



IBM Systems - iSeries
Independent disk pools

Version 5 Release 4





IBM Systems - iSeries
Independent disk pools

Version 5 Release 4

Note

Before using this information and the product it supports read the information in "Notices," on page 29.

Fifth Edition (September 2007)

| This edition applies to version 5, release 4, modification 0 of IBM i5/OS (product number 5722-SS1) and to all
| subsequent releases and modifications until otherwise indicated in new editions. This version does not run on all
| reduced instruction set computer (RISC) models nor does it run on CISC models.

© Copyright International Business Machines Corporation 2002, 2007. All rights reserved.

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

Contents

Independent disk pool examples 1

iSeries Navigator requirements for disk management	1
Accessing disk units in iSeries Navigator	1
Setting up communication	1
Evaluating the current configuration	2
Printing your disk configuration	3
Calculating disk space requirements	4
Planning for independent disk pools	5
Switchable and stand-alone independent disk pools	5
Hardware requirements for independent disk pools	6
Physical planning requirements	7
Software and licensing requirements	9
Communications requirements	9
Cluster requirements	10
Application considerations for independent disk pools	10
Switching independent disk pool between different releases	11
Storing and printing spooled files	12
Object identification	13
Multiple system libraries	13

Supported and unsupported object types	14
Configuring independent disk pools	16
Converting UDFS disk pools	16
Creating a switchable independent disk pool	17
Creating a dedicated independent disk pool	23
Creating a new disk pool group	23
Managing independent disk pools	24
Making a disk pool available	24
Making a disk pool unavailable	25
Backing up and recovering independent disk pools	25
Releasing job reservations on an independent disk pool	26
Switching independent disk pool between different releases	27
Examples: Independent disk pool configurations	27
Examples: Dedicated independent disk pools	27

Appendix. Notices 29

Programming Interface Information	31
Trademarks	31
Terms and conditions	31

Independent disk pool examples

This topic collection provides examples of independent disk pools, from a conceptual explanation to planning, configuring, and managing independent disk pools on your systems.

iSeries Navigator requirements for disk management

There are requirements for disk configuration with iSeries™ Navigator.

Changing the disk configuration of your system is a time-consuming process, so you want to plan carefully to be as efficient as possible. Before you begin disk management using iSeries Navigator, do these procedures to ensure that you are ready.

Accessing disk units in iSeries Navigator

Before you can perform any disk management tasks with iSeries Navigator, you need to install the Configuration and Service component and enable the Disk Units folder.

Follow these procedures to access the Disk Units folder.

1. Install the Configuration and Service component
 - a. From the **File** menu of iSeries Navigator, select **Install Options**, and then click **Selective Setup**.
 - b. Follow the instructions on the resulting dialog box to install the Configuration and Service component.
2. Enable the Disk Units folder
 - a. In iSeries Navigator, right-click the system connection and select **Application Administration**.
 - b. On the resulting window, click **OK**.
 - c. Click the **Host Applications** tab.
 - d. Expand your operating system.
 - e. Select **Disk Units** to have **Default Access** or **All Object Access**.
 - f. Click **OK**.
 - g. Restart iSeries Navigator.
3. Access the Disk Units folder to perform all disk management functions
 - a. In iSeries Navigator, expand **My Connections** (or your active environment).
 - b. Expand the System i™ you want to examine, **Configuration and Service** → **Hardware** → **Disk Units**.

Setting up communication

This topic describes how to configure the service tools system and set up your service IP address.

iSeries Navigator allows you to access the system from your PC through the service tools system to perform disk management functions at two different levels. You can access the system when it is fully started, or you can access the system when it is in dedicated service tools (DST) mode. DST provides some additional functions for disk management that are not available when the system is fully restarted. Before you attempt to use any disk management functions, you must configure the service tools system. If you want to access DST functions, you must also set a service IP address.

1. Configure the service tools system

To access disk management functions in iSeries Navigator, you must first configure the service tools system with DST access and user IDs. Be familiar with Service tool concepts before you start. See [Configure the service tools system and Configure service tools user IDs](#) for instructions.

2. Set the service IP address

To access DST functions on your system from iSeries Navigator, you need to specify a service IP address for the system. The service IP address specifies the TCP/IP address of the system when it is at DST. This address takes the form *xxx.xxx.xxx.xxx* where *xxx* is an integer from 0 to 255. The address can also be a domain name that resolves to an address as previously described. Contact your network administrator for this information. Make sure you have configured the service tools system before you continue with these instructions.

To set the service IP address for your system, follow these steps.

- a. In iSeries Navigator, expand **My Connections** (or your active environment).
- b. Right-click the system for which you want to specify a service IP address, and select **Properties**.
- c. Click the **Service** tab.
- d. If your system is fully restarted, click **Lookup**. The system attempts to locate the correct service IP address. If your system is at DST, specify the service IP address, and click **OK**.

Once the service IP address is set, you can connect to the system when it is in DST mode by using iSeries Navigator. Start iSeries Navigator to connect to the system. iSeries Navigator opens with a subset of functions that you can perform in DST.

Note: If you are unable to configure the service IP address, you can still access DST Disk Management functions. In the Environment tasks window, click Open iSeries Navigator Service Tools window and follow the instructions on the resulting displays.

Evaluating the current configuration

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

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

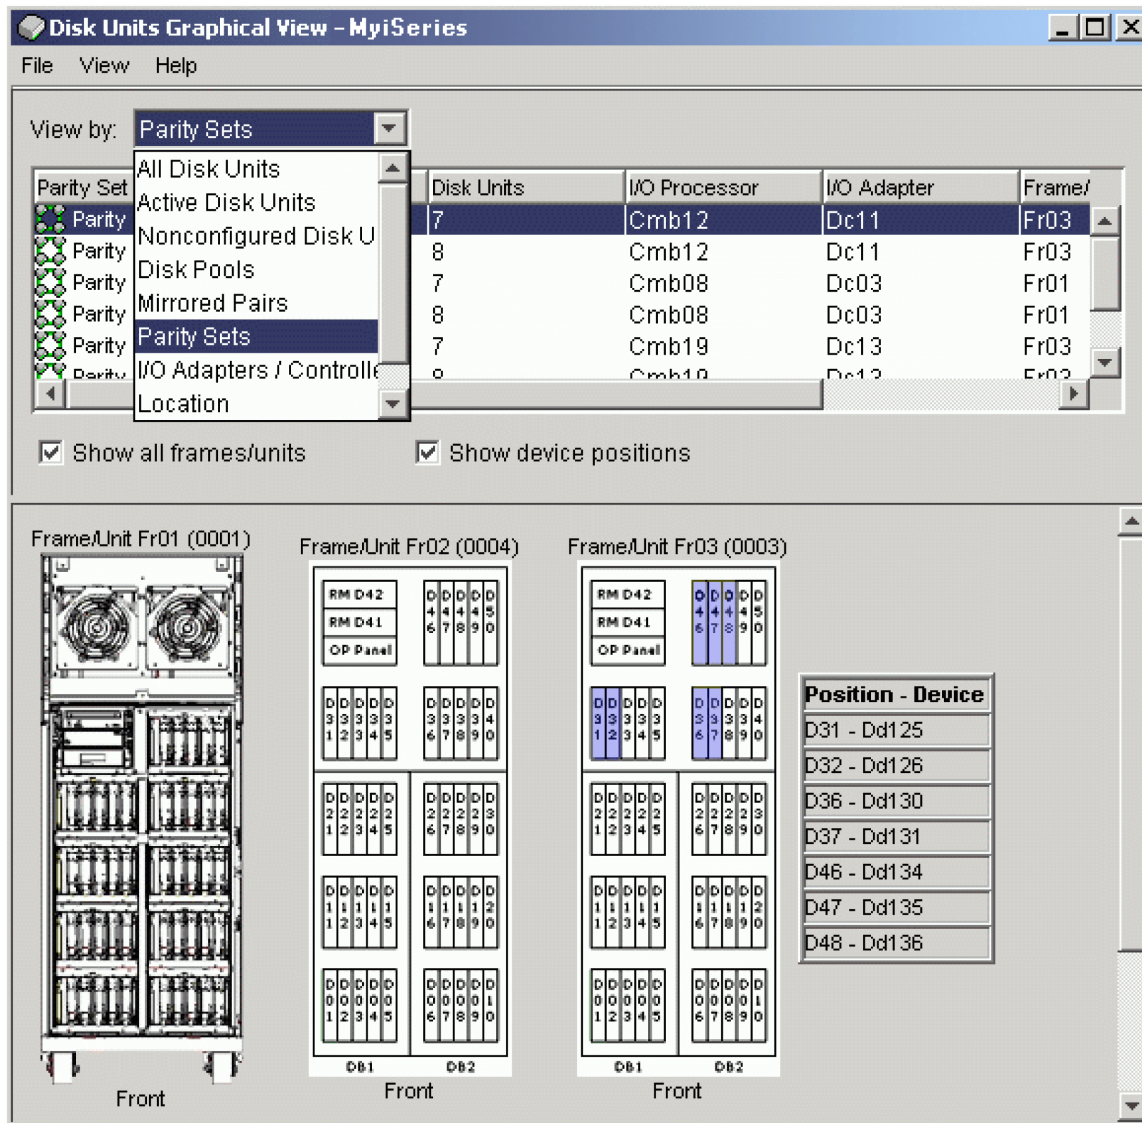
You can choose how to view the hardware in the Disk Unit Graphical View window. For example, you can select to view by disk pools, and then select a disk pool in the list to display only those frames that contain the disk units that make up the selected disk pool. You can select Show all frames to see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

