned to disk are unique. The separate ranges provide additional disk unit numbers. Disk numbering follows these standards:

- Disk drive numbers in IASPs are assigned numbers in the range 4001 through 6047.
- Each node in the cluster can have up to 2047 disk units in its system and user ASPs.
- A device domain can have up to 2047 units in all IASPs within the device domain.
- There can be 128 nodes in a cluster, and each node can be its own device domain. This means that there may be as many as 128 device domains in a cluster.

Figure 6-44 shows the IASP pool number and disk unit numbers assigned by the system to an IASP. The disk units in this example are in the range 4001 through 4004. The IASP numbers are 33 through 36.

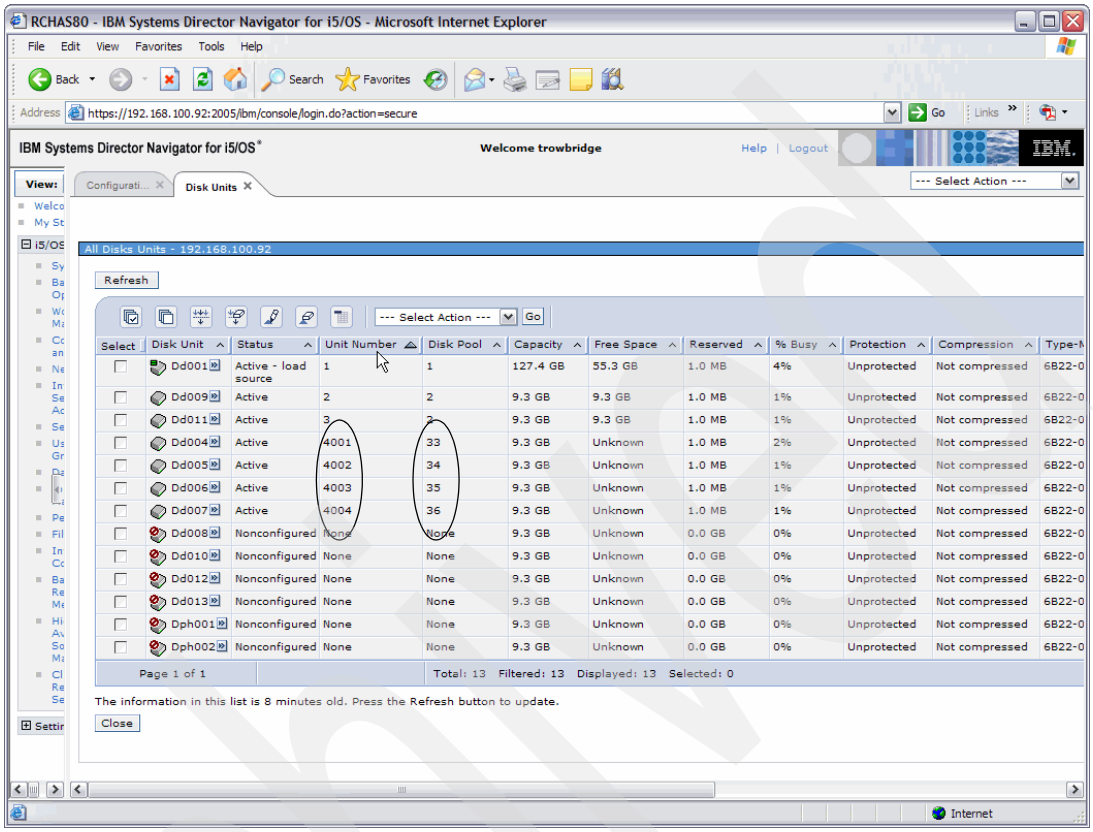


*Figure 6-44 Disk and disk pool numbering example*

## 6.3 Switching independent disk pools

This section discusses the actions necessary to initiate the switching of a cluster resource group. It is not intended as a complete guideline for using cluster resources. For more information about using clustering technology refer to *Implementing PowerHA for IBM i*, SG24-7405.

A planned switch of an independent disk pool to its designated backup can be performed at any time, provided that cluster resource services are available and the cluster nodes are synchronized. Cluster software performs the process of switching independent disk pools between nodes within a domain. If the independent disk pool is to be switched between logical partitions, then the disk units can be either internal or external components, but cannot be based on virtual disk units.

If the IASPs are to be switched between machines (that is, separate central electronic complexes (CECs)), then the disk units must be external, such as those in a switchable tower or a storage area network (SAN).

### 6.3.1 Planned disk pool switch

To perform a planned disk pool switch, follow these steps:

1. Place the job queues on hold for jobs using the disk pool.

   **Note:** Remember that if the job queue object is in an IASP managed by the CRG being switched, the IASP is made unavailable and the job queue contents are discarded.

2. End all jobs using objects within the IASP or that have the IASP in their namespace. These jobs are ended immediately during the switch.
3. Make the disk pool unavailable using Systems Director Navigator, as explained in 6.1.4, “Making an independent disk pool unavailable” on page 85. Cluster resource services will also make the disk pool unavailable before switching or failing it over to the target (new hosting system).
4. Expand **i5/OS Management** → **Cluster Resource Services** in the Systems Director Navigator.
5. Select the **Work with cluster resource groups**, as shown in Figure 6-45.

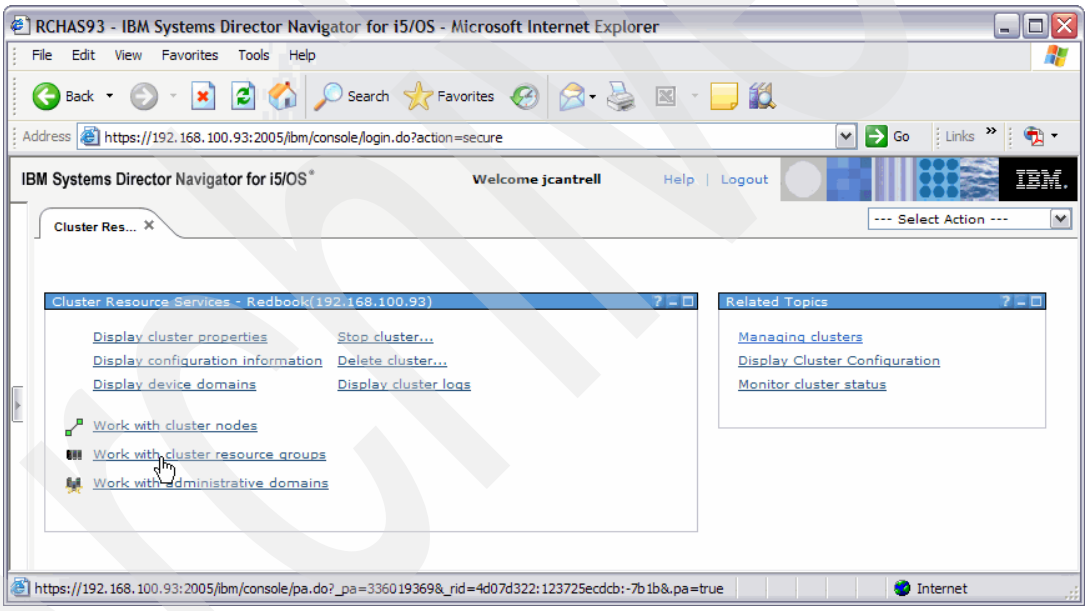


*Figure 6-45 Selecting Work with cluster resource groups*

6. Use the double arrow icon on the Cluster Resource Group entry to open the pop-up menu, then select the **Switch** option, as shown in Figure 6-46.

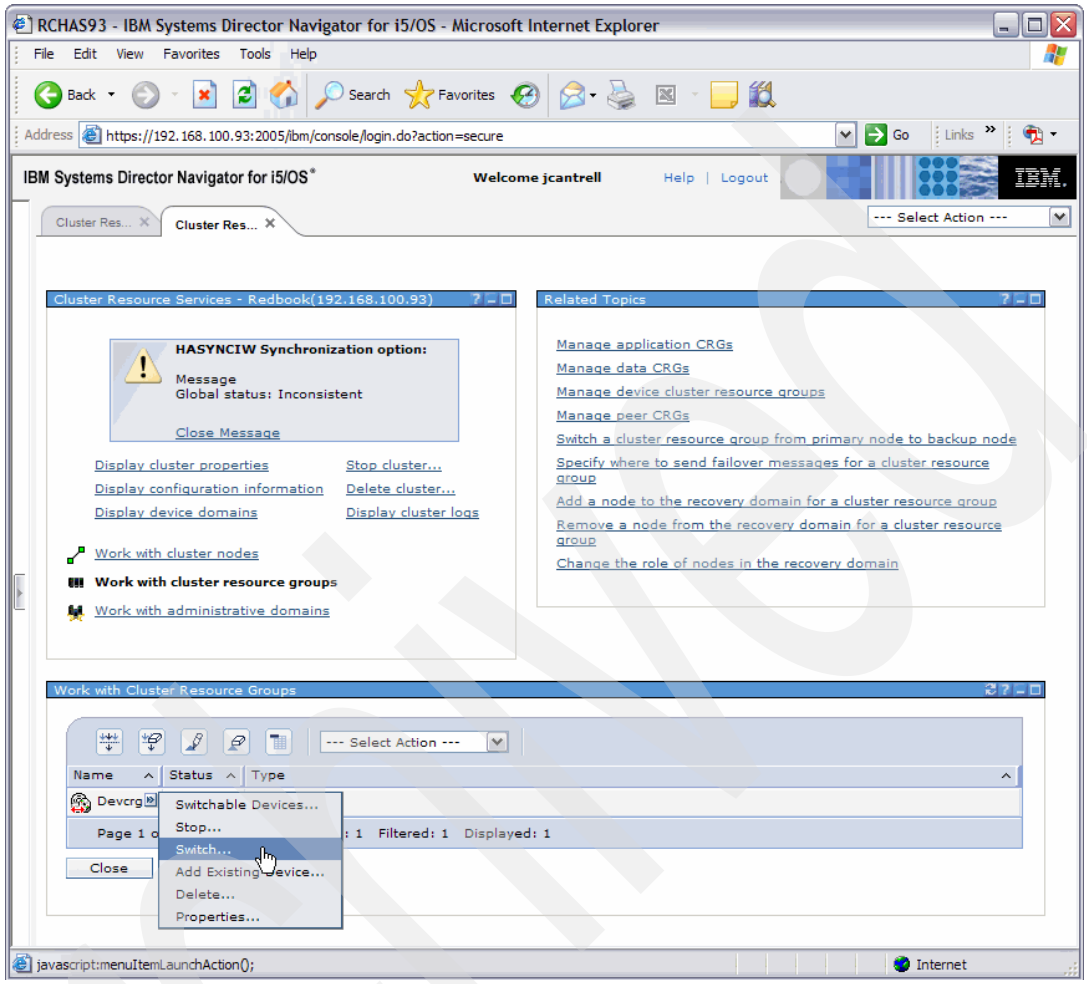


*Figure 6-46 Selecting the Switch option*

7. A confirmation panel is displayed to ensure that the switch is to be run, as shown in Figure 6-47.

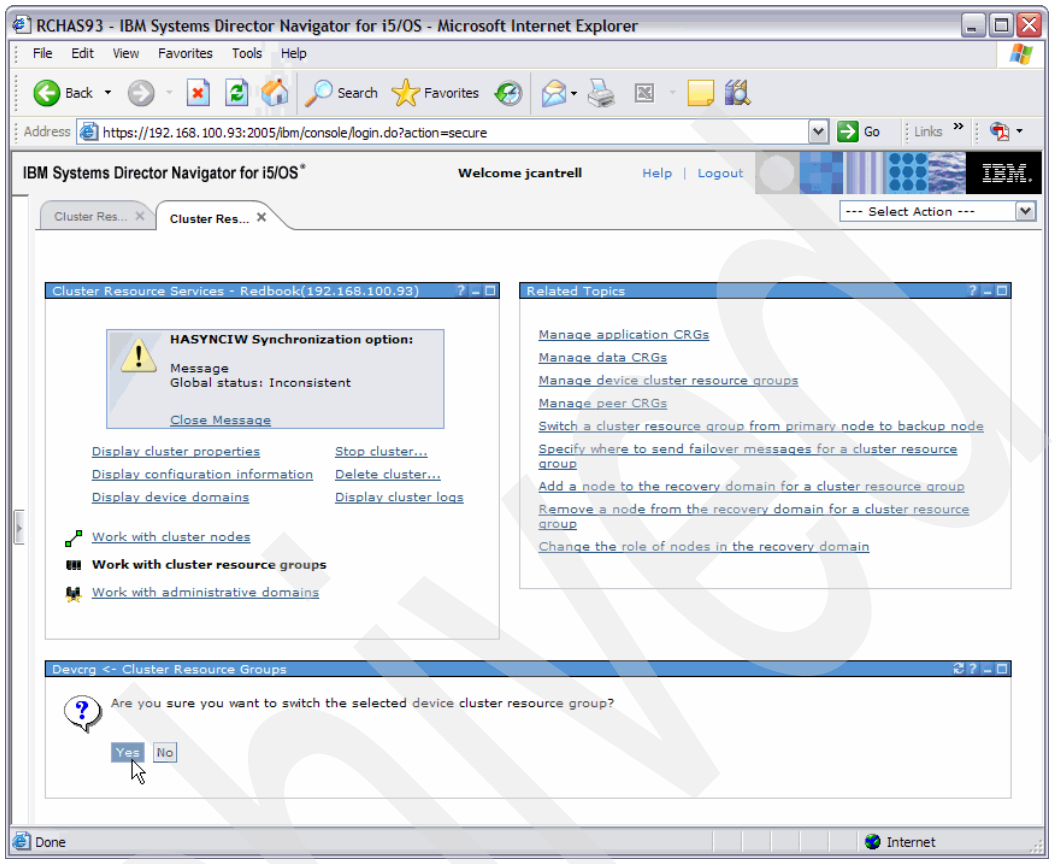


*Figure 6-47 Switch resource group confirmation*

8. The original CRG panel is re-displayed when the switch is complete, as shown in Figure 6-48.

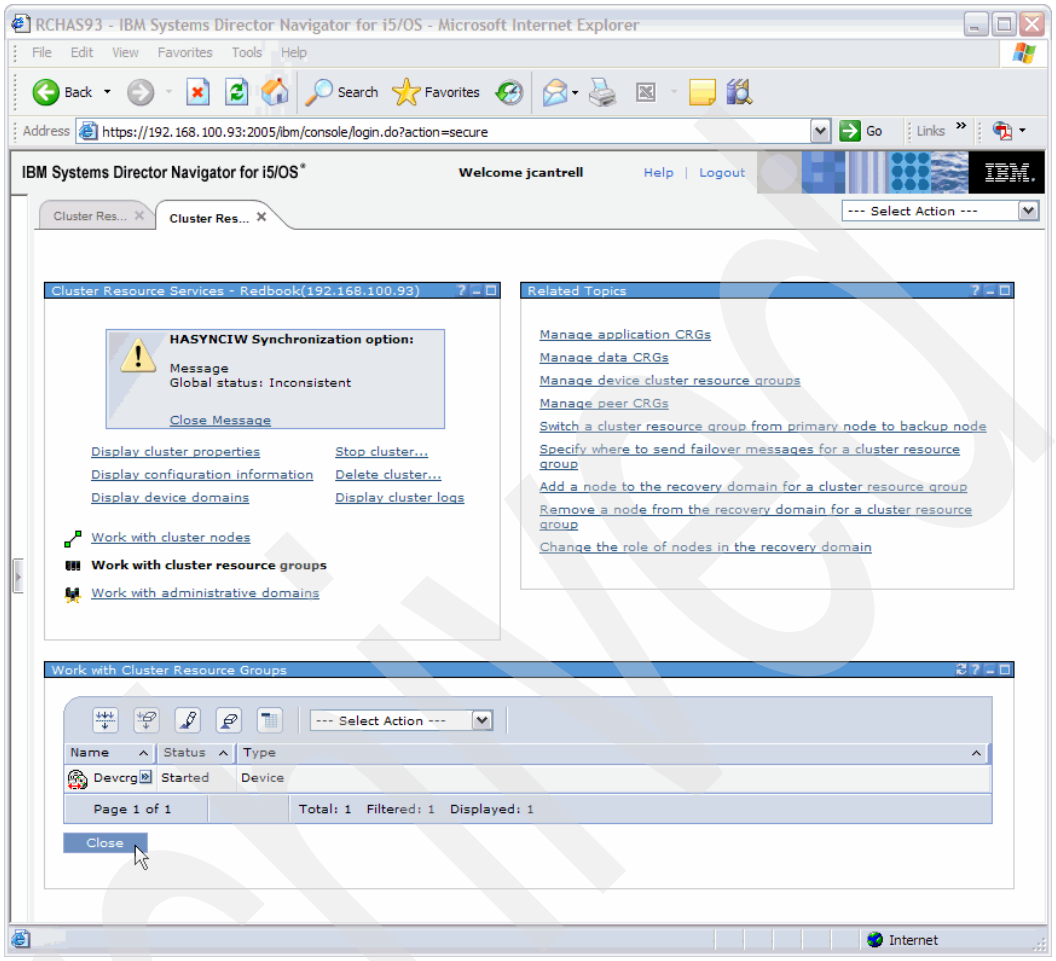


*Figure 6-48 Switch completed display*

9. The availability of the IASP can be determined using the disk pool functions of the System Director Navigator, as described in “Disk pool and disk pool group” on page 84.

Optionally, to observe the actions of the switch process on the target machine, as the ASP device status changes, enter the following command:

```
WRKCFGSTS *DEV *ASP
```

When you see `Available`, the Systems Director Navigator is freed (Figure 6-48) and the IASP is ready for use.

Keep in mind that an online attribute field is associated with each disk pool device description defined in the switchable hardware group (device CRG). This tells Cluster Resource Services whether to vary on (make available) the disk pool on the target node at switch over or fail over. Using the cluster GUI, by default, the online attribute is set, but can be changed. If the online attribute is not set, the disk pool is left in a varied off (unavailable) state.

## 6.3.2 Unplanned IASP switch

A system failure or other major outage may require an unplanned switch of an IASP. This is handled in the same way as a planned switch. However, there are added delay factors due to

the same abnormal IPL considerations for rebuilding database access paths that are encountered during a system IPL. Consider using systems-managed access-path protection (SMAPP) and setting it to the shortest rebuild time that is possible.

7

# Backup and recovery

This chapter describes techniques and strategies for backup and recovery of independent disk pools. When an independent disk pool is added to the system configuration, a plan for the backup and recovery of the user data on these devices should exist because these devices operate differently from the system or basic user auxiliary storage pools (ASPs). These differences require carefully planning the backup strategy to ensure that a complete system backup is obtained.

This chapter explains how to save and restore objects that are stored in independent auxiliary storage pools (IASPs).

## 7.1 Using native IBM i save and restore functions

The native IBM iSAVxxx and RSTxxx commands provide support for IASPs. Using these commands in control language (CL) programs to back up the system is relatively straightforward. In general, these native commands must have access to the namespace where the objects to be saved reside. This can be achieved by using the set auxiliary storage pool group (SETASPGRP) command or by using the ASPDEV parameter on the SAVxxx and RSTxxx commands.

Understanding the concept of using an IASP allows specific libraries or objects in an IASP to be saved or restored. However, if using the save and restore menus to save or restore the system or its components (for example, *NONSYS, *ALLUSR, or *IBM saves) is more familiar, it is important to understand the way in which these are affected by addressability to the IASPs. This is particularly important when using option 21 (entire system).

## 7.2 Saving IASPs

This section looks at various methods of saving IASPs using native IBM i SAVxxx commands. A backup strategy must be decided on. Is the entire system to be saved or is a more granular manner allowed? Why are IASPs being used? Are they switchable or non-switchable? If they are switchable, which node in the cluster will back them up? The answers to these questions can help determine the best method for saving IASPs.

The simplest form of IASP is a stand-alone IASP. This allows multiple namespaces to be created on a single IBM i image. This also creates multiple databases—one for the system ASP and one for each primary IASP. Multiple namespaces allow the same library and object names in different IASPs. Clearly, this has implications for backup and recovery.

Depending on the use of IASPs, it may be easier to back them up separately, especially if each database has different backup characteristics (for example, frequency of save, retention, or media type). If an IASP is used for archival records, a backup may only be required when the archives are made.

## 7.3 Saving your entire system

Let us walk through an example to see what option 21 (save entire system) actually does and what must be done to ensure that it works.

The following example has a system ASP with a library called MYLIBA containing a savefile MYSAVEF, a directory called MYDIRA, and a user ASP with directory UASPDIR. There are three IASPs called IASPA, IASPB, and IASPC. Each has a library called DATALIB and a directory called DATADIR. DATALIB contains objects DLOBJ1, DLOBJ2, and DLSAVEF, while the DATADIR contains files DATA1.gif, DATA2.txt, and DATA3.doc.

This example includes the following actions:

1. Save the entire system. Use option 21 from the GO SAVE menu to complete the following steps:
   a. Vary on (make available) all IASPs.
   b. Put the system in a restricted state.
   c. Save the entire system.

d. When the save is complete, you may vary off (make unavailable) the IASPs.

   The actual commands that are run during these first four steps are:

   ```
   SAVSYS
   SAVLIB SAVLIB(*NONSYS) ACCPTH(*YES) ASPDEV(*SYSBAS)
   SAVDLO DLO(*ALL) FLR(*ANY)
   SAV (('/*') ('/QSYS.LIB' *OMIT) ('/QDLS' *OMIT)) ASPDEV(*SYSBAS)
   ```

2. Build a list of available IASPs.
3. Perform the set ASP group function using one of the following commands:

   ```
   SAVLIB *NONSYS ASPDEV(*CURASPGRP)
   SAV ASPDEV(*CURASPGRP)
   ```

4. The save process is repeated for each IASP:

   ```
   CHKTAP ENDOPT(*UNLOAD)
   ```

## 7.4 Special considerations for save commands

First and foremost, the IASP must be available (varied on) in order to do any form of save of the contents of the IASP.

The ASPDEV parameter allows an IASP to be saved without changing the job thread, by specifying the name in this parameter. However, if saving to save files, this parameter does not affect the DEVICE parameter of the save commands. The SETASPGRP command must be used to save to a save file that exists in an IASP. This also allows you to save to a save file that exists in a different IASP than the one being saved. For example, there is a file called DLOBJ1 and a save file called DLSAVEF in a library called DATALIB in an IASP called IASPA. There also is a save file called MYSAVF in library QGPL.

The following command does *not* work without the SETASPGRP command:

```
SAVOBJ OBJ(DLOBJ1) LIB(DATALIB) DEV(*SAVF) SAVF(DATALIB/DLSAVEF)
```

Even adding ASPDEV as in the following command, it still does not find the save file. Therefore, the following command still will not work if the current namespace does not include IASPA:

```
SAVOBJ OBJ(DLOBJ1) LIB(DATALIB) DEV(*SAVF) SAVF(DATALIB/DLSAVEF) ASPDEV(IASPA)
```

To save the file DLOBJ1 to the save file DLSAVEF, you must use the SETASPGRP command. The correct sequence for using the save file is shown here:

```
SETASPGRP ASPGRP(IASPA)
SAVOBJ OBJ(DLOBJ1) LIB(DATALIB) DEV(*SAVF) SAVF(DATALIB/DLSAVEF)
```

To save an individual object from the integrated file system (IFS) in an IASP, the exact path to that object must be specified in the OBJ parameter of the SAV command, starting with the IASP name. In the example, to save the file DATA1.gif to a savefile in QGPL, use the following command:

```
SAV DEV('/qsys.lib/qgpl.lib/mysavef.file') OBJ(('/IASPA/DATADIR/DATA1.gif'))
```

When using the SAV command and the direct path to an object is specified as in the previous example, it is not necessary to use the SETASPGRP command or the ASPDEV parameter to set the IASP value in the job thread when the target DEV is in *SYSBAS.

To save to a save file in an IASP, you must use the SETASPGRP command before the SAV, otherwise the command will fail on the DEV parameter. In the example, the save file DLSAVEF is in the library DATALIB, which is in the IASP IASPA, and the following command fails without issuing the SETASPGRP command before the SAV:

```
SAV DEV('/iaspa/qsys.lib/datalib.lib/dlsavef.file')
OBJ(('/iaspa/DATADIR/DATA1.gif'))
```

Even specifying the ASPDEV parameter, as shown in the following example, still results in the `File DLSAVEF in library DATALIB not found` error message:

```
SAV DEV('/iaspa/qsys.lib/datalib.lib/dlsavef.file')
OBJ(('/IASPA/DATADIR/DATA1.gif')) aspdev(iaspa)
```

The SETASPGRP must be set before the save can locate the savefile in the IASP. The correct sequence for using the save file is shown here:

```
SETASPGRP ASPGRP(IASPA)
SAV DEV('/iaspa/qsys.lib/datalib.lib/dlsavef.file')
OBJ(('/IASPA/DATADIR/DATA1.gif'))
```

## 7.5 Restoring IASPs

Restoring an entire system that uses independent disk pools becomes a more complicated matter. Restore the system ASP first. Then manually create the independent disk pool or pools using the System Director Navigator. This requires knowledge of the original disk pool sizes and names. Also, if using the restore menu to recover user data, you may want to exclude the restore authority (RSTAUT) command until you recover all of your IASPs.

The recovery of the system involves this procedure:

1. Install the licensed internal code (LIC).
2. Configure and protect your disk drives.
3. Install IBM i as explained in *System i Systems management Recovering your system*, SC41-5304-09 (V6R1).
4. If recovering the system from an option 21 save, perform an option 21 restore at this point using the following commands (prompt for the commands):

   ```
   RSTUSRPRF
   RSTCFG OBJ(*ALL)
   RSTLIB SAVLIB(*NOSYS)
   RSTDLO DLO(*ALL) FLR(*ANY)
   RST ((‘/*’) (‘/QSYS.LIB’ *OMIT) (‘/QDLS’ *OMIT))
   ```

   **Note:** If restoring one or more IASPs, perform the RSTAUT command after all data in all the IASPs are restored.

5. Create the independent storage pools from Systems Director Navigator.
6. Restore the independent ASPs with the following commands:

   ```
   RSTLIB SAVLIB(*NONSYS) ASPDEV(yourIASP)
   RST OBJ(‘/dev/*’)
   ```

7. Repeat step 6 for all independent ASPs.
8. Perform the authority restore:

   ```
   RSTAUT USRPRF(*ALL)
   ```

**Note:** If the RSTAUT command was performed prior to recovering the IASPs, use the RSTUSRPRF command first and then the RSTAUT command.

## 7.5.1 Recovering an IASP after losing the system ASP

When installing LIC on a failed system ASP, the Install Licensed Internal Code display appears, as shown in Figure 7-1.

```
                    Install Licensed Internal Code (LIC)

Disk selected to write the Licensed Internal Code to:
   Serial Number Type Model I/O Bus Controller Device
   xx-xxxxxxx xxxx xxx x x x

Select on of the following:
   1. Restore Licensed Internal Code
   2. Restore Licensed Internal Code and Initialize System
   3. Restore Licensed Internal Code and Recover Configuration
   4. Restore Licensed Internal Code and Restore Disk Unit Data
   5. Restore Licensed Internal Code and Upgrade Load Source
```

*Figure 7-1 Install Licensed Internal Code (LIC) display*

To install LIC and recover an IASP, select option 3 from the Install Licensed Internal Code menu. When installing the LIC using option 3 from the Install Licensed Internal Code menu, the system:

- Clears disk unit 1, which contains information about how all the other disk units on the system are configured.
- Prepares to delete all data in the system ASP. The system ASP is not actually cleared until the initial program load (IPL) after the LIC in installed.

Every disk unit on the system contains information about how it is configured. Dedicated Service Tools (DST) provides an option to recover the disk configuration on the system by using this information. The system reads every disk, assigns it to the correct ASP, and rebuilds the disk configuration information in unit 1. In many cases, the disk configuration can be recovered and the ASPs reloaded. To recover the disk configuration, follow these steps:

1. After installing the LIC, the disk configuration error report display appears on the A or B mode IPL. Type 5 in the option column (OPT) to see the Missing Disk Configuration display. From either display, press F3 (Exit) to use DST.
2. Sign on to DST. The system displays the Use Dedicated Service Tools menu. If using logical partitions, and the primary partition is to be recovered, the LPAR recovery process should cater for the multiple restore operations.
3. From the Use Dedicated Service Tools (DST) menu, select option 4 (Work with disk units).
4. From the Work with Disk Units menu, select option 2 (Work with disk unit recovery).
5. From the Work with Disk Unit Recovery menu, select option 5 (Recover disk configuration).

6. Check the configuration of disk units on the display. The display shows the disk units that are assigned to each user ASP and to the system ASP (ASP 1). The warning message on the display means that the system clears all data on disk units in the system ASP.

   If this configuration *is not correct*, contact a service representative or software support for assistance. Do not proceed further without getting help.

   If the configuration that shown *is correct*, press F10 to confirm the configuration. The system builds the configuration information and returns to the DST menu.

8. Press F12 to cancel the DST menu. The IPL or Install the System menu is displayed.

At this point, continue with the normal recovery of the system ASP, including IBM i, user profiles, configuration, library, documents, and IFS data. Wait until all IASP are recovered before running the RSTAUT command.

After the system ASP is recovered, perform the following steps using Systems Director Navigator:

1. From the System Director Task List, expand **i5/OS Management**.
2. Select **Configuration and Service**.
3. Select the **Disk Pools** link.
4. Open a disk pool. If a problem is detected, `Recover Configuration` or `Recover Unknown Configuration Source` appears in the list. If you see either of these options, select it to continue.
5. Follow the instructions on the window that opens.

**Note:** This is only an example recovery. Your own recovery may differ significantly. You should test your recovery before moving to production.

## 7.5.2 Saving and restoring Linux network storage space (NWSSTG) in an IASP

When a NWSSTG is created on an IASP, it still creates its pointers under the QFPNSSTG directory in the system ASP. This means that to save and restore a network storage space that was created in an IASP, the information under QFPNWSSTG in the system ASP must be saved. The NWSSTG named under the `/dev/IASPname` directory must also be saved. The following command is an example for creating a Linux storage space:

```
CRTNWSSTG NWSSTG(LINUXSTG) NWSSIZE(3000) FORMAT(*OPEN) ASPDEV(IASPA)
```

This command creates a NWSSTG called LINUXSTG in IASPA. The actual storage space resides in `/dev/iaspa/linuxstg.udfs`. It also creates an entry in the `/qfpnwsstg/linuxstg/qfpcontrol` along with a `/mount` directory under `/qfpnwsstg/linuxstg`.

To use a NWSSTG, you must also create a network storage description to link to the storage space. The following commands are examples for creating a Linux network storage description and adding the link to the NWSSTG:

```
CRTNWSD NWSD(LINUXSVR) RSRCNAME(*NONE) TYPE(*GUEST) PARTITION(LINUX) +
ADDNWSSTGL NWSSTG(LINUXSTG) NWSD(LINUXSVR)
```

For more information about creating partitions for Linux on IBM i see *Implementing POWER Linux on IBM System i Platform*, SG24-6388.

If the network storage space was created on a switchable IASP, you can save the pointers and the network server description and restore it to the other system in the cluster. By switching the IASP, Linux can be usable on the other system. The following commands save the objects needed for Linux from System A, place them in a save file in the switchable IASP, restore the objects to system B, and then allow Linux to be active in the partition on system B:

```
CRTLIB LIB(MYLIB) ASP(*ASPDEV) ASPDEV(IASPA)
SETASPGRP ASPGRP(IASPA)
CRTSAVF FILE(MYLIB/SAVEFILE1)
CRTSAVF FILE(MYLIB/SAVEFILE2)
SAV DEV('/IASPA/QSYS.LIB/MYLIB.LIB/SAVEFILE1.FILE') OBJ(('/QFPNWSSTG/LINUXSTG'))
SAVCFG DEV(*SAVF) SAVF(MYLIB/SAVEFILE)
SETASPGRP ASPGRP(*NONE)
```

Note that system B must already have a Linux partition configured, but the network server space and the network server description are restored from system A.

After switching the IASP to system B, you can use the following commands to make Linux usable on system B (assuming that the partition already exits):

```
SETASPGRP ASPGRP(IASPA)
RST DEV('/IASPA/QSYS.LIB/MYLIB.LIB/SAVEFILE1.FILE') OBJ(('/QFPNWSSTG/LINUXSTG'))
RSTCFG OBJ(LINUXSVR) DEV(*SAVF) OBJTYPE(*NWSD) SAVF(MYLIB/SAVFILE2)
SETASPGRP ASPGRP(*NONE)
```

It is important to restore the network server description after the network server storage space, since this automatically links the network server description to the network server storage space. At this point, varying on the network server storage description should bring up the Linux partition.

## 7.6 Using BRMS with IASPs

Auxiliary disk pool devices are specified in IBM i operations by the 10-character device description name rather than by specifying an identifying number. BRMS also uses the name when targeting saves or restores to ASP devices. Refer to the system documentation if you are not familiar with ASP device operations or terminology.

