


IBM

@server

iSeries

Disk protection





@server

iSeries

Disk protection

Contents

Part 1. Disk protection	1
Chapter 1. Choosing disk protection tools	3
Disk pools	3
Decide how to configure user disk pools	4
Consider creating a new disk pool on an active system	7
Make sure that your system has enough working space	7
Device parity protection	12
Planning for device parity protection	13
How device parity protection affects performance	18
Using both device parity protection and mirrored protection	20
Mirrored protection	21
Mirrored protection—benefits	22
Mirrored protection—costs and limitations	22
Planning for mirrored protection	23
Remote DASD mirroring support	36
Chapter 2. Choosing your level of protection	41
Comparison of disk protection options	41
Full mirrored protection versus partial mirrored protection	42
How your system manages auxiliary storage	43
How disks are configured	43
Full protection — single disk pool	45
Full protection — multiple disk pools	45
Partial protection — multiple disk pools	46
Assigning disk units to disk pools	46

Part 1. Disk protection

In addition to having a working backup and recovery strategy, you should also employ some form of data protection on your system. The way to do this is by using disk protection. Disk protection can help prevent data loss, and can keep your system from stopping if you experience a disk failure. There are several disk protection methods that you can use to help protect your data. You can use these methods in different combination with each other.

You can use disk management wizards in iSeries Navigator to help you configure disk pools and protect them with device parity protection or mirrored protection.

Remember: Although disk protection can reduce downtime or make recovery faster, it is **not** a replacement for regular backups. Disk protection cannot help you recover from a complete system loss, a processor failure, or a program failure.

These topics provide information on the different types of disk protection, and using the types with one another:

- Choosing disk protection tools
- Choosing your level of protection

Before proceeding, you may want to review these topics:

- How your system manages auxiliary storage
- How disks are configured

Chapter 1. Choosing disk protection tools

When you think about protecting your system from data loss, you need to think about the following:

Recovery

Can you get back the information that you lost, either by restoring it from backup media or by creating it again?

Availability

Can you reduce or eliminate the amount of time that your system is unavailable after a problem occurs?

Serviceability

Can you service it without affecting the data user?

Your first defense against loss of data is a good backup and recovery strategy. You need a plan for regularly saving the information on your system.


Several disk availability tools are available to reduce or eliminate system downtime and help with data recovery after a disk failure:

- Disk pools
- Device parity protection
- Mirrored protection

Disk pools

A disk pool, also referred to as an auxiliary storage pool (ASP) in the character-based interface, is a software definition of a group of disk units on your system. This means that a disk pool does not necessarily correspond to the physical arrangement of disks. Conceptually, each disk pool on your system is a separate pool of disk units for single-level storage. The system spreads data across the disk units within a disk pool. If a disk failure occurs, you need to recover only the data in the disk pool that contained the failed unit. There are two main categories of disk pools, the system disk pool and user disk pools. There are two types of user disk pools: basic and independent. Independent disk pools are further divided into primary, secondary, and UDFS disk pools. See the following links and the disk pool figure to understand the different types of user disk pools:

- System disk pool
- User disk pools



Your system may have many disk units attached to it for disk pool storage. To your system, they look like a single unit of storage. The system spreads data across all disk units. You can use disk pools to separate your disk units into logical subsets. For more ideas on how to use disk pools on your system, see [Disk pools — example uses](#).

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. There are also additional benefits as well as certain costs and limitations that are inherent in using disk pools.

For more information on user disk pools, see the following topics:

- [Deciding how to configure user disk pools](#)
- [Consider creating a new disk pool on an active system](#)
- [Making sure that your system has enough working space](#)
- [Contrasting basic and independent disk pools](#)

For information on how to implement disk pools in your business, see the [Backup and Recovery Guide](#).



Decide how to configure user disk pools

You can use disk pools for several different purposes, depending on your business needs. Before you configure any user disk pools, examine these topics that describe the various uses.

- [Using disk pools for availability](#)
- [Using disk pools for improved performance](#)

- Using disk pools with document library objects
- Using disk pools with extensive journaling
- Using disk pools with access path journaling

Using disk pools for availability

Different parts of your system may have different requirements for availability and recovery. For example, you may have a large history file that is changed only at the end of the month. The information in the file is useful but not critical. You might put this file in a separate library in a user disk pool that does not have any disk protection (mirrored protection or device parity protection). You could omit this library from your daily save operations. Save it only at the end of the month when it is updated.

Another example would be documents and folders. Some are critical to the organization. Those documents and folders should be protected with device parity protection or mirrored protection. They can be put in a protected user disk pool. Others are kept on the system to provide information but do not change very often. They can be in a different user disk pool, with a different strategy for saving and for protection.

Using disk pools for improved performance


If you are using user disk pools for better system performance, consider dedicating the 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 usually does not improve performance to place a single device-parity protected unit in a user disk pool because the performance of that unit is affected by other disk units in the device parity set.

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

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

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

see the *Work Management*  book. If the units have excessive use, you should consider adding more disk units to the user disk pool.


Using disk pools with document library objects

You can place document library objects (DLOs) in user disk pools. These are the possible advantages of placing DLOs in user disk pools:

- The ability to reduce save times for DLOs and to separate them by their save requirements.
- The ability to separate DLOs by availability requirements. Critical DLOs can be placed in user disk pools that are protected by mirrored protection or device parity protection. DLOs that change infrequently can be placed in unprotected disk pools with slower drives.
- The ability to grow to a larger number of documents.

If you have a current release of the OS/400 licensed program, you can run multiple `SAVDLO` or `RSTDLO` procedures against different disk pools. You can also run multiple `SAVDLO` operations on the same disk pool.

One approach for placing DLOs in user disk pools is to leave only system DLOs (IBM-supplied folders) in the system disk pool. Move other folders to user disk pools. The system folders do not change frequently, so they can be saved infrequently. "How to Transfer a Folder to a Different disk pool" in Backup and

Recovery  , describes the procedure to follow when moving folders from the system disk pool to user disk pools or between user disk pools.

You can specify a disk pool on the SAVDLO command. This allows you to save all the DLOs from a particular disk pool on a given day of the week. For example, you could save DLOs from disk pool 2 on Monday, DLOs from disk pool 3 on Tuesday, and so on. You could save all changed DLOs daily.

The recovery steps if you use this type of save technique would depend on what information was lost. If you lost an entire disk pool, you would restore the last complete saved copy of DLOs from that disk pool. You would then restore the changed DLOs from the daily saves.

When you save DLOs from more than one disk pool in the same operation, a different file and a sequence number will be created on the tape for each disk pool. When you restore, you must specify the correct sequence number. This makes it simple to restore the changed DLOs only to the disk pool that was lost without needing to know all the folder names.


When you specify DLO(*SEARCH) or DLO(*CHG) for the SAVDLO command, specify a disk pool, if possible. Specifying a disk pool saves system resources.

Restrictions for DLOs in User Disk Pools: These restrictions and limitations apply when placing DLOs in user disk pools:

- When using a save file for a save operation, you can save DLOs from only one disk pool.
- If you are saving to a save file and you specify SAVDLO DLO(*SEARCH) or SAVDLO DLO(*CHG), you must also specify a disk pool, even if you know the results of your search are found in a single disk pool.
- Documents that are not in folders must be in the system disk pool.
- Mail can be filed into a folder on a user disk pool. Unfiled mail is in the system disk pool.


Using disk pools with extensive journaling

If journals and objects being journaled are in the same disk pool as the receivers and the disk pool overflows, you must end journaling of all objects and recover from the overflowed condition for the disk

pool. Backup and Recovery  describes how to recover an overflowed disk pool.

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

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

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

Using disk pools with access path journaling

If you plan to use explicit access path journaling, IBM® recommends that you first change the journal to a journal receiver in the system disk pool (disk pool 1) for a few days. Start access path journaling to see

storage requirements for the receiver before you allocate the specific size for a user disk pool. Journal Management, provides more information about how to evaluate the storage requirements for journaling.