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

To activate the graphical view, follow these steps:

1. In iSeries Navigator, expand **My Connections** (or your active environment).
2. Expand the you want to examine, **Configuration and Service** → **Hardware** → **Disk Units**.
3. Right-click **All Disk Units**, and select **Graphical View**.

Here is an example of the graphical view in iSeries Navigator. The View by menu lists several options for viewing disk units.



Printing your disk configuration

Print your disk configuration.

To print your disk configuration for your records, perform these steps:

1. In iSeries Navigator, expand **My Connections** (or your active environment).
2. Expand the System i you want to examine, **Configuration and Service** → **Hardware** → **Disk Units**.
3. Right-click **All Disk Units**, and select **Graphical View**.
4. Select **Show device positions** to associate disk unit names with the device position where they are inserted.
5. On the **Disk Units Graphical View** dialog, select **File** → **Print**.

Find directions to print your disk configuration from the disk units Graphical View in iSeries Navigator.

Calculating disk space requirements

Before you change the disk configuration or the disk protection on your system, you need to calculate the space requirements for the change. This helps you ensure that your system has sufficient disk storage for the changes.

You can use the Disk space calculator to determine if your disk pool contains sufficient storage space to perform changes. To use the calculator, you need to know how much free space and used space exists in the disk pool.

To view your disk pool configuration

1. In iSeries Navigator, expand **My Connections** (or your active environment).
2. Expand the System i you want to examine, **Configuration and Service** → **Hardware** → **Disk Units** → **Disk Pools**.
3. Right-click the source disk pool you want to view and select **Properties**.
4. Select the **Capacity** tab. The **Capacity** tab displays the used space, free space, total capacity, threshold, and percentage of disk space used for the disk pool.
5. Record the used space, free space, and the threshold from the **Capacity** tab.
6. Enter the used-space value and free-space value into the disk space calculator.
7. If you want to use the threshold value, enter the threshold value in the calculator.

The calculator will warn you if your disk use exceeds your threshold.

The calculator uses JavaScript™ to function. Ensure that you are using a browser that supports JavaScript, and that JavaScript is enabled.

Scenario: Calculating disk space when moving a disk unit

Read the scenario to learn how to remove a disk unit from a disk pool.

Before the disk unit is removed from the source disk pool, the data on the disk unit is copied to the other disk units in the source disk pool. You need to ensure that you have sufficient free space on your source disk pool for this data.

Assume that you have 180 GB of used space, 40 GB of free space, the threshold is set to 90%, and the disk unit you are removing from the disk pool has a capacity of 18 GB.

Perform the scenario as follows:

1. Using the disk space calculator on the **Capacity** tab of the **Disk Pools Properties** dialog, enter these values and click **Calculate**.

A graphical representation of the used and free space on your system appears along with the total disk space, the percent used, and your threshold.

2. From the disk space calculator, select **Remove Disk Space from Disk Pool** and enter 18 for the amount. Click **Calculate**.

The graphical representation is redrawn based on the revised values for used and free space after the 18 GB to be removed is removed from the system.

The percentage of disk space used is now 89.1%. This number just fits under your threshold, but not by much.

Planning for independent disk pools

Several requirements must be satisfied to use independent disk pools, particularly if you plan to use switchable independent disk pools. Setting up an environment for switching devices begins with careful planning.

Attention: When you are ready to order a new system or a system upgrade to use clusters, IBM® will assist you in making sure that your cluster requirements are met. See Planning for Clustering.

Creating a stand-alone, or dedicated, independent disk pool does not require as much planning as a switchable independent disk pool. However, you should still take the time to make sure that your future needs will not require you to be able to switch the independent disk pool.

When independent disk pools are used, you should configure a storage pool for the independent disk pools separate from the base storage pool (pool number 2) and separate from storage pools configured for jobs that are not using independent disk pools.

Switchable and stand-alone independent disk pools

There are two basic environments in which you can take advantage of independent disk pools, a multisystem environment managed by a cluster and a single-system environment with a single system.

Independent disk pools in a multisystem clustered environment

| A group of systems in a cluster can take advantage of the switchover capability to move access to the
| independent disk pool from system to system. In this environment, an independent disk pool can be
| switchable when it resides on a switchable device. A switchable device can be an external expansion unit
| (tower), an IOP on the bus shared by logical partitions, or an IOP or without an IOP hardware that is
| assigned to an I/O pool.

Note: Hardware that does not have a physical IOP has a virtual logical representation of the IOP.

A switchable device that contains an independent disk pool can be switched automatically in the case of an unplanned outage, or it can be switched manually by administering a switchover.

Another option that can be leveraged in a multisystem environment is geographic mirroring. Geographic mirroring allows you to maintain two identical copies of an independent disk pool at two sites that are geographically separated. The independent disk pools at the separate sites can be switchable or dedicated.

Dedicated independent disk pools in a single-system environment

An independent disk pool in a single-system environment, with no clustering and no switchable devices, is said to be a dedicated, private, or stand-alone independent disk pool. While you cannot switch the access to the independent disk pool amongst systems in this environment, you can still isolate data in an independent disk pool, keeping it separate from the rest of the disk storage on the system. The independent disk pool can then be made available (brought online) and made unavailable (taken offline) as needed. This might be done, for example, to isolate data associated with a specific application program or to isolate low-use data that is only needed periodically. Dedicated independent disk pools might also be used to consolidate data from several small systems at branch offices to one or more larger systems at a central location, while still keeping the data separate for each branch.

Independent disk pools allow you to isolate certain maintenance functions. Then, when you need to perform disk management functions that normally require the entire system to be at DST, you can perform them by merely varying off the affected independent disk pool.

The following table compares dedicated independent disk pools and independent disk pools in a multisystem environment.

Consideration	Dedicated	Multisystem environment	
	Single system	Multisystem cluster	Logical partitions in a cluster
i5/OS® cluster required	No	Yes	Yes
Connectivity between systems	Not applicable	HSL loop	Virtual OptiConnect
Location of disk units	Any supported internal or external disk units	External expansion unit (tower)	IOP on shared bus
Switchability	No	Yes, between systems	Yes, between partitions
Switchable entity	None	Expansion unit	IOP

In a hardware switching environment, one node in the device domain owns it, and all the other nodes in the device domain show that the independent disk pool exists. In a geographic mirroring environment, one node at each site owns a copy of the independent disk pool. When an independent disk pool is created or deleted, the node creating or deleting the independent disk pool informs all the other nodes in the device domain of the change. If clustering is not active between the nodes, or a node is in the midst of a long running disk pool configuration change, that node will not update and will be inconsistent with the node rest of the nodes. Nodes must be consistent before a failover or switchover. Ending clustering and starting clustering will ensure that the configuration is consistent.

For more on switchable and dedicated independent disk pools, including example configurations for each of these environments, see Examples: Independent disk pool.

Hardware requirements for independent disk pools

Depending on how you plan to use independent disk pools, you must have specific hardware and operating system releases.

