

IBM i
7.4

Systems management
Disk management

IBM

Note

Before using this information and the product it supports, read the information in [“Notices” on page 183.](#)

This edition applies to IBM i 7.4 (product number 5770-SS1) and to all subsequent releases and modifications until otherwise indicated in new editions. This version does not run on all reduced instruction set computer (RISC) models nor does it run on CISC models.

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

© **Copyright International Business Machines Corporation 2004, 2018.**

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 IBM i 7.4.....	1
PDF file for Disk Management.....	1
Getting started with disk management.....	2
Components of disk storage.....	2
Finding the logical address for a disk storage component.....	3
Planning for disk management.....	4
Requirements for disk management.....	4
Enabling and accessing disk units.....	4
Setting up communication.....	4
Evaluating the current configuration.....	5
Printing your disk configuration.....	5
Calculating disk space requirements.....	5
Setting up your disks.....	6
Configuring disks on a new system.....	6
Replacing a disk unit.....	7
Renaming a disk unit.....	7
Formatting a disk unit.....	8
Scanning a disk unit.....	8
Retrieving disk unit logs.....	8
Disk protection.....	9
Comparing disk protection options.....	9
Disk protection types.....	10
Device parity protection.....	10
Device parity protection concepts.....	10
Starting device parity protection.....	21
Starting device parity protection with hot spare protection	21
Managing device parity protection.....	22
Device parity protection examples.....	24
Write cache and auxiliary write cache IOA.....	25
Mirrored protection.....	26
Mirrored protection concepts.....	26
Planning for mirrored protection.....	30
Setting up mirrored protection.....	42
Managing mirrored protection.....	44
Hot spare protection.....	46
Hot spare protection concepts.....	46
Planning for hot spare protection.....	47
Starting hot spare protection.....	47
Managing hot spare protection.....	48
Troubleshooting hot spare protection.....	49
Multipath disk units.....	50
Considerations for multipath disk units.....	50
Disk protection checklists.....	52
Adding a new IOA.....	52
Adding disk units to an existing IOA.....	53
Removing disk units that have device parity protection from a disk pool without mirrored protection.....	54
Removing disk units that have device parity protection from a disk pool with mirrored protection.....	55
Upgrading the load source disk unit with device parity protection.....	56

Adding disk units without disk protection.....	61
Moving disk units between mirrored disk pools.....	62
Upgrading the load source disk unit with local mirroring using non-configured disk units.....	63
Upgrading the load source disk unit with local mirroring using newly installed disk units.....	70
NVMe failure with disk units mirror protected.....	77
Disk pools.....	78
Using disk pool.....	78
Types of disk pools.....	79
System disk pool.....	79
User disk pools.....	80
Basic disk pools.....	80
Independent disk pools.....	80
Primary, Secondary, UDFS disk pools.....	81
Disk pool groups.....	82
Disk pool concepts.....	82
Contrasting basic and independent disk pools.....	82
Disk pool benefits.....	83
Benefits of independent disk pools.....	83
Disk pool costs and limitations.....	84
Planning for disk pools.....	85
Using disk pools for improved performance.....	85
Library user disk pools.....	85
Nonlibrary user disk pools.....	86
System disk pool.....	86
Balancing a disk pool.....	86
Protecting your system disk pool.....	86
Capacity of the system disk pool.....	87
Planning for single system independent disk pools.....	87
Software requirements for single system independent disk pools.....	87
Application considerations for independent disk pools.....	88
Storing and printing spooled files.....	89
Object identification.....	90
Multiple system libraries.....	91
Supported and unsupported object types.....	92
Recommended structure for independent disk pools.....	93
Configuring disk pools.....	94
Creating a disk pool.....	94
Adding a disk unit or disk pool.....	95
Configuring independent disk pools.....	96
Converting UDFS disk pools.....	96
Creating an independent disk pool.....	97
Creating a new disk pool group.....	98
Managing disk pools.....	99
Deleting a disk pool.....	99
Removing a disk unit from a disk pool.....	99
Moving a disk unit from a disk pool.....	100
Clearing data from a disk pool.....	100
Setting the threshold of a disk pool.....	101
Tracing disk pool data.....	101
Balancing disk pool data.....	101
Enabling automatic overflow recovery for a basic disk pool.....	102
Using disk pools with extensive journaling.....	102
Managing independent disk pools.....	102
Making a disk pool available.....	102
Making a disk pool unavailable.....	103
Attach Independent Disk pool.....	104
Backing up and recovering independent disk pools.....	105
Releasing job reservations on an independent disk pool.....	106

Independent disk pool examples.....	107
Examples: Independent disk pool configurations.....	107
Examples: Dedicated independent disk pools.....	107
Example: Making independent disk pool available at startup.....	108
Independent disk pools with distinct databases.....	108
Disk pool checklists.....	109
Moving disk units between nonmirrored disk pools.....	109
Removing disk units without disk protection.....	110
Deleting a disk pool.....	111
Disk encryption.....	112
External load source disk unit.....	113
NVMe.....	113
NVMe concepts.....	113
Managing NVMe.....	114
Display NVM namespaces.....	114
Display NVM devices.....	115
Create NVM namespaces.....	115
Create NVM namespaces to match another NVM device.....	115
Create NVM namespaces to pair with Active mirror protected NVM units.....	115
Delete existing NVM namespaces.....	116
Sanitize/Erase NVM device.....	116
Format NVM device to prepare device for IBM i.....	117
NVMe password protection.....	117
Display the NVMe password protection policy and NVMe device attributes.....	117
Create the NVMe password protection policy.....	118
Change the password on the NVMe password protection policy.....	119
Delete the NVMe password protection policy.....	119
Add an NVMe to the existing NVMe password protection policy.....	120
Remove an NVMe from the existing NVMe password protection policy.....	120
Unlock an NVMe.....	121
Factory Reset an NVMe.....	122
Disk management checklist.....	123
Configuring disks on a new system.....	123
Adding disk units without disk protection.....	123
Adding disk units to an existing IOA.....	125
Adding a new IOA.....	125
Moving disk units between nonmirrored disk pools.....	126
Moving disk units between mirrored disk pools.....	127
Deleting a disk pool.....	128
Removing disk units without disk protection.....	129
Removing disk units that have device parity protection from a disk pool without mirrored protection.....	130
Removing disk units that have device parity protection from a disk pool with mirrored protection.....	131
Upgrading the load source disk unit with device parity protection.....	132
Upgrading the load source disk unit with local mirroring using non-configured disk units.....	137
Upgrading the load source disk unit with local mirroring using newly installed disk units.....	143
Upgrading the load source disk unit without disk protection.....	148
Installing a new partition using a NVMe as the load source disk unit.....	152
Removing an operational NVMe device from a partition.....	154
Upgrading the system ASP to NVMe - mirror protected.....	155
Upgrading the system ASP to NVMe - not mirror protected.....	160
Upgrading a user ASP to NVMe.....	165
Secure Data Erase for the IBM i.....	169
Frequently asked questions.....	177
Related information for Disk management.....	181

Notices.....	183
Programming interface information.....	184
Trademarks.....	184
Terms and conditions.....	185

Disk management

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

A disk pool is also referred to as an Auxiliary Storage Pool or ASP.

What's new for IBM i 7.4



Read about new or significantly changed information for the Disk Management topic collection.

What's new as of October 2019

Information about NVMe has been added. For more information, see [“NVMe” on page 113](#).

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.

In PDF files, you might see revision bars (|) in the left margin of new and changed information.

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

PDF file for Disk Management

You can view and print a PDF file of this information.


To view or download the PDF version of this document, select [Disk management](#).

Saving PDF files

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

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

Downloading Adobe Reader

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

Related reference

[Related information for Disk management](#)

Product manuals, IBM Redbooks (in PDF format), Web sites, and other information center topic collections contain information that relates to the Disk Management topic collection. You can view or print any of the PDF files.

Getting started with disk management

When a new disk unit is attached to the system, the system initially treats it as a non-configured disk. Learn what you can do with your disk units after initial installment.

You can add non-configured disk units to either the system disk pool, basic disk pool, or an independent disk pool. When adding non-configured disk units, use the serial number information that is assigned by the manufacturer to ensure that you are selecting the correct physical disk unit.

When you add a non-configured disk unit to a disk pool, the system assigns a unit number to the disk unit. The unit number can be used instead of the serial number and logical address.

When a disk unit has mirrored protection, the two disk units of the mirrored pair are assigned the same unit number. The serial number and logical address distinguish the two disk 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. If a printer is available, print the display of your disk configuration. If you must verify the unit number assignment, use IBM Navigator for i to display disk unit properties and check the serial number and logical address of each disk unit.

The disk unit that is addressed by the system as disk unit 1, the load source disk unit, is always used by the system to store Licensed Internal Code and data areas. The amount of storage that is used on disk unit 1 can be as large as 2.0 TB and varies depending on the configuration of your system. Disk unit 1 contains a limited amount of user data. Disk unit 1 contains the initial programs and data that is used during an IPL of the system.

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

Related tasks


[Adding a disk unit or disk pool](#)

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

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.



Bus:

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

IOP:

The IOP is attached to the bus. The IOP is used to transfer information between main storage and specific groups of IOA. 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 IOA. Some systems do not have an IOP.

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 disk units. You order hardware at the disk-unit level. Each disk unit has a unique serial number.

The system 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 IOP, an IOA, and a device number.

Finding the logical address for a disk storage component

Use this information to find the logical address for a disk storage component.

To find the logical address for a disk storage component using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Units**.
3. Right-click the disk unit that you want to find the address for.
4. Select **Properties**.

Planning 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.

Requirements for disk management

There are requirements that must be completed before you can configure or work with your disks.

Do these procedures before you begin using disk management. You can use IBM Navigator for i to use disk management.

Enabling and accessing disk units

You must follow these procedures before you can perform any disk management tasks using IBM Navigator for i.

Before you can perform any disk management tasks using IBM Navigator for i you need to set up the proper authorizations for dedicated service tools (DST).

1. The user profile that is used to access disk units in IBM Navigator for i must have at least these authorities:
 - *ALLOBJ – All object authority
 - *SERVICE
2. From the IBM i command line, type **STRSST** and press enter to start System service tools.
3. Sign on to SST with your service tools user ID and password, or use the **QSECOFR** service tools user ID and password.
4. Select option 8 (Work with service tools server security and devices).
5. Select option 4 to change the service tools password level to **PWLVL 2**.
6. Select option 1 (Work with service tools user IDs) to work with service tools user IDs.
7. Create a service tools user ID or change the password of an existing user profile to apply the new password level.
8. Give this service tools user ID at least these authorities:
 - Disk units – operation
 - Disk units – administration
9. Press Enter to enable these changes.

Related information

[Accessing service tools using DST](#)

Setting up communication

This topic describes how to configure the service tools system.

IBM Navigator for i can be accessed from the World Wide Web. Before you use any disk management functions, you must configure the service tools system.

To access disk management functions in IBM Navigator for i you must first configure the service tools system and it must have DST access. You must also configure the user IDs. Before you start, be familiar with the service tools concepts and follow the instructions in [“Enabling and accessing disk units”](#) on page 4. See [Configuring the service tools server](#) and [Managing service tools user IDs](#) for more instructions.

For more information, see [Service tools](#).

Related information

[Service tool concepts](#)

[Configuring the service tools server](#)

Evaluating the current configuration

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

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

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

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

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

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

Printing your disk configuration

Print your disk configuration.

To print the disk configuration for your records using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Units**.
3. From the **Actions** menu, select **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 box, select the **Print Preview** button.
6. On the new popup browser, select **File->Print**.

Note: You must select the browser option that allows printing the page background with colors and pictures.

Calculating disk space requirements

Before you change the disk configuration or the disk protection on your system, you need to calculate the space requirements for the change. This calculation 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.

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

To view the disk pool configuration using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Pools**.

3. Right click the source disk pool that you want to view and select **Properties**.
4. Select the **Capacity** tab. The **Capacity** tab displays the used space, free space, total capacity, threshold, and percentage of disk space that is used for the disk pool.
5. Record the used space, free space, and the threshold from the **Capacity** tab.
6. Enter the used-space value and free-space value into the disk space calculator.
7. If you want to use the threshold value, enter the threshold value in the calculator. The calculator warns you if your disk use exceeds your threshold.

Scenario: Calculating disk space when moving a disk unit

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

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

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

Perform the following scenario:

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.

Setting up your disks

Evaluate and perform initial set up for your disks.

Configuring disks on a new system

This checklist shows the sequence of tasks that you use to configure disks on a new system. Whether you need to perform all the tasks depends on the disk protection that you want on your system.

“Disk protection” on page 9 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 disk unit appear as non-configured.	“Evaluating the current configuration” on page 5

Task	What to do	Where to learn more
2.____	Use the Add Disk Unit wizard to add non-configured 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.	“Adding a disk unit or disk pool” on page 95
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%.	“Setting the threshold of a disk pool” on page 101
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.	“Starting mirrored protection” on page 43
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.	“Evaluating the current configuration” on page 5
7.____	Print your disk configuration to have available in case a recovery situation occurs.	“Printing your disk configuration” on page 5

Replacing a disk unit

If you must 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 failed. A disk unit with device parity protection cannot be replaced with a non-configured disk even if it failed.

To replace a failed disk unit or exchange a suspended mirrored disk unit using IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **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.

Renaming a disk unit

You can change 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 using IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Units**.
3. Right-click the disk unit that you want to rename and select **Rename**.
4. Follow the instruction on the dialog box that is displayed.

Formatting a disk unit

You can select to clear all data from a non-configured disk unit and write the sectors, which prepares the disk unit for use in the system.

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 using IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Units**.
3. Right-click the disk unit that you want to format and select **Format**.
4. Follow the instructions on the dialog box that is displayed.

Scanning 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 using IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Units**.
3. Right-click the disk unit that you want to scan and select **Scan**.
4. Follow the instructions on the dialog box that is displayed.

Retrieving disk unit logs

You can 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 using IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Units**.
3. Select the disk unit that you want to retrieve the log for.
4. Right-click the 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 **Start a service tool**.
3. On the Start a Service Tool display, select **Product activity log**.
4. On the Product Activity Log display, select **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 **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 IBM i Technical Support.

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.

Comparing disk protection options

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

Use this table to determine what factors are important to you, when determining disk protection options.			
	Device parity protection	Device parity protection with auxiliary cache	Mirrored protection
Usable disk capacity	excellent	excellent	good
Redundancy	good	very good	excellent
Cost	excellent	excellent	good
Performance	very good	very good	excellent

This table provides an overview of the hardware that can be used on the system to protect against different types of failure.			
Scope of redundancy	Device parity protection	Device parity protection with auxiliary cache	Mirrored protection
Disk	Yes	Yes	Yes
Input/Output Adapter (IOA) cache	No	Yes	Yes
IOA	No	Yes	Yes
Enclosure	No	Yes	Yes
HSL/loop	No	No	Yes

Hardware requirements

- 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.
- 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, Input/Output Processors (IOPs), and IOAs, depending on the level of protection that you want.
- Hot spare protection requires an extra disk unit that is ready and waiting to be put into action when another disk unit fails.

Disk protection types

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

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.

To protect data, the disk 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.

Hot spare protection can be used with device parity protection. When a valid hot spare disk unit exists in the system and a device parity protected disk unit fails, the hot spare disk unit replaces the failed disk unit.

Important: Device parity protection is not a substitute for a backup and recovery strategy. It is necessary to perform routine saves of your system.

Related concepts

[Hot spare protection](#)

Protect your disk units with hot spare protection.

Device parity protection concepts

Learn more about RAID 5, RAID 6, and RAID 10 protection to decide which level of protection is best for you.

RAID 5 concepts

RAID 5 protection protects data from being lost because of a disk unit failure or because of damage to a disk. RAID 5 protection protects against a one disk unit failure.

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 across the 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.

Note:

1. Systems with IOAs released before 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.
2. 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.

Related concepts

[Mirrored protection](#)

Mirrored protection is beneficial if you have a multibus system or a system with a large single bus. A greater number of disk units provides more opportunity for failure and increased recovery time.

How RAID 5 works

This topic describes how device parity protection works.

Parity set optimization on SCSI adapters

The IOA determines how parity sets are formed. SCSI IOAs allow you to 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 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.

Parity set optimization on SAS adapters

The IOA determines how parity sets are formed. SAS IOAs allow 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 bus failure on the IOA. The parity set is formed from at least three disk units of equal capacity, with no more than two disk units attached to an individual bus on the IOA.

Including additional disk units to a parity set

The capacity of the new disk unit, before it is added to the parity set, must be the same capacity or greater than the capacity of the disk units that are currently in the parity set, before device parity protection was started on them. 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 System i® 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.

When a 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 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 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.


Elements of RAID 5 protection

This topic describes and illustrates RAID 5 protection.

The following diagrams illustrate the elements of a parity set that contains four disk units. Each parity set begins with an IOP that is attached to an 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.

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.



How RAID 5 affects performance

Learn about how performance is affected when using RAID 5 protection.

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

Benefits of RAID 5 protection

There are many benefits of using RAID 5 device parity protection.

- 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.
- Only one disk unit of capacity stores parity data in a parity set.

Costs and limitations of RAID 5 protection

There are costs and limitations when using RAID 5 protection.

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

Parity set optimization for RAID 5 protection

When you select to optimize a parity set, the IOA 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.

Note: SAS IOAs provide optimal performance, capacity, and balance, so selecting any of the parity set optimizations, other than Availability, will not affect the outcome of the parity sets.

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 IOA. For example, if an IOA 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 IOA may generate fewer parity sets with more disk units in each parity set. For example, if an IOA 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 IOA 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 IOA may generate more parity sets with fewer numbers of disk units. For example, if an IOA had 15 disk units and is optimized for performance, the result might be three parity sets with five disk units each.

When in a dual storage IOA configuration, the system attempts to create an even number of parity sets. An even number of parity sets distributes the workload evenly between a pair of adapters which are in a dual storage IOA configuration. This provides the fastest data access since each adapter has a piece of the workload.

Changing parity set optimization

Changing the parity set optimization 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. To change the parity set optimization for all new parity sets that are created refer to [Changing parity set optimization for RAID 5 protection](#).

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 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.

IOA migration

There are considerations and limitations when migrating an IOA.

Before you begin the migration to the new IOA, as with any configuration change, it is important to do a normal system power down. 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 V5R2 is migrated to an IOA that was released after V5R2, your disk units are not protected by device parity protection while parity is being regenerated.

Important: Because your disk units are not parity protected during a migration it is necessary to perform a save.

You cannot migrate a parity set back to adapters released before V5R2 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.

RAID 6 concepts

RAID 6 protection protects data from being lost because of a disk unit failure or because of damage to a disk. RAID 6 protection protects up to two disk unit failures.

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.

The Reed Soloman algorithm and the hardware finite field multiplier, are used to create the stripes of parity data in a RAID set. These features enhance performance and functionality.

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

How RAID 6 protection works

This topic describes how RAID 6 protection works.

The IOA 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 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 bus on the IOA.

It is possible to include additional disk units of the same capacity in a 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 IBM Systems Director Navigator for i5/OS and System i 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* disk unit with a model number of 090 because it is a disk unit that does not store parity data.

When a 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 parity set by stopping parity protection and starting it again.

A write cache is included in the 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.

Related tasks

[Changing parity set optimization for device parity protection](#)

You can choose how you want your parity sets to be optimized.


Elements of RAID 6 protection

This topic describes and illustrates RAID 6 protection.

The following diagrams illustrate the elements of a parity set that contains four disk units. Each parity set begins with an IOP that is attached to an 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.



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 RAID 6 affects performance

This topic describes the performance of using RAID 6 protection.

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.

Benefits of RAID 6 protection

There are many benefits of using RAID 6 parity protection.

- 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.
- Two disk units of capacity are dedicated to storing parity data in a parity set.

Costs and limitations of RAID 6 protection

There are costs and limitations when using RAID 6 protection.

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

Parity set optimization for RAID 6 protection

RAID 6 IOAs provide optimal performance, capacity, and balance, so selecting any of the parity set optimizations, other than Availability, will not affect the outcome of the parity sets.

When you select to optimize a parity set, the IOA will choose disk units for parity sets according to the optimization value you have chosen.

When not in a dual storage IOA configuration, RAID 6 IOAs provide optimal performance, capacity, and balance, so selecting any of the parity set optimizations, other than Availability, will not affect the outcome of the parity sets.

When in a dual storage IOA configuration, using Performance parity set optimization is recommended to provide optimal performance, capacity, balance, and availability.

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 parity set is formed from at least four disk units of equal capacity, with no more than two disk units attached to an individual bus on the IOA.

Performance

When in a dual storage IOA configuration, the system attempts to create an even number of parity sets. An even number of parity sets distributes the workload evenly between a pair of adapters which are in a dual storage IOA configuration. This provides the fastest data access since each adapter has a piece of the workload.

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 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.

IOA migration

There are considerations and limitations when migrating an IOA.

Before you begin the migration to the new IOA, as with any configuration change, it is important to do a normal system power down. 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 V5R2 is migrated to an IOA that was released after V5R2, your disk units are not protected by device parity protection while parity is being regenerated.

Important: Because your disk units are not parity protected during a migration it is necessary to perform a save.

You cannot migrate a parity set back to adapters released before V5R2 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.

RAID 10 concepts

RAID 10 protection protects data from being lost because of a disk unit failure or because of damage to a disk. RAID 10 protection provides a level of protection comparable to device level mirroring.

RAID 10

RAID 10 protects availability due to disk unit failure by pairing disks together into logical mirrors. Each pair of disks is considered a parity set. In addition to being logically mirrored, RAID 10 also uses block-level striping. RAID 10 protection is effectively the combination of RAID 0 (data striping) and RAID 1 (disk mirroring).

A complete restore of ASP data is only required if both disks in the parity set fail. If both disks in the parity set fail, you must restore the data from the backup media.

Input/output adapter caching greatly enhances both the read and write performance characteristics in a RAID 10 configuration. Non-cached read operations are distributed between the two disks in the parity set reducing the workload on each disk. Non-cached write operations (write-cache overruns) require a write operation by the adapter to each disk in the parity set to complete the write operation.

How RAID 10 protection works

This topic describes how RAID 10 protection works.

The IOA determines how parity sets are formed. RAID 10 protection uses two disk units in each parity set. Optimal performance, capacity, and balancing are achieved with any of the parity optimization settings.

It is not possible to include more disk units of the same capacity in a parity set after device parity protection is started. To add disk units to an existing I/O adapter that is running RAID 10, the system requires that you start a new parity set, rather than include them in an existing parity set.

In IBM Navigator for i, 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 might be eligible to start a new parity set. An unprotected unit also has a model number of 050.

A write cache is included in the IOA for each parity set to improve performance of interactive write workloads.

Note: Starting RAID 10 protection before the disk units are added to a disk pool significantly reduces the time that it takes to start device parity protection and configure the disk units.

Elements of RAID 10 protection

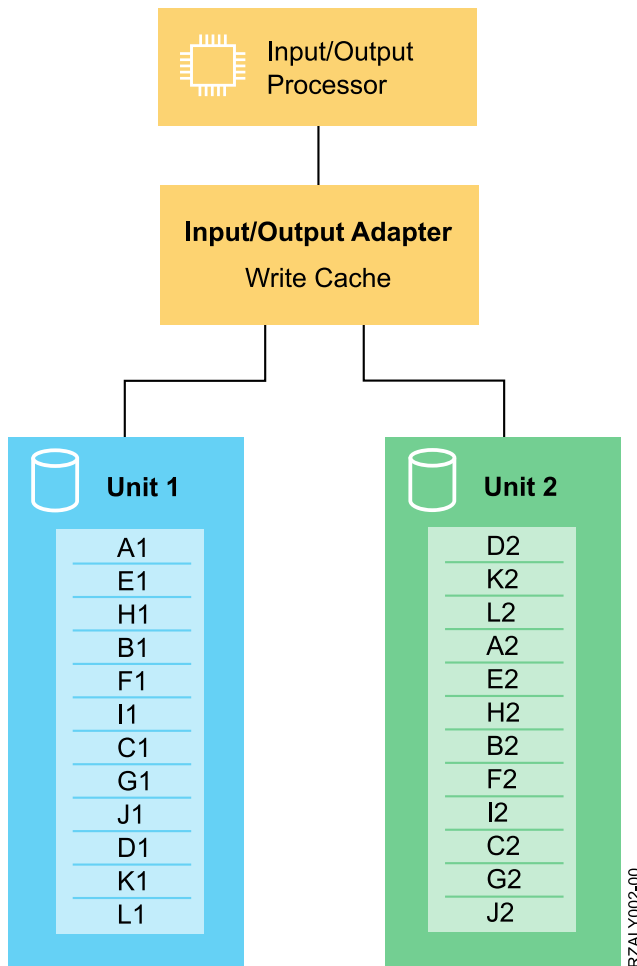
This topic describes and illustrates RAID 10 protection.

The following diagram illustrates the elements of a parity set that contains two disk units. Each parity set begins with a virtual IOP, based in software, that is attached to an IOA, which contains the write cache. The IOA transmits read and write signals to the attached disk units.

Each lettered element provides the name for a particular data stripe.

The associated number indicates the copies of the data.

The IOA provides the stripe allocation size and positions on each unit in the parity set.



How RAID 10 affects performance

This topic describes the performance of using RAID 10 protection.

Because RAID 10 uses two disk units per parity set, RAID 10 behaves effectively as device level mirroring. At the device level, read operations are distributed between the two units, while write operations must occur to each.

Many factors affect the performance of any protection scheme.

- Disk unit capacity.
- Disk unit connection throughput such as 6 GB SAS, and 3 GB SAS.
- Disk unit rotational speed such as 7200 RPM, 10 K RPM, or 15 K RPM.
- I/O adapter connection throughput for the disk unit. The connection is scaled to the fastest throughput supported by all the devices connected. For SAS connections, each connection is separately negotiated by the adapter and the device.
- I/O adapter read and write cache capacity.
- PCIe bus width such as x8 or x16. The connection is scaled to the widest bandwidth supported by both the bus and the adapter connected. For PCIe, a separate connection exists for each I/O adapter installed.
- PCIe bus speed such as PCIe 1, PCIe 2, or PCIe 3. The connection is scaled to the fastest speed supported by both the bus and the devices connected. For PCIe, a separate connection exists for each I/O adapter installed.
- Processor 12X loop loading.
- Processor I/O operation bandwidth.

RAID 10 does not have separate parity stripe data for protection. It uses a second copy of the data. Therefore, fewer I/O operations are done to maintain availability than there is with RAID 5 or RAID 6.

Compared to operating system-based software mirroring, a RAID 10 write operation is only done one time to write-cache within the IOA by the operating system. By writing to cache one time, the amount of memory that is needed by the operating system to support I/O operations is reduced.

Benefits of RAID 10 protection

There are many benefits of using RAID 10 parity protection.

- Lost data is automatically reconstructed by the IOA after a disk failure.
- The system continues to run after a single disk failure in a RAID 10 parity set.
- Failed disk units can be replaced without stopping the system.

Costs and limitations of RAID 10 protection

There are costs and limitations when you use RAID 10 protection.

- Since RAID 10 implements a mirrored disk array, the number of disk units that are required is twice that of an unprotected array.
- Since disk write operations are performed once to each disk unit in the array, heavy write workloads, such as restore operations, might see slower performance, take longer, or be susceptible to increased cache overruns.

Parity set optimization for RAID 10 protection

RAID 10 IOAs provide optimal performance, capacity, and availability. Selection of the performance optimization setting varies based on the IOAs used.

When not in a dual storage IOA configuration, RAID 10 IOAs provide optimal performance, capacity, and balance, so selecting any of the parity set optimizations, other than Availability, does not affect the outcome of the parity sets.

When in a dual storage IOA configuration, use Performance parity set optimization to provide optimal performance, capacity, balance, and availability.

Performance

When in a dual storage IOA configuration, the system attempts to create an even number of parity sets. An even number of parity sets distributes the workload evenly between a pair of adapters that are in a dual storage IOA configuration. An even number of parity sets provides the fastest data access since each adapter has a portion of the workload.

Read operations on a failed RAID 10 disk unit

To access the data that was contained on a failed disk unit in a RAID 10 parity set, device parity protection must read the other disk unit in the parity set that contains the failed disk unit.

The performance impact from a failed disk unit is negligible for a RAID 10 protected system.

RAID 10 IOA migration

There are considerations and limitations when you migrate to a new class of IOA.

Before you begin the migration to the new IOA, as with any configuration change, it is important to do a normal system power down. The power down assures that all of your cache data is written to disk before the power down completes.

When you migrate to a new class of IOA, device parity protection might need to be rebuilt. During the rebuild of device parity protection, disk units are not protected.

Important: Because your disk units are not protected by RAID during a migration, it is necessary to perform a save.

You cannot migrate a parity set to adapters of a previous class of IOA after you migrated to a new adapter. You cannot migrate a parity set 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 between RAID 5, RAID 6, or RAID 10 protection, you must stop and restart the device parity protection.

Note: You cannot migrate RAID 10 protected disk units to an adapter that does not support RAID 10.

Starting device parity protection

The best time to start device parity protection is when you add new or non-configured disk units. The Add a disk unit or disk pool wizard has steps for including disk units in a parity set and starting device parity protection.

Notes:

- All disk units in a parity set must be the same capacity.
- Systems with IOAs released before V5R2 of OS/400, the minimum number of disk units for a RAID 5 parity set is 4. The maximum number of disk units in a parity set is 10.
- Systems with IOAs released after V5R2 can have a minimum number of 3 disk units for a RAID 5 parity set. The maximum number of disk units in a parity set is 18.
- The minimum number of disk units in a RAID 6 parity set is 4. The maximum number of disk units in a parity set is 18.

Follow these steps to start parity protection using IBM Navigator for i.

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **All Tasks > Disk Management > Start Parity**.
3. Follow the instructions on the dialog box that is displayed.

Related concepts

[How RAID 6 protection works](#)

This topic describes how RAID 6 protection works.

Related tasks

[Starting device parity protection with hot spare protection](#)

This topic provides instructions for starting device parity protection with hot spare protection.

Starting device parity protection with hot spare protection

This topic provides instructions for starting device parity protection with hot spare protection.

Note: Before proceeding with this section, read the [“Hot spare protection concepts”](#) on page 46 and [“Planning for hot spare protection”](#) on page 47 sections.

