



IBM Systems - iSeries
Disk management

Version 5 Release 4





IBM Systems - iSeries
Disk management

Version 5 Release 4

Note

Before using this information and the product it supports read the information in "Notices," on page 139.

Fourth Edition (February 2006)

| 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, 2006. All rights reserved.

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

Contents

Disk management	1
What's new for V5R4	1
Printable PDF	1
Disk management concepts	2
Disk storage concepts	2
Disk pools	6
Disk protection	33
External load source unit	49
Plan for disk management	49
iSeries Navigator requirements for disk management	49
Plan for independent disk pools	51
Plan for disk protection	57
Configure your disks	72
Evaluate the current configuration.	72
Calculate disk space requirements.	73
Choose the correct procedure for configuring disks	74
Create a basic disk pool	91
Add a disk unit or disk pool	91

Move and remove disk units	92
Configure independent disk pools.	92
Protect data on disk units	106
Manage your disks	107
Manage disk units.	107
Manage disk pools	109
Manage independent disk pools	110
Work with device parity protection	120
Work with mirrored protection	123
Using independent disk pools.	125
Examples: Independent disk pool configurations	125
Frequently asked questions.	132
Related information for Disk management.	136

Appendix. Notices	139
Programming Interface Information	141
Trademarks	141
Terms and conditions.	141

Disk management

Use the information in this topic to effectively manage your disk units, disk pools, and independent disk pools. Find strategies to help you protect the data on your disk units.



What's new for V5R4

This topic highlights changes to disk management for V5R4.

- | • Checklist 12: Upgrade load source disk unit with local mirroring protection explains the procedure for upgrading the capacity of the load source disk unit.
- | • Use the Start DASD Management Operation (QYASSDMO) API to reset the correct path for missing multipath connections.
- | • An external storage unit located on a storage area network (SAN) can now be configured as the load source unit. For more information, see "External load source unit" on page 49.

How to see what's new or changed

To help you see where technical changes have been made, this information uses:

- The  image to mark where new or changed information begins.
- The  image to mark where new or changed information ends.

To find other information about what's new or changed this release, see the Memo to users.




Printable PDF

View or download a PDF version of this Disk management topic for viewing or printing.

To view or download the PDF version of this document, select Disk management (about 1.5 MB).

Other information

You can also view or print any of the following PDFs:

- Manuals:
 - Clusters
 - Independent disk pools
 - Storage solutions
 - Backup and Recovery manual 
- | • IBM[®] Redbooks[™]:
 - | – Clustering and IASPs for Higher Availability 
 - | – iSeries[™] Independent ASPs - A Guide to Moving Applications to IASPs 


Saving PDF files

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

1. Right-click the PDF in your browser (right-click the link above).
- | 2. Click the option that saves the PDF locally.

3. Navigate to the directory in which you want to save the PDF.
4. Click **Save**.

Downloading Adobe Reader

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

Disk management concepts

Learn about how disk storage works, how you can leverage disk pools in your environment. Protect your data with device parity protection and mirrored protection.

Before you use disk management mechanisms in your environment, it is important to understand some key concepts, including concepts about disk storage, disk pools, device parity protection, and mirrored protection.

Disk storage concepts

Disk units are assigned to a disk pool on a storage unit basis. The system treats each storage unit within a disk unit as a separate unit of auxiliary storage. When a new disk unit is attached to the system, the system initially treats each storage unit within it as nonconfigured. You can add these nonconfigured storage units to either the system disk pool, basic disk pool, or independent disk pool of your choosing. When adding nonconfigured storage units, use the serial number information that is assigned by the manufacturer to ensure that you are selecting the correct physical storage unit. Additionally, the individual storage units within the disk unit can be identified through the address information that can be obtained from the DST Display Disk Configuration display.

When you add a nonconfigured storage unit to a disk pool, the system assigns a unit number to the storage unit. The unit number can be used instead of the serial number and address. The same unit number is used for a specific storage unit even if you connect the disk unit to the system in a different way.

When a unit has mirrored protection, the two storage units of the mirrored pair are assigned the same unit number. The serial number and the address distinguish between the two storage units in a mirrored pair.


To determine which physical disk unit is being identified with each unit number, make note of the unit number assignment to ensure correct identification. If a printer is available, print the display of your disk configuration. If you need to verify the unit number assignment, use iSeries Navigator to display disk unit properties and check the serial numbers and addresses of each unit.

The storage unit that is addressed by the system as unit 1 is always used by the system to store licensed internal code and data areas. The amount of storage that is used on unit 1 is quite large and varies depending on the configuration of your system. Unit 1 contains a limited amount of user data. Because unit 1 contains the initial programs and data that are used during an IPL of the system, it is also known as the loadsource unit.


The system reserves a fixed amount of storage on units other than unit 1. The size of this reserved area is 1.08 MB per unit, reducing the space available on each unit by that amount.

Components of disk storage

The system uses several electronic components to manage the transfer of data from a disk to main storage. Data and programs must be in main storage before they can be used. This figure shows the hardware that is used for data transfer.



| This figure shows the hardware that is used for data transfer with a system that does not have an IOP.



Bus: The bus is the main communications channel for input and output data transfer. A system may have one or more buses.

I/O processor:

The input/output processor (IOP) is attached to the bus. The IOP is used to transfer information between main storage and specific groups of IOAs. Some IOPs are dedicated to specific types of IOAs, such as storage IOAs. Other IOPs can attach to more than one type of IOA, for example, communication IOAs and storage IOAs.

Input-output adapter (IOA):

The IOA attaches to the IOP and handles the information transfer between the IOP and the disk units.

Disk unit:

Disk units are the actual devices that contain the storage units. You order hardware at the disk-unit level. Each disk unit has a unique serial number.

The server accesses a disk unit by way of a logical address. The **logical address** consists of a system bus, a system card, an I/O bus, an I/O processor, an I/O adapter, I/O bus, and a device number.

To find the logical address for a disk storage component:


1. In iSeries Navigator, expand **My Connections** (or your active environment).
2. Expand the iSeries server you want to examine.
3. Expand **Configuration and Service**.
4. Expand **Hardware**.
5. Expand **Disk Units**.
6. Expand **All Disk Units**.
7. Right-click a disk unit and select **Properties**.

Considerations for multipath disk units

Multiple connections to the logical unit number (LUN) from the IOA processors was implemented in V5R3, up to eight connections can be defined from multiple input/output processors on an iSeries server to a single LUN in the Enterprise Disk Storage. The Enterprise Disk Storage concurrently supports diverse host systems over diverse attachment protocols. Data storage is allocated among the attached host systems with the Enterprise Disk Storage Specialist, a Web-based interface. Each connection for a multipath disk unit functions independently. Several connections provide availability by allowing disk storage to be utilized even if a single path fails.

When you use multipath disk units, you must consider the implications of moving IOPs and multipath connections between nodes. You must not split multipath connections between nodes, either by moving IOPs between logical partitions or by switching expansion units between systems. If two different nodes both have connections to the same LUN in the Enterprise Disk Storage, both nodes might potentially overwrite data from the other node.

The following figure shows a logical partition configuration with multiple paths configured correctly. Partition 0 (P0) has multiple paths defined from two different IOPs to one LUN in the Enterprise Disk Storage. Partition 1 (P1) also has multiple paths defined from two different IOPs to a different LUN in the Enterprise Disk Storage. The configuration would be incorrect if each logical partition had defined connections to a single LUN.



The system enforces the following rules when you use multipath disk units in a multiple-system environment:

- If you move an IOP with a multipath connection to a different logical partition, you must also move all other IOPs with connections to the same disk unit to the same logical partition.
- When you make an expansion unit switchable, make sure that all multipath connections to a disk unit will switch with the expansion unit.
- When you configure a switchable independent disk pool, make sure that all of the required IOPs for multipath disk units will switch with the independent disk pool.

If a multipath configuration rule is violated, the system issues warnings or errors to alert you of the condition. It is important to pay attention when disk unit connections are reported missing. You want to prevent a situation where a node might overwrite data on a LUN that belongs to another node.

Disk unit connections might be missing for a variety of reasons, but especially if one of the preceding rules has been violated. If a connection for a multipath disk unit in a system or basic disk pool is found to be missing during an IPL, a message is sent to the QSYSOPR message queue.

If a connection is missing, and you confirm that the connection has been removed, you can update Hardware Service Manager (HSM) to remove that resource. Hardware service manager is a tool for displaying and working with system hardware from both a logical and a packaging viewpoint, an aid for debugging Input/Output (I/O) processors and devices, and for fixing failing and missing hardware. You can access Hardware Service Manager in System Service Tools (SST) and Dedicated Service Tools (DST) by selecting the option to start a service tool.

Note: Multiple connections are not supported for a load source LUN in an Enterprise Disk Storage from an eServer™ i5 model.

| **Change multipath disk unit**

| Use the Start DASD Management Operation (QYASSDMO) API to remove the missing multiple-path disk units, resulting in only one path.

| When there is a missing connection to a disk unit this message Event EV0D0401, Event Code 1E, Message ID CPI096E appears in your QSYSOPR message queue.

Disk pools

A disk pool, also referred to as an auxiliary storage pool (ASP) in the character-based interface, is a software definition of a group of disk units on your system. This means that a disk pool does not necessarily correspond to the physical arrangement of disks. Conceptually, each disk pool on your system is a separate pool of disk units for single-level storage. The system spreads data across the disk units within a disk pool. If a disk failure occurs, you need to recover only the data in the disk pool that contained the failed unit.

Your system may have many disk units attached to it for disk pool storage. To your system, they look like a single unit of storage. The system spreads data across all disk units. You can use disk pools to separate your disk units into logical subsets. When you assign the disk units on your system to more than one disk pool, each disk pool can have different strategies for availability, backup and recovery, and performance.

Disk pools provide a recovery advantage if the system experiences a disk unit failure resulting in data loss. If this occurs, recovery is only required for the objects in the disk pool that contained the failed disk unit. System objects and user objects in other disk pools are protected from the disk failure.

Disk pool benefits

Placing objects in user disk pools, also called auxiliary storage pools (ASPs) in the character-based interface, can provide several advantages. These include the following:

Additional data protection

By separating libraries, documents, or other objects in a user disk pool, you protect them from data loss when a disk unit in the system disk pool or other user disk pool fails. For example, if you have a disk unit failure, and data contained on the system disk pool is lost, objects contained in user disk pools are not affected and can be used to recover objects in the system disk pool. Conversely, if a failure causes data that is contained in a user disk pool to be lost, data in the system disk pool is not affected.

Improved system performance

Using disk pools can also improve system performance. This is because the system dedicates the disk units that are associated with a disk pool to the objects in that disk pool. For example, suppose you are working in an extensive journaling environment. Placing journals and journaled objects in a “Basic disk pools” on page 11 can reduce contention between the receivers and journaled objects if they are in

different disk pools, which improves journaling performance. If you use independent disk pools to reduce contention, place the objects to be journaled in the primary disk pool and journal receivers in one or more secondary disk pools.

Placing many active journal receivers in the same disk pool is not productive. The resulting contention between writing to more than one receiver in the disk pool can slow system performance. For maximum performance, place each active journal receiver in a separate user disk pool.

Separation of objects with different availability and recovery requirements

You can use different disk protection techniques for different disk pools. You can also specify different target times for recovering access paths. You can assign critical or highly used objects to protected, high-performance disk units. You might assign large, low-usage files, like history files, to unprotected, low-performance disk units.

Related information

“Benefits of independent disk pools” on page 17

Disk pool costs and limitations

You may encounter specific limitations when you use disk pools (auxiliary storage pools):

- The system cannot directly recover lost data from a disk unit media failure. This situation requires you to perform recovery operations.
- Using disk pools can require additional disk devices.
- Using disk pools will require you to manage the amount of data in a disk pool and avoid an overflowed disk pool.
- You will need to perform special recovery steps if a basic disk pool overflows.
- Using disk pools requires you to manage related objects. Some related objects, such as journals and journaled objects, must be in the same user disk pool.

Disk pool uses

Disk pools are used to manage system performance and backup requirements, as follows:

- You can create a disk pool to provide dedicated resources for frequently used objects, such as journal receivers.
- You can create a disk pool to hold save files. Objects can be backed up to save files in a different disk pool. It is unlikely that both the disk pool that contains the object and the disk pool that contains the save file will be lost.
- You can create different disk pools for objects with different recovery and availability requirements. For example, you can put critical database files or documents in a disk pool that has mirrored protection or device parity protection.
- You can create a disk pool to place infrequently used objects, such as large history files, on disk units with slower performance.
- You can use disk pools to manage recovery times for access paths for critical and noncritical database files using system-managed access-path protection.
- An independent disk pool can be used to isolate infrequently used data in order to free up system resources to be utilized only when it is needed.
- An independent disk pool in a clustered environment can provide disk storage that is switchable, allowing continuous availability of resources.


Use disk pools for improved performance: If you are using user disk pools for better system performance, consider dedicating a disk pool to one object that is very active. In this case, you can configure the disk pool with only one disk unit.

However, it typically does not improve performance to place a single device parity-protected unit in a user disk pool because the performance of that unit is affected by other disk units in the device parity set.

Allocating one user disk pool exclusively for journal receivers that are attached to the same journal can improve journaling performance. By having the journal and journaled objects in a separate disk pool from the attached journal receivers, there is no contention for journal receiver write operations. The units that are associated with the disk pool do not have to be repositioned before each read or write operation.

The system spreads journal receivers across multiple disk units to improve performance. The journal receiver may be placed on up to ten disk units in a disk pool. If you specify the `RCVSIZOPT(*MAXOPT1)` or `(*MAXOPT2)` journal option, then the system may place the journal receiver on up to 100 disk units in a disk pool. If you add more disk units to the disk pool while the system is active, the system determines whether to use the new disk units for journal receivers the next time the change journal function is performed.

Another way to improve performance is to make sure there are enough storage units in the user disk pool to support the number of physical input and output operations that are done against the objects in the user disk pool. You might have to experiment by moving objects to a different user disk pool and then monitoring performance in the disk pool to see if the storage units are used excessively. For more information about working with disk status (`WRKDSKSTS` command) to determine if the storage units have excessive use, see Work Management. If the units have excessive use, you should consider adding more disk units to the user disk pool.

- | **Use disk pools with extensive journaling:** If journals and objects being journaled are in the same disk pool as the receivers and the disk pool overflows, you must end journaling of all objects and recover the
- | disk pool overflow. Backup and Recovery  describes how to recover a disk pool that is overfilled.

If the journal receiver is in a different disk pool than the journal, and the user disk pool that the receiver is in overflows, do the following:


1. Create a new receiver in a different user disk pool.
2. Change the journal (`CHGJRN` command) to attach the newly created journal receiver.
3. Save the detached receiver.
4. Delete it.
5. Clear the overfilled disk pool without ending journaling.
6. Create a new receiver in the cleared disk pool.
7. Attach the new receiver with the `CHGJRN` command.

Note: The Backup and Recovery book has more information about working with journal receivers when a disk pool overflows.

Types of disk pools

Fundamentally, a disk pool, also referred to as an auxiliary storage pool (ASP), is a software definition of a group of disk units on your system. This means that a disk pool does not necessarily correspond to the physical arrangement of disks. Conceptually, each disk pool on your system is a separate pool of disk units for single-level storage. The system spreads data across the disk units within a disk pool.

There are two main types of disk pools: system disk pools (system ASPs) and user disk pools (user ASPs). The two types of user disk pools are basic disk pools and independent disk pools. Independent disk pools are divided into primary disk pools, secondary disk pools, and UDFS disk pools. The following example and definitions explain the types of disk pools:



System disk pool

One system disk pool exists per iSeries server. The system automatically creates the system disk pool (Disk Pool 1), which contains disk unit 1 and all other configured disks that are not assigned to a basic or independent disk pool. The system disk pool contains all system objects for the i5/OS™ licensed program and all user objects that are not assigned to a basic or independent disk pool.

User disk pool (user ASP)

There are two types of user disk pools: basic disk pools and independent disk pools. You can create a user disk pool by grouping a set of disk units together and assigning that group to a disk pool (ASP).

Basic disk pools

A basic disk pool is used to isolate some objects from the other objects that are stored in the system disk pool. Basic disk pools are defined by the user. Data in a basic user pool is always accessible whenever the server is up and running. You can configure basic disk pools with numbers 2 through 32.

Independent disk pools

An independent disk pool is a disk pool that contains objects, the directories or libraries that contain the objects, and other object attributes such as authorization and ownership attributes. They are numbered 33 through 255. An independent disk pool can be made available (varied on) and made unavailable (varied off) to the server without restarting the system. When an independent disk pool is associated with a switchable hardware group, it becomes a switchable disk pool and can be switched between iSeries servers in a clustered environment. The three types of independent disk pools are user-defined file system, primary, and secondary:

User-defined file system (UDFS)

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 secondary disk pool.

Primary

An independent disk pool that defines a collection of directories and libraries and may have other secondary disk pools associated with it. A primary disk pool also defines a database for itself and other disk pools that may be added in its disk pool group. Primary disk pools can only be implemented on OS/400® V5R2 or i5/OS V5R3 and later.

Secondary

An independent disk pool that defines a collection of directories and libraries and must be associated with a primary disk pool. A possible use for a secondary disk pool might be to store journal receivers for the objects being journaled in the primary disk pool. Secondary disk pools can only be implemented on OS/400 V5R2 or i5/OS V5R3 and later.

Related information

“Independent disk pool terminology” on page 12


“Disk pool groups” on page 26

System disk pool: The system automatically creates the system disk pool (disk pool 1) which contains disk unit 1 and all other configured disks that are not assigned to a “Basic disk pools” on page 11. The system disk pool contains all system objects for the i5/OS licensed program and all user objects that are not assigned to a basic or independent disk pool.

Note: You can have disk units that are attached to your system but are not configured and are not being used. These are called nonconfigured disk units.

Capacity of the system disk pool: If the system disk pool fills to capacity, the system will end normal activities. If this occurs, you must perform an IPL of the system, and take corrective action (such as deleting objects) to prevent this from recurring.

You can also specify a threshold that, when reached, warns the system operator of a potential shortage of space. For example, if you set the threshold value at 80 for the system disk pool, the system operator message queue (QSYSOPR) is notified when the system disk pool is 80% full. A message is sent every hour until the threshold value is changed, or until objects are deleted or transferred out of the system disk pool, or until disk units are added to the system disk pool. If you ignore this message, the system disk pool fills to capacity, and the system ends abnormally.

You can use a third method of preventing the system disk pool from filling to capacity by using the QSTGLOWLMT and QSTGLOWACN system values. For more information, refer to “How to Change the Storage Threshold for the System Auxiliary Storage Pool” in Backup and Recovery book .

- | *Protect your system disk pool:* Use device parity protection or mirrored protection on the system disk pool
- | to reduce the chance that the system disk pool will lose all data. If the system disk pool is lost,
- | addressability to objects in every user disk pool will also be lost.

You can restore the addressability by restoring the entire system or by running the Reclaim Storage (RCLSTG) command. However, the RCLSTG command cannot recover object ownership. After you run the command, the QDFTOWN user profile owns all objects. You can use the Reclaim Document Library Object (RCLDLO) command procedure to recover ownership of document library objects.

Basic disk pools: You can create a user basic pool by grouping a set of disk units together and assigning that group to a disk pool. Basic disk pools can contain libraries, documents, and certain types of objects. Data in a basic user pool is always accessible whenever the server is up and running. You can configure basic disk pools with numbers 2 through 32. When storage for a basic disk pool is exhausted, the data can overflow into the system disk pool. This is different from an independent disk pool, which does not allow data to overflow into the system disk pool.

After you have disk pools configured, you should protect them by using “Work with mirrored protection” on page 106 or “Device parity protection” on page 34. See “Disk protection” on page 33 for more information.

Library user disk pools: Library user disk pools contain libraries and user-defined file systems (UDFS). Library user disk pool steps are easier to recover than nonlibrary user disk pools.

- Do not create system or product libraries (libraries that begin with a Q or #) or folders (folders that begin with a Q) in a user disk pool. Do not restore any of these libraries or folders to a user disk pool. Doing so can cause unpredictable results.
- Library disk pools may contain both libraries and document library objects. The document library for a user disk pool is called QDOCnnnn, where *nnnn* is the number of the disk pool.
- Journals and objects that are being journaled must be in the same disk pool. Place the journal receivers in a different disk pool. This protects against the loss of both the objects and the receivers if a disk media failure occurs.

In order to begin journaling, the journal (object type *JRN) and the object to be journaled must be in the same disk pool. Use the following commands to start journaling:

- Start Journal Physical File (STRJRNPF) command for physical files
- Start Journal Access Path (STRJRNAP) command for access paths
- Start Journal (STRJRN) command for integrated file system objects
- Start Journal Object (STRJRNOBJ) command for other object types

Journaling cannot be started again for an object that is saved and then restored to a different disk pool that does not contain the journal. The journal and the object must be in the same disk pool for journaling to be automatically started again for the object.

- No database network can cross disk pool boundaries. You cannot create a file in one disk pool that depends on a file in a different disk pool. All based-on physical files for a logical file must be in the same disk pool as the logical file. The system builds access paths only for database files in the same disk pool as the based-on physical file (temporary queries are not limited). Access paths are never shared by files in different disk pools. Record formats are not shared between different disk pools. Instead, a format request is ignored and a new record format is created.
- You can place an SQL collection in a user disk pool. You specify the destination disk pool when you create the collection.
- If the library user disk pool does not contain any database files, set the destination access path recovery time for the disk pool to *NONE. This might be true, for example, if the library user disk pool contains only libraries for journal receivers. If you set the access path recovery time to *NONE, this prevents the system from doing unnecessary work for that disk pool. System-managed access-path protection describes how to set access path recovery times.

Nonlibrary user disk pools: Nonlibrary user disk pools contain journals, journal receivers, and save files whose libraries are in the system disk pool.

If you are assigning access path recovery times for individual disk pools, you should set the target recovery time for a nonlibrary user disk pool to *NONE. A nonlibrary user disk pool cannot contain any database files and cannot, therefore, benefit from system-managed access-path protection (SMAPP). If you set an access path recovery time for a nonlibrary user disk pool to a value other than *NONE, this causes the system to do extra work with no possible benefit. System-managed access-path protection describes how to set access path recovery times.

Contrast basic and independent disk pools: Basic disk pools and independent disk pools, also called auxiliary storage pools (ASPs) in the character-based interface, are both useful to group disk units containing certain information together; however, they have some inherent differences:

- When the server performs an IPL, all of the disk units configured to a basic disk pool must be accounted for in order for the server to continue the IPL. Independent disk pools are not included in the IPL. When you vary on the independent disk pool, the node then verifies that all disk units are present.
- When an unprotected disk unit in a disk pool fails, it typically stops all normal processing on the server until it can be repaired. The total loss of a disk unit in a basic disk pool requires lengthy recovery procedures to restore the lost data before the server can IPL and resume normal operations.
- The data in a basic disk pool belongs to the attaching node and can only be directly accessed by that system. In an independent disk pool the data does not belong to the node, but it belongs to the independent disk pool. You can share the data in the independent disk pool between nodes in a cluster by varying it off of one node and varying it on to another node.
- When you create a basic disk pool, you assign the disk pool a number. When you create an independent disk pool, you name the disk pool and the system assigns a number.
- If a basic disk pool fills up, it can overflow excess data into the system disk pool. When this occurs, the disk pool loses the isolation and protection inherent in disk pools. Independent disk pools cannot overflow. If they did, they lose their independence. When the independent disk pool nears its threshold, you need to add more disk units or delete objects to create more storage space.
- When you make restricted changes to disk configuration in a basic disk pool you must have your server restarted to Dedicated Service Tools (DST). In an offline independent disk pool you do not need to have your server in DST mode to start or stop mirroring, start device parity protection, start compression, remove a disk unit, and so on.

Independent disk pools: An independent disk pool contains user data and all of the necessary system information associated with the data. An independent disk pool can be made available (varied on) and made unavailable (varied off) to the server without restarting the system. When an independent disk pool is associated with a switchable hardware group, it becomes a switchable disk pool and can be switched between iSeries servers in a clustered environment. An independent disk pool that is not associated with a cluster resource group is dedicated to one iSeries server. Independent disk pools can also function in conjunction with other independent disk pools in a disk pool group. Independent disk pools are numbered 33 through 255.

Before you use independent disk pools in your environment, it is important to understand some key concepts, including important terminology, as well as how independent disk pools work and how they can be beneficial.

Independent disk pool terminology: As you work with independent disk pools, you need to become familiar with the following terms. For more terms and concepts, you can access the information center glossary.

Important: The terms **independent auxiliary storage pool (ASP)** and **independent disk pool** are synonymous.

active state

In geographic mirroring, pertaining to the configuration state of a mirror copy that indicates geographic mirroring is being performed if the disk pool is available.

asynchronous

In geographic mirroring, pertaining to the mode where the program issuing the update waits until the operation is completed on the production copy and received for processing on the target system.

cluster

A collection of complete systems that work together to provide a single, unified computing

capability. An iSeries cluster is made up of only iSeries servers and is required when implementing switchable independent disk pools.

cluster resource group (CRG)

A collection of related cluster resources that defines actions to be taken during a switchover or failover operation of the access point of resilient resources. These resilient resources include applications, data, and devices. The group describes a recovery domain and supplies the name of the cluster resource group exit program that manages the movement of an access point. A device CRG contains a list of devices, such as independent disk pools. The independent disk pools can reside on a switchable entity that can be either an expansion unit (frame/unit) or an IOP. In iSeries Navigator, a device cluster resource group is referred to as a switchable hardware group.

cross-site mirroring (XSM)

A feature of i5/OS High Available Switchable Resources (Option 41) that provides geographic mirroring and the services to switch over or automatically cause a failover to a mirror copy, potentially at another location, in the event of an outage at the primary location.

data port services

The generic transport mechanism used by geographic mirroring to send updates from the source system that owns the production copy to the target system that owns the mirror copy.

detach In geographic mirroring, to disassociate a mirror copy from the production copy in order to use the mirror copy for a separate operation, such as to save data, to run reports, or to perform data mining. Detaching a mirror copy suspends geographic mirroring.

detached mirror copy

A detached mirror copy of an independent disk pool is a mirror copy that is disassociated from the production copy in a geographic mirroring environment.

device description

An object that contains information describing a particular device or logical unit (LU) that is attached to the system. A device description is a description of the logical connection between two LUs (local and remote locations). The system-recognized identifier for the object type is *DEVVD.

device domain

A device domain is a collection of cluster nodes that share device resources, such as independent disk pools. For independent disk pools, the resources are: virtual addresses, disk pool numbers and disk unit numbers. An independent disk pool can only be accessed by the nodes in one device domain.

disk pool

An auxiliary storage pool that contains only disk units.

disk pool group

Made up of a primary disk pool and zero or more secondary disk pools, each of which are independent in regard to data storage, but combine to act as one entity.

disk unit

A physical enclosure containing one or more disk drives.

expansion unit

A feature that can be connected to a system unit to provide additional storage and processing capacity.

failover

A cluster event where the primary database server or application server switches over to a backup system due to the failure of the primary server

geographic mirroring

A subfunction of cross-site mirroring (XSM) that generates a mirror image of an independent disk pool on a system, which is (optionally) geographically distant from the originating site for availability or protection purposes.

HSL (high-speed link) loop

The system-to-expansion unit connectivity technology that is required to use switchable independent disk pools residing on an expansion unit (frame/unit). The servers and expansion unit in a cluster using resilient devices on an external expansion unit must be on an HSL loop connected with HSL cables.

independent disk pool

Disk pools 33 to 255. One or more storage units that are defined from the disk units or disk-unit subsystems that make up addressable disk storage. An independent disk pool contains objects, the directories and libraries that contain the objects, and other object attributes such as authorization ownership attributes. An independent disk pool can be made available (varied on) and made unavailable (varied off) without restarting the system. An independent disk pool can be either a) privately connected to a single system b) switchable among multiple systems in a clustering environment or c) duplicated at another site by geographic mirroring. Synonymous with *independent auxiliary storage pool (ASP)*.

insync In geographic mirroring, pertaining to the mirror copy data state that indicates that the production and mirror copy have exactly the same contents.

library namespace

An attribute that can be set for the current thread. The library namespace is the set of objects and libraries that can be accessed in any independent disk pools in a disk pool group plus the libraries in the system disk pool and basic user disk pools (ASPs 2-32) using the regular library-qualified object name syntax. The Set Auxiliary Storage Pool Group (SETASPGRP) command sets the auxiliary storage pool (ASP) group for the current thread.

mirror copy

In geographic mirroring, an independent disk pool that is being geographically mirrored so that it is a replica of the production copy of the independent disk pool. If a switchover or failover causes the system that owns the mirror copy to become the current primary node, the mirror copy becomes the production copy of the independent disk pool. The mirror copy has current data only when geographic mirroring is active.

mirror copy state

In geographic mirroring, the geographic mirroring state of the mirror copy; for example, active, resume pending, resuming, and suspended.

mirror copy data state

In geographic mirroring, the current status of the data that is being geographically mirrored; for example, insync, usable, and unusable.

primary disk pool

An independent disk pool that defines a collection of directories and libraries and may have other secondary disk pools associated with it. A primary disk pool also defines a database for itself and other disk pools that may be added in its disk pool group. Primary disk pools can only be implemented on V5R2 or later of OS/400.

production copy

In geographic mirroring, the independent disk pool to which all production operations are directed. All disk write operations are directed here first and are then replicated to the mirror copy of the independent disk pool. The production copy always has current data.

reattach

In geographic mirroring, to reassociate the detached mirror copy with its production copy after user operations on the detached mirror copy are completed. When the detached mirror copy is

reattached, it is automatically synchronized to match the production copy again. All data on the detached mirror copy is cleared before it is reattached to the production copy.

resume

In geographic mirroring, to start performing geographic mirroring again after it is suspended.

resume pending state

In geographic mirroring, pertaining to the configuration state of a mirror copy that indicates that geographic mirroring requires synchronization but that the disk pool is currently unavailable. When the disk pool is made available, the mirror copy will be synchronized with the current data from the production copy.

resuming state

In geographic mirroring, the configuration state of the mirror copy that attempts to perform geographic mirroring and synchronization when the independent disk pool is available. The mirror copy state is resuming when it is not suspended or active.

secondary disk pool

An independent disk pool that defines a collection of directories and libraries and must be associated with a primary disk pool. Secondary disk pools can only be implemented on V5R2 or later of OS/400. See

site In cross-site mirroring, a location containing a node or nodes with access to either the production copy or mirror copy. The sites can be in close proximity or geographically dispersed.

site primary node

In cross-site mirroring, a node that owns the independent disk pool, either the production copy or mirror copy, at a particular site. The production-site primary node is also the primary node for the cluster resource group. The mirror-site primary node is a backup node in a cluster resource group.

source system

The system that currently owns the production copy of an independent disk pool in a cross-site mirroring (XSM) environment. The target system is a backup node in the cluster resource group and is the mirror site primary node. Changes to the production copy of an independent disk pool are replicated to the mirror copy of the independent disk pool that exist on a backup node within the recovery domain.

suspend

In geographic mirroring, to temporarily stop performing geographic mirroring. If the mirror copy contained usable data when suspended, the mirror copy still contains usable, though possibly outdated, data.

suspended state

In geographic mirroring, pertaining to the configuration state of the mirror copy that does not attempt to perform geographic mirroring when the independent disk pool is available. The mirror copy state is suspended when it is not resuming or active.

switchable entity

The physical resource containing the independent disk pools that can be switched between systems in a cluster. This can be an expansion unit containing disk units in a multiple system environment. This might also be an IOP containing disk units in an LPAR environment.

switchover

A cluster event where the primary database server or application server switches over to a backup system due to the manual intervention from the cluster management interface.

| **full synchronization**

| The geographic mirroring processing that copies data from the production copy to the mirror
| copy. During synchronization the mirror copy contains unusable data. When synchronization is
| completed, the mirror copy contains usable data.

| **partial synchronization**

| While the system is in a suspended state, changes made to the production copy are not sent to
| the mirror copy. If the production copy is suspended with tracking, any changes made to the
| production copy are instead tracked. Once geographic mirror is resumed and partial
| synchronization is started, those tracked changes are then sent to the mirror copy.

synchronous

In geographic mirroring, pertaining to the mode of geographic mirroring where the program that issues the update waits until the operation is completed to disk on both the production copy and the mirror copy. This mode ensures that once control is returned to the client, the operation is accurately reflected on both the production copy and the mirror copy.