Consider creating a new disk pool on an active system


Starting with V3R6 of the OS/400 licensed program, you can add disk units while your system is active. When you add disk units to a disk pool that does not currently exist, the system creates a new disk pool. See Add a disk unit or disk pool for steps to configure a disk pool. If you choose to create a new user disk pool while your system is active, be sure you understand these considerations:

- You cannot start mirrored protection for a basic disk pool while the system is active. You can start mirrored protection for an unavailable independent disk pool when the system is active. The new disk pool is not fully protected unless all of the disk units have device parity protection.
- You cannot move existing disk units to a basic disk pool while your system is active. The system must move data when it moves disk units. This can be done only through Dedicated Service Tools (DST). It is not possible to move disk units from an existing disk pool to an independent disk pool.
- The system uses the size of a user disk pool to determine the storage threshold for the journal receivers that are used by system-managed access-path protection (SMAPP). When you create a disk pool while your system is active, the size of the disk units that you specify on the operation that creates the disk pool is considered the size of the disk pool for SMAPP. For example, assume that you add 2 disk units to a new disk pool, disk pool 2. The total capacity of the 2 disk units is 2062MB. Later, you add 2 more disk units to increase the capacity to 4124MB. For purposes of SMAPP, the size of the disk pool remains 2062MB until the next time you perform an IPL or vary on an independent disk pool. This means that the storage threshold of your SMAPP receivers is lower and the system must change receivers more often. Usually, this will not have a significant impact on system performance.

The system determines the capacity of every disk pool when you perform an IPL or vary on an independent disk pool. At that time, the system makes adjustments to its calculations for SMAPP size requirements. See System-managed access-path protection for more information about SMAPP.

Make sure that your system has enough working space

When you make changes to your disk configuration, the system may need working space. This is particularly true if you plan to move disk units from one disk pool to another disk pool. The system needs to move all the data from the disk unit to other disk units before you move it. "How to Calculate Space

Requirements for an Auxiliary Storage Pool" in Backup and Recovery  provides examples of how to determine how much working storage you need for your situation. There are also system limits for the amount of auxiliary storage.

If your system does not have sufficient interim storage, begin by cleaning up your disk storage. Many times, users keep objects on the system, such as old spooled files or documents, when these objects are no longer needed. Consider using the automatic cleanup function of Operational Assistant to free some disk space on your system.

If cleaning up unnecessary objects in auxiliary storage still does not provide sufficient interim disk space, another alternative is to remove objects from your system temporarily. For example, if you plan to move a large library to a new user disk pool, you can save the library and remove it from the system. Then restore the library after you have moved disk units. Here is an example for doing this:

1. Save private authorities for the objects on your system by typing:
`SAVSECDTA DEV(tape-device)`
2. Save the object by using the appropriate SAVxxx command. For example, to save a library use the SAVLIB command. Consider saving the object twice to 2 different tapes.
3. Delete the object from the system by using the appropriate DLTxxx command. For example, to delete a library, use the DLTLIB command.
4. Recalculate your disk capacity to determine whether you have made sufficient interim space available.

5. If you have enough space, perform the disk configuration operations.
6. Restore the objects that you deleted.

Disk pools — example uses

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

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

Disk pools—benefits

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

- **Additional data protection.** By separating libraries, documents, or other objects in a user disk pool, you protect them from data loss when a disk unit in the system disk pool or other user disk pool fails. For example, if you have a disk unit failure, and data contained on the system disk pool is lost, objects contained in user disk pools are not affected and can be used to recover objects in the system disk pool. Conversely, if a failure causes data that is contained in a user disk pool to be lost, data in the system disk pool is not affected.
- **Improved system performance.** Using disk pools can also improve system performance. This is because the system dedicates the disk units that are associated with a disk pool to the objects in that disk pool. For example, suppose you are working in an extensive journaling environment. Placing journals and journaled objects in a user 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.
- **Higher availability and flexibility.** See Benefits of independent disk pools for more benefits that are unique to independent disk pools.

Disk pools— costs and limitations

There are some specific limitations that you may encounter when using 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 manage related objects. Some related objects, such as journals and journaled objects, must be in the same user disk pool.

System disk pool

The system automatically creates the system disk pool (disk pool 1) which contains disk unit 1 and all other configured disks that are not assigned to a user disk pool. The system disk pool contains all system objects for the OS/400 licensed program and all user objects that are not assigned to a basic or independent disk pool.

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

There are additional considerations that you should be aware of regarding the capacity of the system disk pool and protecting your system disk pool.

Capacity of the system disk pool: If the system disk pool fills to capacity, the system will terminate 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 reoccurring.

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) and the system message queue (QSYSMSG) are 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. If you ignore this message, the system disk pool will fill to capacity, and the system will end abnormally.

You can use a third method of preventing the system disk pool from filling to capacity by using the QSTGLOWLMT and QSTGLOWACN system values. For more information, refer to "How to Change the

Storage Threshold for the System Auxiliary Storage Pool" in Backup and Recovery .

Protecting your system disk pool: IBM recommends that you use device parity protection or mirrored protection on the system disk pool. Using disk protection tools reduces 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.

User 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. User disk pools can contain libraries, documents and certain types of objects. User disk pools exist in two forms: basic disk pools and independent disk pools. In a clustered environment independent disk pools can be switched between systems without having to perform an IPL, allowing for continuously available data. You can configure basic disk pools with numbers 2 through 32. Independent disk pools are numbered 33 through 255. To learn more about how basic and independent disk pools differ, see Contrasting basic and independent disk pools.

See the following topics for more information about library and non-library disk pools:

- Library user disk pools
- Non-library user disk pools

Once you have disk pools configured, you should protect them by using mirroring or device parity protection.

Library user disk pools: Library user disk pools, contain libraries and user-defined file systems (UDFS). IBM recommends that you use library user disk pools because the recovery steps are easier than with non-library user disk pools. There are several factors to consider when using library user disk pools.

What You Should Know About Library User disk pools:

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

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

- Start Journal Physical File (STRJRNPf) command for physical files
- Start Journal Access Path (STRJRnAP) command for access paths
- Start Journal (STRJRn) command for integrated file system objects
- Start Journal Object (STRJRnOBJ) command for other object types

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

- No database network can cross disk pool boundaries. You cannot create a file in one disk pool that depends on a file in a different disk pool. All based-on physical files for a logical file must be in the same disk pool as the logical file. The system builds access paths only for database files in the same disk pool as the based-on physical file (temporary queries are not limited). Access paths are never shared by files in different disk pools. Record formats are not shared between different disk pools. Instead, a format request is ignored and a new record format is created.
- You can place an SQL collection in a user disk pool. You specify the target disk pool when you create the collection.
- If the library user disk pool does not contain any database files, set the target access path recovery time for the disk pool to *NONE. This would 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.

Non-library user disk pools: Non-library 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 non-library user disk pool to *NONE. A non-library 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 non-library 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.

Protecting disk pools: Keep the following points in mind regarding disk pool protection:

- All disk pools, including the system disk pool, should have mirrored protection or consist entirely of disk units with device parity protection to ensure that the system continues to run after a disk failure in a disk pool.
- If a disk failure occurs in a disk pool that does not have mirrored protection, the system may not continue to run, depending on the type of disk unit and the error.
- If a disk failure occurs in a disk pool that has mirrored protection, the system continues to run (unless both storage units of a mirrored have failed).
- If a disk unit fails in a disk pool that has device parity protection, the system continues running as long as no other disk unit in the same device parity set fails.

System limits for disk pool storage: During an IPL, the system determines how much auxiliary storage is configured on the system. The total amount is the sum of the capacity of the configured units and their mirrored pairs, if any. Disk units that are not configured are not included. The amount of disk storage is compared to the maximum that is supported for a particular model.

If more than the recommended amount of auxiliary storage is configured, a message (CPI1158) is sent to the system operator's message queue (QSYSOPR) and the QSYSMSG message queue (if it exists on the system). This message indicates that too much auxiliary storage exists on the system. This message is sent once during each IPL as long as the amount of auxiliary storage on the system is more than the maximum amount supported.

Independent disk pools

The terms **independent auxiliary storage pool** and **independent disk pool** are synonymous.

An independent disk pool is a collection of disk units that can be brought online or taken offline independent of the rest of the storage on a system, including the system disk pool, user disk pools, and other independent disk pools. Independent disk pools are useful both in single system and multiple system environments. For related information, see system disk pool and user disk pool.