On the following Start Parity screens, only the device parity sets that can be protected by a hot spare are shown. If you do not see all the parity sets and IOAs expected, the system is unable to create the necessary hot spare devices required. Refer to the [“Troubleshooting hot spare protection”](#) on page 49 section for possible reasons.

To start device parity protection with hot spare protection using IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **All Tasks > Disk Management > Start Parity**.
3. In the **Hot Spare Protection** drop-down list select **Yes**.

To start device parity protection with hot spare protection with the command line, take the following steps:

1. Start System Service Tools (STRSST), and specify the user name and password.
2. On the System Service Tools (SST) display, select **Work with disk units**.
3. On the Work with Disk Units display, select **Work with disk configuration**.
4. On the Work with Disk Configuration display, select **Work with device parity protection**.

5. On the Work with device parity protection display, select **Start device parity protection - RAID 5 with hot spare**, **Start device parity protection - RAID 6 with hot spare**, or **Start device parity protection - RAID 10 with hot spare** depending on the level of parity protection that is desired

Note: If a hot spare disk unit of the correct capacity exists for this IOA, any new device parity set is hot spare protected even if you do not select hot spare protection.

Note: If there is a sufficient number of hot spare disk units already available, the system will not create any additional hot spares.

Note: RAID 10 may leave an unprotected disk on each I/O bus when device parity with hot spare is started due to the requirement of two disks per parity set.

Managing device parity protection

Refer to this topic to see the tasks that you can perform to manage your disk protection activities.

Stopping device parity protection

You can select to stop device parity protection on the displayed disk units.

When stop device parity protection is selected, the system validates the configuration 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 must IPL in manual mode to stop device parity.

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.

To stop device parity protection on the disk units in a parity set using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Units**.
3. Select the disk units for which you want to stop device parity protection.
4. Right-click the disk unit and select **Stop Parity**.
5. From the dialog box that is displayed, select **Stop Parity** to stop device parity.

Related concepts

Mirrored protection

Mirrored protection is beneficial if you have a multibus system or a system with a large single bus. A greater number of disk units provides more opportunity for failure and increased recovery time.

Including disk units in a parity set

When you attach a new disk unit to an existing IOA that has device parity protection, you can include the disk unit into an existing parity set with other disk units.



Attention: Including disk units into an existing parity set with PCIe2, PCIe3, and newer generations of SAS RAID adapters can result in a large performance degradation. The degradation can occur because the adapter uses firmware in place of specialized hardware assists when drives are included into an existing parity set. For this reason, the Include disk operation is no longer recommended with PCIe2, PCIe3, and newer generations of SAS RAID adapters. Systems with POWER8® technology and beyond do not allow the Include disk operation. Alternatives for adding disk units to the system are to stop parity on the existing set and then start parity with the additional disks in the set, or to create a new parity set of three or more disks with the additional disk units.

If you want to include a disk unit in a parity set and the disk units of that parity set are in an independent disk pool, the independent disk pool must be varied off. For all other disk pools, you must restart your system to dedicated service tools (DST) mode before you can include them in a parity set.

Adding three or more disk units requires you to create a new parity set. The include function does not work to create a parity set. To create a parity set, go to Start device parity protection.

For a Small Computer System Interface (SCSI) IOA, the disk units to include must be the same capacity as the other disk units in the parity set. For a Serial Attached SCSI (SAS) IOA, the disk units to include can be the same capacity or larger than the other disk units in the parity set

To include a disk unit in a parity set using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Units**.
3. Select the disk units that you want to include.
4. Right-click the selected disk units and select **Include in Parity Set**.
5. From the dialog box that is displayed, select **Include** to include the selected disk units in a parity set.

Related tasks

Starting device parity protection

The best time to start device parity protection is when you add new or non-configured disk units. The Add a disk unit or disk pool wizard has steps for including disk units in a parity set and starting device parity protection.

Excluding 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.

You can exclude a protected disk unit with a model number of 070 or 090 because these disk units do not store parity data.

When you exclude a disk unit from the parity set, the data on the 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 using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Units**.
3. Select the disk units that you want to exclude.
4. Right-click the selected disk units and select **Exclude from Parity Set**.
5. From the dialog box that is displayed, select **Exclude** to exclude the disk units from a parity set.

Changing parity set optimization for device parity protection

You can choose how you want your parity sets to be optimized.

This change stays in effect until you change it again. If you must start parity, you can also change the parity set optimization as part of the start parity process.

To change the parity set optimization using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **All Tasks > Disk Management**.
3. Select **Change parity set optimization**.

Note: When in a dual storage IOA configuration, select performance parity set optimization. When performance parity set optimization is specified, the system attempts to create an even number of parity sets. An even number of parity sets distributes the workload evenly between a pair of adapters, which are in a dual storage IOA configuration.

Note: When not in a dual storage IOA configuration, select capacity parity set optimization or balance parity set optimization. Either of these selections gives you the optimal capacity and balance of the parity set configuration.

Determining 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.

Determining what disks are in a parity set using the SST menu

Follow these steps to 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.

Determining what disks are in a parity set using IBM Navigator for i

Follow these steps to determine what disks are in a parity set using your IBM Navigator for i.

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **All Tasks > Disk Management**.
3. Select **Parity Sets**.
4. Right-click on the parity set and select **Open** to see the list of disk units that are contained in that set.

Device parity protection examples

Refer to this topic to see examples of device parity protection works when a disk failure occurs.

Example: Writing operations on a failed disk unit

An illustration shows how device parity protection works.

The following figure shows a failed disk unit under an IOA with RAID 5 device parity protection. This parity set has 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 1. 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.

Example: Writing to a failed disk unit

This topic gives you an example of what happens when a write operation fails.

A write operation from a system 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 (number of disk units-2 reads) 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

This topic gives you an example of how to write data to a disk unit that has parity data on a failed disk unit.

The write request from a system 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.

Write cache and auxiliary write cache IOA

This topic describes how write cache and auxiliary write cache works.

The write cache provides greater data integrity and improved performance. When a system sends a write operation, the data is written to the cache. Then, a write-completion message is sent back to the system. 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 system:

1. Data is committed to a nonvolatile battery-backed cache in the IOA.
2. A write completion message is sent from the system.
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.

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 the system 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 system and the data is then written to a disk.

Note: In order for write cache mirroring to occur, the storage IOAs must be connected to supported auxiliary write cache IOA. The storage IOA and the auxiliary write cache IOA must also be in the same expansion unit 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: This IOA is not a failover device that can keep the system operational if the disk IOA, or its cache, fails.

Mirrored protection

Mirrored protection is beneficial if you have a multibus system or a system 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 system and is distinct from cross-site mirroring. Mirrored protection works to prevent outage on the system by keeping a second copy of the data on a mirrored disk unit. If one disk unit fails, the system relies on the mirrored disk unit.

Hot spare protection can be used with mirrored protection. When an appropriate hot spare disk unit exists in the system and a mirrored disk unit is suspended because of a disk failure, the hot spare disk unit replaces the failed subunit.

Related concepts

Hot spare protection

Protect your disk units with hot spare protection.

Mirrored protection concepts

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 separate I/O buses, IOAs, IOPs, 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 IBM i systems and is a part of the Licensed Internal Code.

Remote mirroring support allows you to have one mirrored disk unit within a mirrored pair at the local site, and the second mirrored disk 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.

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, IOP, or IOA for the hardware configuration of the system. When disk 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, IOP, or IOA levels, a failure in one of these hardware components allows the other mirrored disk unit to continue operating.

Any data that is written to a disk unit that is mirrored is written to both disk units of the mirrored pair. When data is read from a disk unit that is mirrored, the read operation can be from either disk unit of the mirrored pair. It is transparent to the user which mirrored disk unit the data is being read from. A user is not aware of the existence of two physical copies of the data.

If one disk unit of a mirrored pair fails, the system suspends mirrored protection to the failed mirrored disk unit. The system continues to operate using the remaining mirrored disk unit. The failing mirrored disk unit can be physically repaired or replaced. When a hot spare disk unit is configured to protect the mirrored unit, the hot spare disk unit automatically replaces the failed mirrored disk unit after the failed mirrored disk unit has been suspended for five minutes. The five minute waiting period avoids consuming the hot spare disk unit if the mirror suspension is a temporary condition.

After the failed mirrored disk unit is repaired or replaced, the system synchronizes the mirrored pair by copying current data from the disk unit that has remained operational to the other disk unit. During synchronization, the mirrored disk 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 disk unit becomes active.

Related concepts

Getting started with disk management

When a new disk unit is attached to the system, the system initially treats it as a non-configured disk. Learn what you can do with your disk units after initial installment.

Mirrored protection and performance

When mirrored protection is started, most systems show little difference in performance. There are some cases where 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 disk 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 disk 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 disk units of each mirrored pair. If the system cannot determine whether the last changes were written to both disk units of the mirrored pair, the system synchronizes the mirrored pair by copying the data in question from one disk unit of each mirrored pair to the other disk 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 disk units of the load source mirrored pair must be powered by the basic uninterruptible power supply.

Mirrored protection benefits

Even if your system is not a large one, mirrored protection can provide valuable protection.

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 IOP, 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.

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

There are costs and limitations when using mirrored protection.

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, IOA, or IOP fails, you need duplicate IOA and IOPs. 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 IOP-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, IOPs, and IOA 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 disk 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 disk unit in the mirrored pair fails before the first failing disk unit is repaired and mirrored protection is resumed. If two failed disk 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 disk 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 disk units in a mirrored pair fail causing data loss, the entire disk pool is lost and all disk 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 disk unit may be moved to another disk unit. The preferred disk unit may no longer exist after mirror protection is started.

Related concepts

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.

Related information

Applying journaled changes

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 disk 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 disk 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 disk unit 1 and its mirrored disk unit to only IOA-level protection. See *Mirrored Protection - Configuration Rules*.

Under some conditions, diagnosis and repair can require active mirrored disk 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.

Related information

Mirrored protection—Configuration rules

Advantages of remote disk unit mirroring

There are two advantages of remote 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 load source disk unit mirroring for primary partitions

There are two disadvantages of remote load source disk unit mirroring for primary partitions.

- A system that uses remote load source 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 load source service procedure is performed.

- When remote load source 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.

Note: The disadvantages listed above are not true for POWER5 and POWER6® systems.

Planning for mirrored protection

Successfully plan for protecting your disks with mirrored protection.

Planning 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.

Deciding which disk pools to protect

Mirrored protection can be used to protect one, some, or all disk pools on a system.

Mirrored protection is configured by disk pools because it is the user's level of control over single-level storage. Multiple disk pools are not required to use 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 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 IOPs. The number of disk units in the disk pool that are attached to each IOP should be the same.

Related concepts

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.

Determining the disk units that are needed for mirrored protection

There are considerations that you must take into account when determining the disk units that are needed for protection.

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 disk units of each size are present. The system assigns mirrored pairs and automatically moves the data as necessary. If a disk pool does not contain sufficient storage capacity, or if disk 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.

Planning for storage capacity

For a new system, your IBM marketing representative can help you analyze your system storage requirements. For an existing system, the planned amount of data in the disk pool 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.

Planning 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.

Hot spare protection can be used in conjunction with mirrored protection. When a hot spare disk unit is configured to protect the mirrored unit, the hot spare disk unit automatically replaces the failed mirrored disk unit. This eliminates the need for customer action or waiting for a service representative before mirrored protection is again active for that pair. The failed unit still needs to be repaired or replaced to allow hot spare protection to be restored.

Planning for total storage capacity

This information helps you plan the total number of disk units needed for your system.

After planning for the number and type of disk units needed for each disk pool on the system, and for any spare disk units, add up the total number of disk units of each disk unit type and model. Remember that the number planned is the number of disk 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 disk units to disk units before ordering hardware.

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.

Determining the level of mirrored protection that you want

Determine the level of mirrored protection that you want with the information provided.

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 IOP fails, all of the disk units attached to the failing IOP must have mirrored disk units attached to different IOP.

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 IOP 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 IOP, you have IOP-level protection. If you duplicate buses, you have bus-level protection. Mirrored disk 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.

Related concepts

[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.

Disk-unit level protection

Mirrored protection always provides disk-unit level protection because the disk units are duplicated.




Figure 2. Disk-unit level protection

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.

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

Input/output bus level protection

Determine whether you want I/O bus level protection.




Figure 3. Input/output bus level protection

The input/output bus level protection has the following benefits:

- Your system is available when an I/O bus fails.
- You can concurrently repair a failed disk unit.

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

The input/output bus level protection figure shows the elements of I/O bus protection: one bus, connected to one IOP, connected to one I/O Adapter (IOA), with two or more I/O buses, each attached to a separate disk unit.

IOA level protection

Determine whether you want IOA level protection.




Figure 4. IOA level protection

The IOA level protection has the following benefits.

- Your system is available when an IOA fails.
- You can 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 disk unit attached to a different IOA. This figure shows IOA-level protection. The two disk units make a mirrored pair. With IOA-level protection, the system can continue to operate if one IOA fails. If the IOP fails, the system cannot access data on either of the disk units, and the system is unusable.

The IOA level protection figure shows the elements of IOA-level protection: one bus, connected to one IOP, connected to two IOAs, which are each attached to a separate disk unit.

IOP-level protection

Determine if you want IOP-level protection.




Figure 5. IOP-level protection

The IOP-level protection has the following benefits:

- Your system is available when an IOP fails.
- Your system is available when the cable attached to the IOP fails.
- You can 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 IOP must have a mirrored disk unit attached to a different IOP. On many systems, IOP-level protection is not possible for the mirrored pair for disk unit 1.

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

Bus-level protection

Determine if you want bus-level protection.




Figure 6. Bus-level protection

The bus-level protection has the following benefits:

- 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 disk unit attached to a different bus. Bus-level protection is not possible for disk unit 1.

The bus-level protection 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 disk units make a mirrored pair.

Expansion-unit level protection

Determine if you want expansion-unit level protection.




Figure 7. Expansion-unit level protection

The expansion-unit level protection has the following benefits:

- 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 disk unit present in another expansion unit. The Expansion-unit level protection 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.




Figure 8. Ring-level protection

The ring-level protection has the following benefits:

- 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 disk units present in another expansion unit in the second HSL. The Ring-level protection 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, IOA, and disk units, respectively.

Remote disk unit mirroring support

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.

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. Use remote disk unit mirroring for a higher level of protection for your system.

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.

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 and then Start site-to-site mirroring.

Related tasks

Preparing your system for remote mirroring

To provide protection against a site disaster, all disk units in all disk pools of the system must be mirrored in local-remote pairs.

Starting remote Load Source mirroring

This topic provides instructions for starting remote load source mirroring.

Comparing standard mirroring and remote mirroring

Compare remote mirroring to disk unit standard 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 disk units, half of the new disk units of each capacity of disk unit should be in the remote group and half in the local group. Single device-parity protected disk 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.

Restoring 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 *Recovering your system*.

Related information

Recovering your system

Determining the hardware that is needed for mirroring

In order to communicate with the rest of the system, disk units are attached to IOAs, which are attached to IOPs, 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 IOAs. The same number of IOAs should be under each IOP for that disk type. The IOPs 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.

Planning 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 configuration limits and help you use them in your planning.

For each disk unit type, first plan for the IOAs that are needed and then for the IOPs that are needed. After planning the number of IOPs that are needed for all disk unit types, use the total number of IOPs to plan for the number of buses that are needed.

Planning additional hardware to achieve the level of protection

This topic provides items for you to consider to ensure adequate protection.

Consider the following information to ensure adequate protection.

Disk-unit level protection

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

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.

IOA-level protection

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.

IOP-level protection

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

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

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

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.

Related concepts

Disk-unit level protection

Mirrored protection always provides disk-unit level protection because the disk units are duplicated.

Input/output bus level protection

Determine whether you want I/O bus level protection.

IOA level protection

Determine whether you want IOA level protection.

IOP-level protection

Determine if you want IOP-level protection.

Bus-level protection

Determine if you want bus-level protection.

Expansion-unit level protection

Determine if you want expansion-unit level protection.

Ring-level protection

Determine if you want ring-level protection.

Determining the extra hardware needed for performance

Use this information to decide how much extra hardware you might need.

Mirrored protection normally requires additional disk units and IOPs. However, in some cases, you might need additional hardware to achieve the level of performance that you want.

Processing disk unit requirements

Mirrored protection causes a minor increase in central processing disk 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.

IOP requirements

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

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

Ordering and installing 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 *Installing IBM i features* for installation instructions.

Setting up mirrored protection

Protect your disk units by setting up your system with mirrored protection.

Preparing your system for remote mirroring

To provide protection against a site disaster, all disk units in all disk pools of the system must be mirrored in local-remote pairs.

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. 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.

Starting mirrored protection

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

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

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

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

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

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

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

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

Starting remote Load Source mirroring

This topic provides instructions for starting remote load source mirroring.

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.

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

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

4. Select **Start Mirroring**.

Managing mirrored protection

Refer to this topic to see the tasks that you can perform to manage your disk protection activities.

Stopping mirrored protection

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

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

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

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

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

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

Suspending mirrored protection

If a disk unit in a mirrored pair fails, you must 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 must restart your system to the dedicated service tools (DST) mode before suspending mirroring.

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

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

1. Select **Configuration and Service** from your Navigator for i.
2. Select **Disk Units**.
3. Right-click the disk unit that you want to suspend.
4. Select **Suspend Mirroring**.

Enabling 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 IOPs 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.

To enable remote load source mirroring using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Pools**.
3. Right-click **Disk Pool 1**.
4. Select **Open**.
5. Right-click the load source disk unit.
6. Select **Enable Remote Load Source Mirroring**.

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

1. From the DST Main Menu, select **Work with disk units**.
2. From the Work with disk units menu, select **Work with disk configuration**.
3. From the Work with disk configuration menu, select **Work with mirrored protection**.
4. From the Work with mirrored protection menu, select **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.

Finding 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.

Changing 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.

Hot spare protection

Protect your disk units with hot spare protection.

Hot spare protection concepts

Hot spare disk units are spare disk units stored on a system to replace failed disks in case a disk failure occurs.

Hot spare disk units can be used to protect both device parity protected disk units and mirror-protected disk units.

A hot spare disk unit is stored on the system as a non-configured disk. When a disk failure occurs, the system exchanges the hot spare disk unit with the failed disk unit.

For device parity protected disk units, the hot spare disk unit must be the same capacity as the failed disk unit for a Small Computer System Interface (SCSI) IOA, or the same or larger capacity for a Serial Attached SCSI (SAS) IOA. The hot spare disk unit must also be under the same IOA in order for the exchange to occur. After the exchange occurs, the system rebuilds the data on the new disk unit.

For mirror-protected disk units, the hot spare disk unit must be the same capacity as the failed disk unit in order for the exchange to occur. The exchange of a mirrored subunit with the hot spare disk unit does not occur until mirror-protection has been suspended for 5 minutes and the replacement disk has been formatted. After the exchange occurs, the system synchronizes the data on the new disk unit. If multiple mirror-protected disk units are suspended, the system will not exchange a mirrored subunit with the hot spare disk unit; the hot spare disk unit will only be used if there is a single suspended mirror-protected disk unit.

A hot spare disk unit can be created manually from non-configured disk units on your system to protect existing or future mirror-protected units and device parity protected units. You can select which disk units and how many become hot spares. [“Starting hot spare protection” on page 47](#) explains how to add hot spare disk units to your system. When this option is selected, you must determine if one or more hot spares are desired and which available units become hot spares. The number is based on the total number of disk units attached to the IOA and their capacity.

When using hot spare protection with device parity protection, the system can automatically select and configure the hot spare disk unit when you initially start device parity protection. [“Starting device parity protection with hot spare protection” on page 21](#) explains how to include hot spare disk units when starting device parity protection. When selecting this option, the system automatically determines if one or two hot spares are created and which disk units are selected based on the total number and capacity of disk units attached to the IOA. When dealing with complex configurations such as those with varying disk unit capacities attached to the same IOA, you can start hot spare manually as described previously instead of allowing the system to select the hot spare disk units.

Note: The hot spare disk units are not designated to any particular parity set. The hot spare disk unit protects the first failed disk unit that has parity protection, is the appropriate capacity for the hot spare disk unit, and is under the same IOA as the hot spare disk unit.

Costs and limitations of hot spare protection

There are costs and limitations to consider when using hot spare protection.

- Hot spare disk units protect only parity sets with the same capacity disk units (for SCSI IOAs), or protect only parity sets with the same or smaller capacity disk units (for SAS IOAs).
- Hot spare disk units protect only mirror-protected units where the capacity of the hot spare is the same capacity as required when replacing a subunit in a mirrored pair.
- In order to create a new RAID 6 parity set with hot spare disk units, a minimum of 5 disk units is needed. If there are not at least 5 disk units, then the system will recommend the creation of a RAID 5 parity set with hot spare instead.

All SAS adapters support hot spare. Some SCSI adapters support hot spare.

<i>Table 1. SCSI adapters supporting hot spare</i>
SCSI adapters (CCIN)
2782
5703
571B
2757
2780
571E
571F
5709
573D

Planning for hot spare protection

Successfully plan how to create hot spare disk units.

In order to use hot spare for your parity protected disk units, the following requirements must be met:

- The disk units must be parity protected.
- The hot spare disk unit must be under the same IOA as the disk units that you want protected.
- The hot spare disk unit must be the same capacity as the failed parity protected disk unit (for SCSI IOAs) or must be the same or larger capacity as the failed parity protected disk unit (for SAS IOAs).
- When an IOA controls the load source disk unit, the hot spare disk unit must be in a valid load source location. This may require additional planning when using a SCSI IOA since the valid load source locations are typically a subset of the possible disk unit locations within the enclosure. The system will not allow the **Start device parity protection with Hot Spare** to occur if this requirement is not met.
- The hot spare disk unit must be a non-configured and unprotected disk unit.

In order to use hot spare for your mirror protected disk units, the following requirements must be met:

- The disk units must be mirror protected.
- The hot spare disk unit is under any IOA that supports hot spare.
- The hot spare disk capacity must meet the same requirement as when replacing a subunit in a mirrored pair. The hot spare capacity must be the same or slightly larger than the suspended disk unit.
- The hot spare disk unit must be a non-configured and unprotected disk unit.

Starting hot spare protection

This topic provides instructions for including hot spare disk units to your system.

To start hot spare protection using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Units**.
3. Right-click the disk unit that you want to work with.
4. Select **Start Hot Spare**.

To start hot spare protection with the command line, take the following steps:

1. Start System Service Tools (STRSST), and specify the user name and password.
2. On the System Service Tools (SST) display, select **Work with disk units**.
3. On the Work with Disk Units display, select **Work with disk configuration**.
4. On the Work with Disk Configuration display, select **Start hot spare**.

Managing hot spare protection

This topic describes how to manage hot spare protection.

Stopping hot spare protection

This topic provides instructions for stopping hot spare on a disk unit.

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

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Units**.
3. Right-click the disk unit that you want to work with.
4. Select **Stop Hot Spare**.

To stop hot spare protection with the command line, follow these steps.

1. Start System Service Tools (STRSST), and specify the user name and password.
2. On the System Service Tools (SST) display, select **Work with disk units**.
3. On the Work with Disk Units display, select **Work with disk configuration**.
4. On the Work with Disk Configuration display, select **Stop hot spare**.

Determining which parity sets are hot spare protected

This topic provides instructions for determining which parity sets have hot spare protection.

To determining which parity sets are hot spare protected with IBM Navigator for i, take the following steps:

1. Select **Configuration and Service** from your IBM Navigator for i window.
2. Select **All Tasks>Disk Management**.
3. Select **Parity Sets**.

The **Hot Spare Protected** column indicates which parity set has hot spare protection.

To determining which parity sets are hot spare protected with the command line, take the following steps:

1. Start System Service Tools (STRSST), and specify the user name and password.
2. On the System Service Tools (SST) display, select **Work with disk units**.
3. On the Work with Disk Units display, select **Work with disk configuration**.
4. On the Work with Disk Configuration display, select **Work with device parity protection**.
5. On the Device parity protection display, select **Display disk configuration**.
6. On the Display disk configuration display, select **Display Device Parity Status**.

Determining which disk units are hot spare protected

This topic provides instructions for determining which mirror protected and parity protected disk units are hot spare protected.

To determine which disk units are hot spare protected with your IBM Navigator for i, take the following steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Units**. The **Hot Spare Protected** column indicates which disk units are hot spare protected.
3. If the **Hot Spare Protected** column is not shown, select **Actions>Columns**. Select **Hot Spare Protected** from the Available columns list box and click either **Add Before** or **Add After**. Click **Ok** to complete the operation.

To determine which disk units are hot spare protected with the command line, take the following steps:

1. Start System Service Tools (STRSST), and specify the user name and password.
2. On the System Service Tools (SST) display, select **Work with Disk Units**.

3. On the Work with Disk Units display, select **Display disk configuration**.
4. On the Display Disk Configuration display, select **Display disk configuration protection**.

Displaying hot spare status

This topic provides instructions for displaying the hot spare status.

To Display the hot spare status using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Units**.

The **Status** column indicates the hot spare status.

To displaying hot spare status using the command line, take the following steps:

1. Start System Service Tools (STRSST), and specify the user name and password.
2. On the System Service Tools (SST) display, select **Work with disk units**.
3. On the Work with Disk Units display, select **Work with disk configuration**.
4. On the Work with Disk Configuration display, select **Display disk configuration**.
5. On the Display disk configuration display, select **Display hot spare disk unit status**.

Troubleshooting hot spare protection

Identify some common conditions associated with hot spare protection and failed drives.

- If you are having trouble starting hot spare protection, verify the following:
 - The disk units to become hot spares are in **Non-configured** status (not part of a disk pool).
 - The disk units to become hot spares are not in a device parity set.
 - If the disk units have been used before, old configuration information has been removed by initializing and formatting the disk units before starting hot spare protection.
 - If the desired IOA is the load source IOA, the disk units to become hot spares are installed in a valid load source location for the specific enclosure. See [Load Source Placement Rules](#) for information on valid load source locations.

Note: If you move a non-configured disk unit to a valid load source location and it is not shown as capable of becoming a hot spare, you need to initialize or format the disk unit in its new location in order for it to become a hot spare.

Note: If the option to start hot spare protection is selected and the previous requirements are not met, the desired disk units are not shown as capable of becoming hot spares.

- If you are having trouble starting device parity protection with hot spare protection, verify the following:
 - The same conditions as listed previously for starting hot spare.
 - There are an adequate number of disk units attached to the desired IOA which could become hot spare disk units. If more than 18 physical disk drives are attached to the IOA, you must have 2 valid hot spare candidates.
 - The disk units to become hot spares are the same or larger capacity as the disk units in the parity sets under the desired IOA.
 - If the option to start device parity protection with hot spare protection is selected and the previous requirements are not met, the desired IOA is not shown as capable of starting device parity protection.
- If you are trying to identify a failed disk unit which was automatically replaced by a hot spare, or if the hot spare disk unit itself has failed, see the following:
 - When a device parity protected disk unit fails and the automatic rebuild to the hot spare has started or completed, the failed disk will have a status of **Not operational** or **Read/Write protected**. Follow normal disk unit service and recovery procedures for replacement of the failed disk unit.

- When a hot spare disk unit fails, the failed disk will have a status of **Not operational** or **Read/Write protected**. Follow normal disk unit service and recovery procedures for replacement of the failed disk unit.

Related concepts

Hot spare protection concepts

Hot spare disk units are spare disk units stored on a system to replace failed disks in case a disk failure occurs.

Multipath disk units

Multiple connections can be defined from multiple Input/Output Adapters (IOAs) on a system to a single logical unit number (LUN) in the Enterprise Disk Storage.


Considerations for multipath disk units

There are many considerations when creating multipath disk units.

Multiple connections to the logical unit number (LUN) from the IOAs were implemented in V5R3. Up to eight connections can be defined from multiple IOAs on a system 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 IOAs and multipath connections between nodes. You must not split multipath connections between nodes, either by moving IOAs 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 IOAs to one LUN in the Enterprise Disk Storage. Partition 1 (P1) also has multiple paths defined from two different IOAs 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 IOA with a multipath connection to a different logical partition, you must also move all other IOAs 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 IOAs 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 IOAs, IOPs, 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.

Resetting multipath

The message *Event EV0D0401, Event Code 1E, Message ID CPI096E* will be in the QSYSOPR message queue when there is a missing connection to a disk unit.

Missing paths can be removed by using the multipath reset option. This option is available in Dedicated Service Tools (DST) under the Hardware Service Manager.

You can also use the Start DASD Management Operation (QYASSDMO) API to remove the missing multipath disk units.

Related information

Start [DASD Management Operation \(QYASSDMO\) API](#)

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

Disk protection checklists

There are checklists that you can use for disk protection tasks.

Adding a new IOA

This checklist shows the sequence of tasks that you use to add a new IOA and new disk units to your system.

Before you begin

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 [“Adding disk units without disk protection” on page 61](#) 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.

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 IOA in the system. This is normally done by a service representative.	“Evaluating the current configuration” on page 5
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.	“Evaluating the current configuration” on page 5

Task	What to do	Where to learn more
3.____	Use the Add Disk Unit wizard to add non-configured disks to the correct disk pools. Use the option to start device parity protection.	“Adding a disk unit or disk pool” on page 95
4.____	If you want a different storage threshold for any disk pool, change it. The default storage threshold for each disk pool is 90%.	“Setting the threshold of a disk pool” on page 101
5.____	You may want to restart to the dedicated service tools (DST) mode in order for the device parity protection to take effect.	“Device parity protection” on page 10
6.____	Wait until your system completely restarts.	
7.____	Verify that your disk configuration is correct.	“Evaluating the current configuration” on page 5
8.____	Print your disk configuration to have available in case a recovery situation occurs.	“Printing your disk configuration” on page 5

Adding disk units to an existing IOA

This checklist shows the sequence of tasks that you use to add one or more disks to an existing IOA 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 [“Adding disk units without disk protection” on page 61](#).