## 7.6.1 Backing up ASP devices

The Edit Backup Control Group Entries display has been updated (as shown in Figure 7-2) to include a new Auxiliary storage pool device field. This field does not appear on some backup items entries. Typically, this occurs for backup items that cannot reside on ASP devices. The auxiliary storage pool device prompt is automatically filled in for entries of your existing backup control groups to reflect the scope of the save across ASP devices. These default values should not affect your current backup strategy and should be consistent with what is saved by the control group.

```
                  Edit Backup Control Group Entries              AAAAS400

Group . . . . ..... . : SYSTEM
Default activity  . . . . . FIIIIII
Text  . . . . . . . . . . . Backs up the entire system

Type information, press Enter.

                        Auxiliary   Weekly    Retain  Save      SWA
      Backup      List  Storage     Activity  Object  While     Message
Seq   Items       Type  Pool Device SMTWTFS   Detail  Active    Queue

____  __________  ____  __________
__10  *EXIT_____                    *******
__20  *SAVSYS___  ____              FIIIIII
__30  *IBM______  ____              FIIIIII   *NO_    *NO_____
__40  *ALLUSR___  ____  *SYSBAS     FIIIIII   *ERR    *NO_____
__50  *ALLDLO___                    FIIIIII   *NO_    *NO_____
__60  *LINK_____  ____  *ALLAVL     FIIIIII   *YES    *NO_____
__70  *EXIT_____                    *******

                                                                  Bottom
F3=Exit                    F5=Refresh              F10=Change item
F11=Display exits          F12=Cancel              F24=More keys
```

*Figure 7-2 Edit Backup Control Group Entries (Display 1)*

The *SYSBAS value on the *ALLUSR backup item saves all user libraries on the system (1) and any basic user (2 to 32) ASPs.

The *ALLAVL value for the *LINK backup items saves the links on the system (1) and any basic user (2 to 32) ASPs, as well as the links on all available ASP devices.

**Note:** When saving the directory and files, you should unmount any mounted user-defined file systems (UDFSs) before the save to ensure that the objects in the mounted over directories are saved. UDFSs are automatically unmounted on ASP devices when the system is in restricted state. UDFSs on the system or basic user ASPs must be explicitly unmounted. You must remount any unmounted UDFSs after the save.

Restrictions of the operating system prevent you from attaching a primary ASP to the current job if the system is in restricted state. The only way that you can save the libraries on ASP devices while in restricted state is to specify the device name in the auxiliary storage pool device prompt for the backup item.

Assume that the system is configured with an ASP group consisting of a device named PAYROLL serving as the primary ASP device and ACCOUNTS serving as the secondary ASP device. Figure 7-3 shows that to perform a save of all user libraries, three separate *ALLUSR backup items are required. One *ALLUSR with *SYSBAS is specified to save the libraries on the system and basic user ASPs, and one *ALLUSR for each of the PAYROLL and ACCOUNTS ASP devices.

```
                  Edit Backup Control Group Entries              AAAAS400

Group . . . . . . . . . . . : SYSTEM
Default activity  . . . . . . FIIIIII
Text  . . . . . . . . . . . . Backs up the entire system

Type information, press Enter.

                        Auxiliary   Weekly    Retain  Save      SWA
      Backup      List  Storage     Activity  Object  While     Message
Seq   Items       Type  Pool Device SMTWTFS   Detail  Active    Queue
____  __________  ____  __________  *******
__10  *EXIT_____                    *******
__20  *SAVSYS___  ____              FIIIIII
__30  *IBM______  ____              FIIIIII   *NO_    *NO_____
__40  *ALLUSR___  ____  *SYSBAS___  FIIIIII   *YES    *NO_____
__50  *ALLDLO___                    FIIIIII   *YES    *NO_____
__60  *LINK_____  ____  *ALLAVL___  FIIIIII   *YES    *NO_____
__70  *ALLUSR___  ____  PAYROLL___  FIIIIII   *YES    *NO_____
__80  *ALLUSR___  ____  ACCOUNTS__  FIIIIII   *YES    *NO_____
__90  *EXIT_____                    *******

                                                                 Bottom
F3=Exit                  F5=Refresh              F10=Change item
F11=Display exits        F12=Cancel              F24=More keys
```

*Figure 7-3 Edited Backup Control Group Entries (display 2)*

You may consider using this approach if the ASP's devices are configured as non-switchable disk pools. Notice that the *ALLUSR backup items directed to the PAYROLL and ACCOUNTS ASP devices are after the *LINK backup item. This was done to optimize the recovery because objects saved from ASP devices are always recovered after the objects on the system and basic user ASPs are recovered.

It may be that the ASP group is configured for use in a clustered environment where the group is automatically switched to an alternate system by the cluster management software when it detects that the primary system is no longer operational. It may be useful in this environment to set up a backup strategy that consists of two backup control groups—one to back up the base system and one to back up the switchable ASP group.

Figure 7-4 shows the backup control group entries that would be used on primary and alternate systems to save only the objects on the system and basic user ASPs. The *SYSBAS special value used for the auxiliary storage pool device parameter of the backup items restricts the scope of the saves to the system and basic user ASPs. The exception is the *SAVSYS backup item that saves the private authorities for all objects on all available ASP devices.

```
                  Edit Backup Control Group Entries              AAAAS400

Group . . . . . . . . . . . : SYSTEM
Default activity  . . . . . FIIIIII
Text  . . . . . . . . . . . Backs up entire system but not switched ASPs

Type information, press Enter.

                        Auxiliary   Weekly    Retain  Save      SWA
      Backup      List  Storage     Activity  Object  While     Message
Seq   Items       Type  Pool Device SMTWTFS   Detail  Active    Queue

____  __________  ____  __________  *******
__10  *EXIT                         *******
__20  *SAVSYS     ____              FIIIIII
__30  *IBM        ____              FIIIIII   *NO_    *NO_____
__40  *ALLUSR     ____  *SYSBAS___  FIIIIII   *YES    *NO_____
__50  *ALLDLO                       FIIIIII   *NO_    *NO_____
__60  *LINK       ____  *SYSBAS___  FIIIIII   *YES    *NO_____
__70  *EXIT                         *******

                                                                  Bottom
F3=Exit                    F5=Refresh              F10=Change item
F11=Display exits          F12=Cancel              F24=More keys
```

*Figure 7-4 Edit Backup Control Group Entries (display 3)*

The entries shown in Figure 7-5 are specified in a second backup control group that is scheduled to run on the primary system after the SYSTEM backup control group is run. It is not scheduled to run on the alternate system unless switchover of these ASP devices occurs. As you can see from the entries, only the user libraries and links on the ASP devices are saved by this backup control group.

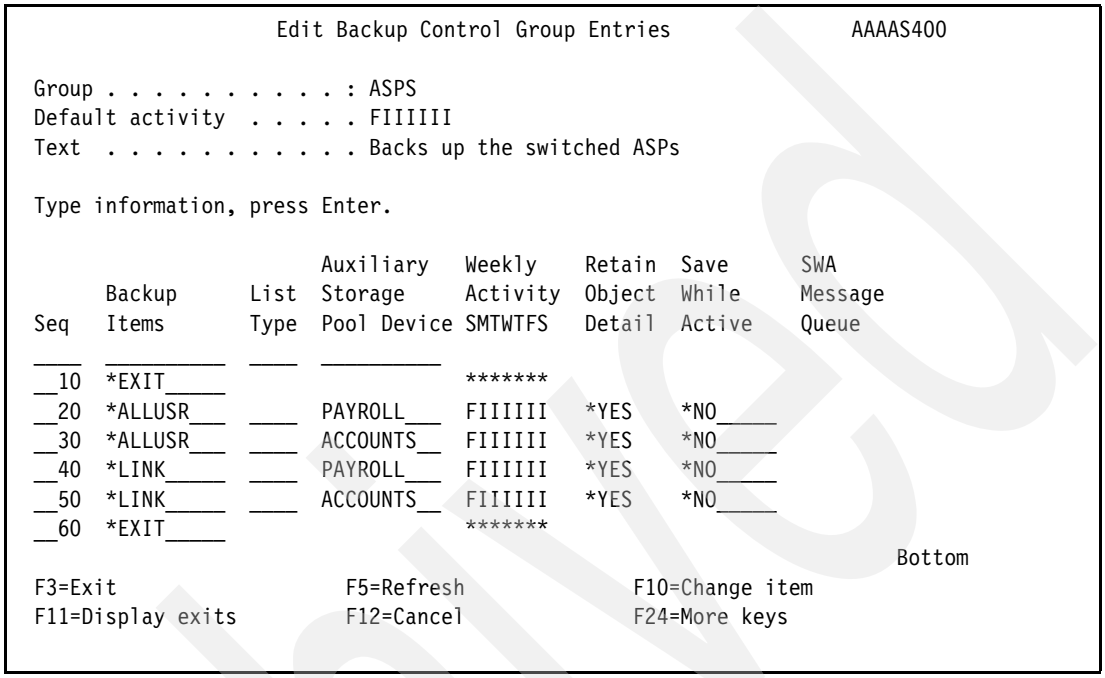


*Figure 7-5 Edit Backup Control Group Entries (display 4)*

**Note:** When using a switched ASP device, you must consider any private authorities that may be defined for objects of these types of devices. All user profiles and all private authorities for objects on ASP devices are saved with the *SAVSYS or *SAVSECDTA backup items. You must consider how to restore these private authorities if this authority changes while the device is switched to an alternate system. You may mean that as part of the recovery, where you restore the associated user profiles from the alternate system to the primary system before you restore any objects to the switched ASP device.

When running multiple backup control groups, you must consider the order in which these control groups run and which control group manages the restart of subsystems to return from the restricted state.

Another less obvious implication of switched ASP devices is when an incremental save is specified for the weekly activity. When the ASP device switches to the alternate system, the alternate system has no history of the switched objects. Therefore, the first save of the objects on the switched ASP devices is a full save rather than an incremental save.

Specifying the individual ASP devices in the previous example can be cumbersome, especially if you have a large number of secondary ASP devices in the group. If the objects on these ASP devices can be saved using save while active, there is no need to perform the save of these objects in restricted state. Also, you can use the *SETASPGRP backup item to attach the device to the job. To do so, add a *SETASPGRP backup item and specify the primary ASP device name in the auxiliary storage pool device prompt for this backup item. Then use the

*CURASPGRP special value for the auxiliary storage pool device prompt on all subsequent backup items to be backed up from that primary ASP device and any of its associated secondary ASP devices. This is shown in the Edit Backup Control Group Entries display in Figure 7-6.

**Notes:** Note the following:

- The *SETASPGRP backup item cannot run while in a restricted state.
- When more than one *SETASPGRP is used as backup items, any previously attached primary ASP is detached from the job when the subsequent *SETASPGRP backup item is processed.
- To explicitly remove an attached primary ASP from the job, use *SETASPGRP with *SYSBAS for the auxiliary storage pool device prompt.

```
                    Edit Backup Control Group Entries              AAAAS400

Group . . . . . . . . . . . : ASPS
Default activity  . . . . . FIIIIII
Text  . . . . . . . . . . . Backs up the switched ASPs using save-while-active

Type information, press Enter.

                         Auxiliary   Weekly    Retain  Save      SWA
      Backup      List   Storage     Activity  Object  While     Message
Seq   Items       Type   Pool Device SMTWTFS   Detail  Active    Queue
____  __________  ____ __________  *******
__10  *EXIT_____                   *******
__20  *SETASPGRP  ____   PAYROLL___  FIIIIII
__30  *ALLUSR___  ____   *CURASPGRP  FIIIIII   *YES    *YES____  *LIB______
__40  *LINK_____  ____   *CURASPGRP  FIIIIII   *YES    *YES____  *LIB______
__50  *SETASPGRP  ____   *SYSBAS___  FIIIIII
__60  *EXIT_____                   *******

                                                                      Bottom
F3=Exit                   F5=Refresh             F10=Change item
F11=Display exits         F12=Cancel             F24=More keys
```

*Figure 7-6 Edit Backup Control Group Entries (display 5)*

## 7.6.2 Recovery of ASP devices

When your system recovery includes objects saved from ASPs devices, additional steps are added to the system recovery report to aid you in recovering your system. These steps appear after the words `Step: Perform IPL` because Systems Director Navigator is used to reconfigure the ASP devices. These steps do not appear if you are recovering by ASP. The assumption is that the ASP that you are recovering is already configured. The step shown in Figure 7-7 provides the guidance to configure the ASP devices during system recovery.

```
********************************************************************************
STEP: Configure auxiliary storage pool devices

   Start date/time _________________  Stop date/time _________________  Duration

         The recovery includes objects saved from auxiliary storage pool
         devices. The auxiliary storage pool devices must be deleted and
         re-configured before recovery can continue.

         Use the following command to view the restored auxiliary storage
         pool device descriptions:

            WRKDEVD DEVD(*ASP)

         Use option 4=Delete to remove the device descriptions from the system.

         Use the "Display ASP Information" report (QP1AASP) to review the
         names of the auxiliary storage pool devices of the saved system.

         --- Attention ------------------------------------------------------
         If the auxiliary storage pool devices are intended to be used as
         switched auxiliary storage pools in a clustered environment, review
         the documentation provided with your cluster management product to
         identify any pre-configuration steps that may be required prior to
         configuring these auxiliary storage pool devices.
         --------------------------------------------------------------------

         Re-configure the auxiliary storage pool devices using iSeries
         Operations Navigator referring to the online disk unit help.
********************************************************************************
```

*Figure 7-7 Configuring ASP devices verification*

The step shown in Figure 7-8 provides the guidance to verify the names of the configured ASP devices to the names of the saved ASP devices. When you configure the ASP devices using Systems Director Navigator, these devices may be assigned the same ASP number depending on the number of disk resources and the order in which they are configured. Since the ASP number may change, it is important to keep the name the same to allow for automated recovery.

If you are restoring objects to a different ASP device than the one from which it was saved, you must change the default recovery options to specify the ASP where the objects are to be restored.

```
********************************************************************************
STEP: Verify auxiliary storage pool device names

   Start date/time ___________________  Stop date/time ___________________  Durati

         Use the following command to view the configured auxiliary storage
         pool device descriptions:

            WRKDEVD DEVD(*ASP)

         Verify the configured auxiliary storage pool device names match the
         names of the auxiliary storage pool devices of the saved system.

         Use option 8=Work with status to verify each auxiliary storage pool
         device status is AVAILABLE.
********************************************************************************
```

*Figure 7-8 Verifying ASP device names*

BRMS saves the history information for ASP devices by the ASP device name. This device name is displayed adjacent to the saved item name, as shown in the sample report in Figure 7-9. The step shown is included in the report if libraries are saved from ASP devices. You may also see a similar section for recovery of any directories and files that may be saved from ASP devices.

```
********************************************************************************
STEP: Recover additional user libraries

   Start date/time _________________  Stop date/time _________________  Duration

You should restore the current version of your user libraries saved
         from auxiliary storage pool devices or saved to TSM servers.

         If you are performing a complete system restore, run the following
         command to continue:

            STRRCYBRM OPTION(*RESUME)

         Otherwise, run the following command(s):

            STRRCYBRM OPTION(*ALLUSR) ACTION(*RESTORE) USEADSM(*YES)
            STRRCYBRM OPTION(*ASP) ACTION(*RESTORE) ASP(auxiliary-storage-pool-n

         Select and recover the following saved item(s) on the "Select Recovery
Items" display using the specified volume(s).

         --- Attention ------------------------------------------------------
         If you have logical files whose based-on physical files are in a
         different library, you must restore all based-on physical files
before you can restore the logical file.
         --------------------------------------------------------------------

         --- Attention ------------------------------------------------------
         If you use journaling, the libraries containing the journals must be
restored before restoring the journaled files.
         --------------------------------------------------------------------
   Saved           ----- ASP ------  Save     Save                  Sequence  C
   Item       Type  Name      Number  Date     Time    Objects Omit  Number    G
   ---------- ----- ---------- ----- -------- -------- ------- ----- --------- -
__ EMPLOYEES  *FULL PAYROLL    00033  8/22/02 13:09:52       6               1 S
__ AUG2002    *FULL ACCOUNTS   00034  8/22/02 13:21:57       6               2 S
********************************************************************************
```

*Figure 7-9 Recovering additional user libraries*

When recovering libraries on ASP devices and recovering these to the same system and the same ASP device, consider these points:

- When you vary on an ASP device, library QSYS200*nnn* (where *nnn* is the ASP number) is created on the ASP device as well as several objects.
- Recover the saved QSYS200nnn library to the ASP device before you recover the other libraries on that ASP device.

- When you recover the saved QSYS200nnn library, use F9 (Recovery defaults) from the Select Recovery Items display. You must also set the Allow object differences prompt on the Restore Command Defaults panel to *ALL.
- When you recover the remaining libraries, use F9 (Recovery defaults) from the Select Recovery Items display. You must also set the Allow object differences prompt on the Restore Command Defaults panel to *FILELVL.

## 7.7 Recovering an independent disk pool

If you experience problems accessing an independent disk pool or making it available, there may be a problem with the disk pool. This section explains a couple of possible problems.

First, the configuration source may be corrupted. When corruption occurs, the independent disk pool appears to have no disk units in it. If this occurs, you can select to recover the configuration information on the configuration source. Recovering the configuration attempts to determine the original configuration and recover it. During this process, the dependent disk pool may need to be cleared, destroying all data on the disk units in the pool. If the disk pool must be cleared, a message appears that warns you of this and allows you to cancel the recovery.

Also, the mirrored disk unit of the configuration source may be damaged. When this happens, the mirrored configuration source becomes unknown. The disk pool is unavailable, and you must recover the configuration information of an unknown configuration source before you make it available. You should only attempt to recover the state of the unknown configuration source when you know that its mirrored disk unit was active before the failures that caused the state to become unknown.

To attempt to recover an independent disk pool, follow these steps:

1. From the System Director Task List, expand **i5/OS Management**.
2. Select **Configuration and Service**.
3. Select the **Disk Pools** link.
4. Open a disk pool. If a problem is detected then `Recover Configuration` or `Recover Unknown Configuration Source` appears in the list. If you see either of these options, select it to continue.
5. Follow the instructions on the window that opens.

8

# Hardware configuration examples

Independent disk pools (independent auxiliary storage pools, IASPs) allow you to configure multiple disk storage units into one group. The group can include:

- Non-switchable independent disk pool
  - Any or all units in the system tower, except the load source
  - Any one or more disks anywhere on the system
- Switchable independent disk pool
  - All the units in one tower
  - A subset of the disk units under a single IOP

To illustrate this concept, this chapter depicts possible independent disk pool configurations.

## 8.1 Protecting independent disk pools

The following protect independent disk pools:

- Device parity protection (RAID)

  RAID is accomplished at the hardware level. When planning the disk units to be used in disk pools, if possible, assign the disk pool within a single RAID set. Thus, only that disk pool is affected by disk failures. Many different configurations are possible, depending on the available non-configured disks. Multiple independent disk pools can reside within one RAID set or one independent disk pool can span multiple RAID sets.

  **Note:** RAID does not provide protection from disk controllers or bus failures.

- Mirroring

  Mirroring the disk in a disk pool provides the most protection from hardware failures. With mirroring, disks with the same capacity are paired. Depending on the available hardware, mirroring can be accomplished up to and including the bus level.

## 8.2 Non-switchable independent disk pools

Non-switchable independent disk pools are available in any system that supports IBM i 6.1. This provides a great deal of flexibility for testing and implementing an independent disk pool.

Figure 8-1 shows two independent disk pools residing in a tower. However, any one or more disks anywhere on the system can be used to create an independent disk pool. This configuration may be used where a separate database is required but the ability to switch is not necessary. IASP #1 and IASP #2 can be varied on/off and accessed independently of each other if they reside in different ASP groups. If, however, they reside within the same ASP group, they are varied on/off and accessed as one entity using the ASP group name.

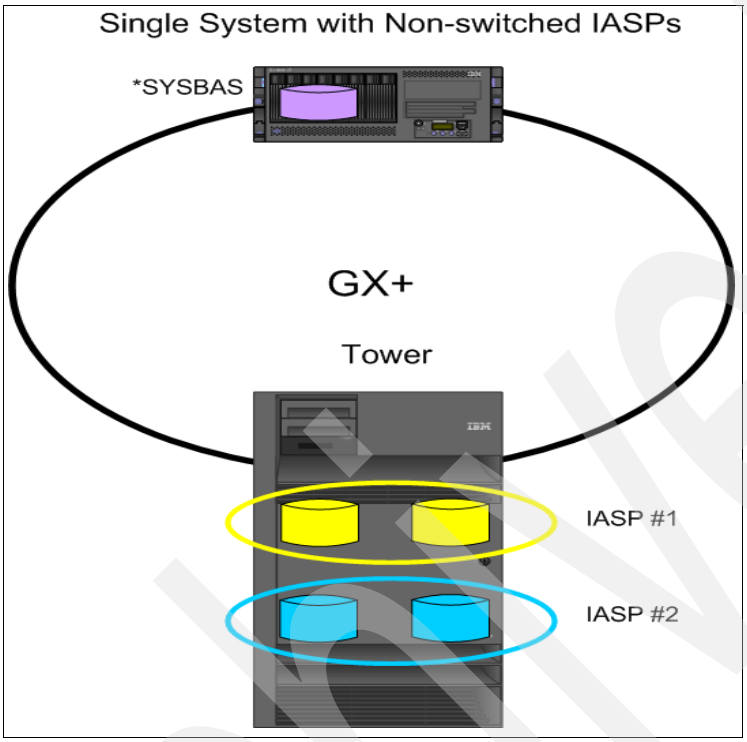


*Figure 8-1 Single system with non-switched disks*

IBM Systems Director for i5/OS provides the ability to protect the disks within the independent disk pools. Mirroring, device parity protection, or no protection options are available. Usually, disks are protected from dedicated service tools. If, however, you have some unprotected disks, you can still create an independent disk pool and protect it without disrupting the normal system operations.

## 8.3 Switchable independent disk pools

Switchable independent disk pools apply to multi-partitioned Power Systems server configurations. The switchable resource can be as small as an IOP and the connected I/O devices or as large as an entire tower. All devices packaged in a single tower, such as workstation controllers, Ethernet cards, tape drives, and CD-ROM drives are also switched. The user is responsible for performing the necessary device configuration and varying on the non-disk devices.

The example shown in Figure 8-2 illustrates two switchable independent disk pools residing in a tower. The Power Systems server has two LPARs. If both independent disk pools are under the same IOP, both are switched at the same time. If they are configured on separate IOPs and in different ASP groups, they can be switched independently of each other.

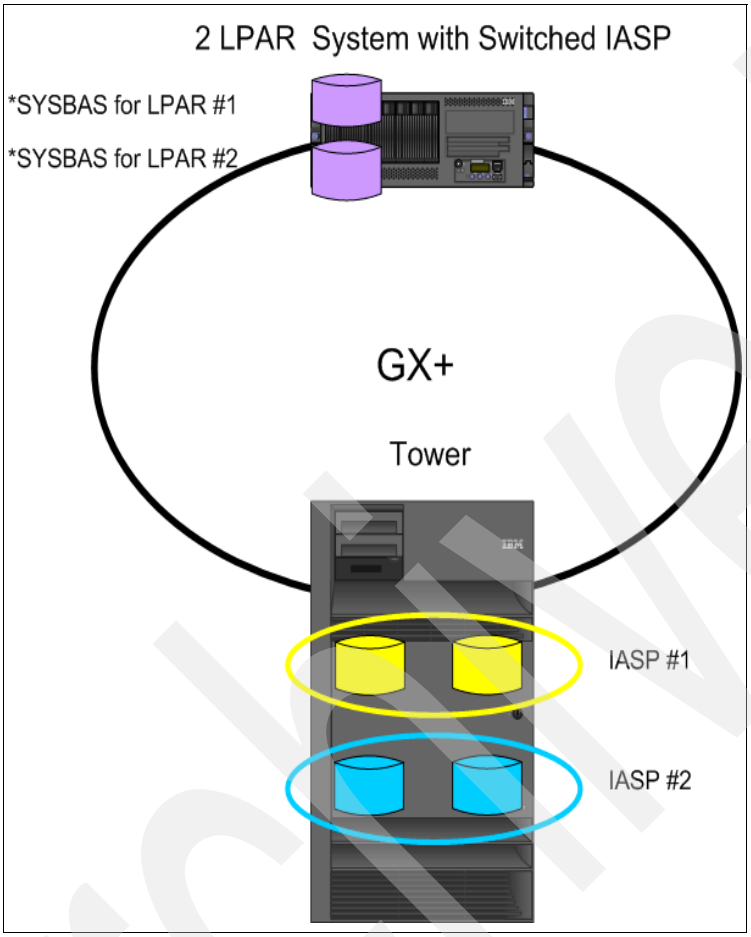


*Figure 8-2 LPAR system with a switched tower*

## 8.4 PowerVM, virtual disk, and independent disk pool considerations

The following items should be considered if you are creating independent disk pools when using virtual disk drives.In some cases Systems Director Navigator for i5/OS will allow you to complete a task. However, the resulting level of protection may not be satisfactory.

### 8.4.1 Creating independent disk pools using virtual disks

All that is needed to implement non-switchable independent disk pools is one or more physical disk units or virtual drives. The disks can reside anywhere within the system. When using virtual disks, however, there are some restrictions when protecting disk pools.

**Note:** Switchable independent disk pools are not supported when using virtual disks.

### 8.4.2 Restrictions when using virtual disks

Protection for disk pools is provided by Systems Director Navigator for i5/OS when the disk pool is created. The options are device parity protection, mirroring, and no protection. Care should be exercised when virtual disks are being used:

- Device parity protection

  Since device parity protection is performed at the hardware level, virtual disks do not provide for device parity protection on the virtual level.

- Mirroring

  One physical disk drive may be divided into multiple virtual disks. If this is the case, care should be taken not to select virtual disks on the same physical disk drive as a mirroring pair, as this would defeat the concept of mirroring.

9

# Creating a switchable disk pool (IASP)

One of the many uses for independent auxiliary storage pools (IASPs) is switching between LPARS. This chapter provides an example of creating a switchable I/O controller between LPARs in a single system, which implies that all disks and other I/O attached to the controller will also be switched.

Refer to the IBM Redbooks publication I*mplementing PowerHA™ for IBM i,* SG24-7405, for a complete discussion of switching and replicating hardware resources.

## 9.1 Prerequisites for creating switchable disk pools (IASPs)

You must complete the following requirements:

- The non-configured disk units that make up the new IASP must be owned by a system.
- You must create the I/O pool using the HMC.
- Virtual OptiConnect (5761-SS1, Option 23) is required for LPAR-to-LPAR switching. Configure using the HMC.
- IBM Systems Director Navigator for i5/OS must have connections to the systems that are using the independent pools.
- Product Option 41 - IBM i - HA Switchable Resources (Licensed Program Product 5722-SS1) is required when you set up simple clustering, for switching between two systems.
- Cluster resource services are used to switch independent disk pools between multiple systems. For a full explanation of clustering, consult *Implementing PowerHA on IBM i,* SG24-7405, or search the iSeries Information Center on the Web at:

  `http://publib.boulder.ibm.com/iseries/`

The steps are:

1. Start the TCP server *inetd on all nodes in the cluster. We recommend making the *inetd TCP/IP server start when TCP is started. To automatically start *inetd when TCP starts, go to the IBM Systems Director Navigator for i5/OS.

   a. Select **TCP/IP Servers**.

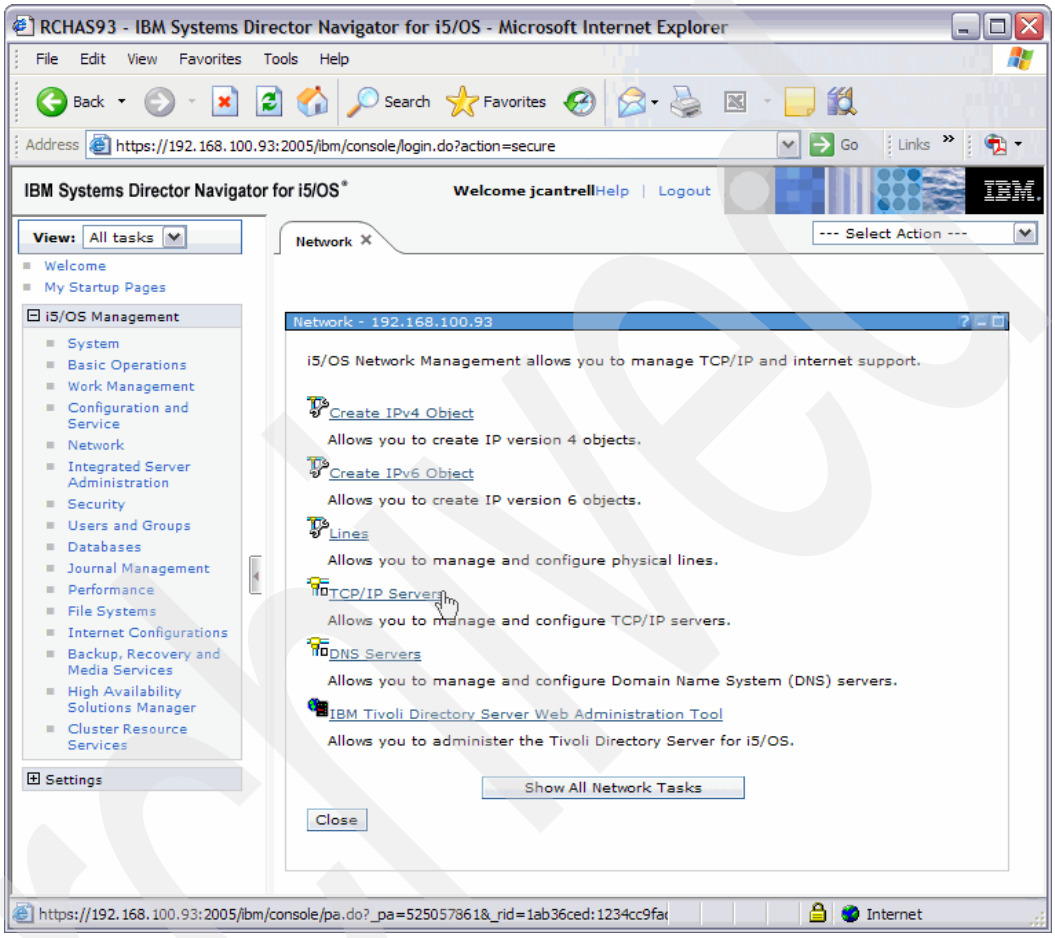


*Figure 9-1 IBM Systems Director Navigator → expand i5/OS management → TCP/IP Servers*

b. Check the properties of the INETD server (Figure 9-2).

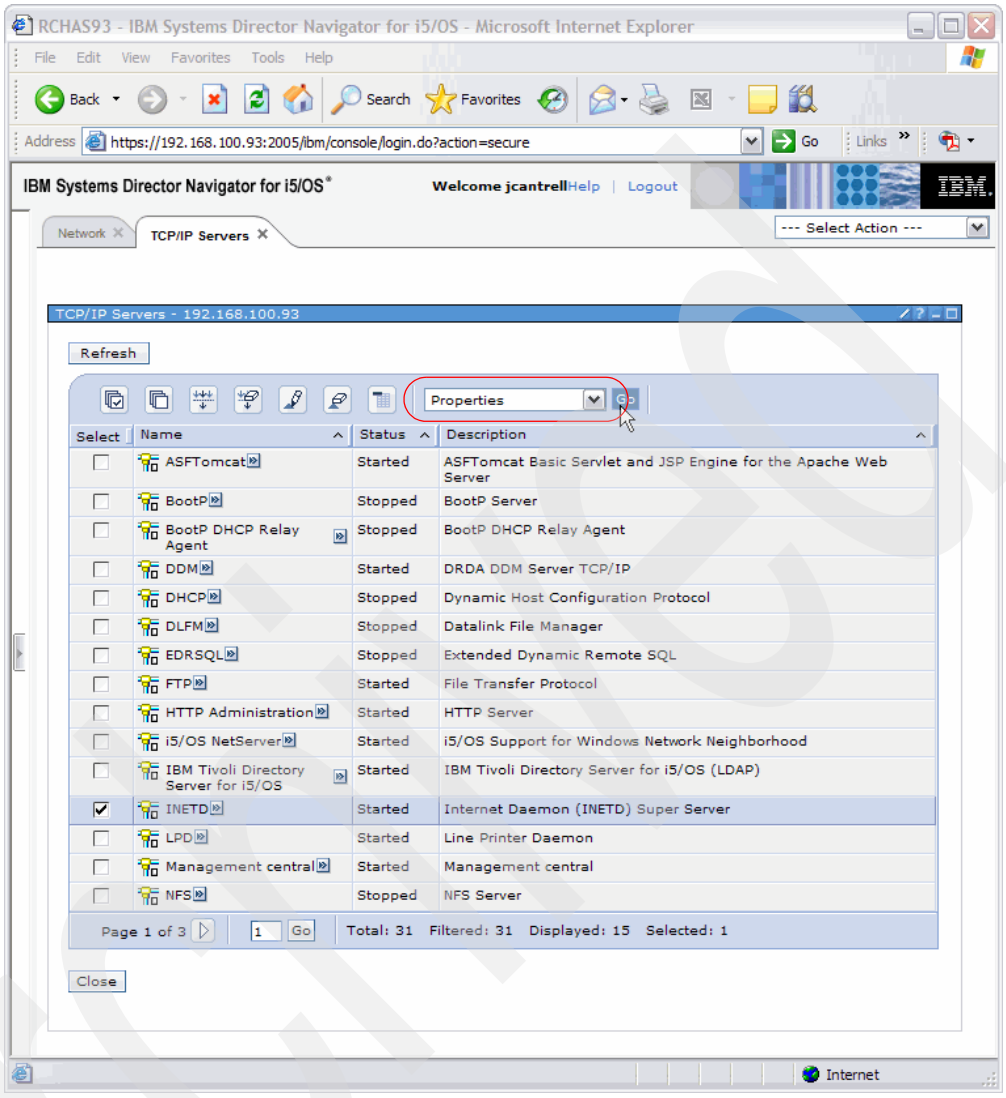


*Figure 9-2 TCP/IP Servers → Select INETD → Properties → Go*

c. Select the box to indicate start when TCP starts (Figure 9-3).