In a single system environment, an independent disk pool can be taken offline independent of other disk pools because the data in the independent disk pool is self-contained, i.e. 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 brought online while the system is active (no IPL required). Using independent disk pools this way can be very useful, for example, if you have large amounts of data that are not needed for normal day-to-day business 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.

In a multi-system environment, the independent disk pool can be switched between systems. A **switchable independent disk pool** is a set of disk units that you can switch between systems so that each system can access the data. Only one system can access the data at a time. As in the single system environment, the independent disk pool can be switched because the independent disk pool is self-contained.

Switchable independent disk pools can help you do the following:

- Keep data available to an application even in the event of a single system outage (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.

For more information, see the Independent Disk Pool topic.

Contrast basic and independent disk pools

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

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

Device parity protection

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

If possible, you should 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 also keep your system operational while a disk unit is being repaired or replaced.

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

Before using device parity protection, you should be aware of the benefits that are associated with it, as well as the costs and limitations.

For additional information on device parity protection, review these topics:

- Planning for device parity protection
- How device parity protection affects performance
- Using both device parity protection and mirrored protection

For information on how to start using device parity protection in your business, see Backup and Recovery.



Planning for device parity protection

If your goal is to have a system with data loss protection and concurrent maintenance repair, plan to use a combination of mirrored protection and device parity protection. For each device parity protection set, the space that is used for parity information is equivalent to one disk unit. Beginning with V5R2 input/output adapters (IOAs), the minimum number of disk units in a parity set is 3; the maximum number of disk units in the parity set is 18. With IOAs developed prior to V5R2, the minimum number of disk units in a parity set is 4; the maximum number of disk units in the parity set is 10. At V5R2 you can optimize your parity sets for capacity, performance, or balanced if you have a V5R2 or later IOA. To learn more about how device parity protection is implemented and how it can be used in conjunction with mirrored protection, see the following topics.

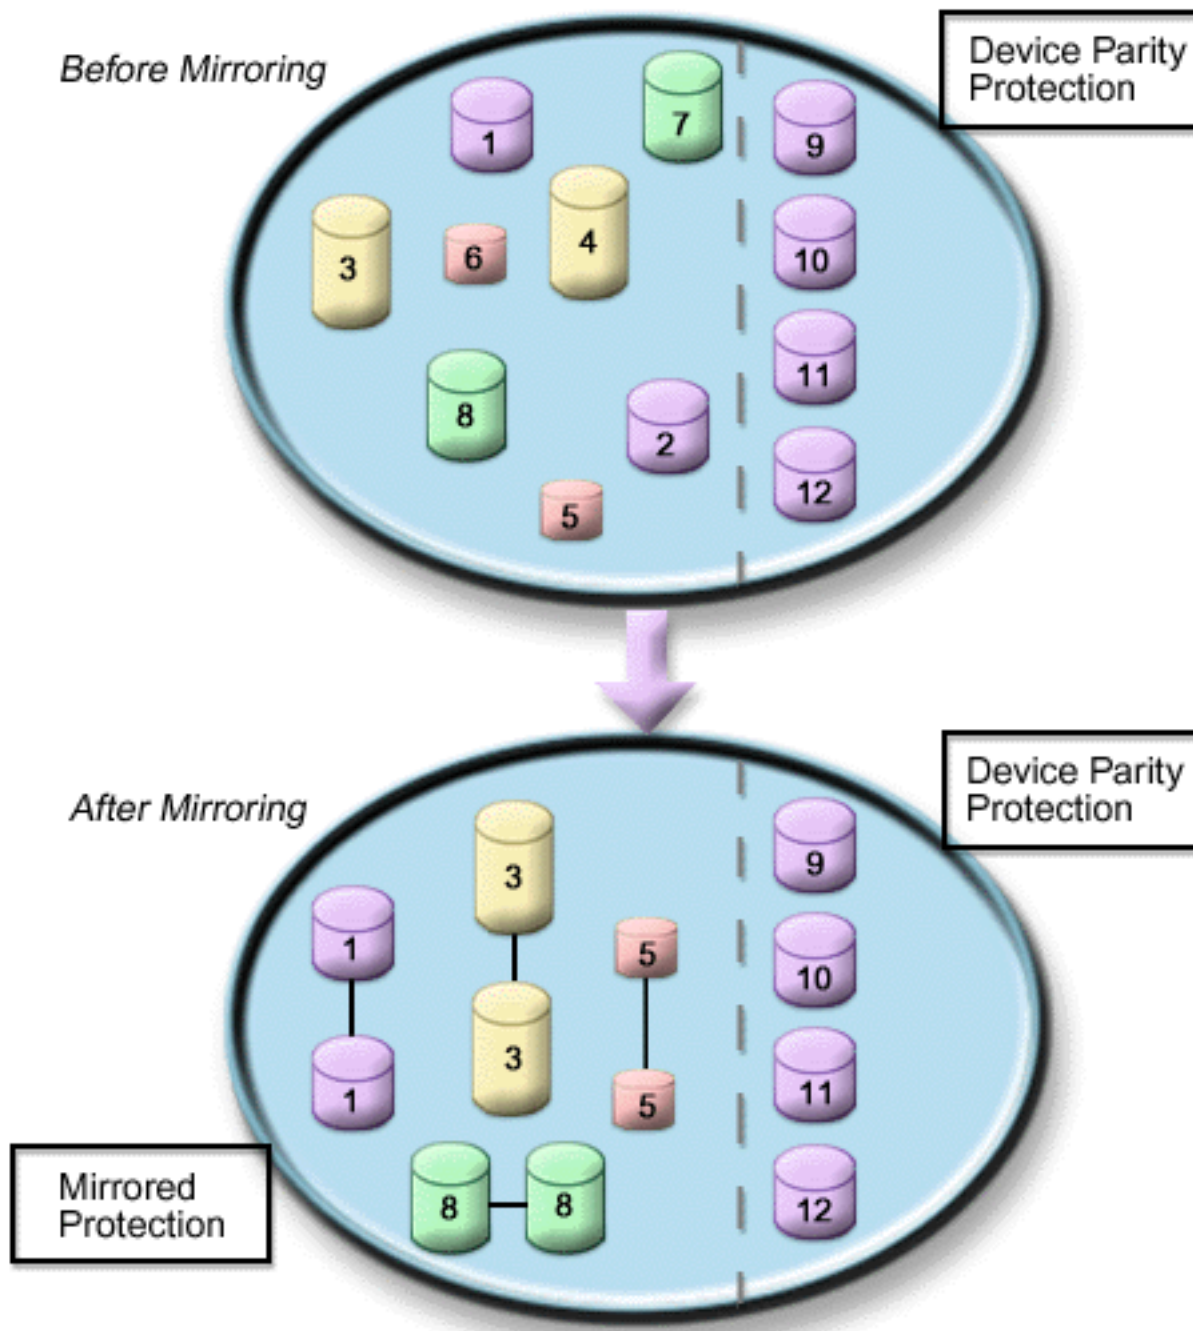
- How device parity protection works
- Examples of device parity and mirrored protection for disk pools

Examples of device parity and mirrored protection for disk pools

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

Here is an example of a system with a single disk pool (auxiliary storage pool) with both mirrored

protection and device parity protection.



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

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

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

Mirrored protection and device parity protection in all disk pools

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

How device parity protection works

When you start parity protection, the IOAs create device parity sets. Beginning with V5R2 input/output adapters (IOAs), the minimum number of disk units in a parity set is 3; the maximum number of disk units in the parity set is 18. With IOAs developed prior to V5R2, the minimum number of disk units in a parity set is 4; the maximum number of disk units in the parity set is 10. A parity set can only tolerate one disk failure. If more than one disk fails, you must restore the data from backup media. Because of the write penalty, 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.

In each parity set, the equivalent of one disk unit is devoted to storing parity data. The number of disk units that actually contain parity data varies according to the number of disk units in the parity set. The following table shows how many disk units in each parity set store parity data:

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

The input/output adapter determines how parity sets are formed. For V5R2 input/output adapters and beyond you do have the ability to choose how you want the parity set to be optimized. You can optimize according to *capacity*, *performance*, or a *balanced* version. 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 may also dedicate slightly more disk capacity to storing parity data.

