



@server

iSeries

Independent disk pools

Version 5 Release 3





@server

iSeries

Independent disk pools

Version 5 Release 3

Note

Before using this information and the product it supports, read the information in "Notices," on page 63 and the manual *IBM eServer Safety Information*.

Third Edition (August 2005)

This edition applies to version 5, release 3, modification 0 of IBM Operating System/400 (product number 5722-SS1) and to all subsequent releases and modifications until otherwise indicated in new editions. This version does not run on all reduced instruction set computer (RISC) models nor does it run on CISC models.

© Copyright International Business Machines Corporation 2002, 2005. All rights reserved.

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

Contents

Using independent disk pools 1

Independent disk pools	1
Independent disk pool terminology	1
Benefits of independent disk pools	5
How independent disk pools work	6
Recommended structure for independent disk pools	12
Switchable and stand-alone independent disk pools	13
Disk pool groups	14
Geographic mirroring	15
Plan for independent disk pools	21
iSeries Navigator requirements for disk management	21
Hardware requirements	22
Physical planning requirements	23
Software and licensing requirements	24
Communications requirements	25
Cluster requirements	26
Application considerations for independent disk pools	26
Configure independent disk pools	27
Create a dedicated independent disk pool	27
Create a switchable independent disk pool	28
Create a new disk pool group	32
Convert UDFS disk pools	32
Configure geographic mirroring with dedicated independent disk pools	33

Configure geographic mirroring with switchable independent disk pools	37
Manage independent disk pools	41
Backup and recovery of independent disk pools	41
Make a disk pool available	42
Make a disk pool unavailable	43
Make your hardware switchable	44
Switch access to backup server	46
Synchronize user profile name, UID, and GID	46
Change the server takeover IP address	46
Manage independent disk pools with geographic mirroring	47
Examples: Independent disk pool configurations	50
Dedicated independent disk pools	51
Scenario: Consolidate servers using switchable independent disk pools	51
Examples: Switchable independent disk pools	52
Example: Independent disk pools with geographic mirroring	57
Code disclaimer information	57
Frequently asked questions	58

Appendix. Notices 63

Programming Interface Information	65
Trademarks	65
Terms and conditions for downloading and printing information	65

Using independent disk pools

If you are interested solely in independent disk pools, use this information to help you plan, configure, and manage independent disk pools.

This topic will provide you with the information you need to use independent disk pools, from a conceptual explanation to planning, configuring, and managing independent disk pools on your servers.

“Plan for independent disk pools” on page 21

“Configure independent disk pools” on page 27

“Manage independent disk pools” on page 41

Note: Read the “Code disclaimer information” on page 57 for important legal information.

Independent disk pools

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

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

Independent disk pool terminology

As you work with independent disk pools, you need to become familiar with the following terms. For more terms and concepts, you can access the Information Center glossary.

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

active state

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

asynchronous

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

cluster

A collection of complete systems that work together to provide a single, unified computing capability. An iSeries cluster is made up of only iSeries servers and is required when implementing switchable independent disk pools.

cluster resource group (CRG)

A collection of related cluster resources that defines actions to be taken during a switchover or

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

cross-site mirroring (XSM)

A feature of OS/400® High Available Switchable Resources (Option 41) that provides geographic mirroring and the services to switch over or automatically cause a failover to a mirror copy, potentially at another location, in the event of an outage at the primary location.

data port services

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

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

detached mirror copy

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

device description

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

device domain

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

disk pool

An auxiliary storage pool that contains only disk units.

disk pool group

Made up of a primary disk pool and zero or more secondary disk pools, each of which are independent in regard to data storage, but combine to act as one entity. See "Disk pool groups" on page 14.

disk unit

A physical enclosure containing one or more disk drives.

expansion unit

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

failover

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

geographic mirroring

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

HSL (high-speed link) loop

The system-to-expansion unit connectivity technology that is required to use switchable

independent disk pools residing on an expansion unit (frame/unit). The servers and expansion unit in a cluster using resilient devices on an external expansion unit must be on an HSL loop connected with HSL cables.

independent disk pool

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

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

library namespace

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

mirror copy

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

mirror copy state

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

mirror copy data state

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

primary disk pool

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

production copy

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

reattach

In geographic mirroring, to reassociate the detached mirror copy with its production copy after user operations on the detached mirror copy are completed. When the detached mirror copy is reattached, it is automatically synchronized to match the production copy again. All data on the detached mirror copy is cleared before it is reattached to the production copy.

resume

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

resume pending state

In geographic mirroring, pertaining to the configuration state of a mirror copy that indicates that

geographic mirroring requires synchronization but that the disk pool is currently unavailable. When the disk pool is made available, the mirror copy will be synchronized with the current data from the production copy.

resuming state

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

secondary disk pool

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

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

site primary node

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

source system

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

suspend

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

suspended state

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

switchable entity

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

switchover

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

synchronization

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

synchronous

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

SYSBAS

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

target system

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

UDFS disk pool

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

unusable

In geographic mirroring, pertaining to the mirror copy data state that indicates that the mirror copy contains incoherent data. This occurs (a) during synchronization because synchronization does not preserve the order of writes and (b) when the system performs geographic mirroring in asynchronous mode.

usable In geographic mirroring, pertaining to the mirror copy data state that indicates that the correct order of updates to the mirror copy from the production copy is being preserved, but the mirror copy may be outdated. The usable mirror copy data state occurs (a) when the system performs geographic mirroring in synchronous mode, (b) after successfully suspending geographic mirroring, and (c) when mirror copy is successfully detached.

vary off

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

vary on

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

Benefits of independent disk pools

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

Multi-system clustered environment

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

Switchable independent disk pools can help you do the following:

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

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

Single-system environment

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

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

- Isolate low-use data with the ability to bring the independent disk pool online only when it is needed.
- Reduce system start time.
- Manage save/restore by independent disk pool.
- Reclaim storage by independent disk pool.
- Divide data between multiple databases.
- Isolate data associated with specific applications or associated with specific groups of users.
- Consolidate data on small systems to independent disk pools on a larger system; for example, in the case of multiple branch offices.
- Perform application maintenance that does not affect the entire system.

How independent disk pools work

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

Independent disk pools are available only when you choose to make them available; they are not made available when you restart your server, unless you include code (“Example: Make independent disk pool available at startup” on page 12) to make them available. When you select to make a disk pool available, the disk pool goes through a process similar to that of restarting the server. While this processing takes place, the disk pool is in an Active state.

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

At this time database cross-reference files will also be updated. The system libraries for the independent disk pool, QSYSnnnnn and QSYS2nnnnn, contain metadata not only for the independent disk pool, but

also for the system disk pool. When the disk pool is made available, database cross-referencing clears the information related to SYSBAS and updates it with current information. The number and complexity of database file objects and SQL packages, procedures, and functions that need to be updated will play a role in the time it takes to make the disk pool available.

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

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

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

IBM® grants you a nonexclusive copyright license to use all programming code examples from which you can generate similar function tailored to your own specific needs.

SUBJECT TO ANY STATUTORY WARRANTIES WHICH CANNOT BE EXCLUDED, IBM, ITS PROGRAM DEVELOPERS AND SUPPLIERS MAKE NO WARRANTIES OR CONDITIONS EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OR CONDITIONS OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT, REGARDING THE PROGRAM OR TECHNICAL SUPPORT, IF ANY.

UNDER NO CIRCUMSTANCES IS IBM, ITS PROGRAM DEVELOPERS OR SUPPLIERS LIABLE FOR ANY OF THE FOLLOWING, EVEN IF INFORMED OF THEIR POSSIBILITY:

1. LOSS OF, OR DAMAGE TO, DATA;
2. SPECIAL, INCIDENTAL, OR INDIRECT DAMAGES, OR FOR ANY ECONOMIC CONSEQUENTIAL DAMAGES; OR
3. LOST PROFITS, BUSINESS, REVENUE, GOODWILL, OR ANTICIPATED SAVINGS.

SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO SOME OR ALL OF THE ABOVE LIMITATIONS OR EXCLUSIONS MAY NOT APPLY TO YOU.