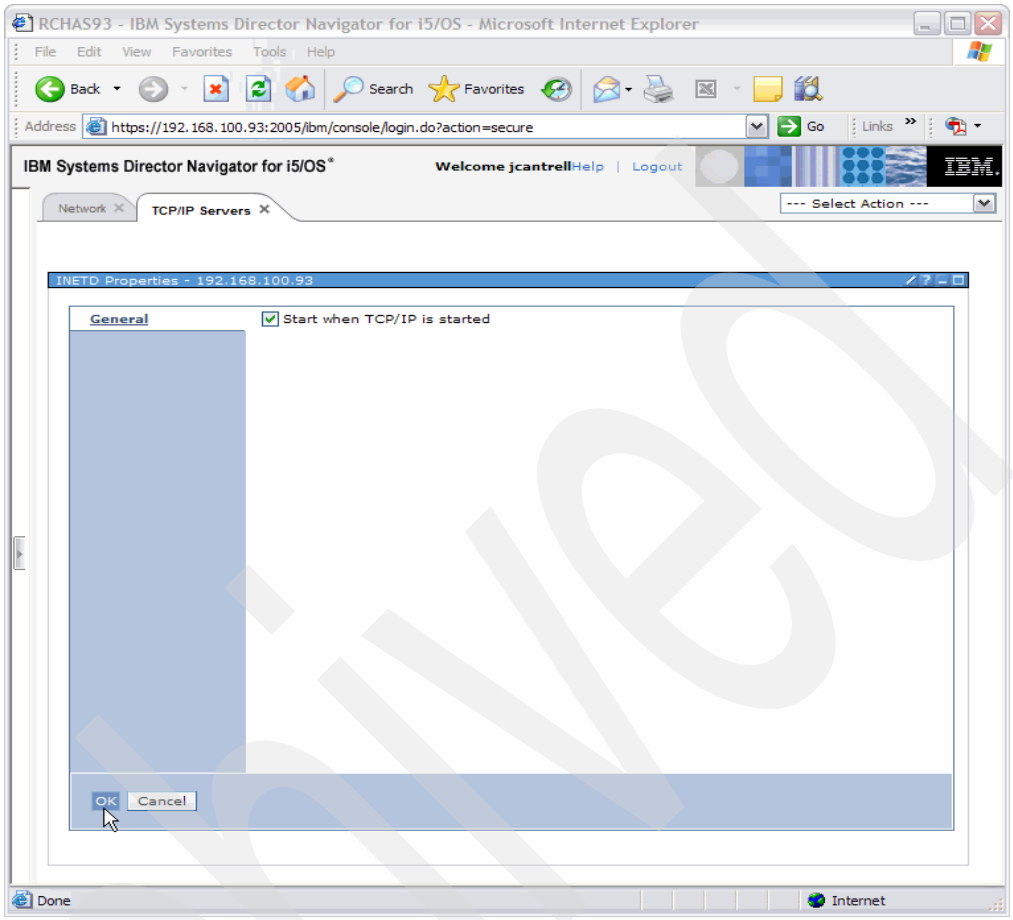


*Figure 9-3 TCP/IP Servers → select INETD → Properties → check start when TCP started → OK*

2. Before you can add a node to a cluster, you must set an appropriate value for the Allow add to cluster (ALWCLUADD) network attribute. DSPNETA shows the current ALWCLUADD parameter. You must have *IOSYSCFG authority to change the attribute. The system is shipped with a default of *NONE. Using one of the following two parameters will allow your system/node to be part of a cluster (Figure 9-4):
   - *ANY will allow any other system to add this system as a node in a cluster.
   - *RQSAUT will allow any other system to add this system as node in a cluster only after the cluster add request has been authenticated.