Environments	independent disk pool use	Requirements
Single system	Stand-alone independent disk pool	One system running OS/400® V5R1M0 ¹ or later.
Multiple systems		Use either multiple systems or multiple partitions, by doing one of the following: <ul style="list-style-type: none"> • Two or more systems. • One system running with logical partitions. Note: The version of i5/OS or OS/400 must be compatible.
	Switchable independent disk pools	One or more switchable device, do one of the following: <ul style="list-style-type: none"> • One or more expansion units (frame/units) residing on a high-speed link (HSL) loop. • One or more IOPs on a shared bus or an IOP that is assigned to an I/O pool².
	Geographic mirroring ³	Two or more systems each with sufficient disk space to create the independent disk pools of similar, but not necessarily matching, capability. Note: Consider including hardware for the IP connections. See Communications requirements for more details.

Notes:

1. OS/400 V5R1M0 only provides independent disk pools containing user-defined file system (UDFS). OS/400 V5R2M0 and later supports library-based objects.
2. In an LPAR environment, you can switch the IOP that contains the independent disk pools between system partitions without having an expansion unit. The IOP must be on the bus shared by multiple partitions or assigned to an I/O pool. All IOAs on the IOP will be switched.
3. OS/400 V5R3M0 and later provides support for geographic mirroring.

Physical planning requirements

Depending on how you plan to use independent disk pools, you must satisfy specific physical planning requirements.



Multisystem clustered environment (for switchable independent disk pools)

High-speed link (HSL) cables must be used to attach the expansion units to the systems in the cluster.

The expansion unit must be physically adjacent in the HSL loop to the alternate system or expansion unit owned by the alternative system. You can include a maximum of two systems (cluster nodes) on each HSL loop, though each system can be connected to multiple HSL loops. You can include a maximum of four expansion units on each HSL loop, though a maximum of three expansion units can be included on each loop segment. On an HSL loop containing two systems, two segments exist, separated by the two systems. All expansion units on one loop segment must be contained in the same device cluster resource group (CRG).

- | In order for an expansion unit to become switchable it must physically be the farthest away from the
- | owning system on the loop segment.

Note: An error will occur if you try to make an expansion unit switchable if there is another expansion unit farther away the owning system that has not become switchable.



The switchable expansion unit must be SPCN-cabled to the system unit that will initially serve as the primary node for the switchable hardware group (device CRG). The primary node might be a primary or

secondary logical partition within the system unit. If using logical partitions, the system buses in the intended expansion unit must be owned and dedicated by the partition involved in the cluster.

Software and licensing requirements

Depending on how you plan to use independent disk pools, you must have specific software and licenses.

If you plan to use switchable independent disk pools or independent disk pools that are geographically mirrored, the following elements are required :

- You need OS/400 V5R1M0¹ or later.

Note: For systems on the same HSL loop see, the High Availability Web site to ensure that you have compatible versions of i5/OS or OS/400.

- iSeries Navigator is the graphical user interface for managing and administering your System i from your Windows® desktop. It is required to perform some of the disk management tasks necessary to use independent disk pools. See for steps to enable iSeries Navigator for disk management.
- You need to install Option 41 HA Switchable Resources. Option 41 gives you the capability to switch independent disk pools between systems. To switch an independent disk pool between systems, the systems must be members of a cluster and the independent disk pool must be associated with a switchable hardware group in that cluster. Option 41 also gives you the capability of using the iSeries Navigator cluster management interface to define and manage a cluster that uses switchable resources.

If you plan to use independent disk pools in a single-system environment, the following elements are required :

- You need OS/400 V5R1M0¹ or later.
- iSeries Navigator is the graphical user interface for managing and administering your system from your Windows desktop. It is required to perform some of the disk management tasks necessary to implementing independent disk pools. See iSeries Navigator requirements for disk management for details.

¹ OS/400 V5R1M0 can be used for implementing independent disk pools containing user-defined file systems (UDFS) only. Support for library-based objects is only available starting with OS/400 V5R2M0. Support for geographic mirroring is available at OS/400 V5R3M0.

Communications requirements

There are specific communication requirements for Multisystem clustered environment.

Depending on how you plan to use independent disk pools, you must satisfy the following communications requirements:

Multisystem clustered environment

Switchable independent disk pools and independent disk pools that are geographically mirrored are configured within a cluster. The communication requirements include the following:

For switchable independent disk pools

At least one TCP/IP communications interface between the systems in the cluster. For redundancy, it is recommended that there be at least two separate interfaces between the systems.

For geographic mirroring

The following is recommended:

- Up to a maximum of four unique TCP/IP addresses, used exclusively for geographic mirroring. Geographic mirroring can generate heavy communications traffic. If geographic mirroring shares the same IP connection with another application, for example clustering, then

geographic mirroring might be suspended which results in synchronization. Likewise, clustering response might be unacceptable which results in partitioned nodes.

- Throughput for each data port connection should match, meaning that the speed and connection type should be the same for all connections between system pairs. If throughput is different, performance will be gated by the slowest connection.

Consider configuring a virtual private network for TCP/IP connections for the following advantages:

- Security of data transmission by encrypting the data
- Increased reliability of data transmission by sending greater redundancy

Single-system environment

There are no communications requirements.

Connections from the production system

Geographic mirroring will establish connections from the production system to each of the data port TCP/IP addresses on the mirror copy. TCP can choose to connect from any available TCP/IP address on the production system according to the TCP routing table. The TCP address used is not limited to the addresses configured for geographic mirroring on the production system. TCP may select the same TCP/IP address on the production system to connect to each TCP/IP address on the mirror copy. To control which TCP/IP address(es) on the production system are used to connect to each address on the mirror copy, TCP/IP routes can be created. This is useful to control which addresses on the production system will be chosen for geographic mirroring. It can also eliminate a single point of failure and potential bottleneck caused when all connections are created from the same TCP/IP address.

Cluster requirements

If you plan to use switchable independent disk pools or geographically mirrored independent disk pools, you need to configure a cluster.

The documentation in these independent disk pools topics guides you through the creation and management of your cluster. However, you may want to prepare your network and system environment in advance.

Use the Cluster configuration checklist to ensure that you are prepared to configure clusters in your environment.

Application considerations for independent disk pools

When you are designing or restructuring your application environment for use with independent disk pools, there are several things you should know.

A few of these considerations include the existence of multiple databases, the objects that can and cannot be created in an independent disk pool, how the library list works, and the placement of programs and data in the correct database.

When a primary independent disk pool is made available for the first time, a new database with the same name is also generated by default. See Independent disk pools with distinct databases for more information. If you write an application to access files and libraries in a disk pool group, you must specify how to access that specific database. Some of your options include:

- Use the Set ASP Group (SETASPGRP) command.
- In an SQL environment, use CONNECT to specify the right database. To achieve the fastest performance, make sure that the database to which you perform an SQL CONNECT operation corresponds with your current library namespace. You may need to use the SETASPGRP command

first to achieve this. If the SQL CONNECT function is not operating within the same library namespace, the application uses Distributed Relational Database Architecture^(TM) support, which can affect performance.

- Use the Change Job Description (CHGJOB) command to set the initial ASP group in the job description for a user profile.

As you write applications that create objects, you must know which objects are supported. See Supported and unsupported object types. If your application uses the Create Library (CRTLIB) command, you must specify CRTLIB ASP(*ASPDEV) ASPDEV(*asp-device-name*). If you do not specify these parameters for CRTLIB, the library is created in the system disk pool by default. However, if you use the SQL statement, CREATE COLLECTION, the default for the IN ASP clause is the current library namespace.

When you are operating in an SQL environment, permanent SQL objects cannot span independent disk pool boundaries. For example, you cannot create a view of an independent disk pool object in the system disk pool. This action fails.

A similar concept is true for commitment control with independent disk pools. If you are connected to an independent disk pool relational database, you cannot make committable changes against objects in any other disk pool. When commitment control is active, you have read-only access. You can make committable changes against QTEMP, but you might receive error messages.

It might also be helpful to understand how the library list works when independent disk pools are implemented. When the library list includes QSYS, QSYS2, or SYSIBM, the Multiple system libraries in the independent disk pool (QSYSnnnnn, QSYS2nnnnn, SYSIBnnnnn) are searched before the libraries in the system disk pool. If the object is found in the independent disk pool, the system disk pool will not be searched. In addition, if you switch to a different disk pool group, any libraries that were in the previous library list are removed from the current library list.