Supported and unsupported OS/400 object types

Objects not supported

The following OS/400 objects are not supported for use in independent disk pools:

*AUTHLR	*DEV D	*JOBQ	*PRDDFN
*AUTL	*DOC	*JOBSCD	*PRDLOD
*CFGL	*DSTMF	*LIND	*RCT
*CNNL	*EDTD	*MODD	*SOCKET
*COSD	*EXITRG	*M36	*SSND
*CRG	*FLR	*M36CFG	*S36
*CSPMAP	*IGCSRT	*NTBD	*USRPRF

*CSPTBL	*IGCTBL	*NWID
*CTLD	*IMGCLG	*NWSD
*DDIR	*IPXD	*PRDAVL

Notes

1. Spooled files are not supported. There is no object type associated with spooled files.
2. *DSTMF is the object type returned for stream files that are being accessed through the QNTC file system from a remote server. So you shouldn't see *DSTMF ever when accessing the IASP directories from the local system.

Supported object types

The following OS/400 objects are supported for use in independent disk pools:

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

Restrictions for supported object types

*ALRTBL

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

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

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

*JOB

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

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

*MSGQ

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

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


*SBSD

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

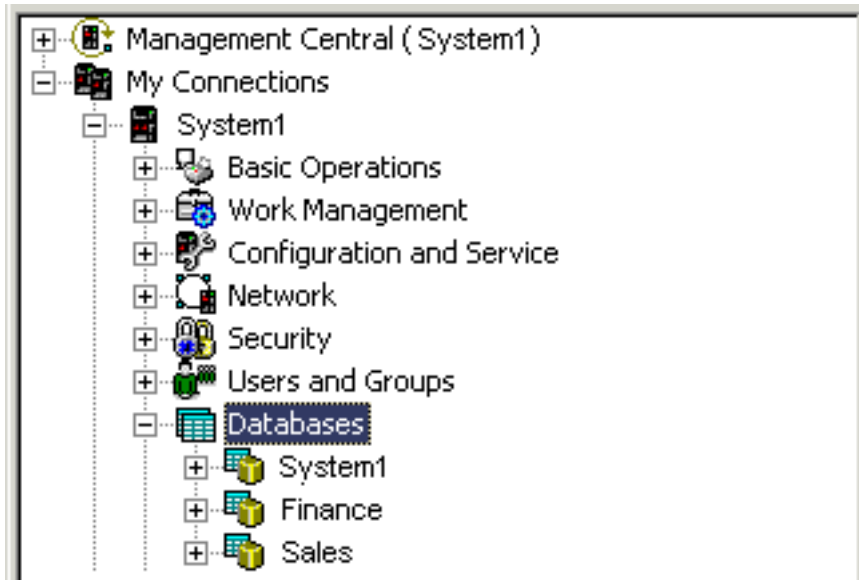
Independent disk pools with distinct databases

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

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



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



See “Object identification” on page 11 for details on identifying objects when independent disk pools exist on your server.

Multiple system libraries

In general, all system libraries continue to exist in the system disk pool. However, to support better isolation and recovery of the independent disk pool group containing system libraries, the following instances of system libraries are also created in the primary disk pool:

1. **QSYSnnnnn**: This contains the database cross reference information for the database represented by the disk pool group. Normally only internal system code creates objects into this library.
2. **QSYS2nnnnn**: This contains the SQL catalogues for the database represented by the disk pool group. Normally only internal system code creates objects into this library.
3. **QRCLnnnnn**: Any recovery object associated with objects within the disk pool group is stored in this library for the primary disk pool for the group. These objects may be needed for recovery when the disk pool group is varied on. The system disk pool equivalent of this library is QRECOVERY .
4. **QRPLnnnnn**: When the reclaim instance is run on the disk pool group, any resulting information normally stored in QRCL is now stored in the QRCL of the primary disk pool for the group. Normally only functions called during reclaim storage processing create objects into this library instance. Also, when reclaim storage recovers the addressability of lost objects, these objects can be inserted into the QRCLnnnnn library. These are user objects that originally existed in another library.
5. **QRPLnnnnn**: Whenever an object contained within the disk pool group is replaced while it is in use, the in-use object is renamed and moved to the QRPLnnnnn library in the primary disk pool for the group. The new object is inserted into the specified library. The system disk pool equivalent of this library is QRPLOBJ. QRPLnnnnn is cleared at vary on.

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

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

The setting of this attribute is as follows:

Library attribute settings

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

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

Object identification

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

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

Control language (CL) commands

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

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

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

Printing considerations

If you choose to store external resources for spooled files, not the spooled files themselves, in a disk pool group, you must be aware of the printing implications. You can store formatting objects such as *FNTRSC, *FORMDF, *OVL, *PAGDFN, and *PAGSEG in a disk pool group. See Spooled file considerations for independent disk pools. In order for the printer writer job to access these objects, you must set the disk pool group so that it exists in the library namespace.

Follow these steps to print the spooled file when external resources are stored in a disk pool group:

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

Switching independent disk pool between different releases

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

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

Example: Make independent disk pool available at startup

If you need your independent disk pools to be made available in most cases when your server is restarted, you might want to consider including the following control language (CL) at the beginning of your Start Up Program (QSTRUP). If you do not want the independent disk pools to be made available when you restart the server, you can either Delete the Data Area (DLTDTAARA) or Rename it (RNMOBJ). However, you need to remember to either Create the Data Area (CRTDTAARA) again or Rename it (RNMOBJ) back to the data area being checked in your Start Up Program. Only the QSYSWRK subsystem should be started before making the independent disk pools available. Then other work will not compete for system resources while your independent disk pools are being made available.

Note: Read the “Code disclaimer information” on page 57 for important legal information.

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

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

Recommended structure for independent disk pools

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

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

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

Structuring disk pool groups

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

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

- If an application consists solely of data in user-defined file systems and the data is not to be journaled, a UDFS disk pool might be the best choice. There is less overhead associated with a UDFS disk pool. There is also less extendibility since the UDFS disk pool cannot contain any library-based objects.
- If you have an application with multiple instances of the application data that you want to keep separate, then you should consider a separate disk pool group for each data instance. See “Dedicated independent disk pools” on page 51 for an example of this scenario.
- If you have multiple applications and the application data is independent, a separate disk pool group for each application might be the appropriate answer. One application’s data is then isolated from other applications and each application is unaffected by actions on others. The application data can therefore be brought online, taken offline, or switched without affecting other applications.
- If you have multiple applications with interdependent data objects, the data for those applications should be combined into a single disk pool group.
- You can use secondary disk pools to separate data objects into different storage domains and thus achieve better performance. The normal use of this is to separate your journal receivers onto different disk units from the data being journaled by placing the journal receivers in a secondary disk pool. However, you might also separate other parts of your application onto different disk units providing that they are in different libraries and the following journaling dependency is satisfied.
- Objects being journaled and the journal for those objects must be on the same disk pool.

Switchable and stand-alone independent disk pools

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

Independent disk pools in a multisystem clustered environment

A group of servers in a cluster can take advantage of the switchover capability within clusters to move access to the independent disk pool from server to server. In this environment, an independent disk pool can be switchable when it resides on a switchable device: an external expansion unit (tower), an input/output processor (IOP) on the bus shared by logical partitions, or an IOP that is assigned to an I/O pool. The server that owns, or is attached to, the switchable device containing the independent disk pool can then be switched, either automatically, in the case of an unplanned outage (failover), or manually by administering a switchover.

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

Dedicated independent disk pools in a single-system environment

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

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

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

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

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

For more on switchable and dedicated independent disk pools, including example configurations for each of these environments, see “Examples: Independent disk pool configurations” on page 50.

Disk pool groups

A disk pool group is made up of a primary disk pool and zero or more secondary disk pools. Each disk pool is independent in regard to data storage, but in the disk pool group they combine to act as one entity. If you make one disk pool available or unavailable, the rest of the disk pools in the group are also made available or unavailable at the same time. Also, in a clustered environment, all of the disk pools in a group switch to another node at the same time.

An example of a practical use for a disk pool group is to isolate journal receivers from the objects for which they contain journal entries. The primary disk pool might contain the libraries, journals, and

objects to be journaled, while the secondary disk pools might contain the associated journal receivers. The journals and journal receivers remain separate for maximum performance and recoverability, but they function together in the disk pool group.

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

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

Geographic mirroring

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

Geographic mirroring is a subfunction of cross-site mirroring (XSM), which is part of OS/400 Option 41, High Available Switchable Resources.