SYSBAS

In the character-based interface, refers to the system disk pool 1 and all configured basic disk pools 2 through 32. Independent disk pools 33 through 255 are not included.

target system

A system that currently owns a mirror copy of an independent ASP in a cross-site mirroring (XSM) environment. Changes to the production copy of an independent disk pool on the source system are replicated to the mirror copy of the independent disk pool that exists on a target system.

| **tracking**

| A process that remembers changes that occur while geographic mirroring is suspended. When
| geographic mirroring resumes the system only synchronizes the tracked changes and does not
| perform a full synchronization.

UDFS disk pool

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 secondary disk pool.

unusable

In geographic mirroring, pertaining to the mirror copy data state that indicates that the mirror copy contains incoherent data. This occurs:

1. During synchronization because synchronization does not preserve the order of writes.
2. When the system performs geographic mirroring in asynchronous mode.

| **Note:** The mirror copy becomes usable during a vary-off of the production copy of the
| independent disk pool.

usable In geographic mirroring, pertaining to the mirror copy data state that indicates that the correct order of updates to the mirror copy from the production copy is being preserved, but the mirror copy may be outdated. The usable mirror copy data state occurs:

1. When the system performs geographic mirroring in synchronous mode.
2. After successfully suspending geographic mirroring.
3. When mirror copy is successfully detached.

Note:

- | 1. The mirror copy becomes usable during a vary-off of the production copy of the
| independent disk pool.
| 2. Number 1 and 2 do not apply during a synchronization.

vary off

To make an independent disk pool unavailable for its normal, intended use. All of the primary and secondary disk pools in a disk pool group will vary off together. Synonymous with *make unavailable*.

vary on

To make an independent disk pool available for its normal, intended use. All of the primary and secondary disk pools in a disk pool group are varied on together. Synonymous with *make available*.

Related information

“Types of disk pools” on page 8

“Disk pool groups” on page 26

Benefits of independent disk pools: There are two environments in which the use of independent disk pools can be beneficial: a multiple-system clustered environment and a single-system environment.

Multiple-system clustered environment

In a multiple-system clustered environment, where the servers are members of an iSeries cluster and an independent disk pool is associated with a switchable device in that cluster, independent disk pools can be switched between systems without having to perform an initial program load (IPL). The independent disk pool can be switched because the independent disk pool is self-contained. This can be a significant advantage because it allows for continuous availability of data, the primary benefit of independent disk pools.

Switchable independent disk pools can help you do the following:

- Keep data available to an application even in the event of a single system outage, either scheduled or unscheduled.
- Eliminate the process of replicating data from one system to another.
- In some situations, isolate disk unit failures within the independent disk pool.
- Achieve high availability and scalability.

A multiple-system environment also gives you the opportunity to perform geographic mirroring. Geographic mirroring allows you to maintain two identical copies of an independent disk pool at two sites that are geographically separated. By having a second copy of critical data at a remote location, you ensure greater protection and availability; for example, in the case of a natural disaster. If you configure the independent disk pools to be switchable, you increase your options to have more backup nodes to allow for failover and switchover of independent disk pools between systems at the same site, in addition to switchover and failover to a system at another site.

Single-system environment

In a single-system environment, where an independent disk pool is privately connected to a single server, independent disk pool or independent disk pool groups can be made unavailable, independent of other disk pools because the data in the independent disk pool or independent disk pool group is self-contained. The independent disk pool or independent disk pool group can also be made available, while the system is active, without having to perform an IPL. Using independent disk pools this way can be useful, for example, if you have large amounts of data that are not needed for normal day-to-day processing. The independent disk pool containing this data can be left offline until it is needed. When large amounts of storage are normally kept offline, you can shorten processing time for operations such as IPL and reclaim storage.

Single-system independent disk pools can help you do the following:

- Isolate low-use data with the ability to bring the independent disk pool online only when it is needed.
- Reduce system start time.
- Manage save and restore by independent disk pool.
- Reclaim storage by independent disk pool.
- Divide data between multiple databases.

- Isolate data associated with specific applications or associated with specific groups of users.
- Consolidate data on small systems to independent disk pools on a larger system. For example, in the case of multiple branch offices.
- Perform application maintenance that does not affect the entire system.

Related information

“Disk pool benefits” on page 6

How independent disk pools work: The key characteristic of an independent disk pool is its ability to be, of course, *independent* of the rest of the storage on a server. It is independent because the data in the independent disk pool is self-contained. This means that all of the necessary system information associated with the data resides within the independent disk pool. The unique qualities of an independent disk pool allow it to be switched in a multisystem environment and to be made available and unavailable in a single-system environment.

Independent disk pools are available only when you choose to make them available; they are not made available when you restart your server, unless you include code (“Example: Make independent disk pool available at startup” on page 23) to make them available. When you select to make a disk pool available, the disk pool goes through a process similar to that of restarting the server. While this processing takes place, the disk pool is in an Active state.

While the disk pool is in Active state, recovery steps are being performed. The disk pool is synchronized with other disk pools that may be in the disk pool group. Also, journaled objects are synchronized with their associated journal. System libraries are created for the primary disk pool: QSYSnnnnn, QSYS2nnnnn, QRCLnnnnn, QRCYnnnnn, QRPLnnnnn, SYSIBnnnnn (where *nnnnn* is the primary disk pool number, right-aligned and padded with zeros). For example, the QSYS library for independent disk pool 33 is QSYS00033.

At this time database cross-reference files will also be updated. The system libraries for the independent disk pool, QSYSnnnnn and QSYS2nnnnn, contain metadata not only for the independent disk pool, but also for the system disk pool. When the disk pool is made available, database cross-referencing clears the information related to SYSBAS and updates it with current information. The number and complexity of database file objects and SQL packages, procedures, and functions that need to be updated will play a role in the time it takes to make the disk pool available.

As the independent disk pool is made available, several server jobs are started to support the independent disk pool. In order for server jobs to remain unique on the server, those that service the independent disk pool are given their own simple job name when the disk pool is made available. The server jobs are essential to the operation of the disk pool; do not tamper with these server jobs. The following is a list of server jobs that are created (nn = number):

1. **QDBXnnnXR** - handles database cross-reference file server functions
2. **QDBXnnnXR2** - handles database cross-reference field (column) information
3. **QDBnnnSV01** - handles database, journal, and commitment control events
4. **QDBnnnSV02 through QDBnnnSVnn** - high priority jobs that service the database
5. **QDBnnnSVnn through QDBnnnSVnn** - low priority jobs that service the database

When the recovery process is completed, the disk pool is in an Available state, ready for you to use. When you make a disk pool group available, you will see a completion messages for each disk pool. If the make available process encounters problems, such as an object not synchronized with a journal, you will need to resolve the issues reported in the error messages. See the job log, the system operator message queue, and the history log to locate problems and to verify the make available process.

Supported and unsupported object types:

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	*LIND	*RCT
*CNNL	*EDTD	*MODD	*SOCKET
*COSD	*EXITRG	*M36	*SSND
*CRG	*FLR	*M36CFG	*S36
*CSPMAP	*IGCSRT	*NTBD	*USRPRF
*CSPTBL	*IGCTBL	*NWID	
*CTLD	*IMGCLG	*NWSD	
*DDIR	*IPXD	*PRDAVL	

Note: *DSTMF is the object type returned for stream files that are being accessed through the QNTC file system from a remote server. 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	*FNTTBL	*NODGRP	*SPLF
*CHRSF	*FORMDF	*NODL	*SQLPKG
*CHTFMT	*FTR	*OUTQ	*SQLUDT
*CLD	*GSS	*OVL	*SRVPGM
*CLS	*IGCDCT	*PAGDFN	*STMF
*CMD	*JOB D	*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 Y	*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 D

If an active subsystem references the job description object, *JOB D 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 SYSLIBLE() 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**

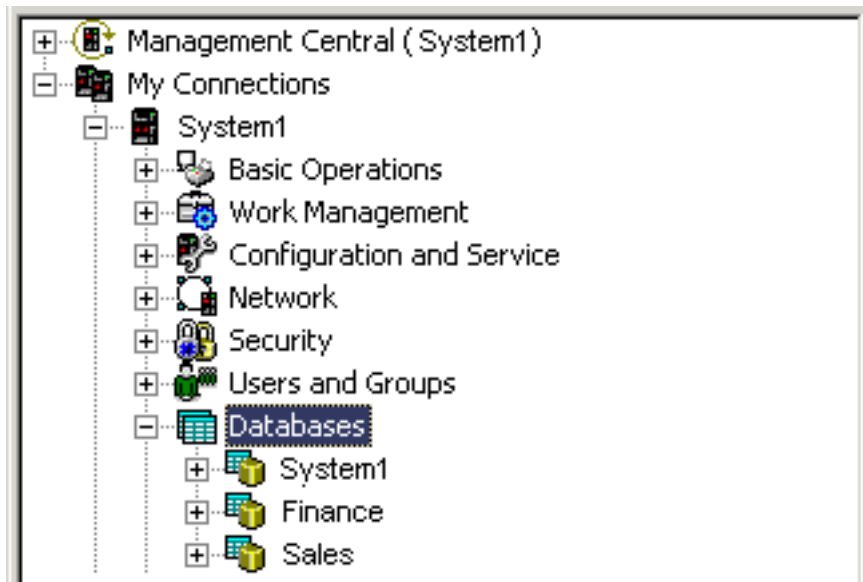
You cannot start a subsystem whose description is located in an independent disk pool.

Independent disk pools with distinct databases: When a primary independent disk pool is configured, a new user database is defined that is separate from the system database. The user database also includes any secondary disk pools that are associated with the primary disk pool. After the primary disk pool is configured, the corresponding user database appears in the Databases folder of iSeries Navigator. By default, the database and the independent disk pool have the same name. You administer the user database with the same functions that you use for the system database. See *Work with multiple databases* for more information.

The figure shows an example of a system with three distinct databases: the System database, the independent disk pool Finance database, and the independent disk pool Sales database.



In the following example, if you expand Databases in iSeries Navigator, you see a list of databases that includes the System database as well as the Finance and Sales user databases. From within a user database (Finance and Sales) you can always access libraries in the System database, but you cannot access libraries in another user database. For example, if you open the Finance database, you can select to display libraries from the System database as well. You cannot display Sales libraries from within the Finance database because Sales is a separate user database.



See “Object identification” on page 22 for details on identifying objects when independent disk pools exist on your server.

Multiple system libraries: In general, all system libraries continue to exist in the system disk pool. However, 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:

1. **QSYSnnnnn:** 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.
2. **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.
3. **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 .
4. **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.
5. **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.

In the above, nnnnn is the independent disk pool number right-aligned and padded with zeros.

One new library attribute, Protected, 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.

Library attribute settings

Library	*SYSBAS library	Protected in independent disk pool	Protected in system disk pool
QSYSnnnnn	QSYS	Yes	No

Library	*SYSBAS library	Protected in independent disk pool	Protected in system disk pool
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.

Object identification: Because the existence of an independent disk pool on a server means that multiple databases exist on a single server, 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.

Printing considerations:

Storing and printing spooled files

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.

| **Avoid duplicated spooled files**

| 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.

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

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

Example: Make independent disk pool available at startup: If you need your independent disk pools to be made available in most cases when your server is restarted, you might want to consider including the following control language (CL) at the beginning of your Start Up Program (QSTRUP). If you do not want the independent disk pools to be made available when you restart the server, you can either Delete the Data Area (DLTDARA) or Rename it (RNMOBJ). However, you need to remember to either Create

the Data Area (CRTDTAARA) again or Rename it (RNMOBJ) back to the data area being checked in your Start Up Program. Only the QSYSWRK subsystem should be started before making the independent disk pools available. Then other work will not compete for system resources while your independent disk pools are being made available.

In this example, the data area VARYONIASP is used. You can name your data area whatever you like. Also, in this example the QRECOVERY library contains the data area; however, you may choose a different library that resides on the system disk pool.

```
MONMSG MSGID(CPF0000)
QSYS/STRSBS SBSDB(QSYSWRK)
QSYS/CHKOBJ OBJ(QRECOVERY/VARYONIASP) OBJTYPE(*DTAARA)
MONMSG MSGID(CPF9801) EXEC(GOTO SKIPVRYCFG)
QSYS/VRYCFG CFGOBJ(IASP1) CFGTYPE(*DEV) STATUS(*ON)
QSYS/VRYCFG CFGOBJ(IASP2) CFGTYPE(*DEV) STATUS(*ON)
SKIPVRYCFG:
```

Recommended structure for independent disk pools: The recommended structure for using independent disk pools is to place the majority of your application data objects into independent disk pools and a minimal number of nonprogram objects in SYSBAS, which is the system disk pool and all configured basic disk pools. The system disk pool and basic user disk pools (SYSBAS) would contain primarily operating system objects, licensed program libraries, and few user libraries. This structure yields the best possible protection and performance. Application data is isolated from unrelated faults and can also be processed independently of other system activity. Vary on and switchover times are optimized with this structure.

Other advantages of this structure are:

- No library in the system disk pool is switchable.
- Since a database network cannot span an independent disk pool boundary, entire database networks are contained within disk pool groups.
- Coding of application transactions are simplified since all data libraries are contained within a single disk pool group.
- Library names can be duplicated across disk pool groups, but not between a disk pool group and the libraries in SYSBAS.

This recommended structure does not exclude other configurations. For example, you might start by migrating only a small portion of your data to a disk pool group and keeping the bulk of your data in SYSBAS. This is certainly supported. However, you should expect longer vary-on and switchover times with this configuration since additional processing is required to merge database cross-reference information into the disk pool group.

Structuring disk pool groups

The iSeries server supports up to 223 independent disk pools, any number of which can be primary, secondary, or user-defined file system (UDFS) disk pools. Therefore, you have significant flexibility in how you place your data into independent disk pools and how you structure disk pool groups. For example, all application data might be placed in a single disk pool group which consists of one primary disk pool and one secondary disk pool. Alternatively, you might create several disk pool groups, some with only a primary disk pool and some with one or more secondary disk pools.

Consider the following factors when planning the placement of your data in disk pools:

- If an application consists solely of data in user-defined file systems and the data is not to be journaled, a UDFS disk pool might be the best choice. There is less overhead associated with a UDFS disk pool. There is also less extensibility since the UDFS disk pool cannot contain any library-based objects.
- If you have an application with multiple instances of the application data that you want to keep separate, then you should consider a separate disk pool group for each data instance. See “Dedicated independent disk pools” on page 126 for an example of this scenario.

- If you have multiple applications and the application data is independent, a separate disk pool group for each application might be the appropriate answer. One application's data is then isolated from other applications and each application is unaffected by actions on others. The application data can therefore be brought online, taken offline, or switched without affecting other applications.
- If you have multiple applications with interdependent data objects, the data for those applications should be combined into a single disk pool group.
- You can use secondary disk pools to separate data objects into different storage domains and thus achieve better performance. The normal use of this is to separate your journal receivers onto different disk units from the data being journaled by placing the journal receivers in a secondary disk pool. However, you might also separate other parts of your application onto different disk units providing that they are in different libraries and the following journaling dependency is satisfied.
- Objects being journaled and the journal for those objects must be on the same disk pool.

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 an iSeries cluster, and a single-system environment with a single iSeries server.

Independent disk pools in a multisystem clustered environment

- | A group of servers in a cluster can take advantage of the switchover capability to move access to the independent disk pool from server to server. 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 input/output processor (IOP) on the bus shared by logical partitions, or an IOP or IOPless 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 servers 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 server. 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 servers at branch offices to one or more larger servers 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
iSeries 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 configurations” on page 125.

Disk pool groups: A disk pool group is made up 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. If you make one disk pool available or unavailable, the rest of the disk pools in the group are also made available or unavailable at the same time. Also, in a clustered environment, all of the disk pools in a group switch to another node at the same time.

An example of a practical use for a disk pool group is to isolate journal receivers from the objects for which they contain journal entries. The primary disk pool might contain the libraries, journals, and objects to be journaled, while the secondary disk pools might contain the associated journal receivers. The journals and journal receivers remain separate for maximum performance and recoverability, but they function together in the disk pool group.

If you delete a disk pool in a disk pool group, be aware of the effects it could have on other disk pools in the group. For example, when the original primary disk pool for a secondary disk pool is deleted, the existing secondary disk pool can be linked to a new primary disk pool only if that primary disk pool has never been made available.

Disk pool groups can only be implemented on OS/400 V5R2 or i5/OS V5R3 and later.

Related information

“Independent disk pool terminology” on page 12

“Types of disk pools” on page 8

Geographic mirroring: Geographic mirroring is a function that keeps two identical copies of an independent disk pool at two sites to provide high availability and disaster recovery. The copy owned by the primary node is the production copy and the copy owned by a backup node at the other site is the mirror copy. User operations and applications access the independent disk pool on the primary node, the node that owns the production copy.

Geographic mirroring is a subfunction of cross-site mirroring (XSM), which is part of i5/OS Option 41, High Available Switchable Resources.

Benefits of geographic mirroring: Geographic mirroring provides these benefits:

- Geographic mirroring provides site disaster protection by keeping a copy of the independent disk pool at another site which can be geographically distant. Having an additional copy at another geographic dispersed site improves availability.
- Geographic mirroring provides high availability with more backup nodes than switchable independent disk pools. In addition to having a production copy and mirrored copy, backup node possibilities are expanded when the independent disk pool is configured as switchable in an expansion unit (frame/unit), on an IOP on a shared bus, or on an IOP that is assigned to an I/O pool.

The geographically mirrored independent disk pool maintains all of the benefits of an independent disk pool, with its ability to be made available or unavailable as well as allow flexibility for the following actions:

- You can protect the production copy and mirror copy with your choice of protection, either disk unit mirroring or device parity protection. The production copy and mirror copy are not required to have the same type of protection.
- You can set the threshold of the disk pool to warn you when storage space is running low. The server sends a message, allowing you the time to add more storage space or to delete unnecessary objects. If the user ignores the warning and the mirror copy disk pool becomes full, geographic mirroring is suspended. If the user ignores the warning and the production disk pool becomes full, the application stops and objects cannot be created.
- The mirror copy can be detached and then separately made available to perform save operations, to create reports, or to perform data mining. When the mirror copy is reattached, it is synchronized with the production copy, and all modifications made to the detached copy are lost. Synchronization can be a lengthy process.
- If you configure the independent disk pools to be switchable, you increase your options to have more backup nodes that allow for failover and switchover methods.

Related concepts

“Device parity protection” on page 34

“Example: Independent disk pools with geographic mirroring” on page 131

Related information

“Mirrored protection” on page 43

“Set the threshold of a disk pool” on page 109

Costs and limitations of geographic mirroring:

Costs

To configure geographic mirroring between two sites, the following items are required:

- At least one iSeries server at each site.
- Sufficient CPU support for the additional CPU capacity required for geographic mirroring. A fraction of a processor for a partition supporting geographic mirroring is not adequate.
- Sufficient disk units at each site for the production and the mirror copy of the geographically mirrored independent disk pools. To avoid disk unit contention, use separate input/output adapters for the production copy on its node and for the mirror copy on its node.
- One TCP/IP connection from each node should connect the two sites. A second TCP/IP connection is strongly recommended to provide redundancy and better performance. You may configure up to four TCP/IP connections. See “Communications requirements” on page 55 for more information.

Limitations

Limitations of geographic mirroring include these constraints:

- When geographic mirroring is being performed, you cannot access the mirror copy; this ensures that the data integrity of the mirror copy is maintained.
- If you detach the mirror copy to perform a save operation, to perform data mining, or to create reports, you must reattach the mirror copy to resume geographic mirroring. The mirror copy must be synchronized with the production copy after it is reattached. Synchronization can be a lengthy process.
- Synchronization can be a lengthy process, especially if geographic mirroring was suspended without tracking.

How geographic mirroring works:

Configure

The nodes participating in geographic mirroring must be in the same cluster, the same device domain, and the same cluster resource group. Before configuring geographic mirroring, you must specify a site name and the TCP/IP address(es) for each node in the recovery domain. If you have more than one node at a site, then the hardware (disk units) you select for the disk pool must be switchable between the nodes at the site. If you only have one node at a site, the hardware does not have to be switchable and should be non-switchable (private).

See “Configure geographic mirroring with dedicated independent disk pools” on page 98 and “Configure geographic mirroring with dedicated independent disk pools” on page 98 for more information.

When geographic mirroring is configured, the mirror copy has the same disk pool number and name as the original disk pool, the production copy. Geographic mirroring is logical mirroring, not physical mirroring. The two disk pools must have similar disk capacities, but the mirror copy may have different numbers and types of disk units as well as different types of disk protection.

Manage

After geographic mirroring is configured, the production copy and mirror copy function as a unit. When the production copy is made available, the mirror copy is brought to a state that allows geographic mirroring to be performed. Synchronization occurs when you make the disk pool available after you configure geographic mirroring. When geographic mirroring is active, changes to the production copy data are transmitted to the mirror copy across TCP/IP connections. Changes can be transmitted either synchronously or asynchronously.

- **Synchronous mode:** The client waits until the operation is complete to disk on both the source and target systems. The mirror copy is always eligible to become the production copy, because the order of writes is preserved on the mirror copy. It is recommended to try synchronous mode first. If your performance remains acceptable, continue to use synchronous mode.
- **Asynchronous mode:** The client must wait only until the operation is complete to disk on the source system and is received for processing on the target system. However, synchronous mode is safer because if the primary node fails or the production copy fails, the mirror copy can become the production copy. In asynchronous mode, the pending updates must be completed before the mirror copy can become the production copy.

To maintain the data integrity of the mirror copy, the user cannot access the mirror copy while geographic mirroring is being performed. The user can detach the mirror copy to perform save operations, to create reports, and to perform data mining. However, the mirror copy must be synchronized with the production copy after it is reattached.

Tracking space

To suspend geographic mirroring with tracking, you set the tracking space when you configure geographic mirroring or change geographic mirroring attributes. Tracking space is allocated in the independent ASPs. The more tracking space you specify, the more changes the system can track. The maximum tracking space allowed is approximately 1% of the independent ASPs capacity.

Suspend geographic mirroring with tracking

If you suspend with tracking, the system will attempt to track changes made to those disk pools. This may reduce the synchronization process by performing partial synchronization when you resume geographic mirroring. If tracking space is exhausted, then when you resume geographic mirroring, complete synchronization is required.

Note: When you resume geographic mirroring, a complete synchronization can be a lengthy process, anywhere from several hours to even longer.

Suspend without tracking

If you suspend geographic mirroring without tracking changes, then when you resume geographic mirroring, a complete synchronization is required between the production and mirror copies. If you suspend geographic mirroring and you do track changes, then only a partial synchronization is required. Complete synchronization can be a very lengthy process, anywhere from one hour, to several hours, to even longer. The length of time it takes to synchronize is dependent on the number and type of disk units as well as how many TCP/IP communication interfaces are dedicated to geographic mirroring.

Synchronization

The production copy can function normally during synchronization, but performance might be negatively affected. During synchronization, the contents of the mirror copy are unusable, and it cannot become the production copy. If the independent disk pool is made unavailable during the synchronization process, synchronization resumes where it left off when the independent disk pool is made available again. Note that the first % complete message (CP1095D), after resuming an interrupted synchronization, shows 0%.

Synchronization type

| These are two types of synchronization:

| Full synchronization

- | • Indicates that a complete synchronization takes place. Changes to the production copy are not tracked to apply to the synchronization.
- | • Deletes all of the data on the mirror copy and copies all of the latest data from the production copy to the mirror copy.

| Partial synchronization

- | • Indicates that changes to the production copy are tracked to apply to the synchronization. This may shorten the synchronization time because a complete synchronization is unnecessary.

Synchronization priority

When you set the attributes for geographic mirroring, you can set the synchronization priority. If synchronization priority is set high, the system uses more resources for synchronization, which results in a sooner completion time. The mirror copy is eligible to become a production copy faster, so you are protected sooner. However, high priority can cause degradation to your application. It is recommended that you try high priority first, so you are protected as soon as possible. If the degradation to your

application performance is not tolerable, then lower the priority.

Recovery timeout

In addition to synchronization priority, you can also set recovery time out. The recovery timeout specifies how long your application can wait when geographic mirroring cannot be performed. When an error, such as IP failure, prevents geographic mirroring, the source system waits and retries for the specified recovery timeout before suspending geographic mirroring which allows your application to continue. The trade-off is between blocking your application or requiring synchronization after suspending geographic mirroring. When your application is blocked for an extended time, other jobs might also be blocked waiting for resources and locks owned by the applications using the geographic mirrored disk pool. When geographic mirroring is suspended, you no longer have the protection of the mirror copy. If your application can tolerate a delay, it is recommended to set recovery timeout from 2 to 5 minutes. If the volume of your data is large (over a terabyte), consider a longer recovery timeout value to reduce the possibility of suspending geographic mirroring. If mirroring is suspended without tracking, the system performs a full synchronization. If geographic mirroring is suspended with tracking, the system performs a partial synchronization.

System roles

When you configure the cluster for geographic mirroring, you have many options for defining the availability and protection of the independent disk pool. When you create the switchable hardware group, you list the order of the backup systems to which the independent disk pool will failover or switch over. If the primary node switches to a backup node at the same site, a hardware switch will occur. If the primary node switches to the other site, the mirror copy on the backup node changes roles to become the production copy. The old primary node becomes the new backup node, and the production copy becomes the mirror copy. The new production copy is now accessible for updates on the remote system. If the independent disk pools are part of a disk pool group, all of the disk pools in the group will switchover together. See “Example: Independent disk pools with geographic mirroring” on page 131.

Requirements for geographic mirroring:

- Geographic mirroring increases CPU load, so there must be sufficient excess CPU capacity. Add processors as needed to increase CPU capacity.
- For optimal performance for geographic mirroring, particularly during synchronization, increase your machine pool size by the amount given by the following formula:

The amount of extra machine pool storage is: $271.5 \text{ MB} + .2\text{MB} * \text{Number of disk units in independent ASPs}$.

The extra machine pool storage is required on the target node. However, because the target node changes when doing switchovers and failovers, you should increase the machine pool on all nodes in the Cluster Resource Group. To prevent the performance adjuster from reducing the machine pool size, you should do one of the following:

1. Set the machine pool minimum size to the calculated amount (the current size plus the extra size for geographic mirroring from the formula) using Work with Shared Storage Pools (WRKSHRPOOL) command or Change Shared Storage Pool (CHGSHRPOOL) command.

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

2. Set QPFRADJ to zero which prohibits the performance adjuster from changing size of the machine pool.
- Configure a separate storage pool for the jobs using geographic mirrored independent disk pools, especially if you specify a long recovery timeout.
 - Geographic mirroring is performed when the disk pool is available. When geographic mirroring is being performed, the system value for the time of day (QTIME) should not be changed.

- Communications requirements for independent disk pools are particularly critical as they affect throughput. See “Communications requirements” on page 55 for more information.
- All independent disk pool requirements must be met. See “Plan for independent disk pools” on page 51 for more information.

Failover and switchover:

Mirror copy failover or switchover

A failover or switchover of the mirror copy when the independent disk pool is online results in a synchronization.

A failover or switchover of the mirror copy to another node at that site when the independent disk pool is online results in a synchronization.

When geographic mirroring is suspended

While geographic mirroring is suspended, a switchover or failover to the mirror copy is prohibited because the mirror copy contains back-level data. However, in the case where the production copy is lost, you can change the order of the recovery domain nodes to convert such a back-level mirror copy into the production copy. Do this by changing the backup node which owns the mirror copy into a primary node. If geographic mirroring is suspended for some of the independent disk pools in the disk pool group, but not all of the independent disk pools in the disk pool group, you cannot convert the mirror copy into a production copy even by changing the order of the recovery domain nodes. If geographic mirroring is suspended for all of the independent disk pools in the group, you can change the order of the recovery domain names. If the independent disk pools were suspended at different times, then the mirror copies are inconsistent and you should not try to convert these inconsistent mirror copies into the production copy.

Examples

The following are examples of failovers and switchovers:

- If the backup node is at the same site as the current primary node, then a failover or switchover of the primary node causes the production copy to switch hardware to that backup node. The former backup node at the same site becomes the primary node. The new primary node performs geographic mirroring to a node at the mirror copy site.
- If the backup node is at the other site, then a failover or switchover of the primary node causes the production copy to swap roles with the mirror copy on the backup node. The former backup node at the other site becomes the primary node. One of the remaining nodes in the recovery domain becomes the backup node at the new mirror copy site.
- If the backup node that owns the mirror copy undergoes a failover or switchover, then the mirror copy moves to the next backup node.
- If the backup node that owns the mirror copy undergoes a failover or switchover and no other backup node is defined, then geographic mirroring is suspended.

Note: A full or partial synchronization is required once geographic mirroring resumes after being in a suspended state.

Ending clustering

Do not end clustering on a node that is performing geographic mirroring. Such nodes own either a production copy or a mirror copy. The following results occur when ending clustering while performing geographic mirroring:

- Ending clustering for the node that owns the production copy when the cluster resource group is active causes failover.

- Ending clustering for the node that owns the mirror copy when the cluster resource group is active causes failover of the mirror copy.
- Ending clustering for the node that owns the mirror copy when failover cannot occur, because the cluster resource group is inactive or because there is no other active node at the mirror copy site, prevents recovery from TCP/IP connection failures.

If you ended clustering inadvertently, you should restart clustering, make the independent disk pools in the cluster resource group unavailable at your first opportunity, then make the Independent ASPs available again. When clustering is ended, geographic mirroring cannot recover from certain communications failures until both clustering and geographic mirroring are restarted.