You also need to carefully consider where you store data, applications, and application exit programs. It is recommended that data should be stored in independent disk pools. If your independent disk pools are dedicated to the system, it might work to store applications and exit programs in the system database so that they are always accessible, regardless of what disk pool group is associated with a job. If you use the independent disk pool in a clustered environment, you must remember that when the disk pool is switched to another system, the exit program must be available there as well. In this case, it may be more appropriate to store the applications and exit programs in the independent disk pool. Remember that the cluster resource group (CRG) exit program cannot exist in an independent disk pool.

If you are using the independent disk pool in a clustered environment, you must also remember that the user profiles are not stored in the independent disk pool. They are kept in the system disk pool. If an independent disk pool fails over or is switched to another node where the user profile does not currently exist, a user profile might be created on the new node. For a user profile to be created, it must own objects on the switchable disk pool, be the primary group of objects on the switchable disk pool, or be privately authorized to objects on the switchable disk pool. The new user profile has no special authorities and the password is set to *NONE.

Authorization lists may also be created. For an authorization list to be created, it must not currently exist on the target system and it must secure an object on the switchable disk unit. When an authorization list is created, the public authority is set to *EXCLUDE, and no users are given private authority to it.

If you are operating in a clustered environment, see Cluster applications for more information about writing and implementing highly available applications within your cluster.

Switching independent disk pool between different releases

There are restrictions to consider when switching independent disk pools between releases.

Once an independent disk pool is made available on a system, it cannot be made available to another system that is running on a previous version of OS/400. However, it is possible to switch a previous release independent disk pool to a system running the current version of i5/OS and make it available. After it is made available on the system running the current version of i5/OS, its internal contents are changed and it cannot be made available to the previous release system again.

Attention: If a current release disk pool is switched to a V5R1 release system, its disk units show up as non-configured on the previous release system. If these disk units are added to another disk pool, the independent disk pool is destroyed.

Storing and printing spooled files

This topic provides instructions for storing and printing spooled files on a disk pool group.

If you choose to store external resources for spooled files in a disk pool group, you must be aware of the printing implications. You can store the external resources such as, Advanced Function Presentation™ (AFP™) objects such as *FNTRSC, *FORMDF, *OVL, *PAGDFN, *PAGSEG, and non-AFP resources in a disk pool group. In order for the printer writer job to access these objects, you must set the disk pool so that it exists in the library namespace with the writer job.

Follow these steps to set the disk pool in the library namespace of the writer job:

1. Make sure that the disk pool group that contains the external resources is available.
2. Set the disk pool group for the current thread using the SETASPGRP (Set ASP Group) command (disk-pool-group-name).
3. Print the spooled file using the STRPRTWTR (Start Printer Writer) command (printer-device-name).

Consider the following environment:

- Two or more systems in a cluster
- A switchable disk pool to be switch between two or more systems in a cluster
- Spooled files using external resources are placed onto the switchable disk pool
- The external resources are in *SYSBAS

To print a spooled file correctly, from any system in the cluster, the external resources must reside in the same libraries on each of the systems in the cluster.

If a spooled file is not with its' external resource in the disk pool that is in a switchable disk environment, then the external resource must exist in the same library on both systems. If the external resource is not in the same disk pool as the spooled file or the external resources are not stored in both libraries on the systems, then the spooled file cannot be read.

Note: For ease of use, it is recommended that the spooled file and the external resources be stored in the same disk pool.

Only one version of a spooled file is allowed to exist in a namespace. A spooled file cannot be:

- Restored into *SYSBAS if it already exists in a disk pool.
- Restored into a disk pool if it already exists in *SYSBAS.
- Moved from disk pool to *SYSBAS if it already exists in another disk pool.

A disk pool can fail to vary on if the disk pool contains the same version of a spooled file that is in *SYSBAS. To correct the problem, do the following steps:

1. View the job log for the job that attempted to vary on the disk pool for a list of duplicate spooled files in *SYSBAS.
2. Use the Delete Spooled File (DLTSPLF) command to delete the duplicate spooled files in *SYSBAS.

3. Vary off the disk pool.
4. Vary on the disk pool.

Object identification

Follow this information to learn about the complexity of using multiple databases on a single system and what different parameters mean to CL commands.

Because the existence of an independent disk pool on a system means that multiple databases exist on a single system, identifying an object is more complex than it is on a system with only a single system database. When multiple databases exist, it is possible to duplicate the names of libraries and objects in separate databases. The library name and object name do not necessarily uniquely identify an object. There will be times when you also need to know the name of the independent disk pool. The name of the independent disk pool and its database are, by default, the same. However, they do not necessarily have to match. A database name can be up to 18 characters long, while an independent disk pool name can be up to 10 characters long.

While the same library name can exist in two different disk pool groups, libraries cannot have the same name in the system disk pool as in an independent disk pool.

Control language (CL) commands

When using control language (CL) commands that support specification of *ALL or *ALLUSR for the libraries to be searched, the system will typically interpret this to mean "all (user) libraries in your current library namespace" rather than "all (user) libraries on the system." Some commands may interpret *ALL or *ALLUSR differently, so it is important to check the command documentation.

If you used the Display Object Description (DSPOBJD) command, certain system libraries, like QSYS, may appear multiple times if the command is showing information for both the system disk pool and one or more independent disk pools.

Note: Most messages that go to the job log (QSYSOPR) or history log do not contain the name of the independent disk pool. They only contain the object name and library. You must determine what, if any, disk pool group the job that issued the message was using to be able to find the object.

Multiple system libraries

All system libraries continue to exist in the system disk pool with additional system libraries.

To support better isolation and recovery of the independent disk pool group containing system libraries, the following instances of system libraries are also created in the primary disk pool:

QSYSnnnnnn

This contains the database cross reference information for the database represented by the disk pool group.

Normally only internal system code creates objects into this library.

QSYS2nnnnn

This contains the SQL catalogues for the database represented by the disk pool group.

Normally only internal system code creates objects into this library.

QRCYnnnnn

Any recovery object associated with objects within the disk pool group is stored in this library for the primary disk pool for the group.

These objects may be needed for recovery when the disk pool group is varied on. The system disk pool equivalent of this library is QRECOVERY.

QRCLnnnnn

When the reclaim instance is run on the disk pool group, any resulting information normally stored in QRCL is now stored in the QRCL of the primary disk pool for the group.

Normally only functions called during reclaim storage processing create objects into this library instance. Also, when reclaim storage recovers the addressability of lost objects, these objects can be inserted into the QRCLnnnnn library. These are user objects that originally existed in another library.

QRPLnnnnn

Whenever an object contained within the disk pool group is replaced while it is in use, the in-use object is renamed and moved to the QRPLnnnnn library in the primary disk pool for the group.

The new object is inserted into the specified library. The system disk pool equivalent of this library is QRPLOBJ. QRPLnnnnn is cleared at vary on.

Note: nnnnn is the independent disk pool number right-aligned and padded with zeros.

Protected, a library attribute, is introduced to support the extended library capability. Since the libraries QSYSnnnnn, QSYS2nnnnn, and SYSIBnnnnn are special versions that correspond to the system libraries, only operating system code can create objects into them. Applications cannot create objects into these libraries.

Table 1. Library attribute settings

Library	*SYSBAS library	Protected in independent disk pool	Protected in system disk pool
QSYSnnnnn	QSYS	Yes	No
QSYS2nnnnn	QSYS2	Yes	No
SYSIBnnnnn	SYSIBM	Yes	No
QRCLnnnnn	QRCL	No	No
QRCYnnnnn	QRECOVERY	No	No
QRPLnnnnn	QRPLOBJ	No	No
All user libs	Not applicable	No	No

The normal search order for objects is to search the libraries based on the user-specified library value, the user's library list, and the namespace in effect for the job. The only exception to this occurs when the user job has a disk pool group in the job's namespace. In this case, aliasing support takes effect for object references to the database control objects in QSYS, QSYS2, and SYSIBM. The objects in the QSYSnnnnn, QSYS2nnnnn, and SYSIBnnnnn libraries are returned so that users are operating on the database control information associated with their extended namespace.

Supported and unsupported object types

Supported and unsupported object types in independent disk pools.