Benefits of geographic mirroring

Geographic mirroring provides these benefits:

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

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

- You can protect the production copy and mirror copy with your choice of protection, either disk unit mirroring or device parity protection. The production copy and mirror copy are not required to have the same type of protection
- You can set the threshold of the disk pool to warn you when storage space is running low. The server sends a message, allowing you the time to add more storage space or to delete unnecessary objects. If the user ignores the warning and the mirror copy disk pool becomes full, geographic mirroring is suspended. If the user ignores the warning and the production disk pool becomes full, the application stops and objects cannot be created.
- The mirror copy can be detached and then separately made available to perform save operations, to create reports, or to perform data mining. When the mirror copy is reattached, it is synchronized with the production copy, and all modifications made to the detached copy are lost. Synchronization can be a lengthy process.
- If you configure the independent disk pools to be switchable, you increase your options to have more backup nodes that allow for failover and switchover methods. See “Example: Independent disk pools with geographic mirroring” on page 57.

Costs and limitations of geographic mirroring

Costs

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

- At least one iSeries server at each site.
- Sufficient CPU support for the additional CPU capacity required for geographic mirroring. A fraction of a processor for a partition supporting geographic mirroring is not adequate.
- Sufficient disk units at each site for the production and the mirror copy of the geographically mirrored independent disk pools. To avoid disk unit contention, use separate input/output adapters for the production copy on its node and for the mirror copy on its node.
- One TCP/IP connection from each node should connect the two sites. A second TCP/IP connection is strongly recommended to provide redundancy and better performance. You may configure up to four TCP/IP connections. See “Communications requirements” on page 25 for more information.

Limitations

Limitations of geographic mirroring include these constraints:

- When geographic mirroring is being performed, you cannot access the mirror copy; this ensures that the data integrity of the mirror copy is maintained.
- If you detach the mirror copy to perform a save operation, to perform data mining, or to create reports, you must reattach the mirror copy to resume geographic mirroring. The mirror copy must be synchronized with the production copy after it is reattached. Synchronization can be a lengthy process.
- If you suspend geographic mirroring, you must resume geographic mirroring to send updates to the mirror copy again. When you resume geographic mirroring, the mirror copy must be synchronized with the production copy. Synchronization can be a lengthy process.

How geographic mirroring works

Configure

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

See “Configure geographic mirroring with dedicated independent disk pools” on page 33 and “Configure geographic mirroring with dedicated independent disk pools” on page 33 for more information.

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

Manage

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

- **Synchronous mode:** The client waits until the operation is complete to disk on both the source and target systems. The mirror copy is always eligible to become the production copy, because the order of

writes is preserved on the mirror copy. It is recommended to try synchronous mode first. If your performance remains acceptable, continue to use synchronous mode.

- **Asynchronous mode:** The client must wait only until the operation is complete to disk on the source system and is received for processing on the target system. However, synchronous mode is safer because if the primary node fails or the production copy fails, the mirror copy can become the production copy. In asynchronous mode, the pending updates must be completed before the mirror copy can become the production copy

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

Synchronization

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

Synchronization priority

When you set the attributes for geographic mirroring, you can set the synchronization priority. If synchronization priority is set high, the system uses more resources for synchronization, which results in a sooner completion time. The mirror copy is eligible to become a production copy faster, so you are protected sooner. However, high priority can cause degradation to your application. It is recommended that you try high priority first, so you are protected as soon as possible. If the degradation to your application performance is not tolerable, then lower the priority.

Recovery timeout

In addition to synchronization priority, you can also set recovery time out. The recovery timeout specifies how long your application can wait when geographic mirroring cannot be performed. When an error, such as IP failure, prevents geographic mirroring, the source system waits and retries for the specified recovery timeout before suspending geographic mirroring which allows your application to continue. The trade-off is between blocking your application or requiring full synchronization after suspending geographic mirroring. When your application is blocked for an extended time, other jobs might also be blocked waiting for resources and locks owned by the applications using the geographic mirrored disk pool. When geographic mirroring is suspended, you no longer have the protection of the mirror copy. If your application can tolerate a delay, it is recommended to set recovery timeout from 2 to 5 minutes. If the volume of your data is large (over a terabyte), consider a longer recovery timeout value to reduce the possibility of suspending geographic mirroring. If mirroring is suspended, the system performs a full synchronization.

System roles

When you configure the cluster for geographic mirroring, you have many options for defining the availability and protection of the independent disk pool. When you create the switchable hardware group, you list the order of the backup systems to which the independent disk pool will failover or switch over. If the primary node switches to a backup node at the same site, a hardware switch will occur. If the primary node switches to the other site, the mirror copy on the backup node changes roles to become the production copy. The old primary node becomes the new backup node, and the production copy becomes the mirror copy. The new production copy is now accessible for updates on the remote

system. If the independent disk pools are part of a disk pool group, all of the disk pools in the group will switchover together. See “Example: Independent disk pools with geographic mirroring” on page 57.

Requirements for geographic mirroring

- Geographic mirroring increases CPU load, so there must be sufficient excess CPU capacity. Add processors as needed to increase CPU capacity.
 - For optimal performance for geographic mirroring, particularly during synchronization, increase your machine pool size by the amount given by the following formula:
The amount of extra machine pool storage is: $271.5 \text{ MB} + .2\text{MB}$.
The extra machine pool storage is required on the target node. However, because the target node changes when doing switchovers and failovers, you should increase the machine pool on all nodes in the Cluster Resource Group. To prevent the performance adjuster from reducing the machine pool size, you should do one of the following:
 1. Set the machine pool minimum size to the calculated amount (the current size plus the extra size for geographic mirroring from the formula) using Work with Shared Storage Pools (WRKSHRPOOL) or Change Shared Storage Pool (CHGSHRPOOL).
- Note:** It is recommended to use this option with the Work with Shared Storage Pools (WRKSHRPOOL) option.
2. Set QPFRADJ to zero which prohibits the performance adjuster from changing size of the machine pool.
- Configure a separate storage pool for the jobs using geographic mirrored independent disk pools, especially if you specify a long recovery timeout.
 - The extra machine pool storage is required on the target node. However, since the target node changes when doing switchovers and failovers, you should increase the machine pool on all nodes in the Cluster Resource Group.
 - Geographic mirroring is performed when the disk pool is available. When geographic mirroring is being performed, the system value for the time of day (QTIME) should not be changed.
 - Communications requirements for independent disk pools are particularly critical as they affect throughput. See “Communications requirements” on page 25 for more information.
 - All independent disk pool requirements must be met. See “Plan for independent disk pools” on page 21 for more information.

Failover and switchover

Mirror copy failover or switchover

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

When geographic mirroring is suspended

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

Examples

The following are examples of failovers and switchovers:

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

Ending clustering

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

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

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

Shut down System

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

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

Note: After suspending geographic mirroring, a full resync will be required when geographic mirroring is resumed.

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

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

Recovery from two production copies

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

Considerations for making a disk pool available at failover or switchover

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

Rolling Upgrades

Upgrades of OS/400 releases made to any nodes involved in geographic mirroring require a rolling upgrade. A rolling upgrade is required because a node at an earlier release might not be able to perform geographic mirroring to a node at a later release and a node at a later release usually cannot perform geographic mirroring to a node at an earlier release. This forces the nodes to be upgraded in an order dictated by the recovery domain order starting with the node that is the last backup. During the rolling upgrade, the production copy and mirroring copy will be moved to their recovery nodes.

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


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

Steps	Recovery domain order			
	Independent disk pool 33		Independent disk pool 34	
	During	After	During	After
9. Switch production copy of independent ASP 33 (A to B)		B, C, D		C, D, B
10. Upgrade node A	B, C, D	B, A, C, D	C, D, B	C, D, B, A
11. Switch production copy of independent ASP 33 (B to A)		A, B, C, D		C, D, B, A
12. Switch mirror copy of independent ASP 34 (B to A)	A, B, C, D			C, D, A, B

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

Plan for independent disk pools

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

Important: When you are ready to order a new server or a server upgrade to use clusters, IBM will assist you in making sure that your cluster requirements are met. See [Planning for Clustering](#) .