Shut down system

If the system owning the mirror copy must be shutdown while performing geographic mirroring, you should do one of the following to avoid causing the application on the production copy to wait for the recovery timeout:

- If another active node is at the mirror copy site, switch the mirror copy to the other node. As part of the switchover, geographic mirroring is suspended, but without the timeout delay.
- If no other active node is at the mirror copy site, suspend geographic mirroring before shutting down the mirror copy system which avoids the recovery timeout delay. Synchronization is required once geographic mirroring is suspended.

| **Note:** After suspending geographic mirroring, a full resynchronization is required when tracking is used
 | or a partial synchronization is required when tracking is not used. Synchronization is required
 | once geographic mirroring is resumed.

Do not shut down the TCP system on a node that is performing geographic mirroring. Such nodes own either a production copy or a mirror copy. The following results occur if the TCP system is shut down:

- If TCP is shutdown on production copy node and cluster resource group is active, failover occurs to the mirror copy.
- If TCP is shutdown on mirror copy node, geographic mirroring is suspended.

Recovery from two production copies

For successive failovers when performing geographic mirroring, the situation can arise that you have two production copies. Ordinarily, the production copy and the mirror copy remain consistent, so the next make available or resume automatically changes the former production copy to become the mirror copy, and the next make available will synchronize the new mirror copy. However, when the two nodes were not communicating, the users may have made both production copies available independently by suspending geographic mirroring. In this case, the system does not know which production copy the user wants. You must resolve the inconsistency by changing the recovery domain order. Once the node to serve as the production copy has been selected, the other production copy node becomes a mirror copy and is synchronized to the production copy.

Considerations for making a disk pool available at failover or switchover

When you specify *ONLINE for the Configuration object online, the system automates the vary-on as part of failover or switchover; therefore, you do not have to issue the vary-on. However, if a geographic mirroring problem occurs during the vary-on, the system suspends geographic mirroring and completes the vary-on. You might prefer to fix the problem and keep geographic mirroring active. Also, if the vary-on fails, the system attempts to go back to the original primary node and vary-on the independent ASP back to the original primary node. You might prefer to fix the problem and vary-on the independent ASP to the new primary node.

Rolling upgrades: Upgrades of i5/OS releases made to any nodes involved in geographic mirroring require a rolling upgrade. The system will perform geographic mirroring from a V5R3M0 node owning the production copy to a V5R4M0 node owning the mirror copy. A rolling upgrade is required because a node at an earlier release might not be able to perform geographic mirroring to a node at a later release and a node at a later release usually cannot perform geographic mirroring to a node at an earlier release. This forces the nodes to be upgraded in an order dictated by the recovery domain order starting with the node that is the last backup. During the rolling upgrade, the production copy and mirroring copy will be moved to their recovery nodes.

In the following example, four nodes at two sites supporting mutual takeover will be upgraded. Nodes A and B are on one site with nodes C and D at another site. Node A owns the production copy of independent disk pool 33, and node C owns the mirror copy of independent disk pool 33. Node C owns the production copy of independent disk pool 34, and node A owns the mirror copy of disk pool 34.

Steps	Recovery domain order			
	Independent disk pool 33		Independent disk pool 34	
	During	After	During	After
Initial		A, B, C, D		C, D, A, B
1. Upgrade node D	A, B, C	A, B, C, D	C, A, B	C, A, B, D
2. Upgrade node B	A, C, D	A, C, B, D	C, A, D	C, A, D, B
3. Switch production copy of independent ASP 34 (C to D)				D
4. Switch mirror copy of independent ASP 34 (A to B)				D, B
5. Switch mirror copy of independent ASP 33 (C to D)		A, B, D		
6. Upgrade node C	A, B, D	A, B, D, C	D, B	D, C, B
7. Switch mirror copy of independent ASP 33 (D to C)		A, B, C, D		D, C, B
8. Switch prod copy of independent ASP 34 (D to C)		A, B, C, D		C, D, B
9. Switch production copy of independent ASP 33 (A to B)		B, C, D		C, D, B
10. Upgrade node A	B, C, D	B, A, C, D	C, D, B	C, D, B, A
11. Switch production copy of independent ASP 33 (B to A)		A, B, C, D		C, D, B, A
12. Switch mirror copy of independent ASP 34 (B to A)	A, B, C, D			C, D, A, B

In step 3 of the table, notice that node A cannot mirror node D because node D is release n+1 while node A is still at release n. Therefore, the mirror copy for Independent ASP 34 is switched to node B which is now at release n+1. Steps 7, 11, and 12 (new numbers after added steps) are not strictly required and can be done later or omitted. They were done here to return the roles to their preferred owners.

Disk protection

It is important to protect all the disk units on your system with either device parity protection or mirrored protection. This prevents the loss of information when a disk failure occurs. In many cases, you can keep your system running while a disk unit is being repaired or replaced. Your system can continue to run in the following scenarios:

- If a disk failure occurs in a disk pool that has mirrored protection.

- If one disk unit in a device parity set fails with RAID 5.
- If two disk units in a device parity set fail with RAID 6.

Device parity protection

Device parity protection is a hardware availability function that protects data from being lost because of disk unit failure or because of damage to a disk. To protect data, the disk input/output adapter (IOA) calculates and saves a parity value for each bit of data. Conceptually, the IOA computes the parity value from the data at the same location on each of the other disk units in the device parity set. When a disk failure occurs, the data can be reconstructed by using the parity value and the values of the bits in the same locations on the other disks. The system continues to run while the data is being reconstructed. The overall goal of device parity protection is to provide high availability and to protect data as inexpensively as possible.

Two types of device parity protection

RAID 5 and RAID 6 are the two types of device parity protection.

RAID 5

If more than one disk fails, you must restore the data from the backup media. Logically, the capacity of one disk unit is dedicated to storing parity data in a parity set. However, in practice the parity data is spread among multiple disk units. Restoring data to a disk pool that has disk units with device parity protection may take longer than a disk pool that contains only unprotected disk units.

Systems with IOAs released after V5R2 can have a minimum number of 3 disk units in a parity set. The maximum number of disk units in a parity set is 18.

Note: Systems with IOAs released before V5R2 of i5/OS, the minimum number of disk units in a parity set is 4. The maximum number of disk units in a parity set is 10.

Number of disk units in a parity set	Number of disk units that store parity
3	2
4-7	4
8-15	8
16-18	16

RAID 6

If more than two disk units fail, you must restore the data from the backup media. Logically, the capacity of two disk units is dedicated to storing parity data in a parity set. However, in practice the parity data is spread among multiple disk units.

The minimum number of disk units in a parity set is 4. The maximum number of disk units in a parity set is 18.

When a RAID 6 parity set is started, all of the disk units contain parity. Restoring data to a disk pool that has disk units with device parity protection may take longer than a disk pool that contains only unprotected disk units.

Note: It is recommended that you use more than four disk units in a RAID 6 device parity set, because the capacity of two disk units is dedicated to storing parity data in a parity set.

| RAID 6 requires a new storage adapter that supports this new function. The 571B is the first adapter to support RAID 6.

Device parity protection is not a substitute for a backup and recovery strategy

Device parity protection is not a substitute for a backup and recovery strategy. Device parity protection can prevent your system from stopping when certain types of failures occur. It can speed up your recovery process for certain types of failures. But device parity protection does not protect you from many types of failures, such as a site disaster or an operator or programmer error. It does not protect against system outages that are caused by failures in other disk-related hardware (such as IOAs, disk I/O processors, or a system bus).

If possible, you should protect all the disk units on your system with either device parity protection or “Work with mirrored protection” on page 106. This prevents the loss of information when disk failure occurs. In many cases, you can also keep your system operational while a disk unit is being repaired or replaced.

For information about how to start using device parity protection, see Backup and Recovery .

Related concepts

“Example: Independent disk pools with geographic mirroring” on page 131

Related information

“Benefits of geographic mirroring” on page 27

“Mirrored protection” on page 43

“Set the threshold of a disk pool” on page 109

Benefits of device parity protection:

RAID 5

- Lost data is automatically reconstructed by the IOA after a disk failure.
- The system continues to run after a single disk failure.
- A failed disk unit can be replaced without stopping the system.
- Device parity protection reduces the potential number of objects that could be damaged when a disk fails.
- Only one disk unit of capacity stores parity data in a parity set.

RAID 6

- Lost data is automatically reconstructed by the IOA after a disk failure.
- The system continues to run after two disks fail.
- Two failed disk units can be replaced without stopping the system.
- Device parity protection reduces the potential number of objects that could be damaged when a disk fails.
- Two disk units of capacity are dedicated to storing parity data in a parity set.

Costs and limitations of device parity protection: Here are the costs and limitations of device parity protection:

- Device parity protection can require additional disk units to prevent slower performance.
- Restore operations can take longer when you use device parity protection.

RAID 5: The system is only capable of handling one disk-unit failure. If more than one disk unit fails, the system may also fail depending on the ASP configuration.

RAID 6: The system is capable of handling up to two disk-unit failures. However, because the amount of parity data is twice as much as parity data in RAID 5, the available storage for user data is reduced. If more than two disk units fail, the system may also fail depending on the ASP configuration.

How device parity protection works: RAID 5

The input/output adapter determines how parity sets are formed. For V5R2 and later input/output adapters, you do have the ability to choose how you want the parity set to be optimized. You can optimize according to *availability*, *capacity*, *performance*, or a *balanced* version. A parity set optimized for availability offers a greater level of protection, because it allows a parity set to remain functional in the event of a single SCSI bus failure on the IOA. The parity set is formed from at least three disk units of equal capacity each attached to a separate SCSI bus on the input/output adapter (IOA). If you optimize by capacity, the IOA tends to create parity sets with a greater number of disk units. You will increase space used for storing user data, but performance may not be as high. If you optimize for performance, the IOA tends to create a parity set with fewer disk units. This should contribute to faster read and write operations, but might also dedicate slightly more disk capacity to storing parity data.

It is possible to include additional disk units of the same capacity in a device parity set after device parity protection is initially started. You can include up to two disk units at the same time; however, if three or more disk units are present and eligible for device parity protection, the system requires that you start a new parity set, rather than include them in an existing parity set. In iSeries Navigator you can view the properties of each disk unit. If the protection status of a disk unit is *unprotected*, it is not protected by device parity protection or mirroring and may be eligible to be included in a parity set or to be started in a new parity set. This will also be indicated by the model number which should be 050 (or 060 if it is a compressed disk unit). You can also exclude disks that do not store parity data from a parity set without stopping device parity protection. You can exclude a *protected* unit with a model number, for example 070 (or 080 if it is a compressed disk unit), because it is a disk unit that does not store parity data.

When a device parity set grows you may want to consider redistributing the parity data. For example you may begin with seven or fewer disk units, but expand to eight or more by including more disk units. When this happens, you can improve the performance on the device parity set by stopping parity protection and starting it again. This redistributes the parity data across eight disks rather than four. In general, spreading the parity data across more disk units improves performance.

A write cache is included in the input/output adapter (IOA) for each parity set to improve performance of interactive write workloads.

Note: If possible, start device parity protection before adding disk units to a disk pool. This significantly reduces the time it takes to start device parity and configure the disk units.

RAID 6

The input/output adapter determines how parity sets are formed. RAID 6 protection gives you optimal performance, capacity, and balance, so selecting any of these parity set optimizations is meaningless and will not affect the outcome of the parity set. If you choose to optimize by availability, a greater level of protection is achieved, because it allows a parity set to remain functional in the event of a single SCSI bus failure on the IOA. The parity set is formed from at least four disk units of equal capacity, with no more than two disk units attached to an individual SCSI bus on the input/output adapter (IOA). Device parity protection reduces the potential number of objects that could be damaged when a disk fails.

It is possible to include additional disk units of the same capacity in a device parity set after device parity protection is initially started. You can include up to two disk units at the same time; however, if three or more disk units are present and eligible for device parity protection, the system requires that you start a new parity set, rather than include them in an existing parity set. In iSeries Navigator you can view the properties of each disk unit. If the protection status of a disk unit is *unprotected*, it is not protected by device parity protection or mirroring and may be eligible to be included in a parity set or to be started in a new parity set. This will also be indicated by the model number which should be 050. You

can also exclude disks that do not store parity data from a parity set without stopping device parity protection. You can exclude a *protected* unit with a model number of 090 because it is a disk unit that does not store parity data.

When a device parity set grows, you may want to consider redistributing the parity data. For example you may begin with seven or fewer disk units, but expand to ten or more by including more disk units. When this happens, you can improve the performance on the device parity set by stopping parity protection and starting it again.

A write cache is included in the input/output adapter (IOA) for each parity set to improve performance of interactive write workloads.

Note: If possible, start device parity protection before adding disk units to a disk pool. This significantly reduces the time it takes to start device parity protection and configure the disk units.

Migrating to a new input/output adapter: RAID 5

Before you begin the migration to the new input/output adapter (IOA), as with any configuration change, it is important to do a normal system turn off. This will assure that all of your cache data is written to disk before the power down completes. When a parity set under an IOA that was released before the V5R2 release is migrated to an IOA that was released after the V5R2 release, your disk units are not protected by device parity protection while parity is being regenerated.

Cannot migrate to old generation adapters

You cannot migrate a parity set back to the old generation of adapters after you have made the change to a new adapter. You cannot migrate a parity set back to the old generation of adapters and keep the data intact. This action requires a save and restore of disk unit data to prevent data loss. To migrate RAID 5 protection to RAID 6 or RAID 6 protection to RAID 5 you must stop and restart the device parity protection.

Note: You cannot migrate RAID 6 to an adapter that does not support RAID 6.

Elements of device parity protection: The following diagrams illustrate the elements of a parity set that contains four disk units. Each parity set begins with an input/output processor (IOP) that is attached to an input/output adapter (IOA), which contains the write cache. The IOA transmits read and write signals to the attached disk units.

P indicates the sections of the disk that contain parity data.

Q indicates the second stripe of parity data.

Note: The second stripe of parity data is only associated with RAID 6 protection.

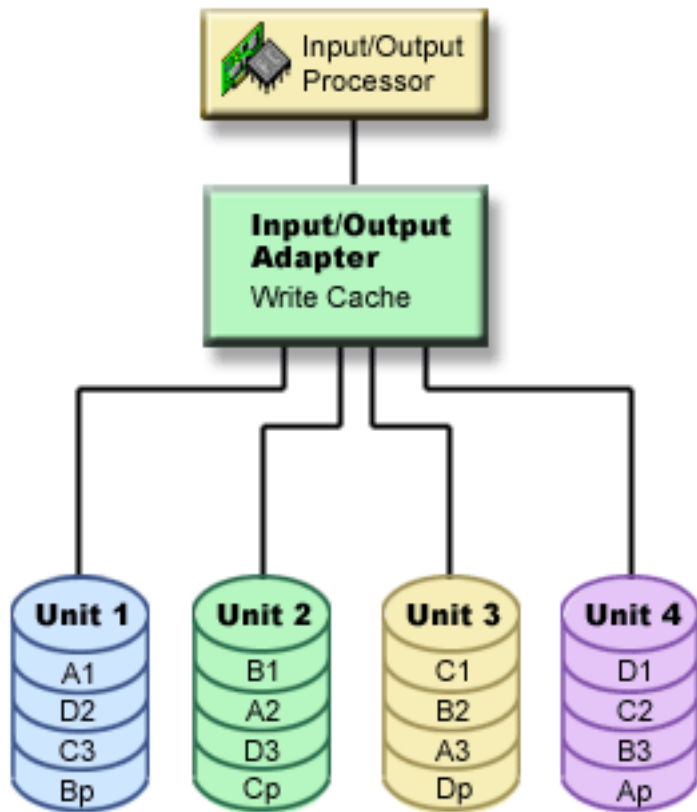


Figure 1. Example of how parity data is distributed with RAID 5 pre-V5R2 IOAs




Figure 2. Example of how parity data is distributed with RAID 5 post-V5R2 IOAs

Performance is improved by spreading the parity data throughout each of the disk units. The device parity protection that is spread throughout the disk units equals one disk unit of memory.




Figure 3. Example of how parity data is distributed with RAID 6

Performance is improved by spreading the parity throughout each of the disk units. The total amount of protection that is spread throughout the disk units equals two disk units of memory.

How device parity protection affects performance: **RAID 5**

Device parity protection requires extra I/O operations to save the parity data. To avoid performance problems, all IOAs contain a nonvolatile write cache that ensures data integrity and provides faster write capability. The system is notified that a write operation is complete as soon as a copy of the data is stored in the write cache. Data is collected in the cache before it gets written to a disk unit. This collection technique reduces the number of physical write operations to the disk unit. Because of the cache, performance is generally about the same on protected and unprotected disk units.

Applications that have many write requests in a short period of time, such as batch programs, can adversely affect performance. Disk unit failures can adversely affect the performance for both read and write operations.

The additional processing that is associated with a disk unit failure in a device parity set can be significant. The decrease in performance is in effect until both the failed unit is repaired (or replaced) and the rebuild process is complete. If device parity protection decreases performance too much, consider using mirrored protection.

RAID 6

Because there is a capacity of two disk units dedicated to storing parity data in a parity set for RAID 6, more I/O operations occur with RAID 6 than RAID 5. This may cause the performance to decrease.

Read operations on a failed disk unit: To access the data that was contained on a failed disk unit, device parity protection must read each disk unit in the device parity set that contains the failed disk unit. Because the read operations can be overlapped, the performance affect may be small.

Because a failed disk unit with device parity protection may contain only a small portion of user data, it is possible that only a few users will be affected by the decrease in performance.

Note: RAID 6 operations are derived from RAID 5, but at a further level of complexity. Because the concept is similar to RAID 5, RAID 6 operations are not described here.

Write operations on a failed disk unit:

Some examples show what happens to write operations when a single disk unit fails in a device parity set with device parity protection. The following figure shows a failed unit under an IOA with device parity protection.

It shows a parity set with four disk units. Each section of the disk unit is marked with a number. Parity sectors are noted with a *p*. Disk unit 3 is failed. Disk unit 1 shows sectors 1, 2, 3, and 4p. Disk unit 2 shows sectors 4, 1, 2, and 3p. Failed disk unit 3 shows sectors 3, 4, 1, and 2p. Disk unit 4 shows sectors 2, 3, 4, and 1p.




Figure 4. Device parity set with failed disk unit.

Note: RAID 6 operations are derived from RAID 5, but at a further level of complexity. Because the concept is similar to RAID 5, RAID 6 operations are not described here.

Example: Writing to a failed disk unit: A write operation from the iSeries server detects that the disk unit that is to contain the data has failed. The write operation is to disk unit 3, sector 1. The following actions occur:

1. The original data is lost on disk unit 3, sector 1, because of the failure.
2. The new parity data is calculated by reading disk unit 1, sector 1, and disk unit 2, sector 1.
3. New parity information is calculated.

4. New data cannot be written to sector 1 on disk unit 3 because of the failure.
5. New parity information is written to parity sector 1 on disk unit 4.
Write operations require multiple read operations (n-2 reads, where n is the number of disk units) and only one write operation for the new parity information. Data from disk unit 3 will be rebuilt during synchronization after disk unit 3 is replaced.

Example: Writing data to a disk unit when its corresponding parity data is on a failed disk unit: The write request from the iSeries server detects a disk failure for the disk unit that contains the corresponding parity data. The write request is to sector 2 on disk unit 4. Parity information for disk unit 4, sector 2, is on failed disk unit 3. The following actions occur:

1. A failure is detected on the disk unit that contains the parity data, disk unit 3.
2. Calculating parity information is not required because it cannot write to parity sector 2 of disk unit 3. Therefore, there is no requirement to read the original data and the parity information.
3. Data is written to disk unit 4, sector 2.
A write operation requires only one write operation for the new data. Parity data for parity sector 2 on disk unit 3 will be rebuilt during synchronization after disk unit 3 is replaced.

Input/output operations during a rebuild process: I/O operations during the rebuild (synchronization) process of the failed disk unit may not require additional disk I/O requests. This depends on where the data is read from or written to on the disk unit that is in the synchronization process. For example:

- A read operation from the disk area that already has been rebuilt requires one read operation.
- A read operation from the disk area that has not been rebuilt is treated as a read operation on a failed disk unit. See “Read operations on a failed disk unit” on page 41 for more information.
- A write operation to the disk area that has already been rebuilt requires normal read and write operations (two read operations and two write operations).
- A write operation to the disk area that has not been rebuilt is treated as a write operation to a failed disk unit. See “Write operations on a failed disk unit” on page 41 for more information.

Note:

1. The rebuild process takes longer when read and write operations to a replaced disk unit are also occurring. Every read request or every write request interrupts the rebuild process to perform the necessary I/O operations.
2. RAID 6 operations are derived from RAID 5, but at a further level of complexity. Because the concept is similar to RAID 5, RAID 6 read operations are not described here

Write cache and auxiliary write cache IOA:

Write cache

The write cache provides greater data integrity and improved performance. When the iSeries(TM) server sends a write operation, the data is written to the cache. Then, a write-completion message is sent back to the server. Later, the data is written to the disk. The cache provides a faster write capability and ensures data integrity.

The following actions occur during a write request from the server:

1. Data is committed to a nonvolatile battery-backed cache in the IOA.
2. A write completion message is sent from the server.
3. The following actions occur after the write completion message is sent.
 - a. A write operation is sent from the IOA cache to the disk unit:
 - A write operation is sent from the IOA cache to the disk unit:
 - Reads the original data.
 - Calculates delta parity by comparing new and original data.

- Writes the new data.
- Write operations for parity data:
 - Reads the original parity information.
 - Calculates the new parity by comparing the delta parity and the original parity.
 - Writes the new parity information.
- b. Data is marked as committed data when it is successfully written to both the data disk unit and the parity disk unit.

The performance for this type of write operation is dependent on disk contention and the time that is needed to calculate the parity information.

Auxiliary cache IOA

The auxiliary cache IOA mirrors the write cache on a storage IOA. Protection is enhanced because two copies of data are stored onto two separate IOAs. If a failure occurs to the write cache, then the auxiliary cache IOA serves as a back up during the recovery of the failed IOA.

When an iSeries(TM) server sends a write operation, the data is written to the write cache on the storage IOA. The storage IOA mirrors the write cache data to the auxiliary write cache IOA. Then, a write-completion message is sent back to the server and the data is then written to a disk.

Note: In order for write cache mirroring to occur, the storage IOAs connected to supported auxiliary write cache IOA. The storage IOA and the auxiliary write cache IOA must also be in the same enclosure and under the same partition.

The auxiliary write cache is an additional IOA that has a one-to-one relationship with a disk IOA. The auxiliary write cache protects against extended outages due to the failure of a disk IOA or its cache by providing a copy of the write cache which can be recovered following the repair of the disk IOA. This avoids a potential system reload and gets the system back on line as soon as the disk IOA is replaced and the recovery procedure completes. Note, however, that this IOA is not a failover device that can keep the system operational if the disk IOA, or it's cache, fails.

For more information about auxiliary write cache see the Planning for IBM i5 Data Protection with Auxiliary Write Cache Solutions redbook.

Mirrored protection

Mirrored protection is a software availability function that protects data from being lost because of failure or because of damage to a disk-related component. Data is protected because the system keeps two copies of data on two separate disk units. When a disk-related component fails, the system may continue to operate without interruption by using the mirrored copy of the data until the failed component is repaired.

- | When you start mirrored protection or add disk units to a disk pool that has mirrored protection, the system creates mirrored pairs using disk units that have similar capacities. The overall goal is to protect as many disk-related components as possible. To provide maximum hardware redundancy and protection, the system attempts to pair disk units that are attached to I/O buses, I/O adapter, I/O processors, buses, and expansion units.

If a disk failure occurs, mirrored protection is intended to prevent data from being lost. Mirrored protection is a software function that uses duplicates of disk-related hardware components to keep your system available if one of the components fails. It can be used on any model of iSeries servers and is a part of the Licensed Internal Code.

Remote mirroring support allows you to have one mirrored unit within a mirrored pair at the local site, and the second mirrored unit at a remote site. For some systems, standard disk unit mirroring will

remain the best choice; for others, remote disk unit mirroring provides important additional capabilities. You must evaluate the uses and needs of your system, consider the advantages and disadvantages of each type of mirroring support, and decide which is best for you.

Related concepts

“Device parity protection” on page 34

“Example: Independent disk pools with geographic mirroring” on page 131

Related information

“Benefits of geographic mirroring” on page 27

“Set the threshold of a disk pool” on page 109

Mirrored protection benefits

With the best possible mirrored protection configuration, the system continues to run after a single disk-related hardware failure. On some system units, the failed hardware can sometimes be repaired or replaced without having to turn off the system. If the failing component is one that cannot be repaired while the system is running, such as a bus or an I/O processor, the system typically continues to run after the failure. Maintenance can be deferred, the system can be shut down normally, and a long recovery time can be avoided.

Even if your system is not a large one, mirrored protection can provide valuable protection. A disk or disk-related hardware failure on an unprotected system leaves your system unusable for several hours. The actual time depends on the kind of failure, the amount of disk storage, your backup strategy, the speed of your tape unit, and the type and amount of processing the system performs. If you or your enterprise cannot tolerate this loss of availability, you should consider mirrored protection for your system, regardless of your system’s size.

Mirrored protection costs and limitations

Costs

The main cost of using mirrored protection is in additional hardware. To achieve high availability and prevent data loss when a disk unit fails, you need mirrored protection for all the disk pools. This normally requires twice as many disk units. If you want continuous operation and prevention of data loss when a disk unit, I/O adapter, or I/O processor fails, you need duplicate I/O adapter and I/O processors. A model upgrade can be done to get nearly continuous operation and to prevent data loss when any of these failures occur, as well as the failure of a bus. If bus 1 fails, the system cannot continue to operate. Because bus failures are rare, and bus-level protection is not significantly greater than I/O processor-level protection, you may not find a model upgrade to be cost-effective for your protection needs.

Mirrored protection has a minimal effect on performance. If the buses, I/O processors, and I/O adapter are no more heavily loaded on a system with mirrored protection than they are on an equivalent system without mirrored protection, then the performance of the two systems should be approximately the same.

In deciding whether to use mirrored protection on your system, you must evaluate the cost of potential downtime against the cost of additional hardware, over the life of the system. The additional cost in performance or system complexity is typically negligible. You should also consider other availability and recovery alternatives, such as device parity protection. Mirrored protection normally requires twice as many storage units. For concurrent maintenance and higher availability on systems with mirrored protection, other disk-related hardware may be required.

Limitations

Although mirrored protection can keep the system available after disk-related hardware failures occur, it is not a replacement for save procedures. There can be multiple types of disk-related hardware failures, or disasters (such as flood or sabotage) that require backup media.

Mirrored protection cannot keep your system available if the remaining storage unit in the mirrored pair fails before the first failing storage unit is repaired and mirrored protection is resumed. If two failed storage units are in different mirrored pairs, the system is still available and normal mirrored protection recovery is done because the mirrored pairs are not dependent on each other for recovery. If a second storage unit of the same mirrored pair fails, the failure may not result in a data loss. If the failure is limited to the disk electronics, or if the service representative can successfully use the save disk unit data function to recover all of the data, no data is lost.

If both storage units in a mirrored pair fail causing data loss, the entire disk pool is lost and all units in the disk pool are cleared. You must be prepared to restore your disk pool from the backup media and apply any journal changes.

When starting the mirrored protection operation, objects that are created on a preferred unit may be moved to another unit. The preferred unit may no longer exist after mirror protection is started.

Mirrored protection and performance

When mirrored protection is started, most systems show little difference in performance; in some cases, mirrored protection can improve performance. Generally, functions that do mostly read operations see equal or better performance with mirrored protection. This is because read operations have a choice of two storage units to read from, and the one with the faster expected response time is selected. Operations that do mostly write operations (such as updating database records) might see slightly reduced performance on a system that has mirrored protection because all changes must be written to both storage units of the mirrored pair. Thus, restore operations are slower.

In some cases, if the system ends abnormally, the system cannot determine whether the last updates were written to both storage units of each mirrored pair. If the system cannot determine whether the last changes were written to both storage units of the mirrored pair, the system synchronizes the mirrored pair by copying the data in question from one storage unit of each mirrored pair to the other storage unit. The synchronization occurs during the IPL that follows the abnormal system end. If the system can save a copy of main storage before it ends, the synchronization process takes just a few minutes. If not, the synchronization process can take much longer. The extreme case might be close to a complete synchronization.

If you have frequent power outages, you might want to consider adding an uninterruptible power supply to your system. Should main power be lost, the uninterruptible power supply allows the system to continue. A basic uninterruptible power supply allows the system time to save a copy of main storage before ending, which avoids long recovery. Both storage units of the load source mirrored pair must be powered by the basic uninterruptible power supply.

How mirrored protection works

Because mirrored protection is configured by disk pool, you can mirror one, some, or all disk pools on the system. By default, every system has a system disk pool. It is not necessary to create user disk pools in order to use mirrored protection. Although mirrored protection is configured by disk pool, all disk pools must be mirrored to provide for maximum system availability. If a disk unit fails in a disk pool that is not mirrored, the system cannot be used until the disk unit is repaired or replaced.

The start mirrored pairing algorithm automatically selects a mirrored configuration that provides the maximum protection at the bus, I/O (input/output) processor, or I/O adapter for the hardware configuration of the system. When storage units of a mirrored pair are on separate buses, they have maximum independence or protection. Because they do not share any resource at the bus, I/O processor, or I/O adapter levels, a failure in one of these hardware components allows the other mirrored unit to continue operating.

Any data that is written to a unit that is mirrored is written to both storage units of the mirrored pair. When data is read from a unit that is mirrored, the read operation can be from either storage unit of the

mirrored pair. It is transparent to the user which mirrored unit the data is being read from. A user is not aware of the existence of two physical copies of the data.

If one storage unit of a mirrored pair fails, the system suspends mirrored protection to the failed mirrored unit. The system continues to operate using the remaining mirrored unit. The failing mirrored unit can be physically repaired or replaced.

After the failed mirrored unit is repaired or replaced, the system synchronizes the mirrored pair by copying current data from the storage unit that has remained operational to the other storage unit. During synchronization, the mirrored unit to which the information is being copied is in the resuming state. Synchronization does not require a dedicated system and runs concurrently with other jobs on the system. System performance is affected during synchronization. When synchronization is complete, the mirrored unit becomes active.


For details on storage on your server, see “Disk storage concepts” on page 2.

Concurrent maintenance: Concurrent maintenance is the process of repairing or replacing a failed disk-related hardware component while the system is being used for normal operations.

On systems without mirrored protection or device parity protection, the system is not available when a disk-related hardware failure occurs and remains unavailable until the failed hardware is repaired or replaced. However, with mirrored protection the failing hardware can often be repaired or replaced while the system is being used.

Concurrent maintenance support is a function of system unit hardware packaging. Mirrored protection only provides concurrent maintenance when the hardware and packaging of the system support it. The best hardware configuration for mirrored protection also provides for the maximum amount of concurrent maintenance.

It is possible for the system to operate successfully through many failures and repair actions. For example, a failure of a disk head assembly does not prevent the system from operating. A replacement of the head assembly and synchronization of the mirrored unit can occur while the system continues to run. The greater your level of protection, the more often concurrent maintenance can be performed.