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

When a device parity set grows you may want to consider redistributing the parity data. For example you may begin with 7 or fewer disk units, but expand to 8 or more by including more disk units. When this happens, you can improve the performance on the device parity set by stopping parity protection and starting it again. This redistributes the parity data across 8 disks rather than 4. In general, spreading the parity data across more disk units improves performance.

A write cache is included in the input/output adapter (IOA) for each parity set to improve performance of interactive write workloads. See Elements of device parity protection to see an example of a parity set with four disk units.

Beginning with V5R2, all input-output adapters (IOAs) are capable of device parity protection. If you have an earlier model adapter, check to see whether it is capable of device parity protection. For information about moving to a newer generation adapter, see *Migrating to a new input/output adapter*.

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

Elements of device parity protection: The following diagrams illustrate the elements of a parity set which contains four disk units. Each parity set begins with an Input/Output Processor (IOP) that is attached to an Input/Output Adapter (IOA) which contains the write cache. The IOA transmits read and write signals to the attached disk units. The first figure shows how parity is distributed with pre-V5R2 adapters. The second figure shows how parity is distributed with V5R2 adapters and beyond.

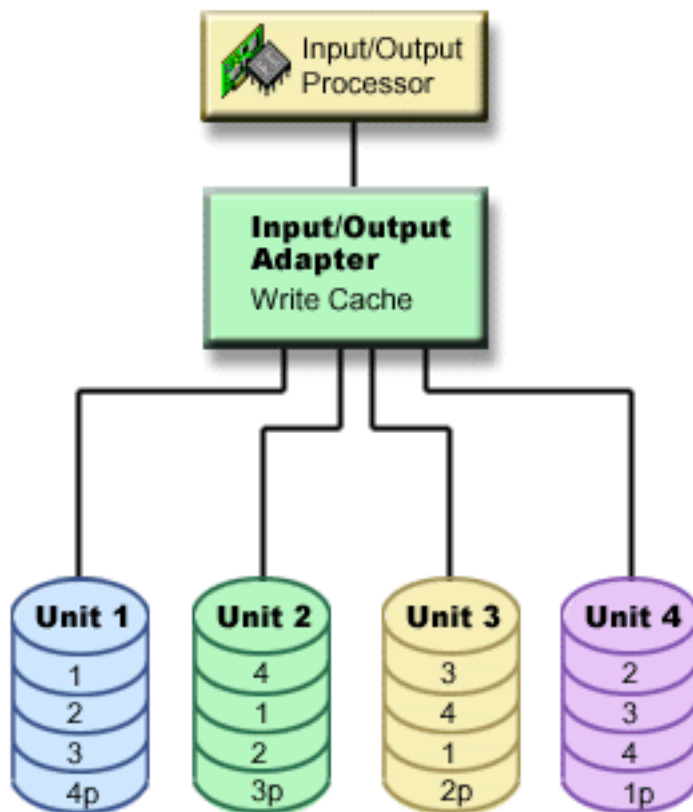


Figure 1. Example of how parity data is distributed with pre-V5R2 IOAs




Figure 2. Example of how parity data is distributed with IOAs at V5R2 and beyond

In the preceding examples, *p* indicates the sections of the disk which contain parity data. The first figure shows an example of a pre-V5R2 IOA in which the parity data is distributed in one large chunk on each disk unit that stores parity data. The second figure shows how V5R2 IOAs and beyond distribute parity data across the disk units in a small number of large chunks. Performance is improved by spreading the parity data throughout each of the disk units.

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

For a more in-depth overview, review the additional information about the write cache pictured above.

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

1. Data is committed to a nonvolatile battery-backed cache in the IOA.

2. A write completion message is sent from the server.

The following actions occur after the write completion message is sent.

1. A write operation is sent from the IOA cache to the disk unit
 - For data:
 - Reads the original data.
 - Calculates delta parity by comparing new and original data.
 - Writes the new data.
 - 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.
2. 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.

Migrating to a new input/output adapter

Before you begin the migration to the new input/output adapter (IOA), as with any configuration change, it is important to do a normal system power down. This will assure that all of your data is saved from the cache. When you migrate an existing parity set from a pre-V5R2 IOA to a V5R2 or later IOA, your disk units will not be protected by device parity protection while parity is being regenerated.

Note:

You cannot migrate the parity set back to the old generation of adapters once you have made the change to the new adapter. If you need to go back, you must stop device parity protection, associate the drives with the old adapter, and restart device parity protection.

Device parity protection—benefits

Here are the benefits of device parity protection:

- Lost data is automatically reconstructed by the disk controller after a disk failure.
- The system continues to run after a single disk failure.
- A failed disk unit can be replaced without stopping the system.
- Device parity protection reduces the number of objects that are damaged when a disk fails.
- Only 1 disk unit of capacity stores parity data in a parity set.

Device parity protection—costs and limitations

Here are the costs and limitations of device parity protection:

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

How device parity protection affects performance

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. A single disk unit failure can adversely affect the performance for both read and write operations.

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

- Disk unit failure in a device parity protection configuration
- Read operations on a failed disk unit
- Write operations on a failed disk unit
- Input-output operations during a rebuild process

Disk unit failure in a device parity protection configuration

If a disk unit fails, the subsystems with device parity protection are considered to be exposed until the synchronization process completes after replacing the failed disk unit. During the time the disk unit is considered exposed, additional I/O operations are required. If a second disk unit fails, you must restore your data from backup media.

Read operations on a failed disk unit

To get the data that was contained on a failed disk unit, device parity protection must read each disk unit in the device parity set that contains the failed disk unit. Because the read operations can be overlapped, the performance impact 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.

Write operations on a failed disk unit

There are some examples available that show what happens to write operations when a single disk unit fails in a device parity set with device parity protection. The figure below shows a failed unit under an IOA with device parity protection. Use the figure for the following examples:

- Example: Writing to a failed disk unit
- Example: Writing data to a disk unit when its corresponding parity data is on a failed disk unit




Figure 3. Device parity set with failed disk unit

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

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

1. The original data is lost on disk unit 3, sector 1, because of the failure.
2. The new parity data is calculated by reading disk unit 1, sector 1; and disk unit 2, sector 1.
3. New parity information is calculated.
4. New data cannot be written to sector 1 on disk unit 3, because of the failure.
5. New parity information is written to parity sector 1 on disk unit 4.

Write operations require multiple reads (N-2 reads, where N is the number of disk units) and only one write operation for the new parity information. Data from disk unit 3 will be rebuilt during synchronization after disk unit 3 is replaced.

Example: Writing data to a disk unit when its corresponding parity data is on a failed disk unit:

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

1. A failure is detected on the disk unit that contains the parity data, disk unit 3.
2. Calculating parity information is not required because it cannot write to parity sector 2 of disk unit 3. Therefore, there is no requirement to read the original data and the parity information.
3. Data is written to disk unit 4, sector 2.

A write operation requires only one write for the new data. Parity data for parity sector 2 on disk unit 3 will be rebuilt during synchronization after disk unit 3 is replaced.

Input-output operations during a rebuild process

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

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

Note: The rebuild process takes longer when read and write operations to a replaced disk unit are also occurring. Every read request or every write request interrupts the rebuild process to perform the necessary I/O operations.

Using both device parity protection and mirrored protection

Device parity protection is a hardware function. Disk pools and mirrored protection are software functions. When you add disk units and start device parity protection, the disk subsystem or IOP is not aware of any software configuration for the disk units. The software that supports disk protection is aware of which units have device parity protection.

These rules and considerations apply when mixing device parity protection with mirrored protection:

- Device parity protection is not implemented on disk pool boundaries.

- Mirrored protection is implemented on disk pool boundaries.
- You can start mirrored protection for a disk pool even if it currently has no units that are available for mirroring because they all have device parity protection. This ensures that the disk pool will always be fully protected, even if you add disks without device parity protection later.
- When a disk unit is added to the system configuration, it may or may not be device parity protected.
- For a fully-protected system, you should entirely protect every disk pool, either by device parity protection or by mirrored protection or both.
- Disk units that are protected by device parity protection can be added to a disk pool that has mirrored protection. The disk units that are protected by device parity protection do not participate in mirrored protection. Hardware protects them already.
- When you add a disk unit that is not protected by device parity protection to a disk pool that has mirrored protection, the new disk unit participates in mirrored protection. Disk units must be added to and removed from a mirrored disk pool in pairs with equal capacities.
- Before you start device parity protection for disk units that are configured (assigned to a disk pool), you must stop mirrored protection for the disk pool.
- Before you stop device parity protection, you must stop mirrored protection for any disk pools that contain affected disk units.
- When you stop mirrored protection, one disk unit from each mirrored pair becomes nonconfigured. You must add the nonconfigured units to the disk pool again before starting mirrored protection.