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

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

iSeries Navigator requirements for disk management

Changing the disk configuration of your server is a time-consuming process, so you want to plan carefully to be as efficient as possible. Before you begin disk management using iSeries Navigator, go through this short list to ensure that you are ready.

Access disk units in iSeries Navigator

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

Install the Configuration and Service component

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

Enable the Disk Units folder

1. In iSeries Navigator, right-click the server connection and select **Application Administration**.
2. On the resulting window, click **OK**.

3. Click the **Host Applications** tab.
4. Expand Operating System/400® > **Service**.
5. Select **Disk Units** to have **Default Access** or **All Object Access**.
6. Click **OK**.
7. Restart iSeries Navigator.

Access the Disk Units folder to perform all disk management functions

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

Set up communication

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

Configuring the service tools server

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

Setting the service IP address

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

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

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

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

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

Hardware requirements

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

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

Note:

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

Physical planning requirements

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

Multisystem clustered environment (for switchable independent disk pools)

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

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

The switchable expansion unit must be SPCN-cabled to the system unit that will initially serve as the primary node for the switchable hardware group (device CRG). The primary node might be a primary or secondary logical partition within the system unit. If using logical partitions, the system buses in the intended expansion unit must be owned and dedicated by the partition involved in the cluster.

This figure shows an HSL loop that is configured correctly, as described previously. There are no physical planning requirements.




Figure 1. Single-system environment

Software and licensing requirements

Depending on how you plan to use independent disk pools, you must have the following software and licenses:

Multisystem clustered environment

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

1. You need OS/400 V5R1M0¹ or later.

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

2. iSeries Navigator is the graphical user interface for managing and administering your iSeries server from your Windows[®] desktop. It is required to perform some of the disk management tasks necessary to use independent disk pools. See "iSeries Navigator requirements for disk management" on page 21 for steps to enable iSeries Navigator for disk management.
3. You need to install Option 41 HA Switchable Resources. Option 41 gives you the capability to switch independent disk pools between systems. To switch an independent disk pool between servers, the servers must be members of a cluster and the independent disk pool must be associated with a switchable hardware group in that cluster. Option 41 also gives you the capability of using the iSeries Navigator cluster management interface to define and manage a cluster that uses switchable resources.

Single-system environment

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

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

Communications requirements

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

Single-system environment

There are no communications requirements.

Multisystem clustered environment

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

- **For switchable independent disk pools**, at least one TCP/IP communications interface between the servers in the cluster. For redundancy, it is recommended that there be at least two separate interfaces between the servers.
- **For geographic mirroring**, the following are recommended:
 - Up to a maximum of four unique TCP/IP addresses, used exclusively for geographic mirroring. Geographic mirroring can generate heavy communications traffic. If geographic mirroring shares the same IP connection with another application, for example clustering, then geographic mirroring might be suspended which results in synchronization. Likewise, clustering response might be unacceptable which results in partitioned nodes.
 - Throughput for each data port connection should match, meaning that the speed and connection type should be the same for all connections between server pairs. If throughput is different, performance will be gated by the slowest connection.
- Consider configuring a virtual private network for TCP/IP connections for the following advantages:
 - Security of data transmission
 - Encrypted data
 - Increased reliability of data transmission
 - Greater redundancy

Connections from the production system

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

Cluster requirements

If you plan to use *switchable* independent disk pools or *geographically mirrored* independent disk pools, you need to configure an iSeries cluster. The documentation in these independent disk pools topics guides you through the creation and management of your cluster. However, you may want to prepare your network and server environment in advance.

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

Application considerations for independent disk pools

When you are designing or restructuring your application environment for use with independent disk pools, there are several things you should be aware of. A few of these considerations include the existence of multiple databases, the objects that can and cannot be created in an independent disk pool, how the library list works, and the placement of programs and data in the correct database.

When a primary independent disk pool is made available for the first time, a new database with the same name is also generated by default. See “Independent disk pools with distinct databases” on page 9 for more information. If you write an application to access files and libraries in a disk pool group, you must specify how to access that specific database. Some of your options include:

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

As you write applications that create objects, you must know which objects are supported. See “Supported and unsupported OS/400 object types” on page 7. If your application uses the Create Library (CRTLIB) command, you must specify CRTLIB ASP(*ASPDEV) ASPDEV(*asp-device-name*). If you do not specify these parameters for CRTLIB, the library is created in the system disk pool by default. However, if you use the SQL statement, CREATE COLLECTION, the default for the IN ASP clause is the current library namespace.

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

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

It might also be helpful to understand how the library list works when independent disk pools are implemented. When the library list includes QSYS, QSYS2, or SYSIBM, the “Multiple system libraries” on page 10 in the independent disk pool (QSYSnnnnn, QSYS2nnnnn, SYSIBnnnnn) are searched before the libraries in the system disk pool. If the object is found in the independent disk pool, the system disk pool will not be searched. In addition, if you switch to a different disk pool group, any libraries that were in the previous library list are removed from the current library list.

You also need to carefully consider where you store data, applications, and application exit programs. It is recommended that data should be stored in independent disk pools. If your independent disk pools

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

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

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

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

Configure independent disk pools

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

“iSeries Navigator requirements for disk management” on page 21

Create a dedicated independent disk pool

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

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

Prerequisite

“iSeries Navigator requirements for disk management” on page 21.

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

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

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

Create a switchable independent disk pool

Before you attempt to do switchable independent disk pools, ensure that you have satisfied the hardware, software, communications, and physical planning requirements. See “Plan for independent disk pools” on page 21.

iSeries Navigator is the recommended interface for creating and managing independent disk pools. Wizards in the clusters and disk management components simplify the tasks and guide you through the process. For some disk management tasks, iSeries Navigator is the only option. Make sure you have fulfilled the “iSeries Navigator requirements for disk management” on page 21.

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

Using iSeries Navigator

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

1. “Create a cluster” on page 29. To use switchable independent disk pools, an iSeries cluster is required.
2. “Make your hardware switchable” on page 44. If you have a stand-alone expansion unit or an IOP that contains disk units that are to be included in an independent disk pool, you must authorize the expansion unit or IOP to grant access to other nodes.
3. “Create a switchable hardware group” on page 30. A switchable hardware group, also known as a device CRG, defines the switchable independent disk pool. This is what manages the switching of the device. This wizard takes you through the steps to create a new switchable hardware group. It will also guide you through the New Disk Pool wizard which will assist you in creating a new disk pool and adding disk units to it for the cluster.

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

4. “Print your disk configuration” on page 31. Print your disk configuration to have in case a recovery situation occurs. Also, record the relationship between the independent disk pool name and number.


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

1. “Start switchable hardware group” on page 31. Start the switchable hardware group to enable device resiliency for the switchable hardware group.
2. “Make a disk pool available” on page 42. To access the disk units in an independent disk pool, you must make the disk pool available (vary on) the disk pool.
3. Perform a test switchover. Before you add data to the disk pool, perform a test switchover on the switchable hardware group you created to ensure the configuration functions as you planned.

Using CL commands and APIs

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

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

1. **Create the cluster.** Create the cluster with required node using the CRTCLU (Create Cluster) Command.
2. **Start the nodes that comprise the cluster.** Start the nodes in the cluster using the STRCLUNOD (Start Cluster Node) Command
3. **Create the device domain.** You must create the device domain for all nodes involved in switching an independent disk pool or set of independent disk pools using the ADDDEVDMNE (Add Device Domain Entry) command.
4. **Create the device descriptions.** Device descriptions must be created on each node that will be in the cluster resource group (CRG). Use the CRTDEVASP (Create Device Description (ASP)) Command. On the command line in the character-based interface, enter CRTDEVASP. In the **Resource Name** and the **Device Description** fields, enter the name of the independent disk pool you plan to create.
5. **Create the cluster resource group.** Create the device CRG with the nodes, their roles in the recovery domain, and independent disk pool device descriptions using the CRTCRG (Create Cluster Resource Group) Command.
6. **“Make your hardware switchable” on page 44.** If you have a stand-alone expansion unit or an IOP that contains disk units that are to be included in an independent disk pool, you must authorize the expansion unit or IOP to grant access to other nodes (**iSeries Navigator required**).
7. . Create the disk pool on the node that owns the disk units using the New Disk Pool wizard when the server is fully restarted. Make sure clustering is active before you start. Name the independent disk pool to match the device description resource name that you gave in step 3. As you add disk units, it is best to localize disk units in the same expansion unit or IOP. Do not spread the disk pool across more device parity sets than necessary.
8. **Print your disk configuration.** Print your disk configuration to have in case of a recovery situation. See How to display your disk configuration in Backup and Recovery.  Also, record the relationship between the independent disk pool name and number.
* You have now created a switchable independent disk pool. The remaining steps are required to prepare it for use.
9. **Start the cluster resource group.** Start the cluster resource group to enable device resiliency using the STRCRG (Start Cluster Resource Group) Command.
10. **Make the disk pool available.** To access the disk units in an independent disk pool you must vary on the disk pool using the VRYCFG (Vary Configuration) command.
11. **Perform a test switchover.** Before you add data to the disk pool, perform a test switchover to ensure the configuration functions as you planned. Use the CHGCRGPRI (Change CRG Primary) command.