On some models, the system restricts the level of protection for unit 1 and its mirrored unit to only IOA-level protection. See “Mirrored Protection - Configuration Rules” in Backup and Recovery  for more information.

Under some conditions, diagnosis and repair can require active mirrored units to be suspended. You may prefer to turn off the system to minimize the exposure of operating with less mirrored protection. Some repair actions require that the system be powered down. Deferred maintenance is the process of waiting to repair or replace a failed disk-related hardware component until the system can be powered down. The system is available, although mirrored protection is reduced by whatever hardware components have failed. Deferred maintenance is only possible with mirrored protection or device parity protection.

Remote disk unit mirroring support

Standard disk unit mirroring support requires that both disk units of the load source mirrored pair (unit 1) are attached to load source IOP. This allows the system to IPL from either load source in the mirrored pair and allows the system to dump main storage to either load source if the system ends abnormally. However, since both load sources must be attached to the same IOP, the best mirroring protection possible for the load source mirrored pair is IOA-level protection. To provide a higher level of protection for your system, you can use remote load source mirroring and remote disk unit mirroring.

Remote disk unit mirroring support, when combined with remote load source mirroring, mirrors the disk unit on local optical buses with the disk unit on optical buses that terminates at a remote location. In this

configuration, the entire system, including the load source, can be protected from a site disaster. If the remote site is lost, the system can continue to run on the disk unit at the local site. If the local disk unit and system unit are lost, a new system unit can be attached to the set of disk units at the remote site, and system processing can be resumed.

Remote disk unit mirroring, like standard disk unit mirroring, supports mixing device-parity-protected disk units in the same disk pool with mirrored disk units; the device parity disk unit can be located at either the local or the remote site. However, if a site disaster occurs at the site containing the device parity disk unit, all data in the disk pools containing the device parity disk unit is lost.

Remote mirroring support makes it possible to divide the disk units on your system into a group of local disk units and a group of remote disk units. The remote disk units are attached to one set of optical buses and the local disk units to another set of buses. The local and remote disk units can be physically separated from one another at different sites by extending the appropriate optical buses to the remote site. The distance between the sites is restricted by the distance that an optical bus may be extended.

If you decide that remote disk unit mirroring is needed for your system, you need to “Prepare your system for remote mirroring” on page 69 and then “Start site-to-site mirroring” on page 106.

Remote load source mirroring: Remote load source mirroring support allows the two disk units of the load source to be on different IOPs or system buses, which provides IOP or bus-level mirrored protection for the load source. However, in such a configuration, the system can only restart from or perform a main storage dump to the load source attached to the Load Source IOP. If the load source on the Load Source IOP fails, the system can continue to run on the other disk unit of the load-source mirrored pair.

Enable remote load source mirroring: Enabling remote load source mirroring makes it possible for the two disk units of the load source mirrored pair to be on different I/O processors or system buses. Remote load source mirroring allows you to protect against a site disaster by dividing the disk storage between the two sites, mirroring one site to another. You must enable remote load source mirroring before starting mirrored protection for disk pool 1. If remote load source mirroring support is enabled after mirrored protection has already been started for disk pool 1, the existing mirrored protection and mirrored pairing of the load source is not changed.

Remote load source mirroring support can be enabled in either the DST or the SST environment in iSeries Navigator or the character-based interface. If you attempt to enable remote load source mirroring and it is currently enabled, the system displays a message that remote load source mirroring is already enabled.

To enable remote load source mirroring, follow these steps:

1. In iSeries Navigator, expand **Disk Units** → **Disk Pools** → **Disk Pool 1**.
2. Right-click the load source disk unit and select **Enable Remote Load Source Mirroring**.

Note: Enabling remote load source mirroring does not start mirrored protection on the disk units. Remote load source mirroring affects only the load source disk units.

To enable remote load source mirroring using the character-based interface, do the following:

1. From the DST Main Menu, select option 4, Work with disk units.
2. From the Work with disk units menu, select option 1, Work with disk configuration.
3. From the Work with disk configuration menu, select option 4, Work with mirrored protection.
4. From the Work with mirrored protection menu, select option 4, Enable remote load source mirroring. This will display an Enable remote load source mirroring confirmation screen.
5. Press Enter at the Enable remote load source mirroring confirmation screen. The Work with mirrored protection screen will be displayed, with a message at the bottom, indicating that remote load source mirroring has been enabled.

Disable remote load-source mirroring: If you want to disable remote load-source mirroring support, you must do one of the following:

- Stop mirrored protection and then disable remote load-source mirroring support. Mirrored protection is local mirroring unlike cross-site mirroring or geographic mirroring.
- Move the remote load source to the Load Source IOP and then disable remote load-source mirroring support.

If the remote load source is moved to the Load Source IOP, the IOP and system might not recognize it because of the different disk unit format sizes that are used by different IOPs. If the remote load source is missing after it has been moved to the Load Source IOP, use the DST replace disk unit function to replace the missing load source with itself. This causes the disk unit to be reformatted so that the Load Source IOP can use it, and then the disk unit is synchronized with the active load source.

Remote load source mirroring may be disabled from either DST or SST. However, disabling remote load source mirroring is not allowed if there is a load source disk unit on the system that is not attached to the Load Source IOP. If you attempt to disable remote load source mirroring support and it is currently disabled, the system will display a message that remote load source mirroring is already disabled.

To disable remote load source mirroring support, do the following:

1. From the DST main menu, select option 4, Work with disk units.
2. From the Work with disk units menu, select option 1, Work with disk configuration
3. From the Work with disk configuration menu, select option 4, Work with mirrored protection
4. From the Work with mirrored protection menu, Select option 5, Disable remote load source mirroring. This will display a Disable remote load source mirroring confirmation screen.
5. Press Enter at the Disable remote load source mirroring confirmation screen. The Work with mirrored protection screen will be displayed, with a message at the bottom, indicating that remote load source mirroring has been disabled.

Remote load source mirroring used with local disk unit: Remote load source mirroring can be used to achieve IOP-level or bus-level protection of the load source mirrored pair, even without remote disk units or buses on the system. There is no special setup required, other than to ensure that a disk unit of the same capacity as the load source is attached to another IOP or bus on the system. If you want to achieve bus-level protection of all mirrored pairs in a disk pool, you should configure your system so that no more than one half of the disk units of any given capacity in that disk pool are attached to any single bus. If you want to achieve IOP-level protection of all mirrored pairs in a disk pool, you must have no more than one half of the disk units of any given capacity in the disk pool attached to any single IOP.

After the system hardware is configured correctly, enable remote load source mirroring and start mirroring for the disk pools you want to protect. Use the normal start mirroring function. There is no special start mirroring function for remote load source support. The system detects that remote load source mirroring is enabled and automatically pairs disk units to provide the best level of protection possible. It is not possible to override or influence the pairing of the disk units other than by changing the way the hardware of the system is connected and configured. Typical mirroring restrictions that concern total disk pool capacity, an even number of disk units of each capacity, and such things, apply.

Advantages of remote disk unit mirroring:

- Remote disk unit mirroring can provide IOP-level or bus-level mirrored protection for the load source.
- Remote disk unit mirroring allows the disk unit to be divided between two sites, mirroring one site to another, to protect against a site disaster.

Disadvantages of remote loadsource disk unit mirroring for primary partitions on an IBM iSeries server:

- A system that uses remote loadsource disk unit mirroring is only able to perform an IPL from the disk unit attached to the load source IOP. If that disk unit fails and cannot be repaired concurrently, the system cannot perform an IPL until the failed load source is fixed and the recover remote loadsource service procedure is performed.
- When remote loadsource disk unit mirroring is active on a system and the load source attached to the Load Source IOP fails, the system cannot perform a main storage dump if the system ends abnormally. This means that the system cannot use the main storage dump to reduce recovery time after a system fails. It also means that the main storage dump is not available to diagnose the problem that caused the system to end abnormally.


Comparison of standard mirroring and remote mirroring: For the most part, the way you manage disk units with remote mirroring is the same as how you manage disk units with standard mirroring. When you add disk units, unprotected disk units must be added in pairs, as with general mirroring. To achieve remote protection of all added units, half of the new units of each capacity of disk unit should be in the remote group and half in the local group. Single device-parity protected units may be added to disk pools using remote mirroring. However, the disk pool is not protected against a site disaster.

You will also see some differences when you Restore remote mirrored protection after a recovery.

Restore remote mirrored protection after a recovery


To restore mirrored protection following the recovery procedures, you need to perform the following steps:

- Obtain and physically attach all required disk units.
- Stop or suspend mirrored protection if it is currently configured on the system.
- Add the new disk units to the correct disk pools.
- Resume mirrored protection

For detailed information about how to recover systems with mirrored protection, see Backup and Recovery .

External load source unit

The load source unit is the storage unit that contains the initial programs and data that are used during an initial program load (IPL) of the system. The load source unit is typically configured as an internal storage unit, but can also be configured as an external storage unit located on a storage area network (SAN).

For information on using an external load source unit that is located on a SAN, see the IBM Redbook iSeries and TotalStorage® , SG24-7120.

Plan for disk management

Depending on how you plan to manage your disks, you must meet certain hardware, software, and communications requirements.

This information will help you manage your disks.

iSeries Navigator requirements for disk management

Changing the disk configuration of your server 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.

Access 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:

Install the Configuration and Service component

1. From the **File** menu of iSeries Navigator, select **Install Options**, and then click **Selective Setup**.
2. Follow the instructions on the resulting dialog box to install the Configuration and Service component.

Enable the Disk Units folder

1. In iSeries Navigator, right-click the server connection and select **Application Administration**.
2. On the resulting window, click **OK**.
3. Click the **Host Applications** tab.
4. Expand your operating system.
5. Select **Disk Units** to have **Default Access** or **All Object Access**.
6. Click **OK**.
7. Restart iSeries Navigator.

Access the Disk Units folder to perform all disk management functions

1. In iSeries Navigator, expand **My Connections** (or your active environment).
2. Expand any iSeries server>**Configuration and Service**>**Hardware**>**Disk Units**.

Set up communication

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

Configuring the service tools server

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

Setting the service IP address

To access DST functions on your server from iSeries Navigator, you need to specify a service IP address for the server. 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 System (DNS) that resolves to an address as previously described. Contact your network administrator for this information. Make sure you have configured the service tools server before you continue with these instructions.

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


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

Plan 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.

Important: When you are ready to order a new server or a server 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.

Hardware requirements

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

Environments	independent disk pool use	Requirements
Single system	Stand-alone independent disk pool	One iSeries server running OS/400 V5R1M0 ¹ or later.
Multiple systems		Use either multiple servers or multiple partitions, by doing one of the following: <ul style="list-style-type: none"> • Two or more iSeries servers. • One iSeries server 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 input/output processors (IOP) on a shared bus or an IOP that is assigned to an I/O pool².
	Geographic mirroring ³	Two or more servers 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.

Note:

1. OS/400 V5R1M0 provides only independent disk pools containing user-defined file system (UDFS) only. OS/400 V5R2M0 or later support library-based objects.

2. In an LPAR environment, you can switch the input/output processor (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 input/output adapters (IOAs) on the IOP will be switched.
3. OS/400 V5R3M0 provides support for geographic mirroring.

Physical planning requirements

Depending on how you plan to use independent disk pools, you must satisfy the following 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 servers 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 servers (cluster nodes) on each HSL loop, though each server 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 servers, two segments exist, separated by the two servers. 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 server 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 server that has not become switchable.
|




Figure 5. These expansion units are all private and are not switchable.




Figure 6. The expansion unit farthest away from the owning server on the loop segment has been made 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 the following software and licenses:

Multiple-system clustered environment

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

1. 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.

2. iSeries Navigator is the graphical user interface for managing and administering your iSeries server from your Windows[®] desktop. It is required to perform some of the disk management tasks necessary to use independent disk pools. See “iSeries Navigator requirements for disk management” on page 49 for steps to enable iSeries Navigator for disk management.
3. 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 servers, the servers 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.

Single-system environment

1. You need OS/400 V5R1M0¹ or later.
2. iSeries Navigator is the graphical user interface for managing and administering your iSeries server 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

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

Single-system environment

There are no communications requirements.

Multisystem clustered environment

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

- **For switchable independent disk pools**, at least one TCP/IP communications interface between the servers in the cluster. For redundancy, it is recommended that there be at least two separate interfaces between the servers.
- **For geographic mirroring**, the following are 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 server 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

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 an iSeries 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 server 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 be aware of. 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” on page 20 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” on page 18. 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” on page 21 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 server, 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 server, 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.

Plan for disk protection

Plan which methods you need to use to protect your data.

Comparison of disk protection options

You should be aware of these considerations when selecting disk protection options:

When using RAID 5 device parity protection, the system continues to run after a single-disk failure. When using RAID 6 device parity protection, the system continues to run after a two-disk failure. With mirrored protection, the system might continue to run after the failure of a disk-related component, such as an IOA or an IOP.

RAID 5 device parity protection requires the capacity of one disk unit that is dedicated to storing parity data in a parity set. RAID 6 device parity protection requires the capacity of two disk units that are dedicated to storing parity data in a parity set. A system with mirrored protection requires twice as much disk capacity as the same system without mirrored protection because all information is stored twice. Mirrored protection might also require more buses, IOPs, and disk IOA, depending on the level of protection that you want. Therefore, mirrored protection is typically more expensive than device parity protection.

Typically, neither device parity protection nor mirrored protection has a noticeable effect on system performance. In some cases, mirrored protection actually improves system performance. The time required to restore data to disk units that are protected by device parity protection is longer than the time required to restore to the same disk devices that do not have device parity protection activated. This is because the parity data must be calculated and written.

This table provides an overview of the availability tools that can be used on the server to protect against different types of failure.

What type of availability is needed?	Device parity protection	Mirrored protection	Basic disk pools	Independent disk pool
Protect from data loss due to disk-related hardware failure	Yes	Yes	See note 2	See note 2
Maintain availability	Yes	Yes	No	Yes ⁴
Help with disk unit recovery	Yes	Yes	Yes ²	Yes ²

What type of availability is needed?	Device parity protection	Mirrored protection	Basic disk pools	Independent disk pool
Maintain availability when input/output adapter (IOA) fails	No	Yes ¹	No	No ⁵
Maintain availability when disk I/O processor fails	No	Yes ¹	No	No ⁵
Maintain availability when system bus fails	No	Yes ¹	No	No ⁵
Site disaster protection	No	Yes ³	No	No ⁵
Ability to switch data between systems	No	No	No	Yes

Notes:

1. Depends on hardware used, configuration, and level of mirrored protection.
2. Configuring disk pools can limit the loss of data and the recovery to a single disk pool.
3. For site disaster protection, remote mirroring is required.
4. In a clustered environment, an independent disk pool can help maintain availability.
5. When using geographic mirroring, independent disk pools can provide site disaster protection.

Plan for device parity protection RAID 5

Systems with IOAs released after V5R2 hold a minimum number of 3 disk units in a parity set; the maximum number of disk units in a parity set is 18.

Note: Systems with IOAs released prior to V5R2 of OS/400, the minimum number of disk units in a parity set is 4. The maximum number of disk units in a parity set is 10.

RAID 6

The minimum number of disk units in a parity set is 4. The maximum number of disk units in a parity set is 18.

To learn more about how device parity protection is implemented, see “How device parity protection works” on page 36. “Examples: Device parity and mirrored protection” on page 70 shows some examples of how device parity protection can be used in conjunction with mirrored protection.

Note: If your goal is to have a system with data loss protection and concurrent maintenance repair, a combination of mirrored protection and device parity protection is recommended.

Plan for mirrored protection

If you have a multibus system or a large single-bus system, you should consider using mirrored protection. The greater the number of disk units that are attached to a system, the more frequent disk-related hardware failures are because there are more individual pieces of hardware that can fail. Therefore, the possibility of data loss or loss of availability as a result of a disk or other hardware failure becomes more likely. Also, as the amount of disk storage on a system increases, the recovery time after a disk storage subsystem hardware failure increases significantly. Downtime becomes more frequent, more lengthy, and more costly.

Decide which disk pools to protect: Mirrored protection is configured by disk pool because it is the user’s level of control over single-level storage. Mirrored protection can be used to protect one, some, or all disk pools on a system. However, multiple disk pools are not required to use mirrored protection. Mirrored protection works well if all disk units on a system are configured into a single disk pool (the default on the iSeries server). In fact, mirroring reduces the need to partition auxiliary storage into disk pools for data protection and recovery. However, disk pools might still be desirable for performance and other reasons.

To provide the best protection and availability for the entire system, all disk pools in the system should have mirrored protection:

- If the system has a mixture of some disk pools with and some disk pools without mirrored protection, a disk unit failure in a disk pool without mirrored protection severely limits the operation of the entire system. Data can be lost in the disk pool in which the failure occurred. A long recovery might be required.
- If a disk fails in a mirrored disk pool, and the system also contains disk pools that are not mirrored, data is not lost. However, in some cases, "Concurrent maintenance" on page 46 might not be possible.

The disk units that are used in disk pools should be selected carefully. For best protection and performance, a disk pool should contain disk units that are attached to several different I/O processors. The number of disk units in the disk pool that are attached to each I/O processor should be the same (that is, balanced).

Determine the disk units that are needed: A mirrored disk pool requires twice as much storage as a disk pool that is not mirrored because the system keeps two copies of all the data in the disk pool. Also, mirrored protection requires an even number of disk units of the same capacity so that disk units can be made into mirrored pairs. On an existing system, it is not necessary to add the same types of disk units already attached to provide the required additional storage capacity. Any new disk units may be added as long as sufficient total storage capacity and an even number of storage units of each size are present. The system assigns mirrored pairs and automatically move the data as necessary. If a disk pool does not contain sufficient storage capacity, or if storage units cannot be paired, mirror protection cannot be started for that disk pool.

The process of determining the disk units that are needed for mirrored protection is similar for existing or new systems. You and your IBM marketing representative should do the following:

1. Plan for storage capacity.
2. Plan a target percent of storage used for the disk pool (how full the disk pool can be).
3. Plan the number and type of disk units needed to provide the storage that is required. For an existing disk pool, you can plan a different type and model of disk unit to provide the required storage. You must ensure an even number of each type of disk unit and model.
4. Plan for disk pools.
5. Plan total storage capacity.

Plan for storage capacity: For a new system, your IBM marketing representative can help you analyze your system storage requirements. For an existing system, the current amount of data in the disk pool that is being planned is a useful starting point. The Display Disk Configuration Capacity option of the dedicated service tools (DST) or system service tools (SST) shows the total size (in millions of bytes) and the percent of storage used for each disk pool on the system. Multiply the size of the disk pools by the percent that is used to calculate the number of megabytes of data currently in the disk pool. In planning future storage requirements for a disk pool, system growth and performance should also be considered.

The planned amount of data and the planned percent of storage used work together to determine the amount of actual auxiliary storage needed for a mirrored disk pool. For example, if a disk pool is to contain 1 GB (GB equals 1 073 741 824 bytes) of actual data, it requires 2 GB of storage for the mirrored copies of the data. If 50% of storage use is planned for that disk pool, the disk pool needs 4 GB of actual storage. If the planned percent of storage that is used is 66%, 3 GB of actual storage are required. One gigabyte of real data (2 GB of mirrored data) in a 5 GB disk pool results in a 40% auxiliary storage utilization.

Plan for spare disk units: Spare disk units can reduce the time the system runs without mirrored protection for a mirrored pair after a disk unit failure. If a disk unit fails and a spare disk unit of similar capacity is available, that spare disk unit can be used to replace the failed disk unit. Using the DST or SST replace option, the user selects the failed disk unit to replace, then selects a spare disk unit to replace

it. The system logically replaces the failed disk unit with the selected spare disk unit, then synchronizes the new disk unit with the remaining good disk unit of the mirrored pair. Mirrored protection for that pair is again active when synchronization is completed (typically, in less than an hour). However, it might take several hours from the time that a service representative is called until the failed disk unit is repaired and synchronized, and mirrored protection is again active for that pair.

Plan for total storage capacity: After planning for the number and type of storage units needed for each disk pool on the system, and for any spare storage units, add up the total number of storage units of each disk unit type and model. Remember that the number planned is the number of storage units of each disk unit type, not the number of disk units. You and your IBM marketing representative need to convert the planned number of storage units to disk units before ordering hardware.

The preceding information helps you plan the total number of disk units needed for your system. If you are planning for a new system, this is the number that needs to be ordered. If you are planning for an existing system, subtract the number of each disk type currently on your system from the number that is planned. This is the number of new disk units that should be ordered.

Determine the level of protection that you want: The level of mirrored protection determines whether the system keeps running when different levels of hardware fail. The level of protection is the amount of duplicate disk-related hardware that you have. The more mirrored pairs that have higher levels of protection, the more often your system will be usable when disk related hardware fails. You may decide that a lower level of protection is more cost effective for your system than a higher level.

When determining what level of protection is adequate, you should consider the relative advantages of each level of protection with respect to the following:

- The ability to keep the system operational during a disk-related hardware failure.
- The ability to perform maintenance concurrently with system operations. To minimize the time that a mirrored pair is unprotected after a failure, you may want to repair failed hardware while the system is operating.

Details: Levels of protection: The level of mirrored protection determines whether the system keeps running when different levels of hardware fail. Mirrored protection always provides disk-unit level protection, which keeps the system available for a single disk-unit failure. To keep the system available for failures of other disk-related hardware requires higher levels of protection. For example, to keep the system available when an I/O processor (IOP) fails, all of the disk units attached to the failing IOP must have mirrored units attached to different IOPs.

The level of mirrored protection also determines whether concurrent maintenance can be done for different types of failures. Certain types of failures require concurrent maintenance to diagnose hardware levels above the failing hardware component. For example, to diagnose a power failure in a disk unit, you need to reset the I/O processor to which the failed disk unit is attached. Therefore, IOP-level protection is required. The higher the level of mirrored protection, the more often concurrent maintenance is possible.


The level of protection you get depends on the hardware you duplicate. If you duplicate disk units, you have disk-unit level protection. If you duplicate IOAs as well, you have IOA-level protection. If you duplicate input/output processors, you have IOP-level protection. If you duplicate buses, you have bus-level protection. Mirrored units always have at least disk-unit level protection.

During the start mirrored protection operation, the system pairs the disk units to provide the maximum level of protection for the system. When disk units are added to a mirrored disk pool, the system pairs only those disk units that are added without rearranging the existing pairs. The hardware configuration includes both the hardware and how the hardware is connected.

Disk-unit level protection:

Mirrored protection always provides disk-unit level protection because the storage units are duplicated. If your main concern is protection of data and not high availability, then disk-unit level protection might be adequate. The disk unit is the most likely hardware component to fail, and disk-unit level protection keeps your system available after a disk unit failure.

Concurrent maintenance is often possible for certain types of disk unit failures that have disk-unit level protection.




This figure shows the elements of disk unit level protection: one bus, connected to one I/O processor, connected to one I/O adapter, which is attached to two separate disk units. The two storage units make a mirrored pair. With disk-unit level protection, the system continues to operate after a disk unit failure. If the I/O adapter or I/O processor fails, the system cannot access data on either of the storage units of the mirrored pair, and the system is unusable.

| *Input/output bus level protection:*

| Determine whether you want I/O bus level protection based on the following:

- | • To keep your system available when an I/O bus fails.
- | • To concurrently repair a failed disk unit.

| To achieve I/O bus protection, all disk units must have a mirrored unit attached to a different I/O bus. This figure shows I/O bus protection. The two storage units make a mirrored pair. With I/O bus protection, the system can continue to operate if one I/O bus fails. If the IOA or the IOP fails, the system cannot access data on either of the disk units, and the system is unusable.




| This figure shows the elements of I/O bus protection: one bus, connected to one I/O processor, connected to one IOA, with two or more I/O buses, each attached to a separate disk unit.

Input/output adapter level protection:

Determine whether you want input/output adapter (IOA) level protection based on the following:

- To keep your system available when an IOA fails.
- To concurrently repair a failed disk unit or IOA. To use problem recovery procedures in preparation for isolating a failing item or to verify a repair action, the IOA must be dedicated to the repair action. If any disk units that are attached to the IOA do not have IOA-level protection, then this part of concurrent maintenance is not possible.

To achieve IOA-level protection, all disk units must have a mirrored unit attached to a different IOA. This figure shows IOA-level protection. The two storage units make a mirrored pair. With IOA-level protection, the system can continue to operate if one IOA fails. If the I/O processor fails, the system cannot access data on either of the disk units, and the system is unusable.




This figure shows the elements of IOA-level protection: one bus, connected to one I/O processor, connected to two IOAs, which are each attached to a separate disk unit.

Input/output processor level protection:

Determine if you want IOP-level protection based on the following:

- To keep your system available when an I/O processor fails.
- To keep your system available when the cable attached to the I/O processor fails.
- To concurrently repair certain types of disk unit failures or cable failures. For these failures, concurrent maintenance needs to reset the IOP. If any disk units that are attached to the IOP do not have IOP-level protection, then concurrent maintenance is not possible.

To achieve IOP-level protection, all disk units that are attached to an I/O processor must have a mirrored unit attached to a different I/O processor. On many systems, IOP-level protection is not possible for the mirrored pair for unit 1.




This figure shows the elements of IOP-level protection: one bus, attached to two IOPs, which are each connected to a separate IOA and a separate disk unit. The two storage units make a mirrored pair. With IOP-level protection, the system can continue to operate if one I/O processor fails. The system becomes unusable only if the bus fails.

Bus-level protection:

Determine if you want Bus-level protection based on the following:

- Bus failures are rare in comparison to other disk-related hardware failures.
- The system can continue to operate after a bus failure.
- The system cannot operate if bus 1 fails.
- If a bus fails, disk I/O operations may continue, but other hardware is lost (such as work stations, printers and communication lines) making the system unusable.
- Concurrent maintenance is not possible for bus failures.

To achieve bus-level protection, all disk units that are attached to a bus must have a mirrored unit attached to a different bus. Bus-level protection is not possible for unit 1.




This figure shows the elements of bus-level protection: one expansion unit that contains two buses attached to separate IOPs, IOAs, and disk units, respectively. The two storage units make a mirrored pair.

Expansion-unit level protection:

Determine if you want Expansion-unit level protection on the following:

- Expansion unit failures are rare compared with other disk-related hardware failures.
- If an expansion unit fails, disk I/O operations may continue, but other hardware is lost (such as work stations, printers and communication lines) making the system unusable.




To achieve expansion-unit level protection, all disk units that are present in the expansion unit must have a mirrored unit present in another expansion unit. The figure shows the elements of expansion-unit level protection: two expansion units that each contain two buses that are attached to separate IOPs, IOAs, and disk units, respectively.

Ring-level protection:

Determine if you want Ring-level protection based on the following:

- HSL failures are rare compared with other disk-related hardware failures.
- If an HSL fails, disk I/O operations might continue, but other hardware is lost (such as work stations, printers and communication lines) making the system unusable.



To achieve ring-level protection, all disk units that are present in the expansion unit in the first HSL must also have a mirrored units present in another expansion unit in the second HSL. The figure shows the elements of ring-level protection: two HSL rings, connected to two expansion units that each contain two buses that are attached to separate IOPs, IOAs, and disk units, respectively.

Determine the hardware that is needed for mirroring: In order to communicate with the rest of the system, disk units are attached to I/O adapters, which are attached to I/O processors, which are attached to buses. The number of each of these types of disk-related hardware that are available on the system directly affects the level of protection that is possible.

To provide the best protection and performance, each level of hardware should be balanced under the next level of hardware. That is, the disk units of each device type and model should be evenly distributed under their I/O adapters. The same number of I/O adapters should be under each I/O processor for that disk type. The I/O processors should be balanced among the available buses.

To plan what disk-related hardware is needed for your mirrored system, you must plan the total number and type of disk units (old and new), that are needed on the system, as well as the level of protection for the system. It is not always possible to plan for and configure a system so that all mirrored pairs meet the planned level of protection. However, it is possible to plan a configuration in which a large percentage of the disk units on the system achieve the required level of protection.

Plan the minimum hardware needed to function: Various rules and limits exist on how storage hardware can be attached together. The limits may be determined by hardware design, architecture restrictions, performance considerations, or support concerns. Your IBM marketing representative can explain these configuration limits and help you use them in your planning.

For each disk unit type, first plan for the I/O adapter that are needed and then for the I/O processors that are needed. After planning the number of I/O processors that are needed for all disk unit types, use the total number of I/O processors to plan for the number of buses that are needed.

Related concepts

Installation, upgrades, and migration

Plan additional hardware to achieve the level of protection: Consider the following information to ensure adequate protection.

- “Disk-unit level protection” on page 61

If you have planned for disk-unit level protection, you do not need to do anything more. All mirrored disk pools have disk-unit level protection if they meet the requirements for starting mirrored protection.

- “Input/output bus level protection” on page 62

If you have planned for I/O bus level protection, you do not need to do anything more. All mirrored disk pools have I/O bus level protection if they meet the requirements for starting mirrored protection.

- “Input/output adapter level protection” on page 63

If your planned disk units do require an IOA, add as many IOAs as possible, keeping within the defined system limits. Then balance the disk units among them according to the standard system configuration rules.

- “Input/output processor level protection” on page 64

If you want IOP-level protection and do not already have the maximum number of IOPs on your system, add as many IOPs as possible, keeping within the defined system limits. Then balance the disk units among them according to the standard system configuration rules. You might need to add additional buses to attach more IOPs.

- “Bus-level protection” on page 65

If you want bus-level protection and already have a multiple-bus system, you do not need to do anything. If your system is configured according to standard configuration rules, the mirrored pairing function pairs storage units to provide bus-level protection for as many mirrored pairs as possible. If you have a single-bus system, you can add additional buses as a feature option.

- “Expansion-unit level protection” on page 66

If your system is configured with an equal number of equal capacity disk units between expansion units, the mirrored pairing function pairs the disk units in different expansion units to provide expansion-unit level protection on as many disk units as possible.

- “Ring-level protection” on page 67

If your system is configured with an equal number of equal capacity disk units between high-speed links (HSL), the mirrored pairing function pairs the disk units in different high-speed link (HSL) configurations to provide ring-level protection on as many disk units as possible.

Determine the extra hardware needed for performance: Mirrored protection normally requires additional disk units and input/output processors. However, in some cases, you might need additional hardware to achieve the level of performance that you want.

Use the following information to decide how much extra hardware you might need:

Processing unit requirements

Mirrored protection causes a minor increase in central processing unit usage (approximately 1% to 2%).

Main storage requirements

If you have mirrored protection, you need to increase the size of your machine pool. Mirrored protection requires storage in the machine pool for general purposes and for each mirrored pair. You should expect to increase your machine pool by approximately 12 KB for each 1 GB of mirrored disk storage (12 KB for 1 GB disk units, 24 KB for 2 GB disk units, and so forth).

During synchronization, mirrored protection uses an additional 512 KB of memory for each mirrored pair that is being synchronized. The system uses the pool with the most storage.

I/O processor requirements

To maintain equivalent performance after starting mirrored protection, your system should have the same ratio of disk units to I/O processors as it did before. To add I/O processors, you might need to upgrade your system for additional buses.