Mirrored protection

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

When you start mirrored protection or add disk units to a disk pool that has mirrored protection, the system creates mirrored pairs using disk units that have identical 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 different controllers, input/output adapter, input/output processors, buses, and towers.

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

Different levels of mirrored protection are possible, depending on what hardware is duplicated. You can duplicate:

- Disk units
- Input/output adapters
- Input/output processors
- Buses
- Towers
- High-speed links

The system remains available during the failure if a failing component and the hardware components that are attached to it are duplicated. For more technical details on your server storage and mirrored protection, see *How the system addresses storage and Mirrored protection—how it works*.

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

best choice; for others, remote DASD 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.

For more information on mirrored protection, see the following topics:

- Mirrored protection—benefits
- Mirrored protection—costs and limitations
- Planning for mirrored protection
- Remote DASD mirroring

For information on how to implement mirrored protection in your business, see Backup and Recovery. 

Mirrored protection—benefits

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

Even if your system is not a large one, mirrored protection can provide you valuable protection. A disk or disk-related hardware failure on an unprotected system leaves your system unusable for several hours. The actual time depends on the kind of failure, the amount of disk storage, your backup strategy, the speed of your tape unit, and the type and amount of processing the system performs. If you or your business 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

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, controller, or I/O processor fails, you need duplicate disk controllers and I/O processors. A model upgrade can be done to get nearly continuous operation and to prevent data loss when any of these failures occur, as well as the failure of a bus. If bus 1 fails, the system cannot continue to operate. Because bus failures are rare, and bus-level protection is not significantly greater than I/O processor-level protection, you may not find a model upgrade to be cost-effective for your protection needs.

Mirrored protection has a minimal effect on performance. If the buses, I/O processors, and controllers 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 or not 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 usually negligible. You should also consider other availability and recovery alternatives, such as device parity protection. Mirrored protection normally requires twice as many storage units. For concurrent maintenance and higher availability on systems with mirrored protection, other disk-related hardware may be required.

Limitations

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

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

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

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

Planning for mirrored protection

If you have a multi-bus 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, simply 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.

When considering mirrored protection, contact your IBM marketing representative to guide you through these planning steps:

1. Decide which disk pool(s) to protect.
2. Determine disk storage capacity requirements.
3. Determine the level of protection you want for each mirrored disk pool.
4. Determine what extra hardware you need for mirrored protection.
5. Determine what extra hardware you need for performance.
6. Order your hardware.
7. Plan the installation of your system and the configuration of new units.
8. Install the new hardware.

For more information on mirrored protection, see the following topics:

- Mirrored protection—benefits
- Mirrored protection—costs and limitations
- Mirrored protection—how it works

Mirrored protection—how it works

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

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

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

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

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

For details on storage on your server, see *How your server addresses storage*.

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

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

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

To determine which physical disk unit is being identified with each unit number, make note of the unit number assignment to ensure correct identification. If a printer is available, print the DST or SST display of your disk configuration. If you need to verify the unit number assignment, use the DST or SST Display Configuration Status display to show the serial numbers and addresses of each unit.

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

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

Remote mirroring: Remote mirroring support makes it possible to divide the disk units on your system into a group of local DASD and a group of remote DASD. The remote DASD are attached to one set of optical buses and the local DASD to another set of buses. The local and remote DASD can be physically separated from one another at different sites by extending the appropriate optical buses to the remote site, thus giving a higher level of protection in the event of a site disaster.


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

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

Concurrent maintenance support is a function of system unit hardware packaging. The entry system (9402) packaging does not support concurrent maintenance. 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 will not prevent the system from operating. A replacement of the head assembly and synchronization of the mirrored unit can occur while the system continues to run. The greater your level of protection, the more often concurrent maintenance can be performed.

On some models, the system restricts the level of protection for unit 1 and its mirrored unit to only

controller-level protection. See "Mirrored Protection - Configuration Rules" in Backup and Recovery.  for more information.

Under some conditions, diagnosis and repair can require active mirrored units to be suspended. You may prefer to power down 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.

Mirrored pair: Two storage units that contain the same data and are referred to by the system as one unit. A **mirrored unit** is a storage unit that is half of a mirrored pair.

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

A **storage unit** is the defined space within a disk unit that is addressed by the system.

A **unit** is the defined division of single-level storage. This space is the smallest disk location addressable by the user. A disk pool is one or more units that are identified by unique unit numbers. A unit in a non-mirrored disk pool is one storage unit. A unit in a mirrored disk pool is a mirrored pair, which is two storage units.

Certain create commands (CRTPF, CRTJRNRCV, etc) can create an object on a specified unit. In the non-mirrored environment this is a single storage unit. In the mirrored environment, the UNIT parameter value means a mirrored pair.

For details on storage on your server, see How the system addresses storage.

Tower: An enclosure that contains storage units and is separately addressable by the system.

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

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

I/O adapter: The input/output adapter (IOA), is attached to the input/output processor (IOP). The input/output adapter transfers information between the IOP and the disk units.

Controller: The disk controller attaches to the IOP and handles the information transfer between the IOP and the disk units. Some disk units have built-in controllers. Others have separate controllers.

Deciding which disk pools to protect

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

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

- If the system has a mixture of some disk pools with and some disk pools without mirrored protection, a disk unit failure in a disk pool without mirrored protection severely limits the operation of the entire system. Data can be lost in the disk pool in which the failure occurred. A long recovery may 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 may not be possible.

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

Determining the disk units that are needed

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

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

1. Plan how much data each disk pool will contain.
2. Plan a target percent of storage used for the disk pool (how full the disk pool will 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.

After planning for all disk pools is completed, plan for spare units, if desired.

Once you know all of this information, you can calculate your total storage needs.

Planning for storage capacity: For a new system, your IBM marketing representative or re-marketing representative can help you analyze your system storage requirements. For an existing system, the current amount of data in the disk pool that is being planned is a useful starting point. The DST or SST Display Disk Configuration Capacity option 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 1GB (GB equals 1 073 741 824 bytes) of actual data, it requires 2GB of storage for the mirrored copies of the data. If 50% full is planned for that disk pool, the disk pool needs 4GB of actual storage. If the planned percent of storage that is used is 66%, 3GB of actual storage are required. One gigabyte of real data (2GB of mirrored data) in a 5GB 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 unit of the capacity is available, that spare unit can be used to replace the failed 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 unit with the selected spare unit, then synchronizes the new unit with the remaining good unit of the mirrored pair. Mirrored protection for that pair is again active when synchronization completes (in usually less than an hour). However, it might take several hours from the time a service representative is called until the failed unit is repaired and synchronized, and mirrored protection is again active for that pair.

To make full use of spare units, you need at least one spare unit of each capacity that you have on your system. This provides a spare for any size of disk unit that may fail. A failed unit must be replaced by a spare of the same capacity.

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

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

Determining the level of protection that you want

The level of mirrored protection determines if 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. The four levels of mirrored protection, in order from lowest to highest, are as follows:

- Disk unit-level protection
- Input/output adapter-level protection
- Input/output processor-level protection
- Bus-level protection
- Tower-level protection
- Ring-level protection

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.

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.

For more information on the levels of protection, see [Levels of protection—more details](#).

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

The level of mirrored protection also determines if 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 requires resetting the I/O processor to which the failed disk unit is attached. Therefore, IOP-level protection is required. The higher the level of mirrored protection, the more often concurrent maintenance is possible.


The level of protection you get depends upon the hardware you duplicate. If you duplicate disk units, you will have disk unit-level protection. If you duplicate disk unit controllers as well, you have controller-level protection. If you duplicate input/output processors, you have IOP-level protection. If you duplicate buses, you have bus-level protection. Mirrored units will always have at least disk unit-level protection. Because most internal disk units have the controller packaged along with the disk unit, they will have at least controller-level protection.