You are now ready to populate the independent disk pool with directories and libraries. Before you do, be sure to read “Independent disk pools with distinct databases” on page 9.

Create a cluster

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

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

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

1. For step-by-step instructions on how to create a cluster, see Create a cluster in the Clusters topic.
2. Verify that all nodes are at the correct potential cluster version. Potential cluster version must be at least 3 for switchable independent disk pools that support libraries. To allow for the V5R3M0 capability of geographic mirroring, the potential cluster version must be set to at least 4. See Adjust the cluster version of a cluster for details.
3. Start all nodes in the cluster, or at least those that will be in the device domains. See Start a cluster node for details.

Create a switchable hardware group

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

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

Using iSeries Navigator (requires Option 41 (OS/400 - HA Switchable Resources))

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

To add a switchable hardware group, follow these steps:

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

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

Using Cluster CL commands and APIs

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

Add Device Domain Entry

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

- ADDDEVDMNE (Add Device Domain Entry) Command

- Add Device Domain Entry (QcstAddDeviceDomainEntry) API

Create Cluster Resource Group

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

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

Start switchable hardware group

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

To start a switchable hardware group, follow these steps:

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

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

Create disk pool

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

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

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

1. In iSeries Navigator, expand **My Connections** (or your active environment).
2. Expand your iSeries server>**Configuration and Service**> **Hardware**>**Disk Units**.
3. Sign on to service tools if the Service Tools Signon dialog box displays. You may be required to configure the service tools server if you have not done so already.
4. Right-click **Disk Pools** and select **New Disk Pool**.
5. Follow the wizard's instructions to add disk units to a new disk pool.

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

Print your disk configuration

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

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

1. In iSeries Navigator, expand **My Connections**.
2. Expand your iSeries server>**Configuration and Service**>**Hardware**>**Disk Units**.
3. Right-click **All Disk Units**, and select **Graphical View**.

4. Select **Show device positions** to associate disk unit names with the device position where they are inserted.
5. On the **Disk Units Graphical View** dialog, select **File> Print**.

Create a new disk pool group

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

You can create a disk pool group and add disk units to the individual disk pools by using the New Disk Pool wizard. If you have existing UDFS disk pools that you would like to include in a disk pool group, see “Convert a UDFS disk pool to primary” or “Convert a UDFS disk pool to secondary” on page 33.

Note: If you want to create a switchable independent disk pool (UDFS, Primary, or Secondary), you must create the cluster first. For more information, see “Create a switchable independent disk pool” on page 28.

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

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

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

Convert UDFS disk pools

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

You must convert UDFS disk pools if you want them to participate in a “Disk pool groups” on page 14. After you convert a UDFS disk pool to a primary or secondary disk pool, you cannot convert it back to a UDFS disk pool. You must create a primary disk pool before you can associate secondary disk pools.

Convert a UDFS disk pool to primary

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

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

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

1. In iSeries Navigator, expand My Connections (or your active environment).
2. Expand your iSeries server>**Configuration and Service>Hardware>Disk Units**.
3. If the Service Tools Signon dialog box displays, sign on to service tools.

4. Select **Disk Pools**.
5. Right-click the required UDFS **Disk Pool** and select **Convert to> Primary disk pool**.
6. On the **Confirm Convert to Primary Disk Pool** dialog box the default for the **Database Name** field is Generated by the system, which means the system generates a database name for you.
7. Click **Convert Disk Pool**.
8. If you need to associate other existing UDFS disk pools with your new primary disk pool in a disk pool group, see “Convert a UDFS disk pool to secondary.”

Convert a UDFS disk pool to secondary

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

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

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

1. In iSeries Navigator, expand My Connections (or your active environment).
2. Expand your iSeries server>**Configuration and Service>Hardware>Disk Units**.
3. If the Service Tools Signon dialog box displays, sign on to service tools.
4. Select **Disk Pools**.
5. In the right pane you may select one or more UDFS disk pools to convert at the same time. Right-click the required **UDFS Disk Pool(s)** and select **Convert to > Secondary disk pool**.
6. On the **Confirm Convert to Secondary Disk Pool** dialog box, select the primary disk pool that you need to associate with the secondary disk pools. The primary disk pool selected should not have been made available prior to being associated with the secondary disk pools. Only the primary disk pools that are currently owned by the system can be selected. You will not be able to change the primary after you perform this action.
7. Click **Convert Disk Pool**.
8. If the disk pool you converted to secondary is in a device cluster resource group, then you must change the Online attribute to *PRIMARY. Use the Change Cluster Resource Group Device Entry CHGCRGDEVE command or the (QcstChangeClusterResourceGroupDev) API to change the Online attribute to *PRIMARY.

Configure geographic mirroring with dedicated independent disk pools

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

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

independent disk pool that is dedicated to Node B.



Prerequisites

Communications requirements

Customize TCP/IP with iSeries Navigator

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

1. Plan and configure your data port TCP/IP routes.
2. Create a cluster.
3. Create the independent disk pool that you want to mirror.
4. Create the device cluster resource group, also known as switchable hardware group, for the independent disk pool that you just created:
 - a. In iSeries Navigator, expand **Management Central**.
 - b. Expand **Clusters**.
 - c. Expand the cluster for which you need to add a switchable hardware group.
 - d. Right-click **Switchable Hardware**, and select **New Group**.
 - e. On the Create new or add existing disk pool dialog box, select **No, add an existing switchable disk pool to the switchable hardware group**.
5. Define your geographic mirroring sites in the recovery domain:
 - a. Right-click the newly created switchable hardware group and select **Properties**.
 - b. Select the **Recovery Domain** tab.
 - c. Select the primary node and click **Edit**.
 - d. In the site name field, specify the primary site for the production copy.
 - e. Click **Add** to specify the data port IP addresses of the primary node.
 - f. On the Edit Node dialog, specify the data port IP addresses for the primary node that you setup in step 1, Plan and configure your TCP/IP, and click **OK**. You can configure up to four data port IP addresses in a one-to-one manner, one port to one port, each port independent of the other. You should consider configuring multiple communication lines to allow for redundancy and the highest throughput. The same number of ports used here should be configured on the mirroring node.
 - g. On the General tab, click **OK**.
 - h. Repeat the previous steps to specify the site name and IP addresses for the mirror copy node.

6. After you have completed the XSM prerequisites, follow these steps to configure geographic mirroring:
 - a. In iSeries Navigator, expand **My Connections** (or your active environment).
 - b. Expand the primary node on the source iSeries server.
 - c. Expand **Configuration and Service>Hardware>Disk Units>Disk Pools**.
 - d. If the Geographic Mirroring columns are not displayed, click the Disk Pool you want to mirror, and select **View>Customize this view>Columns**, then select the desired columns with the suffix Geographic Mirroring from the **Columns available to display list**.
 - e. Right-click the Disk Pool you want to mirror, and select **Geographic Mirroring> Configure Geographic Mirroring**.
 - f. Follow the wizard's instructions to configure geographic mirroring.

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

7. You have now configured geographic mirroring. The remaining steps are required to prepare the independent disk pool for use in this environment. See "Print your disk configuration" on page 31. Print your disk configuration to have in case a recovery situation occurs. Also, record the relationship between the independent disk pool name and number.
 1. "Start switchable hardware group" on page 31. Start the switchable hardware group to enable device resiliency for the switchable hardware group.
 2. "Make a disk pool available" on page 42. To access the disk units in an independent disk pool you must make the disk pool available (vary on) the disk pool.
 3. Wait for resync to complete.
 4. Perform a test switchover. Before you add data to the disk pool, perform a test switchover to ensure the backup node can become the primary node and the primary node can become the backup node.