Because of the limit on buses and I/O processors, you might not be able to maintain the same ratio of disk units to I/O processors. In this case, system performance might be less.

Order and install the new hardware: Your IBM marketing representative will assist you in ordering your new hardware by using the typical order process. That ordering process allows for any other hardware that might be needed as part of your upgrade, such as additional racks and cables.

When your order arrives, see Install iSeries features for installation instructions.

Prepare your system for remote mirroring

When you start remote system mirroring, the local disk unit is mirrored to the remote disk unit. If a site disaster occurs at either the local or remote location, a complete copy of all data on the system still exists, the system configuration can be recovered, and processing can continue. To provide protection against a site disaster, all disk units in all disk pools of the system must be mirrored in local-remote pairs. Follow these steps to prepare your system for remote mirroring:

1. Plan which optical buses will control disk units at the remote site.
 - It is not functionally necessary that the local site and the remote site use the same number of buses. However, it is simplest to configure and understand the system if the number of remote and local buses and disk units are equal.
 - It is functionally necessary that both the local and remote sites have the same number of each capacity of disk units in each disk pool.
2. Plan the distribution of disk units, move disk units if necessary, and verify that half of each capacity of disk units in each disk pool is attached to the local and remote set of buses.
3. Indicate to the system which buses control remote disk units and which buses control local disk units.

Find remote buses: If the buses are not labeled, you might need to manually trace the buses to see which connect to remote locations. You can also use the Hardware Service Manager to determine which buses go to which expansion units. The packaging resources that are associated with a Logical Resource panel displays the frame ID and resource name of the expansion unit that is associated with the bus.

To use the Hardware Service Manager to find the buses that control remote disk units, perform these steps:

1. From the DST Main Menu, select **Start a service tool**.
2. From the Start a Service Tool display, select **Hardware service manager**.

3. From the Hardware Service Manager menu, select **Logical hardware resources**.
4. From the Logical Hardware Resources menu, select **System bus resources**.
5. On the Logical Hardware Resource on System Bus panel, enter option 8 before each bus to display the associated packaging resources.
6. If you need more information to help you find and distinguish the expansion unit in question, enter option 5 for the System expansion unit to display other details about the expansion unit.
7. Record the remote or local location of the bus.
8. Then repeat this procedure for all buses on the system.

Change remote bus resource names: After you determine which buses control remote disk units, use Hardware Service Manager to change the resource names of the remote buses.

To change the resource names of the remote buses, perform these steps:

1. From the DST Main Menu, select **Start a service tool**.
2. From the Start a Service Tool display, select **Hardware service manager**.
3. From the Hardware Service Manager menu, select **Logical hardware resources**.
4. From the Logical Hardware Resources menu, selection **System bus resources**.
5. On the Logical Hardware Resource on System Bus panel, type 2 by the bus whose name you want to change. This displays the Change Logical Hardware Resource Detail panel.
6. On the Change Logical Hardware Resource Detail Panel, on the line labeled New resource name, change the resource name by adding the letter *R* to the beginning of the resource name of the bus. For example, change *BUS08* to *RBUS08*. Press Enter to change the resource name.
7. Repeat this procedure for each remote bus on the system.


Examples: Device parity and mirrored protection

The following examples show different options for protecting disk pools. They include:

- Mirrored protection and device parity protection to protect the system disk pool
- Mirrored protection in the system disk pool and device parity protection in the user disk pools
- Mirrored protection and device parity protection in all disk pools

Mirrored protection and device parity protection to protect the system disk pool

Here is an example of a system with a single disk pool (auxiliary storage pool) with both mirrored protection and device parity protection.



The figure shows a single disk pool with twelve disk units. Disk units 9-12 all have the same capacity and are protected by device parity protection. Disk units 1-8 have varying capacities, but each disk unit can be paired with another disk unit of the same capacity when mirrored protection is started. After mirrored protection is started, the disk units that have been paired together are both identified by the same number. Disk units 1 and 2 are now both named 1, and so forth. The system will fail if more than one disk unit fails for RAID 5 or more than two disk units fail for RAID 6. The failed unit can be repaired concurrently. If one of the mirrored disk units fails, the system continues to run using the operational unit of the mirrored pair.

Mirrored protection in the system disk pool and device parity protection in the user disk pools

Consider device parity protection if you have mirrored protection in the system disk pool and you are going to create basic or independent disk pools. RAID 5 allows for the system to tolerate a failure in one of the disk units in a basic or independent disk pool. RAID 6 allows for the system to tolerate a failure in two disk units. The failure can be repaired while the system continues to run.

Mirrored protection and device parity protection in all disk pools

If you have all disk pools (also known as auxiliary storage pools) protected with mirrored protection and you want to add units to the existing disk pools, consider using device parity protection as well. RAID 5, the system can tolerate a failure in one of the disk units with device parity protection. RAID 6, the system can tolerate two disk unit failures. The failed units can be repaired while the system continues to run. If a failure occurs on a disk unit that has mirrored protection, the system continues to run using the operational unit of the mirrored pair.

Configure your disks

Evaluate and configure your disks.

Evaluate the current configuration

Before you change the disk configuration of your server, it is important to know exactly where the existing disk units are located in relation to disk pools, I/O adapters, and frames. The graphical view of iSeries Navigator eliminates the process of compiling all this information by providing a graphical representation of how your server 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.

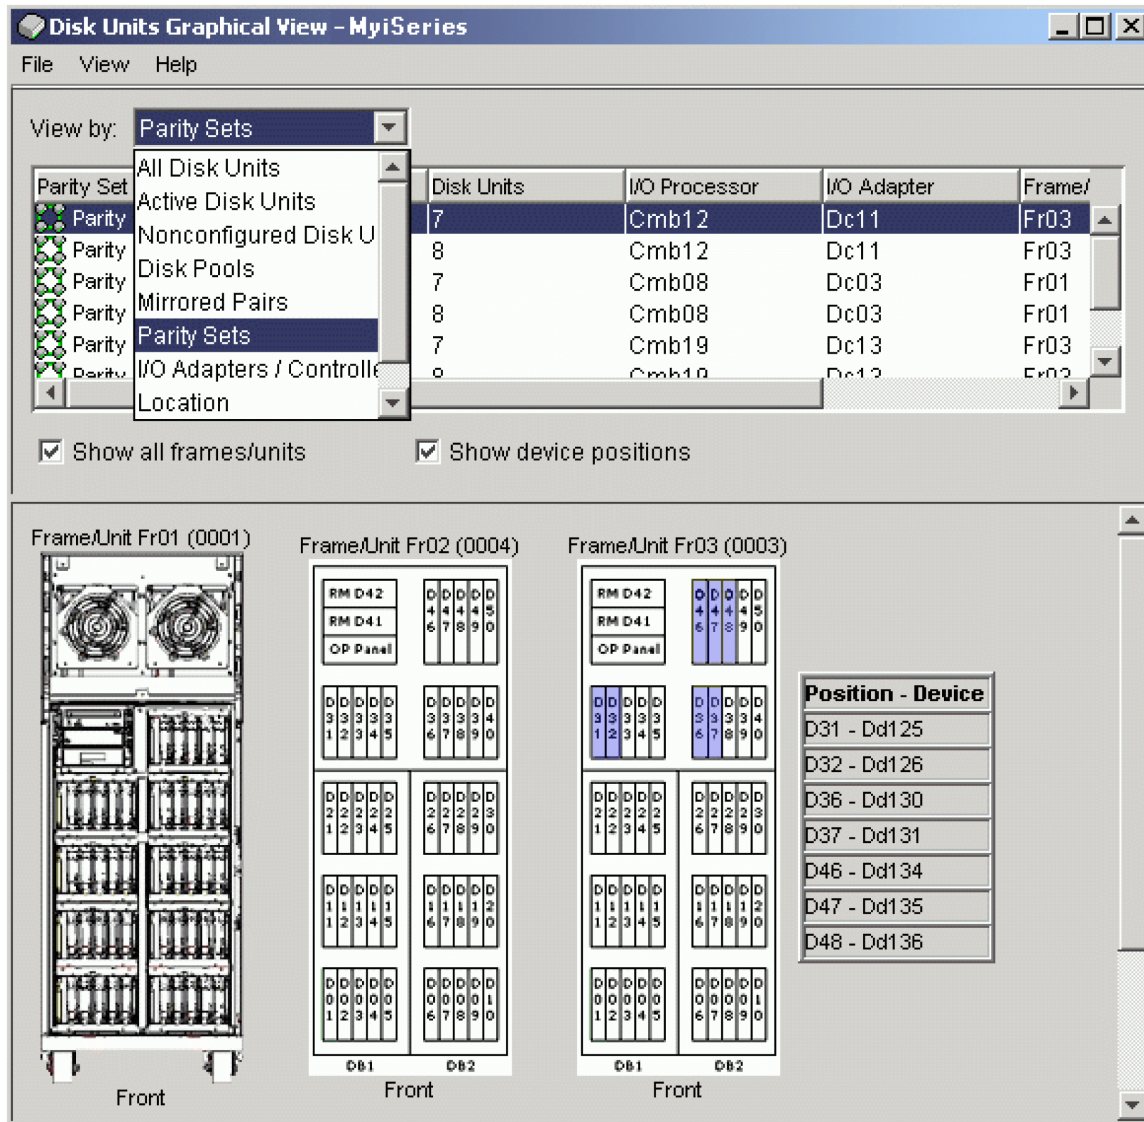
You can choose how to view the hardware in the Disk Unit Graphical View window. For example, you can select to view by disk pools, and then select a disk pool in the list to display only those frames that contain the disk units that make up the selected disk pool. You can select **Show all frames** to 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 the disk unit in a parity set (or exclude it), or rename the disk unit. If the disk unit has mirrored protection (that is, it is one of a mirrored pair), 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**.
2. Expand your iSeries **server** → **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.



When you make changes to your disk unit configuration, print the graphical view for your recovery records. To print the graphical view, on the **Disk Units Graphical View** dialog box, select **File → Print**.

To find more information about the graphical view, consult the online help for disk units.

Calculate 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 **Disk Units → Disk Pools**.

2. Right-click the source disk pool you want to view and select **Properties**.
3. 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.
4. Note the used space, free space, and the threshold from the **Capacity** tab.
5. Enter the used-space value and free-space value in the calculator.
6. 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: Calculate disk space when moving a disk unit

In the following scenario, you are planning 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.

Choose the correct procedure for configuring disks

This topic contains checklists for performing configuration procedures. Use this table to determine which checklist to use for your situation, and to determine whether dedicated service tools (DST) are required.

Task description	Procedure to follow	Requires DST?
Configure your system for the first time.	"Checklist 1: Configure disks on a new system" on page 75	Yes
Add one or more disk units that will not have device parity protection. Use this checklist if you do not plan to start device parity protection for the disks that are capable of device parity protection.	"Checklist 2: Add disk units without device parity protection" on page 76	No
Add one or more disks to an existing input/output adapter (IOA) that has built-in device parity capability. Use this checklist if you plan to protect some or all of the new disks with device parity protection.	"Checklist 3: Add disk units to an existing I/O adapter" on page 77	No
Add a new IOA that has built-in device parity capability. Use this checklist if you plan to protect some or all of the new disks with device parity protection.	"Checklist 4: Add a new I/O adapter" on page 78	Yes

Task description	Procedure to follow	Requires DST?
Move disk units between existing disk pools that do not have mirrored protection.	"Checklist 5: Move disk units between nonmirrored disk pools" on page 79	Yes
Move disk units between existing disk pools that have mirrored protection.	"Checklist 6: Move disk units between mirrored disk pools" on page 79	Yes
Delete a basic disk pool.	"Checklist 7: Delete a disk pool" on page 80	Yes
Remove one or more disk units that do not have device parity protection.	"Checklist 8: Remove disk units without device parity protection" on page 81	Yes ¹
Remove one or more disk units from an IOA. Use this checklist if device parity protection is started for some or all of the disk units that are attached to the IOA and if they are in disk pools without mirrored protection.	"Checklist 9: Remove disk units that have device parity protection from a disk pool without mirrored protection" on page 82	Yes
Remove one or more disk units from an IOA. Use this checklist if device parity protection is started for some or all of the disks units that are attached to the IOA and if they are in disk pools with mirrored protection.	"Checklist 10: Remove disk units that have device parity protection from a disk pool with mirrored protection" on page 83	Yes
Upgrade your load source disk unit while keeping device parity protection active.	Checklist 11: Upgrade load source disk unit with device parity protection	Yes
Upgrade your load source disk unit while keeping device mirrored protection active.	Checklist 12: Upgrade load source disk unit with local mirroring	Yes
¹ Unassigned disk units can be removed from independent disk pools that are varied off without requiring the system to be in DST mode.		

Checklist 1: Configure disks on a new system

This checklist shows the sequence of tasks that you use to configure disks on a new iSeries server. Whether you need to perform all the tasks depends on the disk protection that you want on your system. "Disk protection" on page 33 provides more information about the disk protection that is available.

Attention: When you perform the tasks in this checklist, the system moves large amounts of data. Make sure that you have completely saved your system in the event that you need to recover from an error situation.

Before you begin

Print a copy of this checklist. Check off the configuration tasks as you perform them. This checklist provides an important record of your actions. It might help you diagnose any problems that occur.

Most tasks in the checklist include references to other topics. Refer to these topics if you need more information about how to perform a particular task.

Task	What to do	Where to learn more
1. ___	Display your disk configuration. Currently, all of your disk units except the load source unit appear as nonconfigured.	"Evaluate the current configuration" on page 72

Task	What to do	Where to learn more
2.____	Use the Add Disk Unit wizard to add nonconfigured disks to the correct disk pools. You will have the option to start device parity protection or to start compression if disks are available for these actions.	"Add a disk unit or disk pool" on page 91
3.____	You can change this to a different storage threshold for any disk pool if required. The default storage threshold for each disk pool is 90%.	"Set the threshold of a disk pool" on page 109
4.____	If you chose to create protected disk pools and included pairs of disk units to be mirrored, you might want to restart to the dedicated service tools (DST) mode and start mirroring for those disk pools now.	"Start mirrored protection" on page 124
5.____	If you started mirrored protection for the system disk pool or a basic disk pool, wait until your system completely restarts.	
6.____	Verify that your disk configuration is correct.	"Evaluate the current configuration" on page 72
7.____	Print your disk configuration to have available in case a recovery situation occurs.	"Print your disk configuration" on page 97

Checklist 2: Add disk units without device parity protection

This checklist shows the sequence of tasks that you use to configure disks on a new iSeries server. Whether you need to perform all the tasks depends on the disk protection that you want on your system. "Disk protection" on page 33 provides more information about the disk protection that is available.

Disk pools with mirrored protection

You can add disk units to a disk pool that has mirrored protection without stopping and starting mirrored protection. You must add disk units in pairs with equal capacities. The added units will always be paired with each other. You might want to choose a later time, when your system can be unavailable for several hours, to stop and start mirrored protection. When you start mirrored protection again, the system evaluates the pairing for all disk units on your system. This might provide a higher level of availability for failures that affect an input/output adapter (IOA), an input/output processor (IOP), or a bus.

Attention

When you perform the tasks in this checklist, the system moves large amounts of data. Make sure that you have completely saved your system in the event that you need to recover from an error situation.

Before you begin

Print a copy of this checklist. Check off the configuration tasks as you perform them. This checklist provides an important record of your actions. It might help you diagnose any problems that occur.

Most tasks in the checklist include references to other topics. Refer to these topics if you need more information about how to perform a particular task.

Task	What to do	Where to learn more
1.____	Use the Disk Unit Graphical View window to find empty slots for the disk units you want to install.	"Evaluate the current configuration" on page 72

Task	What to do	Where to learn more
2.____	Right-click an empty slot and start the Install Disk Unit wizard to guide you through the process.	
3.____	Use the Add Disk Unit wizard to add nonconfigured disks to existing or new disk pools. You will have the option to start compression or to add disk units of equal capacity to mirror-protected disk pools if disks are available for these actions.	"Add a disk unit or disk pool" on page 91
4.____	Change the storage threshold for any disk pool if required. The default storage threshold for each disk pool is 90%.	"Set the threshold of a disk pool" on page 109
5.____	If you chose to create protected disk pools and you included pairs of disk units to be mirrored, you might want to restart to the dedicated service tools (DST) mode and start mirroring for those disk pools now.	"Start mirrored protection" on page 124
6.____	If you started mirrored protection for the system disk pool or for a basic disk pool, wait until your system completely restarts.	
7.____	Verify that your disk configuration is correct.	"Evaluate the current configuration" on page 72
8.____	Print your disk configuration to have available in case a recovery situation occurs.	"Print your disk configuration" on page 97

Checklist 3: Add disk units to an existing I/O adapter

This checklist shows the sequence of tasks that you use to add one or more disks to an existing input/output adapter that has built-in device parity protection. Use this checklist if you plan to protect some or all of the new disk units with device parity protection. If you do not plan to protect any of the new disk units, use "Checklist 2: Add disk units without device parity protection" on page 76.

You can use this procedure whether you have mirrored protection on your system because you start device parity protection before you add the disk units to a disk pool.

Attention: When you perform the tasks in this checklist, the system moves large amounts of data. Make sure that you have completely saved your system in the event that you need to recover from an error situation.

Before you begin

Print a copy of this checklist. Check off the configuration tasks as you or the service representative perform them. This checklist provides an important record of your actions. It might help you diagnose any problems that occur.

Most tasks in the checklist include links to other topics. Refer to these topics if you need more information about how to perform a particular task.

Task	What to do	Where to learn more
1.____	Physically attach disk units using the Install Disk Unit wizard available in the Disk Units Graphical View window.	"Evaluate the current configuration" on page 72

Task	What to do	Where to learn more
2.____	Use the Add Disk Unit wizard to add nonconfigured disks to the correct disk pools. The wizard will allow you to include the disk units that you want to protect in device parity protection	"Add a disk unit or disk pool" on page 91
3.____	The default storage threshold for each disk pool is 90%. If you want a different storage threshold for any disk pool, you can change it.	"Set the threshold of a disk pool" on page 109
4.____	Verify that your disk configuration is correct.	"Evaluate the current configuration" on page 72
5.____	Print your disk configuration to have available in case a recovery situation occurs.	"Print your disk configuration" on page 97

Checklist 4: Add a new I/O adapter

This checklist shows the sequence of tasks that you use to add a new input/output adapter (IOA) and new disk units to your system. Use this checklist if you plan to protect some or all of the new disks with device parity protection. You can use this procedure whether you have mirrored protection on your system because you start device parity protection before you add the disk units to a disk pool. If you do have mirrored protection and you are adding disks that do not have device parity protection, you must add them in pairs that have equal capacities.

Note: If you do not plan to start device parity protection for any of the new disks, use the procedure in "Checklist 2: Add disk units without device parity protection" on page 76 to add the new disks.

Attention: When you perform the tasks in this checklist, the system moves large amounts of data. Make sure that you have completely saved your system in the event that you need to recover from an error situation.

Before you begin

Print a copy of this checklist. Check off the configuration tasks as you or the service representative perform them. This checklist provides an important record of your actions. It may help you diagnose any problems that occur.

Most tasks in the checklist include links to other topics. Refer to these topics if you need more information about how to perform a particular task.

Task	What to do	Where to learn more
1.____	Install the new input/output adapter in the server. This is normally done by a service representative.	"Evaluate the current configuration" on page 72
2.____	Physically attach disk units to the new IOA using the Install Disk Unit wizard, which you can access from the Disk Units Graphical View window.	"Evaluate the current configuration" on page 72
3.____	Use the Add Disk Unit wizard to add nonconfigured disks to the correct disk pools. Use the option to start device parity protection.	"Add a disk unit or disk pool" on page 91
4.____	If you want a different storage threshold for any disk pool, change it. The default storage threshold for each disk pool is 90%.	"Set the threshold of a disk pool" on page 109
5.____	You may want to restart to the dedicated service tools (DST) mode in order for the device parity protection to take effect.	"Work with device parity protection" on page 106
6.____	Wait until your system completely restarts.	

Task	What to do	Where to learn more
7.____	Verify that your disk configuration is correct.	"Evaluate the current configuration" on page 72
8.____	Print your disk configuration to have available in case a recovery situation occurs.	"Print your disk configuration" on page 97

Checklist 5: Move disk units between nonmirrored disk pools


This checklist shows the sequence of tasks that you use to move one or more disk units from one basic disk pool to another basic disk pool. Use these tasks when you do not have mirrored protection active for the disk pools. You must restart your server to DST mode to perform the tasks in this checklist.

Attention: When you perform the tasks in this checklist, the system moves large amounts of data. Make sure that you have completely saved your system in the event that you need to recover from an error situation.

Before you begin

Print a copy of this checklist. Check off the configuration tasks as you perform them. This checklist provides an important record of your actions. It may help you diagnose any problems that occur.

Most tasks in the checklist include links to other topics. Refer to these topics if you need more information about how to perform a particular task.

Task	What to do	Where to learn more
1.____	Display your current disk configuration.	"Evaluate the current configuration" on page 72
2.____	Calculate the space requirements for both the source and destination disk pools for the disk units.	"Calculate disk space requirements" on page 73
3.____	Use option 21 from the Save menu to save your entire system.	"Save your server with the GO SAVE command"
4.____	Restart your server and select the option to use dedicated service tools (DST).	"How to start dedicated service tools (DST)" in Backup and Recovery  . From the taskpad in iSeries Navigator, select Open iSeries Navigator Service Tools .
5.____	From the Disk Units Graphical View window, right-click the disk unit you want to move, and select Move.	"Move and remove disk units" on page 92
6.____	Verify that your disk configuration in correct.	"Evaluate the current configuration" on page 72
7.____	Print your disk configuration to have available in case a recovery situation occurs.	"Print your disk configuration" on page 97
8.____	Restart your server.	

Checklist 6: Move disk units between mirrored disk pools


This checklist shows the sequence of tasks that you use to move one or more disk units from one basic disk pool to another basic disk pool. Use these tasks when one or more of the disk pools has mirrored protection. You cannot move disk units when mirrored protection is active. Instead, you remove mirrored pairs from the source disk pool and add them to the destination disk pool. You must restart your server to dedicated service tools (DST) mode to perform the tasks in this checklist.

Attention: When you perform the tasks in this checklist, the system moves large amounts of data. Make sure that you have completely saved your system in the event that you need to recover from an error situation.

Before you begin

Print a copy of this checklist. Check off the configuration tasks as you perform them. This checklist provides an important record of your actions. It may help you diagnose any problems that occur.

Most tasks in the checklist include links to other topics. Refer to these topics if you need more information about how to perform a particular task.

Task	What to do	Where to learn more
1.____	Display your current disk configuration.	"Evaluate the current configuration" on page 72
2.____	Calculate the space requirements for both the source and destination disk pools for the disk units.	"Calculate disk space requirements" on page 73
3.____	Use option 21 from the Save menu to save your entire system.	Save your server with the GO SAVE command
4.____	Restart your server and select the option to use dedicated service tools (DST).	"How to Start Dedicated Service Tools (DST)" in Backup and Recovery  . From the taskpad in iSeries Navigator, select Open iSeries Navigator Service Tools .
5.____	In the disk units graphical view window, filter by mirrored pair, and then hold the Ctrl key until you have selected each disk unit of the mirrored pair. Right-click one of the selected disk units and select Remove .	"Move and remove disk units" on page 92
6.____	Add nonconfigured disk units to the correct disk pools. If you are adding disk units to a protected disk pool and the new disk units do not have device parity protection, you must add pairs of disk units that have identical capacities.	"Add a disk unit or disk pool" on page 91
7.____	If you created a new disk pool when you added disk units, the system set the storage threshold of the disk pool to 90%. If you want a different storage threshold for any disk pool, change it.	"Set the threshold of a disk pool" on page 109
8.____	If you created any new disk pools and you want those disk pools to have mirrored protection, start mirrored protection now.	"Start mirrored protection" on page 124
9.____	Verify that your disk unit configuration is correct.	"Evaluate the current configuration" on page 72
10.____	Print your disk configuration to have available in case a recovery situation occurs.	"Print your disk configuration" on page 97

Checklist 7: Delete a disk pool



This checklist shows the sequence of tasks that you use to delete a basic disk pool or an independent disk pool. If you want to delete or clear an independent disk pool that is unavailable, you can do so when your system is fully restarted. For all other disk pools, you need to restart your system to Dedicated Service Tools (DST) mode before clearing or deleting them.

Attention: Make sure that you have completely saved your system in the event that you need to recover from an error situation. Also note that when a disk pool is deleted, all data remaining in that disk pool is lost.

Before you begin

Print a copy of this checklist. Check off the configuration tasks as you perform them. This checklist provides an important record of your actions. It might help you diagnose any problems that occur.

Most tasks in the checklist include links to other topics. Refer to these topics if you need more information about how to perform a particular task.

Task	What to do	Where to learn more
1. ___	Display your current disk configuration.	"Evaluate the current configuration" on page 72
2. ___	Calculate the space requirements for the remaining disk pools.	"Calculate disk space requirements" on page 73
3. ___	Use option 21 from the Save menu to save your entire system.	Save your server with the GO SAVE command
4. ___	Remove objects from the disk pool that you are deleting, or move the objects to a different disk pool.	Backup and Recovery manual 
5. ___	Restart your server and select the option to use dedicated service tools (DST).	"How to Start Dedicated Service Tools (DST)" in Backup and Recovery  . From the taskpad in iSeries Navigator, select Open iSeries Navigator Service Tools .
6. ___	Delete the disk pool. This procedure places all of the disks that were assigned to the deleted disk pool in nonconfigured status.	"Delete a disk pool" on page 109
7. ___	If you want to add the now nonconfigured disk units to a different disk pool, refer to Checklist 2 or 3.	"Checklist 2: Add disk units without device parity protection" on page 76 or "Checklist 3: Add disk units to an existing I/O adapter" on page 77
8. ___	Verify that your disk unit configuration is correct.	"Evaluate the current configuration" on page 72
9. ___	Print your disk configuration to have available in case a recovery situation occurs.	"Print your disk configuration" on page 97
10. ___	Restart your server.	

Checklist 8: Remove disk units without device parity protection


This checklist shows the sequence of tasks that you use to remove one or more disk units from your system when the disk units do not have device parity protection. Use these tasks when you are permanently removing disk units from your system. Do not use these tasks when you are repairing or replacing a failed disk unit. You must restart your server to dedicated service tools (DST) mode to perform the tasks in this checklist.

Attention: When you perform the tasks in this checklist, the system moves large amounts of data. Make sure that you have completely saved your system in the event that you need to recover from an error situation.

Before you begin

Print a copy of this checklist. Check off the configuration tasks as you perform them. This checklist provides an important record of your actions. It may help you diagnose any problems that occur.

Most tasks in the checklist include links to other topics. Refer to these topics if you need more information about how to perform a particular task.

Task	What to do	Where to learn more
1.____	Display your current disk configuration.	"Evaluate the current configuration" on page 72
2.____	Calculate the space requirements for the disk pools that are involved in disk removal.	"Calculate disk space requirements" on page 73
3.____	Use option 21 from the Save menu to save your entire system.	Save your server with the GO SAVE command
4.____	Restart your server and choose the option to use dedicated service tools (DST).	"How to Start Dedicated Service Tools (DST)" in Backup and Recovery  . From the taskpad in iSeries Navigator, select Open iSeries Navigator Service Tools .
5.____	Remove disk units that you plan to remove from the system.	"Move and remove disk units" on page 92
6.____	Verify that your disk unit configuration is correct.	"Evaluate the current configuration" on page 72
7.____	Print your disk configuration to have available in case a recovery situation occurs.	"Print your disk configuration" on page 97
8.____	Continue restarting your server.	

Checklist 9: Remove disk units that have device parity protection from a disk pool without mirrored protection


This checklist shows the sequence of tasks that you use to remove one or more disk units from an input/output adapter with built-in device parity protection. These tasks apply when the disk pools containing the disk units do not have mirrored protection and when device parity protection is started for the IOA. Use these tasks when you are permanently removing disk units from your system. Do not use these tasks when you are repairing or replacing a failed hard disk. You must restart your server to dedicated service tools (DST) mode to perform the tasks in this checklist.

Attention: When you perform the tasks in this checklist, the system moves large amounts of data. Make sure that you have completely saved your system in the event that you need to recover from an error situation.

Before you begin

Print a copy of this checklist. Check off the configuration tasks as you or the service representative perform them. This checklist provides an important record of your actions. It may help you diagnose any problems that occur.

Most tasks in the checklist include links to other topics. Refer to these topics if you need more information about how to perform a particular task.

Task	What to do	Where to learn more
1.____	Display your current disk configuration.	"Evaluate the current configuration" on page 72
2.____	Calculate the space requirements for the disk pools that are involved in disk removal.	"Calculate disk space requirements" on page 73
3.____	Use option 21 from the Save menu to save your entire system.	Save your server with the GO SAVE command
4.____	Restart your server and select the option to use dedicated service tools (DST).	"How to Start Dedicated Service Tools (DST)" in Backup and Recovery  . From the taskpad in iSeries Navigator, select Open iSeries Navigator Service Tools .

Task	What to do	Where to learn more
5.____	Remove disk units that you plan to remove from the system.	"Move and remove disk units" on page 92
6.____	Exclude the disk units from device parity protection. If you were successful in excluding the disk units, skip to task 8. Otherwise, continue to task 7.	"Exclude disk units from a parity set" on page 123
7.____	Stop device parity protection for all the disk units in the IOP.	"Stop device parity protection" on page 122
8.____	Physically remove disk units. If you stopped device parity protection in task 7, continue with task 9. If you did not stop device parity protection, skip to task 10.	"Move and remove disk units" on page 92
9.____	Start device parity protection again.	"Start device parity protection" on page 122
10.____	Verify that your disk unit configuration is correct.	"Evaluate the current configuration" on page 72
11.____	Print your disk configuration to have available in case a recovery situation occurs.	"Print your disk configuration" on page 97
12.____	Restart your server.	

Checklist 10: Remove disk units that have device parity protection from a disk pool with mirrored protection


This checklist shows the sequence of tasks that you use to remove one or more disk units from an input/output adapter that is capable of device parity protection. These tasks apply when the disk pools that contain the disk units have mirrored protection and when the disk units have device parity protection. Use these tasks when you are permanently removing disk units from your system. Do not use these tasks when you are repairing or replacing a failed disk unit. You must restart your server to the dedicated service tools (DST) mode to perform the tasks in this checklist.

Attention: When you perform the tasks in this checklist, the system moves large amounts of data. Make sure that you have completely saved your system in the event that you need to recover from an error situation.

Before you begin

Print a copy of this checklist. Check off the configuration tasks as you perform them. This checklist provides an important record of your actions. It may help you diagnose any problems that occur.

Most tasks in the checklist include links to other topics. Refer to these topics if you need more information about how to perform a particular task.