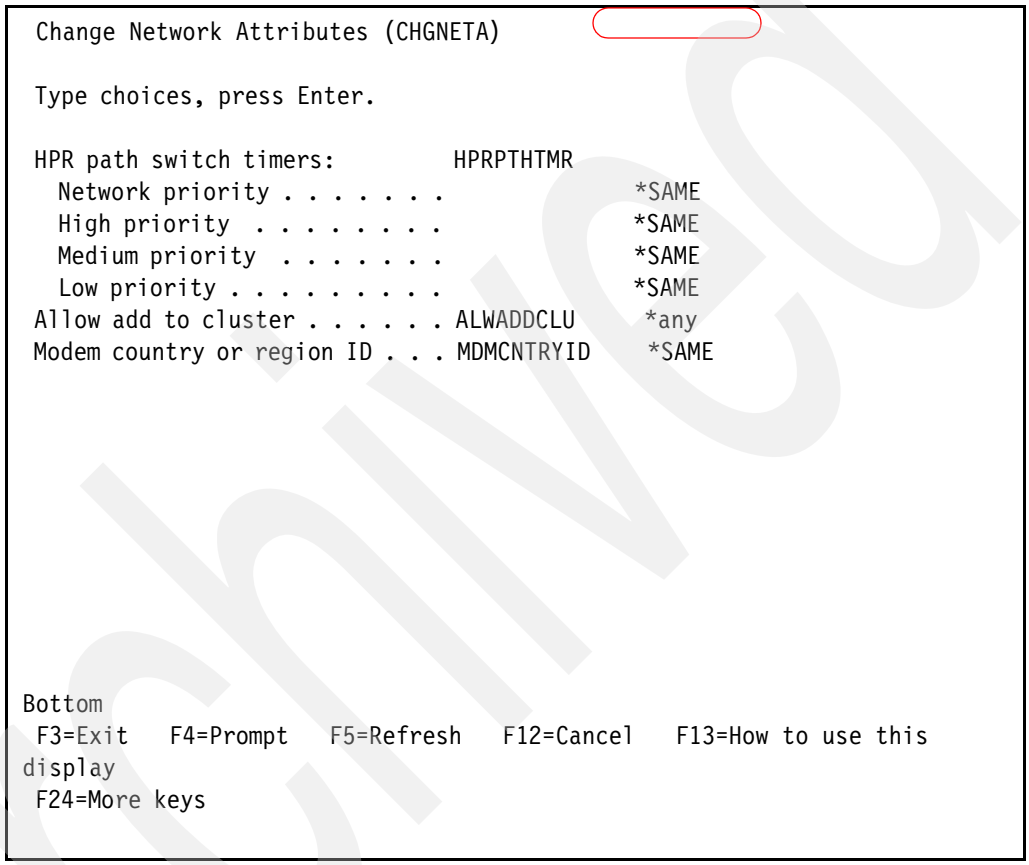
```
 Change Network Attributes (CHGNETA)

 Type choices, press Enter.

 HPR path switch timers:          HPRPTHTMR
   Network priority . . . . . .                *SAME
   High priority  . . . . . .                  *SAME
   Medium priority  . . . . . .                *SAME
   Low priority . . . . . .                    *SAME
 Allow add to cluster . . . . . ALWADDCLU       *any
 Modem country or region ID . . . MDMCNTRYID    *SAME

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

*Figure 9-4 CHGNETA → ALWCLUADD → *ANY or *RQSAUT*

## 9.2 Creating a switchable IASP between LPARs of a system

The multiple LPAR, single-system, switchable IOP configuration is supported for any Power Systems server that allows logical partitioning. We also need option 23 - OptiConnect. Configure (using the HMC) each LPAR to use virtual OptiConnect.

For more information about iSeries LPARs see *LPAR Simplification Tools Handbook*, SG24-7231.

This section shows you how to create a simple two-node (primary/backup) cluster using IBM Director Navigator for i5/OS. The TCP server *INETD must be running, and Option 41 of 5722-SS1 must be installed with its license key applied to continue. Follow these steps:

1. Log into **Director Navigator → i5/OS Management → Cluster Resource Services**. Click **New Cluster** (Figure 9-5).

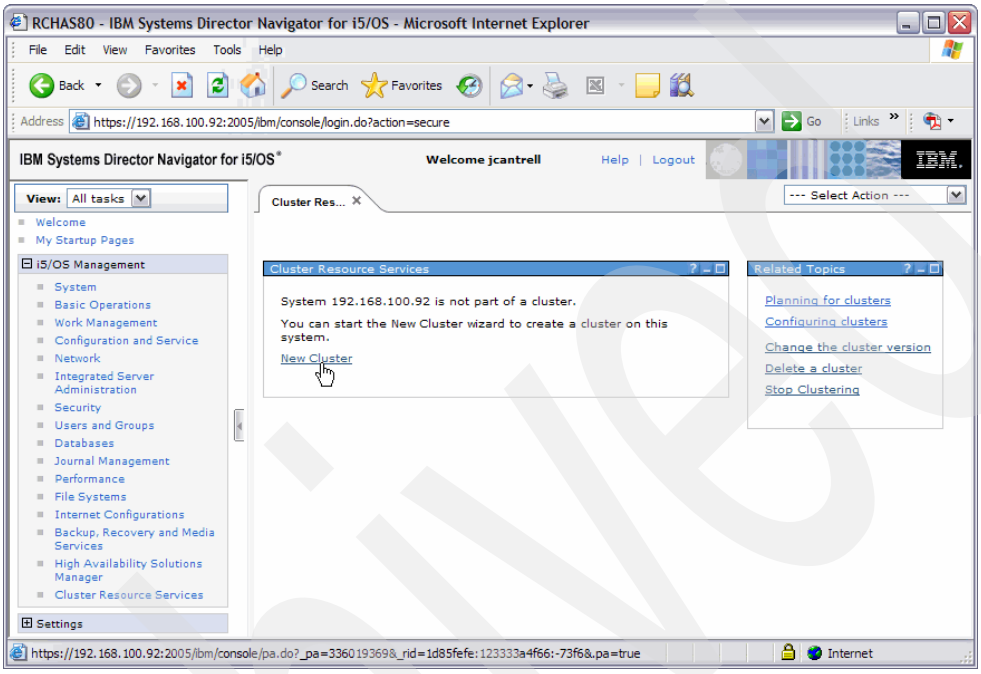


*Figure 9-5 i5/OS Management → Cluster Resource Services → New Cluster*

2. The New Cluster Welcome window appears. Select **Next** (Figure 9-6).

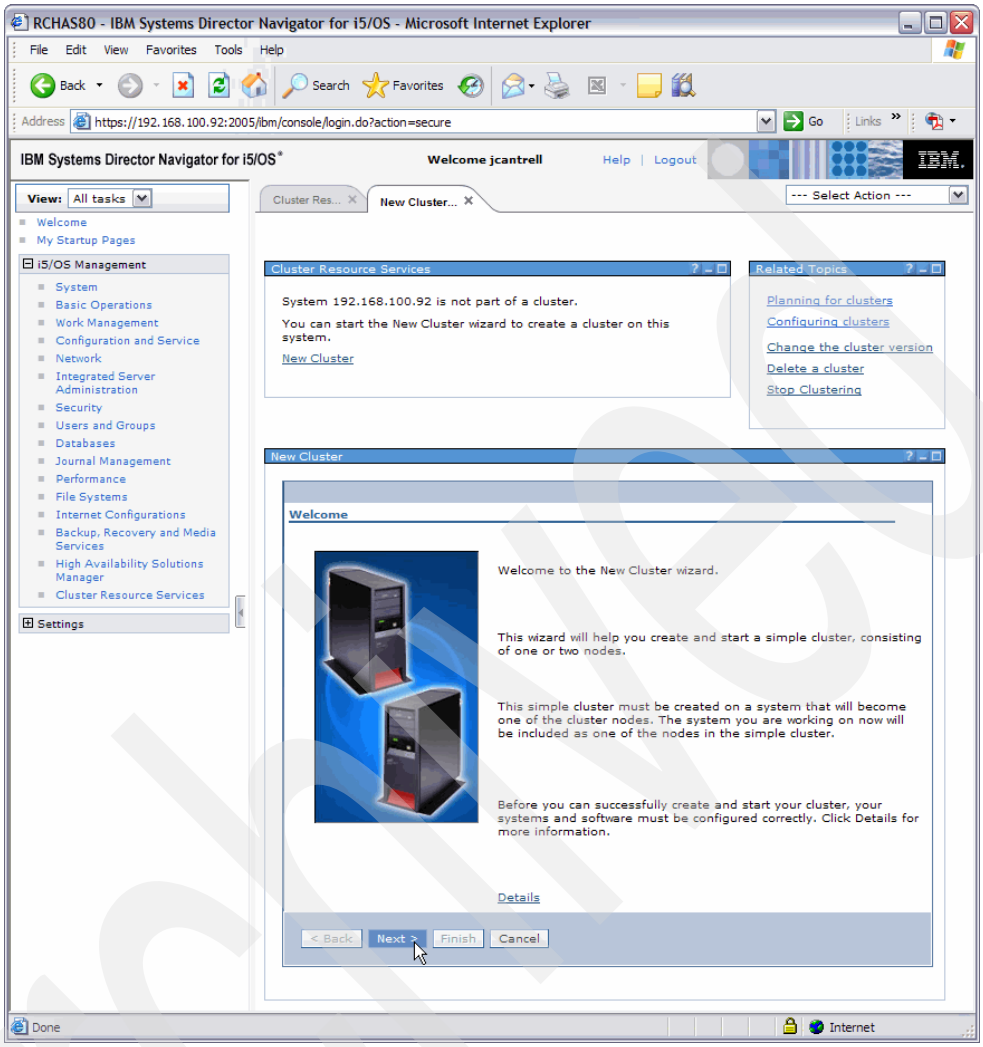


*Figure 9-6 New Cluster → Next*

3. Select **Clusters**, right-click, and select **New Cluster**, as shown in Figure 9-7.

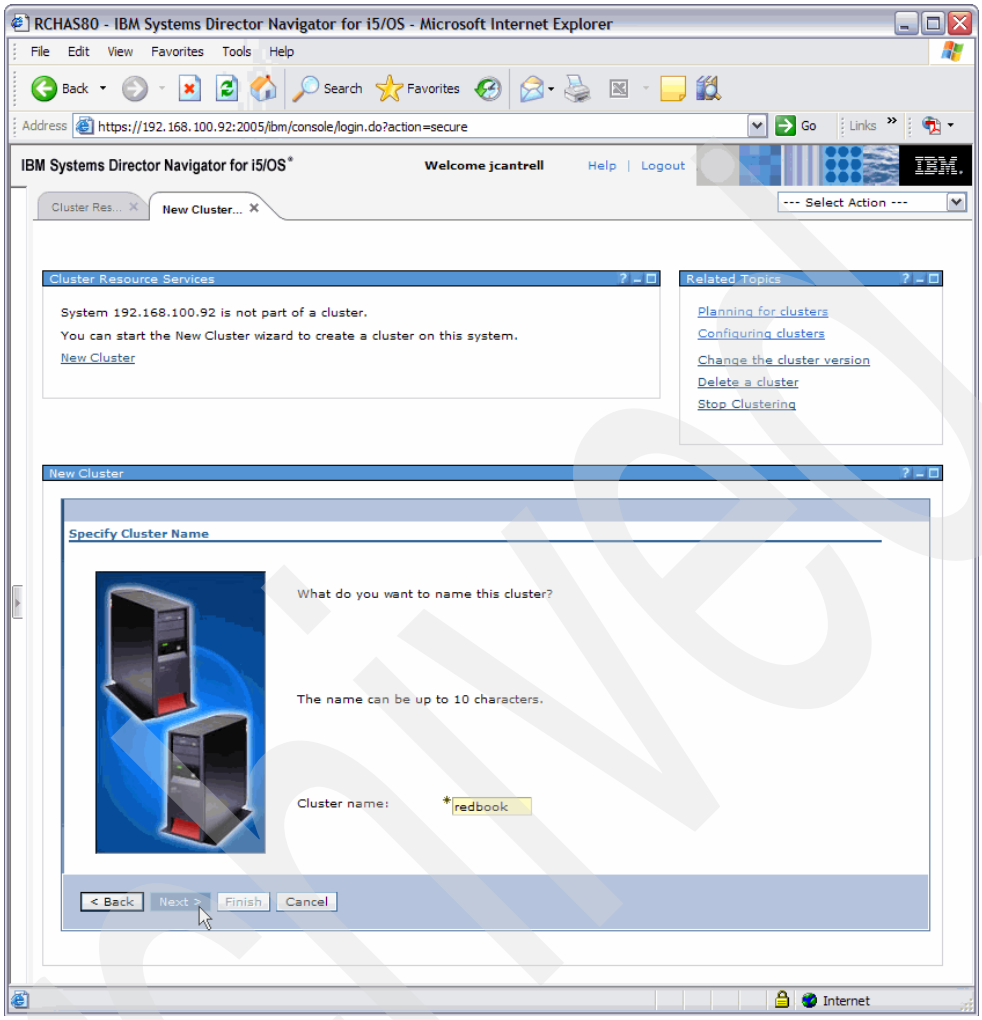


*Figure 9-7 New Cluster → Cluster Name → Next*

4. When specifying the node you must enter the node name (Figure 9-8). This name should match the node name in the network attributes. We recommend that the node name be the same as the system name or the host name.

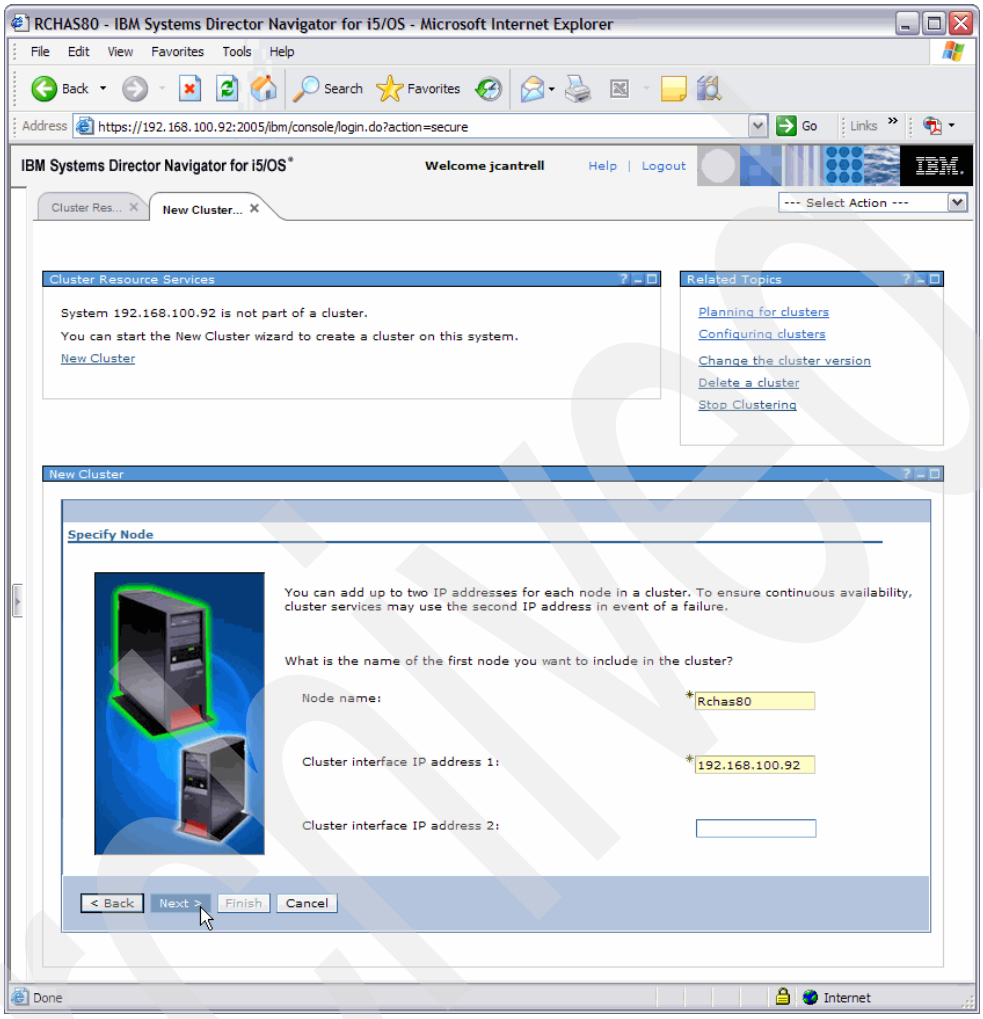


*Figure 9-8 New Cluster → Specify Node → Next*

5. You will be prompted to add an additional node. Enter the name of the node and the appropriate IP address (Figure 9-9).

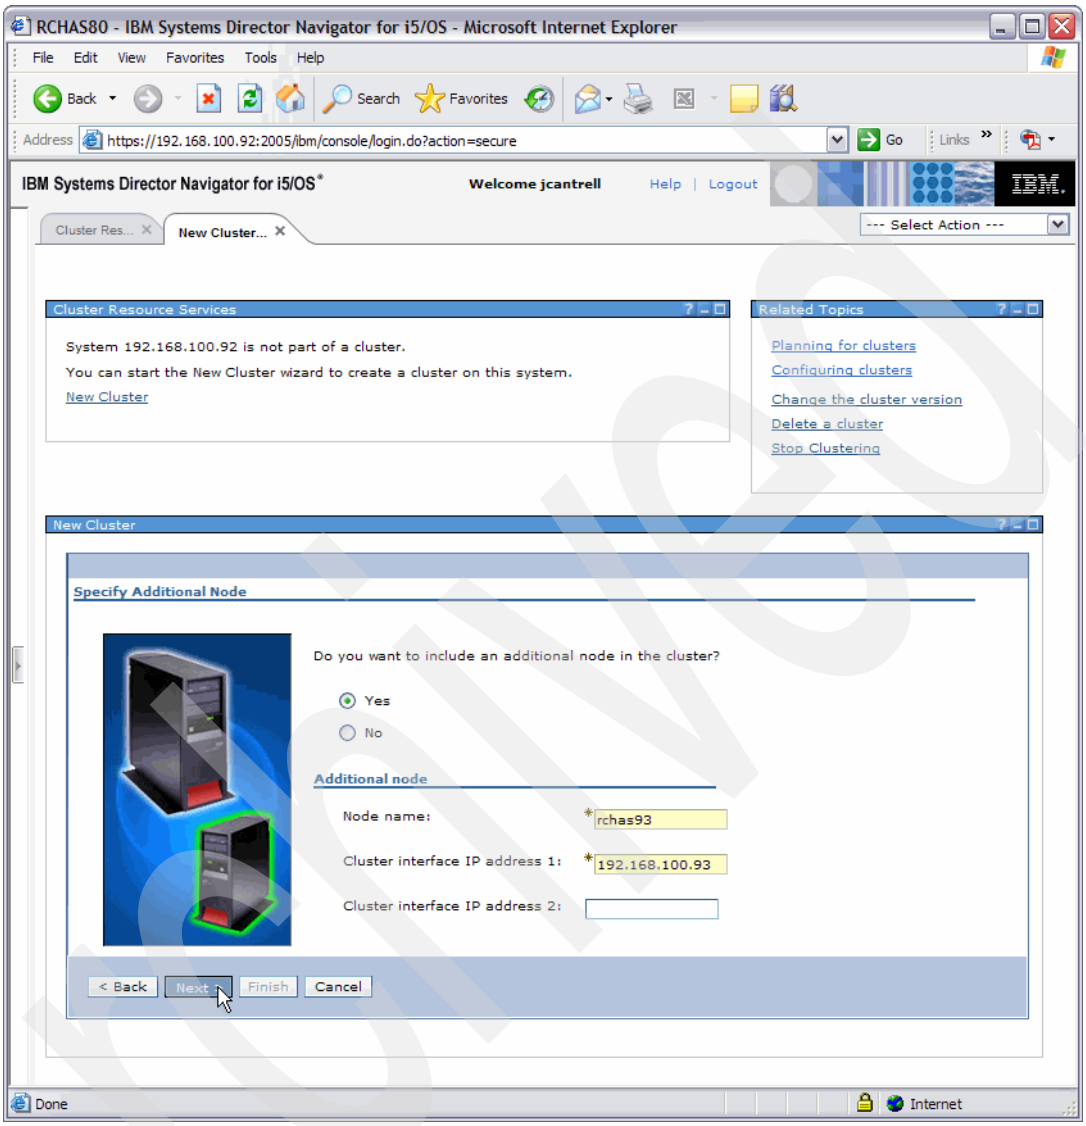


*Figure 9-9 New Cluster → Specify Additional Node*

6. You must enter the user ID and password to sign on to the additional node. Refer to Figure 9-10.

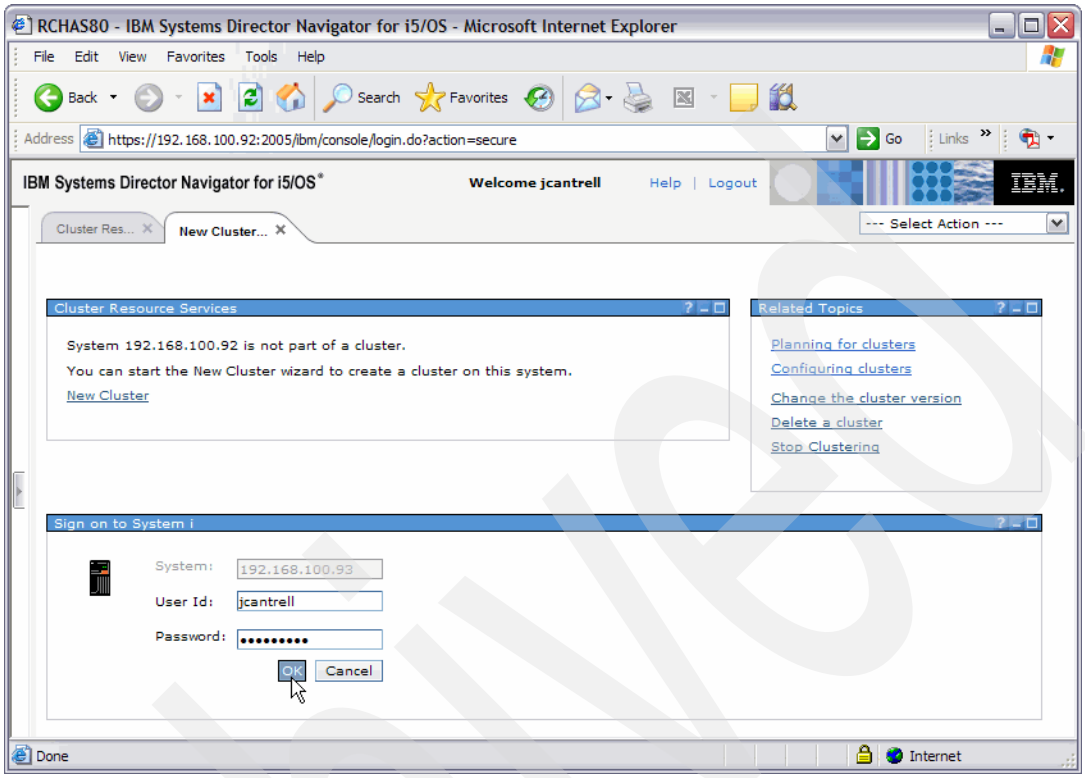


*Figure 9-10 New Cluster → Sign on to System i*

7. Enter the name of the cluster version that you are using (Figure 9-11) and select **Next**.

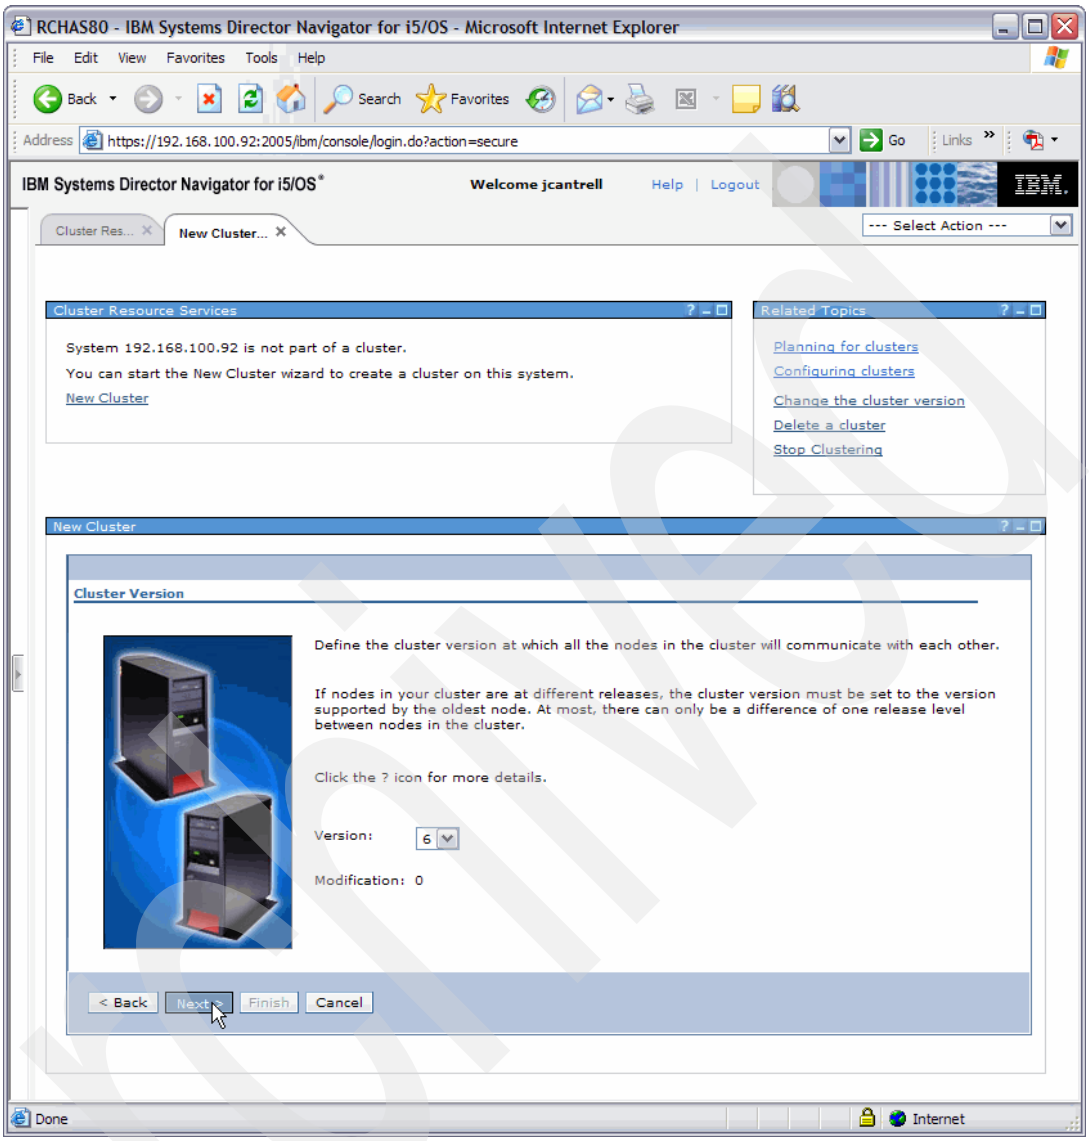


*Figure 9-11 New Cluster → Cluster Versions*

8. The specify Cluster Message Queue window appears. If yes, select **Yes**, then enter the name of the message queue and the library in which it resides. If no, select **No**. Then select **Next** (Figure 9-12).

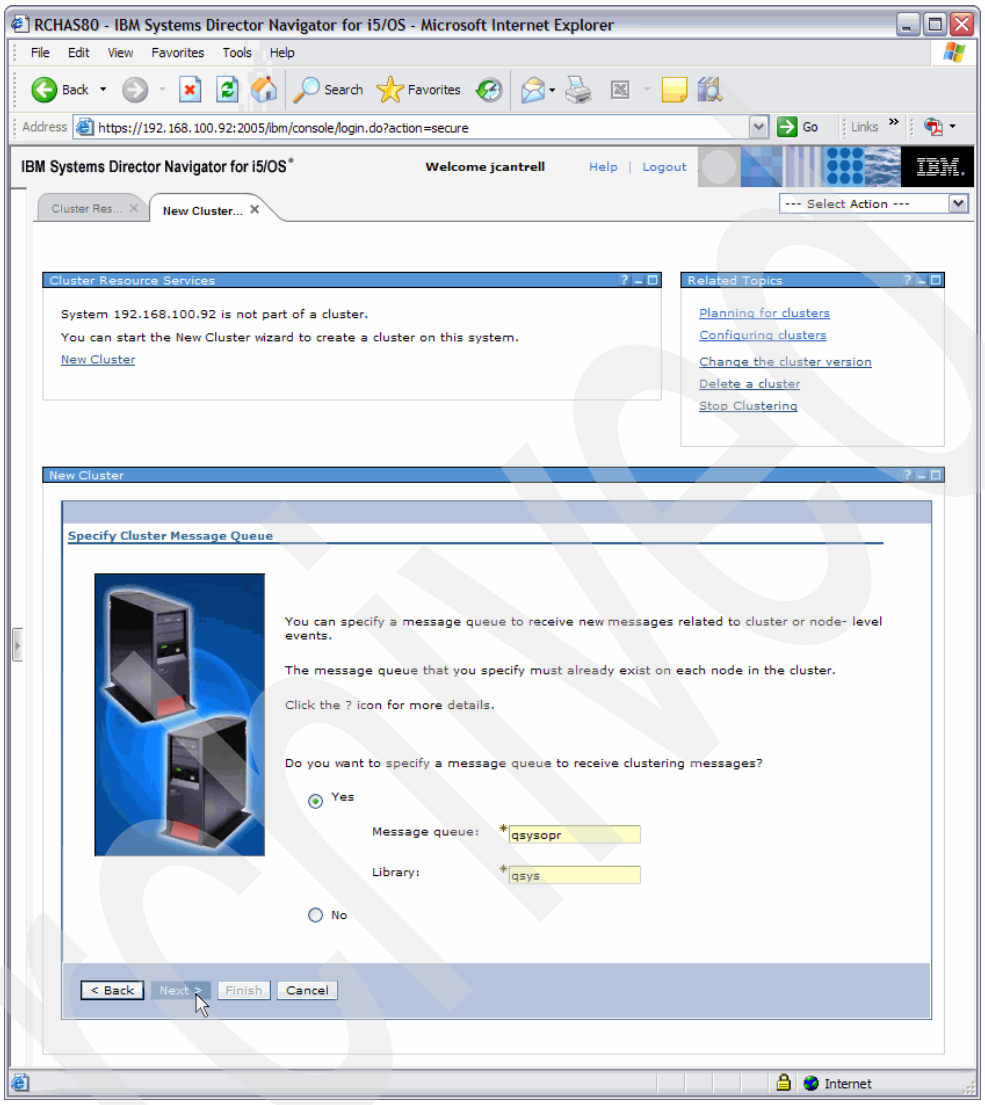


*Figure 9-12 New Cluster → Specify Cluster Message Queue*

9. If you selected yes on the Specify Cluster Message Queue window, you will be prompted for the failover wait time and the rollover default action. Enter your choices and select **Next** (Figure 9-13).