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

Using CL commands and APIs


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

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

1. Plan and configure your data port TCP/IP routes on all nodes in the recommended manner as follows:
 - Node A should have routes to C and D.
 - Node B should have routes to C and D.
 - Node C should have routes to A and B.
 - Node D should have routes to A and B.
2. **Create the cluster.** Create the cluster with desired nodes using the CRTCLU (Create Cluster) command.
3. **Start the nodes that comprise the cluster.** Start the nodes in the cluster using the STRCLUNOD (Start Cluster Node) command
4. **Create the device domain.** You must create the device domain for all nodes involved in switching an independent disk pool using the ADDDEVDMNE (Add Device Domain Entry) command.

5. **Create the device descriptions.** Device descriptions must be created on each node that will be in the cluster resource group (CRG). Use the CRTDEVASP (Create Device Description (ASP)) command. On the command line in the character-based interface, enter CRTDEVASP. In the **Resource Name** and the **Device Description** fields, enter the name of the independent disk pool you plan to create.
6. **Create the cluster resource group.** Create the device CRG with the nodes, their roles in the recovery domain, and the independent disk pool device descriptions. You must also specify a site name and data port IP addresses for each node in the recovery domain. Use the CRTCRG (Create Cluster Resource Group) command.
7. “Create disk pool” on page 31. Create the disk pool on the node that owns the disk units using the New Disk Pool wizard when the server is fully restarted. Make sure clustering is active before you start. Name the independent disk pool to match the device description resource name that you gave in step 3. As you add disk units, it is best to localize disk units in the same expansion unit or IOP. Also, do not spread the disk pool across device parity sets (**iSeries Navigator required**).
8. Follow these steps to configure geographic mirroring:
 - a. In iSeries Navigator, expand **My Connections** (or your active environment).
 - b. Expand any iSeries server.
 - c. Expand **Configuration and Service**.
 - d. Expand **Hardware**.
 - e. Expand **Disk Units**.
 - f. Expand **Disk Pools**.
 - g. Right-click the Disk Pool you want to mirror and select **Geographic Mirroring > Configure Geographic Mirroring**.
 - h. Follow the wizard’s instructions to configure geographic mirroring.

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

9. **Print your disk configuration.** Print your disk configuration to have in case of a recovery situation. See How to display your disk configuration in Backup and Recovery.  Also, record the relationship between the independent disk pool name and number.
You have now configured geographic mirroring. The remaining steps are required to prepare the independent disk pool for use in this environment.

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

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

Configure geographic mirroring with switchable independent disk pools

To configure geographic mirroring you must first configure your cross-site mirroring (XSM) environment and create the independent disk pool that you want to mirror. Before using the iSeries Navigator, you should also define up to four, one-to-one data port TCP/IP routes bidirectionally as part of the connection between all the nodes in the cluster resource group. Geographic mirroring allows you to maintain an exact copy of the independent disk pool on a system at a different location for protection and availability purposes. Configuring your independent disk pool to be switchable between nodes at the same site in the cluster allows for greater availability options. See “Example: Independent disk pools with geographic mirroring” on page 57.

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



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

1. Plan and configure your data port TCP/IP routes. See “Communications requirements” on page 25, and Customize TCP/IP with iSeries Navigator.
2. Create a cluster containing nodes A and B.
3. “Make your hardware switchable” on page 44. If you have stand-alone expansion units or IOPs that contain disk units that are to be included in an independent disk pool, you must authorize the expansion unit or IOP to grant access to other nodes at the same site.

4. “Create a switchable hardware group” on page 30. A switchable hardware group, also known as a device CRG, defines the switchable independent disk pool. This is what manages the switching of the device. This wizard takes you through the steps to create a new switchable hardware group. It will also guide you through the New Disk Pool wizard which assists you in creating a new disk pool and adding disk units to it for the cluster.

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

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

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

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

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

7. Add nodes C and D to the switchable hardware group:
 - a. Right-click the newly created switchable hardware group and select **Properties**.
 - b. Select the **Recovery Domain** tab.
 - c. Click **Add**.
 - d. Select the node and click **OK**. Repeat for each node.
8. Define your geographic mirroring sites in the recovery domain:
 - a. Right-click your switchable hardware group and select **Properties**.
 - b. Select the **Recovery Domain** tab.
 - c. Select the primary node and click **Edit**.
 - d. In the site name field, specify the primary site for the production copy.
 - e. Click **Add** to specify the data port IP addresses of the primary node.
 - f. On the Edit Node dialog box, specify the data port IP addresses for the primary node that you set up in step 1, Plan and configure your TCP/IP routes, and click **OK**. You can configure up to four data port IP addresses. You should consider configuring multiple communication lines to allow for redundancy and the highest throughput. The same number of ports used here should be used on all nodes.
 - g. On the General tab, click **OK**.

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

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

10. "Print your disk configuration" on page 31. Print your disk configuration to have in case a recovery situation occurs. Also, record the relationship between the independent disk pool name and number.

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

1. "Start switchable hardware group" on page 31. Start the switchable hardware group to enable device resiliency for the switchable hardware group.
2. "Make a disk pool available" on page 42. To access the disk units in an independent disk pool you must make the disk pool available (vary on) the disk pool.
3. Wait for resync to complete.
4. Perform a test switchover. Before you add data to the disk pool, perform a test switchover on the switchable hardware group you created to ensure that each node in the recovery domain can become the primary node.

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

Using CL commands and APIs


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

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

1. Plan and configure your TCP/IP routes on all nodes, as follows:
 - Node A should have routes to C and D.
 - Node B should have routes to C and D.
 - Node C should have routes to A and B.
 - Node D should have routes to A and B.
2. **Create the cluster.** Create the cluster with required nodes using the CRTCLU (Create Cluster) command.

3. **Start the nodes that comprise the cluster.** Start the nodes in the cluster using the STRCLUNOD (Start Cluster Node) command
4. **Create the device domain.** You must create the device domain for all nodes involved in switching an independent disk pool using the ADDDEVDMNE (Add Device Domain Entry) command. All nodes must be in the same device domain.
5. **Create the device descriptions.** Device descriptions must be created on all nodes that will be in the cluster resource group (CRG). Use the CRTDEVASP (Create Device Description (ASP)) command. On the command line in the character-based interface, enter CRTDEVASP. In the **Resource Name** and the **Device Description** fields, enter the name of the independent disk pool you plan to create.
6. **Create the cluster resource group.** Create the device CRG with the nodes, their roles in the recovery domain, and independent disk pool device descriptions. You must also specify a site name and up to four data port IP addresses for each node in the recovery domain.
7. **“Make your hardware switchable” on page 44.** If you have stand-alone expansion units or IOPs that contain disk units that are to be included in an independent disk pool, you must authorize the expansion unit or IOP to grant access to other nodes at the same site (**iSeries Navigator required**).
8. **“Create disk pool” on page 31.** Create the disk pool on the node that owns the disk units using the New Disk Pool wizard when the server is fully restarted. Make sure clustering is active before you start. Name the independent disk pool to match the device description resource name that you gave in step 3. As you add disk units, it is best to localize disk units in the same expansion unit or IOP. Also, do not spread the disk pool across device parity sets (**iSeries Navigator required**).
9. Follow these steps to configure geographic mirroring:
 - a. In iSeries Navigator, expand **My Connections** (or your active environment).
 - b. Expand the iSeries server that is the primary node .
 - c. Expand **Configuration and Service**.
 - d. Expand **Hardware**.
 - e. Expand **Disk Units**.
 - f. Expand **Disk Pools**.
 - g. Right-click the Disk Pool you want to mirror and select **Geographic Mirroring> Configure Geographic Mirroring**.
 - h. Follow the wizard’s instructions to configure geographic mirroring.

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

Note: The mirror copy and the production copies must be in different sites. If we have two sites, AB and CD and the production copy is on node A on site AB, the backup copy must be on node C or D on site CD.
10. **Print your disk configuration.** Print your disk configuration to have in case of a recovery situation. See How to display your disk configuration in Backup and Recovery.  Also, record the relationship between the independent disk pool name and number.