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.	“Evaluating the current configuration” on page 5
2.____	Use the Add Disk Unit wizard to add non-configured 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.	“Adding a disk unit or disk pool” on page 95
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.	“Setting the threshold of a disk pool” on page 101

Task	What to do	Where to learn more
4.____	Verify that your disk configuration is correct.	“Evaluating the current configuration” on page 5
5.____	Print your disk configuration to have available in case a recovery situation occurs.	“Printing your disk configuration” on page 5

Removing 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 IOA 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 system 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.	“Evaluating the current configuration” on page 5
2.____	Calculate the space requirements for the disk pools that are involved in disk removal.	“Calculating disk space requirements” on page 5
3.____	Use option 21 from the Save menu to save your entire system.	Save your system with the GO SAVE command
4.____	Restart your system 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 System i Navigator, select OpenSystem i Navigator Service Tools .
5.____	Remove disk units that you plan to remove from the system.	“Removing a disk unit from a disk pool” on page 99
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.	“Excluding disk units from a parity set” on page 23
7.____	Stop device parity protection for all the disk units in the Input/Output Processor (IOP).	“Stopping device parity protection” on page 22

Task	What to do	Where to learn more
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.	“Removing a disk unit from a disk pool” on page 99
9.____	Start device parity protection again.	“Starting device parity protection” on page 21
10.____	Verify that your disk unit configuration is correct.	“Evaluating the current configuration” on page 5
11.____	Print your disk configuration to have available in case a recovery situation occurs.	“Printing your disk configuration” on page 5
12.____	Restart your system.	

Removing 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 IOA 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 system 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.	“Evaluating the current configuration” on page 5
2.____	Calculate the space requirements for the disk pools that are involved in disk removal.	“Calculating disk space requirements” on page 5
3.____	Use option 21 from the Save menu to save your entire system.	Save your system with the GO SAVE command
4.____	Restart your system 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 System i Navigator, select Open System i Navigator Service Tools.
5.____	Remove disk units that you plan to remove from the system.	“Removing a disk unit from a disk pool” on page 99
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.	“Excluding disk units from a parity set” on page 23

Task	What to do	Where to learn more
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 Input/Output Processor (IOP) and have device parity protection.	“Stopping mirrored protection” on page 44
8.____	Stop device parity protection for all the disk units in the IOP.	“Stopping device parity protection” on page 22
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.	“Starting device parity protection” on page 21
11.____	Add unconfigured disk units to the correct disk pools. These disk units became unconfigured when you stopped mirrored protection in task 7.	“Adding a disk unit or disk pool” on page 95
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.	“Setting the threshold of a disk pool” on page 101
13.____	Start mirrored protection for the disk pools that had mirrored protection stopped in task 7.	“Starting mirrored protection” on page 43
14.____	Verify that your disk unit configuration is correct.	“Evaluating the current configuration” on page 5
15.____	Print your disk configuration to have available in case a recovery situation occurs.	“Printing your disk configuration” on page 5

Upgrading the 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 70-GB capacity while keeping device parity active. The disk 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 configuration by following the steps in [“Evaluating the current configuration” on page 5](#), and record your answers. The information that is entered in the table that is labeled Disk configuration questions, is needed to answer the questions in the Load source planning sections.

<i>Table 2. Disk configuration questions</i>	
Disk configuration questions	Disk configuration answers
1. 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.	
2. 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 parity set or optionally mark the disk units in the parity set. Make sure that you can identify the load source disk unit separately from the other disk units.	
3. 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.

<i>Table 3. Load source planning steps</i>	
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 disk units?	Yes / No
Do you know how to physically install and remove disk units from your system? <u>Install</u> and <u>remove</u> replacement disk units on an IBM i.	Yes / No
Do you have IBM Navigator for i 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 system using the <u>GO SAVE</u> command.

Task	What to do
2.____	<p>Turn off your system or logical partition by entering the following command.</p> <pre data-bbox="349 241 722 273">PWRDWN SYS *IMMED RESTART(*NO)</pre> <p>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.</p>
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. 2. Install replacement disk units on your system. <p>Notes:</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 70-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 disk 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.

Task	What to do																				
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 disk 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 12.</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 disk units on the screen and press Enter. 6. Press enter in the Confirm Add Units display. 7. The status display shows how the operation is proceeding. When the function is complete, return to the Use Dedicated Tools (DST) menu. 																				
7.____	<p>Locate the load source disk unit (unit 1) and all other disk units in the parity set that you will be upgrading.</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 Display disk configuration. 4. On the Display Disk Configuration display Select Display device parity status. <p>Note: If the status on the disk units in the parity set does not say Active, do not continue with this procedure. Contact the next level of support.</p> <ol style="list-style-type: none"> 5. Find the parity set that contains the load source disk unit (unit 1). 6. Record the disk units that are in its parity set. Your load source disk unit (unit 1) is already recorded for you. Recording this information is helpful for the next step. <p>Note: There could be up to 18 disk units.</p> <table border="1" data-bbox="370 1203 1128 1283"> <tr> <td>__1__</td><td>_____</td><td>_____</td><td>_____</td><td>_____</td><td>_____</td><td>_____</td><td>_____</td><td>_____</td><td>_____</td> </tr> <tr> <td>_____</td><td>_____</td><td>_____</td><td>_____</td><td>_____</td><td>_____</td><td>_____</td><td>_____</td><td>_____</td><td>_____</td> </tr> </table> <ol style="list-style-type: none"> 7. Return to the Use Dedicated Service Tools (DST) menu. 	__1__	_____	_____	_____	_____	_____	_____	_____	_____	_____	_____	_____	_____	_____	_____	_____	_____	_____	_____	_____
__1__	_____	_____	_____	_____	_____	_____	_____	_____	_____												
_____	_____	_____	_____	_____	_____	_____	_____	_____	_____												
8.____	<p>Remove the disk units from the configuration. You recorded the units in Task 7, Step 6.</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 Remove disk units from configuration. 5. Type a 4 next to each disk unit that you want to remove from the configuration. Select the disk units that were recorded in Task 7, Step 6 and press Enter. <p>Note: You will not be able to select load source disk unit (unit 1). You will remove the load source in Task 14.</p> <ol style="list-style-type: none"> 6. Press Enter in the Confirm Remove disk units display. 7. The status display shows how the operation is proceeding. When the function is complete, return to the Use Dedicated Service Tools (DST) menu. 																				

Task	What to do
9.____	<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 and press Enter. <p>Note: There is no command line because you are using Dedicated Service Tools.</p>
10.____	<p>Use the HMC to change the I/O tag.</p> <ol style="list-style-type: none"> 1. In your HMC, expand Systems Management > Servers. 2. Select the system. 3. Select the check box for the partition. 4. From the menu for the partition, click Configuration > Manage Profiles. 5. In the Managed Profiles pane, click the profile name. 6. In the Logical Partition Profile Properties pane, click the Tagged I/O tab. 7. From the Load source list, select the IOA that you want to control the Load Source disk unit. 8. Click OK. <p>Note: The changes that you make only go into effect when the partition is shut down and then activated again.</p>
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).
12.____	<p>Copy the load source disk unit (unit 1) to one of the replacement disk units that you installed in Task 3.</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. Write down the serial number of the replacement disk unit that is going to replace the load source. _____ 7. Press Enter in the Confirm Copy Disk Unit Data display. 8. The status display shows how the operation is proceeding. When the copy function is complete, the partition will automatically IPL to Dedicated Service Tools (DST).
13.____	<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 and press Enter. <p>Note: There is no command line because you are using Dedicated Service Tools.</p>

Task	What to do
14.____	Physically remove the disk units. <ol style="list-style-type: none"> 1. Physically remove the disk units that you recorded in Task 7, Step 6, and the load source disk unit (unit 1). <p>Note: You can identify which disk units are you removing from printout of the graphical view on your IBM Navigator for i or from the disk units that you marked.</p>
15.____	Work with Dedicated Service Tools (DST). <ol style="list-style-type: none"> 1. Power on your system or logical partition to Dedicated Service Tools (DST).
16.____	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 3. 6. Verify that the load source serial number matches the number that you wrote in Task 12, Step 6.
17.____	IPL your system to IBM i.

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.

Related reference

[Upgrading the load source disk unit with local mirroring using non-configured disk units](#)

This checklist shows the sequence of tasks to upgrade your load source disk unit (unit 1) and the mirrored load source with disk units that have equal capacity, are at least 70-GB, and are mirror protected. The disk units that are replaced in this procedure are discarded.

[Upgrading the load source disk unit with local mirroring using newly installed disk units](#)

This checklist shows the sequence of tasks to upgrade your load source disk unit (unit 1) and the mirrored load source with disk units that have equal capacity, are at least 70-GB, and are mirror protected. The disk units that are replaced in this procedure are discarded.

Adding disk units without disk protection

This checklist shows the sequence of tasks that you use to configure disks on a new system. Whether you need to perform all the tasks depends on the disk protection that you want on your system.

[“Disk protection” on page 9](#) 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 disk 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.	“Evaluating the current configuration” on page 5
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 non-configured 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.	“Adding a disk unit or disk pool” on page 95
4.____	Change the storage threshold for any disk pool if required. The default storage threshold for each disk pool is 90%.	“Setting the threshold of a disk pool” on page 101
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.	“Starting mirrored protection” on page 43
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.	“Evaluating the current configuration” on page 5
8.____	Print your disk configuration to have available in case a recovery situation occurs.	“Printing your disk configuration” on page 5

Moving 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 system 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.	“Evaluating the current configuration” on page 5
2.____	Calculate the space requirements for both the source and destination disk pools for the disk units.	“Calculating disk space requirements” on page 5
3.____	Use option 21 from the Save menu to save your entire system.	Save your system with the GO SAVE command
4.____	Restart your system 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 System i Navigator, select Open System i 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 .	“Removing a disk unit from a disk pool” on page 99
6.____	Add non-configured 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.	“Adding a disk unit or disk pool” on page 95
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.	“Setting the threshold of a disk pool” on page 101
8.____	If you created any new disk pools and you want those disk pools to have mirrored protection, start mirrored protection now.	“Starting mirrored protection” on page 43
9.____	Verify that your disk unit configuration is correct.	“Evaluating the current configuration” on page 5
10._ __	Print your disk configuration to have available in case a recovery situation occurs.	“Printing your disk configuration” on page 5

Upgrading the load source disk unit with local mirroring using non-configured disk units

This checklist shows the sequence of tasks to upgrade your load source disk unit (unit 1) and the mirrored load source with disk units that have equal capacity, are at least 70-GB, and are mirror protected. The disk 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 4. Disk configuration questions

Disk configuration questions	Disk configuration answers
<p>Where is the load source disk unit and the mirrored load source disk unit located? Determine the disk unit locations using Dedicated Service Tools (DST) or System Service Tools (SST).</p> <ol style="list-style-type: none"> From either Use Dedicated Service Tools (DST) or System Service Tools (SST) display, select Work with disk units. From the Work with Disk Units display, select Work with disk configuration. From the Work with Disk Configuration display, select Display disk configuration. From the Display Disk Configuration display, select Display disk configuration status. Write down the resource names (example DD001) of both the load source and load source mirrored unit (unit 1). At this point it is not necessary to know how to determine which resource name is assigned to the load source and the mirrored load source. Resource Name: _____ Resource Name: _____ From the Dedicated Service Tools (DST) or System Service Tools (SST) display, select Start a Service tool. From the Start a Service Tool display, select Hardware service manager. From the Hardware Service Manager, select Locate resource by resource name. Enter the first resource name that was written down in step 5 and press Enter. From the Logical Hardware Resources display, select Associated packaging resource(s). From the Packaging Resources Associated with a Logical Resource display, select Display detail. Write down the Physical location information for the disk unit. Resource Name: _____ Unit ID: _____ Planar ID: _____ Device: _____ Repeat steps 9-12 for the second resource name written down in step 5. Resource Name: _____ Unit ID: _____ Planar ID: _____ Device: _____ The physical locations of the card positions can be found by referencing the map located inside the front door of the IBM i or expansion tower. <p>Tip: Mark these disk units with tape to remember their locations.</p>	


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.

Load source planning steps	Planning requirements answers
Does your load source disk unit (unit 1) have mirrored protection?	Yes / No
Does your system or partition currently have two non-configured disk units with a capacity of at least 70-GB?	Yes / No
Do you know how to physically install and remove disk units from your system or partition? <u>Install and remove</u> replacement disk units on an IBM i.	Yes / No
Were you able to locate the load source and load source mirrored units in the Disk Configuration questions?	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 or partition 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 system using the <u>GO SAVE</u> command from the IBM i command line.
2.	Set the IPL mode to MANUAL.  Attention: Manual IPL is required for Dedicated Service Tools (DST).
3.	Restart your system or partition to Dedicated Service Tools (DST) by entering the following command from the IBM i command line. <pre>PWRDWN SYS *IMMED RESTART(*YES)</pre> Note: Verify that you are in manual mode before you restart your system.
4.	Begin Dedicated Service Tools. 1. From the IPL or Install the System menu, select Use Dedicated Service Tools . 2. Sign on to Dedicated Service Tools (DST).

Task	What to do
5.	<p>Suspend Mirroring on the mirrored load source disk unit (unit 1).</p> <ol style="list-style-type: none"> From the Dedicated Service Tools (DST) menu, select Work with disk units. From the Work with Disk Units display, select Work with disk unit recovery. From the Work with Disk Unit Recovery display, select Suspend Mirrored protection. Write down the resource name (example DD001) of the mirrored load source disk unit (unit 1). _____ Type 1 next to the mirrored load source disk unit (unit 1) and press Enter.
6.	<p>Label the load source and mirrored load source disk units using the physical locations determined in the Disk Configuration questions. If the physical locations were found using Dedicated Service Tools (DST) or System Service Tools (SST), determine the mirrored load source physical location by matching the resource name found in task 5 step 4 with the resource names found in the Disk Configuration questions steps 12-13. Once the mirrored load source disk unit is labeled, label the other disk unit the load source disk unit.</p>
7.	<p>Confirm that non-configured disk units are available as replacements for the load source disk unit and mirrored load source disk unit.</p> <ol style="list-style-type: none"> From the Use Dedicated Service Tools (DST) menu, select Work with disk units. From the Work with Disk Units display, select Work with disk configuration. From the Work with Disk Configuration display, select Display disk configuration. From the Display Disk Configuration display, select Display non-configured disk units. Write down the serial numbers and resource names of the non-configured disk units that are going to replace the load source and mirrored load source. <p>Load Source Replacement</p> <p>Serial Number: _____</p> <p>Resource Name: _____</p> <p>Mirrored Load Source Replacement</p> <p>Serial Number: _____</p> <p>Resource Name: _____</p> <p>Note: Ensure that there are two disk units of equal capacity and are at least 70-GB.</p>

Task	What to do
8.	<p>Determine the physical locations of the non-configured disk units that are replacing the load source and mirrored load source disk units that are documented in task 7 step 5.</p> <p>Determine the disk unit locations using Dedicated Service Tools (DST).</p> <ol style="list-style-type: none"> 1. From Use Dedicated Service Tools or System Service Tools display, select Start a Service Tool. 2. From the Start a Service Tool display, select Hardware service manager. 3. From the Hardware Service Manager, select Locate resource by resource name. 4. Enter the first resource name that was written down in task 7 step 5 and press Enter. 5. From the Logical Hardware Resources display, select Associated packaging resource(s). 6. From the Packaging Resources Associated with a Logical Resource display, select Display detail. 7. Write down the Physical location information for the load source replacement disk unit. Resource Name: _____ Unit ID: _____ Planar ID: _____ Device: _____ 8. Repeat steps 4-7 for the mirrored load source replacement that was written down in task 7 step 5. Resource Name: _____ Unit ID: _____ Planar ID: _____ Device: _____ 9. For internal disk units, the physical locations can be found by referencing the map located inside the front door of the IBM i or expansion tower.
9.	<p>For internal disk units, locate the physical locations of the non-configured disk units listed in task 8 step 7 through 8 and label each.</p> <p>Tip: Mark these disk units with tape to remember their locations.</p>
10.	<p>Power off the system or partition.</p> <ol style="list-style-type: none"> 1. From the Use Dedicated Service Tools (DST) menu, select Start a Service Tool. 2. From 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 (DST).</p>
11.	<p>Remove the mirrored load source disk unit.</p> <ol style="list-style-type: none"> 1. Physically remove the mirrored load source disk unit. <p>Note: This is the mirrored load source disk unit that was marked in task 6.</p>

Task	What to do
12.	<p>Power on your system or partition and begin Dedicated Service Tools (DST).</p> <ol style="list-style-type: none"> 1. Power on your system or partition. 2. From the IPL or Install the System menu, select Use Dedicated Service Tools. 3. Sign on to Dedicated Service Tools.
13.	<p>Copy the load source disk unit to the non-configured load source replacement disk unit written down in task 7 step 5.</p> <ol style="list-style-type: none"> 1. From the Use Dedicated Service Tools (DST) menu, select Work with disk units. 2. From the Work with Disk Units display, select Work with disk unit recovery. 3. From the Work with Disk Unit Recovery display, select Copy disk unit data. 4. Type 1 next to the load source disk unit (unit 1) and press Enter. 5. Type a 1 next to the non-configured disk unit documented as the load source replacement written down in task 7 step 5 and press Enter. 6. Use function F10 to accept the "Other sub-unit will become missing" error. 7. Verify that the correct non-configured disk unit was chosen as the Disk will contain the new copied information, and press Enter on the Confirm Copy Disk Unit Data display. 8. The status display shows how the operation is proceeding. When the function is complete, the partition will automatically IPL to the Use Dedicated Service Tools (DST) menu.
14.	<p>Power off the system or partition.</p> <ol style="list-style-type: none"> 1. From the Use Dedicated Service Tools (DST) menu, select Start a Service Tool. 2. From 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 (DST).</p>
15.	<p>Remove the original load source disk unit.</p> <ol style="list-style-type: none"> 1. Physically remove the original load source disk unit. <p>Note: This is the original load source disk unit that was marked in task 6.</p>
16.	<p>If the system has a Hardware Management Console (HMC), use the HMC to change the I/O tag.</p> <ol style="list-style-type: none"> 1. In your HMC, expand Systems Management > Servers. 2. Select the system. 3. Select the check box for the partition. 4. From the menu for the partition, click Configuration > Manage Profiles. 5. In the Managed Profiles pane, click the profile name. 6. In the Logical Partition Profile Properties pane, click the Tagged I/O tab. 7. From the Load source list, select the IOA that you controls the Load Source disk unit. The Load Source disk unit is the disk that was listed in task 7 step 5. 8. Click OK.
17.	<p>Power on your system or partition and begin Dedicated Service Tools (DST).</p> <ol style="list-style-type: none"> 1. Power on your system or partition. 2. From the IPL or Install the System menu, select Use Dedicated Service Tools. 3. Sign on to Dedicated Service Tools.

Task	What to do
18.	Replace the configured disk unit. <ol style="list-style-type: none"> 1. From the Use Dedicated Service Tools (DST) menu, select Work with disk units. 2. From the Work with Disk Units display, select Work with disk unit recovery. 3. From the Work with Disk Unit Recovery display, select Replace configured disk unit. 4. Type a 1 next to the suspended mirrored load source disk unit (unit 1) and press Enter. 5. Type a 1 next to the non-configured disk unit documented as the mirrored load source replacement in task 7 step 5. 6. Verify the correct non-configured disk unit was chosen, and press Enter on the Confirm Replace Configured Unit display. 7. The status display shows how the operation is proceeding. When the function is complete, return to the Use Dedicated Service Tools (DST) menu.
19.	Examine the configuration. <ol style="list-style-type: none"> 1. From the Use Dedicated Service Tools (DST) menu, select Work with disk units. 2. From the Work with Disk Units display, select Work with disk configuration. 3. From the Work with Disk Configuration display, select Display disk configuration. 4. From 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) and mirrored load source disk units match the serial numbers listed in task 7 step 5.
20.	IPL your system to IBM i.

If this procedure has been completed correctly, your load source will be upgraded. Mirrored protection will be active when the partition IPLs past DST.

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

Related reference

Upgrading the 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 70-GB capacity while keeping device parity active. The disk units that are replaced in this procedure are discarded.

Upgrading the load source disk unit with local mirroring using newly installed disk units

This checklist shows the sequence of tasks to upgrade your load source disk unit (unit 1) and the mirrored load source with disk units that have equal capacity, are at least 70-GB, and are mirror protected. The disk units that are replaced in this procedure are discarded.

Upgrading the load source disk unit with local mirroring using newly installed disk units

This checklist shows the sequence of tasks to upgrade your load source disk unit (unit 1) and the mirrored load source with disk units that have equal capacity, are at least 70-GB, and are mirror protected. The disk units that are replaced in this procedure are discarded.



Attention: These instructions are for all IBM i models.

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 6. Disk configuration questions

Disk configuration questions	Disk configuration answers
<p>Where is the load source disk unit and the mirrored load source disk unit located? Determine the disk unit locations using Dedicated Service Tools (DST) or System Service Tools (SST).</p> <ol style="list-style-type: none"> 1. From either Use Dedicated Service Tools (DST) or System Service Tools (SST) display, select Work with disk units. 2. From the Work with Disk Units display, select Work with disk configuration. 3. From the Work with Disk Configuration display, select Display disk configuration. 4. From the Display Disk Configuration display, select Display disk configuration status. 5. Write down the resource names (example DD001) of both the load source and load source mirrored unit (unit 1). At this point it is not necessary to know how to determine which resource name is assigned to the load source and the mirrored load source. Resource Name: _____ Resource Name: _____ 6. From the Dedicated Service Tools (DST) or System Service Tools (SST) display, select Start a Service tool. 7. From the Start a Service Tool display, select Hardware service manager. 8. From the Hardware Service Manager, select Locate resource by resource name. 9. Enter the first resource name that was written down in step 5 and press Enter. 10. From the Logical Hardware Resources display, select Associated packaging resource(s). 11. From the Packaging Resources Associated with a Logical Resource display, select Display detail. 12. Write down the Physical location information for the disk unit. Resource Name: _____ Unit ID: _____ Planar ID: _____ Device: _____ 13. Repeat steps 9-12 for the second resource name written down in step 5. Resource Name: _____ Unit ID: _____ Planar ID: _____ Device: _____ 14. The physical locations can be found by referencing the map located inside the front door of the IBM i or expansion tower. <p>Tip: Mark these disk units with tape to remember their locations.</p>	

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.

Load source planning steps	Planning requirements answers
Does your load source disk unit (unit 1) have mirrored protection?	Yes / No
Does your system or partition currently have at least one available disk slot to install a load source replacement disk unit or a virtual storage LUN?	Yes / No
Do you know how to physically install and remove disk units from your system? <u>Install and remove</u> replacement disk units on an IBM i.	Yes / No
Where you able to locate the load source and load source mirrored units in the Disk Configuration questions?	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 or partition 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 system using the <u>GO SAVE</u> command from the IBM i command line.
2.	Set the IPL mode to MANUAL. Note: Manual IPL is required for Dedicated Service Tools (DST).
3.	Restart your system or partition to Dedicated Service Tools (DST) by entering the following command from the IBM i command line. <pre>PWRDWN SYS *IMMED RESTART(*YES)</pre> Note: Verify that you are in manual mode before you power down your system.
4.	Begin Dedicated Service Tools. 1. From the IPL or Install the System menu, select Use Dedicated Service Tools . 2. Sign on to Dedicated Service Tools (DST).
5.	Suspend Mirroring on the mirrored load source disk unit (unit 1). 1. From the Dedicated Service Tools (DST) menu, select Work with disk units . 2. From the Work with Disk Units display, select Work with disk unit recovery . 3. From the Work with Disk Unit Recovery display, select Suspend Mirrored protection . 4. Write down the resource name (example DD001) of the mirrored load source disk unit (unit 1) that you are suspending. _____ 5. Type 1 next to the mirrored load source disk unit (unit 1) and press Enter.

Task	What to do
6.	<p>1. Write down the resource name of the load source disk unit that you suspended in task 5. Resource Name: _____</p> <p>2. From the Dedicated Service Tools (DST) display, select Start a Service tool.</p> <p>3. From the Start a Service Tool display, select Hardware service manager.</p> <p>4. From the Hardware Service Manager, select Locate resource by resource name.</p> <p>5. Enter the resource name that was written down in step 1 and press Enter.</p> <p>6. From the Logical Hardware Resources display, select Associated packaging resource(s).</p> <p>7. From the Packaging Resources Associated with a Logical Resource display, select Display detail.</p> <p>8. Write down the Physical location information for the disk unit that was suspended in task 5. Resource Name: _____ Unit ID: _____ Planar ID: _____ Device: _____</p>
7.	<p>Power off your system or partition.</p> <p>1. From the Use Dedicated Service Tools (DST) display, select Start a Service Tool.</p> <p>2. From the Start a Service Tool display, select Operator panel functions.</p> <p>3. Use function F10 to power off the system.</p> <p>4. Press Enter.</p> <p>Note: There is no command line because you are using Dedicated Service Tools.</p>
8.	<p>Install the first replacement disk unit.</p> <p>1. Install the replacement disk unit into the available disk slot or assign a virtual storage LUN.</p> <p>2. Write down the serial number of the replacement disk unit that is to going replace the load source. _____ </p> <p>Note: Ensure that the disk unit has at least a 70-GB capacity.</p>
9.	<p>Power on your partition and begin Dedicated Service Tools (DST).</p> <p>1. Power on your partition.</p> <p>2. Select Use Dedicated Service Tools (DST) from the IPL or Install the System menu.</p> <p>3. Sign on to Dedicated Service Tools (DST).</p>

Task	What to do
10.	<p>Copy the load source disk unit (unit 1) to the replacement disk unit.</p> <ol style="list-style-type: none"> From the Use Dedicated Service Tools (DST) menu, select Work with disk units. From the Work with Disk Units display, select Work with disk unit recovery. From the Work with Disk Unit Recovery display, select Copy disk unit data. Type 1 next to the load source disk unit (unit 1) and press Enter. Write down the resource name of the load source disk unit that you selected: _____ Type 1 next to the replacement disk unit that was installed in task 8 and press Enter. Use function F10 to accept the "Other sub-unit will become missing" error. Verify the correct non-configured disk unit was chosen as the Disk that is copied to, and press Enter on the Confirm Copy Disk Unit Data display. The status display shows how the operation is proceeding. When the function is complete, the partition will automatically IPL to the Use Dedicated Service Tools (DST) menu.
11.	<ol style="list-style-type: none"> From the Dedicated Service Tools (DST) or System Service Tools (SST) display, select Start a Service tool. From the Start a Service Tool display, select Hardware service manager. From the Hardware Service Manager, select Locate resource by resource name. Enter the resource name that was written down in task 10 step 4 and press Enter. From the Logical Hardware Resources display, select Associated packaging resource(s). From the Packaging Resources Associated with a Logical Resource display, select Display detail. Write down the Physical location information for the disk unit. Resource Name: _____ Unit ID: _____ Planar ID: _____ Device: _____
12.	<p>Power off the system or partition.</p> <ol style="list-style-type: none"> From the Use Dedicated Service Tools (DST) menu, select Start a service tool. From the Start a Service Tool display, select Operator panel functions. Use function F10 to power off the system. Press Enter to continue.
13.	<p>Physically remove the original load source disk unit.</p> <ol style="list-style-type: none"> Physically remove the original load source disk unit. <p>Note: This is the original load source disk unit that you copied in task 10, and identified in task 11.</p>
14.	<p>Install the second replacement disk unit that will be the mirror mate to the load source disk unit.</p> <ol style="list-style-type: none"> Install the replacement disk unit into the available disk slot on your system or assign a virtual storage LUN of the same size as the virtual storage LUN in task 8. Write down the serial number of the replacement disk unit that will be the mirror mate to the load source disk unit. _____ <p>Note: Ensure that the disk unit has at least a 70-GB capacity.</p>

Task	What to do
15.	<p>If the system has a Hardware Management Console (HMC), use the HMC to change the I/O tag.</p> <ol style="list-style-type: none"> 1. In your HMC, expand Systems Management > Servers. 2. Select the system. 3. Select the check box for the partition. 4. From the menu for the partition, click Configuration > Manage Profiles. 5. In the Managed Profiles pane, click the profile name. 6. In the Logical Partition Profile Properties pane, click the Tagged I/O tab. 7. From the Load source list, select the IOA that you want to control the Load Source disk unit. 8. Click OK.
16.	<p>Power on your system or partition and begin Dedicated Service Tools (DST).</p> <ol style="list-style-type: none"> 1. Power on your system or partition. 2. Select Use Dedicated Service Tools (DST) from the IPL or Install the System menu. 3. Sign on to Dedicated Service Tools (DST).
17.	<p>Replace the configured disk unit.</p> <ol style="list-style-type: none"> 1. From the Use Dedicated Service Tools (DST) menu, select Work with disk units. 2. From the Work with Disk Units display, select Work with disk unit recovery. 3. From the Work with Disk Unit Recovery display, select Replace configured disk unit. 4. Type a 1 next to the suspended mirrored load source disk unit (unit 1) and press Enter. 5. Type a 1 next to the replacement disk unit, based on the serial number, that was installed in task 14. 6. Verify the correct non-configured disk unit was chosen, and press Enter on the Confirm Replace Configured Unit display. 7. The status display shows how the operation is proceeding. When the function is complete, return to the Use Dedicated Service Tools (DST) menu.
18.	<p>Examine the configuration.</p> <ol style="list-style-type: none"> 1. From the Use Dedicated Service Tools (DST) menu, select Work with disk units. 2. From the Work with Disk Units display, select Work with disk configuration. 3. From the Work with Disk Configuration display, select Display disk configuration. 4. From 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) and mirrored load source disk units match the serial numbers listed in task 8 and task 14.
19.	<p>IPL your system to IBM i.</p>

If this procedure has been completed correctly, your load source will be upgraded. Mirrored protection will be active when the partition IPLs past DST.

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

Related reference

[Upgrading the 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 70-GB capacity while keeping device parity active. The disk units that are replaced in this procedure are discarded.

[Upgrading the load source disk unit with local mirroring using non-configured disk units](#)

This checklist shows the sequence of tasks to upgrade your load source disk unit (unit 1) and the mirrored load source with disk units that have equal capacity, are at least 70-GB, and are mirror protected. The disk units that are replaced in this procedure are discarded.

NVMe failure with disk units mirror protected

When the NVMe card fails, the disk units (namespaces) that are mirror-protected will be suspended. Normal system operations continue on the other NVMe card.