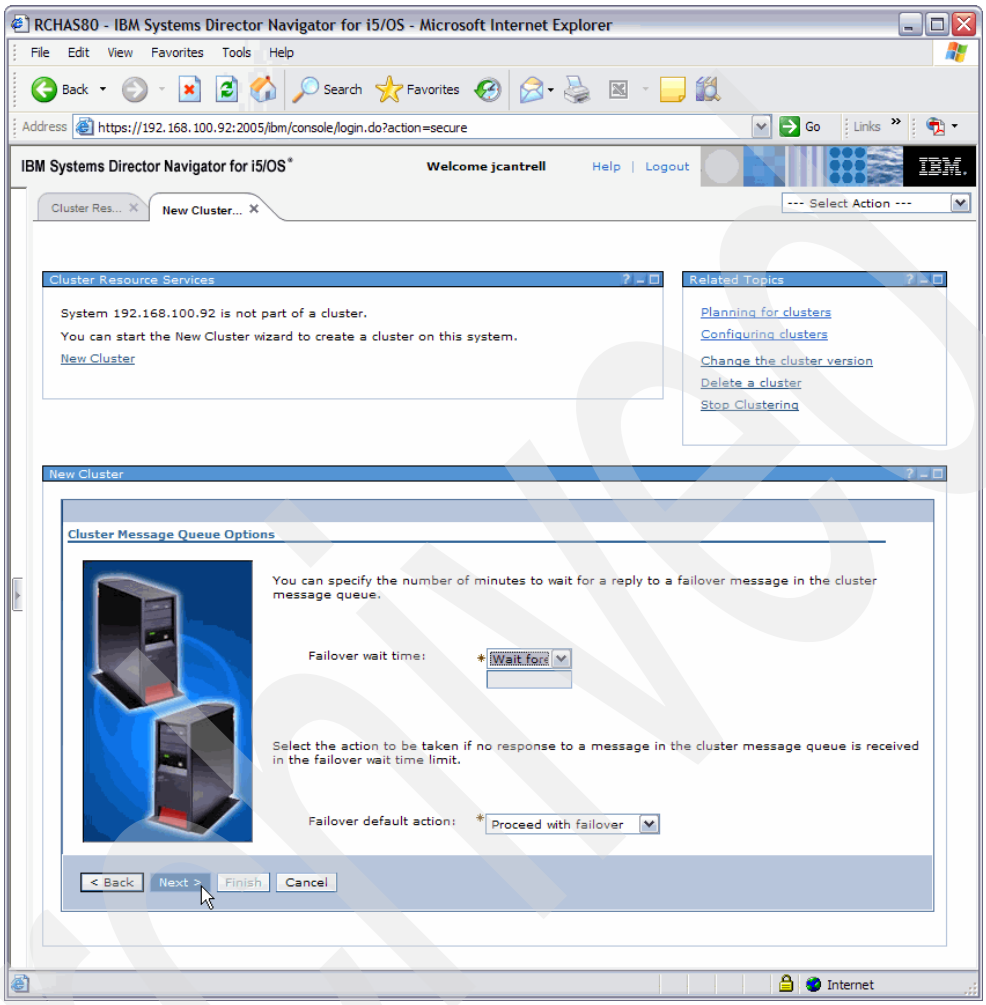


*Figure 9-13 New Cluster → Cluster Message Queue Options*

10. The No Switchable Software Found window appears (Figure 9-14). Click **Next**.

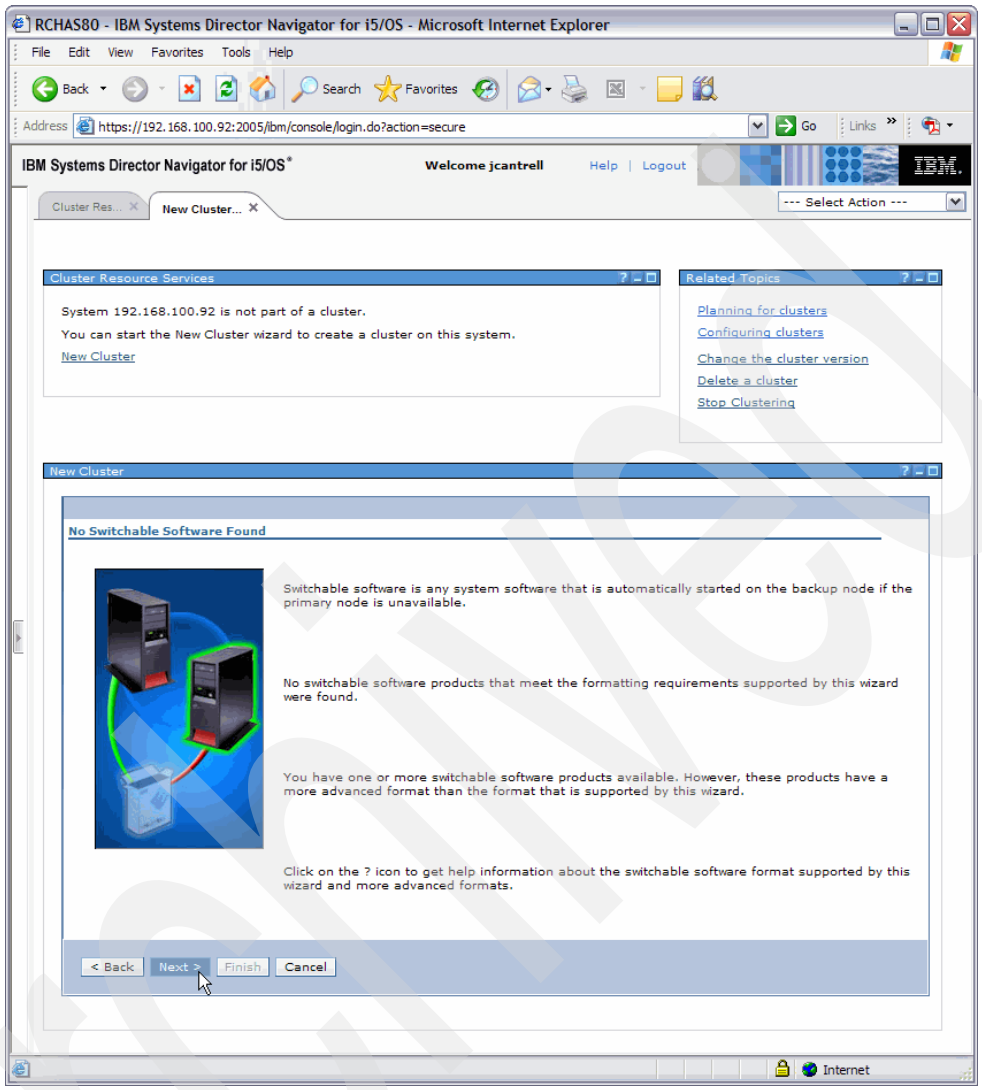


*Figure 9-14 New Cluster → No Switchable Hardware Found*

11. The New Cluster Summary window appears (Figure 9-15). Review the information and click **Finish**.

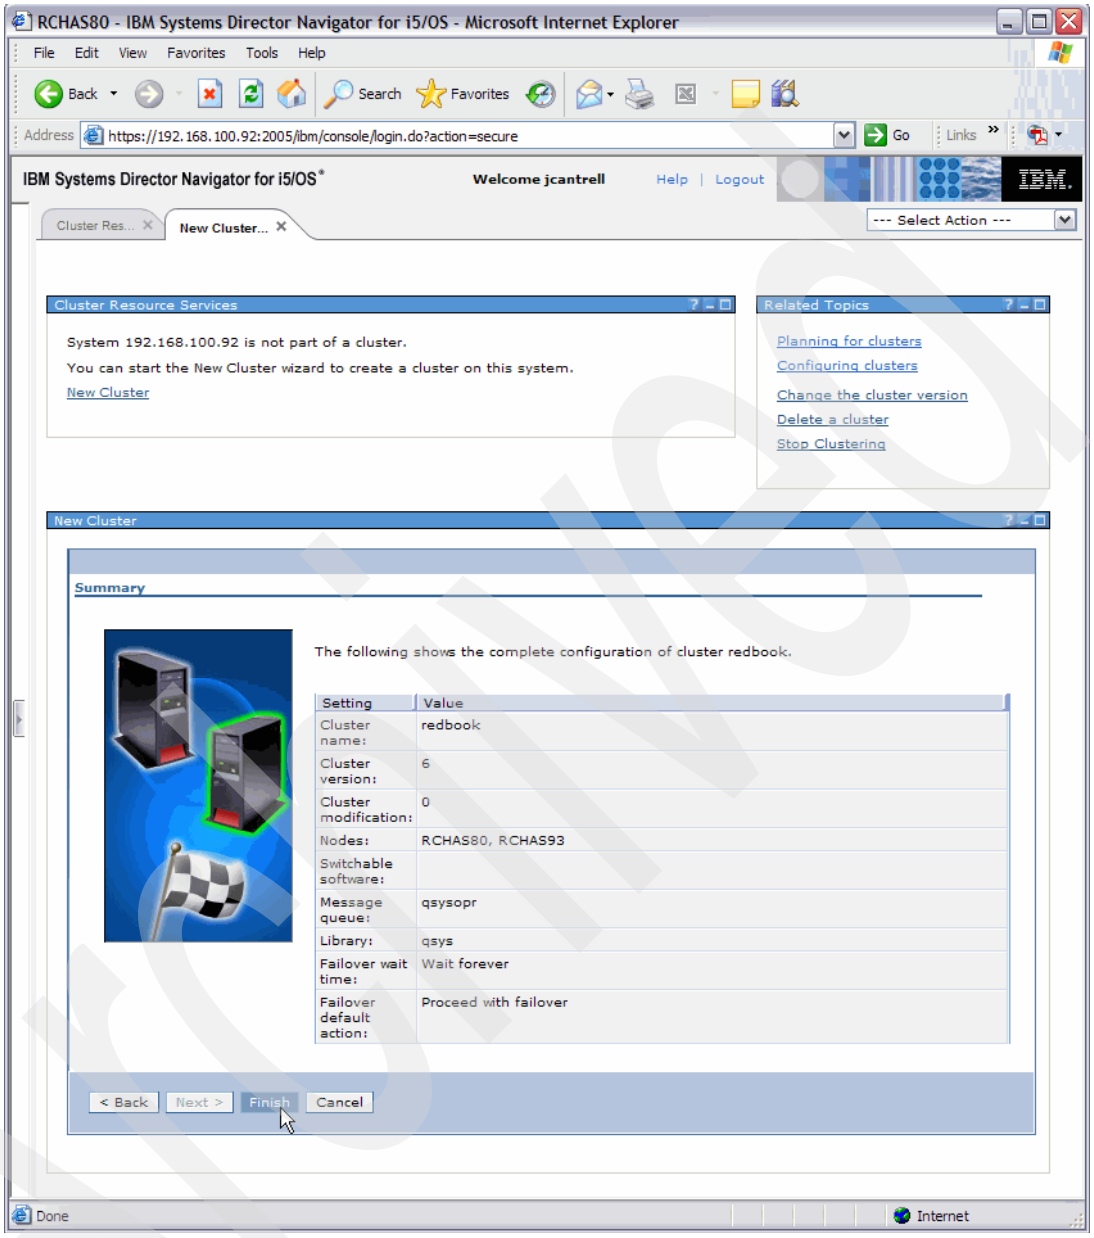


*Figure 9-15 New Cluster → Summary*

12. Next the device domain must be created for each of the nodes in the cluster. On the Cluster Resource Services window for your cluster, select **Work with Cluster Nodes**. See Figure 9-16.

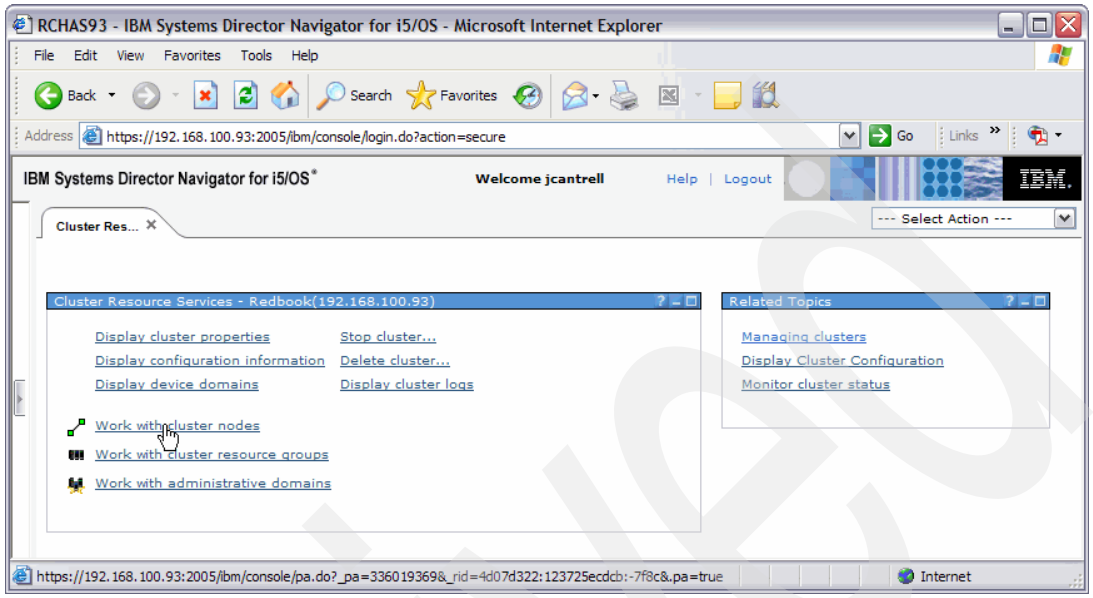


*Figure 9-16 Cluster Resource Services → Work with Cluster Nodes*

13. On the Work with Cluster Nodes window, click the double arrow next to the node and select **Properties** (Figure 9-17).

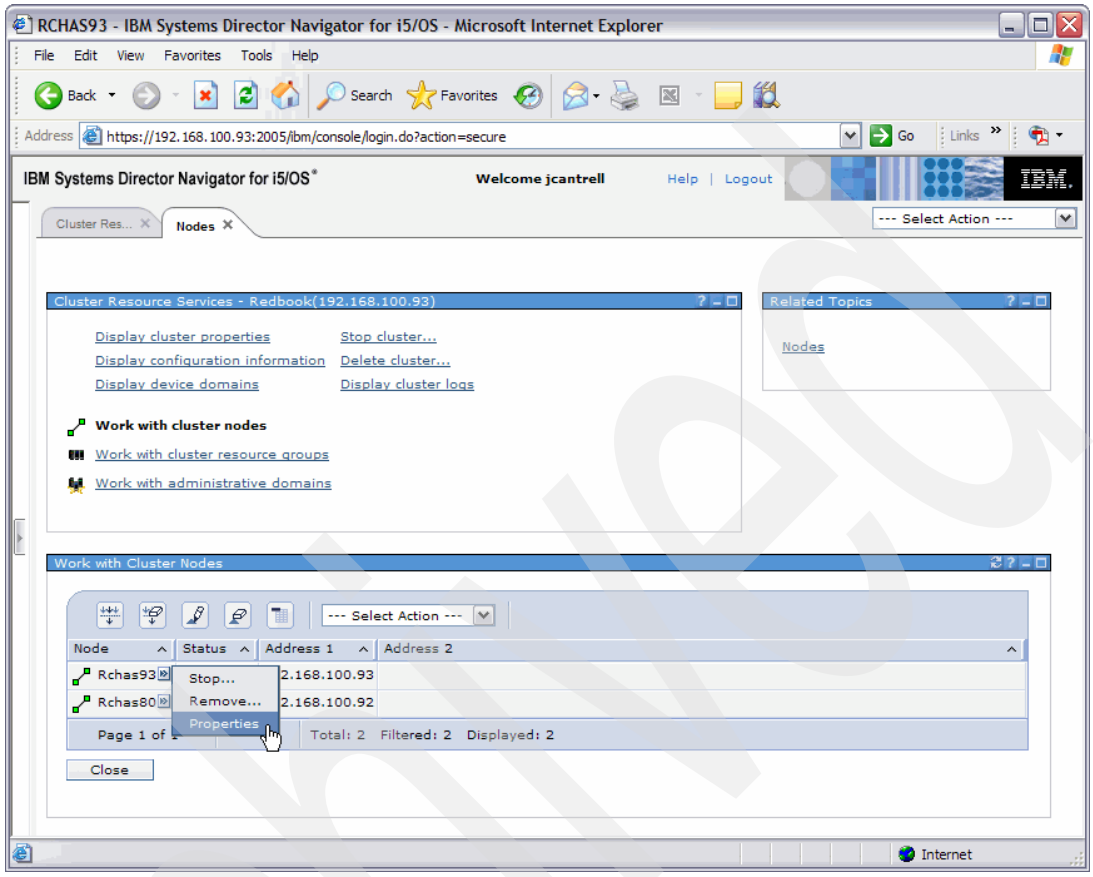


*Figure 9-17 Cluster Resource Services → Work with Cluster Nodes → Properties*

14. On the Properties window, select **Clustering** (Figure 9-18).

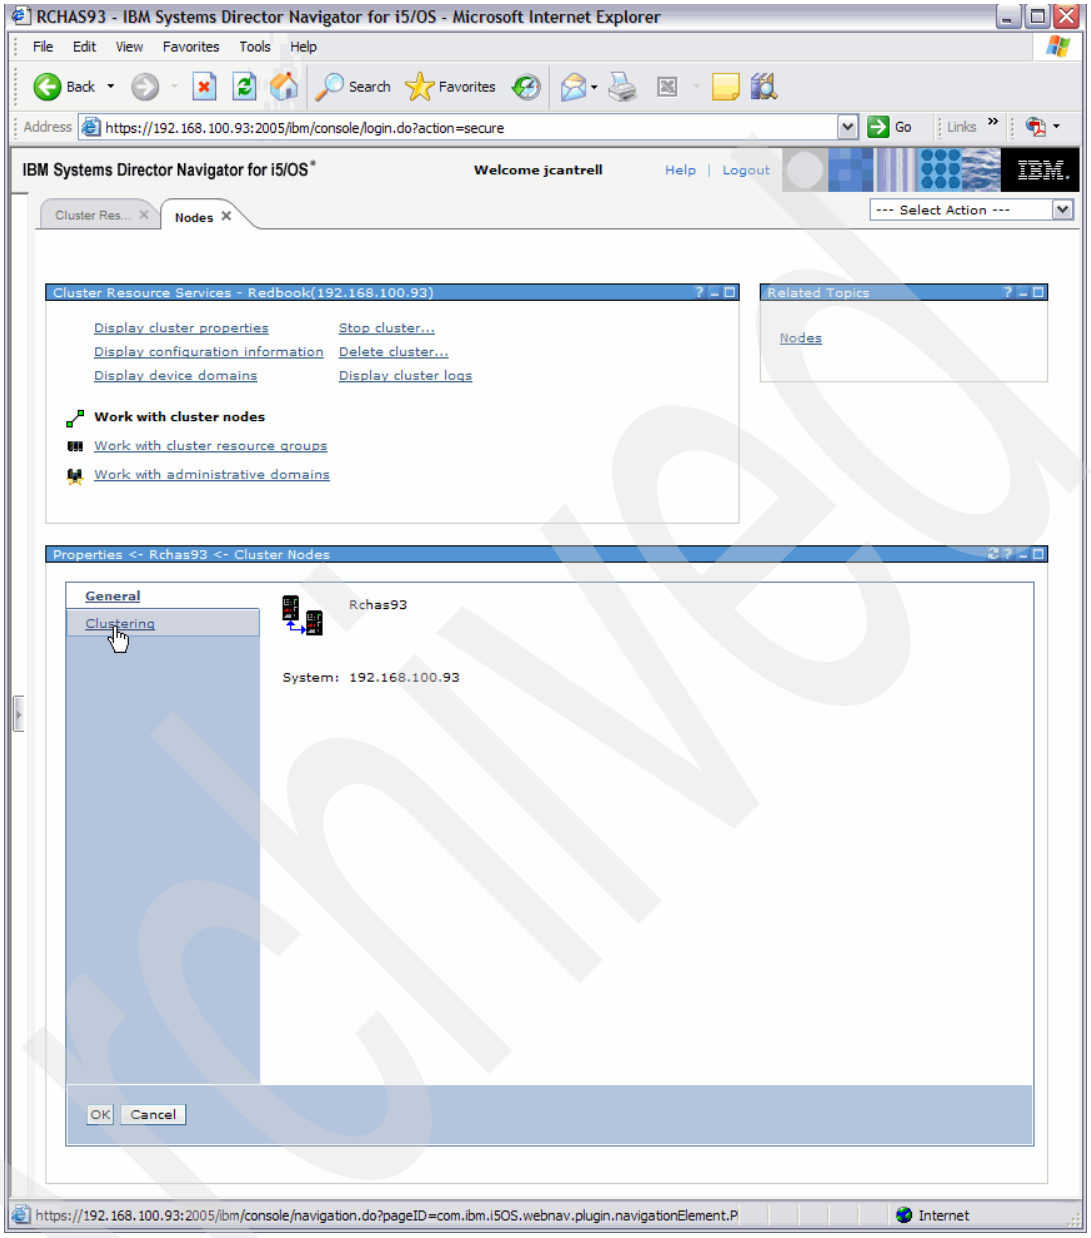


*Figure 9-18 Cluster Resource Group → Properties Cluster Node → Clustering*

15. Enter the device domain name and select **OK** (Figure 9-19).

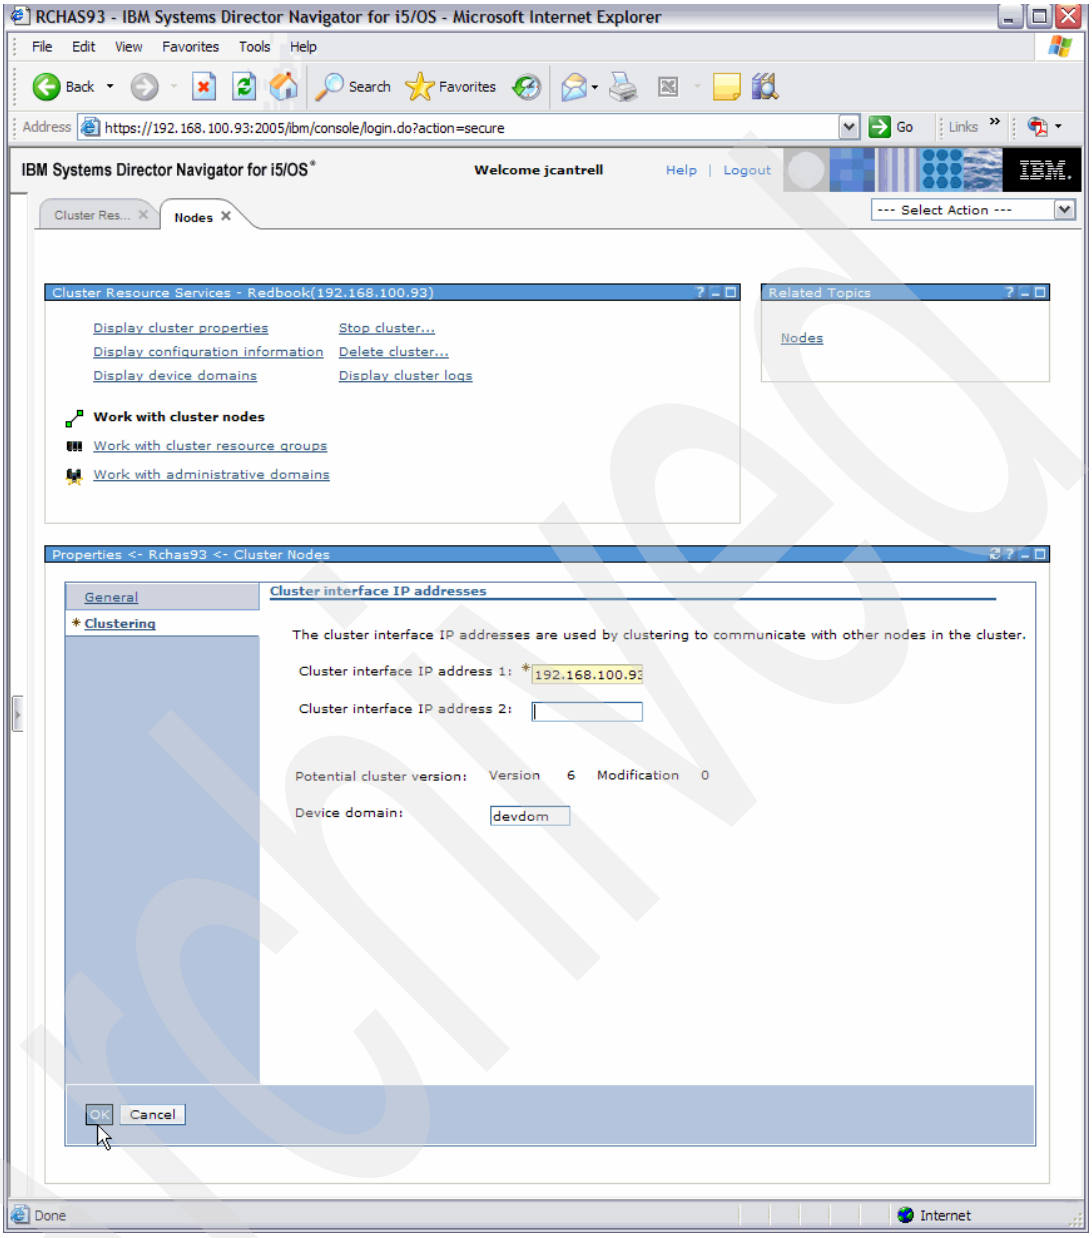


*Figure 9-19 Properties Cluster Node → Device Domain*

16. Next we work with the properties on the second node in the cluster. Click double arrow and select **Properties** (Figure 9-20).

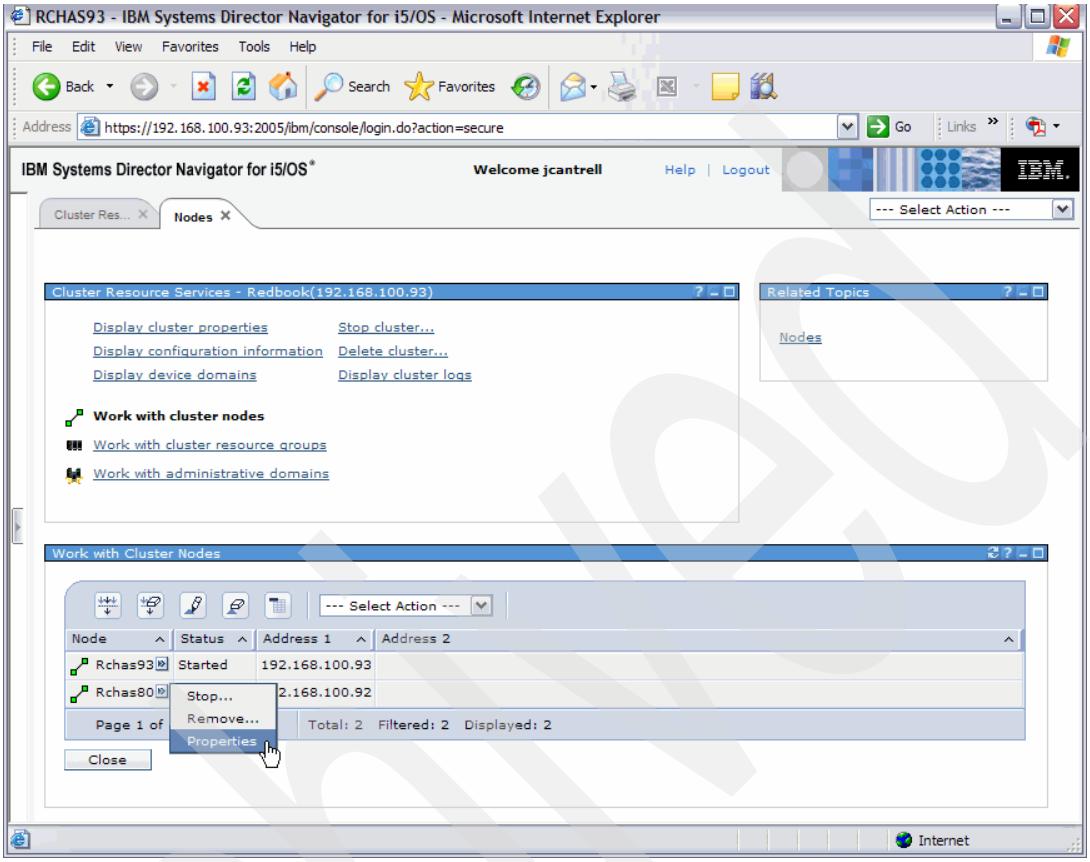


*Figure 9-20 Work with Cluster Nodes → Properties*

17. On the Properties Cluster Nodes window, select **Clustering** (Figure 9-21).

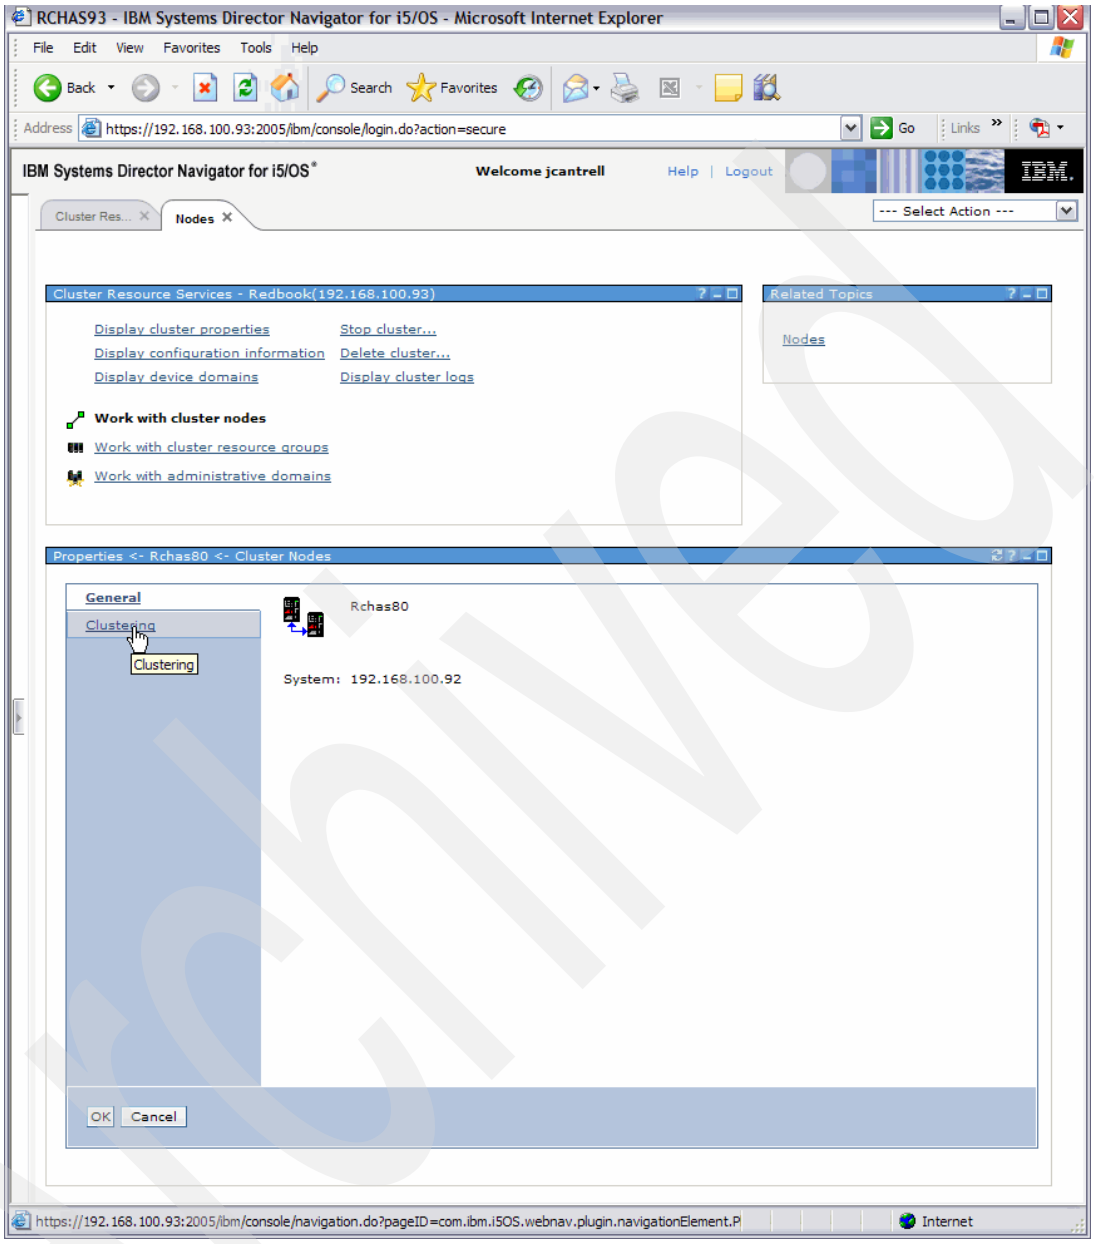


*Figure 9-21 Properties Cluster Nodes → Clustering*

18. Enter the device domain name and select **OK** (Figure 9-22).

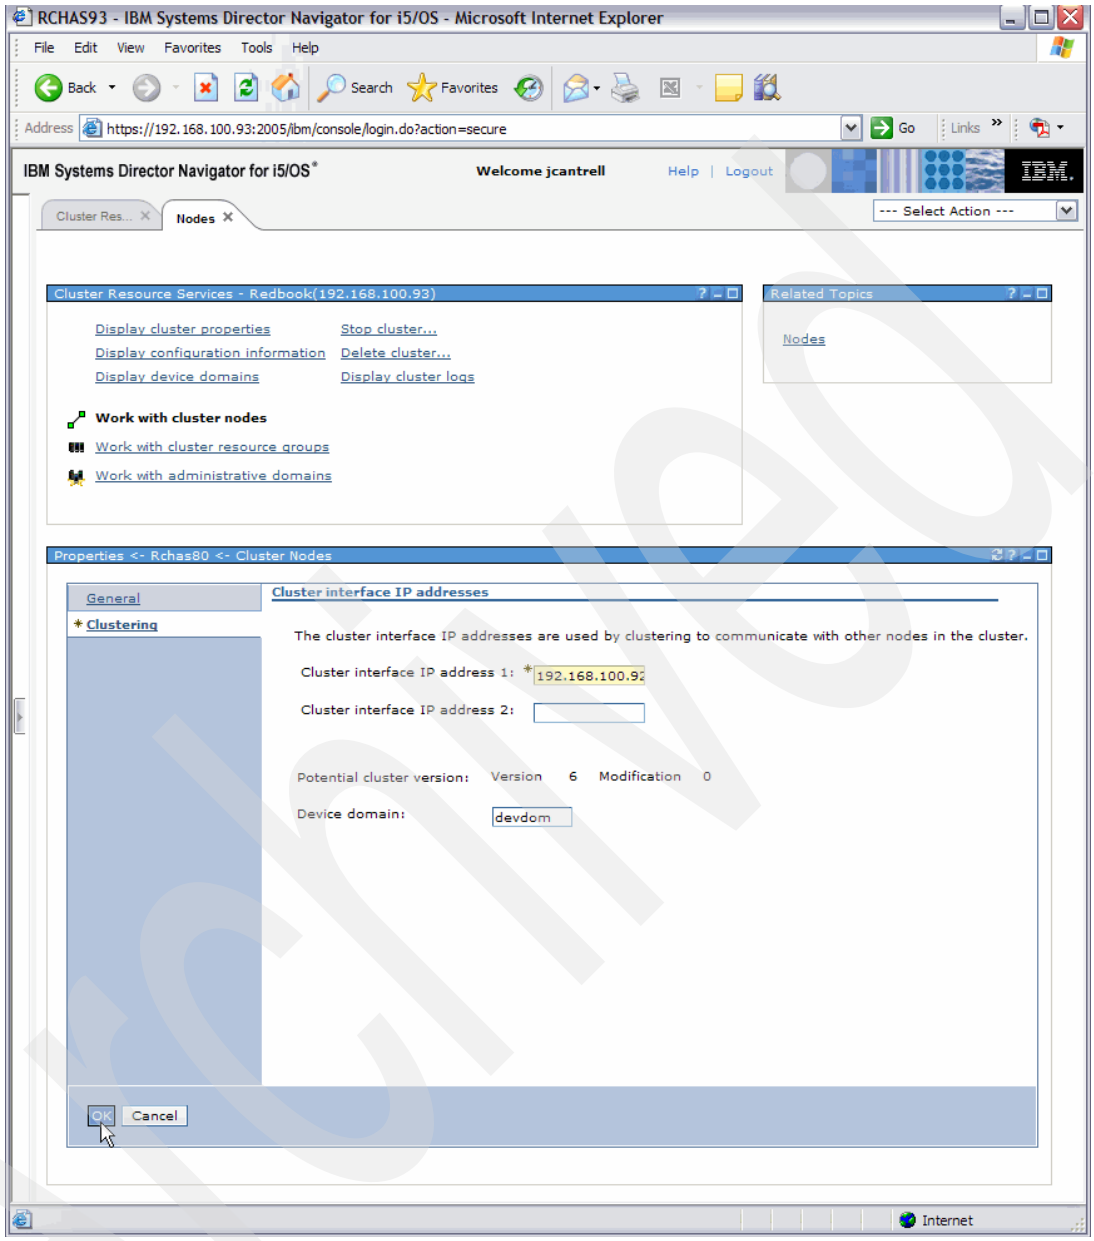


*Figure 9-22 Properties Cluster Node → Device Domain*

19. Upon completion of entering the device domain, select **Close** (Figure 9-23).

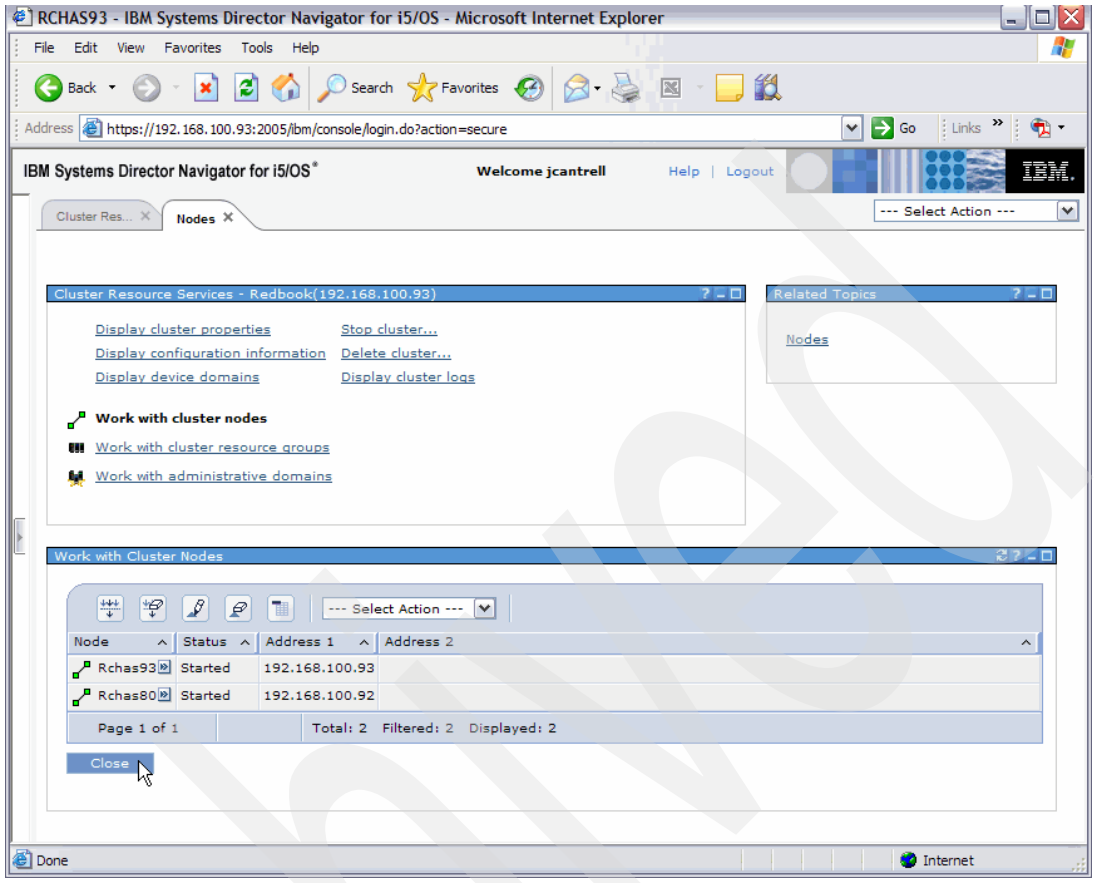


*Figure 9-23 Work with Cluster Nodes → Close*

20. On the second node, you must use the CRTDEVASP command to manually create the device descriptions for all disk pools being switched. The description name and the resource name must be the same names that were used when the IASP was created. See Figure 9-24.

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

 Type choices, press Enter.

 Device description . . . . . > IASPA1        Name
 Resource name  . . . . . > IASPA1        Name
 Relational database  . . . . . .   *GEN
 Message queue  . . . . . .   *SYSOPR       Name
   Library  . . . . . .                    Name, *LIBL, *CURLIB
 Text 'description' . . . . . .   *BLANK


                          Additional Parameters

 Authority  . . . . . .   *CHANGE       Name, *CHANGE, *ALL,
*USE...




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

*Figure 9-24 CRTDEVASP command*

21. Creating the ASP device should automatically create the relational database name. To verify, use the WRKRDBDIRE command (Figure 9-25).