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

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

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


Manage independent disk pools

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

Backup and recovery of independent disk pools

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

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

procedures to restore the data you have saved. See the Backup and Recovery  manual for steps to restore information to the independent disk pools.

If you experience problems accessing an independent disk pool or making it available, the problem might be with the disk pool. The configuration source might be corrupted or the primary and secondary disk pools might need to be re-associated. See the following procedures for steps to recover the disk pools:

Recover an independent disk pool

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

- The configuration source is corrupted. When corruption occurs, the independent disk pool will appear to have no disk units in it. The disk pool may also appear to have no disk units in it if it is switched to another node in a clustered environment. Before you attempt a recovery, make sure that no other system owns the disk pool. If you know the serial numbers of the disk units in the independent disk pool that might need recovery, make sure you are on the system that owns those disk units and that they show as nonconfigured.

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

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

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

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

4. Expand **Hardware**.
5. Expand **Disk Units**.
6. If the Service Tools Signon dialog box displays, sign on to service tools.
7. Select **Disk Pools**.
8. Right-click the problematic disk pool. If iSeries Navigator detects one of the problems listed above, then **Recover Configuration** or **Recover Unknown Configuration Source** appears in the list. If you see either of these options, select it to continue.
9. Follow the instructions on the dialog box displayed.

Recover disk pool group

If the primary disk pool for a secondary disk pool is deleted, or if the primary disk pool is not aware of the secondary disk pool, the secondary disk pool needs to be re-associated with a primary disk pool. You can recover the disk pool group through iSeries Navigator.

To recover a disk pool group, follow these steps:

1. In iSeries Navigator, expand **My Connections** (or your active environment).
2. Expand any iSeries server.
3. Expand **Configuration and Service**.
4. Expand **Hardware**.
5. Expand **Disk Units**.
6. If the Service Tools Signon dialog box displays, sign on to service tools.
7. Select **Disk Pools**.
8. Verify that the required primary disk pool exists. If it does not, you must “Create disk pool” on page 31.
9. In the right pane you can select one or more secondary disk pools that need to be associated with the primary disk pool. Right-click all the secondary **Disk Pools** that need to be associated with a primary disk pool and select **Recover Group**.
10. On the **Confirm Recover Disk Pool Group** dialog box, select the primary disk pool that you want associated with the secondary disk pools. Only the primary disk pools that are currently owned by the system are available to select. You cannot change the primary after you perform this action.
11. Click **Recover Group**.

Make a disk pool available

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

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

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

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

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

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

To make an independent disk pool available:

1. In iSeries Navigator, expand **My Connections** (or your active environment).
2. Expand the primary node iSeries server.
3. Expand **Configuration and Service**.
4. Expand **Hardware**.
5. Expand **Disk Units**.
6. Sign on to service tools if the Service Tools Signon dialog box displays.
7. Expand **Disk Pools**.
8. Right-click the unavailable disk pool and select **Make Available**. You can select multiple disk pools to make available at the same time.
9. From the dialog box displayed, click **Make Available** to make the disk pool available.

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

Make a disk pool unavailable

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

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

Important: Before an independent disk pool can be made unavailable, no jobs can hold reservations on the disk pool. See “Release job reservations on an independent disk pool” on page 44 for details on determining whether jobs are using an independent disk pool and how to release the job reservations.

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

To make an independent disk pool unavailable:

1. In iSeries Navigator, expand **My Connections** (or your active environment).
2. Expand the iSeries server that is the primary node .
3. Expand **Configuration and Service**.
4. Expand **Hardware**.
5. Expand **Disk Units**.
6. Sign on to service tools if the Service Tools Signon dialog box displays.
7. Expand **Disk Pools**.
8. Right-click the disk pool you want to make unavailable and select **Make Unavailable**.
9. From the dialog box that displays, click **Make Unavailable** to make the disk pool unavailable.

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

Release job reservations on an independent disk pool

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

To view the jobs using an independent disk pool:

1. In iSeries Navigator, expand **My Connections** (or your active environment).
2. Expand any iSeries server.
3. Expand **Configuration and Service**.
4. Expand **Hardware**.
5. Expand **Disk Units**.
6. Expand **Disk Pools**.
7. Right-click the disk pool and select **Jobs**.

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

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

Make your hardware switchable

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

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

Make a frame/unit switchable

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

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

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

6. Expand **By Location** and select the frame/units you want to make switchable.
7. Right-click a highlighted frame/unit and select **Make Switchable**.
8. Follow the instructions on the dialog box that displays.

Change a bus ownership type

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

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

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

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

Create an I/O Pool using the Hardware Management Console

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

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

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

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

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

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

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

Switch access to backup server

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

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

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

1. The switchover may have already made the current node unavailable. If not, then from the current node, “Make a disk pool unavailable” on page 43 (vary off).
2. Switch the independent disk pool to the first backup cluster node by performing a switchover in the cluster. See *Perform a switchover* for details.

Synchronize user profile name, UID, and GID

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

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

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

Change the server takeover IP address

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

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

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

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

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

6. Change the server takeover IP address in the **IP address** field.

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

Manage independent disk pools with geographic mirroring

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

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

Some of the tasks that you might need to perform include:

Set the threshold of a disk pool

Work with device parity protection

Work with mirrored protection

Suspend geographic mirroring

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

Note: When you resume geographic mirroring, a complete synchronization is required between the production and mirror copies. Synchronization can be a lengthy process.

To suspend geographic mirroring, follow these steps:

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

Resume geographic mirroring

If you suspend “Geographic mirroring” on page 15, you must resume geographic mirroring in order to reactivate mirroring between the production and mirrored copies again.

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

To resume geographic mirroring, follow these steps:

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

Detach mirror copy

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

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

To detach the mirror copy, follow these steps:

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

Before you make the detached mirror copy available, you should create a second, unique device description for the independent disk pool that differentiates it from the production copy. A separate device description for the mirror copy prevents two instances of the same database in the network. It will also simplify work done outside of iSeries Navigator. Use the detached mirror copy device description to make the detached mirror copy available. When using this other device description, you cannot use make available in iSeries Navigator; instead, you must use the VRYCFG command in the character-based interface.

Reattach mirror copy

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

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

To reattach the mirror copy, follow these steps:

1. In iSeries Navigator, expand **My Connections** (or your active environment).
2. Expand the server that owns the production copy of the disk pool to which you want to reattach the detached mirror copy.
3. Expand **Configuration and Service**.
4. Expand **Hardware**.
5. Expand **Disk Units**.
6. Expand **Disk Pools**.
7. Right-click the production copy of the **Disk Pool** you want to reattach and select **Geographic Mirroring> Reattach Mirror Copy**.

Change disk pool attributes

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

For more information about geographic mirroring attributes, see *How geographic mirroring works*.

To edit the disk pool attributes, follow these steps:

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

Deconfigure geographic mirroring

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

To deconfigure geographic mirroring, follow these steps:

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

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

- Remove the nodes associated with the mirror copy from the switchable hardware group recovery domain.
- Remove the site name and data port IP addresses from the remaining nodes in the cluster.

Messages for geographic mirroring

Geographic mirroring message descriptions and recoveries.

0x00010259

Description: Operation failed because the system did not find the mirror copy.

Recovery: Not all the nodes in the device domain responded. Make sure that clustering is active. If necessary, start clustering (STRCLUNOD). Try the request again. If the problem persists, contact your technical support provider.

0x0001025A

Description: Not all of the disk pools in the disk pool group are geographically mirrored

Recovery: If one disk pool in a disk pool group is geographically mirrored, all of the disk pools in the disk pool group must be geographically mirrored. Take one of the following actions: 1) Configure geographic mirroring for the disk pools which are not geographically mirrored. 2) Deconfigure geographic mirroring for the disk pools that are geographically mirrored.

0x00010265

Description: The detached mirrored copy is available.

Recovery: Make the detached mirrored copy unavailable and then try the reattach operation again.

0x00010380

Description: A disk unit is missing from the configuration of the mirror copy.

Recovery: Find or fix the missing disk unit in the mirror copy. Check the Product Activity Log (PAL) on destination node. Reclaim IOP cache storage.