Task	What to do
1.	<p>Replace the failing card and install a new NVMe device.</p> <ol style="list-style-type: none"> 1. Start System Service Tools (STRSST), and specify the user name and password. 2. On the System Service Tools (SST) display, select Start a Service Tool. 3. On the Start a Service Tool display, select Hardware service manager. 4. On the Hardware Service Manager display, select Packaging hardware resources. 5. On the Packaging Hardware Resources display, select 3=Concurrent maintenance for the Unit ID that contains the failed NVMe device. 6. Press F9 to power off the domain. 7. Physically unplug and remove the failing NVMe device. 8. Physically install the replacement NVMe device. 9. Press F10 to power on the domain.
2.	<p>Prepare the newly installed NVMe card.</p> <ol style="list-style-type: none"> 1. Return to the System Service Tools (SST) display. 2. On the System Service Tools (SST) display, select Work with disk units. 3. On the Work with Disk Units display, select Work with disk configuration. 4. On the Work with Disk Configuration display, select Work with NVM Devices. 5. On the Work with NVM Devices display, select Delete existing NVM Namespaces. 6. Select the device that you installed in task 1 step 8 above. 7. If no namespaces are listed, press F12 and then go to task 3. 8. Select 4=Delete Namespace for each of the listed namespaces. 9. F10 to confirm the delete of the namespaces. 10. F12 to return to SST.

Task	What to do
3.	<p>Resume mirroring.</p> <ol style="list-style-type: none"> 1. Return to the System Service Tools (SST) display. 2. On the System Service Tools (SST) display, select Work with disk units. 3. On the Work with Disk Units display, select Work with disk configuration. 4. On the Work with Disk Configuration display, select Work with NVM Devices. 5. On the Work with NVM Devices display, select Create NVM namespaces to pair with Active mirror protected NVM units. 6. On the Select NVM Device display, select the device that you installed in task 1 step 8 above. 7. On the Create NVM Namespaces - Show the Units Suspended from Mirror Protection display, press F7 to select all of the suspended units. 8. Confirm the creation of the listed namespaces that will be paired with the suspended disk units. 9. Press F10 to confirm the operation. The operation will perform the following steps. <ol style="list-style-type: none"> a. Create new namespaces on the selected NVMe device. b. Pair the new namespaces with the suspended NVMe disk units. c. Start to resume mirror protection on the newly created disk units. 10. Press F12 to return to the System Service Tools (SST) display. 11. This completes the procedure.

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.

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 disk unit.

Your system may have many disk units attached to it for disk pool storage. To your system, they look like a single disk 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.

Using disk pool

Disk pools are used to manage system performance and backup requirements.

System performance and backup requirements:

- 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. The independent disk pool can be utilized only when it is needed.
- An independent disk pool in a clustered environment can provide disk storage that is switchable or replicated, allowing continuous availability of resources.

Types of disk pools

There are several 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 system. 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 IBM i 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 non-configured disk units.

Related concepts

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 system is up and running.

User disk pools

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 system is up and running.

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 system 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 or Device parity protection. See Disk protection for more information.

Related concepts

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.

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.

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 system, unless you include code (Example: Make independent disk pool available at startup) 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 system. 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 system jobs are started to support the independent disk pool. In order for system jobs to remain unique on the system, those that service the independent disk pool are given their own simple job name when the disk pool is made available. The system jobs are essential to the operation of the disk pool. Do not tamper with these system jobs. The following is a list of system jobs that are created (nn = number):

QDBXnnnXR

handles database cross-reference file system functions

QDBXnnnXR2

handles database cross-reference field (column) information

QDBnnnSV01

handles database, journal, and commitment control events

QDBnnnSV02 through QDBnnnSVnn

high priority jobs that service the database

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.

Related reference

[Example: Making independent disk pool available at startup](#)

You can decide whether to have your independent disk pool available when the system is restarted.

Primary, Secondary, UDFS disk pools

An independent disk pool that contains user-defined file systems, directories and libraries, or associated directory and library information.

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 IBM i 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 IBM i V5R3 and later.

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 IBM i V5R3 and later.

Disk pool concepts

A disk pool is a software definition of a group of disk units on your system.

Contrasting basic and independent disk pools

Basic disk pools and independent disk pools have some inherent differences.

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 system performs an IPL, all of the disk units configured to a basic disk pool must be accounted for in order for the system 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 system 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 system 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 system restarted to Dedicated Service Tools (DST). In an offline independent disk pool you do not need to have your system in DST mode to start or stop mirroring, start device parity protection, start compression, remove a disk unit, and so on.

Disk pool benefits

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

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 disk pool 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 concepts

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.

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 system is up and running.

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 (or multiple logical partition) clustered environment, where the systems are members of a 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 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 cross-site mirroring. Cross-site mirroring allows you to maintain two identical copies of an independent disk pool at two sites. By having a second copy of critical data at a second 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.

The use of FlashCopy® is also possible in a multiple-system environment. FlashCopy allows you to make an instant or point-in-time copy of an independent disk pool.

Single-system environment

In a single-system environment, where an independent disk pool is privately connected to a single system, an 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 tasks

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.

Related information

PowerHA data replication technologies

FlashCopy

Disk pool costs and limitations

This topic lists the costs and limitations of using disk pools.

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.

Planning for disk pools

This topic provides instructions about how to create disk pools.

Using disk pools for improved performance

There are several ways that you can use disk pools to improve disk pool 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 disk unit in a user disk pool because the performance of that disk unit is affected by other disk units in the parity set.

Another way to improve performance is to make sure there are enough disk 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 disk units are used excessively. For more information about working with disk status (WRKDSKSTS command) to determine if the disk units have excessive use, see Work Management. If the disk units have excessive use, you should consider adding more disk units to the user disk pool.

Related information

[Work Management](#)

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 (STRJRNP) 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.

Related information

[System-managed access-path protection](#)

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.

Related information

[System-managed access-path protection](#)

System disk pool

You need to balance your system disk pool, protect your system disk pool, and prevent your system disk pool from overfilling.

Balancing a disk pool

You can improve the system 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, see Recovering your system.

You can balance the capacity of disk pools using the Add a disk unit or disk pool wizard.

Related tasks

[Adding a disk unit or disk pool](#)

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

Related information

[Recovering your system](#)

Protecting 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.

Capacity of the system disk pool

There are several ways to prevent the system disk pool from filling to capacity.

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 another method of preventing the system disk pool from filling to capacity by using the QSTGLOWLMT and QSTGLOWACN system values. For more information, refer to [Changing the storage threshold for the system auxiliary storage pool \(ASP\)](#).

Related information

[Changing the storage threshold for the system auxiliary storage pool \(ASP\)](#)

Planning for single system independent disk pools

Several requirements must be satisfied to use single system independent disk pools.

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

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

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.

Software requirements for single system independent disk pools

If you plan to use independent disk pools in a single-system environment, the following software element is required.

The following graphical user interface is required to perform some of the disk management tasks necessary to implement independent disk pools.

- IBM Navigator for i

Application considerations for independent disk pools

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

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

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

- Use the Set ASP Group (SETASPGRP) command.
- In an SQL environment, use CONNECT to specify the right database. To achieve the fastest performance, make sure that the database to which you perform an SQL CONNECT operation corresponds with your current library namespace. You may need to use the SETASPGRP command first to achieve this. If the SQL CONNECT function is not operating within the same library namespace, the application uses Distributed Relational Database Architecture™(TM) support, which can affect performance.
- Use the Change Job Description (CHGJOB) command to set the initial ASP group in the job description for a user profile.

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

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

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

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

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

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

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

Related concepts

[Independent disk pools with distinct databases](#)

Here gives you an example of independent disk pools with distinct databases.

Related tasks

[Multiple system libraries](#)

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

Related reference

[Supported and unsupported object types](#)

Supported and unsupported object types in independent disk pools.

Related information

[\(SETASPGRP\)](#)

[SQL Reference](#)

[Change Job Description \(CHGJOB\) command](#)

[CRTLIB](#)

[Cluster applications](#)

Storing and printing spooled files

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

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

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

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

Consider the following environment:

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

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

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

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

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

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

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

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

Related concepts

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.

Related information

Advanced Function Presentation

*FNTRSC

*OVL

*PAGDFN

SETASPGRP (Set ASP Group) command

STRPRTWTR (Start Printer Writer) command

Delete Spooled File (DLTSPLF)

Object identification

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

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

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

Control language (CL) commands

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

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

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

Related information

Display Object Description (DSPOBJD) command

Multiple system libraries

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

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

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.

QSYS2nnnnn

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

Normally only internal system code creates objects into this library.

QRCYnnnnn

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

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

QRCLnnnnn

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

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

QRPLnnnnn

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

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

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

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

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

The normal search order for objects is to search the libraries based on the user-specified library value, the user's library list, and the namespace in effect for the job. The only exception to this occurs when the

user job has a disk pool group in the job's namespace. In this case, aliasing support takes effect for object references to the database control objects in QSYS, QSYS2, and SYSIBM. The objects in the QSYSnnnnn, QSYS2nnnnn, and SYSIBnnnnn libraries are returned so that users are operating on the database control information associated with their extended namespace.

Supported and unsupported object types

Supported and unsupported object types in independent disk pools.

Objects not supported

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

*AUTHLR	*DEVD	*JOBSCD	*PRDLOD
*AUTL	*DOC	*LIND	*RCT
*CFGL	*DSTMF ¹	*MODD	*SOCKET
*CNL	*EDTD	*M36	*SSND
*COSD	*EXITRG	*M36CFG	*S36
*CRG	*FLR	*NTBD	*USRPRF
*CSPMAP	*IGCSRT	*NWID	
*CSPTBL	*IGCTBL	*NWSD	
*CTLD	*IMGCLG	*PRDAVL	
*DDIR	*IPXD	*PRDDFN	

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

Supported object types

The following objects are supported in independent disk pools:

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

*FIFO

*MGTCOL

*QRYDFN

*WSCST

Restrictions for supported object types

*ALRTBL

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

*CLS

If an active subsystem references the class object, the class must exist in the library name space of the job that is being started by that subsystem. The library name space of a prestart job must be the same as the library name space of the subsystem job.

*FILE

Database files that are either multiple-system database files, or that have DataLink fields that are created as Link Control, cannot be in an independent disk pool. If an active subsystem references the file object, *FILE must exist in the system disk pool; for example, the sign-on display file.

*JOB

If an active subsystem references the job description object, the job description must exist in the system disk pool or in the ASP group parameter of the subsystem description. A subsystem references a job description from an autostart job entry, communication entry, remote location name entry, a workstation entry, or a prestart job entry. When the job description is referenced by a prestart job entry, the INLASGRP parameter for the job description must match the ASPGRP parameter of the subsystem description.

*JOBQ

Jobs that are contained in a job queue on an independent disk pool are ended when the independent disk pool is varied off or if an IPL occurs for the system. Jobs do not switch with the independent disk pool group to another system. To submit a job to a job queue, the job queue must exist in the library name space of the thread that issued the SBMJOB command. For a subsystem to start jobs from a job queue, the job queue must exist in the library name space of the subsystem job.

*LIB

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

*MSGQ

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

*PGM

If an active subsystem references the program object, the program must exist in the library name space of the job that is being started by that subsystem. The library name space of a prestart job must be the same as the library name space of the subsystem job.

*SBSD

The subsystem description can be restored to an Independent Disk Pool and changed while in the Independent disk pool, but it must be copied to the System disk pool to be used. You cannot start a subsystem whose description is in an independent disk pool. You cannot start a subsystem if the ASP group parameter in the subsystem description names an ASP group that is not varied on and available.

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 system disk pool and basic user disk pools, which is the system disk pool and all configured basic disk pools.

The system disk pool and basic user disk pools (SYSBAS) should 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

A system 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 extendibility 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](#) 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.

Related concepts

Examples: [Dedicated independent disk pools](#)

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

Configuring disk pools

This topic provides instructions for configuring your disk pools.

Creating a disk pool

You can create a basic disk pool or independent disk pool with the disk pool wizard.

The 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 system has that capability.

The 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 non-configured 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.

The disk pool wizard also allows you to configure disk encryption. You must set the ASP master key before you can create an encrypted independent disk pool.

To create a disk pool using your IBM Navigator for i, follow these steps:

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

Related concepts

[Requirements for disk management](#)

There are requirements that must be completed before you can configure or work with your disks.

[Disk encryption](#)

Disk encryption allows you to encrypt data that is stored in basic disk pools and independent disk pools.

Related information

[Loading and setting auxiliary storage pool \(ASP\) master key](#)

Adding a disk unit or disk pool

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

The Add Disk Unit and Disk Pool wizards save you time by bundling several time-consuming configuration functions into one efficient process. They also take the guesswork out of disk unit configuration because they understand the capabilities of your system and only offer valid choices. For instance, the wizard does not list the option to start compression unless your system has that capability.

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

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

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

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

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

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

Related concepts

[Requirements for disk management](#)

There are requirements that must be completed before you can configure or work with your disks.

Configuring independent disk pools

After you have satisfied the planning requirements for implementing independent disk pools, you are ready to configure an independent disk pool. You must use the IBM Navigator for i disk management function or the Configure Device ASP (CFGDEVASP) command to configure an independent disk pool.

Related concepts

[Requirements for disk management](#)

There are requirements that must be completed before you can configure or work with your disks.

Converting UDFS disk pools

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

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

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

Related concepts

[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.

Converting a UDFS disk pool to primary

You have the capability to convert UDFS disk pools to library-capable primary and secondary disk pools. Library-capable disk pools support library-based objects. You must convert UDFS disk pools if you want them to participate in a disk pool group.

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

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

To convert a UDFS disk pool to a primary disk pool using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Pools**.
3. Right-click the UDFS disk pool that you want to convert to a primary disk pool and select **Convert to Primary Disk Pool**.
4. 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.
5. Click **Convert Disk Pool**.
6. If you must associate other existing UDFS disk pools with your new primary disk pool in a disk pool group, see [“Converting a UDFS disk pool to secondary”](#) on page 96.

Converting a UDFS disk pool to secondary

You have the capability to convert UDFS disk pools to library-capable primary and secondary disk pools. Library-capable disk pools will support library-based objects. You must convert UDFS disk pools if you want them to participate in a disk pool group.

Before you create a secondary disk pool, you must already have created its primary disk pool.

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

To convert a UDFS disk pool to a secondary disk pool using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.

2. Select **Disk Pools**
3. Right-click the UDFS disk pool that you want to convert to a secondary disk pool and select **Convert to Secondary Disk Pool**.
4. On the **Confirm Convert to Secondary Disk Pool** dialog box, select the primary disk pool that you need to associate with the secondary disk pools.

The primary disk pool that is selected should not be made available before it is associated with the secondary disk pools. Only the primary disk pools that are currently owned by the system can be selected. You will not be able to change the primary after you perform this action.
5. Click **Convert Disk Pool**.
6. If the disk pool you converted to secondary is in a device cluster resource group, then you must change the Online attribute to *PRIMARY.

Use the Change Cluster Resource Group Device Entry (CHGCRGDEVE) command or the Change Cluster Resource Group Device Entry (QcstChgClusterResourceGroupDev) API to change the Online attribute to *PRIMARY.

Related information

[Change CRG Device Entry \(CHGCRGDEVE\) command](#)

[Change Cluster Resource Group Device Entry \(QcstChgClusterResourceGroupDev\) API](#)

Creating an independent disk pool

Create an independent disk pool.

To create an independent disk pool, you can use the New Disk Pool wizard or the Configure Device ASP (CFGDEVASP) command. These assist you in creating a new disk pool and adding disk units to it. The disk pool wizard and the command can be used to include unconfigured disk units in a parity set, and start device parity protection and disk compression. As you add disk units, do not spread disk units that are in same parity set across multiple disk pools, because failure to one parity set would affect multiple disk pools.

The disk pool wizard and the command can also be used to configure disk encryption. You must set the ASP master key before you can create an encrypted independent disk pool.

To use the disk pool wizard to create an independent disk pool using your IBM Navigator for i, follow these steps:

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

Select **Primary** or **Secondary** for the type of disk pool.
5. Print your disk configuration to have available in a recovery situation.
6. Record the relationship between the independent disk pool name and number.

To use the Configure Device ASP command to create the independent disk pool, follow these steps:

1. Type the CFGDEVASP command and press F4 to prompt.
2. Select the *CREATE option.
3. Select protection and encryption.
4. By using *SELECT for the UNITS parameter, a list of available disk units to choose is displayed.
5. Print your disk configuration to have available in a recovery situation.
6. Record the relationship between the independent disk pool name and number.

Note: Add independent disk pools when your system is fully restarted. If you must use the disk pool wizard in the dedicated service tools (DST) mode, you need to create an associated device description for the independent disk pool when the system is fully restarted. Use the Create Device Description (ASP) (CRTDEVASP) command to create the device description. Name the device description and resource

name the same as you name the independent disk pool. You can use the Work with Device Descriptions (WRKDEVDD) command to verify that the device description and independent disk pool name match.

Related concepts

[Requirements for disk management](#)

There are requirements that must be completed before you can configure or work with your disks.

[Disk encryption](#)

Disk encryption allows you to encrypt data that is stored in basic disk pools and independent disk pools.

[Primary, Secondary, UDFS disk pools](#)

An independent disk pool that contains user-defined file systems, directories and libraries, or associated directory and library information.

Related information

[Creating device description \(ASP\)](#)

[Working with device descriptions](#)

[Loading and setting auxiliary storage pool \(ASP\) master key](#)

Creating a new disk pool group

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

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

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

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

To create a new disk pool group using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Pools**.
3. From the **Actions** menu select **New Disk Pool**.
4. On the New Disk Pool dialog box that is displayed, 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.

Related tasks

[Converting a UDFS disk pool to primary](#)

You have the capability to convert UDFS disk pools to library-capable primary and secondary disk pools. Library-capable disk pools support library-based objects. You must convert UDFS disk pools if you want them to participate in a disk pool group.

[Converting a UDFS disk pool to secondary](#)

You have the capability to convert UDFS disk pools to library-capable primary and secondary disk pools. Library-capable disk pools will support library-based objects. You must convert UDFS disk pools if you want them to participate in a disk pool group.

Managing disk pools

This topic provides instructions for managing your disk pool activities.

Deleting 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.

To guarantee that there are no dependencies on objects in the disk pool being deleted, delete all objects in the disk pool before removing it. Check failure messages for indications that objects outside of the disk pool might be impacted.

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 that you remove it from the CRG afterward.

To delete a disk pool using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Pools**.
3. Right-click the disk pools that you want to delete and select **Delete**.
4. Follow the instructions on the dialog box that is displayed.

To delete an independent disk pool using the Configure Device ASP (CFGDEVASP) command, follow these steps:

1. Type in the Configure Device ASP command and press F4 to prompt.
2. Type the independent ASP name and use the *DELETE option.

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

Related information

[Remove Cluster Resource Group Device Entry \(RMVCRGDEVE\) command](#)

[End Cluster Resource Group \(ENDCRG\) command](#)

Removing a disk unit from a disk pool

As your storage needs change, you can select to remove a disk unit from a disk pool.

When you remove a disk unit, the system redistributes the data on that disk unit to other disk units in the disk pool. The removal is not allowed if there is not enough disk capacity in the disk pool to store all the data from the removed disk unit. You cannot remove all the disk units from a disk pool.

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

To remove a disk unit from a disk pool using System Service Tools (SST), follow these steps:

1. Start SST (**STRSST**).

2. Select **Work with disk units**.
3. Select **Work with disk configuration**.
4. Select **Work with removing units from configuration**.
5. Select **Remove units from configuration**
6. Select the disk unit that you want to remove from the configuration.
7. Confirm the function.

IBM Navigator for i can be used to remove a disk unit from a varied off independent disk pool.

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Units**.
3. Right-click the disk unit that you want to remove and select **Remove**.
4. Follow the instructions on the dialog box that is displayed.

IBM Navigator for i can also be used to remove a disk unit from system disk pools, basic disk pools, and independent disk pools that are varied on.

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Pools**.
3. Right-click the disk pool that you want to remove disk units from and select **Remove Disk Units**.
4. Follow the instructions on the dialog box that is displayed.

Moving a disk unit from a disk pool

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 system first moves all the data on that disk unit to other disk units in the original disk pool. The move is not allowed if there is not enough disk capacity in the disk pool to store all the data from the disk unit to move. You cannot move disk units to or from an independent disk pool. For disk units in system disk pools and basic disk pools, you must restart your system to dedicated service tools (DST) mode before moving them.

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

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

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Units**.
3. Right-click the disk unit that you want to move and select **Move**.
4. Follow the instructions on the dialog box that is displayed.

Clearing 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 using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Pools**.
3. Right-click the disk pools that you want to clear and select **Clear**.
4. Follow the instructions on the dialog box that is displayed.

Setting 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 system sends a message so that you have time to add more storage space or delete unnecessary objects.

To change the threshold of a disk pool using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Pools**.
3. Right-click the disk pool that you want to change the threshold for and select **Properties**.
4. From the **Threshold** tab, increase or decrease the threshold for the disk pool.

Tracing disk pool data

You can trace the data in your disk pool to collect usage statistics about the data.

You can trace the data in disk pools. When trace is started on the disk pool, the system collects statistics on the usage of the data in the disk pool. The trace function monitors the frequency that the data on the units in the specified disk pool is accessed. The data that is accessed frequently and the data that is accessed infrequently are identified.

To start tracing on a disk pool using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Pools**.
3. Right-click the disk pool that you want to start tracing on and select **Start Tracing**.
4. Follow the instructions on the dialog box that is displayed.

You can also stop tracing data in disk pools. When trace is stopped on the disk pool, the collecting of the usage statistics about the data in your disk pool is stopped. The trace can be stopped at any time and restarted later. The statistics that are collected are cumulative. For example, if the trace is started and ended and then restarted without clearing the statistics, the second collection of statistics is added to the first collection.

To stop tracing on a disk pool using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Pools**.
3. Right-click the disk pool that you want to stop tracing on and select **Stop Tracing**.
4. Follow the instructions on the dialog box that is displayed.

Balancing disk pool data

You can balance the data in your disk pool using various balance options.

You can start or stop the balancing of data in disk pools by selecting one of several balance options.

- ***CAPACITY** moves the data in the disk pool so that the disk units in the disk pool are approximately the same percent full.

Note: If there are Solid State Disks (SSD) and High Density Disk (HDD) disk types in the disk pool, the data on the SSDs is balanced among the SSD disks, and the data on the HDDs is balanced among the HDD disks.

- ***HSM** moves the data in the disk pool so that the frequently used data is moved to the SSD units, and the infrequently used data is moved off the SSD disks onto the HDD disks.
- ***USAGE** moves the data in the disk pool so that the frequently used data is evenly spread over all of the disk units in the disk pool.

To start balance on a disk pool using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Pools**.
3. Right-click the disk pool that you want to start balance on and select **Start Balance**.
4. Select the balance type and start balance.

To stop balance on a disk pool using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i
2. Select **Disk Pools**.
3. Right-click the disk pool that you want to stop balance on and select **Stop Balance**.
4. Follow the instructions on the dialog box that is displayed.

Enabling automatic overflow recovery for a basic disk pool

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 created sufficient space in the basic disk pool, the system copies the overflow data to the disk pool from the system disk pool. If automatic overflow recovery is disabled, you must manually restart your system to the dedicated service tools (DST) mode and recover the overflow data using DST from the command prompt.

To enable automatic overflow recovery using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Pools**.
3. Select the disk pools that you want to enable automatic overflow recovery for.
4. Right-click the disk pool that you want to enable overflow recovery for and select **Enable Overflow Recovery**.
5. Follow the instructions on the dialog box that is displayed.

Using disk pools with extensive journaling

This topic provides information on working with journal receivers when a disk pool overflows.

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. See [Moving journal receivers from an overflowed basic user auxiliary storage pool \(ASP\)](#) and [Resetting a journal with a status of overflowed](#) for information on how to recover a disk pool that is overfilled.

Related information

[Moving journal receivers from an overflowed basic user auxiliary storage pool \(ASP\)](#)
[Resetting a journal with a status of overflowed](#)

Managing independent disk pools

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

Make sure you can Access disk management functions.

Related concepts

[Requirements for disk management](#)

There are requirements that must be completed before you can configure or work with your disks.

Making a disk pool available

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

To access the disk units in an independent disk pool and the objects in the corresponding database, you must make the disk pool available (vary it on). If you are using geographic mirroring, you must make the

production copy of the disk pool available. You can only make the mirror copy available if it is detached. For a geographically mirrored disk pool, you must also make sure that the switchable hardware group is started before you attempt to make the disk pool available unless geographic mirroring is suspended.

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

If the name of a library or spool file within the independent disk pool is a duplicate of one in the system disk pool, the vary-on process is paused and an inquiry message is sent to the QSYSOPR message queue. This message identifies the duplicate library or spool file. You can then correct the problem and continue the vary-on process.

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

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 using the PowerHA® graphical interface, follow these steps:

1. In a web browser, enter **http://mysystem:2001**, where **mysystem** is the host name of the system.
2. Log on to the system using your user profile and password.
3. Click **PowerHA** in the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, click **Show All Others...** to show any independent disk pools that are not currently highly available.
6. Select **Vary On** from the context menu of the independent disk pool that you want to make available.
7. A progress panel is displayed that shows the current step in the process.

To make an independent disk pool available using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Pools**.
3. Right-click the unavailable disk pool and select **Make Available**.

You can select multiple disk pools to make available at the same time.

4. Click **Make Available** on the dialog box that is displayed to make the disk pool available.

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

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

Related information

[Vary Configuration \(VRYCFG\)](#)

Making a disk pool unavailable

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

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

To make an independent disk pool unavailable using the PowerHA graphical interface, follow these steps:

1. In a web browser, enter **http://mysystem:2001**, where **mysystem** is the host name of the system.
2. Log on to the system using your user profile and password.
3. Click **PowerHA** in the IBM Navigator for i window.
4. On the **PowerHA** page, click **Independent ASPs**.
5. On the **Independent ASPs** page, click **Show All Others...** to show any independent disk pools that are not currently highly available.
6. Select **Vary Off** from the context menu of the independent disk pool that you want to make unavailable.
7. A progress panel is displayed that shows the current step in the process.

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

To make an independent disk pool unavailable using your IBM Navigator for i, follow these steps::

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Pools**.
3. Right-click the disk pool that you want to make unavailable and select **Make Unavailable**.
4. Click **Make Unavailable** on the dialog box that is displayed to make the disk pool unavailable.

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

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

Use the Control ASP Access (QYASPCTLAA) API to restrict the processes that have access to the ASP.

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

Related tasks

[Releasing job reservations on an independent disk pool](#)

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

Related information

[Vary Configuration \(VRYCFG\)](#)

[Display ASP Status \(DSPASPSTS\)](#)

[Control ASP Access \(QYASPCTLAA\) API](#)

[Start DASD Management Operation \(QYASSDMO\) API](#)

Attach Independent Disk pool

An independent disk pool can be moved from one partition and attached to another partition that is in a single-system environment.

A partition in a single system environment is a stand-alone partition. The attach independent disk pool operation locates disk units from an independent disk pool that are connected to a stand-alone partition, and updates the partition to accept the independent disk pool into the configuration. The independent disk pool can then be varied on which allows programs to access the data in the independent disk pool

Restriction: The attach independent disk pool operation enforces the following restrictions:

- The partition must be in a single system environment. That means the partition cannot be a node in a device domain.
- The partition cannot have any independent disk pools that are configured when the attach independent disk pool operation is performed.

- Multiple independent disk pools cannot be attached to the partition unless it is a disk pool group. A disk pool group is a single primary disk pool and all of the secondary disk pools that are associated with the primary.

Copy an independent disk pool and attach it to a stand-alone partition

To use the data in an independent disk pool that is configured on a node in a cluster without corrupting or changing the data in the production disk pool, the disk pool can be copied and attached to a stand-alone partition.

1. Use the copy-services utility, or the external storage tools/functions, to copy all of the disk units in the independent disk pool.
2. On the stand-alone partition that receives the copy of the independent disk pool, run the Configure Device Auxiliary Storage Pool (**CFGDEVASP ASPDEV(*ALL) ACTION(*PREPARE)**) command to prepare the system configuration for the attachment of the disk units that contain the copy of the independent disk pool. This command deletes any independent disk pools that are currently configured.
3. IPL the system to complete the preparation.
4. Assign the FlashCopy disk units to the stand-alone node.
5. Run the attach independent disk pool operation from Dedicated Service Tools (DST) or System Service Tools (SST).
 - Select Work with disk units.
 - Select Work with disk unit recovery.
 - Select **Detect attached IASP disk units**. The Confirm attach IASP disk units screen is displayed. This screen shows the independent disk pool and the units in the disk pool.
 - If the disk pool and units that are displayed are correct, press enter to run the function.
6. Vary-on the independent disk pool.
7. Run jobs that use the data in the independent disk pool.

Backing up and recovering independent disk pools

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

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

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

Related information

[Recovering your system](#)

[Backup, Recovery and Media Services \(BRMS\)](#)

[Saving independent ASPs](#)

Recovering an independent disk pool

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

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

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

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

If the configuration source is corrupted, you can select to recover the configuration information about the configuration source. Recovering the configuration attempts to determine the original configuration and recover it. During this process, the independent disk pool might need to be cleared, destroying all data on the disk units in the pool. If the disk pool needs to be cleared, a message is displayed to inform you. You can cancel the recovery if you do not want the disk pool to be cleared.

- 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 that its mirrored disk unit was active before the failures that caused the state to become unknown.

To attempt to recover an independent disk pool using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Pools**.
3. Right-click the disk pool that is having problems. If IBM Navigator for i detects a problem with the configuration source, the **Recover Configuration** or **Recover Unknown Configuration Source** option is displayed. If either of these options are displayed, select it to continue.
4. Follow the instructions on the dialog box that is displayed.

Recovering a disk pool group

If the primary disk pool for a secondary disk pool is deleted, or if the primary disk pool is not aware of the secondary disk pool, the secondary disk pool needs to be reassociated with a primary disk pool. You can recover the disk pool group through IBM Navigator for i.

To recover a disk pool group using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Pools**.
3. Verify that the required primary disk pool exists. If it does not, you must create a disk pool.
4. Select all the secondary disk pools that need to be associated with a primary disk pool.
5. Right-click the disk pools and select **Recover Group**.
6. On the **Confirm Recover Disk Pool Group** dialog box that is displayed, 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.
7. Click **Recover Group**.

Releasing job reservations on an independent disk pool

If jobs are currently using an independent disk pool, you cannot make the disk pool unavailable (vary off). To make the independent disk pool unavailable, all the jobs using the disk pool must 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 using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Pools**.
3. Right-click the disk pool and select **Jobs**.

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

- 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.

- Let the job run to completion. If the job appears to be processing normally, you can wait until the job runs to completion.
- 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.
- 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.