Objects not supported

The following objects are not supported for use in independent disk pools:

*AUTHLR	*DEV D	*JOBQ	*PRDDFN
*AUTL	*DOC	*JOBSCD	*PRDLOD
*CFGL	*DSTMF ¹	*LIND	*RCT
*CNL	*EDTD	*MODD	*SOCKET
*COSD	*EXITRG	*M36	*SSND
*CRG	*FLR	*M36CFG	*S36

*CSPMAP	*IGCSRT	*NTBD	*USRPRF
*CSPTBL	*IGCTBL	*NWID	
*CTLD	*IMGCLG	*NWSO	
*DDIR	*IPXD	*PRDAVL	

¹ *DSTMF is the object type returned for stream files that are being accessed through the QNTC file system from a remote system. So you shouldn't see *DSTMF ever when accessing the IASP directories from the local system.

Supported object types

The following objects are supported for use in independent disk pools:

*ALRTBL	*FILE	*MSGF	*SCHIDX
*BLKSF	*FNTRSC	*MSGQ	*SPADCT
*BNDDIR	*FNNTBL	*NODGRP	*SPLF
*CHRSF	*FORMDF	*NODL	*SQLPKG
*CHTFMT	*FTR	*OUTQ	*SQLUDT
*CLD	*GSS	*OVL	*SRVPGM
*CLS	*IGCDCT	*PAGDFN	*STMF
*CMD	*JOB	*PAGSEG	*SVRSTG
*CRQD	*JRN	*PDG	*SYMLNK
*CSI	*JRNRCV	*PGM	*TBL
*DIR	*LIB	*PNLGRP	*USRIDX
*DTAARA	*LOCALE	*PSFCFG	*USRQ
*DTADCT	*MEDDFN	*QMFORM	*USRSPC
*DTAQ	*MENU	*QMQR	*VLDL
*FCT	*MGTCOL	*QRYDFN	*WSCST
*FIFO	*MODULE	*SBSD	

Restrictions for supported object types

*ALRTBL

If network attributes reference the alert table, this object needs to exist in the system disk pool.

***CLS** If an active subsystem references the class object, *CLS must exist in the system disk pool.

***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.

*JOB

If an active subsystem references the job description object, *JOB must exist in the system disk pool; for example, autostart job entry, communication entry, remote location name entry, or workstation entry.

***LIB** The library that is specified by CRTSBSD SYSLIBL() must exist in the system disk pool.

*MSGQ

If network attributes reference the message queue, *MSGQ needs to exist in the system disk pool.

***PGM** If an active subsystem references the program object, *PGM must exist in the system disk pool; for example, routing entries and prestart job entries.

*SBSD

You cannot start a subsystem whose description is located in an independent disk pool.

Configuring independent disk pools

After you have satisfied the planning requirements for implementing independent disk pools, you are ready to configure an independent disk pool. You need to use the iSeries Navigator disk management function to configure an independent disk pool.

Converting UDFS disk pools

If you have existing user-defined file system (UDFS) disk pools on your system, you can convert them to primary and secondary disk pools. This allows them to support library-based objects

Support for library-based objects through the use of primary and secondary disk pools was introduced at V5R2.

You must convert UDFS disk pools if you want them to participate in disk pool groups. After you convert a UDFS disk pool to a primary or secondary disk pool, you cannot convert it back to a UDFS disk pool. You must create a primary disk pool before you can associate secondary disk pools.

Converting a UDFS disk pool to primary

You have the capability to convert UDFS disk pools to library-capable primary and 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.

You must create a primary disk pool before you can associate secondary disk pools.

Note: After you convert a UDFS disk pool to a primary or secondary disk pool, you cannot convert it back to a UDFS disk pool.

To convert a UDFS disk pool to a primary disk pool, follow these steps:

1. In iSeries Navigator, expand **My Connections** (or your active environment).
2. Expand the System i you want to examine, **Configuration and Service** → **Hardware** → **Disk Units**.
3. If the Service Tools Signon dialog box is displayed, sign on to service tools.
4. Select **Disk Pools**.
5. Right-click the required UDFS **Disk Pool** and select **Confirm Convert to Primary Disk Pool**.
6. On the **Confirm Convert to Primary Disk Pool** dialog box the default for the **Database Name** field is Generated by the system, which means the system generates a database name for you.
7. Click **Convert Disk Pool**.
8. If you need to associate other existing UDFS disk pools with your new primary disk pool in a disk pool group, see "Converting a UDFS disk pool to secondary."

Converting a UDFS disk pool to secondary

You have the capability to convert UDFS disk pools to library-capable primary and secondary disk pools. Library-capable disk pools will 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 already have created its primary disk pool.

Note: After you convert a UDFS disk pool to a primary or secondary disk pool, you cannot convert it back to a UDFS disk pool.

To convert a UDFS disk pool to a secondary disk pool, follow these steps:

1. In iSeries Navigator, expand **My Connections** (or your active environment).
2. Expand the System i you want to examine, **Configuration and Service** → **Hardware** → **Disk Units** → **Disk Pools**.

3. In the right pane you may select one or more UDFS disk pools to convert at the same time. Right-click the required **UDFS Disk Pool(s)** and select **Convert to Secondary disk pool**.
4. On the **Confirm Convert to Secondary Disk Pool** dialog box, select the primary disk pool that you need to associate with the secondary disk pools. The primary disk pool selected should not have been made available before being associated with the secondary disk pools. Only the primary disk pools that are currently owned by the system can be selected. You will not be able to change the primary after you perform this action.
5. Click **Convert Disk Pool**.
6. If the disk pool you converted to secondary is in a device cluster resource group, then you must change the Online attribute to *PRIMARY. Use the Change Cluster Resource Group Device Entry CHGCRGDEVE command or the (QcstChangeClusterResourceGroupDev) API to change the Online attribute to *PRIMARY.

Creating a switchable independent disk pool

Before you attempt to do switchable independent disk pools, ensure that you have satisfied the hardware, software, communications, and physical planning requirements.

iSeries Navigator is the recommended interface for creating and managing independent disk pools. Wizards in the clusters and disk management components simplify the tasks and guide you through the process. For some disk management tasks, iSeries Navigator is the only option. Make sure you have fulfilled the iSeries Navigator requirements for disk management.

After an expansion unit (frame/unit) is configured as switchable, the disk units in these resources cannot be used in independent disk pools that span different Cluster Resource Groups (CRGs). Even when only one node is defined and no switching will actually be enabled, configuring that expansion unit as switchable is enough to cause this constraint to be enforced.

To create a switchable independent disk pool using iSeries Navigator, do the following steps:

1. Create a cluster. To use switchable independent disk pools, a cluster is required.
2. Make your hardware switchable. If you have a stand-alone expansion unit or an IOP that contains disk units that are to be included in an independent disk pool, you must authorize the expansion unit or IOP to grant access to other nodes.
3. Create a switchable hardware group. A switchable hardware group, also known as a device CRG, defines the switchable independent disk pool. This is what manages the switching of the device. This wizard takes you through the steps to create a new switchable hardware group. It will also guide you through the New Disk Pool wizard which will assist you in creating a new disk pool and adding disk units to it for the cluster.

Note: If you had switchable software products which conform to specific iSeries Navigator cluster guidelines installed when you ran the New Cluster wizard in step 1, the New Cluster wizard may have already prompted you to create a switchable hardware group. If the New Cluster wizard did not detect that a switchable software product was installed, then you have not created the switchable hardware group.

4. Print your disk configurations. Print your disk configuration to have in case a recovery situation occurs. Also, record the relationship between the independent disk pool name and number.

You have now created a switchable independent disk pool. To prepare it for use, do the following steps:

1. Start the switchable hardware group to enable device resiliency for the switchable hardware group.
2. Make a disk pool available. To access the disk units in an independent disk pool you must make the disk pool available (vary on) the disk pool.
3. Perform a test switchover. Before you add data to the disk pool, perform a test switchover on the switchable hardware group you created to ensure the configuration functions as you planned.