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

Disk unit-level protection: Mirrored protection always provides disk unit-level protection because the storage units are duplicated. If your main concern is protection of data and not high availability, then disk unit-level protection may 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 with disk unit-level protection.

This 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 storage units make a mirrored pair. With disk unit-level protection, the system continues to operate after a disk unit failure. If the controller or I/O processor fails, the system cannot access data on either of the storage units of the mirrored pair, and the system is unusable.

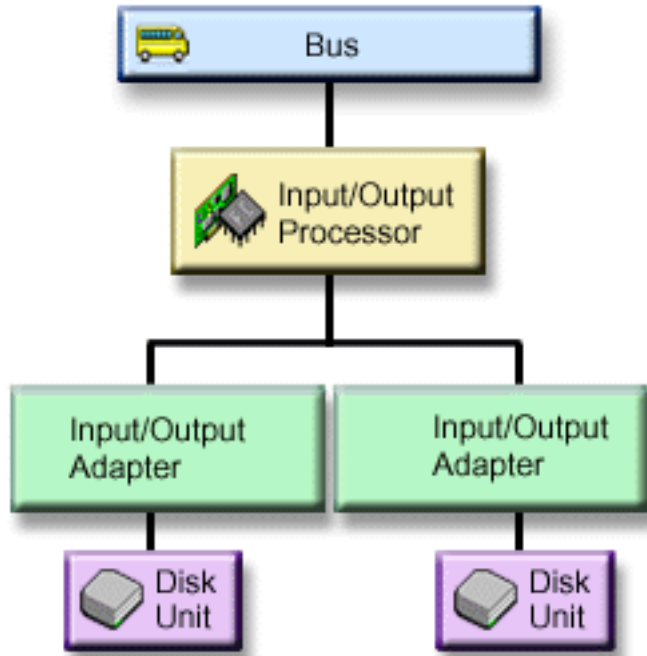


Input/output adapter-level protection: Determine if you want Input/output adapter (IOA)-level protection based on the following:

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

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

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




Input/output processor-level protection: Determine if you want IOP-level protection based on the following:

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

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

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




Bus-level protection: Bus-level protection may allow the system to run when a bus fails. However, bus-level protection is often not cost-effective because of the following:

- If bus 1 fails, the system is not usable.
- If a bus fails, disk I/O operations may continue, but so much other hardware is lost, such as work stations, printers and communication lines, that from a practical standpoint, the system is not usable.
- Bus failures are rare compared with other disk-related hardware failures.
- Concurrent maintenance is not possible for bus failures.

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


This figure shows the elements of bus-level protection: one tower that contains two buses attached to separate IOPs, IOAs, and disk units, respectively. The two storage units make a mirrored pair. With bus-level protection, the system can continue to operate after a bus failure. However, the system cannot continue to operate if bus 1 fails.



Tower-level protection: Tower-level protection may allow the system to run when a tower fails. However, tower-level protection is often not cost-effective because of the following:

- If a tower fails, disk I/O operations may continue, but so much other hardware is lost, such as work stations, printers and communication lines, that from a practical standpoint, the system is not usable.
- Tower failures are rare compared with other disk-related hardware failures.


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



Ring-level protection: Ring-level protection may allow the system to run when a high-speed link (HSL) fails. However, ring-level protection is often not cost-effective because of the following:

- If an HSL fails, disk I/O operations may continue, but so much other hardware is lost, such as work stations, printers and communication lines, that from a practical standpoint, the system is not usable.
- HSL failures are rare compared with other disk-related hardware failures.

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



Determining the hardware that is needed for mirroring

In order to communicate with the rest of the system, disk units are attached to controllers, which are attached to I/O processors, which are attached to buses. The number of each of these types of disk-related hardware 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 controllers. The same number of controllers should be under each I/O processor for that disk type. The I/O processors should be balanced among the available buses.

To plan what disk-related hardware is needed for your mirrored system, you must plan the total number and type of disk units (old and new), that will be 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 very large percentage of the disk units on the system achieve the desired level of protection.

When planning for additional disk-related hardware, you need to do the following:

1. Determine the minimum hardware that is needed for the planned disk units to function. Plan for one disk unit size at a time.

2. Plan the additional hardware needed to provide the desired level of protection for each disk unit type

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 these configuration limits and help you use them in your planning. For a listing of the configuration limits and rules, see Installation, upgrades, and migration.

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

Planning additional hardware to achieve the level of 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 a minimum of disk unit-level protection if they meet the requirements for starting mirrored protection.
- **Controller-level protection**
If the planned disk units do not require a separate controller, you will already have controller-level protection for as many units as possible and you do not need to do anything more. If your planned disk units do require a separate controller, add as many controllers as possible, keeping within the defined system limits. Then balance the disk units among them according to the standard system configuration rules.
- **Input/Output processor-level protection**
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 may 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 need to do nothing. If your system is configured according to standard configuration rules, the mirrored pairing function pairs up 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.
- **Tower-level protection**
If your system is configured with an equal number of equal capacity disk units between towers, the mirrored pairing function will pair up the disk units in different towers to provide tower-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 will pair up the disk units in different high-speed link (HSL) configurations to provide ring-level protection on as many disk units as possible.

Determine the extra hardware needed for performance

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

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

- **Processing unit requirements**
Mirrored protection causes a minor increase in central processing unit usage (approximately 1% to 2%).
- **Main storage requirements**
If you have mirrored protection, you need to increase the size of your machine pool. Mirrored protection requires storage in the machine pool for general purposes and for each mirrored pair. You should expect to increase your machine pool by approximately 12KB for each 1GB of mirrored disk storage (12KB for 1GB DASD, 24KB for 2GB DASD, etc.).

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

- **I/O processor requirements**

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

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

For more information on the effect that mirroring has on performance, see *Mirroring and performance*.

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

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

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

Order the new hardware

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


Planning your installation

You must work with your IBM marketing representative to plan for the installation of mirrored protection on your system. The marketing representative will help you determine whether your system is balanced and meets standard configuration rules, as defined in *Installation, upgrades, and migration*. The system must be configured according to the standard rules in order for the mirrored pairing function to pair up storage units to provide the best protection possible from the hardware that is available. Your marketing representative will also help you plan for the new units that are needed to add for each disk pool.

If you are planning to start mirrored protection on a new system, that system is already configured according to standard configuration rules. If you are using an older system, it may not follow the standard rules. However, wait until after attempting to start mirrored protection before reconfiguring any hardware.

For more information on how to plan your disk pools, see *Planning what disk pools to create*.

Planning what disk pools to create: Plan the user disk pools that will have mirrored protection and

determine what units to add to the disk pools. Backup and Recovery  contains information on how to assign disk units to add to disk pools.

In general, the units in a disk pool should be balanced across several I/O processors, rather than all being attached to the same I/O processor. This provides better protection and performance.

Installing new hardware

When the hardware arrives, your service representative will install the hardware. After the hardware is installed, see Add a disk unit or disk pool for information on how to add new units and start mirrored protection.

Remote DASD mirroring support

Standard DASD mirroring support requires that both disk units of the load source mirrored pair (unit 1) are attached to the Multi-function I/O Processor (MFIOP). 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 I/O Processor (IOP), the best mirroring protection possible for the load source mirrored pair is controller-level protection. To provide a higher level of protection for your system, you can use remote load source mirroring and remote DASD mirroring.

Remote DASD mirroring support, when combined with remote load source mirroring, mirrors the DASD on local optical buses with the DASD on optical buses that terminate 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 DASD at the local site. If the local DASD and system unit are lost, a new system unit can be attached to the set of DASD at the remote site, and system processing can be resumed.

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

Remote mirroring support makes it possible to divide the disk units on your system into a group of local DASD and a group of remote DASD. The remote DASD are attached to one set of optical buses and the local DASD to another set of buses. The local and remote DASD 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 an optical bus may be extended.

For more information on remote DASD mirroring, see the following topics:

- Remote DASD mirroring—advantages

- Remote DASD mirroring—disadvantages

- Comparison of standard and remote mirroring

If you decide that remote DASD mirroring is right for your system, you will need to prepare your system and then start site-to-site mirroring.