Related information

[Ending a job](#)

[Work management troubleshooting](#)

[Work with ASP Jobs \(WRKASPJOB\)](#)

Independent disk pool examples

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

Related concepts

[Configuring independent disk pools](#)

After you have satisfied the planning requirements for implementing independent disk pools, you are ready to configure an independent disk pool. You must use the IBM Navigator for i disk management function or the Configure Device ASP (CFGDEVASP) command to configure an independent disk pool.

[Managing independent disk pools](#)

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

Examples: Independent disk pool configurations


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

Examples: Dedicated independent disk pools

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

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

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



Example: Making independent disk pool available at startup

You can decide whether to have your independent disk pool available when the system is restarted.

If you need your independent disk pools to be made available in most cases when your system 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 system, you can either Delete the Data Area (DLTDTAARA) 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 SBS(D(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:

```

Related information

[DLTDTAARA](#)

[RNMOBJ](#)

[CRTDTAARA](#)


Independent disk pools with distinct databases

Here gives you an example of 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 IBM Navigator for i. 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.

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.



If you expand Databases in IBM Navigator for i, 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.

Related concepts

Object identification

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

Disk pool checklists

Move, remove, and delete your disk pools with the help of the disk pool checklists.

Moving 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 system 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.	“Evaluating the current configuration” on page 5
2.____	Calculate the space requirements for both the source and destination disk pools for the disk units.	“Calculating disk space requirements” on page 5
3.____	Use option 21 from the Save menu to save your entire system.	Save your system with the GO SAVE command
4.____	Restart your system 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 System i Navigator, select Open System i Navigator Service Tools .
5.____	From the Disk Units Graphical View window, right-click the disk unit you want to move, and select Move.	“Removing a disk unit from a disk pool” on page 99
6.____	Verify that your disk configuration is correct.	“Evaluating the current configuration” on page 5
7.____	Print your disk configuration to have available in case a recovery situation occurs.	“Printing your disk configuration” on page 5
8.____	Restart your system.	

Removing disk units without disk 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 system 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.	“Evaluating the current configuration” on page 5
2.____	Calculate the space requirements for the disk pools that are involved in disk removal.	“Calculating disk space requirements” on page 5
3.____	Use option 21 from the Save menu to save your entire system.	Save your system with the GO SAVE command
4.____	Restart your system 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 System i Navigator, select Open System i Navigator Service Tools .
5.____	Remove disk units that you plan to remove from the system.	“Removing a disk unit from a disk pool” on page 99
6.____	Verify that your disk unit configuration is correct.	“Evaluating the current configuration” on page 5
7.____	Print your disk configuration to have available in case a recovery situation occurs.	“Printing your disk configuration” on page 5
8.____	Continue restarting your system.	

Deleting 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.	“Evaluating the current configuration” on page 5
2.____	Calculate the space requirements for the remaining disk pools.	“Calculating disk space requirements” on page 5
3.____	Use option 21 from the Save menu to save your entire system.	Save your system 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

Task	What to do	Where to learn more
5.____	Restart your system 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 System i Navigator, select Open System i Navigator Service Tools .
6.____	Delete the disk pool. This procedure places all of the disks that were assigned to the deleted disk pool in non-configured status.	"Deleting a disk pool" on page 99
7.____	If you want to add the now non-configured disk units to a different disk pool, refer to Checklist 2 or 3.	"Adding disk units without disk protection" on page 61 or "Adding disk units to an existing IOA" on page 53
8.____	Verify that your disk unit configuration is correct.	"Evaluating the current configuration" on page 5
9.____	Print your disk configuration to have available in case a recovery situation occurs.	"Printing your disk configuration" on page 5
10.____	Restart your system.	

Disk encryption

Disk encryption allows you to encrypt data that is stored in basic disk pools and independent disk pools.

Disk encryption protects data from a number of different threats:

- Protects data transmission to and from the disk drive (important in a SAN environment).
- Protects data transmission in the cross site mirroring environment (only when the data being mirrored is on an encrypted independent disk pool).
- Protects data in the case of theft of the disk drive.

To use disk encryption, you must have 5770-SS1 Option 45 - Encrypted ASP Enablement installed. The option to enable encryption is available when you create a disk pool or independent disk pool.

If disk encryption is used in a clustering environment, you must set the master key manually on each system within the device domain. Independent disk pools must be created using IBM Navigator for i. The character-based interface cannot be used.

Disk encryption can be used to encrypt existing disk pools or independent disk pools. Starting disk encryption on an existing disk pool might take an extended amount of time to encrypt the data in the disk pool, potentially affecting system performance.

To start disk encryption on a disk pool using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Pools**.
3. Right-click the disk pool that you want to start encryption on and select **Start Encryption**.
4. Follow the instructions on the dialog box that is displayed.

Disk encryption can be turned off on an encrypted disk pool. Turning off disk encryption takes an extended amount of time to decrypt all the data in the disk pool, potentially affecting system performance.

To stop encryption on a disk pool using your IBM Navigator for i, follow these steps:

1. Select **Configuration and Service** from your IBM Navigator for i.
2. Select **Disk Pools**.
3. Right-click the disk pool that you want to stop encryption on and select **Stop Encryption**.
4. Follow the instructions on the dialog box that is displayed.

Related information

[Loading and setting auxiliary storage pool \(ASP\) master key](#)

External load source disk unit

The load source disk unit is the disk unit that contains the initial programs and data that are used during an initial program load (IPL) of the system.

The load source disk unit is typically configured as an internal disk unit, but can also be configured as an external disk unit located on a storage area network (SAN).

For information on using an external load source disk unit that is located on a SAN, see the IBM Redbooks publication [IBM i and IBM Storwize® Family: A Practical Guide to Usage Scenarios, SG24-8197](#).

Related information

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

NVMe

NVMe (non-volatile memory express) is a host controller interface and storage protocol created to accelerate the transfer of data between enterprise and client systems and solid-state drives (SSD) over a computer's high-speed Peripheral Component Interconnect Express (PCIe) bus.

Related reference

[NVMe failure with disk units mirror protected](#)

When the NVMe card fails, the disk units (namespaces) that are mirror-protected will be suspended. Normal system operations continue on the other NVMe card.

[Installing a new partition using a NVMe as the load source disk unit](#)

To use a namespace on a NVMe device as a load source disk unit, follow this procedure.

[Removing an operational NVMe device from a partition](#)

To remove an operational NVMe device from an IBM i partition, follow this procedure.

[Upgrading the system ASP to NVMe - mirror protected](#)

To upgrade the system ASP which contains all SAS or SAN LUNs to use all NVMe namespaces, where the current system ASP has mirror protection, follow this procedure.

[Upgrading the system ASP to NVMe - not mirror protected](#)

To upgrade the system ASP which contains all SAS or SAN LUNs to use all NVMe namespaces, where the current system ASP is not mirror protected, follow this procedure.

[Upgrading a user ASP to NVMe](#)

To upgrade a user ASP which contains all SAS or SAN LUNs to use all NVMe namespaces, follow this procedure.

Related information

[What is flash storage?](#)

[IBM Storage FlashSystem](#)

NVMe concepts

NVMe (non-volatile memory express) is a host controller interface and storage protocol created to accelerate the transfer of data between enterprise and client systems and solid-state drives (SSD) over a computer's high-speed Peripheral Component Interconnect Express (PCIe) bus.

The NVMe device can be sub-divided up into multiple namespaces. A namespace is a quantity of non-volatile memory that may be formatted into logical blocks. IBM i sees each namespace as a disk unit.

All of the NVMe devices that IBM i supports are considered Self Encrypting Drives (SED). This means that the data is encrypted at rest. However, the key that is used to encrypt and decrypt the data on the NVMe is not protected; thus any actor (once in possession of the device) can access all the data.

Some NVMe devices support the [TCG Storage Security Subsystem Class \(SSC\): Opal](#) design architecture.

This design architecture uses a client-supplied password to protect the data on an NVMe. If the NVMe is moved to another logical partition or loses power, the data on that NVMe becomes read/write protected.

There exists a rare scenario where the NVMe might not be locked even when the NVMe is moved to a different logical partition:

1. The NVMe belongs to the IBM i logical partition NVMe password protection policy.
2. Then the IBM i logical partition terminates abnormally, or the Hardware Management Console (HMC) is used to power off the IBM i logical partition.
3. Then Dynamic Logical Partitioning (DLPAR) is used to move the NVMe to a different logical partition, all while the NVMe remains powered on.

The requirements to enable an NVMe password protection policy on an IBM i logical partition are:

1. Platform Key Storage (PKS) is enabled under the HMC.
2. Logical partition is running on a Power-10 processor.
3. Logical partition firmware is at least FW1030 or higher.
4. The NVMe supports the design architecture.
5. Running the IBM i 7.4 release or higher with a specific technology refresh.

There can be only a single NVMe password protection policy on each IBM i logical partition.

Note: The NVMe password protection policy is not related to disk encryption.

The NVMe devices are already Self Encrypting Devices (SED), so there is no performance penalty on I/O operations to NVMe under the NVMe password protection policy.

Disk encryption is a mechanism that encrypts data that is stored in an auxiliary storage pool (ASP). Basic disk pools and independent disk pools can be encrypted. To use disk encryption, the logical partition must have 5770-SS1 Option 45 Encrypted ASP Enablement installed.

The disk units in an encrypted ASP can be of any type: SSD, External LUNs, Virtual LUNs, NVMe namespaces, etc.

Disk encryption is controlled by the operating system. Disk encryption passwords are not stored in PKS; a different key storage mechanism is used.

The I/O operations to a disk unit protected by disk encryption are summarized below:

- Write: The operating system encrypts the data before the data is transferred to the disk unit to be written to the disk.
- Read: After data is read from a disk unit in an encrypted ASP, the operating system decrypts the data.

There is a significant performance penalty associated with disk encryption; every disk I/O operation requires the data to be encrypted or decrypted by the operating system.

Managing NVMe

This topic describes how to manage your NVMe devices and namespaces.

Display NVM namespaces

Select this function to show all NVMe devices and namespaces that exist on the IBM i partition.

For each NVMe device, the type, serial number, and resource name are shown.

For each namespace that exists on a NVMe device, the type, model, serial number, resource name, ASP number, unit number, and capacity are shown. Namespaces that are non-configured disk units show '*' in the ASP and Unit fields.

1. Select **Work with disk units** on the Use System Service Tools (SST) display.
2. Select **Work with disk configuration** on the Work with Disk Units display.
3. Select **Work with NVM Devices** on the Work with Disk Configuration display.

4. Select **Display NVM namespaces** on the Work with NVM Devices display.

Display NVM devices

Select this function to show all NVMe devices that exist on the IBM i partition.

For each NVMe device, the type, serial number, resource name, the used/available/total capacity, and used/available/total namespaces are shown.

1. Select **Work with disk units** on the Use System Service Tools (SST) display.
2. Select **Work with disk configuration** on the Work with Disk Units display.
3. Select **Work with NVM Devices** on the Work with Disk Configuration display.
4. Select **Display NVM devices** on the Work with NVM Devices display.

Create NVM namespaces

Select this function to allocate storage and create namespaces on a NVMe device.

A NVMe namespace is a quantity of memory that is formatted into logical blocks. A NVMe namespace becomes a disk unit on the IBM i partition.

1. Select **Work with disk units** on the Use System Service Tools (SST) display.
2. Select **Work with disk configuration** on the Work with Disk Units display.
3. Select **Work with NVM Devices** on the Work with Disk Configuration display.
4. Select **Create NVM name spaces** on the Work with NVM Devices display.
5. Select the NVMe device.
6. Select the quantity and capacity of the namespaces you want to create on the NVMe device.
7. Press F10 to confirm your selection.

Create NVM namespaces to match another NVM device

Select this function to create namespaces on a NVMe device that match the capacity of namespaces on another NVMe device.

1. Select **Work with disk units** on the Use System Service Tools (SST) display.
2. Select **Work with disk configuration** on the Work with Disk Units display.
3. Select **Work with NVM Devices** on the Work with Disk Configuration display.
4. Select **Create NVM namespaces to match another NVM device** on the Work with NVM Devices display.
5. Select the source NVMe device that contains the namespaces that you want to match.
6. On the Display Namespaces on the Selected NVM Device display, press Enter to continue.
7. Select the target NVMe device.
8. Press F10 to confirm your selection and create the namespaces.

Create NVM namespaces to pair with Active mirror protected NVM units

Select this function to create namespaces on a NVMe device to pair with active mirror protected NVMe units.

If a NVMe device fails and the NVMe device contains namespaces that are mirror-protected disk units in an auxiliary storage pool, all of the IBM i disk units on the failed NVMe device transition to a 'suspended' mirror state.

Before you use this function, you or your service provider need to physically remove the failing NVMe device and physically install a new NVMe device. See [“NVMe failure with disk units mirror protected” on page 77](#) for more information.

This function performs the following steps.

1. Examine the auxiliary storage pools to locate the NVMe-type mirror-protected disk units that are suspended from mirror protection.
2. Create new namespaces on the newly installed selected NVMe device.
3. Pair the new namespaces with the suspended NVMe disk units.
4. Start to resume mirror protection on the newly created disk units.

Related reference

NVMe failure with disk units mirror protected

When the NVMe card fails, the disk units (namespaces) that are mirror-protected will be suspended. Normal system operations continue on the other NVMe card.

Delete existing NVM namespaces

Select this function to destroy the IBM i disk unit associated with the NVMe namespace and return the namespace storage to the available pool in the NVMe device.

Note:

This option is available only on namespaces with disk units that are non-configured or suspended from mirrored protection.

1. Select **Work with disk units** on the Use System Service Tools (SST) display.
2. Select **Work with disk configuration** on the Work with Disk Units display.
3. Select **Work with NVM Devices** on the Work with Disk Configuration display.
4. Select **Delete existing NVM Namespaces** on the Work with NVM Devices display.
5. Select the NVMe device.
6. Select the namespaces on the NVMe device that you want to delete.
7. Press F10 to confirm your selection.

Sanitize/Erase NVM device

Select this function to remove and destroy all data from the NVMe media and the cache.

There are five types of sanitize:

User data secure erase

All user data is erased. The content of the user data after the erase is indeterminate (For example, the user data might be zero filled, one filled, etc.). The controller might perform a cryptographic erase when a User Data Erase is requested if all user data is encrypted.

Format cryptographic erase

All user data is erased cryptographically. This is accomplished by deleting the encryption key.

Sanitize overwrite

User data is altered by writing a fixed data pattern or related patterns to all locations on the media within the NVM subsystem in which user data may be stored one or more times.

Sanitize block erase

User data is altered with a low-level block erase method that is specific to the media for all locations on the media within the NVM subsystem in which user data may be stored.

Sanitize cryptographic erase

User data is altered by changing the media encryption keys for all locations on the media within the NVM subsystem in which user data may be stored.

Notes:

- Each type of NVMe device might implement only a subset of the Sanitize functions described above.
- The sanitize function destroys all data in the namespaces which exist on the selected NVMe device.
- The sanitize function is available only for NVMe devices that only contain namespaces which represent non-configured disk units (disk units which do not belong to an auxiliary storage pool).

1. Select **Work with disk units** on the Use System Service Tools (SST) display.
2. Select **Work with disk configuration** on the Work with Disk Units display.
3. Select **Work with NVM Devices** on the Work with Disk Configuration display.
4. Select **Sanitize/Erase NVM device** on the Work with NVM Devices display.
5. Select the NVMe device.
6. Select the type of sanitize function that you want to perform.
7. Press F10 to confirm your selection.

Format NVM device to prepare device for IBM i

Select this function to change the block size of the NVMe device to a format that can be used on the IBM i partition.

A disk unit that is type 6600 has a block size that the IBM i cannot use. This function should only be performed on disk units with a type of 6600. This function will change the disk unit type to be the correct value for that device.

Note: This function destroys all data in the namespaces that exist on the selected NVMe device. If there are no namespaces on the NVMe device, the Format command will not be issued.

1. Select **Work with disk units** on the Use System Service Tools (SST) display.
2. Select **Work with disk configuration** on the Work with Disk Units display.
3. Select **Work with NVM Devices** on the Work with Disk Configuration display.
4. Select **Format NVM device to prepare device for IBM i** on the Work with NVM Devices display.
5. Select the NVMe device.
6. Press F10 to confirm your selection.

NVMe password protection

Select this function to use the features of the NVMe password protection policy.

Prior to using any of the following procedures listed below, the Platform KeyStore must be enabled through the **HMC's Partition Properties -> Advanced Settings -> Platform KeyStore Size** value must be set to a non-zero value.

Display the NVMe password protection policy and NVMe device attributes

This section describes how to display the NVMe password protection policy and the NVMe device attributes.

To display the NVMe password protection policy and NVMe device attributes:

1. On the Use System Service Tools (SST) display, select **Start a service tool**
2. On the Start a Service Tool display, select **Display/Alter/Dump**
3. On the Display/Alter/Dump Output Device display, select **Display/Alter Storage**
4. On the Select Data display, select **Licensed Internal Code (LIC) data**
5. On the Select LIC Data display, select **Advanced Analysis** (option 14)
6. On the Select Advanced Analysis Command display:
 - a. Type 1 in the Option field
 - b. Type NVMEDISPLAYLP in the Command field
 - c. Press the Enter key
7. On the Specify Advanced Analysis Option display, Enter key

A summary display shows the NVMe password protection policy information and the attributes of each NVMe on the IBM i logical partition.

An example of that display is shown here:

Primary policy exists
No secondary policy. RC = 1

NVMe Device Attributes

Serial Number	Resource Name	Type	Support Locking	Locked	Under Policy	Eligible To Add	Exceed Limit	FixUp Req
Y0YA87YEL049	DC02	5947	Yes	No	No	Yes	No	N/A
Y0YAF4YFM028	DC05	5B50	Yes	No	No	Yes	No	N/A
Y0YAF4YFM01U	DC06	5B50	Yes	No	Yes	No	No	No
Y11STSJWRSYJ	DC07	5B51	Yes	No	Yes	No	No	No

The description of each column is:

1. **Serial Number:** The number assigned by the device manufacturer to identify a specific NVMe.
2. **Resource Name:** The IBM i assigned name of the NVMe.
3. **Type:** The type of NVMe.
4. **Support Locking:** 'Yes' if the NVMe supports the architecture that is used to password protect the NVMe, 'No' otherwise.
5. **Locked:** Shows if the NVMe is currently locked. 'Yes' if the NVMe is currently locked, 'No' if it is not locked. If an NVMe is locked, all of the namespaces within that device are read/write protected.
6. **Under Policy:** 'Yes' if the NVMe belongs to the NVMe password protection policy for the IBM i logical partition, 'No' otherwise.
7. **Eligible to Add:** 'Yes' if the NVMe is eligible to be added to the NVMe password protection policy for the IBM i logical partition, 'No' otherwise.
8. **Exceed Limit:** Describes an error state in the NVMe, where multiple attempts have been made to access the NVMe with an incorrect password. The NVMe must be power cycled to get the NVMe out of this error state.
9. **FixUp Req:** Describes an error state in the NVMe, where the internal stored passwords are inconsistent.

Create the NVMe password protection policy

This section describes how to create the NVMe password protection policy on the IBM i logical partition.

Before creating the NVMe password protection policy on the logical partition, it is recommended that all of the system and user data be saved via the [GO SAVE: Option 21](#) (saving the entire system) procedure.

Follow these steps to create the NVMe password protection policy:

1. On the Use System Service Tools (SST) display, select **Start a service tool**
2. On the Start a Service Tool display, select **Display/Alter/Dump**
3. On the Display/Alter/Dump Output Device display, select **Display/Alter Storage**
4. On the Select Data display, select **Licensed Internal Code (LIC) data**
5. On the Select LIC Data display, select **Advanced Analysis** (option 14)
6. On the Select Advanced Analysis Command display:
 - a. Type 1 in the Option field
 - b. Type NVMECREATELP in the Command field
 - c. Press the Enter key
7. On the Options field of the Specify Advanced Analysis Option display, specify the password, and confirm the password value
 - a. Type -P "<password>" on the first Options line and type -C "<password>" on the second Options line

- i) Each instance of the password must be enclosed in double quotation marks (“ ”).
- ii) The password must be at least 8 and not more than 32 characters in length
- iii) Sample usage: NVMECREATELP
 - -P “Policy Password”
 - -C “PolicyPassword”

b. Press the Enter key

After the above steps have been completed the NVMe password protection policy is started and the user-supplied password is stored in PKS on the IBM i logical partition.

Each NVMe which contains a namespace that represents a configured disk unit is automatically added to the NVMe password protection policy; the password is stored in the NVMe, and the NVMe is then password protected.

Change the password on the NVMe password protection policy

This topic describes how to change the NVMe password protection policy on the IBM i logical partition.

Before changing the password on the NVMe password protection policy, it is recommended that all of the system and user data be saved via the GO SAVE: Option 21 (saving the entire system) procedure.

Follow these steps to change the password:

1. On the Use System Service Tools (SST) display, select **Start a service tool**
2. On the Start a Service Tool display, select **Display/Alter/Dump**
3. On the Display/Alter/Dump Output Device display, select **Display/Alter Storage**
4. On the Select Data display, select **Licensed Internal Code (LIC) data**
5. On the Select LIC Data display, select **Advanced Analysis** (option 14)
6. On the Select Advanced Analysis Command display:
 - a. Type 1 in the Option field
 - b. Type NVMECHANGE PWLP in the Command field
 - c. Press the Enter key
7. On the Specify Advanced Analysis Option display, specify the current password on the first Options line, and put the new password, and a confirmation of the new password on the second Options line:
 - a. Type -P “<currentpassword>” -N “<newpassword>” -C “<newpassword>”
 - i) Each instance of the password must be enclosed in double quotation marks (“ ”).
 - ii) The new password must be at least 8 and not more than 32 characters in length.
 - iii) Sample usage: NVMECHANGE PWLP
 - -P “Policy Password”
 - -N “NewPolicyPassword” -C “NewPolicyPassword”
 - b. Press the Enter key

After the above steps have been completed the new password is stored in the PKS of the IBM i logical partition and each NVMe that belongs to the NVMe password protection policy is updated with the new password.

Delete the NVMe password protection policy

This section describes how to delete the NVMe password protection policy on the IBM i logical partition.

Follow these steps to delete the NVMe password protection policy:

1. On the Use System Service Tools (SST) display, select **Start a service tool**
2. On the Start a Service Tool display, select **Display/Alter/Dump**
3. On the Display/Alter/Dump Output Device display, select **Display/Alter Storage**

4. On the Select Data display, select **Licensed Internal Code (LIC) data**
5. On the Select LIC Data display, select **Advanced Analysis** (option 14)
6. On the Select Advanced Analysis Command display:
 - a. Type 1 in the Option field
 - b. Type NVMEDELETELP in the Command field
 - c. Press the Enter key
7. On the Options field of the Specify Advanced Analysis Option display
 - a. Type -P "<password>"
 - i) The password must be enclosed in double quotation marks (" ").
 - ii) The password is the NVMe password protection policy password
 - iii) Sample usage: NVMEDELETELP -P "PolicyPassword"
 - b. Press the Enter key

After the above steps have been completed each NVMe that belongs to the NVMe password protection policy is updated to remove itself from the policy and PKS is updated to remove the NVMe password protection policy on the IBM i logical partition.

Add an NVMe to the existing NVMe password protection policy

This section describes how to add an unprotected NVMe to the NVMe password protection policy on the IBM i logical partition.

Follow these steps to add an NVMe to the existing NVMe password protection policy:

1. On the Use System Service Tools (SST) display, select **Start a service tool**
2. On the Start a Service Tool display, select **Display/Alter/Dump**
3. On the Display/Alter/Dump Output Device display, select **Display/Alter Storage**
4. On the Select Data display, select **Licensed Internal Code (LIC) data**
5. On the Select LIC Data display, select **Advanced Analysis** (option 14)
6. On the Select Advanced Analysis Command display:
 - a. Type 1 in the Option field
 - b. Type NVMEADDDEVLP in the Command field
 - c. Press the Enter key
7. On the Options field of the Specify Advanced Analysis Option display
 - a. Type -P "<password>" -D DCXX
 - i) The password must be enclosed in double quotation marks (" ").
 - ii) DCXX is the resource name of the NVMe to be added.
 - iii) The password is the NVMe password protection policy password.
 - iv) Sample usage: NVMEADDDEVLP -P "PolicyPassword" -D DC05
 - b. Press the Enter key

After the above steps have been completed the password of the NVMe password protection policy is stored in the specified NVMe. That NVMe is now password protected.

Remove an NVMe from the existing NVMe password protection policy

This section describes how to remove a protected NVMe from the NVMe password protection policy on the IBM i logical partition.

There are two scenarios where this function is used:

1. The NVMe is to be removed from the current NVMe password protection policy.

2. The NVMe is in a locked state because the password in the NVMe does not match the password stored in the PKS of the IBM i logical partition. As a result, all of the namespaces under the NVMe are read/write protected.

Follow these steps to remove an NVMe from the existing NVMe password protection policy:

1. On the Use System Service Tools (SST) display, select **Start a service tool**
2. On the Start a Service Tool display, select **Display/Alter/Dump**
3. On the Display/Alter/Dump Output Device display, select **Display/Alter Storage**
4. On the Select Data display, select **Licensed Internal Code (LIC) data**
5. On the Select LIC Data display, select **Advanced Analysis** (option 14)
6. On the Select Advanced Analysis Command display:
 - a. Type 1 in the Option field
 - b. Type NVMEREMOVEDEVLP in the Command field
 - c. Press the Enter key
7. On the Options field of the Specify Advanced Analysis Option display
 - a. Type -P "<password>" -D DCXX
 - i) The password must be enclosed in double quotation marks (" ").
 - ii) DCXX is the resource name of the NVMe to be removed from the NVMe password protection policy.
 - If NVMe is in a locked status, the password that is typed in the command is the NVMe password protection policy password on the IBM i logical partition where the NVMe came from.
 - If the NVMe is not in a locked status, the password that is typed in the command is the NVMe password protection policy password.
 - iii) Sample usage: NVMEREMOVEDEVLP -P "PolicyPassword" -D DC05
 - b. Press the Enter key

After the above steps have been completed the password for the NVMe password protection policy is removed from the specified NVMe. That NVMe is no longer password protected.

Unlock an NVMe

This section describes how to unlock an NVMe that is locked (read/write protected) on an IBM i logical partition.

The NVMe would be in a locked state if it belonged to an NVMe password protection policy on one IBM i logical partition and then the NVMe was moved to a different IBM i logical partition. Thus, the password in the NVMe does not match the password stored in the PKS of the logical partition where the NVMe currently resides; the NVMe is locked and all of the namespaces under that NVMe are read/write protected.

If the password of the NVMe password protection policy on the IBM i logical partition where the NVMe came from is known, the NVMe can be unlocked and the data can be made accessible.

Follow these steps to unlock an NVMe:

1. On the Use System Service Tools (SST) display, select **Start a service tool**
2. On the Start a Service Tool display, select **Display/Alter/Dump**
3. On the Display/Alter/Dump Output Device display, select **Display/Alter Storage**
4. On the Select Data display, select **Licensed Internal Code (LIC) data**
5. On the Select LIC Data display, select **Advanced Analysis** (option 14)
6. On the Select Advanced Analysis Command display:
 - a. Type 1 in the Option field

- b. Type NVMEUNLOCKDEVLP in the Command field
 - c. Press the Enter key
7. On the Options field of the Specify Advanced Analysis Option display, you specify the password
- a. Type -P "<password>" -D DCXX
 - i) The password must be enclosed in double quotation marks (" ").
 - ii) DCXX is the resource name of the NVMe to be unlocked.
 - iii) The password that is typed in the command is the password of the NVMe password protection policy on the IBM i logical partition where the NVMe came from.
 - iv) Sample usage: NVMEUNLOCKDEVLP -P "Password" -D DC05
 - b. Press the Enter key

After the above steps have been completed the specified NVMe is unlocked and is no longer read/write protected; however, the NVMe will become locked again if the device is power cycled.

Factory Reset an NVMe

This section describes how to perform a factory reset operation on a locked NVMe.

The NVMe is in a locked state if it was moved to a different IBM i logical partition. The password stored in the NVMe would not match the password stored in the PKS of the logical partition where the NVMe currently resides; the NVMe is locked and all of the namespaces under that NVMe are read/write protected.

If the password to unlock the NVMe is not known, the only way to get the NVMe out of the locked state is to perform a factory reset.

Note: The Factory Reset will destroy all of the data in all of the namespaces on the NVMe.

Follow these steps to factory reset an NVMe:

1. On the Use System Service Tools (SST) display, select **Start a service tool**
2. On the Start a Service Tool display, select **Display/Alter/Dump**
3. On the Display/Alter/Dump Output Device display, select **Display/Alter Storage**
4. On the Select Data display, select **Licensed Internal Code (LIC) data**
5. On the Select LIC Data display, select **Advanced Analysis** (option 14)
6. On the Select Advanced Analysis Command display:
 - a. Type 1 in the Option field
 - b. Type NVMEFACTORYRESETDEVLP in the Command field
 - c. Press the Enter key
7. On the Options field of the Specify Advanced Analysis Option display
 - a. Type -PSID "<externalLabel>" -D DCXX
 - i) The PSID (Physical Security Identifier) is printed on the external label of the NVMe.
 - ii) The PSID must be enclosed in double quotation marks (" ").
 - iii) DCXX is the resource name of the NVMe to be factory reset.
 - iv) Sample usage: NVMEFACTORYRESETDEVLP -P "PSID_VALUE_FROM_EXTERNAL_LABEL" -D DC05
 - b. Press the Enter key

After the above steps have been completed all of the data in all of the namespaces in the specified NVMe is destroyed. The specified NVMe is unlocked, and the password is removed from the NVMe.

Disk management checklist

This topic contains checklists for performing configuration procedures. Use this information to determine which checklist to use for your situation.

Configuring disks on a new system

This checklist shows the sequence of tasks that you use to configure disks on a new system. Whether you need to perform all the tasks depends on the disk protection that you want on your system.

[“Disk protection” on page 9](#) 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 disk unit appear as non-configured.	“Evaluating the current configuration” on page 5
2.____	Use the Add Disk Unit wizard to add non-configured 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.	“Adding a disk unit or disk pool” on page 95
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%.	“Setting the threshold of a disk pool” on page 101
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.	“Starting mirrored protection” on page 43