```
 Work with Relational Database Directory Entries

 Position to  . . . . . .

 Type options, press Enter.
   1=Add   2=Change   4=Remove   5=Display details   6=Print details

                              Remote
 Option  Entry                Location                  Text

         IASPA1               LOOPBACK
         IASPA1S              LOOPBACK
         RCHAS80              *LOCAL




Bottom
 F3=Exit   F5=Refresh   F6=Print list   F12=Cancel   F22=Display entire
field
 (C) COPYRIGHT IBM CORP. 1980, 2007.
```

*Figure 9-25 WRKRDBDIRE command*

Next you must create the cluster resource group.

## 9.3 Creating a cluster resource group

You must create a definition in the cluster to make both systems *aware* when using switchable devices. This is called the *cluster resource group* (CRG). For information regarding the device cluster resource group, refer to *Implementing PowerHA for IBM i*, SG24-7405.

1. On the Cluster Resource window select **Work with Cluster Resource Groups** (Figure 9-26).

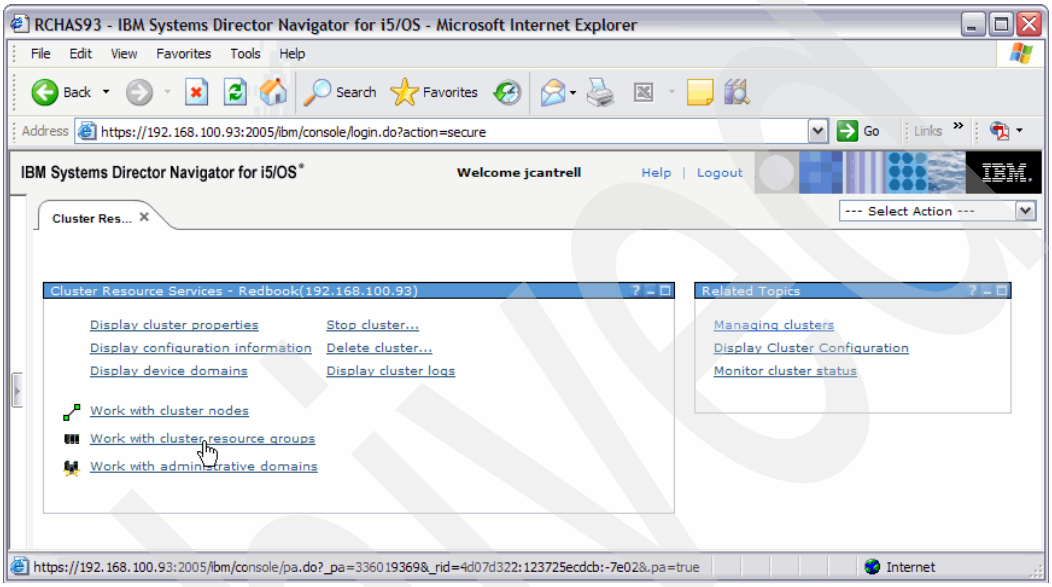


*Figure 9-26 Cluster Resources → Work with Cluster Resource Groups*

2. Under Work with Cluster Resource Groups, use the drop-down to select **New Device CRG** (Figure 9-27).

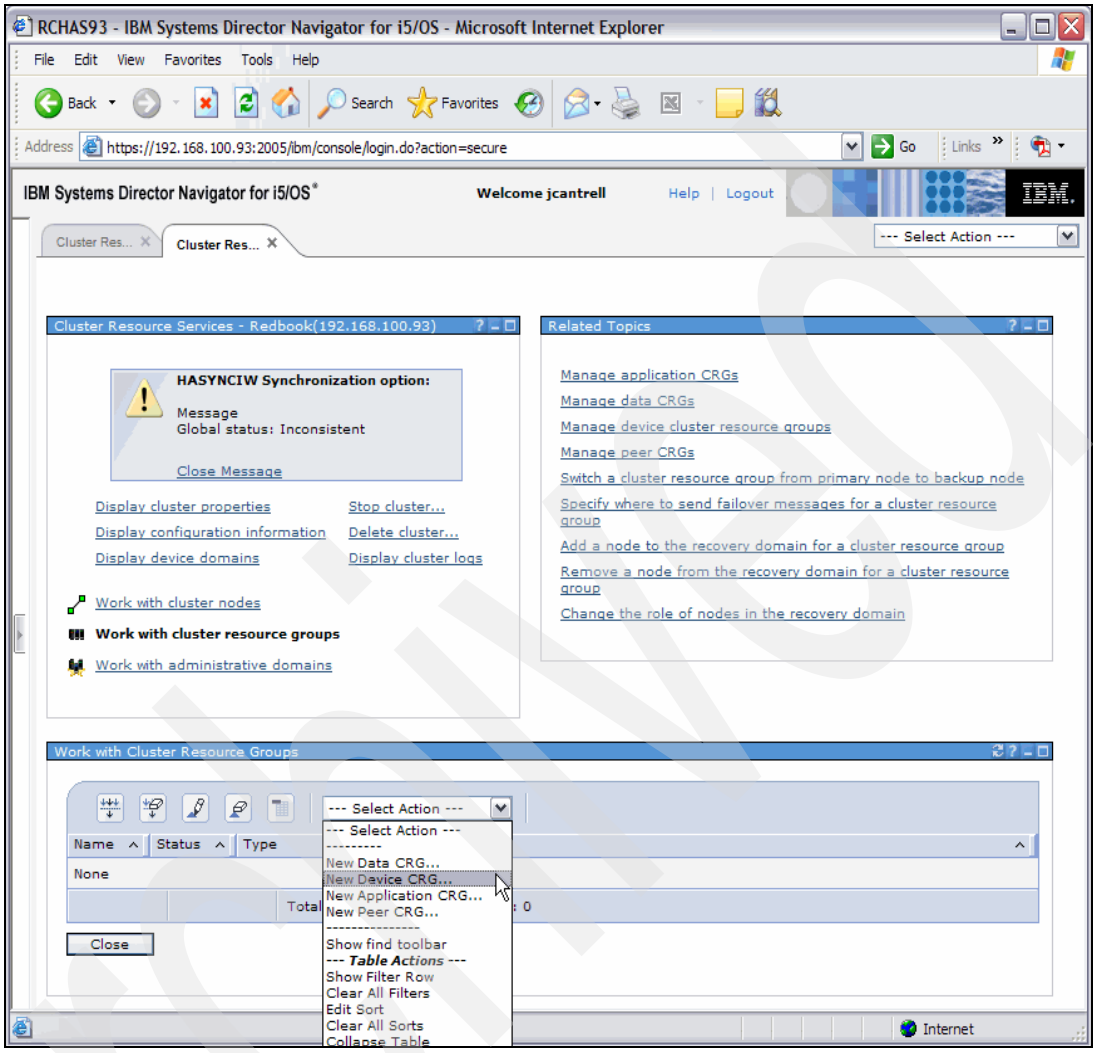


*Figure 9-27 Work with Cluster Resource Groups → New Device CRG*

3. The Welcome window appears. Select **Next** (Figure 9-28).

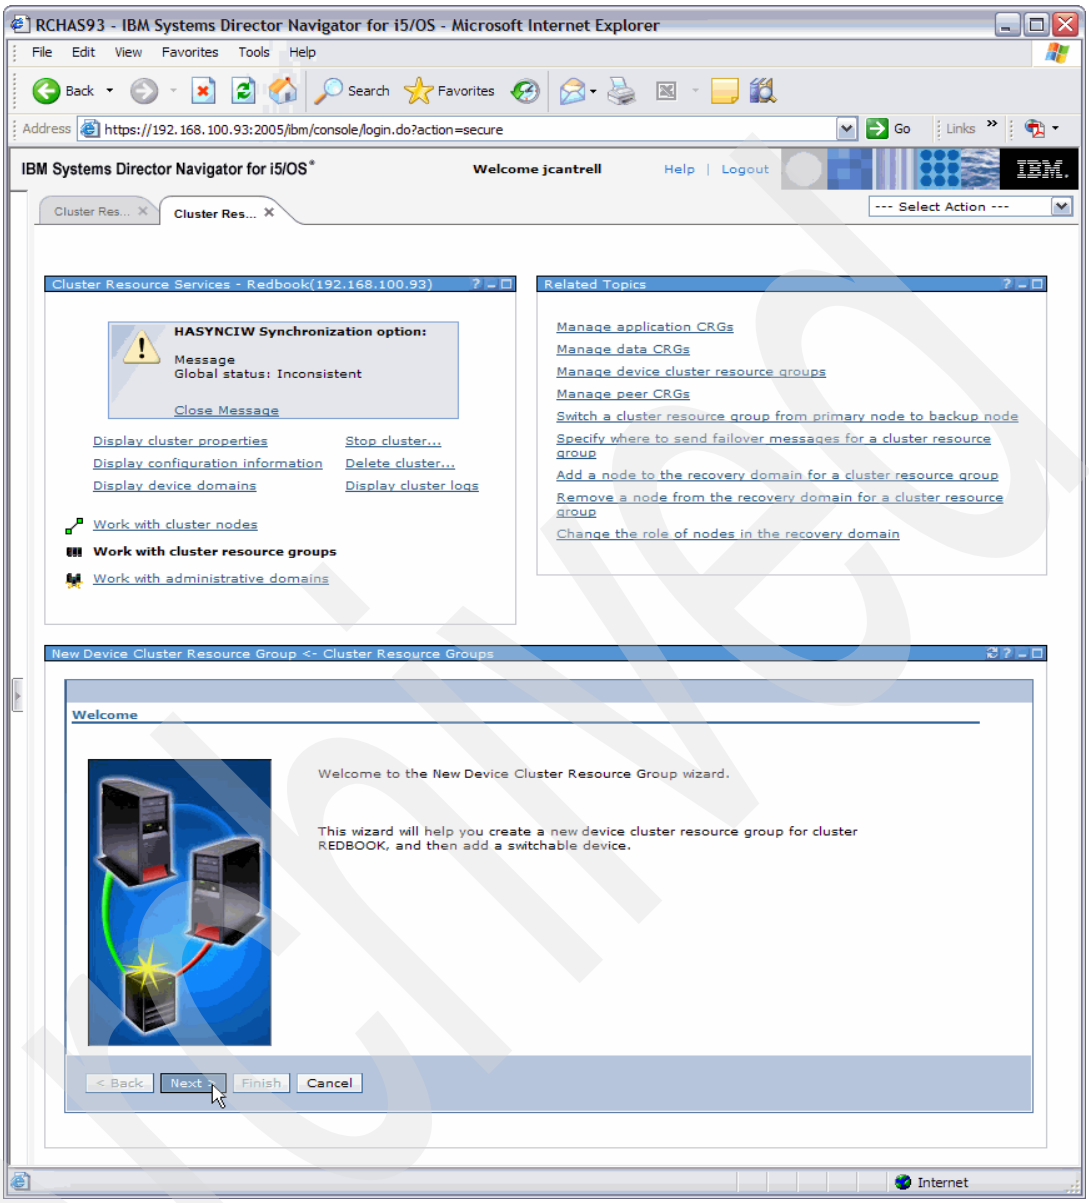


*Figure 9-28 New Device Cluster Resource Group → Welcome*

4. Enter the name of the primary node and select **Next** (Figure 9-29).

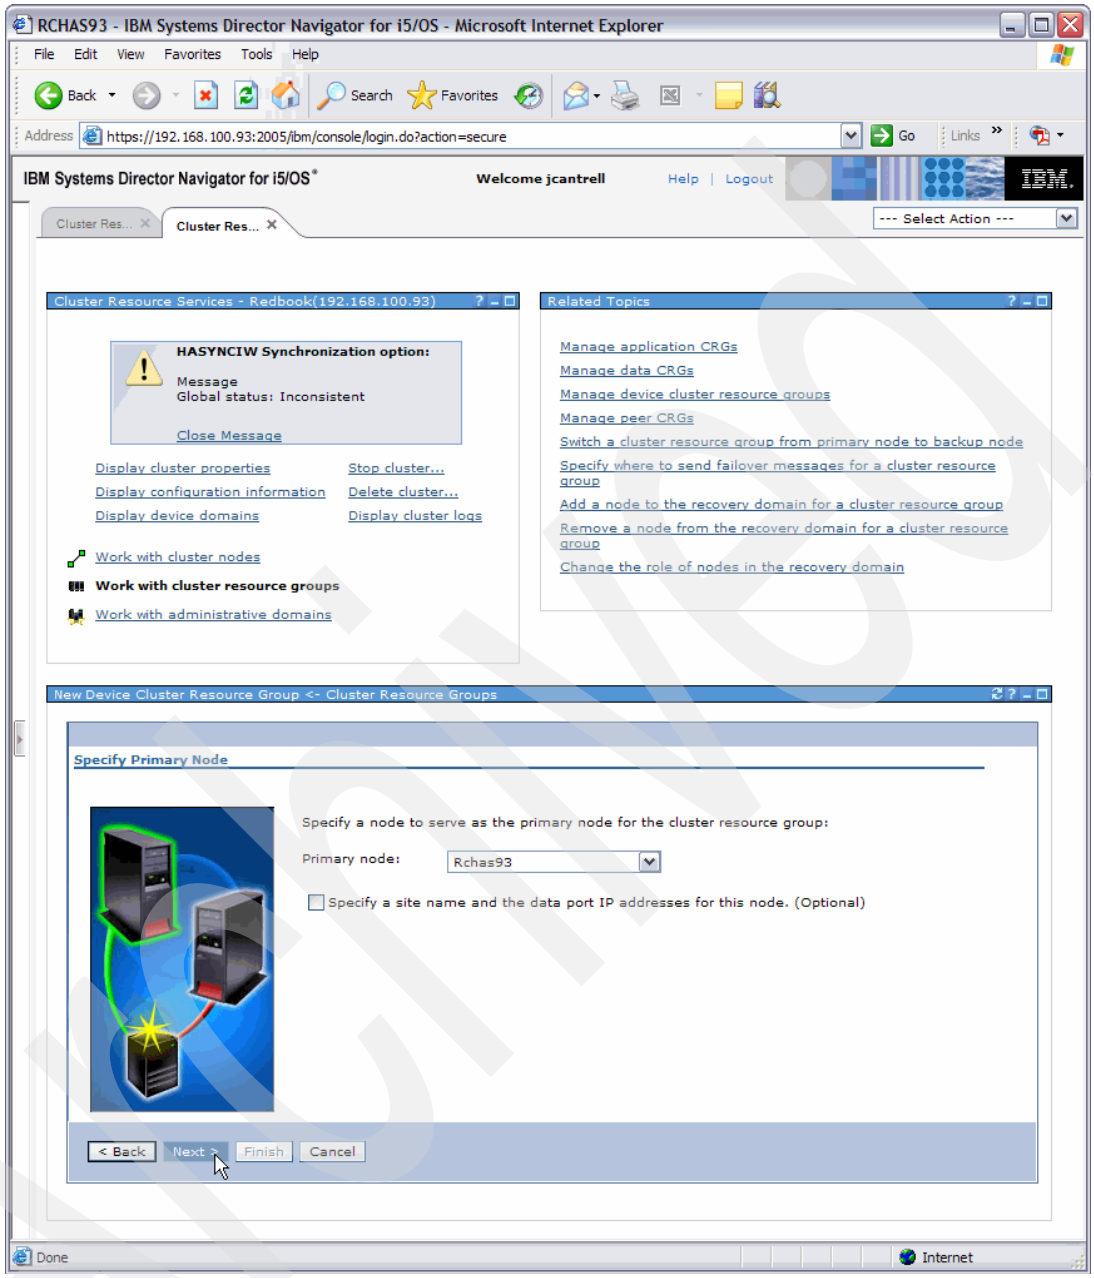


*Figure 9-29 Specify Primary Node*

5. The Specify Additional Node window appears. Enter the name of the additional node and select **Next** (Figure 9-30).

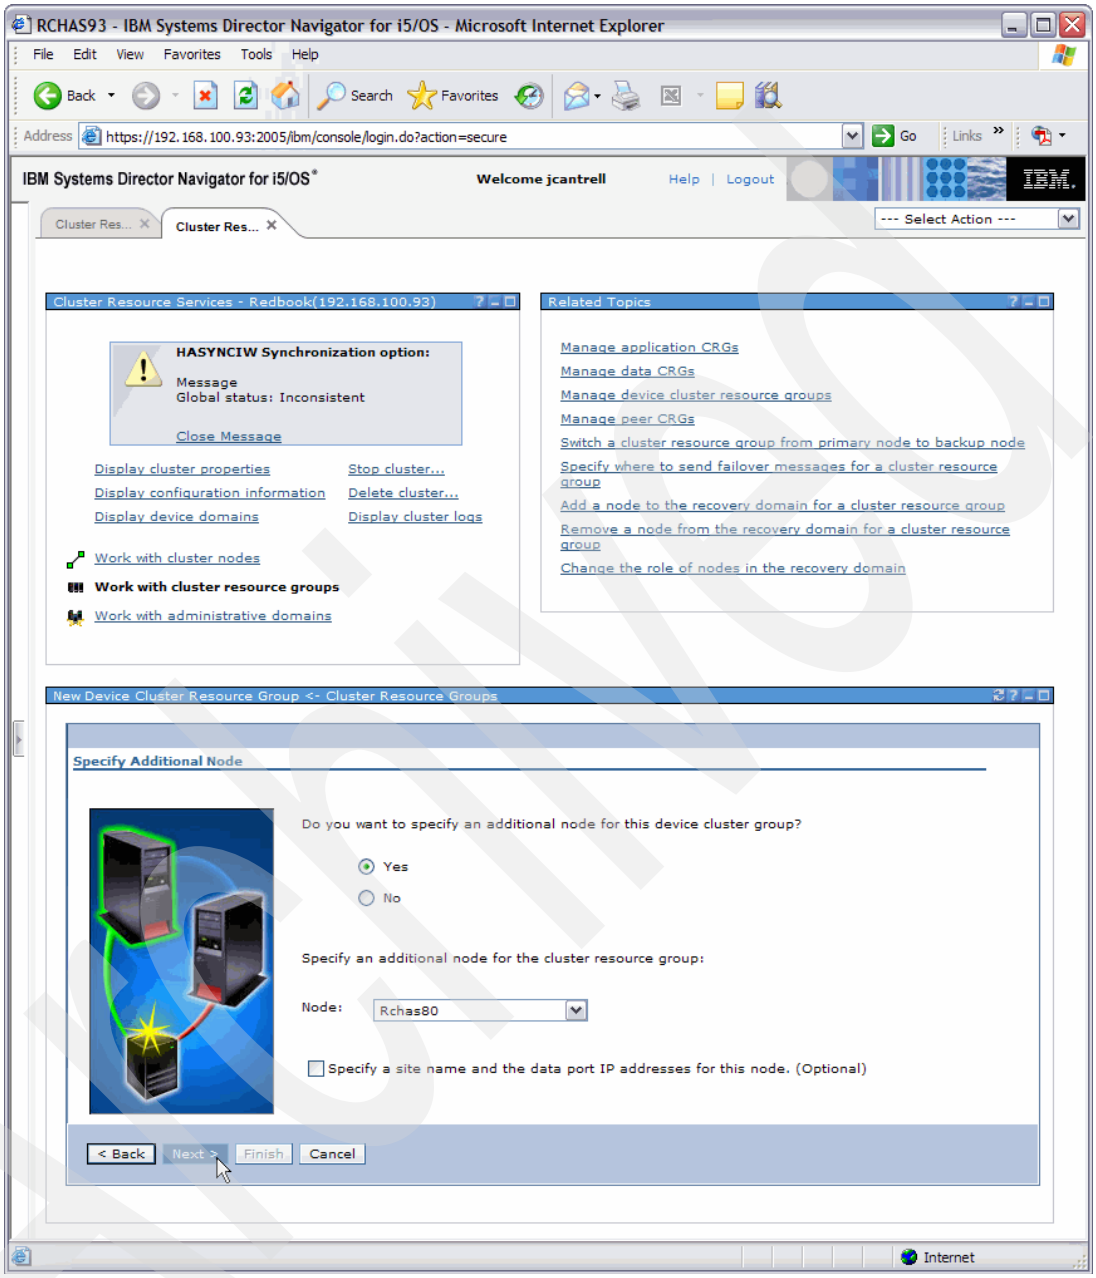


*Figure 9-30 Specify Additional Node*

6. The Specify Name window appears. Enter the name of the cluster resource group and a description, then select **Next** (Figure 9-31).


*Figure 9-31 Specify Name*

7. Next is the Specify Failover Message Queue. If you wish to define it, do so here and select **Next** (Figure 9-32).

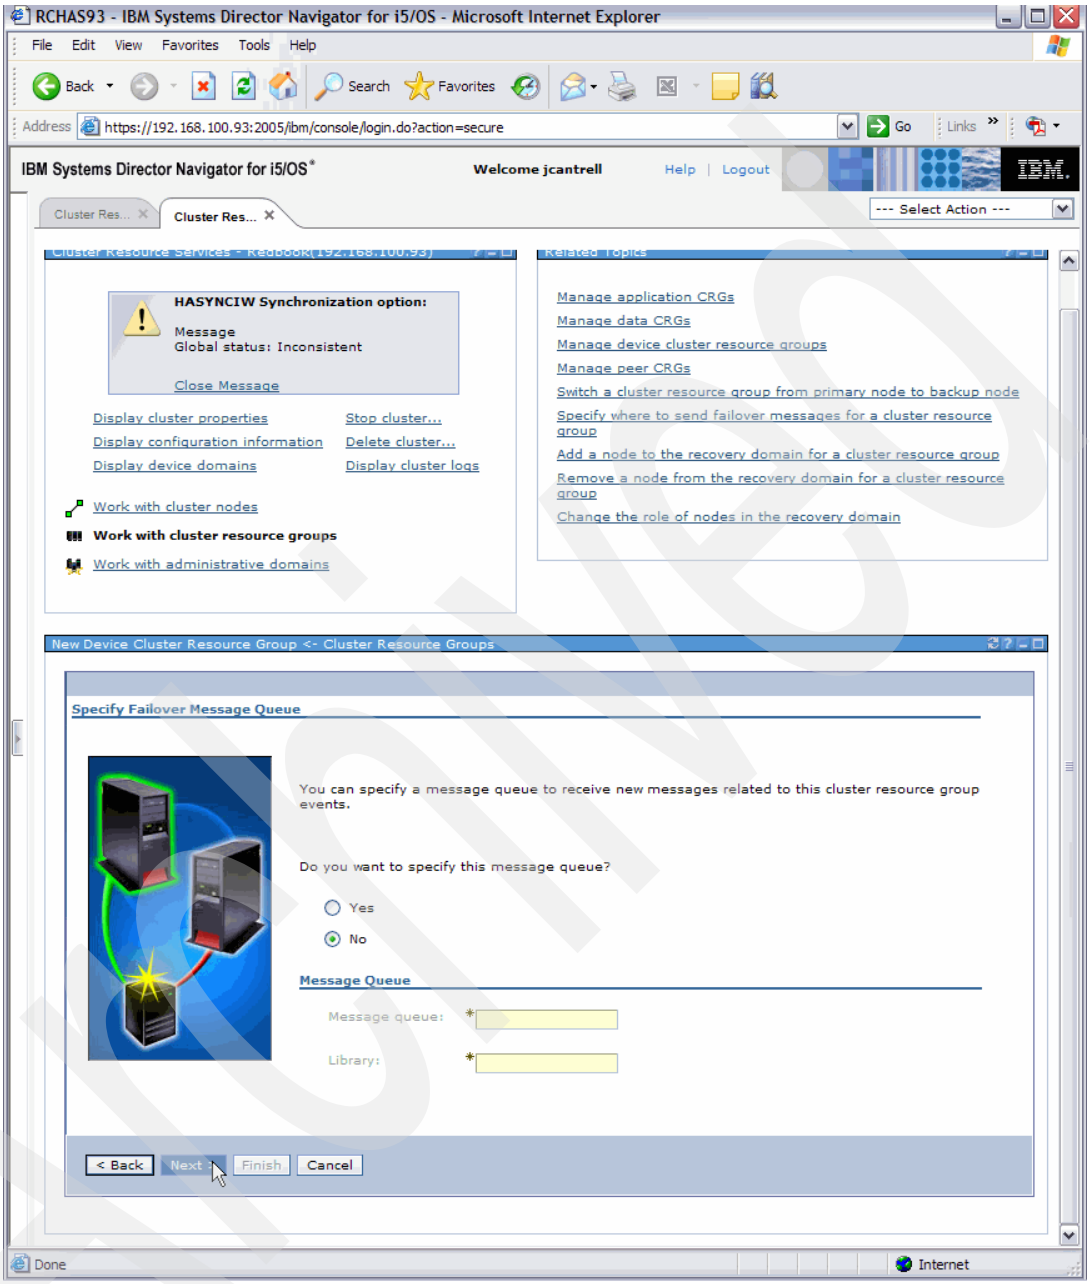


*Figure 9-32 Specify Failover Message Queue*

8. The Configuration Object window uses a drop-down to specify the configuration object type and device type. Our configuration object type is Device description and the device type is AS - Auxiliary storage pool. Select **Next**. See Figure 9-33.

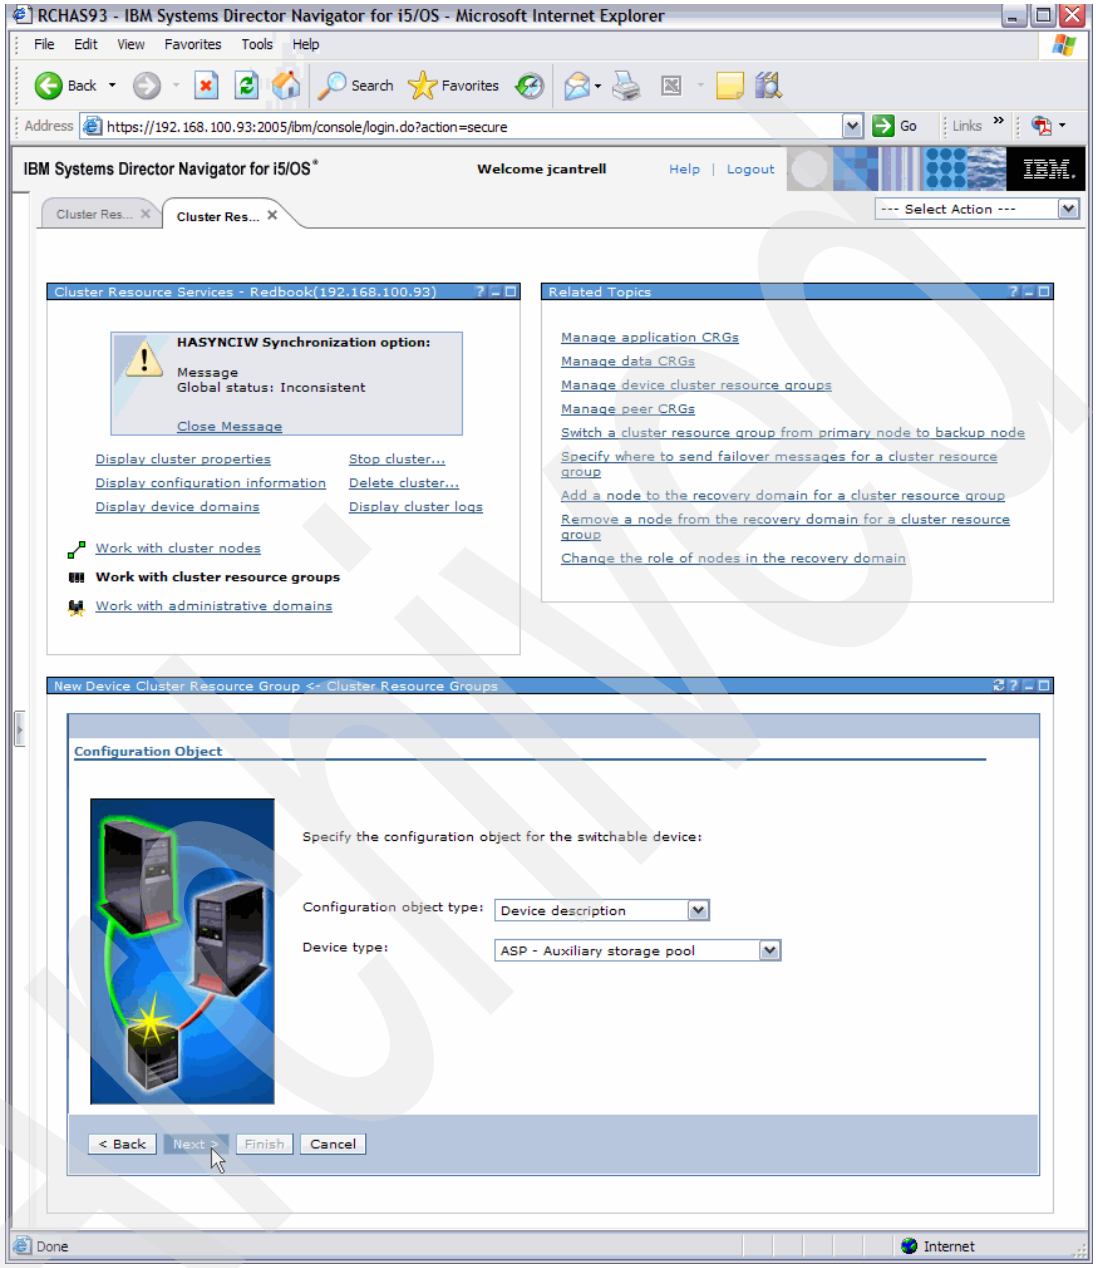


*Figure 9-33 Configuration Object*

9. You have an opportunity on the next window to create new disk pools or add existing disk pools. In our example we add the disk pools that we created. Click **No** and enter the name of the disk pool being added to the cluster. Refer to Figure 9-34.

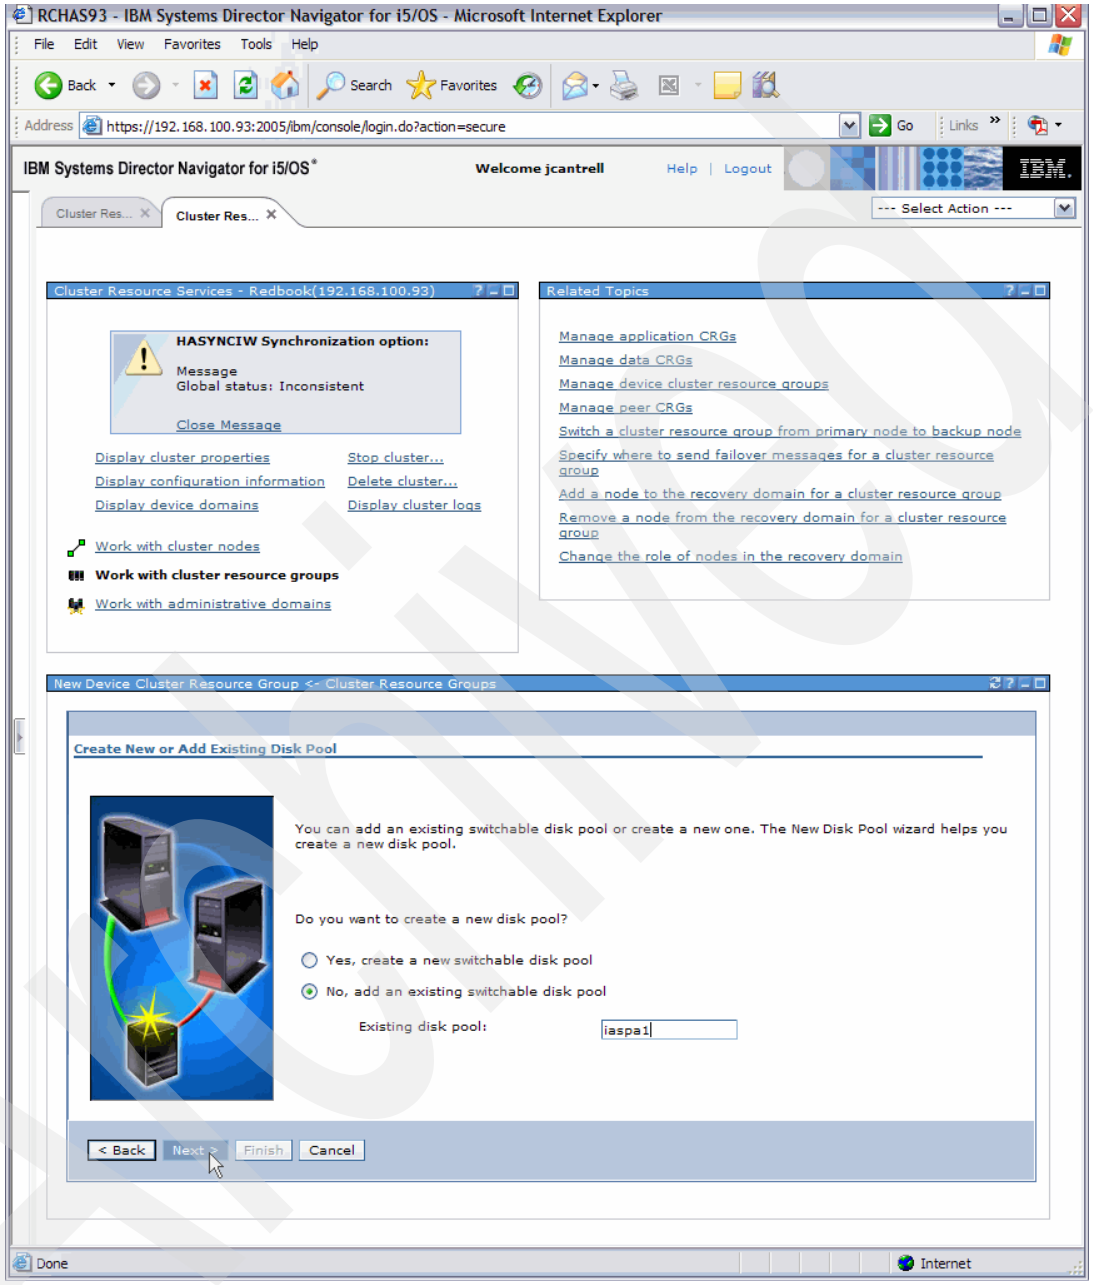


*Figure 9-34 Create New or Add Existing Disk Pool*

10. Because the IASPA1 disk pool contains a secondary disk pool, the warning message shown in Figure 9-35 appears. Click **Yes**.

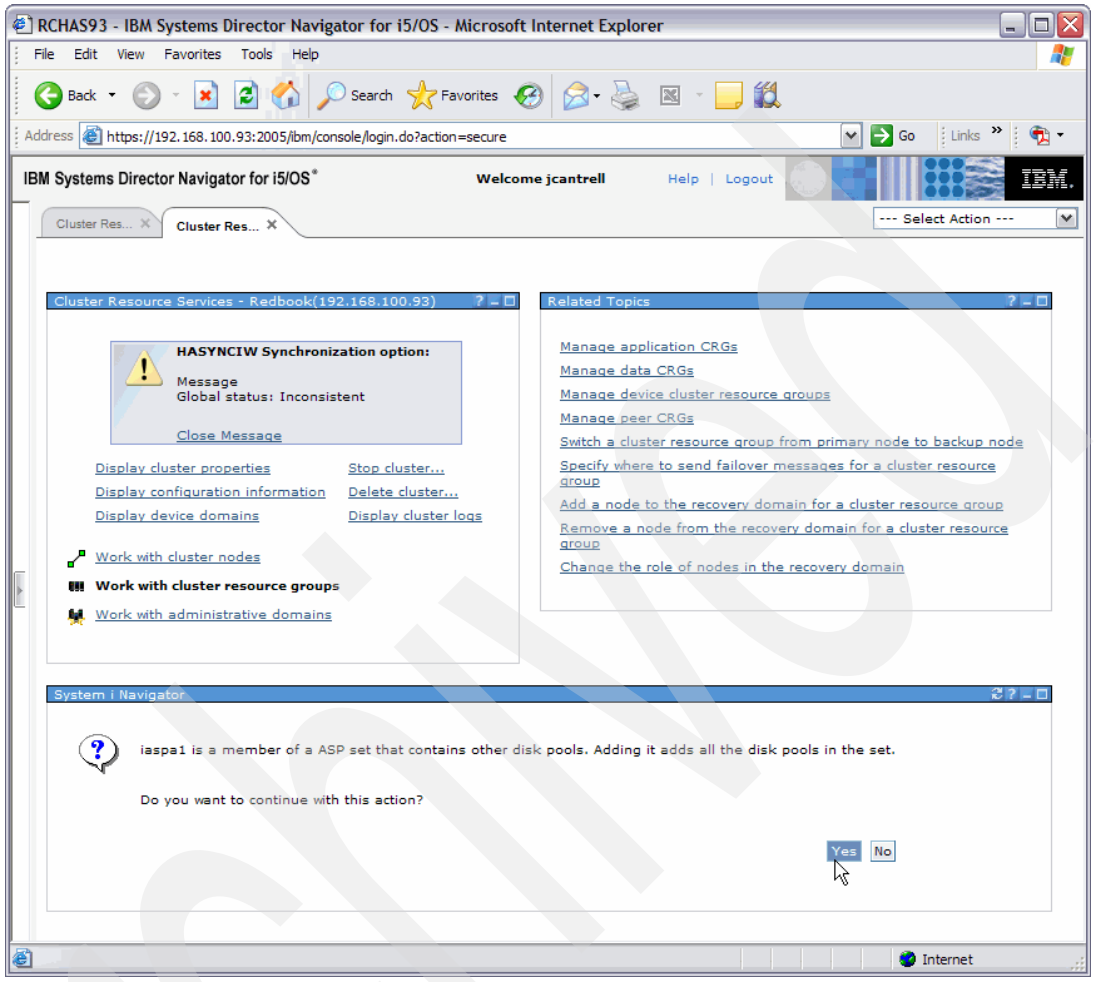


*Figure 9-35 Warning Message*

11. The Cluster Resource Group Summary is displayed (Figure 9-36). Click **Finish**.

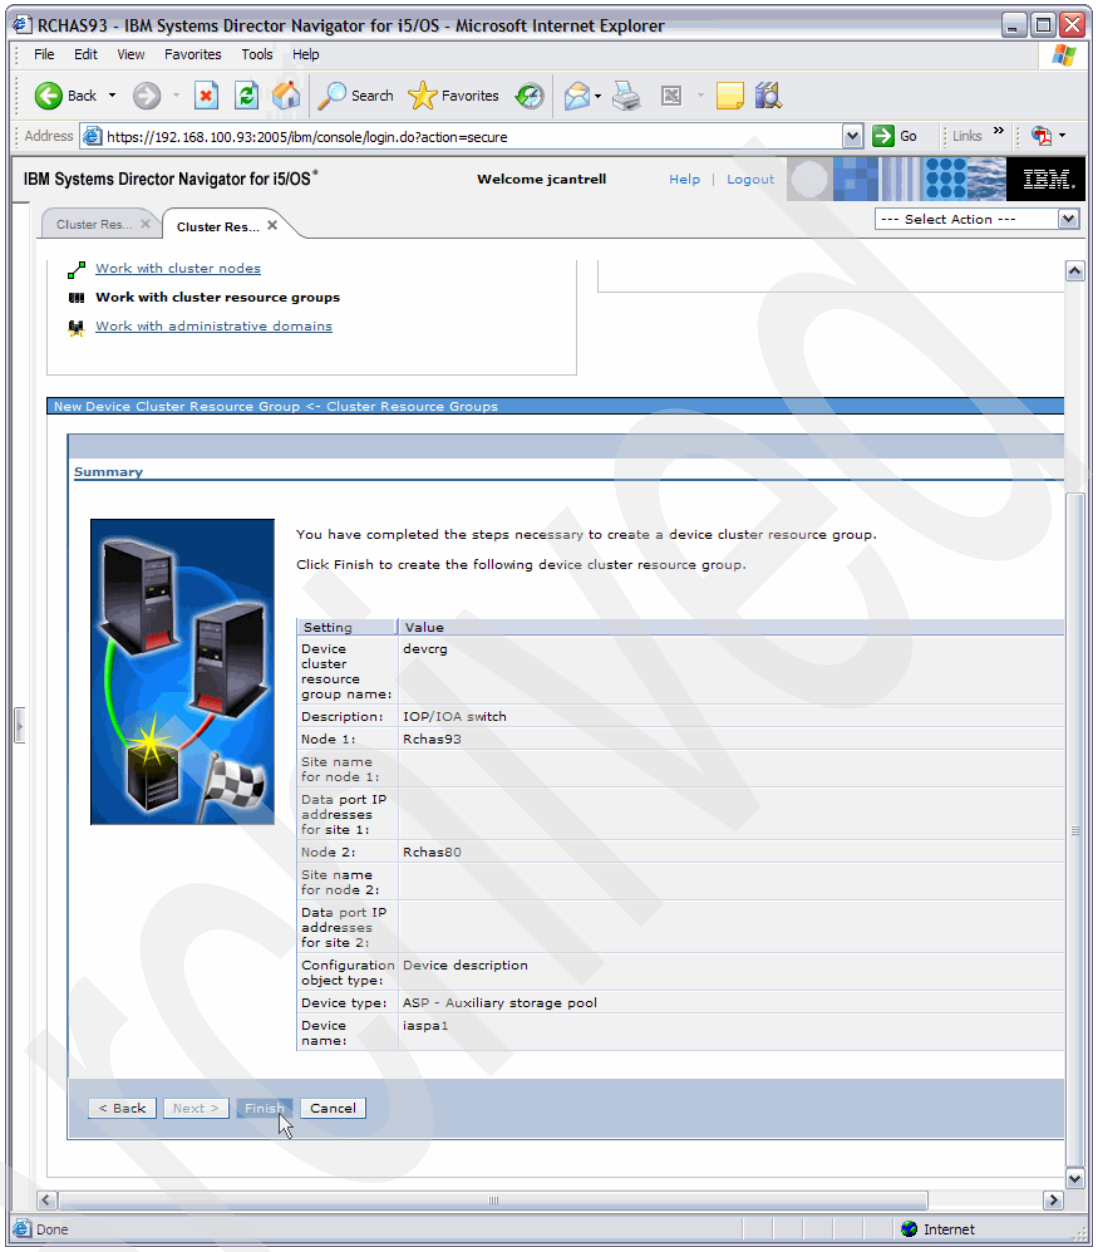


*Figure 9-36 Cluster Resource Groups → Summary*

12. Start the cluster resource group. Using the double arrow next to the CRG name, select **Start** (Figure 9-37).

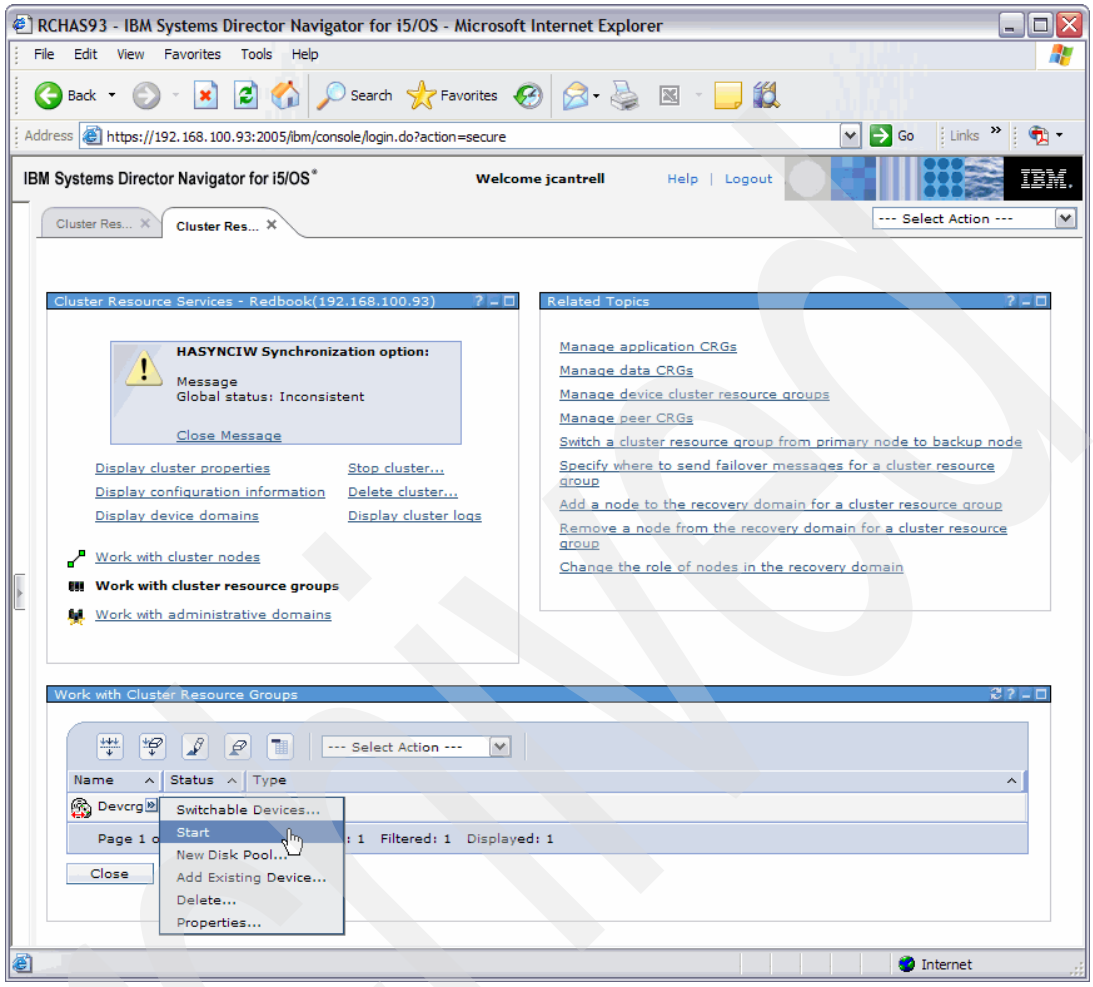


*Figure 9-37 Work with Cluster Resource Groups → Start*

A switchable disk pool can be made available on any node in the recovery domain of the cluster resource group if it is switched to that node first.

A

# Prerequisite steps

Before you can create independent auxiliary storage pools (IASPs), you must perform these simple prerequisite steps. These steps have to do with ensuring that Dedicated Service Tools (DST) is set up properly and can be accessed using IBM Systems Director Navigator for i5/OS.

# Enabling and accessing disk units

You must follow these procedures before you can perform any disk pool and disk management tasks using IBM Systems Director Navigator for i5/OS. The following steps will assist you in setting the proper authorizations for DST.

1. Ensure that the IBM i user profile that you are using to access disk units has these authorities:
   - *ALLOBJ: all object authority
   - *SERVICE
2. Start DST. Refer to the information about accessing service tools using DST.
3. Sign on to DST using your service tools user ID and password.
4. When the Use Dedicated Service Tools (DST) display is shown, select option 5 (Work with DST environment) and press Enter. The Work with DST Environment display is shown.
5. At the Work with DST Environment menu, select option 6 (Service tools security data).
6. At the Work with Service Tools Security Data menu, select option 6 (Change password level). Make sure that the password level is set to Secure Hash Algorithm (SHA) encryption or password level 2, and press F12.

   **Tip:** Steps 7–11 may be done in SST.

7. At the Work with DST Environment display, select option 3 (Service tools user IDs) to work with service tools user IDs.
8. Create a service tools user ID that matches the IBM i user profile and that also has the same password in uppercase. The service tools user ID and password must match the IBM i user profile and password of the user using IBM Systems Director Navigator for IBM i. For example, if the user profile and password combination is BOB and my1pass, then the DST user ID and password combination must be BOB and MY1PASS.
9. Give this service tools user ID these authorities:
   - Disk units: operation
   - Disk units: administration
10. Press Enter to enable these changes.
11. Exit DST and start Systems Director Navigator:
    a. Open your browser.
    b. Go to:

       `http://ipaddressofyouribmi:2001`

B

# Application object planning checklist

When moving an application into independent disk pools, use this checklist for planning the various pieces of the independent disk pool puzzle. After you decide where an object type will reside within in the system, fill in the appropriate boxes. The object types listed here are the supported independent auxiliary storage pool (IASP) objects as of IBM i 6.1. See Appendix D, “Supported and unsupported objects” on page 197, for the complete listing of independent disk pool supported and unsupported object types.

# Application object location checklist

Table B-1 provides a checklist to assign the ASP location of objects.

*Table B-1 Checklist to assign the ASP location of objects*

| Object type | Attribute | Locate in *SYSBAS | Pool group name (primary pool) | Secondary pool name | Library name |
|---|---|---|---|---|---|
| *ALRTBL | | | | | |
| *BLKSF | | | | | |
| *BNDDIR | | | | | |
| *CHRSF | | | | | |
| *CHTFMT | | | | | |
| *CLD | | | | | |
| *CLS | | | | | |
| *CMD | | | | | |
| *CRQD | | | | | |
| *CSI | | | | | |
| *DIR[a] | | | | | |
| *DTAARA[a] | | | | | |
| *DTADCT | | | | | |
| *DTAQ[a] | | | | | |
| *FCT | | | | | |
| *FIFO | | | | | |
| *FILE | | | | | |
| | DFU | | | | |
| | DSPF | | | | |
| | LF | | | | |
| | PF-DTA[a] | | | | |
| | PF-SRC | | | | |
| | | | | | |
| *FNTRSC | | | | | |
| *FNTTBL | | | | | |
| *FORMDF | | | | | |
| *FTR | | | | | |
| *GSS | | | | | |
| *IGCDCT | | | | | |
| *JOBD | | | | | |

| Object type | Attribute | Locate in *SYSBAS | Pool group name (primary pool) | Secondary pool name | Library name |
|---|---|---|---|---|---|
| *JOBQ | | | | | |
| *JRN[a] | | | | | |
| *JRNRCV[b] | | | | | |
| *LIB | | | | | |
| *LOCALE | | | | | |
| *MEDDFN | | | | | |
| *MENU | | | | | |
| *MGTCOL | | | | | |
| *MODULE | | | | | |
| *MSGF | | | | | |
| *MSGQ | | | | | |
| *NODGRP | | | | | |
| *NODL | | | | | |
| *OUTQ | | | | | |
| *OVL | | | | | |
| *PAGDFN | | | | | |
| *PAGSEG | | | | | |
| *PDG | | | | | |
| *PGM | | | | | |
| *PNLGRP | | | | | |
| *PSFCFG | | | | | |
| *QMFORM | | | | | |
| *QMQRY | | | | | |
| *QRYDFN | | | | | |
| *SBSD | | | | | |
| *SCHIDX | | | | | |
| *SPADCT | | | | | |
| *SPLF | | | | | |
| *SQLPKG | | | | | |
| *SQLUDT | | | | | |
| *SRVPGM | | | | | |
| *STMF[a] | | | | | |
| *SVRSTG | | | | | |

| Object type | Attribute | Locate in *SYSBAS | Pool group name (primary pool) | Secondary pool name | Library name |
|---|---|---|---|---|---|
| *SYMLNK | | | | | |
| *TBL | | | | | |
| *USRIDX | | | | | |
| *USRQ | | | | | |
| *USRSPC | | | | | |
| *VLDL | | | | | |
| *WSCST | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |

a. When journaling, journaled objects must reside in the same independent disk pool as the journal.
b. Journal receivers may reside in the same pool group as the journaled objects, but reside in the secondary pools.

C

# The project checklist

Table C-1 provides a sample checklist for the independent auxiliary storage pool (IASP) project. This checklist is meant to be a supplement to the reading material and may not provide all the needed steps for your project.

*Table C-1 Project outline*

| **IASP enablement project outline** | | | |
|---|---|---|---|
| PLANNING | | | |
| Business needs | Document/communicate | | |
| Disk pools | | **Name:** | **Disk units:** |
| | Primary IASP[1] | | |
| | Secondary IASP[1] | | |
| | RDB[2] | | |
| Identify objects/locate[4] | | **Name:** | **IASP name:** |
| | Application | | |
| | Library[3] | | |
| | Objects[4] | | |
| | IFS | | |
| | Evaluate commands using the ASP parameter. See Appendix J, “IASP command reference” on page 231 | | |
| Database | | Join logicals[5] | |
| | | Commitment control[5] | |
| | | Triggers[5] | |
| | | Views[5] | |

| IASP enablement project outline | | | |
|---|---|---|---|
| | | Tables[5] | |
| | | DDM file[6] | |
| | | SQL[7] | |
| | | JDBC[7] | |
| | | ODBC[7] | |
| | | FTP | |
| | | QS36F | |
| | | Journals/receivers[9] | |
| Work management | | | |
| | JOBD[10] | | |
| | Library list | | |
| | | | |
| | Determine application access: | | |
| | | JOBD[10] | |
| | | SETASPGRP | |
| | | SQL[7] | |
| | Modify/create JOBDs for user profiles using INLASPGRP[10,11] | | |
| | | | |
| | Evaluate subsystems | | |
| | | Autostart job entries (AJE) | |
| | | Prestart job entries (PJE) | |
| | | Routing entries *CLS | |
| | | Job queue entries | |
| | | Communication entries (JOBD) | |
| | | Remote location name entries (JOBD) | |
| | Spool files | | |
| | | | |
| Other | SMAPP settings | | |
| | Changing a namespace[8] | Connect, SETASPGRP or JOBD | |
| | Exit programs | | |

| IASP enablement project outline | | | |
|---|---|---|---|
| | Evaluate commands that can use aspdev parm | | |
| | Journaling | | |
| | IBM profiles owned objects - identify | | |
| Education | Administrative staff | Development | |
| | | | |
| TESTING | | | |
| | Set DST to password level 2 | | |
| | Set up user ID and password (all caps) in SST | | |
| | Create the primary disk pool | | |
| | Create the secondary disk pool | | |
| | Make available | | |
| | Modify the startup program to vary on IASP at IPL | | |
| | Verify that IFS exists /diskpoolname. Check directory for Qibm And Qsys.lib | | |
| | Create new library for split content | | |
| | Create new JOBD using appropriate INLASPGRP[10,11] and library list | | |
| | Update user profiles with new JOBD | | |
| | UID/GID changes | | |
| | Change commands that use IASP parameters | | |
| | Save/restore | | |
| | Migrate IFS | | |
| | Create symbolic link to IFS | | |
| SETUP | | | |
| | Set Dst to password level 2 | | |
| | Setup user ID and password (all caps) in Sst | | |
| | Create the primary disk pool | | |
| | Create the secondary disk pool | | |
| | Make available | | |
| | Modify startup program to vary on IASP at IPL | | |
| | Verify IFS directory /diskpoolname. Check directory for QIBM and QSYS.LIB | | |
| | Create new library for split content | | |

| IASP enablement project outline | | | |
|---|---|---|---|
| | Create new JOBD using appropriate INLASPGRP[10,11] and INLLIBL parameters | | |
| | Update user profiles with new JOBD | | |
| | UID/GID changes (same across LPARS) | | |
| | Change commands with an ASP parameter | | |
| | Save/restore | | |
| | Migrate IFS | | |
| | Create symbolic link to IFS | | |

[1]Check IFS. Should be no existing link with this name.

[2]Will default to IASP primary name.

[3]Consider new lib for split content between *sysbas/app.

[4]Refer to table application object location checklist (shows supported objects).

[5]Do not cross IASP boundary.

[6]Use RDB name.

[7]Evaluate connect statements.

[8]Use the JOBD user profile where possible.

[9]Cannot cross ASP group boundaries.

[10]Do not use QDFTJOBD.

[11]Do not modify QDTFJOBD.

# D

# Supported and unsupported objects

Table D-1 lists all object types on IBM i 6.1 and whether they are supported in an independent disk pool. This table is in object type sequence.

As a general rule, no object can be referenced unless the independent disk pool is varied on and the SETASPGRP command is run to ensure that the ASP device is in the job's thread.

*Table D-1 Supported/unsupported object types in independent disk pools*

| Supported yes/no | Object type | Description | Qualifications |
|---|---|---|---|
| Yes | *ALRTBL | Alert table | If network attributes reference the alert table, this object must exist in the system disk pool. |
| No | *AUTHLR | Authority holder | |
| No | *AUTL | Authorization list | |
| Yes | *BLKSF | Block special file | |
| Yes | *BNDDIR | Binding directory | |
| No | *CFGL | Configuration list | |
| Yes | *CHRSF | Character special file | |
| Yes | *CHTFMT | Chart format | |
| Yes | *CLD | C locale description | |
| Yes | *CLS | Class | If an active subsystem references the class object, *cls must exist in the system disk pool. |
| Yes | *CMD | Command | |
| No | *CNNL | Connection list | |

| Supported yes/no | Object type | Description | Qualifications |
|---|---|---|---|
| No | *COSD | Class-of-service description | |
| No | *CRG | Cluster resource group | |
| Yes | *CRQD | Change request description | |
| Yes | *CSI | Communication-side information object | |
| No | *CSPMAP | Cross-system product map | |
| No | *CSPTBL | Cross-system product table | |
| No | *CTLD | Controller description | |
| No | *DDIR | Distributed directory | |
| No | *DEVD | Device description | |
| Yes | *DIR | Directory (qdls) | |
| No | *DOC | Document | |
| No | *DSTMF | Distributed stream file | |
| Yes | *DTAARA | Data area | |
| Yes | *DTADCT | Data dictionary | |
| Yes | *DTAQ | Data queue | |
| No | *EDTD | Edit description | |
| No | *EXITRG | Exit registration | |
| Yes | *FCT | Forms control table | |
| Yes | *FIFO | First-in-first-out special file | |
| Yes | *FILE | File | Database files that are either multiple-system database files or that have DataLink fields that are created as link control cannot be located in an independent disk pool. If an active subsystem references the file object, *FILE must exist in the system disk pool, for example, the sign-on display file. |
| No | *FLR | Folder | |
| Yes | *FNTRSC | Font resource | |
| Yes | *FNTTBL | Font mapping table | |
| Yes | *FORMDF | Forms definition | |
| Yes | *FTR | Filter | |
| Yes | *GSS | Graphics symbol set | |
| Yes | *IGCDCT | DBCS conversion dictionary | |
| No | *IGCSRT | DBCS sort table | |
| No | *IGCTBL | DBCS font table | |

| Supported yes/no | Object type | Description | Qualifications |
|---|---|---|---|
| No | *IMGCLG | Image catalog | |
| No | *IPXD | Internet packet exchange description | |
| Yes | *JOBD | Job description | If an active subsystem references the job description object, *JOBD must exist in the system disk pool, for example, autostart job entry, communication entry, remote location name entry, or workstation entry. |
| Yes | *JOBQ | Job queue | Not persistent with independent disk pool vary on/off. |
| No | *JOBSCD | Job schedule | |
| Yes | *JRN | Journal | |
| Yes | *JRNRCV | Journal receiver | |
| Yes | *LIB | Library | The library that is specified by CRTSBSD SYSLIBLE() must exist in the system disk pool. |
| No | *LIND | Line description | |
| Yes | *LOCALE | Locale | |
| No | *M36 | System 36 machine | |
| No | *M36CFG | System 36 machine configuration | |
| Yes | *MEDDFN | Media definition | |
| Yes | *MENU | Menu | Give special thought to where initial menus in user profiles are located, especially for QSECOFR. |
| Yes | *MGTCOL | Management collection | |
| No | *MODD | Mode description | |
| Yes | *MODULE | Module | |
| Yes | *MSGF | Message file | |
| Yes | *MSGQ | Message queue | If network attributes reference the message queue, *MSGQ must exist in the system disk pool. |
| Yes | *NODGRP | Node group | |
| Yes | *NODL | Node list | |
| No | *NTBD | NetBIOS description | |
| No | *NWID | Network identifier | |
| No | *NWSD | Network server description | |
| Yes | *OUTQ | Output queue | |
| Yes | *OVL | Overlay | |

| Supported yes/no | Object type | Description | Qualifications |
|---|---|---|---|
| Yes | *PAGDFN | Page definition | |
| Yes | *PAGSEG | Page segment | |
| Yes | *PDG | Printer description group | |
| Yes | *PGM | Program | If an active subsystem references the program object, *PGM must exist in the system disk pool, for example, routing entries and prestart job entries. |
| Yes | *PNLGRP | Panel group | |
| No | *PRDAVL | Product availability | |
| No | *PRDDFN | Product definition | |
| No | *PRDLOD | Product load | |
| Yes | *PSFCFG | Printer services facility configuration | |
| Yes | *QMFORM | Query form | |
| Yes | *QMQRY | Query manager query | |
| Yes | *QRYDFN | Query definition | |
| No | *RCT | Reference code translation table | |
| No | *S36 | System 36 machine description | |
| Yes | *SBSD | Subsystem description | You cannot start a subsystem whose description is located in an independent disk pool. |
| Yes | *SCHIDX | Search index | |
| No | *SOCKET | Socket | |
| Yes | *SPADCT | Spelling aid dictionary | |
| Yes | *SPLF | S | |
| Yes | *SQLPKG | SQL package | |
| Yes | *SQLUDT | User defined sql type | |
| Yes | *SRVPGM | Service program | |
| No | *SSND | Session description | |
| Yes | *STMF | Stream file | |
| Yes | *SVRSTG | Server storage | |
| Yes | *SYMLNK | Symbolic link | |
| Yes | *TBL | Table | |
| Yes | *USRIDX | User-defined index | |
| No | *USRPRF | User profile | |
| Yes | *USRQ | User queue | |

| Supported yes/no | Object type | Description | Qualifications |
|---|---|---|---|
| Yes | *USRSPC | User-defined space | |
| Yes | *VLDL | Validation list | |
| Yes | *WSCST | Workstation customization table | |

E

# SETASPGRP command

This appendix describes the set auxiliary storage pool group (SETASPGRP) command. It contains help text from IBM i 6.1 and syntax based on the IBM i 6.1 Information Center. For more information go to the Information Center and select the supported IBM i version. Type `SETASPGRP` in the Search field. You can locate the Information Center on the Web at:

`http://publib.boulder.ibm.com/iseries/`

# Details of the SETASPGRP command

The SETASPGRP command sets the auxiliary storage pool (ASP) group for the current thread. This command allows you to change the libraries in the library list for the current thread. If an ASP group was already set, this command removes the old ASP group from the current thread and sets the specified ASP group for the current thread. After the specified ASP group is set for the current thread, all libraries in the independent ASPs in the ASP group are accessible. Objects in those libraries can be referenced using regular library-qualified object name syntax. The libraries in the independent ASPs in the ASP group plus the libraries in the system ASP (ASP number 1) and basic user ASPs (ASP numbers 2 to 32) form the *library namespace* for the thread. All libraries in the library list must be in the new library namespace or the library list is not changed and the new ASP group is not set.

Restrictions:

- You must have *USE authority to all ASP device descriptions in the ASP group and to all the specified libraries in the library list before the library namespace and the library list are changed. If you are not authorized to an ASP device description or to one of the libraries, the ASP group is not set and the library list is not changed.
- When *CURUSR is specified for the ASPGRP or USRLIBL parameter, you must have *READ authority to the job description (JOBD) listed in your user profile and *EXECUTE authority to the library containing that JOBD.
- The SETASPGRP command is not allowed in the following:
  - System jobs QPFRADJ, QJOBSCD, QSYSARB, QSYSARB3, QSYSARB4, and QLUS.
  - All subsystem monitor jobs.
  - DDM, DRDA, database host server, and SQL server jobs once the initial namespace has been established.
  - Receive journal entry (RCVJRNE) and delete journal receiver (DLTJRNRCV) command exit programs.
  - Management central registered inventory gathering service (RIGS) exit programs (exit point QIBM_QYIV_INVGTRSRV).
  - Trigger or format selector programs that run as part of database I/O operations.
  - Attention programs (the PGM parameter of the SETATNPGM command).
  - Break handling programs (the PGM parameter of the CHGMSGQ command).
  - Programming development manager (PDM) functions.

Table E-1 shows the SETASPGRP parameters.

*Table E-1 SETASPGRP parameters*

| Keyword | Description | Choices | Notes |
|---|---|---|---|
| **ASPGRP** | ASP Group | *Name*, *CURUSR, *NONE | Required, position 1 |
| **SYSLIBL** | System library list | **CURSYSBAS*, *SYSVAL | Optional |
| **CURLIB** | Current library | *Name*, **CURSYSBAS*, *CURUSR, *CRTDFT | Optional |
| **USRLIBL** | Libraries for current thread | Single values: **CURSYSBAS*, *CURUSR, *SYSVAL, *NONE, Other values (up to 250 repetitions): *Name* | Optional |

The parameters include:

- ASP group (ASPGRP)

  This specifies the name of the ASP group to set for the current thread. The ASP group name is the name of the primary ASP device within the ASP group. The possible values are:

  **Name** This value specifies the name of the primary ASP in the ASP group to be set for the current thread. All libraries from all ASPs in this ASP group are included in the library namespace.

  ***CURUSR** The ASP group is set to the value defined for the initial ASP group in the default job description of the user profile under which the thread is currently running.

  ***NONE** This value specifies for the current thread to have no ASP group. The library namespace does not include libraries from any ASP group. Only the libraries in the system ASP and any basic user ASPs are in the library namespace.

- System library list (SYSLIBL)

  This parameter specifies the system part of the library list for the thread in which the command is entered. The possible values are:

  ***CURSYSBAS** The libraries in the system part of the library list of the current thread that are found in the system ASP (ASP number 1) or any configured basic user ASP (ASP numbers 2 to 32) are used as the new system part of the library list.

  ***SYSVAL** The system part of the library list is set from the current value of system value QSYSLIBL.

- Current library (CURLIB)

  This parameter specifies the library to be used in the current library entry of the library list for the thread. If *CURUSR or a library name is specified and the library cannot be found in the new library namespace, an error message is sent and the library list and ASP group are not changed. The possible values are:

  ***CURSYSBAS** The library name in the current library entry of the library list is used as the new current library if the library is found in the system ASP (ASP number 1) or any configured basic user ASP (ASP numbers 2 to 32). If the library name in the current entry is not found in the system ASP or any basic user ASP, the current library entry is removed from the library list.

  ***CURUSR** The current library is set to the value defined for current library in the user profile under which the thread is currently running.

  ***CRTDFT** This value changes the library list to remove any name from the current library entry. If objects are created specifying *CURLIB for the library name, library QGPL is used.

  **Name** This value specifies the name of the library that replaces the current library entry in the library list.

- Libraries for current job (USRLIBL)

  This parameter specifies the libraries that are placed in the user part of the library list. If *CURUSR or a list of library names is specified and any of these libraries cannot be found in the new library namespace, an error message is sent and the library list and ASP group are not changed. The possible values are:

  ***CURSYSBAS** The libraries in the user part of the library list of the current thread that are found in the system ASP (ASP number 1) or any configured basic user ASP (ASP numbers 2 to 32) are used as the new user part of the library list.

  ***CURUSR** The user part of the library list for the thread is set to the value defined for the Initial library list in the default job description of the user profile under which the thread is currently running.

  ***SYSVAL** The user part of the library list is set from the current value of system value QUSRLIBL.

  ***NONE** This value changes the user part of the library list to remove all library names.

  **Name** This value specifies the names of the libraries to be used as the user part of the library list, in the order in which they are to be searched. The number of libraries that can be specified ranges from 1 through 250.

# Examples of using the SETASPGRP command

The following sections provide some examples of using the SETASPGRP command.

## Example 1: set new ASP group

The command in Figure E-1 will set the ASP group for the thread in which the command runs to be WAREHOUSE1. This changes the library namespace for the thread to include all libraries in any of the independent ASPs in the ASP group identified by the independent ASP device named WAREHOUSE1.

The system part of the library list is set from the system value QSYSLIBL. The current library entry of the library list is set from the Current library value defined in the user profile under which the thread is currently running. The user part of the library list is set using the current user part of the library list and by removing any libraries that are not found in the system ASP or configured basic user ASPs.