Remote load source mirroring

Remote load source mirroring support allows the two disk units of the load source to be on different IOPs or system buses, which provides IOP- or bus-level mirrored protection for the load source. However, in such a configuration the system can only IPL from or perform a main storage dump to the load source attached to the MFIOP. If the load source on the MFIOP fails, the system can continue to run on the other disk unit of the load source mirrored pair, but the system will not be able to IPL or perform a main storage dump until the load source attached to the MFIOP is repaired and usable.

For more information on remote load source mirroring, see the following topics:

- Enabling remote load source mirroring
- Disabling remote load source mirroring
- Using remote load source mirroring with local DASD

Enabling remote load source mirroring: To use remote load source mirroring support, remote load source mirroring must first be enabled. Then mirrored protection must be started 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 will not be changed.

Remote load source mirroring support can be enabled in either the DST or the SST environment in iSeries Navigator or the character-based interface. If you attempt to enable remote load source mirroring and it is currently enabled, the system will display a message that remote load source mirroring is already enabled. There are no other errors or warnings for enabling remote load source mirroring support.

To enable remote load source mirroring, do the following:

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

Disabling remote load source mirroring: If you wish to disable remote load source mirroring support, you must either:

- Stop mirrored protection and then disable remote load source mirroring support.

or

- Move the remote load source to the MFIOP and then disable remote load source mirroring support.

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

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

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

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

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

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

Remote DASD mirroring—advantages

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

Remote DASD mirroring—disadvantages

- A system that uses Remote DASD Mirroring is only able to IPL from one DASD of the load source mirrored pair. If that DASD fails and cannot be repaired concurrently, the system cannot IPL until the failed load source is fixed and the remote load source recovery procedure is performed.
- When Remote DASD Mirroring is active on a system and the one load source the system can use to IPL 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 or continuously-powered mainstore (CPM) to reduce recovery time after a system crash. It also means that the main storage dump is not available to diagnose the problem that cause the system to end abnormally.

Comparison of DASD management with standard mirroring and remote mirroring

For the most part, the way you manage DASD with remote mirroring is the same as how you manage DASD with standard mirroring. The differences are in how you add disk units and how you restore mirrored protection after a recovery.

Adding disk units: Unprotected disk units must be added in pairs, as with general mirroring. To achieve remote protection of all added units, half of the new units of each capacity of DASD should be in the remote group and half in the local group. Single device-parity protected units may be added to disk pools using remote mirroring. However, the disk pool will not be protected against a site disaster.

Restoring remote mirrored protection after a recovery: To restore mirrored protection following the recovery procedures, you will need to perform the following steps:

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

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



Preparing your system for remote mirroring

When you start remote system mirroring, the local DASD is mirrored to the remote DASD. 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. In order to provide protection against a site disaster, all DASD in all disk pools of the system must be mirrored in local-remote pairs. Follow these steps to prepare your system for remote mirroring:

1. Plan which optical buses will drive DASD 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 DASD are equal.
 - It is functionally necessary that both the local and remote sites have the same number of each capacity of DASD in each disk pool.
2. Plan the distribution of DASD, move DASD if necessary, and verify that half of each capacity of DASD in each disk pool are attached to the local and remote set of buses.
3. Indicate to the system which buses drive remote DASD and which buses drive local DASD. To do this, you must first find which buses drive remote DASD and record those bus numbers. Then, you have to change the system resource IDs of the remote buses so that they begin with an *R*.
For example, if you determine that BUS11 drives remote DASD, then you would change the system resource ID of that bus to *RBUS11*

Finding remote buses: If the buses are not labelled, you may have to trace the buses by hand to see which go to remote locations. You can also use the Hardware Service Manager to determine which buses go to which expansion units.

To use the Hardware Service Manager to find the buses that drive remote DASD, perform these steps:

1. From the DST Main Menu, select option 7 (Start a service tool).
2. From the Start a Service Tool display, select option 4 (Hardware service manager).
3. From the Hardware Service Manager menu, select option 2, Logical hardware resources.
4. From the Logical hardware resources menu, selection option 1, System bus resources.
5. On the Logical hardware resource on system bus screen, enter option 8 before each bus to display associated packaging resources.
6. The Packaging resources that are associated with a logical resource screen displays the frame ID and resource name of the expansion unit that is associated with the bus. 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.

Record the remote or local location of the bus. Then repeat this procedure for all buses on the system.

Changing remote bus resource names: Once you know which buses drive remote DASD, 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 option 7 (Start a service tool).
2. From the Start a Service Tool display, select option 4 (Hardware service manager).
3. From the Hardware Service Manager menu, select option 2, Logical hardware resources.
4. From the Logical hardware resources menu, selection option 1, System bus resources.
5. On the Logical hardware resource on system bus screen, select with the number 2 the bus whose name you wish to change. This will display the Change logical hardware resource detail screen.
6. On the Change logical hardware resource detail screen, on the line labelled 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.

Repeat this procedure for each remote bus on the system.

Starting site-to-site mirroring

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

1. Enable remote load source mirroring. This enables you to have a load source as part of the remote group of DASD.
2. Start mirroring using the normal start mirroring function.

When mirroring is started the system will use the resource name to recognize the remote buses and will attempt to pair the DASD on the remote buses with the DASD on the local buses. Because remote load source mirroring is enabled, the system will also pair the load source with a remote DASD. Normal mirroring restrictions that concern total disk pool capacity, an even number of disk units of each capacity, and such things, apply.

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

Chapter 2. Choosing your level of protection

There are several different ways to configure your system in order take advantage of the disk protection features. Before selecting the disk protection options that you want to use, compare the extent of protection that each one provides.

- Comparison of disk protection options
- Full mirrored protection versus partial mirrored protection

After comparing the disk protection options, select one of these methods of using the options:

- Full Protection — Single disk pool
- Full Protection — Multiple disk pools
- Partial Protection — Multiple disk pools
- “Assigning disk units to disk pools” on page 46

Comparison of disk protection options

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

- With both device parity protection and mirrored protection, the system continues to run after a single disk failure. With mirrored protection, the system may continue to run after the failure of a disk-related component, such as a controller or an IOP.
- If a second disk failure occurs such that the system has two failed disks, the system is more likely to continue to run with mirrored protection than with device parity protection. With device parity protection, the probability of the system failing on the second disk failure can be expressed as P out of n . This is where P is the total number of disks on the system and n is the number of disks in the device parity set that had the first disk failure. With mirrored protection, the probability of the system failing on the second disk failure is 1 out of n .
- Device parity protection requires one disk of existing disk capacity per parity set for storage of parity information. A system with mirrored protection requires twice as much disk capacity as the same system without mirrored protection because all information is stored twice. Mirrored protection may also require more buses, IOPs, and disk controllers, depending on the level of protection that you want. Therefore, mirrored protection is typically a more expensive solution than device parity protection.
- Usually, neither device parity protection nor mirrored protection has a noticeable effect on system performance. In some cases, mirrored protection actually improves system performance.
- The time required to restore data to disk units protected by device parity protection is longer than the restore time to the same disk devices without device parity protection activated, because the parity data must be calculated and written.

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

What type of availability is needed?	Device Parity Protection	Mirrored Protection	Basic disk pools	Independent disk pool
Protect from data loss due to disk-related hardware failure	Yes	Yes	See note ²	See note ²
Maintain availability	Yes	Yes	No	Yes ⁴
Help with disk unit recovery	Yes	Yes	Yes ²	Yes ²
Maintain availability when input-output adapter (IOA) fails	No	Yes ¹	No	No
Maintain availability when disk I/O processor fails	No	Yes ¹	No	No
Maintain availability when system bus fails	No	Yes ¹	No	No
Site disaster protection	No	Yes ³	No	No

What type of availability is needed?	Device Parity Protection	Mirrored Protection	Basic disk pools	Independent disk pool
Ability to switch data between systems	No	No	No	Yes
Notes:				
1	Depends on hardware used, configuration, and level of mirrored protection.			
2	Configuring disk pools can limit the loss of data and the recovery to a single disk pool.			
3	For site disaster protection, remote mirroring is required.			
4	In a clustered environment an independent disk pool can help maintain availability .			

See also:

- “How your system manages auxiliary storage” on page 43
- “How disks are configured” on page 43

Full mirrored protection versus partial mirrored protection

Full mirrored protection and partial mirrored protection do not provide the same availability results. These two implementations of mirrored protection are quite different. The scenarios of a disk unit on the iSeries server for each of these mirroring methods requires different user responses.