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.	“Evaluating the current configuration” on page 5
7.____	Print your disk configuration to have available in case a recovery situation occurs.	“Printing your disk configuration” on page 5

Adding disk units without disk protection

This checklist shows the sequence of tasks that you use to configure disks on a new system. Whether you need to perform all the tasks depends on the disk protection that you want on your system.

[“Disk protection” on page 9](#) 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 disk 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.	“Evaluating the current configuration” on page 5
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 non-configured 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.	“Adding a disk unit or disk pool” on page 95
4.____	Change the storage threshold for any disk pool if required. The default storage threshold for each disk pool is 90%.	“Setting the threshold of a disk pool” on page 101
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.	“Starting mirrored protection” on page 43
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.	“Evaluating the current configuration” on page 5
8.____	Print your disk configuration to have available in case a recovery situation occurs.	“Printing your disk configuration” on page 5

Adding disk units to an existing IOA

This checklist shows the sequence of tasks that you use to add one or more disks to an existing IOA 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 [“Adding disk units without disk protection” on page 61](#).

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.	“Evaluating the current configuration” on page 5
2.____	Use the Add Disk Unit wizard to add non-configured 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.	“Adding a disk unit or disk pool” on page 95
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.	“Setting the threshold of a disk pool” on page 101
4.____	Verify that your disk configuration is correct.	“Evaluating the current configuration” on page 5
5.____	Print your disk configuration to have available in case a recovery situation occurs.	“Printing your disk configuration” on page 5

Adding a new IOA

This checklist shows the sequence of tasks that you use to add a new IOA and new disk units to your system.

Before you begin

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 [“Adding disk units without disk protection” on page 61](#) 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.

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 IOA in the system. This is normally done by a service representative.	“Evaluating the current configuration” on page 5
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.	“Evaluating the current configuration” on page 5
3.____	Use the Add Disk Unit wizard to add non-configured disks to the correct disk pools. Use the option to start device parity protection.	“Adding a disk unit or disk pool” on page 95
4.____	If you want a different storage threshold for any disk pool, change it. The default storage threshold for each disk pool is 90%.	“Setting the threshold of a disk pool” on page 101
5.____	You may want to restart to the dedicated service tools (DST) mode in order for the device parity protection to take effect.	“Device parity protection” on page 10
6.____	Wait until your system completely restarts.	
7.____	Verify that your disk configuration is correct.	“Evaluating the current configuration” on page 5
8.____	Print your disk configuration to have available in case a recovery situation occurs.	“Printing your disk configuration” on page 5

Moving 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 system 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.	“Evaluating the current configuration” on page 5
2.____	Calculate the space requirements for both the source and destination disk pools for the disk units.	“Calculating disk space requirements” on page 5
3.____	Use option 21 from the Save menu to save your entire system.	Save your system with the GO SAVE command
4.____	Restart your system 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 System i Navigator, select Open System i Navigator Service Tools .
5.____	From the Disk Units Graphical View window, right-click the disk unit you want to move, and select Move.	“Removing a disk unit from a disk pool” on page 99
6.____	Verify that your disk configuration is correct.	“Evaluating the current configuration” on page 5
7.____	Print your disk configuration to have available in case a recovery situation occurs.	“Printing your disk configuration” on page 5
8.____	Restart your system.	

Moving 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 system 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.	“Evaluating the current configuration” on page 5
2.____	Calculate the space requirements for both the source and destination disk pools for the disk units.	“Calculating disk space requirements” on page 5
3.____	Use option 21 from the Save menu to save your entire system.	Save your system with the GO SAVE command

Task	What to do	Where to learn more
4.____	Restart your system 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 System i Navigator, select Open System i 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 .	"Removing a disk unit from a disk pool" on page 99
6.____	Add non-configured 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.	"Adding a disk unit or disk pool" on page 95
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.	"Setting the threshold of a disk pool" on page 101
8.____	If you created any new disk pools and you want those disk pools to have mirrored protection, start mirrored protection now.	"Starting mirrored protection" on page 43
9.____	Verify that your disk unit configuration is correct.	"Evaluating the current configuration" on page 5
10._ __	Print your disk configuration to have available in case a recovery situation occurs.	"Printing your disk configuration" on page 5

Deleting 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.	"Evaluating the current configuration" on page 5

Task	What to do	Where to learn more
2.____	Calculate the space requirements for the remaining disk pools.	“Calculating disk space requirements” on page 5
3.____	Use option 21 from the Save menu to save your entire system.	Save your system 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
5.____	Restart your system 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 System i Navigator, select Open System i Navigator Service Tools .
6.____	Delete the disk pool. This procedure places all of the disks that were assigned to the deleted disk pool in non-configured status.	“Deleting a disk pool” on page 99
7.____	If you want to add the now non-configured disk units to a different disk pool, refer to Checklist 2 or 3.	“Adding disk units without disk protection” on page 61 or “Adding disk units to an existing IOA” on page 53
8.____	Verify that your disk unit configuration is correct.	“Evaluating the current configuration” on page 5
9.____	Print your disk configuration to have available in case a recovery situation occurs.	“Printing your disk configuration” on page 5
10.____	Restart your system.	

Removing disk units without disk 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 system 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.	“Evaluating the current configuration” on page 5
2.____	Calculate the space requirements for the disk pools that are involved in disk removal.	“Calculating disk space requirements” on page 5
3.____	Use option 21 from the Save menu to save your entire system.	Save your system with the GO SAVE command

Task	What to do	Where to learn more
4.____	Restart your system 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 System i Navigator, select Open System i Navigator Service Tools .
5.____	Remove disk units that you plan to remove from the system.	"Removing a disk unit from a disk pool" on page 99
6.____	Verify that your disk unit configuration is correct.	"Evaluating the current configuration" on page 5
7.____	Print your disk configuration to have available in case a recovery situation occurs.	"Printing your disk configuration" on page 5
8.____	Continue restarting your system.	

Removing 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 IOA 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 system 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.	"Evaluating the current configuration" on page 5
2.____	Calculate the space requirements for the disk pools that are involved in disk removal.	"Calculating disk space requirements" on page 5
3.____	Use option 21 from the Save menu to save your entire system.	Save your system with the GO SAVE command
4.____	Restart your system 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 System i Navigator, select OpenSystem i Navigator Service Tools .
5.____	Remove disk units that you plan to remove from the system.	"Removing a disk unit from a disk pool" on page 99

Task	What to do	Where to learn more
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.	“Excluding disk units from a parity set” on page 23
7.____	Stop device parity protection for all the disk units in the Input/Output Processor (IOP).	“Stopping device parity protection” on page 22
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.	“Removing a disk unit from a disk pool” on page 99
9.____	Start device parity protection again.	“Starting device parity protection” on page 21
10.____	Verify that your disk unit configuration is correct.	“Evaluating the current configuration” on page 5
11.____	Print your disk configuration to have available in case a recovery situation occurs.	“Printing your disk configuration” on page 5
12.____	Restart your system.	

Removing 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 IOA 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 system 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.	“Evaluating the current configuration” on page 5
2.____	Calculate the space requirements for the disk pools that are involved in disk removal.	“Calculating disk space requirements” on page 5
3.____	Use option 21 from the Save menu to save your entire system.	Save your system with the GO SAVE command
4.____	Restart your system 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 System i Navigator, select Open System i Navigator Service Tools .

Task	What to do	Where to learn more
5.____	Remove disk units that you plan to remove from the system.	“Removing a disk unit from a disk pool” on page 99
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.	“Excluding disk units from a parity set” on page 23
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 Input/Output Processor (IOP) and have device parity protection.	“Stopping mirrored protection” on page 44
8.____	Stop device parity protection for all the disk units in the IOP.	“Stopping device parity protection” on page 22
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.	“Starting device parity protection” on page 21
11.____	Add unconfigured disk units to the correct disk pools. These disk units became unconfigured when you stopped mirrored protection in task 7.	“Adding a disk unit or disk pool” on page 95
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.	“Setting the threshold of a disk pool” on page 101
13.____	Start mirrored protection for the disk pools that had mirrored protection stopped in task 7.	“Starting mirrored protection” on page 43
14.____	Verify that your disk unit configuration is correct.	“Evaluating the current configuration” on page 5
15.____	Print your disk configuration to have available in case a recovery situation occurs.	“Printing your disk configuration” on page 5

Upgrading the 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 70-GB capacity while keeping device parity active. The disk 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 configuration by following the steps in “Evaluating the current configuration” on page 5, and record your answers. The information that is entered in the table that is labeled Disk configuration questions, is needed to answer the questions in the Load source planning sections.

Disk configuration questions	Disk configuration answers
1. 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.	
2. 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 parity set or optionally mark the disk units in the parity set. Make sure that you can identify the load source disk unit separately from the other disk units.	
3. 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.

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 disk units?	Yes / No
Do you know how to physically install and remove disk units from your system? <u>Install</u> and <u>remove</u> replacement disk units on an IBM i.	Yes / No
Do you have IBM Navigator for i 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 system using the <u>GO SAVE</u> command.

Task	What to do
2.____	<p>Turn off your system or logical partition by entering the following command.</p> <pre data-bbox="347 233 716 268">PWRDWN SYS *IMMED RESTART(*NO)</pre> <p>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.</p>
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. 2. Install replacement disk units on your system. <p>Notes:</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 70-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 disk 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.

Task	What to do																																									
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 disk 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 12.</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 disk units on the screen and press Enter. 6. Press enter in the Confirm Add Units display. 7. The status display shows how the operation is proceeding. When the function is complete, return to the Use Dedicated Tools (DST) menu. 																																									
7.____	<p>Locate the load source disk unit (unit 1) and all other disk units in the parity set that you will be upgrading.</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 Display disk configuration. 4. On the Display Disk Configuration display Select Display device parity status. <p>Note: If the status on the disk units in the parity set does not say Active, do not continue with this procedure. Contact the next level of support.</p> <ol style="list-style-type: none"> 5. Find the parity set that contains the load source disk unit (unit 1). 6. Record the disk units that are in its parity set. Your load source disk unit (unit 1) is already recorded for you. Recording this information is helpful for the next step. <p>Note: There could be up to 18 disk units.</p> <table border="1" data-bbox="370 1203 1128 1283"> <tr> <td>__1__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td> </tr> <tr> <td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td><td>__</td> </tr> </table> <ol style="list-style-type: none"> 7. Return to the Use Dedicated Service Tools (DST) menu. 	__1__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__
__1__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__																							
__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__	__																						
8.____	<p>Remove the disk units from the configuration. You recorded the units in Task 7, Step 6.</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 Remove disk units from configuration. 5. Type a 4 next to each disk unit that you want to remove from the configuration. Select the disk units that were recorded in Task 7, Step 6 and press Enter. <p>Note: You will not be able to select load source disk unit (unit 1). You will remove the load source in Task 14.</p> <ol style="list-style-type: none"> 6. Press Enter in the Confirm Remove disk units display. 7. The status display shows how the operation is proceeding. When the function is complete, return to the Use Dedicated Service Tools (DST) menu. 																																									

Task	What to do
9.____	<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 and press Enter. <p>Note: There is no command line because you are using Dedicated Service Tools.</p>
10.____	<p>Use the HMC to change the I/O tag.</p> <ol style="list-style-type: none"> 1. In your HMC, expand Systems Management > Servers. 2. Select the system. 3. Select the check box for the partition. 4. From the menu for the partition, click Configuration > Manage Profiles. 5. In the Managed Profiles pane, click the profile name. 6. In the Logical Partition Profile Properties pane, click the Tagged I/O tab. 7. From the Load source list, select the IOA that you want to control the Load Source disk unit. 8. Click OK. <p>Note: The changes that you make only go into effect when the partition is shut down and then activated again.</p>
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).
12.____	<p>Copy the load source disk unit (unit 1) to one of the replacement disk units that you installed in Task 3.</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. Write down the serial number of the replacement disk unit that is going to replace the load source. _____ 7. Press Enter in the Confirm Copy Disk Unit Data display. 8. The status display shows how the operation is proceeding. When the copy function is complete, the partition will automatically IPL to Dedicated Service Tools (DST).
13.____	<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 and press Enter. <p>Note: There is no command line because you are using Dedicated Service Tools.</p>

Task	What to do
14.____	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 IBM Navigator for i or from the disk units that you marked.
15.____	Work with Dedicated Service Tools (DST). 1. Power on your system or logical partition to Dedicated Service Tools (DST).
16.____	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 12, Step 6 .
17.____	IPL your system to IBM i.

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.

Upgrading the load source disk unit with local mirroring using non-configured disk units

This checklist shows the sequence of tasks to upgrade your load source disk unit (unit 1) and the mirrored load source with disk units that have equal capacity, are at least 70-GB, and are mirror protected. The disk 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 11. Disk configuration questions

Disk configuration questions	Disk configuration answers
<p>Where is the load source disk unit and the mirrored load source disk unit located? Determine the disk unit locations using Dedicated Service Tools (DST) or System Service Tools (SST).</p> <ol style="list-style-type: none"> From either Use Dedicated Service Tools (DST) or System Service Tools (SST) display, select Work with disk units. From the Work with Disk Units display, select Work with disk configuration. From the Work with Disk Configuration display, select Display disk configuration. From the Display Disk Configuration display, select Display disk configuration status. Write down the resource names (example DD001) of both the load source and load source mirrored unit (unit 1). At this point it is not necessary to know how to determine which resource name is assigned to the load source and the mirrored load source. Resource Name: _____ Resource Name: _____ From the Dedicated Service Tools (DST) or System Service Tools (SST) display, select Start a Service tool. From the Start a Service Tool display, select Hardware service manager. From the Hardware Service Manager, select Locate resource by resource name. Enter the first resource name that was written down in step 5 and press Enter. From the Logical Hardware Resources display, select Associated packaging resource(s). From the Packaging Resources Associated with a Logical Resource display, select Display detail. Write down the Physical location information for the disk unit. Resource Name: _____ Unit ID: _____ Planar ID: _____ Device: _____ Repeat steps 9-12 for the second resource name written down in step 5. Resource Name: _____ Unit ID: _____ Planar ID: _____ Device: _____ The physical locations of the card positions can be found by referencing the map located inside the front door of the IBM i or expansion tower. <p>Tip: Mark these disk units with tape to remember their locations.</p>	


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.

Load source planning steps	Planning requirements answers
Does your load source disk unit (unit 1) have mirrored protection?	Yes / No
Does your system or partition currently have two non-configured disk units with a capacity of at least 70-GB?	Yes / No
Do you know how to physically install and remove disk units from your system or partition? <u>Install and remove</u> replacement disk units on an IBM i.	Yes / No
Were you able to locate the load source and load source mirrored units in the Disk Configuration questions?	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 or partition 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 system using the <u>GO SAVE</u> command from the IBM i command line.
2.	Set the IPL mode to MANUAL.  Attention: Manual IPL is required for Dedicated Service Tools (DST).
3.	Restart your system or partition to Dedicated Service Tools (DST) by entering the following command from the IBM i command line. <pre>PWRDWN SYS *IMMED RESTART(*YES)</pre> Note: Verify that you are in manual mode before you restart your system.
4.	Begin Dedicated Service Tools. 1. From the IPL or Install the System menu, select Use Dedicated Service Tools . 2. Sign on to Dedicated Service Tools (DST).

Task	What to do
5.	<p>Suspend Mirroring on the mirrored load source disk unit (unit 1).</p> <ol style="list-style-type: none"> 1. From the Dedicated Service Tools (DST) menu, select Work with disk units. 2. From the Work with Disk Units display, select Work with disk unit recovery. 3. From the Work with Disk Unit Recovery display, select Suspend Mirrored protection. 4. Write down the resource name (example DD001) of the mirrored load source disk unit (unit 1). _____ 5. Type 1 next to the mirrored load source disk unit (unit 1) and press Enter.
6.	<p>Label the load source and mirrored load source disk units using the physical locations determined in the Disk Configuration questions. If the physical locations were found using Dedicated Service Tools (DST) or System Service Tools (SST), determine the mirrored load source physical location by matching the resource name found in task 5 step 4 with the resource names found in the Disk Configuration questions steps 12-13. Once the mirrored load source disk unit is labeled, label the other disk unit the load source disk unit.</p>
7.	<p>Confirm that non-configured disk units are available as replacements for the load source disk unit and mirrored load source disk unit.</p> <ol style="list-style-type: none"> 1. From the Use Dedicated Service Tools (DST) menu, select Work with disk units. 2. From the Work with Disk Units display, select Work with disk configuration. 3. From the Work with Disk Configuration display, select Display disk configuration. 4. From the Display Disk Configuration display, select Display non-configured disk units. 5. Write down the serial numbers and resource names of the non-configured disk units that are going to replace the load source and mirrored load source. <p>Load Source Replacement</p> <p>Serial Number: _____</p> <p>Resource Name: _____</p> <p>Mirrored Load Source Replacement</p> <p>Serial Number: _____</p> <p>Resource Name: _____</p> <p>Note: Ensure that there are two disk units of equal capacity and are at least 70-GB.</p>

Task	What to do
8.	<p>Determine the physical locations of the non-configured disk units that are replacing the load source and mirrored load source disk units that are documented in task 7 step 5.</p> <p>Determine the disk unit locations using Dedicated Service Tools (DST).</p> <ol style="list-style-type: none"> 1. From Use Dedicated Service Tools or System Service Tools display, select Start a Service Tool. 2. From the Start a Service Tool display, select Hardware service manager. 3. From the Hardware Service Manager, select Locate resource by resource name. 4. Enter the first resource name that was written down in task 7 step 5 and press Enter. 5. From the Logical Hardware Resources display, select Associated packaging resource(s). 6. From the Packaging Resources Associated with a Logical Resource display, select Display detail. 7. Write down the Physical location information for the load source replacement disk unit. Resource Name: _____ Unit ID: _____ Planar ID: _____ Device: _____ 8. Repeat steps 4-7 for the mirrored load source replacement that was written down in task 7 step 5. Resource Name: _____ Unit ID: _____ Planar ID: _____ Device: _____ 9. For internal disk units, the physical locations can be found by referencing the map located inside the front door of the IBM i or expansion tower.
9.	<p>For internal disk units, locate the physical locations of the non-configured disk units listed in task 8 step 7 through 8 and label each.</p> <p>Tip: Mark these disk units with tape to remember their locations.</p>
10.	<p>Power off the system or partition.</p> <ol style="list-style-type: none"> 1. From the Use Dedicated Service Tools (DST) menu, select Start a Service Tool. 2. From 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 (DST).</p>
11.	<p>Remove the mirrored load source disk unit.</p> <ol style="list-style-type: none"> 1. Physically remove the mirrored load source disk unit. <p>Note: This is the mirrored load source disk unit that was marked in task 6.</p>

Task	What to do
12.	<p>Power on your system or partition and begin Dedicated Service Tools (DST).</p> <ol style="list-style-type: none"> 1. Power on your system or partition. 2. From the IPL or Install the System menu, select Use Dedicated Service Tools. 3. Sign on to Dedicated Service Tools.
13.	<p>Copy the load source disk unit to the non-configured load source replacement disk unit written down in task 7 step 5.</p> <ol style="list-style-type: none"> 1. From the Use Dedicated Service Tools (DST) menu, select Work with disk units. 2. From the Work with Disk Units display, select Work with disk unit recovery. 3. From the Work with Disk Unit Recovery display, select Copy disk unit data. 4. Type 1 next to the load source disk unit (unit 1) and press Enter. 5. Type a 1 next to the non-configured disk unit documented as the load source replacement written down in task 7 step 5 and press Enter. 6. Use function F10 to accept the "Other sub-unit will become missing" error. 7. Verify that the correct non-configured disk unit was chosen as the Disk will contain the new copied information, and press Enter on the Confirm Copy Disk Unit Data display. 8. The status display shows how the operation is proceeding. When the function is complete, the partition will automatically IPL to the Use Dedicated Service Tools (DST) menu.
14.	<p>Power off the system or partition.</p> <ol style="list-style-type: none"> 1. From the Use Dedicated Service Tools (DST) menu, select Start a Service Tool. 2. From 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 (DST).</p>
15.	<p>Remove the original load source disk unit.</p> <ol style="list-style-type: none"> 1. Physically remove the original load source disk unit. <p>Note: This is the original load source disk unit that was marked in task 6.</p>
16.	<p>If the system has a Hardware Management Console (HMC), use the HMC to change the I/O tag.</p> <ol style="list-style-type: none"> 1. In your HMC, expand Systems Management > Servers. 2. Select the system. 3. Select the check box for the partition. 4. From the menu for the partition, click Configuration > Manage Profiles. 5. In the Managed Profiles pane, click the profile name. 6. In the Logical Partition Profile Properties pane, click the Tagged I/O tab. 7. From the Load source list, select the IOA that you controls the Load Source disk unit. The Load Source disk unit is the disk that was listed in task 7 step 5. 8. Click OK.
17.	<p>Power on your system or partition and begin Dedicated Service Tools (DST).</p> <ol style="list-style-type: none"> 1. Power on your system or partition. 2. From the IPL or Install the System menu, select Use Dedicated Service Tools. 3. Sign on to Dedicated Service Tools.

Task	What to do
18.	Replace the configured disk unit. <ol style="list-style-type: none"> 1. From the Use Dedicated Service Tools (DST) menu, select Work with disk units. 2. From the Work with Disk Units display, select Work with disk unit recovery. 3. From the Work with Disk Unit Recovery display, select Replace configured disk unit. 4. Type a 1 next to the suspended mirrored load source disk unit (unit 1) and press Enter. 5. Type a 1 next to the non-configured disk unit documented as the mirrored load source replacement in task 7 step 5. 6. Verify the correct non-configured disk unit was chosen, and press Enter on the Confirm Replace Configured Unit display. 7. The status display shows how the operation is proceeding. When the function is complete, return to the Use Dedicated Service Tools (DST) menu.
19.	Examine the configuration. <ol style="list-style-type: none"> 1. From the Use Dedicated Service Tools (DST) menu, select Work with disk units. 2. From the Work with Disk Units display, select Work with disk configuration. 3. From the Work with Disk Configuration display, select Display disk configuration. 4. From 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) and mirrored load source disk units match the serial numbers listed in task 7 step 5.
20.	IPL your system to IBM i.

If this procedure has been completed correctly, your load source will be upgraded. Mirrored protection will be active when the partition IPLs past DST.

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

Upgrading the load source disk unit with local mirroring using newly installed disk units

This checklist shows the sequence of tasks to upgrade your load source disk unit (unit 1) and the mirrored load source with disk units that have equal capacity, are at least 70-GB, and are mirror protected. The disk units that are replaced in this procedure are discarded.



Attention: These instructions are for all IBM i models.

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 13. Disk configuration questions

Disk configuration questions	Disk configuration answers
<p>Where is the load source disk unit and the mirrored load source disk unit located? Determine the disk unit locations using Dedicated Service Tools (DST) or System Service Tools (SST).</p> <ol style="list-style-type: none"> 1. From either Use Dedicated Service Tools (DST) or System Service Tools (SST) display, select Work with disk units. 2. From the Work with Disk Units display, select Work with disk configuration. 3. From the Work with Disk Configuration display, select Display disk configuration. 4. From the Display Disk Configuration display, select Display disk configuration status. 5. Write down the resource names (example DD001) of both the load source and load source mirrored unit (unit 1). At this point it is not necessary to know how to determine which resource name is assigned to the load source and the mirrored load source. Resource Name: _____ Resource Name: _____ 6. From the Dedicated Service Tools (DST) or System Service Tools (SST) display, select Start a Service tool. 7. From the Start a Service Tool display, select Hardware service manager. 8. From the Hardware Service Manager, select Locate resource by resource name. 9. Enter the first resource name that was written down in step 5 and press Enter. 10. From the Logical Hardware Resources display, select Associated packaging resource(s). 11. From the Packaging Resources Associated with a Logical Resource display, select Display detail. 12. Write down the Physical location information for the disk unit. Resource Name: _____ Unit ID: _____ Planar ID: _____ Device: _____ 13. Repeat steps 9-12 for the second resource name written down in step 5. Resource Name: _____ Unit ID: _____ Planar ID: _____ Device: _____ 14. The physical locations can be found by referencing the map located inside the front door of the IBM i or expansion tower. <p>Tip: Mark these disk units with tape to remember their locations.</p>	

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.

Load source planning steps	Planning requirements answers
Does your load source disk unit (unit 1) have mirrored protection?	Yes / No
Does your system or partition currently have at least one available disk slot to install a load source replacement disk unit or a virtual storage LUN?	Yes / No
Do you know how to physically install and remove disk units from your system? <u>Install and remove</u> replacement disk units on an IBM i.	Yes / No
Where you able to locate the load source and load source mirrored units in the Disk Configuration questions?	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 or partition 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 system using the <u>GO SAVE</u> command from the IBM i command line.
2.	Set the IPL mode to MANUAL. Note: Manual IPL is required for Dedicated Service Tools (DST).
3.	Restart your system or partition to Dedicated Service Tools (DST) by entering the following command from the IBM i command line. <pre>PWRDWN SYS *IMMED RESTART(*YES)</pre> Note: Verify that you are in manual mode before you power down your system.
4.	Begin Dedicated Service Tools. 1. From the IPL or Install the System menu, select Use Dedicated Service Tools . 2. Sign on to Dedicated Service Tools (DST).
5.	Suspend Mirroring on the mirrored load source disk unit (unit 1). 1. From the Dedicated Service Tools (DST) menu, select Work with disk units . 2. From the Work with Disk Units display, select Work with disk unit recovery . 3. From the Work with Disk Unit Recovery display, select Suspend Mirrored protection . 4. Write down the resource name (example DD001) of the mirrored load source disk unit (unit 1) that you are suspending. _____ 5. Type 1 next to the mirrored load source disk unit (unit 1) and press Enter.

Task	What to do
6.	<p>1. Write down the resource name of the load source disk unit that you suspended in task 5. Resource Name: _____</p> <p>2. From the Dedicated Service Tools (DST) display, select Start a Service tool.</p> <p>3. From the Start a Service Tool display, select Hardware service manager.</p> <p>4. From the Hardware Service Manager, select Locate resource by resource name.</p> <p>5. Enter the resource name that was written down in step 1 and press Enter.</p> <p>6. From the Logical Hardware Resources display, select Associated packaging resource(s).</p> <p>7. From the Packaging Resources Associated with a Logical Resource display, select Display detail.</p> <p>8. Write down the Physical location information for the disk unit that was suspended in task 5. Resource Name: _____ Unit ID: _____ Planar ID: _____ Device: _____</p>
7.	<p>Power off your system or partition.</p> <p>1. From the Use Dedicated Service Tools (DST) display, select Start a Service Tool.</p> <p>2. From the Start a Service Tool display, select Operator panel functions.</p> <p>3. Use function F10 to power off the system.</p> <p>4. Press Enter.</p> <p>Note: There is no command line because you are using Dedicated Service Tools.</p>
8.	<p>Install the first replacement disk unit.</p> <p>1. Install the replacement disk unit into the available disk slot or assign a virtual storage LUN.</p> <p>2. Write down the serial number of the replacement disk unit that is to going replace the load source. _____ </p> <p>Note: Ensure that the disk unit has at least a 70-GB capacity.</p>
9.	<p>Power on your partition and begin Dedicated Service Tools (DST).</p> <p>1. Power on your partition.</p> <p>2. Select Use Dedicated Service Tools (DST) from the IPL or Install the System menu.</p> <p>3. Sign on to Dedicated Service Tools (DST).</p>

Task	What to do
10.	<p>Copy the load source disk unit (unit 1) to the replacement disk unit.</p> <ol style="list-style-type: none"> From the Use Dedicated Service Tools (DST) menu, select Work with disk units. From the Work with Disk Units display, select Work with disk unit recovery. From the Work with Disk Unit Recovery display, select Copy disk unit data. Type 1 next to the load source disk unit (unit 1) and press Enter. Write down the resource name of the load source disk unit that you selected: _____ Type 1 next to the replacement disk unit that was installed in task 8 and press Enter. Use function F10 to accept the "Other sub-unit will become missing" error. Verify the correct non-configured disk unit was chosen as the Disk that is copied to, and press Enter on the Confirm Copy Disk Unit Data display. The status display shows how the operation is proceeding. When the function is complete, the partition will automatically IPL to the Use Dedicated Service Tools (DST) menu.