You can also use CL commands and APIs for creating a switchable independent disk pool, however, there are some tasks that require that you use iSeries Navigator. To configure your geographic mirroring with the CL commands and APIs, follow these steps:

1. Create the cluster with the nodes that you want using the CRTCLU (Create Cluster) command.
2. Start the nodes in the cluster using the STRCLUNOD (Start Cluster Node) command.
3. Create the device domain. You must create the device domain for all nodes involved in switching an independent disk pool using the ADDDEVDMNE (Add Device Domain Entry) command.
4. Create the device descriptions. Device descriptions must be created on each node that will be in the cluster resource group (CRG). Use the CRTDEVASP (Create Device Description (ASP) command. On the command line in the character-based interface, enter CRTDEVASP. In the **Resource Name** and the **Device Description** fields, enter the name of the independent disk pool you plan to create.
5. Create the cluster resource group. Create the device CRG with the nodes, their roles in the recovery domain, and the independent disk pool device descriptions. You must also specify a site name and data port IP addresses for each node in the recovery domain. Use the CRTCRG (Create Cluster Resource Group) command.
6. Make your hardware switchable. If you have a stand-alone expansion unit or an IOP that contains disk units that are to be included in an independent disk pool, you must authorize the expansion unit or IOP to grant access to other nodes (**iSeries Navigator required**).
7. Create a disk pool on the node that owns the disk units using the New Disk Pool wizard when the system is fully restarted. Make sure clustering is active before you start. Name the independent disk pool to match the device description resource name. As you add disk units, it is best to localize disk units in the same expansion unit or IOP. Also, do not spread the disk pool across parity sets than necessary.
8. Print your disk configuration in case of a recovery situation. See how to display your disk configuration in Backup and Recovery. Also, record the relationship between the independent disk pool name and number.
9. Start the cluster resource group to enable device resiliency using the STRCRG (Start Cluster Resource Group) command.
10. To access the disk units in an independent disk pool, you must vary on the disk pool using the VRYCFG (Vary Configuration) command.
11. Before you add data to the disk pool, use the CHGCRGPRI (Change CRG Primary) command to perform a test switchover. Use the test switchover to ensure the backup node can become the primary node and the primary node can become the backup node. .

| **Specifying a recovery time for the independent disk pool**

- | To improve the vary on performance after an abnormal vary off, consider specifying a private customized
| access path recovery time specifically for that independent disk pool via the Change Recovery for Access
| Paths (CHGRCYAP) command rather than relying upon the overall system-wide access path recovery
| time. This will limit the amount of time spent rebuilding access paths during the vary on.

Creating a cluster

For an independent disk pool to be switchable among systems or for geographic mirroring to be enabled, a cluster is required.

A cluster is a collection or group of one or more systems that work together as a single system. For complete documentation on clusters and how they work, see Clusters.

There are multiple ways to create and manage a cluster. You can use iSeries Navigator to create a cluster, a cluster middleware business partner solution, or IBM cluster commands and APIs. See Solutions for configuring clusters for a complete look at the options for configuring and managing clusters.

To create a cluster for use with switchable independent disk pools:

1. Create a cluster.
2. Verify that all nodes are at the correct potential cluster version. Potential cluster version must be at least 3 for switchable independent disk pools that support libraries. To allow for the V5R3M0 capability of geographic mirroring, the potential cluster version must be set to at least 4.
3. Start all nodes in the cluster, or at least those that will be in the device domains.

Creating a switchable hardware group

Create a switchable hardware group, also known as a device cluster resource group (CRG).

The switchable hardware group contains a list of switchable devices. Each device in the list identifies a switchable independent disk pool. The entire collection of devices are switched to the backup node when an outage, planned or unplanned, occurs. Optionally, the devices can also be made available (varied on) as part of the switchover or failover process.

A switchable hardware group identifies a device domain. A device domain is a subset of cluster nodes that share a set of resilient devices. The device domain is created automatically when you use the iSeries Navigator wizard to create a cluster. If you are using cluster CL commands and APIs, you must add each node that you want to be switchable to the device domain.

Using iSeries Navigator (requires Option 41 (i5/OS - HA Switchable Resources))

The New Switchable Hardware Group wizard will take you through the steps to create a new switchable hardware group and add a disk pool to it for the cluster.

To add a switchable hardware group, follow these steps:

1. In iSeries Navigator, expand **Management Central**.
2. Expand **Clusters**.
3. Expand the cluster for which you need to add a switchable hardware group.
4. Right-click **Switchable Hardware**, and select **New Group**.
5. By default the New Disk Pool wizard creates a protected disk pool that allows you to choose how you want to protect the disk units. You can use device parity protection, mirrored protection, or a combination of both. After the disk pool is created, you are prompted to start disk unit mirroring. This ensures that if you make changes to your disk pool configuration, it will remain protected. You can also create an unprotected disk pool by unchecking the protection option.

Note: Make sure that all the nodes in the recovery domain are started.

Using Cluster CL commands and APIs

You can also use the following to add a device domain entry and create a device cluster resource group:

Add Device Domain Entry

Adds a node to a device domain membership list so that it can participate in recovery actions for resilient devices. The addition of the first node to a device domain has the effect of creating that device domain.

- ADDDEVDMNE (Add Device Domain Entry) Command
- Add Device Domain Entry (QcstAddDeviceDomainEntry) API

Create Cluster Resource Group

Creates a cluster resource group object. The cluster resource group object identifies a recovery domain, which is a set of nodes in the cluster that will play a role in recovery.

- CRTCRG (Create Cluster Resource Group) Command
- Create Cluster Resource Group (QcstCreateClusterResourceGroup) API

Starting a switchable hardware group

To enable device resiliency for the switchable hardware group, you must start the switchable hardware group. Use the steps to start a switchable hardware group.

1. In iSeries Navigator, expand **Management Central**.
2. Expand **Clusters**.
3. Expand the cluster that contains the switchable hardware you need to start.
4. Click **Switchable Devices**.
5. Right-click the switchable devices group you need to start, and select **Start**.

You can also use the Start Cluster Resource Group (STRCRG) command in the character-based interface to start the switchable hardware group.

Creating an independent disk pool

You can create a new independent disk pool and add disk units to it. You can also add disk units to an existing disk pool.

By default the New Disk Pool wizard creates a protected disk pool that allows you to choose how you want to protect the disk units. You can use device parity protection, mirrored protection, or a combination of both. After the disk pool is created, you will be prompted to start mirroring. This ensures that if you make changes to your disk pool configuration, it will remain protected. You can also create an unprotected disk pool by unchecking the protection option.

Note: If you are creating a switchable disk pool, make sure you have completed the previous steps in **Create a switchable independent disk pool**.

To create a new disk pool and add disk units to it, follow these steps:

1. In iSeries Navigator, expand **My Connections** (or your active environment).
2. Expand the System i you want to examine, **Configuration and Service** → **Hardware** → **Disk Units**.
3. Right-click **Disk Pools** and select **New Disk Pool**.
4. Follow the wizard's instructions to add disk units to a new disk pool.

Note: If you are creating a switchable independent disk pool, use the same name for the disk pool that you used when creating the device descriptions.

Making your hardware switchable

In a multisystem environment, you must make an external expansion unit switchable.

When you are using independent disk pools in a switchable environment, the associated hardware must be authorized to switch as well. Depending on your environment, this can include frame/units or IOPs and their associated resources. Refer to the following steps that apply to your switchable environment.

An independent disk pool can contain disk units within several expansion units. If you have a stand-alone expansion unit that contains disk units included in an independent disk pool, you must authorize the expansion unit to grant access to other systems. This is called making an expansion unit switchable. If you do not want other systems to be able to access the stand-alone expansion unit, you must make the expansion unit private.

To make a frame/unit switchable, follow these steps:

1. In iSeries Navigator, expand **My Connections** (or your active environment).
2. Expand the System i you want to examine, **Configuration and Service** → **Hardware** → **Disk Units** → **By Location** and select the frame or disk unit that you want to make switchable.
3. Right-click a highlighted frame or disk unit and select **Make Switchable**.
4. Follow the instructions on the dialog box that is displayed.

To allow an IOP to be switched, the bus containing the IOP that controls the disk units to be switched must be *owned shared* by the primary node. The bus must also be *use bus shared* by the backup node. See Dynamically switching IOPs between partitions for more information.

To complete this task, you need a Service Tools user profile with administration authority to the System Partitions function in dedicated service tools (DST). For more information about obtaining logical partition privileges, refer to Logical partition authority.

To change the ownership type for a bus using Management Central, follow these steps:

1. In iSeries Navigator, expand **My Connections** (or your active environment).
2. Select the primary partition of the system.
3. Expand **Configuration and Service** and select **Logical Partitions**.
4. Right-click the **Logical Partition** and select **Configure Partitions**.
5. In the Configure Logical Partitions window, right-click the bus for which you want to change ownership and select **Properties**.
6. Select the **Partitions** page.
7. Select the partition that owns the bus in **Owning logical partition**, and then select the ownership type in **Sharing**. If the ownership type is shared, the partitions that share the bus appear in the list. Click **Help** if you need more information about these options.
8. Click **OK**.