It does not matter if you are using just the system disk pool (disk pool 1) or multiple user disk pools (2 through 255), full mirrored protection protects all disk units in the iSeries server. Partial mirrored protection protects only a portion of the disk units designated by one or more disk pool s. However, not all storage units in the disk configuration are protected. Therefore, the planning of the disk unit placement and what disk pools are selected for mirrored protection becomes more difficult.

Besides the planning of disk pools, the significant difference between the two mirrored protection methods regards availability. With full mirrored protection, you maximize the availability of the iSeries server when a disk subsystem failure occurs. With this method of mirrored protection, it does not matter which disk pool has the failure. With partial mirrored protection, the system continues to run while reporting the failed storage unit to the system operator (QSYSOPR) message queue. However, if the disk failure occurs in a disk pool that does not have mirrored protection, SRC A6xx 0266 is sent when that disk pool is accessed by any job on the system. Because the storage units in the disk pool do not have mirrored units, the storage management directory becomes unusable and all input and output operations to the disk pool are suspended.

The disk attention SRC does not mean that the system has ended. All input and output operations are queued to allow the service representative to investigate the cause of the disk failure. If the problem is not with the disk media, the failing cards are replaced, the failed disk unit is powered on, and the system continues from the point that the equipment error occurred. All queued input and output operations resume. However, if a disk media failure occurs, the service representative performs a main storage dump to minimize the time for the next IPL to OS/400®, and allows the system to end processing.

With full mirrored protection, the operation of the system is not interrupted while diagnostics and most repairs to resolve the disk subsystem failure problem are taking place. With I/O processor-level protection, the maximum concurrent maintenance is possible, depending on the error. In any case, the user has complete control over the shutdown of the system should a power-down be required to service the disk problem; the system does not end abnormally.

Although critical data is protected with partial mirrored protection, and a restore operation is not required for the data in the protected disk pool, you do not have the maximum availability that is provided by full mirrored protection because of the exposure of the unprotected disk pool . If your availability requirements

state that your system must be in operation within minutes following a failure or remain active during your business hours, partial mirrored protection is not an option in most cases.

How your system manages auxiliary storage

To understand the availability option on your server, you need a basic understanding of how your iSeries server manages disk storage. On your server, main memory is called **main storage**. Disk storage is called **auxiliary storage**. You may also hear disk storage referred to as **DASD (direct access storage device)**.

Many other computer systems require you to take responsibility for how information is stored on disks. When you create a new file, you must tell the system where to put the file and how big to make it. You must balance files across different disk units to provide good system performance. If you discover later that a file needs to be larger, you need to copy it to a location on disk that has enough space for the new, larger file. You may need to move files between disk units to maintain system performance.


The iSeries server is different in that it takes responsibility for managing the information in auxiliary storage. When you create a file, you estimate how many records it should have. The system places the file in the best location for good performance. In fact, it may spread the data in the file across multiple disk units. When you add more records to the file, the system assigns additional space on one or more disk units.

Single-level storage is the unique architecture of the iSeries server that allows main storage and auxiliary storage to work together accurately and efficiently. With single-level storage, programs and system users ask for data by name, not by where the data is physically located. The system keeps track of where the most current copy of any piece of information is located in main storage or in auxiliary storage.

How disks are configured

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 picture 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 buses.

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

Input-output adapter (IOA): The IOA attaches to the IOP and handles the information transfer between the IOP and the disk units.

Disk unit: Disk units are the actual devices that contain the storage units. You order hardware at the disk-unit level. Each disk unit has a unique serial number. Additional information is available on how the your server addresses individual storage units.

How the System Addresses Individual Storage Units

To move data to and from auxiliary storage, the system needs a way to identify a single storage unit. Every hardware component (bus, I/O processor, controller, and storage unit) has a unique address.

The address of a storage unit consists of system bus, system board, system card, I/O bus, controller, and device numbers.

Disk Unit Hardware Resource Information Details

```
Type.....: 6603
Model.....: 030
Serial number....: 00-0109928
Resource name....: DD002

SPD bus
System bus.....: 1
System board....: 0
System card.....: 1

Storage
I/O bus.....: 0
Controller.....: 1
Device.....: 0
```

Full protection — single disk pool

A simplest way to manage and protect your auxiliary storage is to do the following:

- Assign all disk units to a single disk pool (the system disk pool).
- Use device parity protection for all disk units that have the hardware capability.
- Use mirrored protection for the remaining disk units on the system.

With this method, your system continues to run if a single disk unit fails. When the failed disk is replaced, the system reconstructs the information so that no data is lost. The system may also continue to run when a disk-related hardware component fails. Whether your system continues to run depends on your configuration. For example, the system will continue to run if an IOP fails and all of the attached disk units have mirrored pairs that are attached to a different IOP.

When you use a combination of mirrored protection and device parity protection to fully protect your system, you increase your disk capacity requirements. Device parity protection requires up to 25% of the space on your disk units to store parity information. Mirrored protection doubles the disk requirement for all disks that do not have the capability for device parity protection.

Full protection — multiple disk pools

You may want to divide your disk units into several disk pools (auxiliary storage pools). Sometimes, your overall system performance may improve by having user disk pools. For example, you can isolate journal receivers in a basic or secondary disk pool . Or, you can place history files or documents that seldom change in a disk pool that has lower performance disk units.

You can fully protect a system with multiple disk pools by doing the following:

- Use device parity protection for all disk units that have the hardware capability.
- Set up mirrored protection for every disk pool on the system. You can set up mirrored protection even for a disk pool that has only disk units with device parity protection. That way, if you add units that do not have device parity protection in the future, those units are automatically mirrored.

Note: For mirrored protection you must add new units in pairs of units with equal capacity.

Before configuring this level of protection, be sure that you know how to assign disk units to disk pools.

Partial protection — multiple disk pools

Sometimes, full protection (using a combination of device parity protection and mirrored protection) may be too costly. If this happens, you need to develop a strategy to protect the critical information on your system. Your objectives should be to minimize the loss of data and to reduce the amount of time that critical applications are not available. Your strategy will probably involve dividing your system into basic or independent disk pools and protecting only certain disk pools. Note, however, that if the system is not fully protected and an unprotected disk unit fails, serious problems can occur. The entire system can become unusable, end abnormally, require a long recovery, and data in the disk pool that contains the failed unit will have to be restored.

Before configuring this level of protection, be sure that you know how to assign disk units to disk pools.

The following list has suggestions for developing your strategy:

- If you protect the system disk pool with a combination of mirrored protection and device parity protection, you can reduce or eliminate recovery time. The system disk pool, and particularly the load source unit, contain information that is critical to keeping your system operational. For example, the system disk pool has security information, configuration information, and addresses for all the libraries on the system.
- Consider how you can recover object information. If you have on-line applications and your objects change constantly, consider using journaling and placing journal receivers in a protected user disk pool.
- Think about what information does not need protection, probably because it changes infrequently. For example, history files may need to be on-line for reference, but the data in the history files may not change except at the end of the month. You might place those files in a separate disk pool that does not have any disk protection. If a failure occurs, the system will become unusable, but the files can be restored without any loss of data. The same may be true for documents.
- Consider other information that may not need disk protection. For example, your application programs may be in a separate library from the application data. Probably, the programs change infrequently. The program libraries might be placed in a basic disk pool that is not protected. If a failure occurs, the system will become unusable, but the programs can be restored.


Two simple guidelines can summarize the previous list:

1. To reduce recovery time, protect the system disk pool.
2. To reduce loss of data, make conscious decisions about which libraries and objects to protect.

Assigning disk units to disk pools

If you decide that you want more than one disk pool, also called an auxiliary storage pool (ASP) in the character-based interface, you need to determine the following for each disk pool:

- How much storage you need.
- What disk protection, if any, to use.
- Which disk units to assign.
- Which objects to place in the disk pool.

The Workstation Customization Programming  book provides information to help you with these decisions.

When you work with disk configuration, you may find it helpful to begin by printing your current system configuration. You can obtain this information from the Hardware Service Manager in system service tools (SST) or from the Disk Units folder of iSeries Navigator.



Printed in U.S.A.