Task	What to do	Where to learn more
1.____	Display your current disk configuration.	"Evaluate the current configuration" on page 72
2.____	Calculate the space requirements for the disk pools that are involved in disk removal.	"Calculate disk space requirements" on page 73
3.____	Use option 21 from the Save menu to save your entire system.	Save your server with the GO SAVE command
4.____	Restart your server and select the option to use dedicated service tools (DST).	"How to Start Dedicated Service Tools (DST)" in Backup and Recovery  . From the taskpad in iSeries Navigator, select Open iSeries Navigator Service Tools .

Task	What to do	Where to learn more
5.____	Remove disk units that you plan to remove from the system.	"Move and remove disk units" on page 92
6.____	Exclude the disk units from device parity protection. If you were successful in excluding the disk units, skip to task 9. Otherwise, continue to task 7.	"Exclude disk units from a parity set" on page 123
7.____	Stop mirrored protection for the disk pools that will have disk units removed. When you stop mirrored protection, one disk unit from each mirrored pair becomes unconfigured. You need to stop mirrored protection only if the disk pool contains other disk units that are attached to the IOP and have device parity protection.	"Stop mirrored protection" on page 124
8.____	Stop device parity protection for all the disk units in the IOP.	"Stop device parity protection" on page 122
9.____	Physically remove disk units. This is normally done by a service representative. If you stopped device parity protection in task 8, continue with task 10. If you did not stop device parity protection, skip to task 14.	
10.____	Start device parity protection again.	"Start device parity protection" on page 122
11.____	Add unconfigured disk units to the correct disk pools. These disk units became unconfigured when you stopped mirrored protection in task 7.	"Add a disk unit or disk pool" on page 91
12.____	If you created a new disk pool on your system when you added disk units, the system set the storage threshold of the disk pool to 90%. If you want a different storage threshold for any disk pool, you can change it.	"Set the threshold of a disk pool" on page 109
13.____	Start mirrored protection for the disk pools that had mirrored protection stopped in task 7.	"Start mirrored protection" on page 124
14.____	Verify that your disk unit configuration is correct.	"Evaluate the current configuration" on page 72
15.____	Print your disk configuration to have available in case a recovery situation occurs.	"Print your disk configuration" on page 97

Checklist 11: Upgrade load source disk unit with device parity protection

This checklist shows the sequence of tasks to upgrade your load source disk unit (unit 1) with a disk unit that has at least a 17-GB capacity while keeping device parity active. The units that are replaced in this procedure are discarded.

Print a copy of this checklist. Check off the configuration tasks as you perform them. This checklist provides an important record of your actions. It may help you diagnose any problems that occur.

Before you begin

Evaluate your disk configuration and record your answers. The information entered in the Before you begin table is needed to answer the questions in the Load source planning sections.

Table 1. Disk configuration questions

Disk configuration questions	Disk configuration answers
How many disk units are in the parity set that contain the load source disk unit? Note: The parity set will contain between 3-18 disk units.	
Where are the disk units in the parity set that contain the load source disk unit located? Note: It is recommended that you print the graphical view of the device parity set or optionally mark the disk units in the parity set. Make sure that you can identify the load source unit separately from the other disk units.	
How many replacement disk units do you have? Note: You need a minimum of three disk units with the same capacity.	

Load source planning steps

Answer the questions below. If you answer Yes to all of these questions, you can perform the Load source disk unit upgrade. However, if you answer No to any of the questions, call your next level of support to perform this upgrade.

Table 2. Load source planning steps

Load source planning steps	Planning requirements answers
Does your load source disk unit (unit 1) have device parity protection?	Yes / No
Are there enough open slots to install your replacement units? Note: The number of slots must be at least the same number of open slots as the number of disk units that you have in the parity set that contain the load source. It must also be under the IOA that contains your load source disk unit (unit 1).	Yes / No
Do you have equal or more replacement disk units than the number of disk units in the device parity set that contain the load source disk unit?	Yes / No
Do you know how to physically install and remove disk units from your system? Install and remove replacement units on an IBM iSeries server. Install and remove replacement units on an IBM eServer i5.	Yes / No
Do you have iSeries Navigator or know how to find the physical location of disk units on your system? You will need to know this information in several tasks below.	Yes / No
The load source upgrade will require multiple hours to complete. Make sure that you schedule the upgrade during a time frame when the system can be unavailable to perform normal system activities. Do you have time to perform the load source upgrade?	Yes / No

Load source disk unit upgrade

Note: If problems occur that are not described in this procedure, contact your next level of support.

Task	What to do
1. ___	Back up your server using the GO SAVE command.
2. ___	Turn off your system or logical partition by entering the following command. PWRDOWNSYS *IMMED RESTART(*NO) The Power Down System (PWRDOWNSYS) command prepares the system to end. Use the immediate (*IMMED) value to end all active jobs immediately and use the RESTART (*NO) value to power off the system so that you can install disk units in the next task.

Task	What to do
3.____	<p>Install the replacement disk units.</p> <ol style="list-style-type: none"> 1. Ensure that there are enough open disk slots under the IOA that contain the load source disk unit (unit 1). 2. Install replacement units on your system. <p>Note:</p> <ol style="list-style-type: none"> a. Ensure that the disk units being installed have the same capacity as one another and that each disk unit has at least a 17-GB capacity. b. It is recommended that you mark these disk units with tape to remember their location. Mark these disk units differently than how you marked the existing disk units.
4.____	<p>Work with Dedicated Service Tools (DST).</p> <ol style="list-style-type: none"> 1. Power on your system or logical partition to Dedicated Service Tools (DST). <p>Note: Verify that you are in manual mode before you power on your system.</p>
5.____	<p>Start device parity on the replacement disk units, and select the type of RAID protection that you want.</p> <ol style="list-style-type: none"> 1. On the Use Dedicated Service Tools (DST) menu, select Work with disk units 2. On the Work with Disk Units display, select Work with disk configuration. 3. On the Work with Disk Configuration display, select Work with device parity protection. 4. On the Work with Device Parity Protection display, select Start device parity protection. You may select RAID 5 or RAID 6 protection. <p>Note: RAID 6 protection requires special hardware. If your hardware does not meet the requirements, it will default to RAID 5 protection.</p> <ol style="list-style-type: none"> 5. Type a 1 in the Option column of the storage subsystems that will have device parity protection. 6. You are shown the Confirm Starting Device Parity Protection display. The display shows all the disk subsystems that you have selected and the individual disk units that are eligible to be started. Disk units that have an asterisk (*) in the ASP and Unit columns are not yet configured. Verify that these disk units are the ones installed in task 3 and that the ASP and unit number all have an asterisk (*). 7. Press the Enter key to continue. 8. After you have verified that these disk units are the replacement units that you installed in task 3, press the Enter key, to start device parity protection. This procedure continues to run until it is complete. 9. The status display shows how the operation is proceeding. When the function is complete, return to the Use Dedicated Service Tools (DST) menu.
6.____	<p>Add the non-configured disk units into the auxiliary storage pools.</p> <ol style="list-style-type: none"> 1. On the Use Dedicated Service Tools (DST) menu, select Work with disk units. 2. On the Work with Disk Units display, select Work with disk configuration. 3. On the Work with Disk Configuration display, select Work with ASP configuration. 4. On the Work with ASP Configuration display, select Add units to ASPs. <p>Note: Add all but one of your disk units. The disk unit that remains unconfigured will be used as your load source in task 9.</p> <ol style="list-style-type: none"> 5. Decide which ASP will include the new disk units. Type this ASP number beside each of the new units on the screen and press Enter. 6. On the Confirm Add Units display, press Enter. 7. The status display shows how the operation is proceeding. When the function is complete, return to the Use Dedicated Tools (DST) menu.

Task	What to do
10.__	Turn off the system or logical partition. Note: This task is critical. Follow the directions thoroughly. 1. On the Use Dedicated Service Tools (DST) menu, select Start a service tool . 2. On the Start a Service Tool display, select Operator panel functions . 3. Use function F10 to power off the system and press Enter. Note: There is no command line because you are using Dedicated Service Tools.
11.__	Physically remove the disk units. 1. Physically remove the disk units that you recorded in task 7, step 6, and the load source disk unit (unit 1). Note: You can identify which disk units are you removing from printout of the graphical view on your iSeries navigator or from the disk units that you marked.
12.__	Move the replacement disk unit. 1. Move the replacement disk unit that contains the load source information into the slot where the old load source disk unit (unit 1) originally resided.
13.__	Work with Dedicated Service Tools (DST). 1. Power on your system or logical partition to Dedicated Service Tools (DST).
14.__	Examine the configuration. 1. On the Use Dedicated Service Tools (DST) menu, select Work with disk units . 2. On the Work with Disk Units display, select Work with disk configuration . 3. On the Work with Disk Configuration display, select Display disk configuration . 4. On the Display Disk Configuration display, select Display disk configuration status . 5. Examine the configuration information to ensure that the load source disk unit (unit 1) is one of the replacement disk units that you installed in task 3. 6. Verify that the load source serial number matches the number that you wrote in task 9, step 7.
15.__	IPL your system to i5/OS.

If this procedure has been completed correctly, your load source will be upgraded and device parity protection will be active.

| If this procedure has not been completed correctly, contact your next level of support.

| **Checklist 12: Upgrade load source disk unit with local mirroring**

| This checklist shows the sequence of tasks to upgrade your load source disk unit (unit 1) with a disk unit that has at least a 17-GB capacity while keeping mirrored protection active. The units that are replaced in this procedure are discarded.

| Print a copy of this checklist. Check off the configuration tasks as you perform them. This checklist provides an important record of your actions. It may help you diagnose any problems that occur.

| **Before you begin**

| Evaluate your disk configuration and record your answers. The information entered in the Before you begin table is needed to answer the questions in the Load source planning section.

| *Table 3. Disk configuration questions*

Disk configuration questions	Disk configuration answers
Where is the load source disk unit and the mirrored load source disk unit located? Note: It is recommended that you print the graphical view of the load source disk unit and the mirrored load source disk unit or optionally mark them.	

Table 3. Disk configuration questions (continued)

Disk configuration questions	Disk configuration answers
How many replacement disk units do you have? Note: You need two disk units of the same capacity.	

Load source planning steps

Answer the questions below. If you answer Yes to all of these questions, you can perform the load source disk unit upgrade. However, if you answer No to any of the questions, call your next level of support to perform this upgrade.

Table 4. Load source planning steps

Load source planning steps	Planning requirements answers
Does your load source disk unit (unit 1) have mirrored protection?	Yes / No
Are there enough open slots to install your replacement units? Note: There must be at least two open slots for the replacement disk units.	Yes / No
Do you know how to physically install and remove disk units from your system? Install and remove replacement units on an IBM iSeries server. Install and remove replacement units on an IBM eServer i5.	Yes / No
Do you have iSeries Navigator or know how to find the physical location of disk units on your system? You will need to know this information in several tasks below.	Yes / No
The load source upgrade will require multiple hours to complete. Make sure that you schedule the upgrade during a time frame when the system can be unavailable to perform normal system activities. Do you have time to perform the load source upgrade?	Yes / No

Load source disk unit upgrade

Note: If problems occur that are not described in this procedure contact your next level of support.

Task	What to do
1.____	Back up your server using the GO SAVE command.
2.____	Use iSeries Navigator to find the physical location of the load source disk unit (unit 1) and the mirrored load source disk unit. 1. Open the iSeries Navigator Service Tools Window. 2. Right-click All Disk Units and select Graphical View. Note: It is recommended that you mark these disk units with tape to remember their location.
3.____	Turn off your system or logical partition by entering the following command. <code>PWRDWN SYS *IMMED RESTART(*NO)</code> The Power Down System (PWRDWN SYS) command prepares the system to end. Use the immediate (*IMMED) value to end all active jobs immediately and use the RESTART (*NO) value to power off the system so that you can install disk units in the next task.

Task	What to do
4.____	<p>Install one of the replacement disk units.</p> <ol style="list-style-type: none"> 1. Install replacement unit on your system. <p>Note:</p> <ol style="list-style-type: none"> 1. Ensure that the disk unit has at least a 17-GB capacity. 2. It is recommended that you mark these disk units with tape to remember their location. Ensure that you mark them differently than how they are marked in task 2.
5.____	<p>Work with Dedicated Service Tools (DST).</p> <ol style="list-style-type: none"> 1. Power on your system or logical partition to Dedicated Service Tools (DST). Note: Verify that you are in manual mode before you power on your system.
6.____	<p>Copy the load source disk unit (unit 1) to the replacement disk unit.</p> <ol style="list-style-type: none"> 1. On the Use Dedicated Service Tools (DST) menu, select Work with disk units. 2. On the Work with Disk Units display, select Work with disk unit recovery. 3. On the Work with Disk Unit Recovery display, select Copy disk unit data. 4. Type 1 next to load source disk unit (unit 1) and press Enter. 5. Type 1 next to one of the disk units that you installed in task 3. 6. Determine the location of the replacement load source unit <ol style="list-style-type: none"> a. Open the iSeries Navigator Service Tools Window. b. Right-click All Disk Units and choose Graphical View. c. Find the disk unit with the serial number listed above, and note the location of that unit. 7. Write down the serial number of the replacement unit that is going replace the load source. _____ Note: This information is available in iSeries Navigator and is needed for task 9. 8. On the Confirm Copy Disk Unit Data display, press Enter. 9. The status display shows how the operation is proceeding. When the function is complete, return to the Use Dedicated Service Tools (DST) menu.
7.____	<p>Turn off the system or logical partition.</p> <p>Note: This task is critical. Follow the directions thoroughly.</p> <ol style="list-style-type: none"> 1. On the Use Dedicated Service Tools (DST) menu, select Start a service tool. 2. On the Start a Service Tool display, select Operator panel functions. 3. Use function F10 to power off the system. 4. Press Enter. <p>Note: There is no command line because you are using Dedicated Service Tools.</p>
8.____	<p>Physically remove the old load source disk unit.</p> <ol style="list-style-type: none"> 1. Physically remove the old load source disk unit. Note: This is the original load source disk unit that was marked in task 2.
9.____	<p>Move the replacement disk unit.</p> <ol style="list-style-type: none"> 1. Move the replacement disk unit that now contains the load source information into the slot where the old load source disk unit (unit 1) originally resided.
10.____	<p>Replace the mirrored load source disk unit.</p> <ol style="list-style-type: none"> 1. Physically remove the mirrored load source disk unit. 2. Install the second replacement unit into the slot that the mirrored load source originally resided.
11.____	<p>Work with Dedicated Service Tools (DST).</p> <ol style="list-style-type: none"> 1. Power on your system or logical partition to Dedicated Service Tools (DST).

Task	What to do
12.__	Replace the configured disk unit. <ol style="list-style-type: none"> 1. On the Use Dedicated Service Tools (DST) menu, select Work with disk units. 2. On the Work with Disk Units display, select Work with disk unit recovery. 3. On the Work with Disk Unit Recovery display, select Replace configured unit. 4. Type 1 next to the suspended unit 1 disk unit and press Enter. 5. Type a 1 next to the newly installed disk unit and press Enter. 6. On the Confirm Replace Configured Unit display, press Enter 7. The status display shows how the operation is proceeding. When the function is complete, return to the Use Dedicated Service Tools (DST) menu.
13.__	Examine the configuration. <ol style="list-style-type: none"> 1. On the Use Dedicated Service Tools (DST) menu, select Work with disk units. 2. On the Work with Disk Units display, select Work with disk configuration. 3. On the Work with Disk Configuration display, select Display disk configuration. 4. On the Display Disk Configuration display, select Display disk configuration status. 5. Examine the configuration information to ensure that the load source disk unit (unit 1) is one of the replacement disk units that you installed in task 4. 6. Verify that the load source serial number matches the number that you wrote in task 6, step 7.
14.__	IPL your system to i5/OS.

If this procedure has been completed correctly, your load source will be upgraded and mirrored protection will be active.

If this procedure has not been completed correctly, contact your next level of support.

Create a basic disk pool

The New Disk Pool wizard saves you time by bundling several time-consuming configuration functions into one efficient process. It also takes the guesswork out of disk unit configuration because it understands the capabilities of your system and only offers valid choices. For instance, the wizard does not list the option to start compression unless your server has that capability.

The New Disk Pool wizard allows you to create a basic disk pool or independent disk pool or to use an existing disk pool to add new or nonconfigured disk units. When you choose to create a *protected* disk pool, the wizard forces you to include the disk units in Device parity protection or to add enough disk units of the same capacity to start Mirrored protection. The wizard also gives you the option of balancing data across the disk pool or starting disk compression if these are permissible actions for your system configuration. You decide which options to choose so that the operation is tailored to your system.

Prerequisites

“iSeries Navigator requirements for disk management” on page 49

To use the New Disk Pool wizards, follow these steps:

1. In iSeries Navigator, expand **Disk Units**.
2. To create a new disk pool, right-click **Disk Pools** and select **New Disk Pool**.
3. Follow the instructions in the wizard to complete the task.

Add a disk unit or disk pool

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

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

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

Add Disk Unit wizard

“iSeries Navigator requirements for disk management” on page 49.

To use the Add Disk Unit wizard, follow these steps:

1. In iSeries Navigator, expand **Disk Units**.
2. To add disk units, right-click **All Disk Units** and select **Add Disk Unit**.
3. Follow the instructions in the wizard to complete the task.

Move and remove disk units

As your storage needs change, you can select to move a disk unit from one disk pool to another disk pool. When you move a disk unit, the server first moves all of the data on that disk unit to other disk units in the original disk pool. You cannot move disk units to or from an independent disk pool. For disk units in system disk pools and basic disk pools, you need to restart your system to dedicated service tools (DST) mode before moving them.

When you remove a disk unit, the server redistributes the data on that disk unit to other disk units in the disk pool. If you want to remove a disk unit from an independent disk pool that is unavailable, you can do so when your system is fully restarted. For all other disk pools, you need to restart your system to DST before removing them.

Depending on disk unit capacity and performance, the move or remove process can take from several minutes to over an hour to complete, potentially affecting system performance.

To move or remove a disk unit from a disk pool, follow these steps:

1. In iSeries Navigator, expand **Disk Units**.
2. Right-click the disk unit you want to move and select **Move** or **Remove**.
3. Follow the instructions on the resulting dialog box.

Configure 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.

“iSeries Navigator requirements for disk management” on page 49

Create 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 device 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.

Create a dedicated independent disk pool

“iSeries Navigator requirements for disk management” on page 49.

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 your iSeries **server** → **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 server 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 server 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 (WRKDEVDD) command to verify that the device description and independent disk pool name match.

Create 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. See “Plan for independent disk pools” on page 51.

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” on page 49.

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.

Using iSeries Navigator

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

1. “Create a cluster” on page 95. To use switchable independent disk pools, an iSeries cluster is required.
2. “Make your hardware switchable” on page 114. 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” on page 95. 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 configuration" on page 97. 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:

1. "Start switchable hardware group" on page 96. Start the switchable hardware group to enable device resiliency for the switchable hardware group.
2. "Make a disk pool available" on page 112. 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.

Using CL commands and APIs

To create a switchable independent disk pool using CL commands and APIs, do the following:

You can use CL commands and APIs for creating a switchable independent disk pool, however there are some tasks that require that you use iSeries Navigator.

1. **Create the cluster.** Create the cluster with required node using the CRTCLU (Create Cluster) Command.
2. **Start the nodes that comprise the cluster.** 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 or set of independent disk pools 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 independent disk pool device descriptions using the CRTCRG (Create Cluster Resource Group) Command.
6. **"Make your hardware switchable" on page 114.** 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 the disk pool on the node that owns the disk units using the New Disk Pool wizard when the server is fully restarted. Make sure clustering is active before you start. Name the independent disk pool to match the device description resource name that you gave in step 3. As you add disk units, it is best to localize disk units in the same expansion unit or IOP. Do not spread the disk pool across more device parity sets than necessary.
8. **Print your disk configuration.** Print your disk configuration to have 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.

You have now created a switchable independent disk pool. The remaining steps are required to prepare it for use.

9. **Start the cluster resource group.** Start the cluster resource group to enable device resiliency using the STRCRG (Start Cluster Resource Group) Command.

10. **Make the disk pool available.** To access the disk units in an independent disk pool you must vary on the disk pool using the VRYCFG (Vary Configuration) command.
11. **Perform a test switchover.** Before you add data to the disk pool, perform a test switchover to ensure the configuration functions as you planned. Use the CHGCRGPRI (Change CRG Primary) command.

You are now ready to populate the independent disk pool with directories and libraries. Before you do, be sure to read “Independent disk pools with distinct databases” on page 20.

Create a cluster: For an independent disk pool to be switchable among servers or for geographic mirroring to be enabled, an iSeries cluster is required. An iSeries cluster is a collection or group of one or more servers that work together as a single server. For complete documentation on clusters and how they work, see Clusters.

There are multiple way 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. For step-by-step instructions on how to create a cluster, see Create a cluster in the Clusters topic.
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. See Adjust the cluster version of a cluster for details.
3. Start all nodes in the cluster, or at least those that will be in the device domains. See Start a cluster node for details.

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

Start switchable hardware group: To enable device resiliency for the switchable hardware group, you must start the switchable hardware group.

To start a switchable hardware group, follow these steps:

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 Hardware**.
5. Right-click the switchable hardware 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.

Create a 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 your **iSeries server** → **Configuration and Service** → **Hardware** → **Disk Units** .
3. Sign on to service tools if the Service Tools Signon dialog box is displayed. You may be required to configure the service tools server if you have not done so already.
4. Right-click **Disk Pools** and select **New Disk Pool**.
5. 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.

Print your disk configuration:

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

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

1. In iSeries Navigator, expand **My Connections**.
2. Expand your iSeries **server** → **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**.

Create a new disk pool group

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.

You can create a disk pool group and add disk units to the individual disk pools by using the New Disk Pool wizard. If you have existing UDFS disk pools that you would like to include in a disk pool group, see “Convert a UDFS disk pool to primary” or “Convert a UDFS disk pool to secondary” on page 98.

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 “Create a switchable independent disk pool” on page 93.

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

1. In iSeries Navigator, expand **My Connections** (or your active environment).
2. Expand any iSeries **server** → **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.

Convert UDFS disk pools

Support for library-based objects through the use of primary and secondary disk pools was introduced at V5R2. If you have existing user-defined file system (UDFS) disk pools on your server, you can convert them to primary and secondary disk pools. This allows them to support library-based objects.

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.

Convert 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 your iSeries **server** → **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 “Convert a UDFS disk pool to secondary.”

Convert 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:

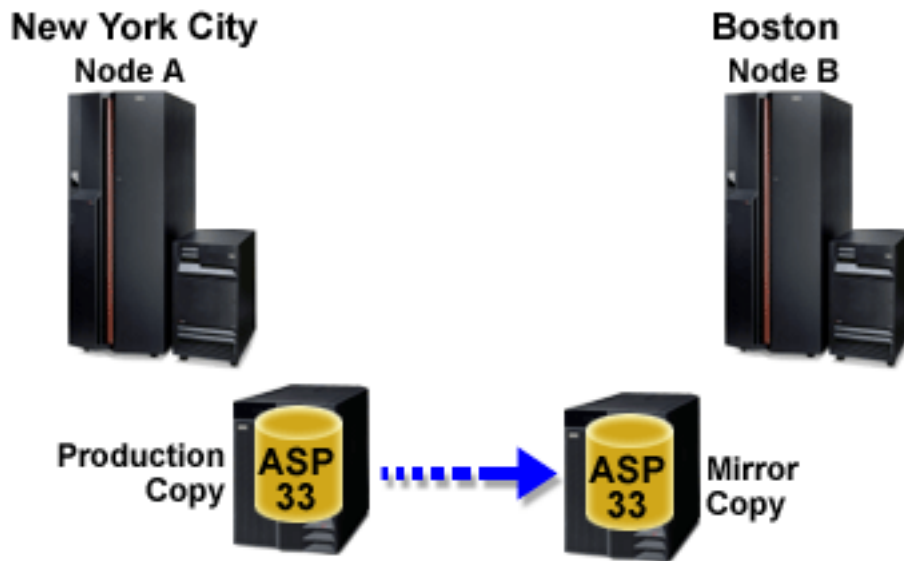
1. In iSeries Navigator, expand My Connections (or your active environment).
2. Expand your iSeries **server** → **Configuration and Service** → **Hardware** → **Disk Units**.
3. If the Service Tools Signon dialog box displays, sign on to service tools.
4. Select **Disk Pools**.
5. In the right pane 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**.
6. 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.
7. Click **Convert Disk Pool**.
8. 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.

Configure geographic mirroring with dedicated independent disk pools

To configure geographic mirroring, you must first configure your cross-site mirroring (XSM) environment and create the independent disk pool that you want to mirror. This includes defining your primary and backup nodes within the recovery domain. Before using iSeries Navigator, you should also define at least one and as many as four data port TCP/IP addresses, which would constitute one or more one-to-one bidirectional routes as part of the connection between the production copy nodes and the mirror copy nodes. Geographic mirroring will allow you to maintain an exact copy of the independent disk pool on a system at a different location for protection and availability purposes.

The following figure shows an example configuration for geographic mirroring. Primary Node A in New York City is the source system for the production copy of the independent disk pool that is dedicated to

Node A. Node B is the backup system in Boston that is the target node for the mirror copy of the independent disk pool that is dedicated to Node B.



Configured geographic mirroring

Communications requirements

Customize TCP/IP with iSeries Navigator

To configure your geographic mirroring with the iSeries Navigator, follow these steps:

1. Plan and configure your data port TCP/IP routes.
2. Create a cluster.
3. Create the independent disk pool that you want to mirror.
4. Create the device cluster resource group, also known as switchable hardware group, for the independent disk pool that you just created:
 - a. In iSeries Navigator, expand **Management Central**.
 - b. Expand **Clusters**.
 - c. Expand the cluster for which you need to add a switchable hardware group.
 - d. Right-click **Switchable Hardware**, and select **New Group**.
 - e. On the Create new or add existing disk pool dialog box, select **No, add an existing switchable disk pool to the switchable hardware group**.
5. Define your geographic mirroring sites in the recovery domain:
 - a. Right-click the newly created switchable hardware group and select **Properties**.
 - b. Select the **Recovery Domain** tab.
 - c. Select the primary node and click **Edit**.
 - d. In the site name field, specify the primary site for the production copy. Since this example contains two nodes, each of the nodes must have a different site name, with one node per site. In environments with more than two nodes, the site with more than one node must have an independent disk pool in either a hardware tower or IOP which can be switched between nodes within the same site. In this case, nodes within that site would have the same site name.

Note: The site name for the production and mirror copy can not be the same.

 - e. Click **Add** to specify the data port IP addresses of the primary node.

- f. On the Edit Node dialog, specify the data port IP addresses for the primary node that you setup in step 1, Plan and configure your TCP/IP, and click **OK**. You can configure up to four data port IP addresses in a one-to-one manner, one port to one port, each port independent of the other. You should consider configuring multiple communication lines to allow for redundancy and the highest throughput. The same number of ports used here should be configured on the mirroring node.
 - g. On the General tab, click **OK**.
 - h. Repeat the previous steps to specify the site name and IP addresses for the mirror copy node.
6. After you have completed the XSM prerequisites, follow these steps to configure geographic mirroring:
- a. In iSeries Navigator, expand **My Connections** (or your active environment).
 - b. Expand the primary node on the source iSeries server.
 - c. Expand **Configuration and Service>Hardware>Disk Units>Disk Pools**.
 - d. If the Geographic Mirroring columns are not displayed, click the Disk Pool you want to mirror, and select **View>Customize this view>Columns**, then select the columns with the suffix Geographic Mirroring from the **Columns available to display list**.
 - e. Right-click the Disk Pool you want to mirror, and select **Geographic Mirroring> Configure Geographic Mirroring**.
 - f. Follow the wizard's instructions to configure geographic mirroring.

Note: The disk pools you select to geographically mirror must be in the same switchable hardware group. If you want to geographically mirror disk pools in more than one switchable hardware group, you need to complete the wizard one time for each switchable hardware group.

7. You have now configured geographic mirroring. The remaining steps are required to prepare the independent disk pool for use in this environment. See "Print your disk configuration" on page 97. Print your disk configuration to have in case a recovery situation occurs. Also, record the relationship between the independent disk pool name and number.
 1. "Start switchable hardware group" on page 96. Start the switchable hardware group to enable device resiliency for the switchable hardware group.
 2. "Make a disk pool available" on page 112. To access the disk units in an independent disk pool you must make the disk pool available (vary on) the disk pool.
 3. Wait for resync to complete.
 4. Perform a test switchover. Before you add data to the disk pool, perform a test switchover to ensure the backup node can become the primary node and the primary node can become the backup node.

Note: If you remove a node from a device domain after you configure geographic mirroring, the removed node takes any production copies or mirror copies that it owns. These nodes are changed to nongeographic mirrored disk pools.

Using CL commands and APIs


To configure your geographic mirroring with the CL commands and APIs, follow these steps:

You can use CL commands and APIs for creating a switchable independent disk pool, however there are some tasks that require that you use iSeries Navigator.

1. Plan and configure your data port TCP/IP routes on all nodes in the recommended manner as follows:
 - Node A should have routes to C and D.
 - Node B should have routes to C and D.
 - Node C should have routes to A and B.
 - Node D should have routes to A and B.

2. **Create the cluster.** Create the cluster with the nodes that you want using the CRTCLU (Create Cluster) command.
3. **Start the nodes that comprise the cluster.** Start the nodes in the cluster using the STRCLUNOD (Start Cluster Node) command
4. **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.
5. **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.
6. **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.
7. "Create a disk pool" on page 96. Create the disk pool on the node that owns the disk units using the New Disk Pool wizard when the server is fully restarted. Make sure clustering is active before you start. Name the independent disk pool to match the device description resource name that you gave in step 3. 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 device parity sets (**iSeries Navigator required**).
8. Follow these steps to configure geographic mirroring:
 - a. In iSeries Navigator, expand **My Connections** (or your active environment).
 - b. Expand any iSeries server.
 - c. Expand **Configuration and Service**.
 - d. Expand **Hardware**.
 - e. Expand **Disk Units**.
 - f. Expand **Disk Pools**.
 - g. Right-click the Disk Pool you want to mirror and select **Geographic Mirroring > Configure Geographic Mirroring**.
 - h. Follow the wizard's instructions to configure geographic mirroring.

Note: The disk pools you select to geographically mirror must be in the same switchable hardware group. If you want to geographically mirror disk pools in more than one switchable hardware group, you will need to complete the wizard one time for each switchable hardware group.

9. **Print your disk configuration.** Print your disk configuration to have 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.
You have now configured geographic mirroring. The remaining steps are required to prepare the independent disk pool for use in this environment.

You have now configured geographic mirroring. The remaining steps are required to prepare the independent disk pool for use in this environment.

1. **Start the cluster resource group.** Start the cluster resource group to enable device resiliency using the STRCRG (Start Cluster Resource Group) command.
2. **Make the disk pool available.** To access the disk units in an independent disk pool you must vary on the disk pool using the VRYCFG (Vary Configuration) command. Vary on will also reconnect connections, so that any new route definitions can take effect.

3. **Perform a test switchover.** Before you add data to the disk pool, perform a test switchover to ensure the backup node can become the primary node and the primary node can become the backup node. Use the CHGCRGPRI (Change CRG Primary) command.