If you are using the Hardware Management Console to manage your logical partitions, you must create an I/O pool that includes the IOP, IOA, and all attached resources to allow an independent disk pool to be switchable between partitions. You must grant access to each partition that you want to own the independent disk pool by assigning the I/O pool in each partition profile.

To create an I/O pool that can be switched between partitions, follow these steps:

1. Open the Logical Partition Profile Properties window to change partition profile properties and assign resources to an I/O pool.
2. Click the **Physical I/O** tab.
3. In the **Profile I/O devices** column, expand the bus that contains the IOP that you want to make switchable.
4. Select the IOP that you want to assign to an I/O pool. The IOP must be *desired* (no check mark in the **Required** column).
5. Click the I/O Pool column so that the cursor appears in the row of the IOP you want to assign to an I/O Pool, and type the number for the I/O pool.
6. Repeat these steps to add each IOA and resource under the control of the IOP to the I/O pool.
7. Click **OK**.

After you have added the resources to the I/O pool, complete the following steps to associate the I/O pool with each additional partition that you want to be able to own the independent disk pool in the switchable environment.

1. Open the Logical Partition Profile Properties window to change partition profile properties for each additional partition that needs to access the independent disk pool.
2. Click the **Physical I/O** tab.
3. Click **Advanced**.
4. On the I/O Pools window, in the **I/O pools to add** field, type the number of the I/O pool to which you assigned the resources that you want to switch with the independent disk pool.
5. Click **Add** → **OK**.

For the I/O pool changes to take effect, complete the following steps for each partition whose partition profile was changed:

1. Shut down the partition. See *Restarting and shutting down i5/OS in a logical partition*.
2. Start the logical partition by activating the partition profile to reflect the changes.

Synchronizing user profile name, UID, and GID

Synchronize user profiles across your cluster to reduce the amount of processing required when you make a disk pool available.

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

Management Central provides a method for administrators to edit user profiles across multiple systems. See *Manage users and groups with Management Central* for details.

Switching access to backup system

Perform a cluster switchover when you want a backup system to access the switchable device containing an independent disk pool.

In a multisystem clustered environment that uses switchable independent disk pools, an independent disk pool can only be accessed by one node at a time. Up to date access to a switchable independent disk pool is managed through the switchover function within the cluster.

To switch access from the current node in the cluster to the first backup node:

1. The switchover may have already made the current node unavailable. If not, then from the current node, make a disk pool unavailable (vary off).
2. Switch the independent disk pool to the first backup cluster node by performing a switchover in the cluster.

Changing the system takeover IP address

Change the IP address for a system associated with a relational database in a clustered, switchable environment.

The system takeover IP address is associated with a primary disk pool in a clustered, switchable environment. Specifically, it is the IP address for a system associated with the relational database name in the device description for a switchable independent disk pool. The specified address must exist on all nodes in the recovery domain if the cluster resource group is active.

To change the system takeover IP address for a primary disk pool, follow these steps:

1. In iSeries Navigator, expand **Management Central** → **Clusters**.
2. Expand the cluster that contains the switchable hardware group.
3. Expand **Switchable Hardware**.
4. Click the switchable hardware group and then right-click the required primary disk pool and select **Properties**.

Note: The system takeover IP address can only be associated with a primary switchable independent disk pool.

5. Change the system takeover IP address in the **IP address** field.

You can also use the CHGCRGDEVE (Change Cluster Resource Group Device Entry) command in the character-based interface to change the system takeover IP address.

Creating a dedicated independent disk pool

Creating a dedicated (or stand-alone) independent disk pool does not require as much planning and configuration as a switchable independent disk pool requires. However, you should still take the time to make sure that your future needs will not require you to be able to switch the independent disk pool.

To create a dedicated independent disk pool, you can use the New Disk Pool wizard in iSeries Navigator. This will assist you in creating a new disk pool and adding disk units to it. The New Disk Pool wizard also allows you to include unconfigured disk units in a parity set, and start device parity protection and disk compression. As you add disk units, do not spread disk units that are in same parity set across multiple disk pools, because failure to one parity set would affect multiple disk pools.

To use the New Disk Pool wizard to create a dedicated independent disk pool, follow these steps:

1. In iSeries Navigator, expand **My Connections** (or your active environment).
2. Expand the System i you want to examine, **Configuration and Service** → **Hardware** → **Disk Units**.
3. Right-click **Disk Pools** and select **New Disk Pool**.
4. Follow the wizard's instructions to add disk units to a new disk pool.
5. Print your disk configuration to have available in case of a recovery situation.
6. Record the relationship between the independent disk pool name and number.

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

Creating a new disk pool group

You can create a disk pool group and add disk units to the individual disk pools by using the Disk Pool wizard.

A disk pool group is made up of a primary disk pool and zero or more secondary disk pools. A practical use of a disk pool group is to isolate journal receivers, which might reside in one or more secondary disk pools, from the objects for which they contain journal entries, which reside in the primary disk pool.

If you have existing UDFS disk pools that you would like to include in a disk pool group, see [Converting a UDFS disk pool to primary or Convert a UDFS disk pool to secondary](#).

Note: If you want to create a switchable independent disk pool (UDFS, primary, or secondary), you must create the cluster first. For more information, see [Creating a switchable independent disk pool](#).

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

1. In iSeries Navigator, expand **My Connections** (or your active environment).
2. Expand the System i you want to examine, **Configuration and Service** → **Hardware** → **Disk Units**.
3. Right-click **Disk Pools** and select **New Disk Pool**.
4. On the resulting New Disk Pool dialog box, select **Primary** for the Type of Disk Pool field and complete the required information.

Note: If you have already created a primary disk pool with which you want to associate one or more secondary disk pools in a disk pool group, you can skip this step. After you have created the primary disk pool, click New Disk Pool if you want to create a secondary disk pool to associate with the primary disk pool. From the resulting dialog box, select Secondary for the Type of Disk Pool field and complete the required information. Repeat this step for each secondary disk pool you want to create. Follow the wizard's instructions to add disk units to the new disk pools.

Managing independent disk pools

After you configure an independent disk pool, you can perform management tasks using iSeries Navigator.

Make sure you can Access disk management functions.

Making a disk pool available

To access the disk units in an independent disk pool, you must make the disk pool available (vary it on).

To access the disk units in an independent disk pool and the objects in the corresponding database, you must make the disk pool available (vary it on). If you are using geographic mirroring, you must make the production copy of the disk pool available. You can only make the mirror copy available if it is detached. For a geographically mirrored disk pool, you must also make sure that the switchable hardware group is started before attempting to make the disk pool available unless geographic mirroring is suspended.

In a multisystem clustered environment, you can make the disk pool available to the current node or to another node in the cluster. The independent disk pool can only be varied on for one node at a time. When you want to access the independent disk pool from a different node, you must switch the independent disk pool to the backup cluster node. See Performing a switchover for details on switching a device CRG (referred to as a switchable hardware group in iSeries Navigator) to the backup node.

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.

When you make a disk pool available or perform disk configuration changes on an independent disk pool, processing can seem to stop. If you are doing other device description activities, then make available and disk configuration changes will wait.

Failures early in make available processing of a geographically mirrored disk pool might cause a full synchronization on the next make available or resume.

To make an independent disk pool available:

1. In iSeries Navigator, expand **My Connections** (or your active environment).
2. Expand the you want to examine, **Configuration and Service** → **Hardware** → **Disk Units**.
3. Expand **Disk Pools**.
4. Right-click the unavailable disk pool and select **Make Available**. You can select multiple disk pools to make available at the same time.
5. From the dialog box displayed, click **Make Available** to make the disk pool available.

You can use the Vary Configuration (VRYCFG) command in the character-based interface to make the disk pool available.

- | Use the Display ASP Status (DSPASPSTS) command to identify where a step is in the process.

Making a disk pool unavailable

You can select an independent disk pool to make it unavailable (vary it off). You cannot access any of the disk units or objects in the independent disk pool or its corresponding database until it is made available (varied on) again. The pool can be made available again on the same system or another system in the recovery domain of the cluster resource group.