11.	<ol style="list-style-type: none"> From the Dedicated Service Tools (DST) or System Service Tools (SST) display, select Start a Service tool. From the Start a Service Tool display, select Hardware service manager. From the Hardware Service Manager, select Locate resource by resource name. Enter the resource name that was written down in task 10 step 4 and press Enter. From the Logical Hardware Resources display, select Associated packaging resource(s). From the Packaging Resources Associated with a Logical Resource display, select Display detail. Write down the Physical location information for the disk unit. Resource Name: _____ Unit ID: _____ Planar ID: _____ Device: _____
12.	<p>Power off the system or partition.</p> <ol style="list-style-type: none"> From the Use Dedicated Service Tools (DST) menu, select Start a service tool. From the Start a Service Tool display, select Operator panel functions. Use function F10 to power off the system. Press Enter to continue.
13.	<p>Physically remove the original load source disk unit.</p> <ol style="list-style-type: none"> Physically remove the original load source disk unit. <p>Note: This is the original load source disk unit that you copied in task 10, and identified in task 11.</p>
14.	<p>Install the second replacement disk unit that will be the mirror mate to the load source disk unit.</p> <ol style="list-style-type: none"> Install the replacement disk unit into the available disk slot on your system or assign a virtual storage LUN of the same size as the virtual storage LUN in task 8. Write down the serial number of the replacement disk unit that will be the mirror mate to the load source disk unit. _____ <p>Note: Ensure that the disk unit has at least a 70-GB capacity.</p>

Task	What to do
15.	<p>If the system has a Hardware Management Console (HMC), use the HMC to change the I/O tag.</p> <ol style="list-style-type: none"> 1. In your HMC, expand Systems Management > Servers. 2. Select the system. 3. Select the check box for the partition. 4. From the menu for the partition, click Configuration > Manage Profiles. 5. In the Managed Profiles pane, click the profile name. 6. In the Logical Partition Profile Properties pane, click the Tagged I/O tab. 7. From the Load source list, select the IOA that you want to control the Load Source disk unit. 8. Click OK.
16.	<p>Power on your system or partition and begin Dedicated Service Tools (DST).</p> <ol style="list-style-type: none"> 1. Power on your system or partition. 2. Select Use Dedicated Service Tools (DST) from the IPL or Install the System menu. 3. Sign on to Dedicated Service Tools (DST).
17.	<p>Replace the configured disk unit.</p> <ol style="list-style-type: none"> 1. From the Use Dedicated Service Tools (DST) menu, select Work with disk units. 2. From the Work with Disk Units display, select Work with disk unit recovery. 3. From the Work with Disk Unit Recovery display, select Replace configured disk unit. 4. Type a 1 next to the suspended mirrored load source disk unit (unit 1) and press Enter. 5. Type a 1 next to the replacement disk unit, based on the serial number, that was installed in task 14. 6. Verify the correct non-configured disk unit was chosen, and press Enter on the Confirm Replace Configured Unit display. 7. The status display shows how the operation is proceeding. When the function is complete, return to the Use Dedicated Service Tools (DST) menu.
18.	<p>Examine the configuration.</p> <ol style="list-style-type: none"> 1. From the Use Dedicated Service Tools (DST) menu, select Work with disk units. 2. From the Work with Disk Units display, select Work with disk configuration. 3. From the Work with Disk Configuration display, select Display disk configuration. 4. From 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) and mirrored load source disk units match the serial numbers listed in task 8 and task 14.
19.	<p>IPL your system to IBM i.</p>

If this procedure has been completed correctly, your load source will be upgraded. Mirrored protection will be active when the partition IPLs past DST.

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

Upgrading the load source disk unit without disk 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 70-GB capacity. The disk 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.

<i>Table 15. Disk configuration questions</i>	
Disk configuration questions	Disk configuration answers
<p>Where is the load source disk unit located?</p> <p>Determine the disk unit locations using Dedicated Service Tools (DST) or System Service Tools (SST).</p> <ol style="list-style-type: none"> 1. From either Use Dedicated Service Tools (DST) or System Service Tools (SST) display, select Work with disk units. 2. From the Work with Disk Units display, select Work with disk configuration. 3. From the Work with Disk Configuration display, select Display disk configuration. 4. From the Display Disk Configuration display, select Display disk configuration status. 5. Write down the resource name (example DD001) of the load source disk unit (unit 1). Resource Name: _____ 6. From the Dedicated Service Tools (DST) or System Service Tools (SST) display, select Start a Service tool. 7. From the Start a Service Tool display, select Hardware service manager. 8. From the Hardware Service Manager, select Locate resource by resource name. 9. Enter the first resource name that was written down in step 5 and press Enter. 10. From the Logical Hardware Resources display, select Associated packaging resource(s). 11. From the Packaging Resources Associated with a Logical Resource display, select Display detail. 12. Write down the Physical location information for the load source disk unit. Resource Name: _____ Unit ID: _____ Planar ID: _____ Device: _____ 13. The physical location can be found by referencing the map located inside the front door of the IBM i or expansion tower. <p>Tip: Mark this disk unit with tape to remember its location.</p>	

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 16. Load source planning steps

Load source planning steps	Planning requirements answers
Is your loadsource unprotected?	Yes / No
Do you know how to physically install and remove disk units from your system? <u>Install</u> and <u>remove</u> replacement disk units on an IBM i.	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 system using the <u>GO SAVE</u> command.
2.____	<p>Turn off your system or logical partition by entering the following command.</p> <pre>PWRDWN SYS *IMMED RESTART(*NO)</pre> <p>The Power Down System (PWRDWN SYS) command prepares the partition to end. Use the immediate (*IMMED) value to end all active jobs immediately and use the RESTART (*NO) value to power off the partition so that you can install disk units in the next task.</p>
3.____	<p>Install the replacement disk unit or assign a virtual storage LUN.</p> <ol style="list-style-type: none"> If you are not using a virtual storage LUN for the replacement unit, ensure that there is an open disk slot. Install the replacement disk unit on your system. <p>Notes:</p> <ol style="list-style-type: none"> Ensure that the disk unit being installed has at least a 70-GB capacity and has the same or greater capacity than the existing load source disk unit. It is recommended that you mark this disk unit with tape to remember its location. Mark this disk unit differently than how you marked the existing load source disk unit.
4.____	<p>Work with Dedicated Service Tools (DST).</p> <ol style="list-style-type: none"> 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 partition.</p>

Task	What to do
5._	<p>Copy the load source disk unit (unit 1) to the designated 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 the disk unit that you installed in task 3. 6. Write down the serial number of the replacement disk unit that is going replace the load source. _____ 7. Press Enter in the Confirm Copy Disk Unit Data display. 8. The status display shows how the operation is proceeding. When the function is complete, the partition will automatically IPL to the Use Dedicated Service Tools (DST) menu.
6.____	<p>Turn off the system or logical partition.</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 partition and press Enter. <p>Note: There is no command line because you are using Dedicated Service Tools.</p>
7.____	<p>Physically remove the disk unit.</p> <ol style="list-style-type: none"> 1. Physically remove the original load source disk unit (unit 1) that you identified in the Disk configuration questions.
8.____	<p>If the system has a Hardware Management Console (HMC) and the disk unit that you installed in task 3 is under a different IOA than the original load source disk unit, use the HMC to change the I/O tag.</p> <ol style="list-style-type: none"> 1. In your HMC, expand Systems Management > Servers. 2. Select the system. 3. Select the check box for the partition. 4. From the menu for the partition, click Configuration > Manage Profiles. 5. In the Managed Profiles pane, click the profile name. 6. In the Logical Partition Profile Properties pane, click the Tagged I/O tab. 7. From the Load source list, select the IOA that you want to control the Load Source disk unit. 8. Click OK.
9.____	<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
10.____	<p>Examine the configuration.</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 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 the replacement disk unit that you installed in task 3. 6. Verify the load source serial number.
11.____	IPL your system to IBM i.

If this procedure has been completed correctly, your load source will be upgraded.

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

Related reference

Upgrading the 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 70-GB capacity while keeping device parity active. The disk units that are replaced in this procedure are discarded.

Upgrading the load source disk unit with local mirroring using non-configured disk units

This checklist shows the sequence of tasks to upgrade your load source disk unit (unit 1) and the mirrored load source with disk units that have equal capacity, are at least 70-GB, and are mirror protected. The disk units that are replaced in this procedure are discarded.

Upgrading the load source disk unit with local mirroring using newly installed disk units

This checklist shows the sequence of tasks to upgrade your load source disk unit (unit 1) and the mirrored load source with disk units that have equal capacity, are at least 70-GB, and are mirror protected. The disk units that are replaced in this procedure are discarded.

Installing a new partition using a NVMe as the load source disk unit

To use a namespace on a NVMe device as a load source disk unit, follow this procedure.

Task	What to do
1.	Power off the partition.
2.	<p>If the system has a Hardware Management Console (HMC), use the HMC to change the I/O tag.</p> <ol style="list-style-type: none"> 1. In your HMC, expand Systems Management > Servers. 2. Select the system. 3. Select the check box for the partition. 4. From the menu for the partition, click Configuration > Manage Profiles. 5. In the Managed Profiles pane, click the profile name. 6. In the Logical Partition Profile Properties pane, click the Tagged I/O tab. 7. From the list, select the NVMe device. 8. Click OK. 9. Activate the partition; IPL with IPL Source D and the Key Lock position in Manual.

Task	What to do
3.	<p>Prepare the NVM device.</p> <ol style="list-style-type: none"> 1. On the Install Licensed Internal Code display, select Work with Dedicated Service Tools (DST). 2. Sign on. 3. On the Use Dedicated Service Tools (DST) display, select Work with disk units. 4. On the Work with Disk Units display, select Work with NVM Devices. 5. On the Work with NVM Devices display, select Delete existing NVM Namespaces. 6. Select the NVMe device. 7. If there are no namespaces listed, press F12 and go to step 10. 8. Select the 4=Delete Namespace option on each of the listed namespaces. 9. Press F10 to confirm the delete of the namespaces. 10. On the Work with NVM Devices display, select Create NVM namespaces. 11. Select the NVMe device. 12. Enter the quantity and capacity of the namespaces you want on the NVMe device. Select the capacity of the namespace. The LoadSource disk unit on the IBM i must be at least 70 GB and less than 2645 GB. Select a capacity within that range. Record the quantity and capacity values that you select; these will be needed in a later step. 13. Press F10 to confirm the creation of the namespace. 14. Press F12 to return to the Use Dedicated Service Tools (DST) display. 15. On the Install Licensed Internal Code menu, select Install Licensed Internal Code.
4.	<p>Install Licensed Internal Code.</p> <ol style="list-style-type: none"> 1. On the Install Licensed Internal Code display, select Install Licensed Internal Code and Initialize System. 2. Select the disk unit that is in the NVMe device. 3. Confirm the selection. 4. The NVMe disk unit is zeroed, the Licensed Internal Code is installed, and the partition will then IPL to DST.

Task	What to do
5.	<p>Add units to the ASP.</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 NVM Devices. 4. On the Work with NVM Devices display, select Create NVM namespaces. 5. Select the NVMe device which is not the same NVMe device that contains the Load Source. 6. Enter the same quantity and capacity of the namespaces that you used in task 3 step 12 above. 7. Press F10 to confirm the creation of the name space. 8. Press F12 to return to the Work with Disk Units display. 9. On the Work with Disk Units display, select Work with ASP configuration. 10. On the Work with ASP Configuration display, select Add units to ASPs. 11. On the Add units to ASPs display, select Add units to existing ASPs. 12. You will see a list of disk units. Type 1 in the Specify Asp column for each of the NVMe disk units. 13. Confirm the add of the disk units to the system ASP. 14. Press F12 to return to the Work with Disk Configuration display.
6.	<p>Start mirrored protection.</p> <ol style="list-style-type: none"> 1. On the Work with Disk Configuration display, select Work with mirrored protection. 2. On the Work with Mirrored Protection display, select Start mirrored protection. 3. Select ASP 1. 4. Confirm the start of mirrored protection. 5. The partition will update the configuration and IPL to DST.

Removing an operational NVMe device from a partition

To remove an operational NVMe device from an IBM i partition, follow this procedure.

Task	What to do
1.	<p>Identify the namespaces on the NVMe device.</p> <ol style="list-style-type: none"> 1. Start System Service Tools (STRSST), and specify the user name and password. 2. On the System Service Tools (SST) display, select Work with disk units. 3. On the Work with Disk Units display, select Display disk configuration. 4. On the Display Disk Configuration display, select Display NVM namespaces. 5. Locate the NVMe device that you want to remove. Write down the Serial Number and Resource Name of the NVMe device. You will use this information in future steps. 6. If there are no namespaces associated with the NVMe device, go to task 4. 7. If every namespace in the NVMe device that you located in step 5 has an asterisk (*) shown for the ASP and Unit number, go to task 3. 8. Write down the ASP, Unit, and Resource Name of each namespace listed under the NVMe device that you located in step 5. The namespaces are IBM i disk units.

Task	What to do
2.	<p>Suspend units from mirror protection.</p> <ol style="list-style-type: none"> 1. Press F12 to return to the Work with Disk Units display. 2. On the Work with Disk Units display, select Work with disk unit recovery. 3. On the Work with Disk Recovery display, select Suspend mirrored protection. 4. Locate each disk unit that you wrote down in task 1 step 8 that has an ASP and Unit number. Select 1=Suspend Mirrored Protection on each located unit. Press Enter to suspend mirror protection on the units. 5. Press F12 to return to the Work with Disk Units display.
3.	<p>Delete the namespaces from the NVMe device.</p> <ol style="list-style-type: none"> 1. On the Work with Disk Units display, select Work with disk configuration. 2. On the Work with Disk Configuration display, select Work with NVM Devices. 3. On the Work with NVM Device display, select Delete existing NVM Namespaces. 4. Select the NVMe device. 5. Select 4=Delete Namespace for each namespace listed for the NVM device. 6. Press F10 to confirm the delete of the name spaces. 7. Press F12 to return to the System Service Tools (SST) display.
4.	<p>Remove the NVMe device.</p> <ol style="list-style-type: none"> 1. On the System Service Tools (SST) display, select Start a service tool. 2. On the Start a Service Tool display, select Hardware service manager. 3. On the Hardware Service Manager display, select Packaging hardware resources. 4. On the Packaging Hardware Resources display, select 3=Concurrent maintenance for the Unit ID that contains the NVMe device. 5. Press F9 to power off the domain. 6. Physically remove the NVMe device. 7. Press F10 to power on the domain.

Upgrading the system ASP to NVMe - mirror protected

To upgrade the system ASP which contains all SAS or SAN LUNs to use all NVMe namespaces, where the current system ASP has mirror protection, follow this procedure.

Task	What to do
1.	<p>Determine if the system ASP is mirror protected.</p> <ol style="list-style-type: none"> 1. Start System Service Tools (STRSST), and specify the user name and password. 2. On the System Service Tools (SST) display, select Work with disk units. 3. On the Work with Disk Units display, select Work with disk configuration. 4. On the Work with Disk Configuration display, select Display disk configuration. 5. On the Display Disk Configuration display, select Display disk configuration protection. 6. Examine the column called 'Protection'. If the value in that column for ASP 1 is not 'Mirrored', then use the procedure “Upgrading the system ASP to NVMe - not mirror protected” on page 160.

Task	What to do
2.	<p>Display the current disk configuration and calculate the space requirements for the system ASP (ASP 1).</p> <ol style="list-style-type: none"> 1. Press F12 to return to the Display Disk Configuration display. 2. On the Display Disk Configuration display, select Display disk configuration capacity. 3. Examine the entry for the system ASP (ASP 1). 4. Write down the Size value in the column titled '---Unprotected---' for ASP 1. Write down the Size value in the column titled '----Protected-----' for ASP 1. 5. Add the two numbers you wrote down in step 4. You need enough NVMe devices to contain twice that capacity. IBM recommends that all NVMe units have mirror protection. 6. Ensure that you have the correct quantity of NVMe devices, and that the capacity of the devices is enough to hold the data in the ASP.
3.	<p>Backup the data on the partition.</p> <ol style="list-style-type: none"> 1. Use the GO SAVE command, and use option 21 from the Save menu to save the data.
4.	<p>Physically Install the NVMe devices.</p> <ol style="list-style-type: none"> 1. Install the NVMe devices.
5.	<p>Prepare the NVMe devices.</p> <ol style="list-style-type: none"> 1. Start System Service Tools (STRSST), and specify the user name and password. 2. On the System Service Tools (SST) display, select Work with disk units. 3. On the Work with Disk Units display, select Work with disk configuration. 4. On the Work with Disk Configuration display, select Work with NVM Devices. 5. Repeat these steps for each NVMe device that is installed on the partition. <ol style="list-style-type: none"> a. Select a NVMe device. b. Select Delete existing NVM namespaces. c. Select the NVMe device. d. If there are no namespaces listed for the NVMe device, press F12 and go to the next NVM device. e. Select the option 4=Delete Namespace on each of the listed namespaces. f. Press F10 to confirm the delete of the namespaces. 6. When all of the namespaces have been deleted on all of the NVMe devices, press F12 to return to the System Service Tools (SST) display.

Task	What to do
6.	<p>Create NVMe namespaces.</p> <ol style="list-style-type: none"> 1. On the System Service Tools (SST) display, 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 NVM Devices. 4. Repeat these steps for each NVMe device that is installed on the partition. <ol style="list-style-type: none"> a. Select Create NVM namespaces. b. Select the NVM device. c. Enter the quantity and capacity of the namespaces you want on the NVMe device. It is recommended that a capacity of 393 GB be used. d. Press F10 to confirm the creation of the namespaces. 5. When all of the namespaces have been created on all of the NVMe devices, press F12 to return to the System Service Tools (SST) display.
7.	<p>Identify two NVMe disk units which will become the Load Source and the mirror mate of the Load Source.</p> <ol style="list-style-type: none"> 1. On the System Service Tools (SST) display, select Work with disk units. 2. On the Work with Disk Units display, select Display disk configuration. 3. On the Display Disk Configuration display, select Display non-configured units. 4. Write down the serial number and resource name of a NVMe disk unit. This unit will become the Load Source disk unit in task 13. Serial number: _____ Resource name: _____ 5. Write down the serial number and resource name of a NVMe disk unit which is on a NVMe device different than the unit you found in step 4. This unit will become the mirror mate of the Load Source disk unit in task 14. Serial number: _____ Resource name: _____ 6. Press F12 to return to the System Service Tools (SST) display. 7. On the System Service Tools (SST) display, select Start a service tool. 8. On the Start a Service Tool display, select Hardware service manager. 9. On the Hardware Service Manager display, select Locate resource by resource name. 10. Type the resource name that you wrote down in step 4 and press Enter. 11. Select the option 8=Associated packaging resource(s) and press Enter. 12. Select the option 5=Display Detail and press Enter. 13. Write down the Unit ID, Planar ID, and Card values. This information is needed in task 12. Unit ID: _____ Planar ID: _____ Card: _____ 14. Press F12 to return to the System Service Tools (SST) display.

Task	What to do
8.	<p>Add NVMe units to the system ASP (ASP 1).</p> <ol style="list-style-type: none"> 1. On the System Service Tools (SST) display, 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 Add units to ASPs. 4. On the Add Units to ASPs display, select Add units to existing ASPs. 5. On the Specify ASPs to Add Units to display, you will see a list of non-configured NVMe disk units displayed. Type 1 on the Specify ASP field for the NVMe disk units to be added to the system ASP. <p>Notes: The System ASP is mirror protected, so you must select an even quantity of NVMe disk units to add to the System ASP. Two NVMe disks that reside in the same NVMe device cannot be used to construct a mirror protected unit.</p> <p>For example, consider this configuration: A NVMe device has three namespaces (disk units) and another NVMe device has 5 namespaces (disk units). You should select 3 disk units under each NVMe device to add to the System ASP. This ensures that the IBM i operating system can match the NVMe disk units under two different NVMe devices to construct a mirror protected unit.</p> <p>Ensure that the serial number and resource name you wrote down in task 7 step 4 is not in the list of units that you select.</p> <p>Ensure that the serial number and resource name you wrote down in task 7 step 5 is not in the list of units that you select.</p> <ol style="list-style-type: none"> 6. Confirm the add of the disk unit to ASP 1. 7. The status display shows the progress of the Add operation. When the Add operation is complete, press F12 to return to the Work with Disk Configuration display.
9.	<p>Remove units from the system ASP (ASP 1).</p> <p>Note: This task could take a significant amount of time to complete</p> <ol style="list-style-type: none"> 1. On the Work with Disk Configuration display, select Work with removing units from configuration. 2. On the Work with Removing Units From Configuration display, select Remove units from configuration. 3. On the Remove Units from Configuration display, Select the units in ASP 1 which are not NVMe units. You added the NVMe units to the system ASP in task 8. 4. Press Enter to confirm the remove of the non-NVMe disk units from the system ASP. 5. If any errors occur, correct those errors before proceeding. Contact your next level of support if needed. 6. Press F12 to return to the Work with Disk Units display.
10.	<p>Move data from the Load Source disk unit.</p> <p>Note: This task will take up to 2 hours to complete. Data is moved off of the Load Source disk unit, so that task 13 will run faster.</p> <ol style="list-style-type: none"> 1. Type STRASPBAL TYPE(*ENDALC) UNIT(1) and press Enter. 2. Type STRASPBAL TYPE(*MOVDTA) TIMLMT(120) and press Enter. 3. When the move is complete, type STRASPBAL TYPE(*RSMALC) UNIT(1) and press Enter.

Task	What to do
11.	<p>Suspend the Load Source unit from mirror protection.</p> <ol style="list-style-type: none"> 1. If the Load Source disk unit is not mirror protected, go to task 12. 2. On the Work with Disk Units display, select Work with disk unit recovery. 3. On the Work with Disk Unit Recovery display, select Suspend mirrored protection. 4. Select 1=Suspend Mirrored Protection on the Load Source disk unit (unit 1). 5. Press F12 to return to the Work with Disk Units display.
12.	<p>If the system has a Hardware Management Console (HMC), use the HMC to change the I/O tag.</p> <ol style="list-style-type: none"> 1. Ensure that all users have signed off the partition. 2. Type PWRDWN SYS*IMMED RESTART (*NO) to shut down the partition. 3. In your HMC, expand Systems Management > Servers. 4. Select the system. 5. Select the check box for the partition. 6. From the menu for the partition, click Configuration > Manage Profiles. 7. In the Managed Profiles pane, click the profile name. 8. In the Logical Partition Profile Properties pane, click the Tagged I/O tab. 9. From the Load source pulldown list, select the IOA with 'Non-Volatile Memory Controller' at the address you identified in task 7 step 13. 10. Click OK. 11. Expand Properties > General Properties. 12. Set the IPL Source to B. 13. Set the Key Lock position to Manual. 14. Click Save. 15. Activate the partition.
13.	<p>Copy disk unit data to the NVMe device.</p> <ol style="list-style-type: none"> 1. Sign on to Dedicated Service Tools (DST). 2. On the Use Dedicated Service Tools (DST) display, select Work with disk units. 3. On the Work with Disk Units display, select Work with disk unit recovery. 4. On the Work with Disk Unit Recovery display, select Copy disk unit data. 5. Type 1 next to the Load Source disk unit (unit 1) and press Enter. 6. Type 1 next to the NVMe disk unit that has the serial number and resource name you wrote down in task 7 step 4, and press Enter. 7. On the Confirm Copy Disk Unit Data display, press Enter. 8. The status display shows progress of the operation. When the copy function is complete, the partition will IPL.

Task	What to do
14.	<p>Create the mirror mate for the NVMe Load Source disk unit (unit 1).</p> <ol style="list-style-type: none"> 1. Sign on to Dedicated Service Tools (DST). 2. On the Use Dedicated Service Tools (DST) display, select Work with disk units. 3. On the Work with Disk Units display, select Work with disk unit recovery. 4. On the Work with Disk Unit Recovery display, select Replace configured unit. 5. On the Select Configured Unit to Replace display, type 1 for Unit 1. 6. Select the non-configured NVMe unit that has the serial number and resource name that you wrote down in task 7 step 5. 7. On the Confirm the Replace of Configured Unit Data display, press Enter. This will create a mirror mate for the Load Source disk unit.
15.	<p>IPL the partition to IBM i and examine the configuration.</p> <ol style="list-style-type: none"> 1. IPL to the IBM i Main Menu. 2. Start System Service Tools (STRSST), and specify the user name and password. 3. On the System Service Tools (SST) display, select Work with disk units. 4. On the Work with Disk Units display, select Work with disk configuration. 5. On the Work with Disk Configuration display, select Display disk configuration. 6. On the Display Disk Configuration display, select Display disk configuration status. 7. Examine the configuration information to ensure that all of the units in ASP 1 are NVME type of disks, and that all units in ASP 1 are mirror protected. <p>If the procedure has not been completed correctly, contact your next level of support.</p>

Upgrading the system ASP to NVMe - not mirror protected

To upgrade the system ASP which contains all SAS or SAN LUNs to use all NVMe namespaces, where the current system ASP is not mirror protected, follow this procedure.

Task	What to do
1.	<p>Determine if the system ASP is mirror protected.</p> <ol style="list-style-type: none"> 1. Start System Service Tools (STRSST), and specify the user name and password. 2. On the System Service Tools (SST) display, select Work with disk units. 3. On the Work with Disk Units display, select Work with disk configuration. 4. On the Work with Disk Configuration display, select Display disk configuration. 5. On the Display Disk Configuration display, select Display disk configuration protection. 6. Examine the column called 'Protection'. If the value in that column for ASP 1 is 'Mirrored', then use the procedure “Upgrading the system ASP to NVMe - mirror protected” on page 155.

Task	What to do
2.	<p>Display the current disk configuration and calculate the space requirements for the system ASP (ASP 1).</p> <ol style="list-style-type: none"> 1. Press F12 to return to the Display Disk Configuration display. 2. On the Display Disk Configuration display, select Display disk configuration capacity. 3. Examine the entry for the system ASP (ASP 1). 4. Write down the Size value in the column titled '---Unprotected---' for ASP 1. Write down the Size value in the column titled '----Protected-----' for ASP 1. 5. Add the two numbers you wrote down in step 4. You need enough NVMe devices to contain twice that capacity. IBM recommends that all NVMe units have mirror protection. 6. Ensure that you have the correct quantity of NVMe devices, and that the capacity of the devices is enough to hold the data in the ASP.
3.	<p>Backup the data on the partition.</p> <ol style="list-style-type: none"> 1. Use the GO SAVE command, and use option 21 from the Save menu to save the data.
4.	<p>Physically Install the NVMe devices.</p> <ol style="list-style-type: none"> 1. Install the NVMe devices.
5.	<p>Prepare the NVMe devices.</p> <ol style="list-style-type: none"> 1. Start System Service Tools (STRSST), and specify the user name and password. 2. On the System Service Tools (SST) display, select Work with disk units. 3. On the Work with Disk Units display, select Work with disk configuration. 4. On the Work with Disk Configuration display, select Work with NVM Devices 5. Repeat these steps for each NVMe device that is installed on the partition. <ol style="list-style-type: none"> a. Select a NVMe device. b. Select Delete existing NVM namespaces. c. Select the NVMe device. d. If there are no namespaces listed for the NVMe device, press F12 and go to the next NVM device. e. Select the option 4=Delete Namespace on each of the listed namespaces. f. Press F10 to confirm the delete of the namespaces. 6. When all of the namespaces have been deleted on all of the NVMe devices, press F12 to return to the System Service Tools (SST) display.

Task	What to do
6.	<p>Create NVMe namespaces.</p> <ol style="list-style-type: none"> 1. On the System Service Tools (SST) display, 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 NVM Devices. 4. Repeat these steps for each NVMe device that is installed on the partition. <ol style="list-style-type: none"> a. Select Create NVM namespaces. b. Select the NVM device. c. Enter the quantity and capacity of the namespaces you want on the NVMe device. It is recommended that a capacity of 393 GB be used. d. Press F10 to confirm the creation of the namespaces. 5. When all of the namespaces have been created on all of the NVMe devices, press F12 to return to the System Service Tools (SST) display.
7.	<p>Identify the NVMe namespace which will become the Load Source.</p> <ol style="list-style-type: none"> 1. On the System Service Tools (SST) display, select Work with disk units. 2. On the Work with Disk Units display, select Display disk configuration. 3. On the Display Disk Configuration display, select Display non-configured units. 4. Write down the serial number and resource name of a NVMe disk unit. This unit will become the Load Source disk unit in task 12. Serial number: _____ Resource name: _____ 5. Press F12 to return to the System Service Tools (SST) display. 6. On the System Service Tools (SST) display, select Start a service tool. 7. On the Start a Service Tool display, select Hardware service manager. 8. On the Hardware Service Manager display, select Locate resource by resource name. 9. Type the resource name that you wrote down in step 4 and press Enter. 10. Select the option 8=Associated packaging resource(s) and press Enter. 11. Select the option 5=Display Detail and press Enter. 12. Write down the Unit ID, Planar ID, and Card values. This information is needed in task 11. Unit ID: _____ Planar ID: _____ Card: _____ 13. Press F12 to return to the System Service Tools (SST) display.

Task	What to do
8.	<p>Add NVMe units to the system ASP (ASP 1).</p> <ol style="list-style-type: none"> 1. On the System Service Tools (SST) display, 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 Add units to ASPs. 4. On the Add Units to ASPs display, select Add units to existing ASPs. 5. On the Specify ASPs to Add Units to display, you will see a list of non-configured NVMe disk units displayed. Type 1 on the Specify ASP field for the NVMe disk units to be added to the system ASP. <p>Notes: The System ASP will be mirror protected in task 13, so you must select an even quantity of NVMe disk units to add to the System ASP. Two NVMe disks that reside in the same NVMe device cannot be used to construct a mirror protected unit.</p> <p>For example, consider this configuration: A NVMe device has three namespaces (disk units) and another NVMe device has 5 namespaces (disk units). You should select 3 disk units under each NVMe device to add to the System ASP. This ensures that the IBM i operating system can match the NVMe disk units under two different NVMe devices to construct a mirror protected unit.</p> <p>Ensure that the serial number and resource name you wrote down in task 7 step 4 is not in the list of units that you select.</p> <ol style="list-style-type: none"> 6. Confirm the add of the disk unit to ASP 1. 7. The status display shows the progress of the Add operation. When the Add operation is complete, press F12 to return to the Work with Disk Configuration display.
9.	<p>Remove units from the system ASP (ASP 1).</p> <p>Note: This task could take a significant amount of time to complete</p> <ol style="list-style-type: none"> 1. On the Work with Disk Configuration display, select Work with removing units from configuration. 2. On the Work with Removing Units From Configuration display, select Remove units from configuration. 3. On the Remove Units from Configuration display, Select the units in ASP 1 which are not NVMe units. You added the NVMe units to the system ASP in task 8. 4. Press Enter to confirm the remove of the non-NVMe disk units from the system ASP. 5. If any errors occur, correct those errors before proceeding. Contact your next level of support if needed. 6. Press F12 to return to the Work with Disk Units display.
10.	<p>Move data from the Load Source disk unit.</p> <p>Note: This task will take up to 2 hours to complete. Data is moved off of the Load Source disk unit, so that task 12 will run faster.</p> <ol style="list-style-type: none"> 1. Type STRASPBAL TYPE(*ENDALC) UNIT(1) and press Enter. 2. Type STRASPBAL TYPE(*MOVDTA) TIMLMT(120) and press Enter. 3. When the move is complete, type STRASPBAL TYPE(*RSMALC) UNIT(1) and press Enter.

Task	What to do
11.	<p>If the system has a Hardware Management Console (HMC), use the HMC to change the I/O tag.</p> <ol style="list-style-type: none"> 1. Ensure that all users have signed off the partition. 2. Type <code>PWRDWN SYS*IMMED RESTART(*NO)</code> to shut down the partition. 3. In your HMC, expand Systems Management > Servers. 4. Select the system. 5. Select the check box for the partition. 6. From the menu for the partition, click Configuration > Manage Profiles. 7. In the Managed Profiles pane, click the profile name. 8. In the Logical Partition Profile Properties pane, click the Tagged I/O tab. 9. From the Load source pulldown list, select the IOA with 'Non-Volatile Memory Controller' at the address you identified in task 7 step 12. 10. Click OK. 11. Expand Properties > General Properties. 12. Set the IPL Source to B. 13. Set the Key Lock position to Manual. 14. Click Save. 15. Activate the partition.