Configure geographic mirroring with switchable independent disk pools

To configure geographic mirroring you must first configure your cross-site mirroring (XSM) environment and create the independent disk pool that you want to mirror. Before using the iSeries Navigator, you should also define up to four, one-to-one data port TCP/IP routes bidirectionally as part of the connection between all the nodes in the cluster resource group. Geographic mirroring allows you to maintain an exact copy of the independent disk pool on a system at a different location for protection and availability purposes. Configuring your independent disk pool to be switchable between nodes at the same site in the cluster allows for greater availability options. See “Example: Independent disk pools with geographic mirroring” on page 131.

The following example shows geographic mirroring between sites and both sites using switchable independent disk pools. These configuration steps correlate to the graphic. You might also configure one site to contain switchable independent disk pools, while the other site uses a dedicated independent disk pool. If this is the case, change the instructions to fit your specific environment.



To configure geographic mirroring with switchable independent disk pools using the iSeries Navigator, follow these steps:

1. Plan and configure your data port TCP/IP routes. See “Communications requirements” on page 55 and Customize TCP/IP with iSeries Navigator.
2. Create a cluster containing nodes A and B.

3. “Make your hardware switchable” on page 114. If you have stand-alone expansion units or IOPs that contain 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 at the same site.
4. “Create a switchable hardware group” on page 95. 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 assists 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 might 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.

5. Add nodes C and D to the cluster and to the same device domain nodes A and B are in. This will enable independent disk pool to switching (swap roles) between nodes at both sites:
 - a. In iSeries Navigator, expand **Management Central**.
 - b. Expand **Clusters**.
 - c. Expand the cluster for which you need to add a node.
 - d. Right-click Nodes, and select **Add Node**.

Note: Clusters configured through iSeries Navigator can be made up of a maximum of four nodes. If four nodes already exist in the cluster, the **Add Node** option is disabled. If your clustering needs to extend beyond four nodes, you can use cluster resource services application programming interfaces (API) and CL commands to support up to 128 nodes. However, only four nodes are supported through the iSeries Navigator interface.

6. Add nodes C and D to the device domain:
 - a. In iSeries Navigator, expand **Management Central**.
 - b. Expand **Clusters**.
 - c. Expand the cluster containing the node you want to add to the device domain.
 - d. Click **Nodes**.
 - e. In the right pane, right-click the required node (node C) and select **Properties**.
 - f. On the **Clustering** page, in the **Device Domain** field, enter the name of the device domain that node A and node B exist in and click **OK**.

Repeat this process to add node D to the same device domain as nodes A, B, and C.

7. Add nodes C and D to the switchable hardware group:
 - a. Right-click the newly created switchable hardware group and select **Properties**.
 - b. Select the **Recovery Domain** tab.
 - c. Click **Add**.
 - d. Select the node and click **OK**. Repeat for each node.
8. Define your geographic mirroring sites in the recovery domain:
 - a. Right-click your switchable hardware group and select **Properties**.
 - b. Select the **Recovery Domain** tab.
 - c. Select the primary node and click **Edit**.
 - d. In the site name field, specify the primary site for the production copy.
 - e. Click **Add** to specify the data port IP addresses of the primary node.
 - f. On the Edit Node dialog box, specify the data port IP addresses for the primary node that you set up in step 1, Plan and configure your TCP/IP routes, and click **OK**. You can configure up to four

data port IP addresses. You should consider configuring multiple communication lines to allow for redundancy and the highest throughput. The same number of ports used here should be used on all nodes.

- g. On the General tab, click **OK**.
 - h. Repeat the previous steps to specify the site name and IP address for all other nodes in the switchable hardware group.
9. After you have completed the XSM prerequisites, follow these steps to configure geographic mirroring:
- a. In iSeries Navigator, expand **My Connections** (or your active environment).
 - b. Expand your iSeries **server** → **Configuration and Service** → **Hardware** → **Disk Units** → **Disk Pools**.
 - c. If the Geographic Mirroring columns are not displayed, click the Disk Pool you want to mirror, and select **View** → **Customize this view** → **Columns**, then select the columns with the suffix "- Geographic Mirroring" from the **Columns available to display list** .
 - d. Right-click the disk pool you want to mirror, and select **Geographic Mirroring** → **Configure Geographic Mirroring**.
 - e. Follow the wizard's instructions to configure geographic mirroring.

Note: The disk pools you select to geographically mirror must be in the same switchable hardware group. If you want to geographically mirror disk pools in more than one switchable hardware group, you will need to complete the wizard one time for each switchable hardware group.

10. "Print your disk configuration" on page 97. 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 configured geographic mirroring . The remaining steps are required to prepare the independent disk pool for use in this environment.

1. "Start switchable hardware group" on page 96. Start the switchable hardware group to enable device resiliency for the switchable hardware group.
2. "Make a disk pool available" on page 112. To access the disk units in an independent disk pool you must make the disk pool available (vary on) the disk pool.
3. Wait for resync to complete.
4. 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 that each node in the recovery domain can become the primary node.

Note: If you remove a node from a device domain after you configure geographic mirroring, the removed node takes any production copies or mirror copies that it owns. These are changed to non-geographic mirrored disk pools.

Using CL commands and APIs


To configure geographic mirroring with switchable independent disk pools using CL commands and APIs, follow these steps:

You can use CL commands and APIs for creating a switchable independent disk pool, however there are some tasks that require that you use iSeries Navigator.

1. Plan and configure your TCP/IP routes on all nodes, as follows:
 - Node A should have routes to C and D.
 - Node B should have routes to C and D.
 - Node C should have routes to A and B.
 - Node D should have routes to A and B.

2. **Create the cluster.** Create the cluster with required nodes using the CRTCLU (Create Cluster) command.
3. **Start the nodes that comprise the cluster.** Start the nodes in the cluster using the STRCLUNOD (Start Cluster Node) command
4. **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. All nodes must be in the same device domain.
5. **Create the device descriptions.** Device descriptions must be created on all nodes 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.
6. **Create the cluster resource group.** Create the device CRG with the nodes, their roles in the recovery domain, and independent disk pool device descriptions You must also specify a site name and up to four data port IP addresses for each node in the recovery domain.
7. **“Make your hardware switchable” on page 114.** If you have stand-alone expansion units or IOPs that contain 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 at the same site (**iSeries Navigator required**).
8. **“Create a disk pool” on page 96.** Create the disk pool on the node that owns the disk units using the New Disk Pool wizard when the server is fully restarted. Make sure clustering is active before you start. Name the independent disk pool to match the device description resource name that you gave in step 3. 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 device parity sets (**iSeries Navigator required**).
9. Follow these steps to configure geographic mirroring:
 - a. In iSeries Navigator, expand **My Connections** (or your active environment).
 - b. Expand the iSeries server that is the primary node .
 - c. Expand **Configuration and Service**.
 - d. Expand **Hardware**.
 - e. Expand **Disk Units**.
 - f. Expand **Disk Pools**.
 - g. Right-click the Disk Pool you want to mirror and select **Geographic Mirroring → Configure Geographic Mirroring**.
 - h. Follow the wizard’s instructions to configure geographic mirroring.

Note: The disk pools you select to geographically mirror must be in the same switchable hardware group. If you want to geographically mirror disk pools in more than one switchable hardware group, you will need to complete the wizard one time for each switchable hardware group.

Note: The mirror copy and the production copies must be in different sites. If we have two sites, AB and CD and the production copy is on node A on site AB, the backup copy must be on node C or D on site CD.
10. **Print your disk configuration.** Print your disk configuration to have 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.

You have now configured geographic mirroring. The remaining steps are required to prepare the independent disk pool for use in this environment.

1. **Start the cluster resource group.** Start the cluster resource group to enable device resiliency using the STRCRG (Start Cluster Resource Group) command.
2. **Make the disk pool available.** To access the disk units in an independent disk pool you must vary on the disk pool using the VRYCFG (Vary Configuration) command. Vary on will also reconnect connections, so that any new route definitions can take effect.
3. Wait for resync to complete.
4. **Perform a test switchover.** Before you add data to the disk pool, perform test switchovers on the switchable hardware group you created to ensure that each node in the recovery domain can become the primary node. Use the CHGCRGPRI (Change CRG Primary) command.

Protect data on disk units

To obtain optimal data protection, use iSeries Navigator to protect all the disk units on your system with either device parity protection or mirrored protection. This prevents the loss of information when a disk failure occurs. In many cases, you can keep your system running while a disk unit is being repaired or replaced.

Work with device parity protection

Device parity protection uses a data redundancy technique that protects data by spreading the parity data across multiple disk units in the parity set. When a failure occurs on a disk unit that has device parity protection, the data is reconstructed.

Related information

“Manage independent disk pools with geographic mirroring” on page 116

Find instructions to suspend and resume geographic mirror, detach and reattach the mirror copy, and delete the geographic mirroring configuration entirely.

“Set the threshold of a disk pool” on page 109

“Work with mirrored protection”

Work with mirrored protection

Mirrored protection is beneficial if you have a multibus server or a server with a large single bus. A greater number of disk units provides more opportunity for failure and increased recovery time. Mirrored protection is local to a single server and is distinct from cross-site mirroring or geographic mirroring. Mirrored protection works to prevent outage on the server by keeping a second copy of the data on a mirrored disk unit. If one disk unit fails, the server relies on the mirrored disk unit.

“Disable remote load-source mirroring” on page 48

Related information

“Manage independent disk pools with geographic mirroring” on page 116

Find instructions to suspend and resume geographic mirror, detach and reattach the mirror copy, and delete the geographic mirroring configuration entirely.

“Set the threshold of a disk pool” on page 109

“Work with device parity protection”

Start site-to-site mirroring

After you have prepared your system for mirroring, follow these steps to start remote mirroring:

1. “Enable remote load source mirroring” on page 47. This enables you to have a load source as part of the remote group of disk units.
2. Start mirroring using the start mirroring function.

When mirroring is started the system uses the resource name to recognize the remote buses and attempts to pair the disk unit on the remote buses with the disk unit on the local buses. Because remote load source mirroring is enabled, the system also pairs the load source with a remote disk unit. Mirroring restrictions that concern total disk pool capacity, an even number of disk units of each capacity, and so forth, apply.

3. On the confirmation panel for start mirroring, verify that all mirrored pairs have a level of protection of *Remote Bus*. If they do not, press F12 to cancel start mirroring, determine why some units have a lower level of protection than expected, fix the problem, and attempt to start mirroring again.

Manage your disks

Locate instructions for managing disk units and disk pools, managing independent disk pools, and keeping track of disk protection.

Manage disk units

iSeries Navigator gives you the flexibility to move disk units to other disk pools or to replace a failed disk unit in an existing disk pool. You can also rename, format, and scan disk units.

Note: Before changing the disk configuration of your server, you should have already read “iSeries Navigator requirements for disk management” on page 49. Planning is necessary to determine which procedural checklist to use and to calculate disk pool space requirements.

Replace a disk unit

If you need to replace a failed disk unit or exchange a disk unit to prevent failure, the Replace Disk Unit wizard makes the process a simple task. The disk unit to be replaced or exchanged must be running with either mirrored protection or device parity protection. To replace a mirrored disk unit, you must first suspend mirroring. A disk unit that is running with device parity protection can be exchanged only if it has failed. A disk unit running with device parity protection cannot be replaced with a nonconfigured disk even if it has failed.

To replace a failed disk unit or exchange a suspended mirrored unit, follow these steps:

1. In iSeries Navigator, expand **Disk Units**.
2. Select **All Disk Units**.
3. Right-click the disk unit that you want to replace and select **Replace Disk Unit**.
4. Follow the wizard’s instructions to replace the failed disk unit.

Rename a disk unit

iSeries Navigator gives you the option of changing the default disk unit name to something more meaningful to you. For instance, you can change Dd001 to LoadSource. You cannot specify names with spaces in them.

To rename the disk unit, follow these steps:

1. In iSeries Navigator, expand **Disk Units**.
2. Select the disk unit you want to rename.
3. Right-click the disk unit and select **Rename**.
4. Follow the instructions on the resulting dialog box.

Format a disk unit

You can select to clear all data from a nonconfigured disk unit and write the sectors, which prepares the disk unit for use in an iSeries server. Depending on disk unit capacity and performance, formatting a disk unit can take from several minutes to over an hour to complete, potentially affecting system performance.

To format a disk unit, follow these steps:

1. In iSeries Navigator, expand **Disk Units**.
2. Right-click the disk unit you want to format and select **Format**.
3. Follow the instructions on the resulting dialog.

Scan a disk unit

You can select to scan a disk unit to check the surface of the disk units and correct any sectors with errors. Depending on disk unit capacity and performance, scanning a disk unit can take anywhere from several minutes to over an hour to complete, potentially affecting system performance.

To scan a disk unit, follow these steps:

1. In iSeries Navigator, expand **Disk Units**.
2. Right-click the disk unit you want to scan and select **Scan**.
3. Follow the instructions on the resulting dialog box.

Start disk compression

Compression increases the apparent capacity of disk units by encoding the data to take up less physical storage space. Compression does affect performance because of the overhead required to compress and decompress the data. You may choose to compress data that you access infrequently or that does not require high I/O performance rates. If you want to compress a nonconfigured disk unit or a disk unit in an independent disk pool that is unavailable, you can do so when your system is fully restarted. For all other disk pools, you need to restart your server to the DST mode before compressing them.

Note: Disk compression is only capable on systems with IOAs released before V5R2.

To start disk compression, follow these steps:

1. Expand **All Disk Units**.
2. Select the disk units that you want to compress.
3. Right-click a selected disk unit and select **Start Compression**.
4. Follow the instructions on the resulting dialog box to start compression on the selected disk units.

Retrieve disk unit logs

You can use iSeries Navigator to gather information about a specific disk unit. Only newer generation disk units return meaningful logs. This function should be used under the direction of your next level of support during maintenance activities. To retrieve a disk unit log, follow these steps:

1. In iSeries Navigator, expand **My Connections**.
2. Expand any iSeries server.
3. Expand **Configuration and Service**.
4. Expand **Hardware**.
5. Expand **Disk Units**.
6. Select **All Disk Units**.
7. Right-click a specific disk unit and select **Retrieve Disk Log**.

If you want to analyze the device log, complete the following steps to package the information in a spooled file to send it electronically.

1. Start System Service Tools (STRSST), and specify the user name and password.
2. On the System Service Tools (SST) display, select option 1 (Start a service tool).
3. On the Start a Service Tool display, select option 1 (Product activity log).
4. On the Product Activity Log display, select option 1 (Analyze log).
5. On the Select Subsystem Data display, select 1 for the Log field to include all logs. Specify the date and time information in the From and To fields.
6. On the Select Analysis Report Options display, select option 3 (Print options) for the Report type field. In the Reference codes field, specify 5505.
7. On the Select Options for Printed Report display, select option 4 in the Report type field to print the full report. In the Include hexadecimal data field, select Y (Yes).

8. The device log information is stored in a spooled file, which can be electronically sent to iSeries Technical Support.

Manage disk pools

Here you can find overviews and procedures for functions that can help you manage your disk pools. The functions are available through iSeries Navigator.

Delete a disk pool

If you never need to access the data in a disk pool again, you can choose to delete the disk pool. All data on the disk units in the disk pool is destroyed. If you delete the disk pool all disk units are removed and you can no longer access the disk pool. If you want to delete an independent disk pool that is unavailable, you can do so when your system is fully restarted. For all other disk pools, you need to restart your system to DST mode before clearing or deleting them.

If you delete an independent disk pool that is participating in a clustered environment, it is recommended that you first remove the disk pool from the cluster resource group (CRG) using the Remove Cluster Resource Group Device Entry (RMVCRGDEVE) command. Under certain circumstances, you must end the CRG first; for example, if you plan to remove a subset of an independent disk pool group or remove the last independent disk pool in the CRG, use the End Cluster Resource Group (ENDCRG) command first. If you must delete the independent disk pool first, make sure you remove it from the CRG afterward.

To delete a disk pool, follow these steps:

1. In iSeries Navigator, expand **Disk Units**.
2. Expand **Disk Pools** and select the disk pools you want to delete.
3. Right-click a selected disk pool and select **Delete**.
4. Follow the instructions on the dialog box that is displayed.

Note: To delete a geographically mirrored independent disk pool, you must delete the mirror copy before the production copy.

Clear data from a disk pool

If you never need to access the data in a disk pool again, you can choose to clear the disk pool. All data on the disk units in the disk pool is destroyed, but the disk units are still available for new data storage. If you want to clear an independent disk pool that is unavailable, you can do so when your system is fully restarted. For all other disk pools, you need to restart your system to DST mode before clearing or deleting them.

To clear a disk pool, follow these steps:

1. In iSeries Navigator, expand **Disk Units**.
2. Expand **Disk Pools** and select the disk pools you want to clear.
3. Right-click a selected disk pool and select **Clear**.
4. Follow the instructions on the dialog box that is displayed.

Set the threshold of a disk pool

You can eliminate recovery problems that occur when a disk pool overflows by setting a disk pool threshold. When the data stored in the disk pool exceeds the specified threshold, the server sends a message, allowing you time to add more storage space or to delete unnecessary objects. The default threshold of a disk pool is set to 90%. You can change this value by dragging a pointer up or down the threshold scale.

To change the threshold of a disk pool, follow these steps:

1. Expand **Disk Pools**.

2. Right-click the disk pool for which you want to change the threshold and select **Properties**.
3. From the **Threshold** tab, increase or decrease the threshold for the disk pool.

Related concepts

“Device parity protection” on page 34

“Example: Independent disk pools with geographic mirroring” on page 131

Related information

“Benefits of geographic mirroring” on page 27

“Mirrored protection” on page 43

“Manage independent disk pools with geographic mirroring” on page 116

Find instructions to suspend and resume geographic mirror, detach and reattach the mirror copy, and delete the geographic mirroring configuration entirely.

“Work with device parity protection” on page 106

“Work with mirrored protection” on page 106

Enable automatic overflow recovery

If a basic disk pool fills up and the data in the basic disk pool overflows into the system disk pool, the basic disk pool is said to be in an overflow state. If you enable automatic overflow recovery, you can recover the overflow data to the system disk pool by restarting your system. If you have created sufficient space in the basic disk pool, the system will copy the overflow data to the disk pool from the system disk pool. If automatic overflow recovery is disabled, you need to manually restart your server to the dedicated service tools (DST) mode and recover the overflow data using DST from the command prompt.

To enable automatic overflow recovery, follow these steps:

1. Expand **Disk Pools** and select the disk pools on which you want to enable automatic overflow recovery.
2. Right-click a selected disk pool and select **Enable Overflow Recovery**.
3. Follow the instructions on the resulting dialog box.

Balance a disk pool

You can improve the server performance by ensuring that the disk units in the disk pool have equal percentages of data residing on them. The capacity balancing function ensures that the disk units in the disk pool are balanced. You can balance a disk pool when you use the Add Disk Unit or New Disk Pool wizards. To customize your system with usage balancing or hierarchical storage management (HSM)

balancing, consult Backup and Recovery .

You can balance the capacity of disk pools using the “Add a disk unit or disk pool” on page 91 wizard.


Manage 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.

Backup and recovery of 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; see Save independent ASPs for more

information. 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 re-associated. See the following procedures for steps to recover the disk pools:

Recover an independent disk pool: 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 nonconfigured.

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 any iSeries server.
3. Expand **Configuration and Service**.
4. Expand **Hardware**.
5. Expand **Disk Units**.
6. If the Service Tools Signon dialog box displays, sign on to service tools.
7. Select **Disk Pools**.
8. 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.
9. Follow the instructions on the dialog box displayed.

Recover 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 re-associated 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 any iSeries server.
3. Expand **Configuration and Service**.
4. Expand **Hardware**.
5. Expand **Disk Units**.
6. If the Service Tools Signon dialog box displays, sign on to service tools.

7. Select **Disk Pools**.
8. Verify that the required primary disk pool exists. If it does not, you must “Create a disk pool” on page 96.
9. 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**.
10. 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.
11. Click **Recover Group**.

Make 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 Perform 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 activity, 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 primary node iSeries server.
3. Expand **Configuration and Service**.
4. Expand **Hardware**.
5. Expand **Disk Units**.
6. Sign on to service tools if the Service Tools Signon dialog box displays.
7. Expand **Disk Pools**.
8. Right-click the unavailable disk pool and select **Make Available**. You can select multiple disk pools to make available at the same time.
9. 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.

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

Make a disk pool unavailable

You can select an independent disk pool to make it unavailable (vary it off).

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 pool" 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 iSeries server that is the primary node .
3. Expand **Configuration and Service**.
4. Expand **Hardware**.
5. Expand **Disk Units**.
6. Sign on to service tools if the Service Tools Signon dialog box is displayed.
7. Expand **Disk Pools**.
8. Right-click the disk pool you want to make unavailable and select **Make Unavailable**.
9. 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.

Release 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 any iSeries server.
3. Expand **Configuration and Service** → **Hardware** → **Disk Units** → **Disk Pools**.
4. 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.

Make 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 input/output processors (IOPs) and their associated resources. Refer to the following steps that apply to your switchable environment.

Make a frame and unit switchable

An independent disk pool can contain disk units within several expansion units (frame/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 servers. This is called making an expansion unit switchable. If you do not want other servers 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 any iSeries server.
3. Expand **Configuration and Service** → **Hardware** → **Disk UnitsBy Location** and select the frame/units you want to make switchable.
4. Right-click a highlighted frame/unit and select **Make Switchable**.
5. Follow the instructions on the dialog box that is displayed.

Change a bus ownership type

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**.
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**.

Create an I/O Pool using the Hardware Management Console

If you are using the Hardware Management Console to manage your logical partitions, you must create an I/O pool that includes the IOP, input/output adapter (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**.
6. Click **OK**.

For the I/O pool changes to take effect, you must 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.

Switch access to backup server

Perform a cluster switchover when you want a backup server 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. Current 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. See Perform a switchover for details.

Synchronize 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 servers 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 server and then made available on a different server, 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.

Change the server takeover IP address

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

The server takeover IP address is associated with a primary disk pool in a clustered, switchable environment. Specifically, it is the IP address for a server 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 server takeover IP address for a primary disk pool, follow these steps:

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

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

6. Change the server 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 server takeover IP address.

Manage independent disk pools with geographic mirroring

Find instructions to suspend and resume geographic mirror, detach and reattach the mirror copy, and delete the geographic mirroring configuration entirely.

After you have configured geographic mirroring, you can perform management tasks using iSeries Navigator.

There may be additional tasks that you might need to perform include.

Related information

“Set the threshold of a disk pool” on page 109

“Work with device parity protection” on page 106

“Work with mirrored protection” on page 106

Suspend geographic mirroring: You can choose to temporarily stop geographic mirroring by selecting to suspend geographic mirroring. Any changes made on the production copy of the independent disk pool are not be transmitted to the mirror copy.

Note: When you resume geographic mirroring, synchronization is required between the production and mirror copies. If geographic mirroring was suspended without tracking, then full synchronization is required. This can be a lengthy process.

To suspend geographic mirroring, follow these steps:

1. In iSeries Navigator, expand **My Connections** (or your active environment).
2. Expand the server that owns the production copy of the geographically mirrored disk pool that you want to suspend.
3. Expand **Configuration and Service**.
4. Expand **Hardware**.
5. Expand **Disk Units**.
6. Expand **Disk Pools**.
7. Right-click the production copy of the **Disk Pool** you want to suspend and select **Geographic Mirroring** → **Suspend Geographic Mirroring**.

| **Suspend with tracking**

| If you suspend with tracking, the system will attempt to track changes made to those disk pools. This may shorten the length of the synchronization process by performing partial synchronization when you resume geographic mirroring. If tracking space is exhausted, then when you resume geographic mirroring, complete synchronization is required.

| **Suspend without tracking**

| If you suspend without tracking, then when you resume geographic mirroring, complete synchronization occurs.

| **Note:** If you suspend geographic mirroring without tracking changes, then when you resume geographic mirroring, a complete synchronization is required between the production and mirror copies. If you suspend geographic mirroring and you do track changes, then only a partial synchronization is required. Complete synchronization can be a very lengthy process, anywhere from one to several hours or longer. The length of time it takes to synchronize is dependent on the number and type of disk units as well as how many TCP/IP communication interfaces are dedicated to geographic mirroring.

| Use the Start DASD Management Operation (QYASSDMO) API to reduce the amount of time it takes to make a disk pool unavailable.

Resume geographic mirroring: If you suspend “Geographic mirroring” on page 26, you must resume geographic mirroring in order to reactivate mirroring between the production and mirrored copies again.

Note: When you resume geographic mirroring, the production and mirror copies are synchronized concurrent with performing geographic mirroring. Synchronization can be a lengthy process. If a disk pool becoming unavailable interrupts synchronization, then synchronization will continue from where it was interrupted when the disk pool becomes available again. When an interrupted synchronization is continued, the first message (CPI0985D) states that the synchronization is 0% complete.

To resume geographic mirroring, follow these steps:

1. In iSeries Navigator, expand **My Connections** (or your active environment).
2. Expand the server that owns the production copy of the disk pool for which you want to resume geographic mirroring.
3. Expand **Configuration and Service**.
4. Expand **Hardware**.
5. Expand **Disk Units**.
6. Expand **Disk Pools**.
7. Right-click the **Disk Pool** you want to resume and select **Geographic Mirroring** → **Resume Geographic Mirroring**.

| Use the Start DASD Management Operation (QYASSDMO) API to reduce the amount of time it takes to
| make a disk pool unavailable.

Detach mirror copy: If you are using geographic mirroring and want to access the mirror copy to perform save operations or data mining, or to create reports, you must detach the mirror copy from the production copy. You detach the mirror copy by accessing the production copy of the disk pool.

Note: When you reattach the detached mirror copy, a complete synchronization is required between the production and the mirror copies. Synchronization can be a lengthy process.

To detach the mirror copy, follow these steps:

1. It is recommended that you make the independent disk pool unavailable to ensure the production copy is not altered while the detach is being performed. See *Make the independent disk pool unavailable*.
2. In iSeries Navigator, expand **My Connections** (or your active environment).
3. Expand the server that owns the production copy of the disk pool from which you want to detach the mirror copy.
4. Expand **Configuration and Service** → **Hardware** → **Disk Units** → **Disk Pools**.
5. Right-click the production copy of the **Disk Pool** you want to detach and select **Geographic Mirroring** → **Detach Mirror Copy**.
6. If **Geographic Mirroring** → **Detach Mirror Copy** cannot be clicked, because it is disabled. The mirror copy is not in sync with the production copy, so geographic mirroring must be resumed, the disk pool varied on, and production and mirror copies synchronized before the mirror copy can be detached.

Before you make the detached mirror copy available, you should create a second, unique device description for the independent disk pool that differentiates it from the production copy. A separate device description for the mirror copy prevents two instances of the same database in the network. It will also simplify work done outside of iSeries Navigator. Use the detached mirror copy device description to make the detached mirror copy available.

Reattach mirror copy: If you detached the mirror copy and have completed your work with the detached mirror copy, you must reattach the detached mirror copy in order to resume using geographic mirroring. You reattach the detached mirror copy by accessing the production copy of the disk pool. The detached mirror copy must be unavailable when you reattach it to the production copy.

Note: When you reattach the detached mirror copy, a complete synchronization is required between the production copy and the mirror copy. Synchronization can be a lengthy process.

To reattach the mirror copy, follow these steps:

1. In iSeries Navigator, expand **My Connections** (or your active environment).

2. Expand the server that owns the production copy of the disk pool to which you want to reattach the detached mirror copy.
3. Expand **Configuration and Service**.
4. Expand **Hardware**.
5. Expand **Disk Units**.
6. Expand **Disk Pools**.
7. Right-click the production copy of the **Disk Pool** you want to reattach and select **Geographic Mirroring** → **Reattach Mirror Copy**.

Change disk pool attributes: You can change the geographic mirroring attributes for a disk pool from the production copy when the disk pool is unavailable. The values specified for the primary disk pool for synchronous or asynchronous performance mode and for recovery timeout are used for each disk pool in the disk pool group.

For more information about geographic mirroring attributes, see [How geographic mirroring works](#).

To edit the disk pool attributes, follow these steps:

1. In iSeries Navigator, expand **My Connections** (or your active environment).
2. Expand the iSeries server that owns the production copy of the geographically mirrored disk pool for which you want to edit the attributes.
3. Expand **Configuration and Service** → **Hardware** → **Disk Units** **Disk Pools**.
4. Right-click the production copy of the **Disk Pool** for which you want to edit the attributes and select **Geographic Mirroring** → **Change Attributes** .

Geographic mirroring attributes

You can change the geographic mirroring attributes for a disk pool from the production copy when the disk pool is unavailable. The values specified for the primary disk pool for synchronous or asynchronous performance mode and for recover timeout are used for each disk pool in the disk pool group.

Deconfigure geographic mirroring: If you no longer want the capability to use geographic mirroring for a specific disk pool or disk pool group, you can select to **Deconfigure Geographic Mirroring**. If you deconfigure geographic mirroring, the system stops geographic mirroring and deletes the mirror copy of the disk pools on the nodes in the mirror copy site. The disk pool must be offline to deconfigure geographic mirroring.

To deconfigure geographic mirroring, follow these steps:

1. In iSeries Navigator, expand **My Connections** (or your active environment).
2. Expand the iSeries server that owns the production copy of the disk pool for which you want to deconfigure geographic mirroring.
3. Expand **Configuration and Service**.
4. Expand **Hardware**.
5. Expand **Disk Units**.
6. Expand **Disk Pools**.
7. Right-click the production copy of the **Disk Pool** you want to deconfigure and select **Geographic Mirroring** → **Deconfigure Geographic Mirroring**.

To complete the process, update your cluster configuration, as follows:

- Remove the nodes associated with the mirror copy from the switchable hardware group recovery domain.
- Remove the site name and data port IP addresses from the remaining nodes in the cluster.

Messages for geographic mirroring: Geographic mirroring message descriptions and recoveries.

0x00010259

Description: Operation failed because the system did not find the mirror copy.

Recovery: Not all the nodes in the device domain responded. Make sure that clustering is active. If necessary, start clustering (STRCLUNOD). Try the request again. If the problem persists, contact your technical support provider.

0x0001025A

Description: Not all of the disk pools in the disk pool group are geographically mirrored

Recovery: If one disk pool in a disk pool group is geographically mirrored, all of the disk pools in the disk pool group must be geographically mirrored. Take one of the following actions: 1) Configure geographic mirroring for the disk pools which are not geographically mirrored. 2) Deconfigure geographic mirroring for the disk pools that are geographically mirrored.

0x00010265

Description: The detached mirrored copy is available.

Recovery: Make the detached mirrored copy unavailable and then try the reattach operation again.

0x00010380

Description: A disk unit is missing from the configuration of the mirror copy.

Recovery: Find or fix the missing disk unit in the mirror copy. Check the Product Activity Log (PAL) on destination node. Reclaim IOP cache storage.

0x00011210

Description: The proposed secondary disk pool for the disk pool group is not geographically mirrored

Recovery: If one disk pool in a disk pool group is geographically mirrored, all of the disk pools in the disk pool group must be geographically mirrored. You must configure geographic mirroring for the proposed secondary disk pool which is not geographically mirrored, either now or after completing this operation.