```
SETASPGRP   ASPGRP(WAREHOUSE1) SYSLIBL(*SYSVAL)  CURLIB(*CURUSR)
            USRLIBL(*CURSYSBAS)
```

*Figure E-1 Setting a new ASP group*

### Example 2: set to no ASP group

The command in Figure E-2 removes any ASP group for the thread in which the command runs. This changes the library namespace for the thread to include only those libraries in the system ASP (ASP number 1) and basic user ASPs (ASP numbers 2 to 32).

```
SETASPGRP   ASPGRP(*NONE)  SYSLIBL(*CURSYSBAS)  CURLIB(*CRTDFT)
            USRLIBL(*NONE)
```

*Figure E-2 Set to no ASP group*

The system part of the library list is set using the current system part of the library list and by removing any libraries that are not found in the system ASP or configured basic user ASPs. The current library entry of the library list is changed to be empty. This causes library QGPL to be used as the current library. The user part of the library list is changed to be empty.

## Error messages

The SETASPGRP command can receive *Escape message `CPFB8E9 ASP group &1 not set for thread &2.`

F

# Migrating Integrated File Systems (IFS)

Migrating IFS can be done using the IBM Systems Director Navigator for i5/OS. When an independent auxiliary storage pool (IASP) is made available (varied on) an IFS directory is created in the `/root` using the disk pool name. We are moving the `/flights` directory from `/root` to the `/IASPA1` directory (our primary IASP).

To do this:

1. Log into Director Navigator and select **File Systems** (Figure F-1).

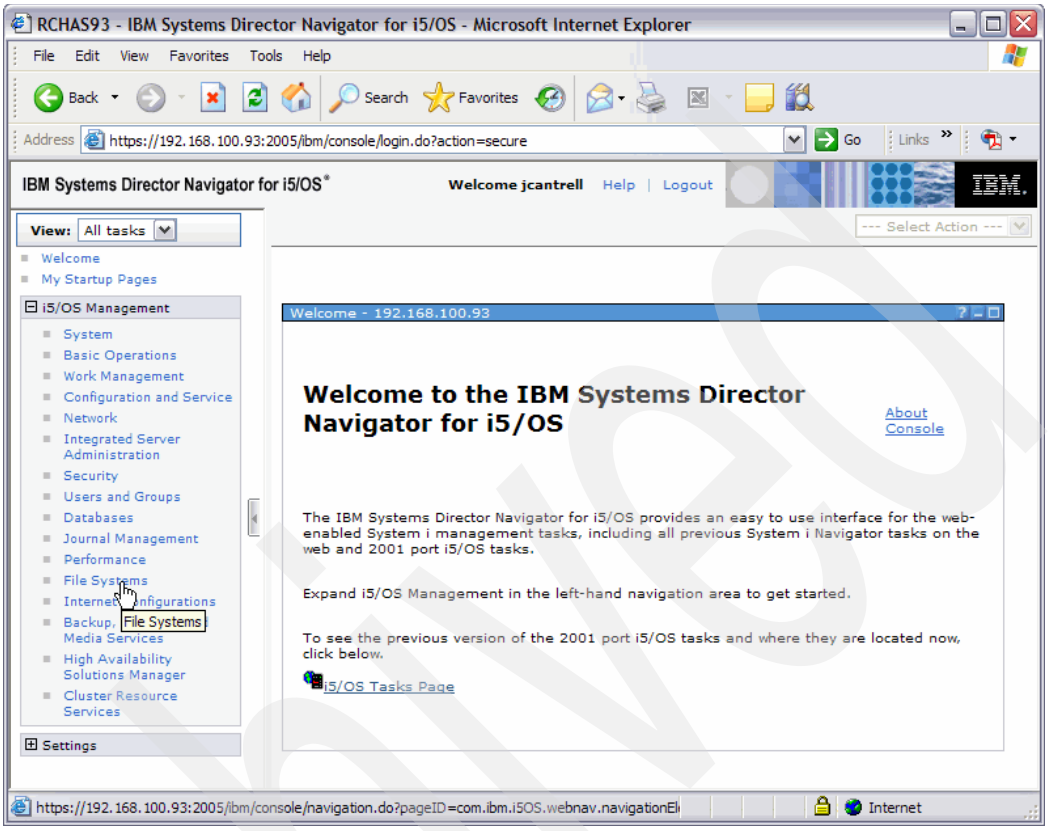


*Figure F-1 Director Navigator → i5/OS Management → File Systems*

2. Click **Integrated File Systems** (Figure F-2).

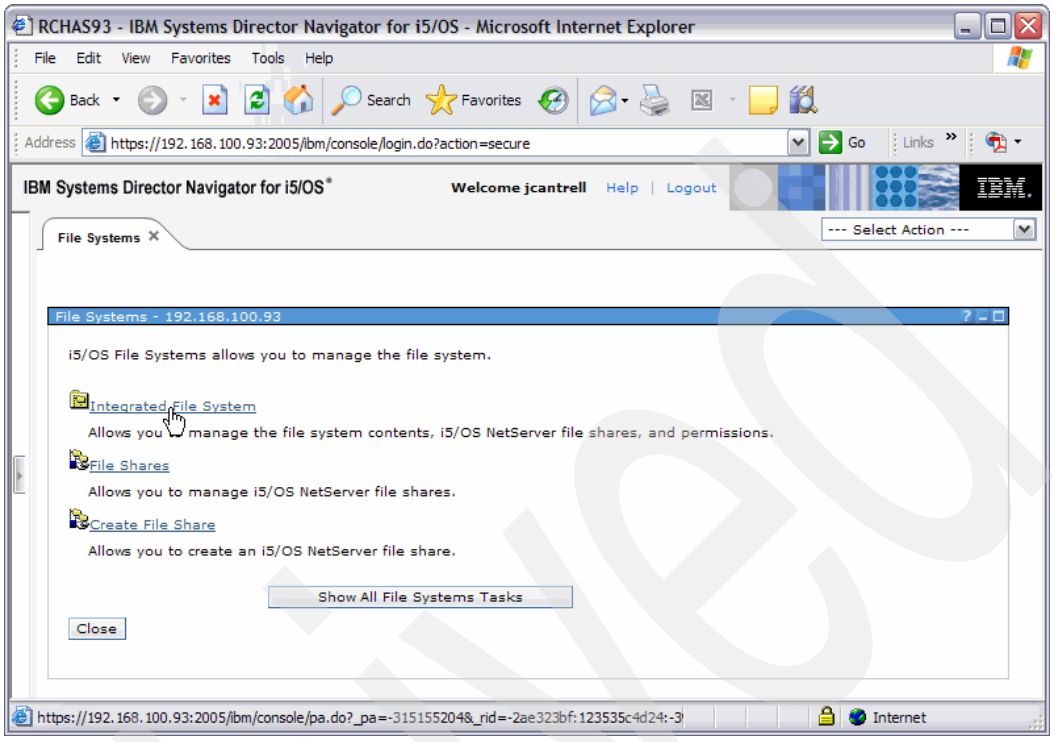


*Figure F-2 Files Systems → Integrated File Systems*

3. Select and open the `root` directory (Figure F-3).

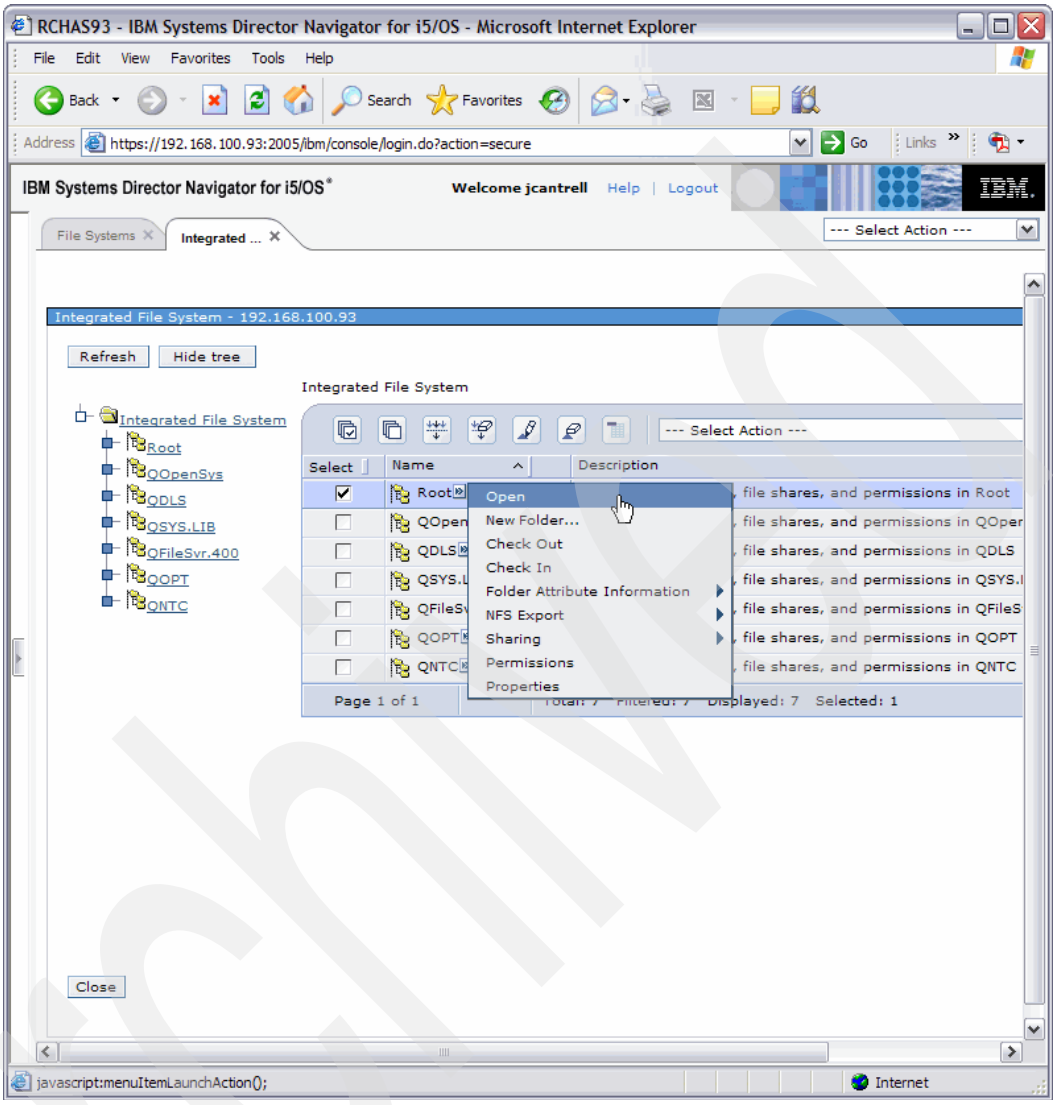


*Figure F-3 File Systems → Integrated File Systems (IFS) → Select Root → Open*

4. To copy the flights directory, select **flights** and open the pull-down menu. On the pull-down select **Copy** (Figure F-4).

> **Tip:** Do not use the move function. A move changes the pointers, not the disk location of the IFS. Using a copy will physically move the location of the IFS to the IASP.

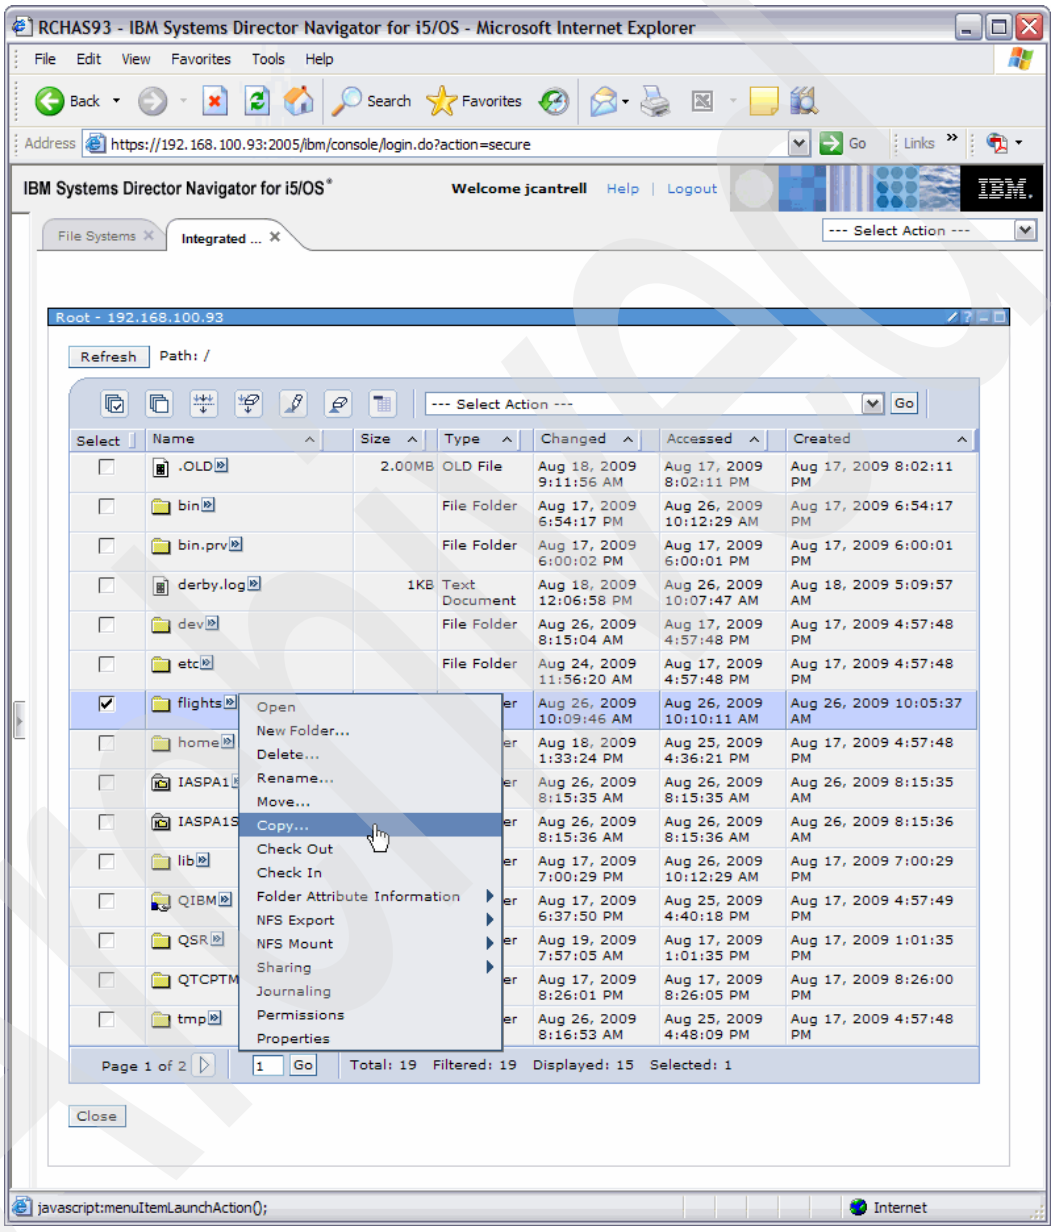


*Figure F-4 File Systems → IFS → Select Flights → Copy*

5. The copy window appears. Enter your copy to the location. We are copying flights to disk pool IASPA1. Enter `/IASPA1` and click **Copy** (Figure F-5).

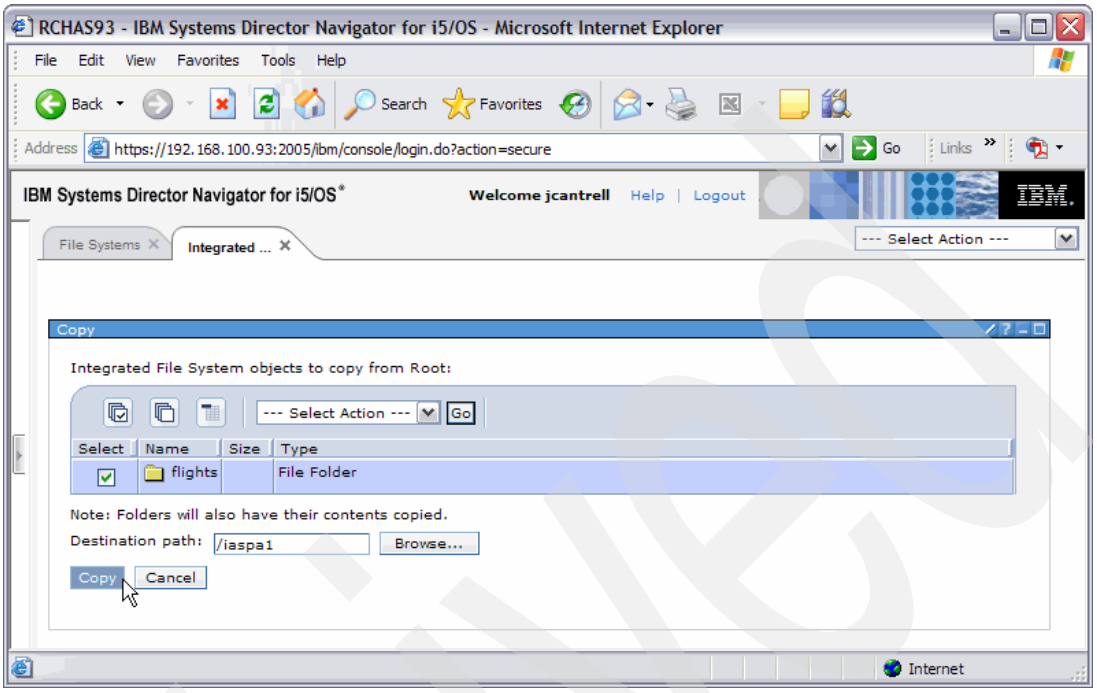


*Figure F-5 File Systems → IFS → Copy*

6. After the copy completes, open the IASP directory (Figure F-6).

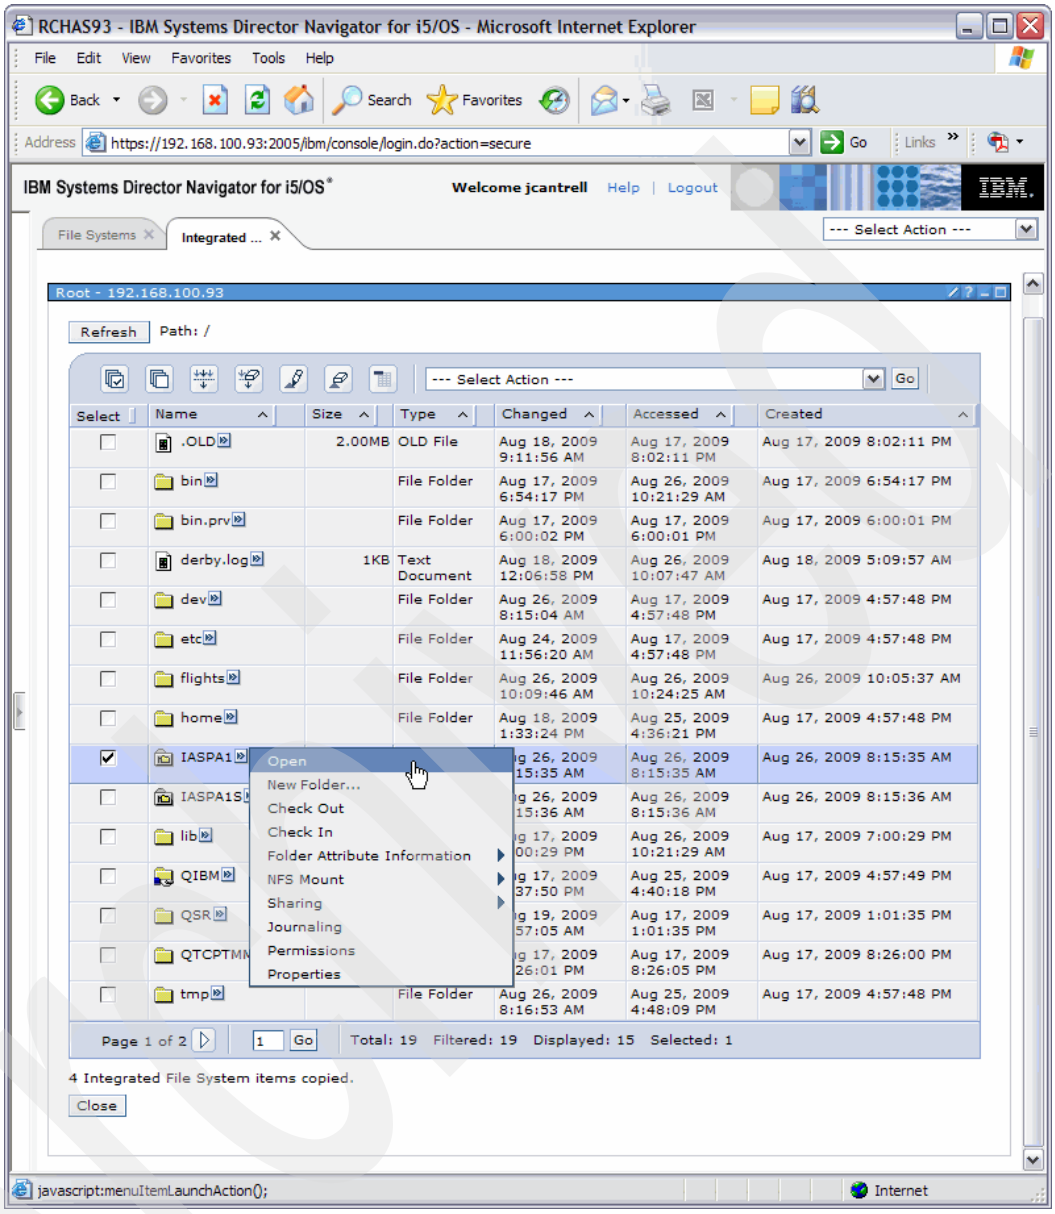


*Figure F-6 File Systems → IFS → Path → Open*

7. Verify that the copied directory does exist in the IASP directory (Figure F-7).

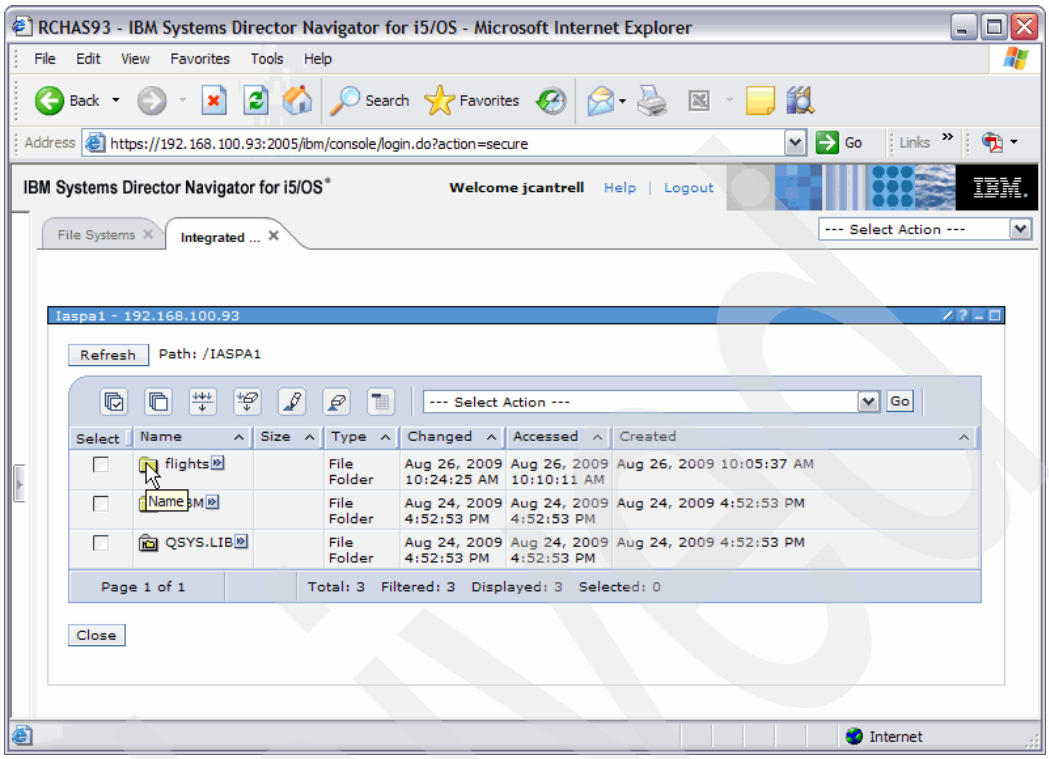


*Figure F-7 File Systems → IFS → Path → Flights*

8. Once you verify the proper placement of the directory, delete the directory in the system ASP (Figure F-8).

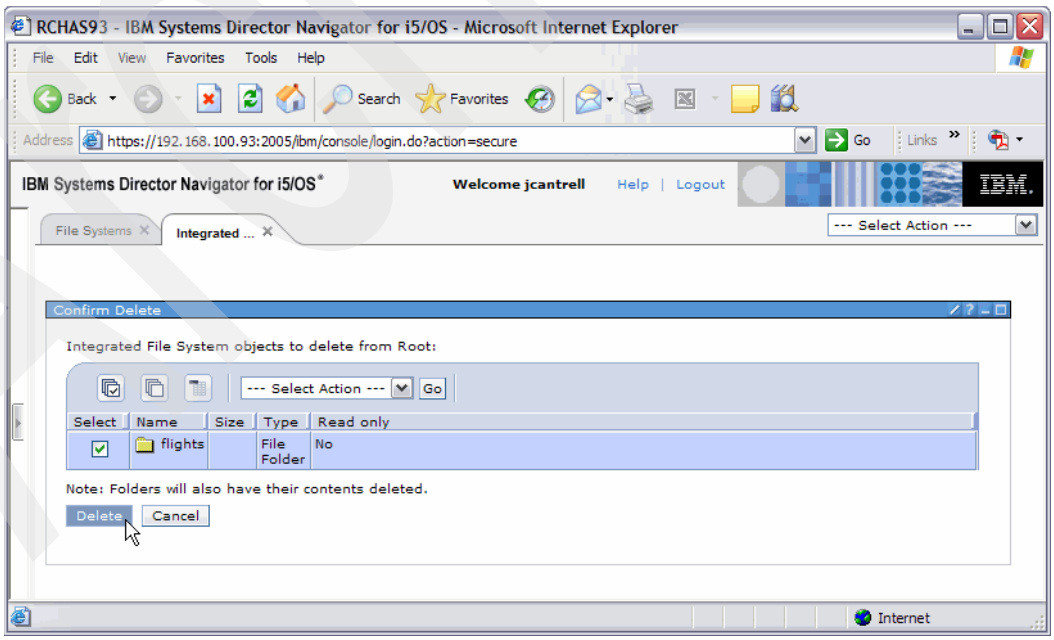


*Figure F-8 Files Systems → IFS → Path → Confirm Delete*

9. You can create a symbolic link using the ADDLNK command. The object parameter is the IASP path. The new link is the old path (Figure F-9).