12.	<p>Copy disk unit data to the NVMe device.</p> <ol style="list-style-type: none"> 1. Sign on to Dedicated Service Tools (DST). 2. On the Use Dedicated Service Tools (DST) display, select Work with disk units. 3. On the Work with Disk Units display, select Work with disk unit recovery. 4. On the Work with Disk Unit Recovery display, select Copy disk unit data. 5. Type 1 next to the Load Source disk unit (unit 1) and press Enter. 6. Type 1 next to the NVMe disk unit that has the serial number and resource name you wrote down in task 7 step 4, and press Enter. 7. On the Confirm Copy Disk Unit Data display, press Enter. 8. The status display shows progress of the operation. When the copy function is complete, the partition will IPL.

Task	What to do
13.	<p>Start Mirror Protection on the System ASP (ASP 1).</p> <p>Note: IBM recommends that the NVMe devices be mirror protected.</p> <ol style="list-style-type: none"> 1. Sign on to Dedicated Service Tools (DST). 2. On the Use Dedicated Service Tools (DST) display, select Work with disk units. 3. On the Work with Disk Units display, select Work with disk configuration. 4. On the Work with Disk Configuration display, select Work with mirrored protection. 5. On the Work with Mirrored Protection display, select Start mirrored protection. 6. Select ASP 1. 7. Confirm the start of mirror protection on ASP 1. 8. If any errors occur during the start of mirror protection on ASP 1, contact your next level of support for assistance. 9. The disk configuration is updated and the partition IPLs to Dedicated Service Tools (DST). 10. Sign on to Dedicated Service Tools (DST).
14.	<p>IPL the partition to IBM i and examine the configuration.</p> <ol style="list-style-type: none"> 1. IPL to the IBM i Main Menu. 2. Start System Service Tools (STRSST), and specify the user name and password. 3. On the System Service Tools (SST) display, select Work with disk units. 4. On the Work with Disk Units display, select Work with disk configuration. 5. On the Work with Disk Configuration display, select Display disk configuration. 6. On the Display Disk Configuration display, select Display disk configuration status. 7. Examine the configuration information to ensure that all of the units in ASP 1 are NVMe type of disks, and that all units in ASP 1 are mirror protected. <p>If the procedure has not been completed correctly, contact your next level of support.</p>

Upgrading a user ASP to NVMe

To upgrade a user ASP which contains all SAS or SAN LUNs to use all NVMe namespaces, follow this procedure.

Note: This procedure is for a user ASP. If you want to upgrade the system ASP (ASP 1), use [“Upgrading the system ASP to NVMe - mirror protected” on page 155](#) or [“Upgrading the system ASP to NVMe - not mirror protected” on page 160](#).

Task	What to do
1.	<p>Display the current disk configuration and calculate the space requirements for the ASP.</p> <ol style="list-style-type: none"> 1. Start System Service Tools (STRSST), and specify the user name and password. 2. On the System Service Tools (SST) display, select Work with disk units. 3. On the Work with Disk Units display, select Display disk configuration. 4. On the Display Disk Configuration display, select Display disk configuration capacity. 5. Examine the entry for the user ASP. 6. Write down the Size value in the column titled '---Unprotected---' for the user ASP. Write down the Size value in the column titled '---Protected---' for the user ASP. 7. Add the two numbers you wrote down in step 6. You need enough NVMe devices to contain twice that capacity. IBM recommends that all NVMe units have mirror protection. 8. Ensure that you have the correct quantity of NVMe devices, and that the capacity of the devices is enough to hold the data in the ASP.
2.	<p>Backup the data on the partition.</p> <ol style="list-style-type: none"> 1. Use the GO SAVE command, and use option 21 from the Save menu to save the data.
3.	<p>Physically Install the NVMe devices.</p> <ol style="list-style-type: none"> 1. Install the NVMe devices.
4.	<p>Prepare the NVMe devices.</p> <ol style="list-style-type: none"> 1. Start System Service Tools (STRSST), and specify the user name and password. 2. On the System Service Tools (SST) display, select Work with disk units. 3. On the Work with Disk Units display, select Work with disk configuration. 4. On the Work with Disk Configuration display, select Work with NVM Devices. 5. Repeat these steps for each NVMe device that you installed on the partition in task 3. <ol style="list-style-type: none"> a. Select Delete existing NVM namespaces. b. Select the NVMe device. c. If there are no namespaces listed for the NVMe device, press F12 and go to the next NVM device. d. Select the option 4=Delete Namespace on each of the listed namespaces. e. Press F10 to confirm the delete of the namespaces. 6. When all of the namespaces have been deleted on all of the NVMe devices, press F12 to return to the System Service Tools (SST) display.

Task	What to do
5.	<p>Create NVMe namespaces.</p> <ol style="list-style-type: none"> 1. On the System Service Tools (SST) display, 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 NVM Devices. 4. Repeat these steps for each NVMe device that is installed on the partition. <ol style="list-style-type: none"> a. Select Create NVM namespaces. b. Select the NVM device. c. Enter the quantity and capacity of the namespaces you want on the NVMe device. It is recommended that a capacity of 393 GB be used. d. Press F10 to confirm the creation of the namespaces. 5. When all of the namespaces have been created on all of the NVMe devices, press F12 to return to the System Service Tools (SST) display.
6.	<p>Add NVMe units to the user ASP.</p> <ol style="list-style-type: none"> 1. On the System Service Tools (SST) display, 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 Add units to ASPs. 4. On the Add Units to ASPs display, select Add units to existing ASPs. 5. On the Specify ASPs to Add Units to display, you will see a list of non-configured NVMe disk units displayed. Type the user ASP number on the Specify ASP field for the NVMe disk units to be added to the user ASP. <p>Note: The User ASP will be mirror protected in task 9, so you must select an even quantity of NVMe disk units to add to the User ASP. Two NVMe disks that reside in the same NVMe device cannot be used to construct a mirror protected unit.</p> <p>For example, consider this configuration: A NVMe device has three namespaces (disk units) and another NVMe device has 5 namespaces (disk units). You should select 3 disk units under each NVMe device to add to the user ASP. This ensures that the IBM i operating system can match the NVMe disk units under two different NVMe devices to construct a mirror protected unit.</p> 6. Confirm the add of the disk unit to the ASP. 7. The status display shows the progress of the Add operation. When the Add operation is complete, press F12 to return to the Work with Disk Units display.

Task	What to do
7.	<p>Remove units from the user ASP.</p> <p>Note: This task could take a significant amount of time to complete</p> <ol style="list-style-type: none"> 1. On the Work with Disk Units display, select Work with disk configuration. 2. On the Work with Disk Configuration display, select Work with removing units from configuration. 3. On the Work with Removing Units From Configuration display, select Remove units from configuration. 4. On the Remove Units from Configuration display, select the option 4=Remove unit from configuration on the units in the user ASP which are not NVMe units. You added the NVMe units to the user ASP in task 6. 5. Press Enter to confirm the remove of the non-NVMe disk units from the user ASP. 6. If any errors occur, correct those errors before proceeding. Contact your next level of support if needed. 7. Press F12 to return to the Work with Disk Units display.
8.	<p>Determine if the user ASP is mirror protected.</p> <ol style="list-style-type: none"> 1. On the Work with Disk Units display, select Work with disk configuration. 2. On the Work with Disk Configuration display, select Display disk configuration. 3. On the Display Disk Configuration display, select Display disk configuration protection. 4. Examine the column called 'Protection'. If the value in that column for the user ASP is 'Mirrored', then the ASP is already mirror protected and you can proceed to task 10.
9.	<p>Start Mirror Protection on the user ASP, if needed.</p> <p>Note: IBM recommends that the NVMe devices be mirror protected.</p> <ol style="list-style-type: none"> 1. Ensure that all normal activities are ended and all users have signed off the logical partition. 2. IPL to DST by setting the keylock to Manual Mode and then type PWRDWNSYS *IMMED RESTART (*YES). 3. Sign on to Dedicated Service Tools (DST). 4. On the Use Dedicated Service Tools (DST) display, select Work with disk units. 5. On the Work with Disk Units display, select Work with disk configuration. 6. On the Work with Disk Configuration display, select Work with mirrored protection. 7. On the Work with Mirrored Protection display, select Start mirrored protection. 8. Select the user ASP. 9. Confirm the start of mirror protection on the ASP. 10. If any errors occur during the start of mirror protection on the ASP, contact your next level of support for assistance. 11. The disk configuration is updated and the partition IPLs to Dedicated Service Tools (DST). 12. Sign on to Dedicated Service Tools (DST). 13. IPL to IBM i.

Task	What to do
10.	<p>Examine the configuration.</p> <ol style="list-style-type: none"> 1. Start System Service Tools (STRSST), and specify the user name and password. 2. On the System Service Tools (SST) display, select Work with disk units. 3. On the Work with Disk Units display, select Work with disk configuration. 4. On the Work with Disk Configuration display, select Display disk configuration. 5. On the Display Disk Configuration display, select Display disk configuration status. 6. Examine the configuration information to ensure that all of the units in the user ASP are NVMe type of disks, and that all units in the user ASP are mirror protected. <p>If the procedure has not been completed correctly, contact your next level of support.</p>

Secure Data Erase for the IBM i

Use this procedure to complete the secure data deletion for the SAS attached disk drives, SAS attached solid-state drives, and PCIe attached NVMe devices connected to an IBM i partition.

Ensure that the firmware of the PCIe3 SAS adapter that controls the SAS attached disk drive or solid-state drive has firmware level 19512900, or later.

- If the partition is using IBM i 7.2, the adapter firmware level 19512900 is available in PTF MF67086.
- If the partition is using IBM i 7.3, the adapter firmware level 19512900 is available in PTF MF67087.
- If the partition is using IBM i 7.4, the adapter firmware level 19512900 is available in PTF MF67089.

Ensure that the IBM i Licensed Internal Code has the code level shown below, or later.

- The partition using IBM i 7.2 needs PTF MF99109.
- The partition using IBM i 7.3 needs PTF MF99208.
- The partition using IBM i 7.4 needs PTF MF99302.

To Display the PTF information on the partition: https://www.ibm.com/support/knowledgecenter/ssw_ibm_i_74/rzam8/rzam8fixinfostatus.htm

Task	What to do
1.	<p>Delete the data from all Auxiliary Storage Pools (ASPs).</p> <ol style="list-style-type: none"> 1. On the Use Dedicated Service Tools (DST) display, 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 Delete ASP data. 5. On the Select ASP to Delete Data From, select 4=Delete ASP Data for every ASP listed. 6. On the Confirm Delete ASP Data display, press F10 to confirm.
2.	<p>Delete the User Auxiliary Storage Pools (User ASPs).</p> <ol style="list-style-type: none"> 1. Press F3 to return to the Work with disk units display. 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 Delete user ASP. 5. If an error message says that there is no user ASP, go to Task 3. 6. On the Delete User ASP display, select 4=Delete for every ASP listed. 7. On the Confirm Delete of User ASP display, press F10 to confirm.

Task	What to do
3.	<p>Remove all units from the System Auxiliary Storage Pool (ASP 1).</p> <ol style="list-style-type: none"> 1. Press F3 to return to the Work with disk units display. 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 Work with removing units from configuration. 5. On the Work with Removing Units from Configuration display, select Remove units from configuration. 6. If an error message says there are No units available to be removed, go to Task 4. 7. On the Remove Units from Configuration display, select 4=Remove unit from configuration for every unit listed. 8. On the Confirm Remove Disk Units display, press Enter to confirm.
4.	<p>Stop mirrored protection.</p> <ol style="list-style-type: none"> 1. Press F3 to return to the Work with disk units display. 2. On the Work with Disk Units display, select Work with disk configuration. 3. On the Work with Disk Configuration display, select Work with mirrored protection. 4. On the Work with Mirrored Protection display, select Stop mirrored protection. 5. If an error message says there are no ASPs with mirrored protection, go to Task 5. 6. On the Select ASP to Stop Mirrored Protection display, select 1=Select to stop mirrored protection. 7. On the Confirm Stop Mirrored Protection display, press Enter to confirm. 8. The partition automatically IPLs to DST. 9. Sign on to DST. 10. On the Use Dedicated Service Tools (DST) display, select Work with disk units.
5.	<p>Stop device parity protection.</p> <ol style="list-style-type: none"> 1. On the Work with Disk Units display, select Work with disk configuration. 2. On the Work with Disk Configuration display, select Work with device parity protection. 3. On the Work with Device Parity Protection display, select Stop device parity protection. 4. If an error screen says there are no disk units eligible for the selected operation, go to Task 6. 5. On the Stop Device Parity Protection display, select 1=Stop device parity protection for every parity set listed. 6. On the Confirm Stop Device Parity Protection display, press F10 to confirm. 7. A status screen shows the progress of the stop device parity function.
6.	<p>Stop Hot Spare.</p> <ol style="list-style-type: none"> 1. On the Work with Disk Units display, select Work with disk configuration. 2. On the Work with Disk Configuration display, select Work with Hot Spare configuration. 3. On the Work with Hot Spare display, select Stop Hot Spare. 4. On the Stop Hot Spare display, select 1=Stop Hot Spare for every disk unit listed. If there are no Hot Spare units, go to Task 7. 5. On the Confirm Stop Hot Spare display, press Enter to confirm.

Task	What to do
7.	<p>If the partition is running IBM i 7.3 or IBM i 7.2, skip this task and go to Task 8.</p> <p>Delete NVMe namespaces.</p> <ol style="list-style-type: none"> 1. On the Work with Disk Configuration display, select Work with NVM Devices. 2. If there are no NVMe devices, go to Task 8. 3. For each NVM Device, do the following: <ol style="list-style-type: none"> a. On the Work with NVM Device display, select Delete existing NVM Namespaces. b. Select the NVM device. c. On the Delete Existing NVM Namespaces display, select 4=Delete Namespace for every namespaces listed. d. On the Confirm Delete Existing NVM Namespaces display, press F10 to confirm the delete of the namespaces.
8.	<p>Initialize and format the disk units.</p> <ol style="list-style-type: none"> 1. Press F3 to return to the Work with disk units display. 2. On the Work with disk units display, select Work with disk configuration. 3. On the Work with Disk Configuration display, select Work with disk unit recovery. 4. On the Work with Disk Unit Recovery display, select Disk unit problem recovery procedures. 5. On the Disk Unit Problem Recovery Procedures display, select Initialize and format disk unit. 6. On the Select Disk Units for Initialize and Format display, select 1=Select for every disk unit listed. 7. You may see a Problem Report display informing of you of disk units that are possibly configured. Press F10 to Ignore problems and continue. 8. On the Confirm Initialize and Format Disk Unit display, press F10 to confirm. 9. A status screen shows the progress of the Initialize function.

Task	What to do																																				
9.	<p>If the partition is running IBM i 7.3 or IBM i 7.2, go to Task 17.</p> <p>Record the information about the NVMe devices.</p> <ol style="list-style-type: none"> 1. Press F3 to return to the Work with disk units display. 2. On the Work with disk units display, select Work with disk configuration. 3. On the Work with Disk Configuration display, select Work with NVM Devices. 4. On the Work with NVM Device display, select Display NVM devices. 5. If there are no NVMe devices listed, go to Task 17. 6. Using the information shown on the Display NVM Devices display, update the appropriate table below to record the Serial Number, Resource Name, and Type of each NVM device. <div data-bbox="396 569 1468 835" style="border: 1px solid black; padding: 5px;"> <p><i>Table 17. Use this table to record information about NVMe devices with a type of 58FC, 58FD, or 58FE</i></p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="width: 33%;">Serial Number</th> <th style="width: 33%;">Resource Name</th> <th style="width: 33%;">Type</th> </tr> </thead> <tbody> <tr><td> </td><td> </td><td> </td></tr> <tr><td> </td><td> </td><td> </td></tr> <tr><td> </td><td> </td><td> </td></tr> <tr><td> </td><td> </td><td> </td></tr> <tr><td> </td><td> </td><td> </td></tr> </tbody> </table> </div> <ol style="list-style-type: none"> 7. <i>Table 18. Use this table to record information about NVMe devices with a type of 594A, 594B, 594C, 59B8, 59B9, or 59BA</i> <div data-bbox="396 926 1468 1108" style="border: 1px solid black; padding: 5px;"> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="width: 33%;">Serial Number</th> <th style="width: 33%;">Resource Name</th> <th style="width: 33%;">Type</th> </tr> </thead> <tbody> <tr><td> </td><td> </td><td> </td></tr> <tr><td> </td><td> </td><td> </td></tr> <tr><td> </td><td> </td><td> </td></tr> <tr><td> </td><td> </td><td> </td></tr> <tr><td> </td><td> </td><td> </td></tr> </tbody> </table> </div>	Serial Number	Resource Name	Type																Serial Number	Resource Name	Type															
Serial Number	Resource Name	Type																																			
Serial Number	Resource Name	Type																																			
10.	<p>Determine if the LoadSource disk unit is controlled by an NVMe .</p> <p>To find if the LoadSource disk unit is controlled by a NVMe, do the following:</p> <ol style="list-style-type: none"> 1. On the Work with NVM Device display, select Display NVM namespaces. 2. If there are no NVMe devices listed, go to Task 17. 3. Using the information shown on the Display NVM Namespaces display, look for a namespace that has a value of 1 in the ASP column and a value of 1 in the Unit column; that namespace is the LoadSource disk unit. <p>If one of the disk units under a NVMe is the LoadSource disk unit, record the LoadSource disk unit Serial Number here:</p> <p>LoadSource disk Serial Number: _____</p> <p>Record the serial number of the NMVe Device here:</p> <p>Serial Number of NVMe Device that contains the Load Source disk: _____</p>																																				

Task	What to do
11.	<p>Sanitize Block Erase each NVMe device listed in Table 18 on page 172 on Task 9.</p> <p>1. For each NVMe Device in Table 18 on page 172 in Task 9, do the following:</p> <p>Note: The NVMe device that contains the LoadSource disk unit will not be displayed on the selection screen.</p> <ol style="list-style-type: none"> On the Work with Disk Configuration display, select Work with NVM Devices. On the Work with NVM Device display, select Sanitize/Erase NVM device. On the Select NVM Device display, select the NVM device. On the Select Type of Sanitize display, select Sanitize Block Erase. On the Confirm Sanitize/Erase NVM Device display, press F10 to confirm your choice to Sanitize Block Erase the selected NVMe device.
12.	<p>Format Cryptographic Erase each NVMe device listed in Table 17 on page 172 on Task 9.</p> <p>1. For each NVMe Device in Table 17 on page 172 in Task 9, do the following:</p> <p>Note: Do not do this step for the NVMe device that contains the LoadSource disk unit. If a NVMe device contains the LoadSource disk unit, Task 10 has the serial number of the NVMe device.</p> <ol style="list-style-type: none"> On the Work with Disk Configuration display, select Work with NVM Devices. On the Work with NVM Device display, select Create NVM namespaces. On the Select NVM Device display, select the NVM device. On the Create NVM Namespaces display, enter the value 32 for the Quantity of namespaces to create and enter the value 32 for the Capacity of each namespace. On the Confirm Create NVM Namespaces display, Press F10 to confirm the choice to create namespaces. <p>2. For each NVMe Device in Table 17 on page 172 in Task 9, do the following:</p> <p>Note: The NVMe device that contains the LoadSource disk unit will not be displayed on the selection screen.</p> <ol style="list-style-type: none"> On the Work with NVM Device display, select Sanitize/Erase NVM device. On the Select NVM Device display, select the NVM device. On the Select Type of Sanitize display, select Format Cryptographic Erase. On the Confirm Sanitize/Erase NVM Device display, press F10 to confirm your choice to Format Cryptographic Erase the selected NVMe device.

Task	What to do
13.	<p>IPL the partition in D-mode to process the NVMe device that contains the LoadSource disk unit.</p> <ol style="list-style-type: none"> 1. Make sure you can access either the Hardware Management Console (HMC) or the Integrated Virtualization Manager to control your new logical partition. 2. If you are using Operations Console, see the Getting Started document in the Documentation directory of the IBM i Access Client Solutions CD. Section 9.9 Establishing a Console Connection to IBM i provides instruction on connecting a LAN console. 3. Verify that this logical partition has an alternate IPL resource (the resource for loading) assigned. For more information about alternate IPL resources, see the I/O devices topic. 4. Load the installation media volume that contains the Licensed Internal Code into the installation resource defined for this logical partition. 5. Use the control panel to verify or to set the mode selection to Manual and the IPL source to D. 6. Use the control panel to activate the logical partition. 7. Wait for the display to appear that has the language feature to select. (This could take several minutes). 8. The Select a Language Group display shows the service tools language that is currently installed on the partition. Press Enter to select the language. 9. On the Confirm Language Group display, press Enter to confirm the language group.

Task	What to do
14.	<p>Delete namespace and then Format Cryptographic Erase the NVMe device that contains the LoadSource disk unit.</p> <p>In Task 10, you identified the NVMe device that contains the LoadSource disk unit.</p> <p>If the NVMe device that controls the LoadSource is from Table 18 on page 172 in Task 9, go to Task 15.</p> <p>If the NVMe device that controls the LoadSource is from Table 17 on page 172 in Task 9, perform this task.</p> <ol style="list-style-type: none"> 1. Select Work with disk units. 2. On the Work with Disk Units display, select Work with NVM Devices. 3. On the Work with NVM Device display, select Delete Existing NVM Namespaces. 4. On the Select NVM Device display, select the NVMe device that you identified in Task 10. 5. Select the NVMe namespace with the serial number you recorded in Task 10. 6. Find the NVMe device whose serial number matches the information you recorded in Task 10. <p>Note: The NVMe Resource name shown on the screen might not match the information you recorded for the NVMe device in Table 17 on page 172 in Task 9.</p> <ol style="list-style-type: none"> 7. A 'Possibly Configured Units' warning screen may be presented. Press Enter to accept the warning and continue. 8. On the Confirm Delete Existing NVM Namespaces display, press F10 to confirm the delete of the namespace. 9. When the delete of the namespaces completes, Press F12 to return to the Work with Disk Units display. 10. On the Work with Disk Units display, select Work with NVM Devices. 11. On the Work with NVM Device display, select Create NVM namespaces. 12. On the Select NVM Device display, select the NVMe device that you identified in Task 10. 13. Find the NVMe device whose serial number matches the information you recorded in Task 10. <p>Note: The NVMe Resource name shown on the screen might not match the information you recorded for the NVMe device in Table 17 on page 172 in Task 9.</p> <ol style="list-style-type: none"> 14. On the Create NVM Namespaces display, enter the value 32 for the Quantity of namespaces to create and enter the value 32 for the Capacity of each namespace. 15. On the Confirm Create NVM Namespaces display, press F10 to confirm the creation of the namespaces. 16. On the Work with NVM Device display, select Sanitize/Erase NVM device. 17. On the Select NVM Device display, select the NVM device you identified in Task 10. 18. Find the NVMe device whose serial number matches the information you recorded in Task 10. <p>Note: The NVMe Resource name shown on the screen might not match the information you recorded for the NVMe device in Table 17 on page 172 in Task 9.</p> <ol style="list-style-type: none"> 19. On the Select Type of Sanitize display, select Format Cryptographic Erase. 20. On the Confirm Sanitize/Erase NVM Device display, press F10 to confirm the Format Cryptographic Erase of the selected NVMe device. 21. Go to Task 16 to power off the partition.

Task	What to do
15.	<p>Sanitize Block Erase of the NVMe device that contains the LoadSource disk unit.</p> <ol style="list-style-type: none"> 1. If the NVMe device is from Table 18 on page 172 in Task 9, do the following: <ol style="list-style-type: none"> a. Select Work with disk units. b. On the Work with Disk Units display, select Work with NVM Devices. c. On the Work with NVM Device display, select Sanitize/Erase NVM device. d. On the Select NVM Device display, select the NVM device that you identified in Task 10. e. Find the NVMe device whose serial number matches the information you recorded in Task 10. <p>Note: The Resource Name of the NVMe device might not match the information for the NVMe device that you recorded in Task 9.</p> f. On the Select type of Sanitize display, select Sanitize Block Erase. g. On the Confirm Sanitize/Erase NVM Device display, press F10 to confirm your choice to Sanitize Block Erase the selected NVMe device.
16.	<p>Power off the partition.</p> <ol style="list-style-type: none"> 1. F3 to return to Use Dedicated Service Tools (DST) display. 2. On the Use Dedicated Service Tools (DST) display, select Start a service tool. 3. On the Start a Service Tool display, select Operator panel functions. 4. On the Operator Panel Functions display, enter the value 3 for IPL source: and the value 1 for IPL Mode: 5. Press F10 to power off the partition. 6. Press Enter to confirm the power down of the partition <p>The Secure Data Erase procedure is complete.</p>

Task	What to do
17.	<p>Erase data on the LoadSource disk unit.</p> <ol style="list-style-type: none"> 1. Make sure you can access either the Hardware Management Console (HMC) or the Integrated Virtualization Manager to control your new logical partition. 2. If you are using Operations Console, see the Getting Started document in the Documentation directory of the IBM i Access Client Solutions CD. Section 9.9 Establishing a Console Connection to IBM i provides instruction on connecting a LAN console. 3. Verify that this logical partition has an alternate IPL resource (the resource for loading) assigned. For more information about alternate IPL resources, see the I/O devices topic. 4. Load the installation media volume that contains the Licensed Internal Code into the installation resource defined for this logical partition. 5. Use the control panel to verify or to set the mode selection to Manual and the IPL source to D. 6. Use the control panel to activate the logical partition. 7. Wait for the display to appear that has the language feature to select. (This could take several minutes) 8. The Select a Language Group display shows the service tools language that is currently installed on the partition. Press Enter to select the language. 9. On the Confirm Language Group display, press Enter to confirm the language group. 10. On the Install Licensed Internal Code display, select Install Licensed Internal Code. 11. On the Install Licensed Internal Code (LIC) display, select Install Licensed Internal Code and Initialize system. 12. On the Install LIC and Initialize System - Confirmation display, press F10 to Confirm. 13. A status screen shows the progress of the operation. 14. The partition automatically IPLs to DST when the function is complete. <p>The Secure Data Erase procedure is complete.</p>

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?](#)

IBM Navigator for i graphical user interface

1. [How do I access the IBM Navigator for i disk management function?](#)
2. [What is the difference between the disk management functions in IBM Navigator for i and the character-based command interface?](#)
3. [What is the service tools system?](#)
4. [Why does the data I see in IBM Navigator for i appear to be out of date?](#)
5. [Why can't I connect to the service tools system after I add the service table entry?](#)

Configuring

1. [How do I create a new independent disk pool?](#)
2. [How do I create a disk pool group?](#)

Performance

1. [Why is performance slow?](#)
2. [Why is performance slow for cross-site 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?](#)
7. [What is a CPF2204 message, and how can it be fixed?](#)
8. [What is a CPF22E2 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 system. 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 [“Independent disk pools” on page 80](#) 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 IBM i cluster, and a single-system environment with a single IBM i. See [“Benefits of independent disk pools” on page 83](#) for details.

[Back to questions](#)

How should I structure my independent disk pools?

There are recommendations for structuring and populating your independent disk pools. See [“Recommended structure for independent disk pools” on page 93](#) 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 82](#) for details.

[Back to questions](#)

[Back to questions](#)

IBM Navigator for i graphical user interface

How do I access the IBM Navigator for i disk management function?

Before you can access disk management functions, you must complete some setup tasks. See [“Enabling and accessing disk units” on page 4](#) for details.

[Back to questions](#)

What is the difference between the disk management functions in IBM Navigator for i and in the character-based (command) interface?

Support for many independent disk pool tasks are only available through IBM Navigator for i. Almost all disk management functions that are available from the system service tools (SST) mode are available through IBM Navigator for i.

[Back to questions](#)

[Back to questions](#)

What is the service tools system?

The service tools system 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 system. See [“Setting up communication”](#) on page 4 for details.

[Back to questions](#)

Why does the data I see in the IBM Navigator for i window appear to be out of date?

The disk management function in IBM Navigator for i caches information, and therefore needs to be refreshed to have the most current data visible. After you make a configuration change, IBM Navigator for i should refresh itself. If it does not, however, you can manually refresh it by clicking the Refresh button on the IBM Navigator for i. 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 system 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 system field. Make sure that TCP is lowercase. If it is not, then remove the entry, and recreate 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?

See [“Configuring independent disk pools”](#) on page 96 for details.

[Back to questions](#)

How do I create a disk pool group?

See [“Creating a new disk pool group”](#) on page 98 for details.

[Back to questions](#)

[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.

[Back to questions](#)

Why is performance slow for cross-site mirroring?

For cross-site 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. The volume of disk unit writes that your applications require also plays a role in the performance of your cross-site 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 IBM Navigator for i failed or was closed while the disk pool was being created, you might need to [“Converting UDFS disk pools” on page 96](#) 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 [“Converting UDFS disk pools” on page 96](#) 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](#)

What is a CPF2204 message, and how can it be fixed?

If you are trying to access disk units using IBM Navigator for i and you get informational message CPF2204 - User profile *SST not found, it is most likely that you have not created the service tools user ID. See [“Enabling and accessing disk units” on page 4](#).

What is a CPF22E2 message, and how can it be fixed?


If you are trying to access disk units using IBM Navigator for i and you get informational message CPF22E2 - Password not correct for user profile *SST, it is most likely that you have not set the service tools user ID password to password level 2, or you did not change the SST password to apply the new password level. See [“Enabling and accessing disk units”](#) on page 4.

[Back to questions](#)

Related information for Disk management

Product manuals, IBM Redbooks (in PDF format), Web sites, and other information center topic collections contain information that relates to the Disk Management topic collection. You can view or print any of the PDF files.

IBM Redbooks

- [IBM i and IBM Storwize Family: A Practical Guide to Usage Scenarios](#)  This IBM Redbooks publication gives a broad overview of the IBM Storwize family products and their features and functions. It describes the setup that is required on the storage side for attaching IBM Storwize family products to an IBM i environment.

Other information

- [Availability roadmap](#)
- [High Availability](#)

Notices

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

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

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

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

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

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan

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

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

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

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

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

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

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

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. _enter the year or years_.

Programming interface information

This Disk management publication documents intended Programming Interfaces that allow the customer to write programs to obtain the services of IBM i.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "[Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml)" at www.ibm.com/legal/copytrade.shtml.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Linux® is a registered 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 product and service names might be trademarks of IBM or other companies.

Terms and conditions

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

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

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

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

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

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

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



Product Number: 5770-SS1