0x00011211

Description: Duplicate mirror copies exist.

Recovery: Check for locally mirrored disk units that may exist on two systems, Enterprise Storage Server[®] FlashCopy[®], or back level independent disk pool copies. See the Product Activity Log (PAL) on the mirror copy node for more information. Eliminate duplication and try the request again. If the problem persists, contact your technical support provider, or see iSeries and AS/400[®] Technical Support for information about IBM support and services.

Work with device parity protection

Device parity protection uses a data redundancy technique that protects data by spreading the parity data across multiple disk units in the parity set. When a failure occurs on a disk unit that has device parity protection, the data is reconstructed.

Related information

“Manage independent disk pools with geographic mirroring” on page 116

Find instructions to suspend and resume geographic mirror, detach and reattach the mirror copy, and delete the geographic mirroring configuration entirely.

“Set the threshold of a disk pool” on page 109

“Work with mirrored protection” on page 106

Change parity set optimization

If you use a V5R2 input/output adapter (IOA) and OS/400 V5R2 or later, you can now choose how you want your parity sets to be optimized. When you select to optimize a parity set, the I/O adapter will choose disk units for parity sets according to the optimization value you have chosen. Depending on your configuration, different parity set optimizations might generate the same parity sets. You have several options for parity set optimization:

Availability: A parity set optimized for availability offers a greater level of protection because it allows a parity set to remain functional in the event of a I/O bus failure. The availability optimization value ensures that a parity set is formed from at least three disk units of equal capacity each attached to a separate bus on the input/output adapter (IOA). For example, if an I/O adapter had 15 disk units and was optimized for availability, the result might be five parity sets with three disk units each attached to separate I/O buses on the adapter. OS/400 V5R3 is required to optimize for availability.

Capacity: A parity set optimized for capacity stores the most data possible. The I/O adapter may generate fewer parity sets with more disk units in each parity set. For example, if an I/O adapter has 15 disk units and is optimized for capacity, the result might be one parity set containing 15 disk units.

Balanced: A balanced parity set compromises between the ability to store large amounts of data and also provide fast access to data. For example, if an I/O adapter has 15 disk units and you choose balanced parity optimization, the result might be two parity sets, one with nine disk units and one with six disk units.

Performance: Parity sets optimized for performance provide the fastest data access. The I/O adapter may generate more parity sets with fewer numbers of disk units. For example, if an I/O adapter had 15 disk units and is optimized for performance, the result might be three parity sets with five disk units each.

Steps to change parity set optimization: To change the parity set optimization for all new parity sets that are created, use the following steps. This change stays in effect until you change it again. If you need to start parity, you can also change the parity set optimization as part of the start parity process.

1. Expand **Disk Units**.
2. Right-click **Parity Sets** and select **Change Optimization**.

Note: RAID 6 protection gives you optimal performance, capacity, and balance, so selecting any of the parity set optimizations does not affect the outcome of the parity set.

| Determine what disks are in a parity set using the DST menu

- | Follow these steps to find the disk units in a parity set using the DST menu.
- | 1. Select **Work with disk units** on the Use Dedicated Service Tools (DST) menu.
- | 2. Select **Work with disk configuration** on the Work with Disk Units display.
- | 3. Select **Display disk configuration** on the Work with Disk Configuration display.
- | 4. Select **Display device parity status** on the Display Disk Configuration display.

| Determine what disks are in a parity set using the SST menu

- | 1. Select **Work with disk units** on the Use System Service Tools (SST) menu.
- | 2. Select **Display disk configuration** from the Work with Disk Configuration display.
- | 3. Select **Display device parity status** on the Display Disk Configuration display.

| Determine what disks are in a parity set using iSeries Navigator

- | 1. In the iSeries Navigator display, click the plus sign next to your system.
- | 2. Click on the plus sign next to Configuration and Service.
- | 3. Click on the plus sign next to Hardware.

- | 4. Click on the plus sign next to **Disk Units**.
- | 5. Log into Service Tools Click on **Parity Sets**.
- | 6. Click on the plus sign of each parity set to see the list of disk units contained in that set.

Start device parity protection

The best time to start device parity protection is when you add new or nonconfigured disk units. The “Add a disk unit or disk pool” on page 91 has steps for including disk units in a parity set and starting device parity protection. It is also possible to start device parity protection at a later time.

RAID 5 parity set

- Systems with IOAs released after V5R2 hold a minimum number of 3 disk units in a parity set. The maximum number of disk units in a parity set is 18.

Note: For systems with IOAs released after OS/400 V5R2, the minimum number of disk units in a parity set is 4. The maximum number of disk units in a parity set is 10.

- All devices in a parity set must be the same capacity.

RAID 6 parity set

The minimum number of disk units in a parity set is 4. The maximum number of disk units in a parity set is 18.

To learn more about how device parity protection is implemented, see How device parity protection works. Examples: Device parity and mirrored protection shows some examples of how device parity protection can be used in conjunction with mirrored protection.

Start device parity protection

1. In iSeries Navigator, expand **Disk Units**.
2. Select the disk units for which you want to start device parity protection.
3. Right-click a selected disk unit and select **Start Parity**.
- | 4. Select the level of RAID protection you want.
5. From the resulting window, click **Start Parity** to start device parity protection on the displayed disk units.

Stop device parity protection

You can select to stop device parity protection on the displayed disk units. The list displays all of the disk units in the parity set. When preparing to stop device parity protection, the system performs validity checking to make sure that stopping device parity protection does not leave the system in a configuration that is not supported. Depending on disk unit capacity and performance, stopping device parity protection can take from several minutes to over an hour to complete, potentially affecting system performance.

You cannot stop device parity protection on a disk unit that is in a mirrored disk pool. To stop device parity protection, you must first “Work with mirrored protection” on page 106.

To stop device parity protection on the disk units in a parity set, follow these steps:

1. In iSeries Navigator, expand **Disk Units**.
2. Select the disk units for which you want to stop device parity protection.
3. Right-click a selected disk unit and select **Stop Parity**.
4. From the resulting dialog box, click **Stop Parity** to stop device parity protection.

Include disk units in a parity set

You can select which disk units you want to include in a parity set. When you attach a new disk unit to an existing I/O processor that has device parity protection, you can include the disk unit in a device parity set with other disk units of similar capacity.

If you want to include a disk unit in an independent disk pool that is unavailable, you must first IPL your system. For all other disk pools, you need to restart your system to dedicated service tools (DST) mode before including them in a parity set.

To include a disk unit in a parity set, follow these steps:

1. Expand **Disk Units**.
2. Select the disk units that you want to include.
3. Right-click a selected disk unit and select **Include in Parity Set**.
4. From the resulting dialog box, click **Include** to include the selected disk units in a parity set.

Adding three or more disk units requires you to create a new parity set. The include function does not work when creating a parity set. To create a parity set, go to Start device parity protection.

Exclude disk units from a parity set

You can select which disk units you want to exclude from the parity set as long as they do not contain parity data.

- | For RAID 5 protection, you can exclude a protection disk unit with a model number of 070, or 080 if it is compressed, because these disk units do not store parity data.
- | For RAID 6 protection, you can exclude a protected disk unit with a model number of 090 because it is a disk unit that does not store parity data.

When you exclude a disk unit from the parity set, the data on this disk unit remains there but is no longer protected by device parity protection. If the disk pool is protected, you are not allowed to exclude a disk unit that belongs to that disk pool from a parity set. The system does not allow unprotected disk units to reside in a protected disk pool.

If you want to exclude disk units from an independent disk pool that is unavailable, you can do so when your system is fully restarted. For all other disk pools, you need to restart your system to dedicated service tools (DST) mode before excluding them from a parity set.

- | **Note:** Not all disk units in a parity protected set are eligible to be excluded. To be eligible, the parity set must contain at least four devices with RAID 5 protection and at least five devices for RAID 6 protection, and the candidate devices cannot contain parity data.

To exclude a disk unit from a parity set, follow these steps:

1. Expand **Disk Units**.
2. Select the disk units that you want to exclude.
3. Right-click a selected disk unit and select **Exclude from Parity Set**.
4. From the resulting dialog box, click **Exclude** to exclude the disk units from a parity set.

Work with mirrored protection

Mirrored protection is beneficial if you have a multibus server or a server with a large single bus. A greater number of disk units provides more opportunity for failure and increased recovery time. Mirrored protection is local to a single server and is distinct from cross-site mirroring or geographic mirroring. Mirrored protection works to prevent outage on the server by keeping a second copy of the data on a mirrored disk unit. If one disk unit fails, the server relies on the mirrored disk unit.

“Disable remote load-source mirroring” on page 48

Related information

“Manage independent disk pools with geographic mirroring” on page 116

Find instructions to suspend and resume geographic mirror, detach and reattach the mirror copy, and delete the geographic mirroring configuration entirely.

“Set the threshold of a disk pool” on page 109

“Work with device parity protection” on page 106

Start mirrored protection

The Add Disk Unit and New Disk Pool wizards guide you through the process of adding pairs of similar capacity disk units to a protected disk pool. When you have your disks configured correctly, you are ready to start mirroring for mirrored protection. Mirrored protection is local to a single server and is distinct from cross-site mirroring or geographic mirroring. If you want to start mirroring on an “Independent disk pools” on page 12 that is unavailable, you can do so when your system is fully restarted. For all other disk pools, you need to restart your system to the dedicated service tools (DST) mode before starting mirrored protection.

To start mirroring, follow these steps:

1. In iSeries Navigator, expand **Disk Units**.
2. Expand **Disk Pools**.
3. Right-click the disk pools you want to mirror, and select **Start Mirroring**.

Stop mirrored protection

When you stop mirrored protection, one unit from each mirrored pair is unconfigured. Before you can stop mirrored protection for a disk pool, at least one unit in each mirrored pair in that disk pool must be present and active. To control which mirrored unit of each pair is unconfigured, you may suspend the storage units that you want to become unconfigured. For units that are not suspended, the selection is automatic.

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

Mirrored protection is local to a single server and is distinct from cross-site mirroring or geographic mirroring.

To stop mirrored protection, follow these steps:

1. In iSeries Navigator, expand **Disk Units**.
2. Expand **Disk Pools**.
3. Select the disk pools you want to mirror.
4. Right-click any selected disk pool, and select **Stop Mirroring**.
5. Click **Stop Mirroring** from the resulting confirmation dialog box.

Suspend mirrored protection

If a disk unit in a mirrored pair fails, you need to suspend mirroring to repair or replace it. If you want to suspend mirroring on an independent disk pool that is unavailable, you can do so when your system is fully restarted. For all other disk pools, you need to restart your system to the dedicated service tools (DST) mode before suspending mirroring.

Mirrored protection is local to a single server and is distinct from cross-site mirroring or geographic mirroring.

To suspend mirrored protection, follow these steps:

1. In iSeries Navigator, expand **Disk Units**.
2. Double-click **All Disk Units**.
3. Select the disk unit for which you want to suspend mirrored protection.
4. Right-click the selected disk unit and select **Suspend Mirroring**.

Enable remote load source mirroring

Enabling remote load source mirroring makes it possible for the two disk units of the load source mirrored pair to be on different I/O processors or system buses. Remote load source mirroring allows you to protect against a site disaster by dividing the disk storage between the two sites, mirroring one site to another. You must enable remote load source mirroring before starting mirrored protection for disk pool 1. If remote load source mirroring support is enabled after mirrored protection has already been started for disk pool 1, the existing mirrored protection and mirrored pairing of the load source is not changed.

Remote load source mirroring support can be enabled in either the DST or the SST environment in iSeries Navigator or the character-based interface. If you attempt to enable remote load source mirroring and it is currently enabled, the system displays a message that remote load source mirroring is already enabled.

To enable remote load source mirroring, follow these steps:

1. In iSeries Navigator, expand **Disk Units** → **Disk Pools** → **Disk Pool 1**.
2. Right-click the load source disk unit and select **Enable Remote Load Source Mirroring**.

Note: Enabling remote load source mirroring does not start mirrored protection on the disk units. Remote load source mirroring affects only the load source disk units.

To enable remote load source mirroring using the character-based interface, do the following:

1. From the DST Main Menu, select option 4, Work with disk units.
2. From the Work with disk units menu, select option 1, Work with disk configuration.
3. From the Work with disk configuration menu, select option 4, Work with mirrored protection.
4. From the Work with mirrored protection menu, select option 4, Enable remote load source mirroring. This will display an Enable remote load source mirroring confirmation screen.
5. Press Enter at the Enable remote load source mirroring confirmation screen. The Work with mirrored protection screen will be displayed, with a message at the bottom, indicating that remote load source mirroring has been enabled.

Using independent disk pools

If you are interested solely in independent disk pools, use this information to help you plan, configure, and manage independent disk pools.

This topic will provide you with the information you need to use independent disk pools, from a conceptual explanation to planning, configuring, and managing independent disk pools on your servers.

“Plan for independent disk pools” on page 51

“Configure independent disk pools” on page 92

“Manage independent disk pools” on page 110


Examples: Independent disk pool configurations

Independent disk pools can be switchable among a group of servers 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 server, independent of the rest of the storage on the server.

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.



Examples: Switchable independent disk pools

In a multisystem environment, an independent disk pool can be switched between servers in a cluster. A switchable independent disk pool is a set of disk units that you can switch between servers so that each server can access the data. Only one system can access the data at a time.

Switchable independent disk pools can reside on one of two types of switchable hardware devices:

External expansion unit

The switchable device can be an external expansion unit connected to the clustered servers on the same high-speed link (HSL) loop.

Input/output processor (IOP) in a logical partition

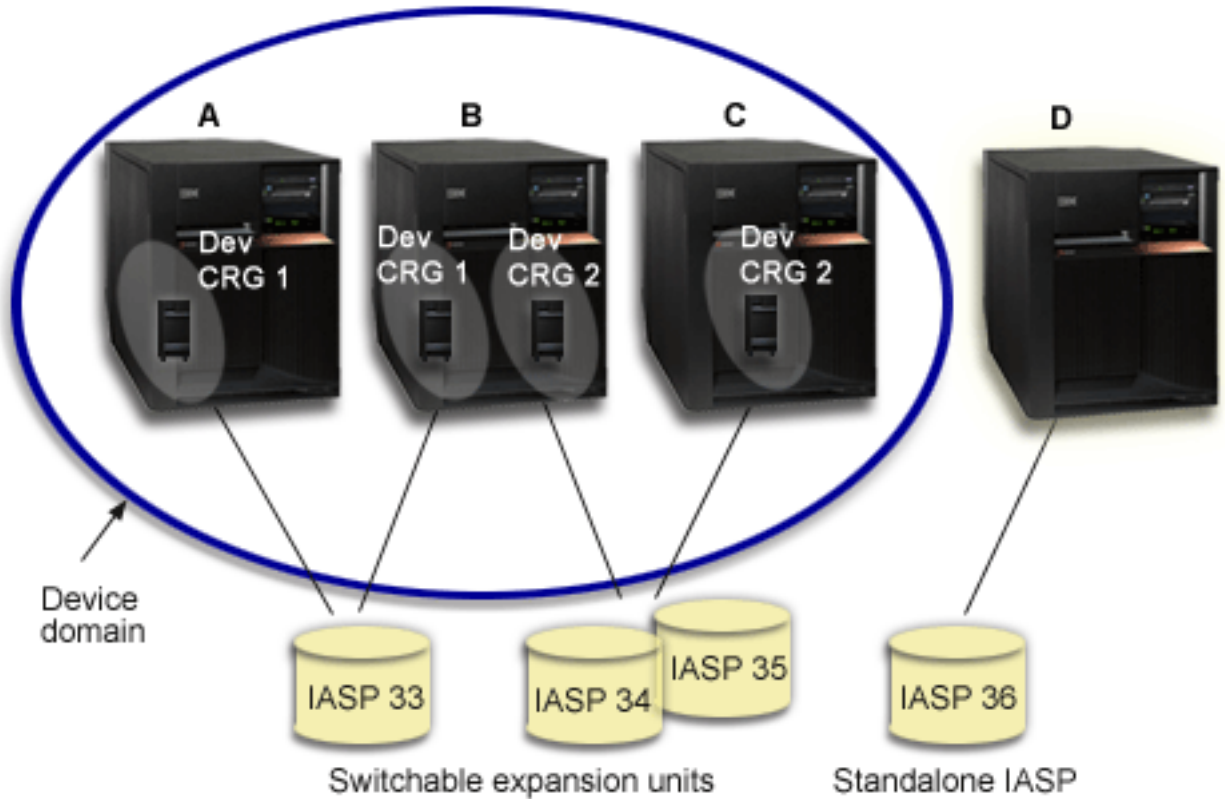
In an LPAR environment, the switchable device can be an IOP on the bus shared by the partitions or in an I/O pool.

The entity that switches is actually the expansion unit or the IOP containing the independent disk pool. When an expansion unit or IOP is switched, all of the hardware attached to the switchable entity is moved to the backup system.

The following example configurations and scenario illustrate some typical switchable independent disk pools implementations:


Example: Switchable expansion unit:

In this example, the following figure shows a cluster consisting of four nodes. Nodes named A, B, and C are defined to be in the same device domain. There are two switchable expansion units - one contains IASP33 and the other contains IASP34 and IASP35. The expansion unit containing IASP33 is on an HSL loop that also contains nodes A and B. This first expansion unit can be switched between nodes A and B. The expansion unit containing IASP34 and IASP35 can be on another HSL loop that also contains nodes B and C. This second expansion unit can be switched between nodes B and C. Node D is contained in the cluster, but is not a member of the device domain and therefore can only access IASP36, a stand-alone (or dedicated) independent disk pool.




Example: Switchable IOP with logical partitions:

In this logical partition example, the following figure shows a cluster consisting of four logical partitions on a single iSeries server. All four nodes belong to the same device domain. IASP36 is composed of disk units accessible through IOP Y. IOP Y is on the shared bus so it can be switched between all of the nodes in the cluster: A, B, C, and D. When the IOP is switched, everything that is physically connected to that IOP is also moved to the new primary node.



Example: Switchable expansion unit with logical partitions:

The example, shown in the figure below, depicts a combination of the previous two examples. IASP36 is composed of disk units contained in a switchable expansion unit. The expansion unit is on the same HSL loop as two systems, one of which is made up of four logical partitions. Assume that nodes C and D and the second server, node E, are defined to be in the same device domain, and that the independent disk pool can be switched between those three nodes.



Scenario: Consolidate servers using switchable independent disk pools:

Situation

Your company's network currently uses 30 small servers distributed within a single region, all in the same time zone, using the same language, and running the same release of the operating system and programming code. The amount of time and effort you spend maintaining the small systems and keeping them at the same operating system and application release levels is significant.

Objectives

To reduce the resource required to maintain and administer your servers, you want to consolidate by reducing the number of servers in your network.

The objectives of this scenario are as follows:

- To consolidate from 30 small servers to one larger server at a central location
- To maintain data independence for each geographic region

Details

None of the 30 small servers in your network require more than four disk units.

Prerequisites and assumptions

A potential consolidation answer for your network is logical partitioning (LPAR). However, in your scenario, consolidating the 30 locations with logical partitioning is not ideal because:

- The effort required to manage the partitions is approximately the same as managing 30 distributed systems.
- Each partition requires an IOP in order to support a load source for the partition. As a result, 30 IOPs are required for the consolidated system.
- Additional expansion units are required to hold the IOPs needed for the 30 partitions. Since each location uses only a few disk units, the expansion units might be nearly empty.

As a result, the LPAR solution is not justifiable from an economic point of view for your scenario.

A better way to solve your particular scenario is to use switchable independent disk pools to provide server consolidation. By creating one switchable independent disk pool for each of the 30 branch offices, you will be able to reduce the number of IOPs from 30 to 7, while requiring just two expansion units. This is an economically attractive alternative.

Design

To understand how to use switchable independent disk pools, see “Create a switchable independent disk pool” on page 93. In addition to the planning and configuration steps for implementing switchable independent disk pools, the following strategies can be used to ensure that your users at the respective branch offices can seamlessly access data:

- To ensure that users receive access to the correct set of data, your run-time environment can be changed to make sure that users from different branch offices connect to their data in the corresponding independent disk pool. This can be accomplished through a simple adjustment to user profiles and to the job descriptions that are specified by user profiles.

All user profiles from a particular branch office will use one job description. The job description will specify the independent disk pool that contains the user’s data, and create the library list that each job will use. With these simple changes, the task of getting each user to the correct set of data is completed.

- Another run-time problem to be pointed out is the resolution of duplicate subsystems and job queues. Each branch office uses a cloned subsystem description to run batch jobs. Each of the subsystems uses job queues that have the same name on each of the branch office subsystems. If a single subsystem and a single set of job queues are used in the consolidated environment, jobs submitted by users from different branch offices will all be placed on the same set of queues and initiated by a single subsystem. This results in work flow that is inconsistent with the run-time environment of the distributed systems.

To resolve this problem, the subsystems will be given unique names. Then, a command to start all of the subsystems will be added to the startup program. Finally, each of the job queues used by the subsystem will be moved into a library that is unique to each of the job descriptions that are used by the branch offices. As a result, any application that submits a job will require no changes in order to submit batch jobs to its unique queue.

Example: Independent disk pools with geographic mirroring:

The following example shows one way that geographic mirroring can be configured. Node A and Node B are located in New York City. Node C and Node D are located in Boston. All four nodes are configured in the same recovery domain. The production copy can be switched between nodes A and B. The mirror copy can be switched between nodes C and D. Because all of the nodes are in the same recovery domain, the source system in New York can also exchange roles with the target system in Boston, allowing Boston to host the production copy.



This company has defined the following roles for the nodes in the recovery domain:

Node	Role
Node A	Primary
Node B	Backup 1
Node C	Backup 2
Node D	Backup 3

In the event of a natural disaster in New York, Node C in Boston becomes the primary node by upgrading its mirror copy to a production copy. Node C becomes the source system for geographic mirroring, although geographic mirroring will be suspended because there is no target node because of the natural disaster in New York. When the New York site recovers, Node A becomes a backup node and its previous production copy becomes the mirror copy.

Related concepts

“Device parity protection” on page 34

Related information

“Benefits of geographic mirroring” on page 27

“Mirrored protection” on page 43

“Set the threshold of a disk pool” on page 109

Frequently asked questions

Here is a list of independent disk pool questions and answers. If you have a question that is not on this page, please contact us.

General

1. How do independent disk pools work?
2. How can independent disk pools be implemented in my environment?
3. How should I structure my independent disk pools?
4. What is a disk pool group?
5. What is geographic mirroring?

iSeries Navigator graphical user interface

1. How do I access the iSeries Navigator disk management function?
2. What is the difference between the disk management functions in iSeries Navigator and the character-based command interface?
3. How do I access the disk management function when the system is in dedicated service tools (DST) mode?
4. What is the service tools server?
5. Why does the data I see in iSeries Navigator appear to be out of date?
6. Why can't I connect to the service tools server after I add the service table entry?

Configuring

1. How do I create a new disk pool or independent disk pool?
2. How do I create a disk pool group?
3. How do I configure geographic mirroring?

Performance

1. Why is performance slow?
2. Why is performance slow for geographic mirroring?

Troubleshooting

1. Why do no disk units appear as eligible to be added to my disk pool?
2. Why doesn't the device description get deleted when I delete the disk pool?
3. Why do I get a message saying the device description is already created?
4. Why does the primary or secondary disk pool I try to create appear to be a UDFS disk pool?
5. Why do I get a message that says my disk pool is not the right type when I try to create a library in it?
6. What is a CPDB716 message, and how can it be fixed?

General

How do independent disk pools work?

The key characteristic of an independent disk pool is its ability to be, of course, independent of the rest of the storage on a server. It is independent because the data in the independent disk pool is self-contained.

This means that all of the necessary system information associated with the data resides within the independent disk pool. See “How independent disk pools work” on page 18 for details.

[Back to questions](#)

How can independent disk pools be implemented in my environment?

There are two basic environments in which you can take advantage of independent disk pools: a multi-system environment managed by an iSeries cluster, and a single-system environment with a single iSeries server. See “Switchable and stand-alone independent disk pools” on page 25 for details.

[Back to questions](#)

How should I structure my independent disk pools?

IBM provides some recommendations for structuring and populating your independent disk pools. See “Recommended structure for independent disk pools” on page 24 for details.

[Back to questions](#)

What is a disk pool group?

A disk pool group is made up 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. See “Disk pool groups” on page 26 for details.

[Back to questions](#)

What is geographic mirroring?

Geographic mirroring is a function that generates a mirror copy of an independent disk pool on a system that is (optionally) geographically distant from the originating site for availability or protection purposes. See “Geographic mirroring” on page 26 for details.

[Back to questions](#)

iSeries Navigator graphical user interface

How do I access the iSeries Navigator disk management function?

Before you can access disk management functions in iSeries Navigator, you must complete some setup tasks. See [Access disk management functions](#) for details.

[Back to questions](#)

What is the difference between the disk management functions in iSeries Navigator and in the character-based (command) interface?

Support for many independent disk pool tasks are only available through iSeries Navigator. Almost all disk management functions that are available from the system service tools (SST) mode are available through iSeries Navigator. A number of disk management functions that are only available from the dedicated service tools (DST) mode are also available.

[Back to questions](#)

How do I access the disk management function when the system is in dedicated service tools (DST) mode?

Beginning with V5R1, the Disk Units folder in iSeries Navigator is available when the system is in dedicated service tools (DST) mode.

[Back to questions](#)

What is the service tools server?

The service tools server allows you to use your PC to perform service tools functions through TCP/IP. Before you attempt to use any disk management functions, you must configure the service tools server. See “Set up communication” on page 50 for details.

[Back to questions](#)

Why does the data I see in the iSeries Navigator window appear to be out of date?

The disk management function in iSeries Navigator caches information, and therefore needs to be refreshed to have the most current data visible. After you make a configuration change, iSeries Navigator should refresh itself. If it does not, however, you can manually refresh it by clicking the Refresh button on the iSeries Navigator toolbar. You can also set iSeries Navigator to refresh periodically. Depending on the size of your server, however, you might not want to do this. Disk unit configuration data tends to be fairly static and does not need to be refreshed often. If your system is large, it can take a significant amount of time to download all information.

[Back to questions](#)

Why can't I connect to the service tools server after I add the service table entry?

The Add Service Table Entry (ADDSRVTBLE) command is case sensitive. In particular, it is important to ensure that the Protocol = 'tcp', and not 'TCP'. To ensure this is the case, use the Work with Service Table Entry (WRKSRVTBLE) command, and check the as-sts server field. Make sure that TCP is lowercase. If it is not, then remove the entry, and re-create it by issuing the following command exactly as shown:

```
ADDSRVTBLE SERVICE('as-sts') PORT(3000) PROTOCOL('tcp') TEXT('Service Tools Server')  
ALIAS('AS-STs')
```

[Back to questions](#)

Configuring

How do I create a new independent disk pool?

You can create an independent disk pool in a clustered, multisystem environment or on a single system. See the following topics for details:

- “Create a switchable independent disk pool” on page 93
- “Create a dedicated independent disk pool” on page 92

[Back to questions](#)

How do I create a disk pool group?

See “Create a new disk pool group” on page 97 for details.

[Back to questions](#)

How do I configure geographic mirroring?

You can configure geographic mirroring for independent disk pools that are dedicated or switchable between systems. See the following topics for details:

- “Configure geographic mirroring with dedicated independent disk pools” on page 98
- “Configure geographic mirroring with dedicated independent disk pools” on page 98

[Back to questions](#)

Performance

Why is performance slow?

There are several factors that can influence performance. Make sure your PC's TCP/IP settings are configured correctly. Specifically, make sure that you do not have an incorrect secondary gateway. If you do have a secondary gateway, remove it. This should provide a significant increase in performance. See Requirements for geographic mirroring for more detailed information.

[Back to questions](#)

Why is performance slow for geographic mirroring?

For geographic mirroring you should consider the distance that the independent disk pool is being mirrored. The type and number of communication lines as well as their bandwidth have an effect on performance. You can configure up to four TCP/IP communication interfaces on multiple adapters. You should consider configuring multiple communication lines to allow for the highest performance. The volume of disk unit writes that your applications require also plays a role in the performance of your geographically mirrored independent disk pool.

[Back to questions](#)

Troubleshooting

Why do no disk units appear as eligible to be added to my disk pool?

There are a number of possible reasons for this. First, you must have an unconfigured disk unit to add. If the disk pool is protected, you can only add parity disks, or disks in pairs, so that they can be mirrored.

If your system is in a clustered environment, disk unit eligibility is more complex. Each disk unit is assigned a rank, which indicates its eligibility to be added to a particular disk pool. If the rank of the disk unit is above 300, then the disk is ineligible. A complete list of the ranks, and what they mean, is available in the disk management online help.

[Back to questions](#)

Why doesn't the device description get deleted when I delete the disk pool?

Because the device description does not always get created by the disk management function, it might not be deleted when the disk pool gets deleted. You need to manually delete it using the Delete Device Description (DLTDEVD) command.

[Back to questions](#)

Why do I get a message saying the device description is already created?

When you create a new independent disk pool, an attempt is made to create an associated device description. If a device description of the same name as the disk pool already exists, you will receive a warning message, and the existing device description will not be changed. Most of the time, this is not a problem. However, if the device description's name and associated resource do not match, this becomes a problem, and this is why you receive the warning message.

Back to questions

Why does the primary or secondary disk pool that I try to create appear to be a UDFS disk pool?

If iSeries Navigator failed or was closed while the disk pool was being created, you might need to "Convert UDFS disk pools" on page 97 to a primary or secondary disk pool.

Back to questions

Why do I get a message that says my disk pool is not the right type when I try to create a library in it?

Make sure that the disk pool you are trying to create a library in is a primary or secondary disk pool, not a UDFS disk pool. If the disk pool is a UDFS disk pool and you want to create a library in it, you need to "Convert UDFS disk pools" on page 97 to a primary or secondary disk pool.

| What is a CPDB716 message, and how can it be fixed?


When a copy from an original ASP is made and there is an attempt to vary on the copy on the system that had used the original ASP, message CPDB716 appears. Before the system accepts the copy, an IPL must be performed.

Back to questions



Related information for Disk management

Listed here are the iSeries manuals and IBM Redbooks (in PDF format), Web sites, and information Center topics that relate to the disk management topic. You can view or print any of the PDFs.



Manuals

- You can refer to Backup and Recovery  (4 MB) for more information about disk configuration when using the character-based interface.

IBM Redbooks

- Clustering and IASPs for Higher Availability  (6.4 MB)
- iSeries Independent ASPs: A Guide to Moving Applications to IASPs  (3.4 MB)

Web sites


- High Availability and Clusters  (www.ibm.com/servers/eserver/series/ha/) This is the IBM site for High Availability and Clusters.
- Learning Services US  (www.ibm.com/services/learning/us/) This is the IBM site for IT product training, custom solutions, and e-Learning. You can search for courses offered on clustering and independent disk pools.

Saving PDF files

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

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

Downloading Adobe Reader

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

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