Important: Before an independent disk pool can be made unavailable, no jobs can hold reservations on the disk pool. See Release job reservations on an independent disk for details on determining whether jobs are using an independent disk pool and how to release the job reservations.

When making a UDFS disk pool unavailable using iSeries Navigator, messages might be generated that require a response in the character-based interface. iSeries Navigator will not provide any indication that a message is waiting.

To make an independent disk pool unavailable:

1. In iSeries Navigator, expand **My Connections** (or your active environment).
2. Expand the System i you want to examine, **Configuration and Service** → **Hardware** → **Disk Units**.
3. Expand **Disk Pools**.
4. Right-click the disk pool you want to make unavailable and select **Make Unavailable**.
5. From the dialog box that displays, click **Make Unavailable** to make the disk pool unavailable.

You can use the Vary Configuration (VRYCFG) command in the character-based interface to make the disk pool unavailable.

- | Use the Display ASP Status (DSPASPSTS) command to identify where a step is in the process.
- | Use the Control ASP Access (QYASPCTLAA) API to restrict the processes that have access to the ASP.
- | Use the Start DASD Management Operation (QYASSDMO) API to reduce the amount of time it takes to make a disk pool unavailable.

Backing up and recovering independent disk pools

Be sure to consider a save strategy for your independent disk pools.

A good save strategy is just as important for independent disk pools as it is with the rest of your system information. If you use independent disk pools, it is recommended that you use Backup, Recovery and Media Services (BRMS) to save your independent disk pool data. If you need to perform a recovery, BRMS simplifies the process. However, BRMS is not required. In the case of disk failures or a complete system loss, you may need to follow recovery procedures to restore the data you have saved. See the Backup and Recovery manual for steps to restore information to the independent disk pools.

If you experience problems accessing an independent disk pool or making it available, the problem might be with the disk pool. The configuration source might be corrupted or the primary and secondary disk pools might need to be reassociated.

Recovering an independent disk pool

Read the problems that can occur with a disk pool and how to recover it.

If you are experiencing problems accessing an independent disk pool or making it available, there may be a problem with the disk pool. Possible problems include:

- The configuration source is corrupted. When corruption occurs, the independent disk pool will appear to have no disk units in it. The disk pool may also appear to have no disk units in it if it is switched to another node in a clustered environment. Before you attempt a recovery, make sure that no other

system owns the disk pool. If you know the serial numbers of the disk units in the independent disk pool that might need recovery, make sure you are on the system that owns those disk units and that they show as non-configured.

If the configuration source is corrupted, you can select to recover the configuration information about the configuration source. Recovering the configuration attempts to determine the original configuration and recover it. During this process, the independent disk pool might need to be cleared, destroying all data on the disk units in the pool. If the disk pool needs to be cleared, a message is displayed telling you of this and allowing you to cancel the recovery.

- The mirrored disk unit of the configuration source is damaged. When this happens, the mirrored configuration source becomes unknown. The disk pool will be unavailable, and you must recover the configuration information of an unknown configuration source before making it available. You should only attempt to recover the state of the unknown configuration source when you know 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. In iSeries Navigator, expand **My Connections** (or your active environment).
2. Expand the System i you want to examine, **Configuration and Service** → **Hardware** → **Disk Units** → **Disk Pools**.
3. Right-click the problematic disk pool. If iSeries Navigator detects one of the problems listed above, then **Recover Configuration** or **Recover Unknown Configuration Source** appears in the list. If you see either of these options, select it to continue.
4. Follow the instructions on the dialog box displayed.

Recovering 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 needs to be reassociated with a primary disk pool. You can recover the disk pool group through iSeries Navigator.

To recover a disk pool group, follow these steps:

1. In iSeries Navigator, expand **My Connections** (or your active environment).
2. Expand the System i you want to examine, **Configuration and Service** → **Hardware** → **Disk Units** → **Disk Pools**.
3. Verify that the required primary disk pool exists. If it does not, you must Create a disk pool.
4. In the right pane you can select one or more secondary disk pools that need to be associated with the primary disk pool. Right-click all the secondary **Disk Pools** that need to be associated with a primary disk pool and select **Recover Group**.
5. On the **Confirm Recover Disk Pool Group** dialog box, select the primary disk pool that you want associated with the secondary disk pools. Only the primary disk pools that are currently owned by the system are available to select. You cannot change the primary after you perform this action.
6. Click **Recover Group**.

Releasing job reservations on an independent disk pool

If jobs are currently using an independent disk pool, you cannot make the disk pool unavailable (vary off). In order to make the independent disk pool unavailable, all the jobs using the disk pool need to release their reservation on the disk pool. To determine how to handle this situation, you must first view the jobs.

To view the jobs using an independent disk pool:

1. In iSeries Navigator, expand **My Connections** (or your active environment).
2. Expand the System i you want to examine, **Configuration and Service** → **Hardware** → **Disk Units** → **Disk Pools**.
3. Right-click the disk pool and select **Jobs**.

After you have identified the jobs using an independent disk pool that you want to make unavailable, you have several options for each job:

1. End the job. Be sure to consider the affect this action may have before ending a job. This may not be a good idea in some cases. See *Ending a job* for details.
2. Let the job run to completion. If the job appears to be processing normally, you can wait until the job runs to completion.
3. Diagnose and recover a poorly performing or hung job. If the job does not appear to be running, determine if the job is hung or is performing slowly. See the *Work management troubleshooting* topic for tips on diagnosing and handling problematic jobs.
4. Release held jobs or threads.

| Use the `Work with ASP Jobs (WRKASPJOB)` command when an independent disk pool is unavailable to
| identify any jobs that hold locks on objects in the disk pool.

Switching independent disk pool between different releases

There are restrictions to consider when switching independent disk pools between releases.

Once an independent disk pool is made available on a system, it cannot be made available to another system that is running on a previous version of OS/400. However, it is possible to switch a previous release independent disk pool to a system running the current version of i5/OS and make it available. After it is made available on the system running the current version of i5/OS, its internal contents are changed and it cannot be made available to the previous release system again.

Attention: If a current release disk pool is switched to a V5R1 release system, its disk units show up as non-configured on the previous release system. If these disk units are added to another disk pool, the independent disk pool is destroyed.

Examples: Independent disk pool configurations


Independent disk pools can be switchable among a group of systems in a cluster, providing the benefits of continuous availability of the disk units they contain. Or they can be stand-alone (or dedicated) on a single system, independent of the rest of the storage on the system.

Examples: Dedicated independent disk pools

In a single-system environment, a dedicated (or stand-alone), independent disk pool can be varied off independent of other disk pools because the data in the independent disk pool is self-contained.

That is, all of the necessary system information associated with the independent disk pool's data is contained within the independent disk pool. The independent disk pool can also be varied on while the system is active; that is, no initial program load (IPL) is required. Using independent disk pools this way can be useful, for example, if you have large amounts of data that are not needed for day-to-day business processing. The independent disk pool containing this data can be left varied off until it is needed. When large amounts of storage are routinely kept varied off, you can shorten processing time for operations such as IPL and reclaim storage.

In the figure, the user has five independent disk pools. They can represent three different applications where the third application might have archived data. The system automatically creates the system disk pool (referred to as *Disk Pool 1* or *ASP 1*), which contains all system programs and system data.



Appendix. Notices

This information was developed for products and services offered in the U.S.A.

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

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

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

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

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation

Software Interoperability Coordinator, Department YBWA
3605 Highway 52 N
Rochester, MN 55901
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

- | The licensed program described in this information and all licensed material available for it are provided
- | by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement,
- | IBM License Agreement for Machine Code, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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

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

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

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

COPYRIGHT LICENSE:

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

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_. All rights reserved.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Programming Interface Information

This (ADD NAME OF PUBLICATION HERE) publication documents intended Programming Interfaces that allow the customer to write programs to obtain the services of (ADD PRODUCT NAME HERE).

Trademarks

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

- | Advanced Function Presentation
- | AFP
- | AS/400
- | Enterprise Storage Server
- | eServer
- | FlashCopy
- | i5/OS
- | IBM
- | iSeries
- | OS/400
- | Redbooks
- | TotalStorage

- | Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

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

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

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

UNIX is a registered trademark of The Open Group in the United States and other countries.

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

Terms and conditions

Permissions for the use of these publications is granted subject to the following terms and conditions.

Personal Use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of IBM.

Commercial Use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.



Printed in USA