```
 Add Link (ADDLNK)

 Type choices, press Enter.

 Object . . . . . .   /iaspa1/flights

 New link . . . . . .   /flights

 Link type  . . . . . .   *SYMBOLIC     *SYMBOLIC, *HARD









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

*Figure F-9 ADDLNK command*

G

# Space and timing considerations

If the maximum allowed storage (MAXSTG) parameter in the user profile is set to *NOMAX, then no further action is required. The system automatically expands storage as necessary. If this parameter in the user profiles has a value in it, then you must make considerations for expansion of that size.

For each user profile on the system, an amount of storage can be set. If a value is set for that parameter, that same value can be used in each independent disk pool as it is created or made available on the system.

If the system ASP has a large number of user profiles, or the system has a large number of objects, consider space for this storage when configuring the independent disk pool. The vary on time is affected also.

## Accounting for space used by user profiles and authorization lists

The total storage attributed to an individual user profile is the amount of the storage allowed in each online independent disk pool and by the user in the local system and basic user ASPs. Use the display user profile (DSPUSRPRF) command to display the value of this storage. If the system has *n* independent disk pools, then the total storage that a user profile is allowed to use is the result of:

`((`*n*`+1) *MAXSTG)`

To prevent disruptions to a system, the limit for the maximum amount of auxiliary storage that can be allocated by a user profile in an independent disk pool is not enforced when bringing that independent disk pool online. That is, if bringing an independent disk pool online causes a user profile to exceed its storage limit for that independent disk pool, the storage limit is allowed to exceed the specified amount. Subsequent requests for additional storage cause a `User Profile Storage Limit Exceeded` exception message.

## Timing considerations

This section outlines the timing considerations for independent disk pools:

- The time to vary on an independent disk pool during the switching process depends on the number of objects on the *SYSBAS, not the size of the database objects. If possible, keep the number of objects small.
- For a shorter vary on or off, keep the UID and GID of user profiles that own objects on the independent disk pool the same between nodes of the cluster. Having different UIDs lengthens the vary-on time.

  **Important:** Having different UIDs between systems lengthens the vary-on time substantially because UIDs that do not exist must be created.

- All the disk units within a tower are switched as an entity. When a tower containing the independent disk pool is switched, all other IOPs and devices on that tower (if any), such as tape drive, CD ROM, printers, and so on, are also switched to the other node.
- The number of devices in a tower affects the switch-over time. The devices in a switch-over are reset as part of the switch-over process.
- The devices within a tower that are switched can auto-configure. This can result in different resource names on the switched-to node. Manually configure the devices on the node that the tower is switched to so that the resource names match on both systems.

H

# Using virtual tape on IBM i

Virtual tape supports all save/restore commands and APIs (except SAVSTG). Using virtual tapes and media can be faster than saving directly to tape because it offers similar performance to saving to save files. The best performance results when the virtual tape is configured in a separate ASP or independent auxiliary storage pool (IASP).

## Key advantages of using virtual tape

Using virtual tapes eliminates the following save file limitations:

- One library per save file
- SAVSYS not supported on save file
- Parallel saves not supported on save file
- 1 TB size limitation on save file

Using virtual tape also eliminates the following media error limitations:

- Saves ending because of tape device or media errors.
- Save-while-active checkpoint restriction.
- Once the checkpoint is reached, saves cannot be restarted.

Duplicate saves to media (DUPTAP or DUPMEDBRM) can occur:

- When tape devices are available
- At your convenience

On-site and off-site storage is available to:

- Keep virtual volumes on systems, as needed
- Keep duplicated volumes off-site

## Consider before using virtual tape

Consider the following:

- Additional DASD is required. If doing a full system save to virtual tape on your 100 GB system, 45% of free DASD is required to do the save.

  **Note:** For ASPs smaller than 100 GB, virtual tape will not allocate more than 95% of the ASP storage. For ASPs larger than 100 GB, virtual tape will not allocate additional storage once the amount of free space reaches 5 GB. When unsure if there is enough free storage to support the image size, we recommend that *IMGSIZ be used. In summary, the lesser of 5% or 5 GB of DASD storage remains unallocated.

- Your system cannot be installed from a SAVSYS on virtual media D-IPL and can be done only from physical media (tape or CD/DVD).
- If using ASPs and UDFS, make sure that the UDFS is mounted when the image catalog is created.
- Using virtual tape may not always be faster because of:
  - Tape technology
  - System configuration and environment
- Data compaction is not supported on virtual media.

  **Note:** When running the DUPTAP or DUPMEDBRM command, compaction does not occur. This is because the default for data compaction on the DUPTAP command is set to *FROMFILE. BRMS does not currently have this option.

- If using data compression, this causes a significant performance impact. The DTACPR parameter on the save commands default is *DEV or *NO. If this value is changed to *YES, SNA low data compression is used.
- Up to 256 1,000,000 MB virtual tape volumes can be stored in a single image catalog.
- There can be up to 35 virtual tape devices on a system or partition.
- Virtual tape volumes are stored as Integrated File System objects on the System i system.
- When the following PTFs or their superseding PTFs are applied, the behavior of the DUPTAP command when the source device is virtual tape and COMPACT(*FROMFILE) is specified will be changed. COMPACT(*FROMFILE) would be treated the same as COMPACT(*YES) when the input device is virtual tape. COMPACT(*NO) could still be used to turn off compaction.
- The PTFs affecting this and DUPTAP performance at V5R4M0 are:
  - SI24891
  - MF39598
  - MF40282
  - MF39601
- The reason for the change is that without compaction there is no noticeable performance improvement with the PTF and we do not want to require that customers using virtual tape to change the way that they do things before they get a benefit from the PTF.
- When any image catalog is in a READY status and the WRKIMGCLG command and a *LINK, SAVBRM, or SAV is run for that directory, the stream files are not saved. If the image catalog is in a not ready status, the image catalog entries will be saved and this could result in a significant increase in the amount of data being saved. After creating virtual tape images, they are automatically included when performing a full system save using GO SAVE, Option 21. The virtual tape images could significantly increase the time that it takes to complete the Option 21 save operation, even if the image catalog entries do not contain data. To exclude the virtual images from a full system save, use one of the following strategies:
  - Use the change attribute (CHGATR) command to mark the image catalog directory as non-saveable, for example, CHGATR OBJ('/Catalog-Path') ATR(*ALWSAV) VALUE(*NO).
  - Use the load image catalog (LODIMGCLG) command to make the image catalog ready. Image catalogs with a ready status are omitted from the save. In an attended save, omitting the image catalog directories can be specified on the save object (SAV) command.
- We recommend that PTF MF41110 be loaded and applied before using virtual tape. More information can be found in APAR MA34516. The initialize tape (INZTAP) command makes any pre-existing data on the virtual tape volume inaccessible by the virtual tape device. Running the INZTAP command and specifying CLEAR(*YES) can be used to delete any existing data in a virtual tape volume. However, this should be used only if there are security concerns with the existing data because this operation can take a long time and uses significant system resources.
- We recommend that PTF MF41110 be loaded and applied before using virtual tape. More information can be found in APAR MA34516.

## Using virtual tape from other partitions or servers

The host partition must mount the virtual tape volume before it can be used by another partition or iSCSI attached sever. (This is for Linux guest partitions and is not supported on AIX® guest partitions.) To use virtual tape from other partitions, do the following:

1. Ensure that the virtual tape device description has the unload device at vary off parameter set to the value *NO.
2. Mount the virtual volume in the virtual tape device using the image catalog commands.
3. Make sure that the virtual volume density is *VRT256K.
4. Vary off the virtual tape device.

> **Note:** The virtual volume is still mounted. The other partition or the iSCSI-attached server can now use the previously mounted virtual tape volume through virtual I/O to the virtual tape device. Only the mounted volume can be used. Volume spanning is not supported.

The virtual tape device is configured as a 3580 model 002 tape device to the other partitions or iSCSI-attached servers.

> **Note:** The image catalog commands should not be used to remove or change the virtual tape volumes while they are being used by other partitions or iSCSi-attached servers. Only the host operating system partition can unload (eject) the mounted volume.

## Density

The following is a discussion about the tape density:

- The density (format) parameter limits the block size that can be written to a virtual tape volume so that the volume will be compatible for duplication to your physical tape device.
- Volumes with a density of *VRT256K use an optimum block size of 256 KB and are compatible with 35xx type devices and the newer QIC tape devices.
- Volumes with a density of *VRT240K use an optimum block size of 240 KB and are compatible with VXA and 8 mm technology devices and the 35xx and newer QIC devices.
- Volumes with a density of *VRT64K use an optimum block size of 64 KB and are compatible with 3490F Model 18 track media, VXA and 8 mm technology devices, and with the 35xx and newer QIC devices.
- Volumes with a density of *VRT32K do not use an optimum block size and are compatible with all devices.
- The virtual tape media cannot be copied to physical media if you choose an incompatible block size. Ensure a virtual tape density with a block size that is compatible with the physical tape devices on your system.

## BRMS

BRMS does not support:

- Dependent catalogs.
- Auto creation of volumes because the volume must be enrolled in the BRMS inventory before it can be used. Message CPF3742 will be posted and the job will end.

Consider the following when using virtual media in BRMS:

- A create virtual tape devices outside of BRMS must be created using the CRTDEVTAP command.
- An image catalog and virtual volumes outside of BRMS must be created using the following commands:
  - CRTIMGCLG
  - ADDIMGCLGE

## Disaster recovery

For disaster recovery, do the following:

1. Duplicate virtual tape images to physical Media (DUPMEDBRM). When this is done, the physical tapes become the original saved volume and the virtual volume becomes the duplicate saved volume.
2. We recommended that you specify append to media *NO when saving to virtual media.
3. To make recovery faster and easier, and using parallel saves, we recommend that you keep virtual tape size such that it will fit to one physical tape, so you would get a one-to-one relationship when a DUPMEDBRM is run.
4. If you used parallel saves, need to recover from physical media, and do not have the same amount of drives used during the save, a lot of tape changing will need to be done.
5. Virtual media volumes cannot normally be shared in a BRMS network. However, it is possible using iASPs to share the volumes if the iASP is switched from one system to the other.

## Messages related to virtual tape

The following messages are related to virtual tape:

| | |
|---|---|
| **CPF41B0** | Incorrect image catalog name specified. |
| **CPF41B3** | No more volumes to mount from catalog. You specified VOL(*MOUNTED) and the last accessible virtual tape volume in the catalog was already used and unloaded. |
| **CPF41B4** | Virtual tape volume not available. The specified volume is in unload status. |
| **CPF41B5** | Virtual tape volume not found. |
| **CPF6760** | Device &1 is not ready. This error is reported when the requested virtual volume could not be mounted. Typically, this error occurs when the requested virtual volume is already mounted in a different virtual |

device by a reference or a dependent image catalog and the mount is not allowed.

**CPF67F5** Duplicate cartridge or virtual volume name found.

**CPF4373** End of media on device. This message is sent when the storage threshold of the ASP containing the virtual tape volume exceeds the maximum allowed storage for virtual tape, or when the maximum storage allowed threshold of the user profile that owns the virtual tape is exceeded and additional storage cannot be allocated to continue the output operation. The maximum allowed storage for virtual tape is the greater of 95% or 5 GB of free space remaining in the ASP.

**CPF4371** Device not operational. This message is sent when an unexpected error occurs. Possible causes are:

– The virtual volume stream file was removed while in use.

– A force vary off was performed on an IASP containing a virtual tape volume that was in use.

It is necessary to vary the virtual tape device description off and back on with RESET(*YES) to clear the error.

A diagnostic message related to virtual tape and media is:

CPDBC04 Error on command &3 during virtual tape function &2. Reason code 26. Reason code 26 is sent when a command to mount or change a virtual volume could not be processed because the volume is currently mounted in a device. The volume will need to be unloaded or unmounted from the device it is currently in before the command can be completed.

**Note:** If the virtual volume was used by another partition and left mounted, there may not be any image catalogs that show that the volume is in mounted status. The CHKTAP command with ENDOPT(*UNLOAD) can be used to unload the virtual tape volume.

I

# APIs pertaining to independent disk pools

This appendix lists the APIs that you can use to query or manipulate independent disk pool information.

## QYASPOL

The open list of ASPs (QYASPOL) API generates a list of ASPs or information about an ASP. The information may include:

- Identification of all ASPs configured to a system
- Attributes of an ASP
- Unassigned disk units or disk units assigned to an ASP
- Hardware problems during vary-on of an independent ASP
- Current vary-on activity
- Jobs using an independent ASP

Upon successful completion of this API, a handle is returned in the list information parameter. You may use this handle on subsequent calls to the following APIs:

- Get list entries (QGYGTLE)
- Close list (QGYCLST)

This API resides in IBM i 6.1, 5761-SS1, BOSS Option 12 (Host Servers) in library QGY.

## QGYCLST

The close list (QGYCLST) API closes a previously opened list. Any internal storage associated with that list is freed. The handle specified on the call to this API is no longer valid after the call completes. The handle is generated by one of the following list APIs:

- Open list of jog log messages (QGYOLJBL)
- Open list of messages (QGYOLMSG)
- Open list of objects (QGYOLOBJ)
- Open list of printers (QGYRPRTL)
- Open list of spooled files (QGYOLSPL)
- Open list of ASPs (QYASPOL)

## QGYGTLE

The get list entries (QGYGTLE) API allows requests to get entries from previously opened lists on the iSeries server. A list exists if an initial request is already made and the list is not closed using the close list (QGYCLST) API.

Initial requests are made by calling the following APIs:

- Open list of job log messages (QGYOLJBL)
- Open list of messages (QGYOLMSG)
- Open list of objects (QGYOLOBJ)
- Open list of printers (QGYRPRTL)
- Open list of spooled files (QGYOLSPL)
- Open list of user certificates (QSYOLUC)
- Open list of validation list entries (QSYOLVLE)
- Open list of ASPs (QYASPOL)
- Retrieve objects secured by authorization list (QGYRATLO)

## QHSMMOVL

The move library to ASP (QHSMMOVL) API moves a library and its contents from its existing ASP to the specified target ASP through a save and restore process. The API, however, preserves private authorities to the objects that are normally lost with a save and restore operation.

Program (*PGM) objects in the library will be placed in library QRPLOBJ (or library QRPL*xxxxx* if the library is in a primary or secondary ASP (where *xxxxx* is the ASP number of the primary ASP of the ASP group)) and a copy of each *PGM object will be moved with the library to the target ASP. After a library has been moved, the following attributes are changed for every object in the library and for the library:

- The date last used will be set to blank.
- The change date and time will be set to the current date and time.
- The days used count will be set to zero.
- The date use count reset will be set to blank.
- The restore date and time will be set to the current date and time.

The move library to ASP (QHSMMOVL) API has the following restrictions:

- Data queue entries are not moved and will be lost.
- The QSYSWRK subsystem must be active.
- A library cannot be moved in the following cases:
  - The library is considered a system library. The restricted libraries are:
    - Any library with a name that begins with the letter Q.
    - Any library with a name that begins with the symbol # that is not considered a user library and the target ASP is not a basic user ASP.
    - Libraries SYSIBM, SYSIBMADM, SYSPROC, and SYSTOOLS.
    - Library SYSIB*xxxxx* (where *xxxxx* is the number of a primary ASP).
  - The library cannot be renamed. See the rename object (RNMOBJ) command for the restrictions on renaming a library.
  - The library contains *JRN or *JRNRCV objects, objects that are journaled, or the library itself is journaled.
  - The library contains files with database dependencies outside the library.
  - 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 target ASP does not have enough space for the library and its objects.
  - The target ASP is not the system ASP (ASP 1), a basic user ASP (ASPs 2–32), or a primary or secondary ASP in the range f 3–255. A library object (*LIB) must be able to be created in the target ASP, which means that the ASP cannot be a UDFS ASP and it cannot contain a journal, journal receiver, or save file object where the object's ASP is a basic user ASP and the object's library is in the system ASP.
  - When the target ASP is a primary or secondary ASP, the library contains a job queue object (*JOBQ).
  - When the target ASP is a primary or secondary ASP, the library must contain only object types that can reside in an ASP.

J

# IASP command reference

Table J-1 on page 232 shows all commands where the ASPGRP, ASPDEV, and INLASPGRP parameters are used. There are numerous commands using the ASP parameter. These can also be used in the independent auxiliary storage pool (IASP) environment (see tip below). Commands looking for a specific ASP number relate to a user ASP (UDFS) and are not included.

**Tip:** Using the SETASPGRP command or JOBs will set the proper namespace for each user.

*Table J-1 Commands using various ASP parameters*

| Function | IBM i 6.1 command | Parameter |
|---|---|---|
| **AJS/JS functions** | **JS commands** | |
| The add job using job scheduler command allows you to schedule batch jobs by adding an entry to the job schedule. | ADDJOBJS | *INLASPGRP |
| The change job using job scheduler command allows you to change an entry in the jobs schedule. | CHGJOBJS | INLASPGRP |
| The submit job using job scheduler allows you to submit a batch job through IBM Advanced Job Scheduler. | SBMJOBJS | INLASPGRP |
| **ASP functions** | **ASP commands** | |
| The change auxiliary storage pool attributes command allows you to change attributes that control the behavior of an auxiliary storage pool (ASP). | CHGASPA | ASPDEV |
| The change ASP activity command can be used to suspend database transactions and database and IFS file change operations for the system and configured ASPs or IASPs. | CHGASPACT | ASPDEV |
| The change device description command changes the device description for an ASP device. | CHGDEVASP | |
| The check ASP balance command allows you to check which ASP balance function is currently active and which units have been marked to not allow new allocations (*ENDALC). | CHKASPBAL | |
| The create device description command creates a device description for an ASP device. | CRTDEVASP | |
| The display ASP status command shows the vary progress status of an ASP device when bringing the device online or taking it offline. | DSPASPSTS | ASPDEV |
| The end ASP balance command allows you to end the ASP balance function that was started using the start ASP balance CL command. | ENDASPBAL | ASPDEV |
| The set auxiliary storage pool group command sets the ASP group for the current thread. This command also allows you to change the libraries in the library list for the current thread. | SETASPGRP | ASPGRP |
| The start ASP balance command allows you to start the ASP balancing function for one or more ASPs. | STRASPBAL | ASPDEV |
| The trace ASP balance command controls the function that gathers the ASP usage statistics. | TRCASPBAL | ASPDEV |
| The work with ASP descriptions command takes you to the work with ASP descriptions display or produces the ASP descriptions report. | WRKASPBRM | |
| The work with ASP jobs command allows you to work with a list of jobs that are using an ASP. | WRKASPJOB | ASPDEV |
| **AUT functions** | **AUT commands** | |
| The restore authority command restores the private authorities to user profiles. | RSTAUT | ASPDEV (SAV/RST) |

| Function | IBM i 6.1 command | Parameter |
|---|---|---|
| **DSK functions** | **DSK commands** | |
| The print disk information command is used to print disk space information that was stored in database file QAEZDISK or QAEZD*nnnnn* by the retrieve disk information command, where *nnnnn* is the ASP number of the independent ASP (IASP) for which disk space information was retrieved. | PRTDSKINF | ASPDEV |
| The retrieve disk information command is used to collect disk space information. | RTVDSKINF | ASPDEV |
| The start disk reorganization command allows you to start the disk reorganization function for one or more ASPs. | STRDSKRGZ | ASPDEV |
| **JOB functions** | **JOB commands** | |
| The retrieve job attributes command is used in a CL program to retrieve the values of one or more job attributes and place the values in the specified CL variable. | RTVJOBA | ASPGRP |
| The submit job command allows a job that is running to submit another job to a job queue to be run later as a batch job. | SBMJOB | INLASPGRP |
| **JOBD functions** | **JOBD commands** | |
| The change job description command changes the job-related attributes specified for a job description object through the create job description command. | CHGJOBD | INLASPGRP |
| The create job description command creates a job description object that contains a specific set of job-related attributes that can be used by one or more jobs. | CRTJOBD | INLASPGRP |
| **JRN functions** | **JRN commands** | |
| The create journal command creates a journal as a local journal with the specified attributes and attaches the specified journal receiver to the journal. | CRTJRN | *use SETASPGRP* |
| The create journal receiver command creates a journal receiver. After a journal receiver is attached to a journal, journal entries can be placed in it. | CRTJRNRCV | *use SETASPGRP* |
| **LIB functions** | **LIB commands** | |
| The clear library command deletes all of the objects from the specified library that you have the authority to delete. | CLRLIB | ASPDEV |
| The create library command adds a new library to the system. | CRTLIB | ASPDEV |
| The delete library command deletes a specified library from the system after all objects in the library are deleted. If a library that is deleted contains objects, this command first deletes all of the objects and then deletes the library. | DLTLIB | ASPDEV |
| The display library command displays the contents of one or more specified libraries. That is, it displays a list of the names and types of all objects contained in each library, regardless of the authorization on each object. | DSPLIB | ASPDEV |

| Function | IBM i 6.1 command | Parameter |
|---|---|---|
| The retrieve library description command is used to retrieve the description of a library. | RTVLIBD | ASPDEV, ASPGRP |
| The save library command allows you to save a copy of one or more libraries. | SAVLIB | ASPDEV |
| The save/restore library command allows you to save and restore a copy of one or more libraries to another system. | SAVRSTLIB | |
| The work with libraries command shows a list of libraries and allows you to copy, delete, display, print, save, restore, change, and clear specified libraries. | WRKLIB | ASPDEV |
| **NWS function** | **NWS command** | |
| The create network server description command creates a description for a network server. | CRTNWSD | STGASPDEV,POOL |
| The create network server storage space command creates a storage space used by a network server. | CRTNWSSTG | ASPDEV |
| **Object functions** | **OBJ commands** | |
| The analyze user objects command collects or reports information for user-created objects on the system. | ANZUSROBJ | ASPDEV |
| The display object description command shows the names and attributes of specified objects in the specified library or in the libraries of the thread's library list. | DSPOBJD | ASPDEV |
| The change object auditing command allows users with *AUDIT special authority to set up auditing on an object. | CHGOBJAUD | ASPDEV |
| The change object owner command transfers object ownership from one user to another. | CHGOBJOWN | ASPDEV |
| The change object primary group command changes the object's primary group from one user to another. | CHGOBJPGP | ASPDEV |
| The create duplicate object command copies a single object or a group of objects. | CRTDUPOBJ | ASPDEV, TOASPDEV |
| The display object description command shows the names and attributes of specified objects in the specified library or in the libraries of the thread's library list. | DSPOBJD | ASPDEV |
| The display object authority command displays the list of authorized users of an object and their assigned authority. | DSPOBJAUT | ASPDEV |
| The edit object authority command displays the list of authorized users of an object and their associated user authorities. | EDTOBJAUT | ASPDEV |
| The grant object authority command grants specific authority for the objects named in the command. | GRTOBJAUT | ASPDEV, REFASPDEV |
| The move object command removes an object from its currently assigned library and places it in a different library. | MOVOBJ | ASPDEV, TOASPDEV |
| The rename object command changes the name of an object in a library. | RNMOBJ | ASPDEV |

| Function | IBM i 6.1 command | Parameter |
|---|---|---|
| The retrieve object description command returns the description of a specific object to a CL program or REXX procedure. | RTVOBJD | ASPDEV |
| The revoke object authority command is used to take away specific (or all) authority for the named objects from one or more users named in the command, or to remove the authority of an authorization list for the named objects. | RVKOBJAUT | ASPDEV |
| The save changed object command saves a copy of each changed object or group of objects located in the same library. | SAVCHGOBJ | ASPDEV |
| The save object command saves a copy of a single object or a group of objects located in the same library. | SAVOBJ | ASPDEV |
| The save/restore object command saves and restores a single object, or a group of objects located in the same library, to another system. | SAVRSTOBJ | ASPDEV, RSTASPDEV |
| The save/restore changed object command saves and concurrently restores a copy of each changed object, or group of objects located in the same library, to another system. | SAVRSTCHG | ASPDEV, RSTASPDEV |
| The restore object command restores to the system a single object or a group of objects in a single library that were saved on diskette, tape, optical volume, or in a save file using a single command. | RSTOBJ | RSTASPDEV |
| The work with object locks command allows you to work with the object lock requests in the system for a specified object. | WRKOBJLCK | ASPDEV |
| The work with objects by owner command is used to manage objects for any user profile. | WRKOBJOWN | *see SETASPGRP* |
| **RCY functions** | **RCY commands** | |
| The change recovery for access paths command is used to change the target access path recovery time for the system or for one or more ASPs. | CHGRCYAP | |
| The Display Recovery for Access Paths display shows a list of access path recovery times for the system and for the ASPs that are currently on the system. | DSPRCYAP | |
| The Edit Recovery for Access Paths display shows a list of access path recovery times for the system and for ASPs that are currently active on the system. | EDTRCYAP | |
| **Miscellaneous functions** | **Miscellaneous commands** | |
| The create save file command creates a save file. | CRTSAVF | Use SETASPGRP |
| The install Windows Server command installs the Windows server base operating system on an Integrated xSeries® Server. | INSWNTSVR | STGASPDEV |
| The save security data command saves all security information without requiring a system in a restricted state. | SAVSECDTA | ASPDEV |

| Function | IBM i 6.1 command | Parameter |
|---|---|---|
| The save command saves a copy of one or more objects that can be used in the integrated file system. | SAV | ASPDEV |
| The work with device descriptions command is used to display and to work with device description functions through the Work with Device Descriptions display. | WRKDEVD | |

# Related publications

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

## IBM Redbooks publications

For information about ordering these publications, see “How to get Redbooks publications” on page 237. Note that some of the documents referenced here may be available in softcopy only.

- *Implementing PowerHA for IBM i*, SG24-7405
- *i5/OS V5R4 Virtual Tape: A Guide to Planning and Implementation*, SG24-7164
- *Implementing POWER Linux on IBM System i Platform*, SG24-6388
- *LPAR Simplification Tools Handbook*, SG24-7231
- *i5/OS V5R4 Virtual Tape: A Guide to Planning and Implementation*, SG24-7164

## Online resources

The IBM System i and IBM i Information Center Web site is also relevant as a further information source:

`http://publib.boulder.ibm.com/iseries`

## How to get Redbooks publications

You can search for, view, or download Redbooks publications, Redpapers publications, Technotes, draft publications and Additional materials, as well as order hardcopy Redbooks publications, at this Web site:

`ibm.com/redbooks`

## Help from IBM

IBM Support and downloads

`ibm.com/support`

IBM Global Services

`ibm.com/services`

# Index

## D

## E



## F



## G



## H



## I

## J



## L



## M



## N



## O

## P

## Q



## R



## S

## T



## U

## V



## W



## X

# IBM i 6.1 Independent ASPs: A Guide to Quick Implementation of Independent ASPs

**Take advantage of the IBM i 6.1 support for independent ASPs**

**Improve your uptime by using IASPs**

**Easily install and configure IASPs for your environment**

This IBM Redbooks publication explains how to configure and manage independent disk pool (IASP) functionality of IBM i 6.1. It is designed to help IBM technical professionals, business partners, and customers understand and implement independent disk pools in the IBM i 6.1.

In addition, this book provides the background information that is necessary to plan, implement, and customize this functionality to your particular environment. It provides guidance on running user applications with either application data or most application objects residing in an independent disk pool. Considering that you can also use independent disk pools in a cluster environment, this book shows you the basic steps to make your independent disk pool switchable between two Power Systems servers or a single server with multiple LPARs.

Independent auxiliary storage pools have many business and technical advantages for Power Systems using IBM i. Not only are independent auxiliary storage pools easy to create and maintain, most applications can use them with simple work management changes. IASPs can provide immediate benefits to your enterprise.

**INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION**

**BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE**

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

**For more information:**
**ibm.com**/redbooks

SG24-7811-00 ISBN 0738433683