0x00011210

Description: The proposed secondary disk pool for the disk pool group is not geographically mirrored

Recovery: If one disk pool in a disk pool group is geographically mirrored, all of the disk pools in the disk pool group must be geographically mirrored. You must configure geographic mirroring for the proposed secondary disk pool which is not geographically mirrored, either now or after completing this operation.

0x00011211

Description: Duplicate mirror copies exist.

Recovery: Check for locally mirrored disk units that may exist on two systems, Enterprise Storage Server[®] FlashCopy[®], or back level independent disk pool copies. See the Product Activity Log (PAL) on the mirror copy node for more information. Eliminate duplication and try the request again. If the problem persists, contact your technical support provider, or see iSeries and AS/400[®] Technical Support for information about IBM support and services.


Examples: Independent disk pool configurations

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

Dedicated independent disk pools

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

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



Scenario: Consolidate servers using switchable independent disk pools

Situation

Your company's network currently uses 30 small servers distributed within a single country, all in the same time zone, using the same language, and running the same release of the operating system and programming code. The amount of time and effort you spend maintaining the small systems and keeping them at the same operating system and application release levels is significant.

Objectives

To reduce the resource required to maintain and administer your servers, you want to consolidate by reducing the number of servers in your network.

The objectives of this scenario are as follows:

- To consolidate from 30 small servers to one larger server at a central location
- To maintain data independence for each geographic region

Details

None of the 30 small servers in your network require more than four disk units.

Prerequisites and assumptions

A potential consolidation answer for your network is logical partitioning (LPAR). However, in your scenario, consolidating the 30 locations with logical partitioning is not ideal because:

- The effort required to manage the partitions is approximately the same as managing 30 distributed systems.
- Each partition requires an IOP in order to support a load source for the partition. As a result, 30 IOPs are required for the consolidated system.
- Additional expansion units are required to hold the IOPs needed for the 30 partitions. Since each location uses only a few disk units, the expansion units might be nearly empty.

As a result, the LPAR solution is not justifiable from an economic point of view for your scenario.

A better solution for your particular scenario is to use switchable independent disk pools to provide server consolidation. By creating one switchable independent disk pool for each of the 30 branch offices, you will be able to reduce the number of IOPs from 30 to 7, while requiring just two expansion units. This is an economically attractive alternative.

Design

To understand how to use switchable independent disk pools, see “Create a switchable independent disk pool” on page 28. In addition to the planning and configuration steps for implementing switchable independent disk pools, the following strategies can be used to ensure that your users at the respective branch offices can seamlessly access data:

- To ensure that users receive access to the correct set of data, your run-time environment can be changed to make sure that users from different branch offices connect to their data in the corresponding independent disk pool. This can be accomplished through a simple adjustment to user profiles and to the job descriptions that are specified by user profiles.

All user profiles from a particular branch office will use one job description. The job description will specify the independent disk pool that contains the user’s data, and create the library list that each job will use. With these simple changes, the task of getting each user to the correct set of data is completed.

- Another run-time problem to be addressed is the resolution of duplicate subsystems and job queues. Each branch office uses a cloned subsystem description to run batch jobs. Each of the subsystems uses job queues that have the same name on each of the branch office subsystems. If a single subsystem and a single set of job queues are used in the consolidated environment, jobs submitted by users from different branch offices will all be placed on the same set of queues and initiated by a single subsystem. This results in work flow that is inconsistent with the run-time environment of the distributed systems.

To resolve this problem, the subsystems will be given unique names. Then, a command to start all of the subsystems will be added to the startup program. Finally, each of the job queues used by the subsystem will be moved into a library that is unique to each of the job descriptions that are used by the branch offices. As a result, any application that submits a job will require no changes in order to submit batch jobs to its unique queue.

Examples: Switchable independent disk pools

In a multisystem environment, an independent disk pool can be switched between servers in a cluster. A switchable independent disk pool is a set of disk units that you can switch between servers so that each server can access the data. Only one system can access the data at a time.

Switchable independent disk pools can reside on one of two types of switchable hardware devices:

External expansion unit

The switchable device can be an external expansion unit connected to the clustered servers on the same high-speed link (HSL) loop.

Input/output processor (IOP) in a logical partition

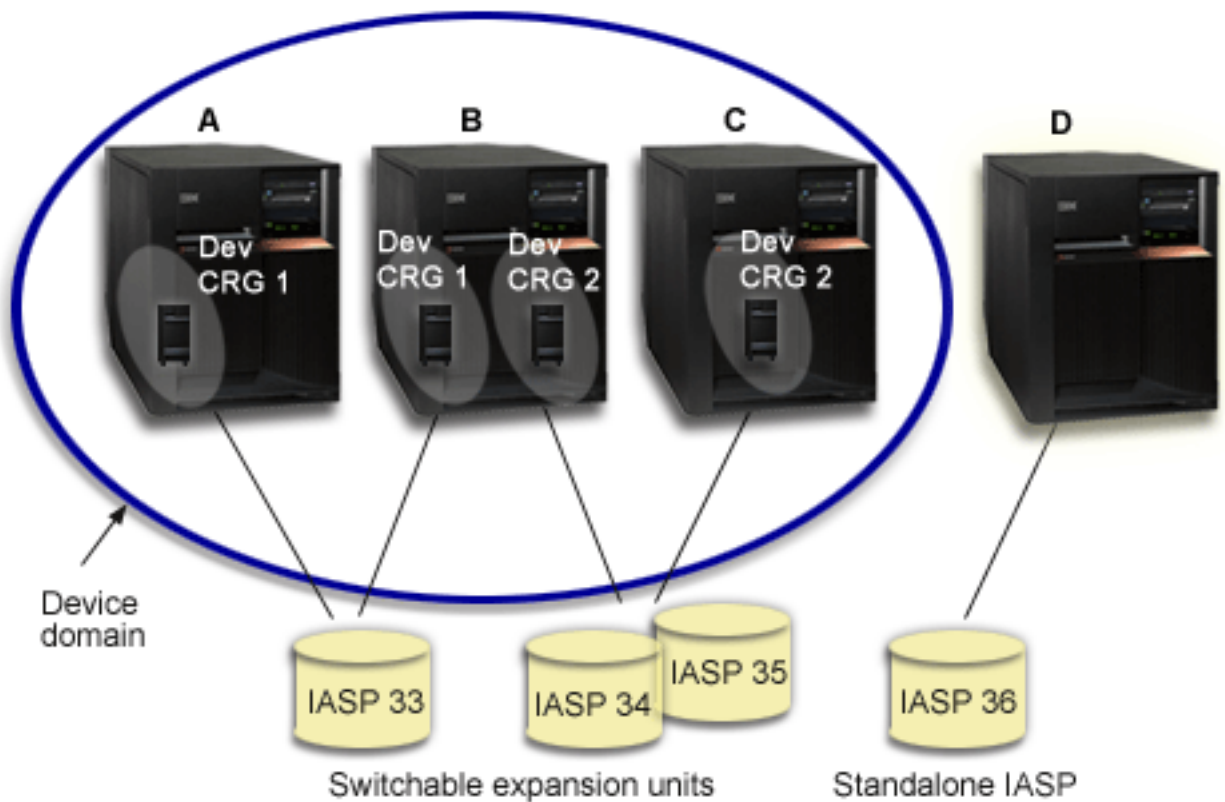
In an LPAR environment, the switchable device can be an IOP on the bus shared by the partitions or in an I/O pool.

The entity that switches is actually the expansion unit or the IOP containing the independent disk pool. When an expansion unit or IOP is switched, all of the hardware attached to the switchable entity is moved to the backup system.

The following example configurations illustrate some typical switchable independent disk pools implementations:


Example: Switchable expansion unit

In this example, the following figure shows a cluster consisting of four nodes. Nodes named A, B, and C are defined to be in the same device domain. There are two switchable expansion units - one contains IASP33 and the other contains IASP34 and IASP35. The expansion unit containing IASP33 is on an HSL loop that also contains nodes A and B. This first expansion unit can be switched between nodes A and B. The expansion unit containing IASP34 and IASP35 can be on another HSL loop that also contains nodes B and C. This second expansion unit can be switched between nodes B and C. Node D is contained in the cluster, but is not a member of the device domain and therefore can only access IASP36, a stand-alone (or dedicated) independent disk pool.



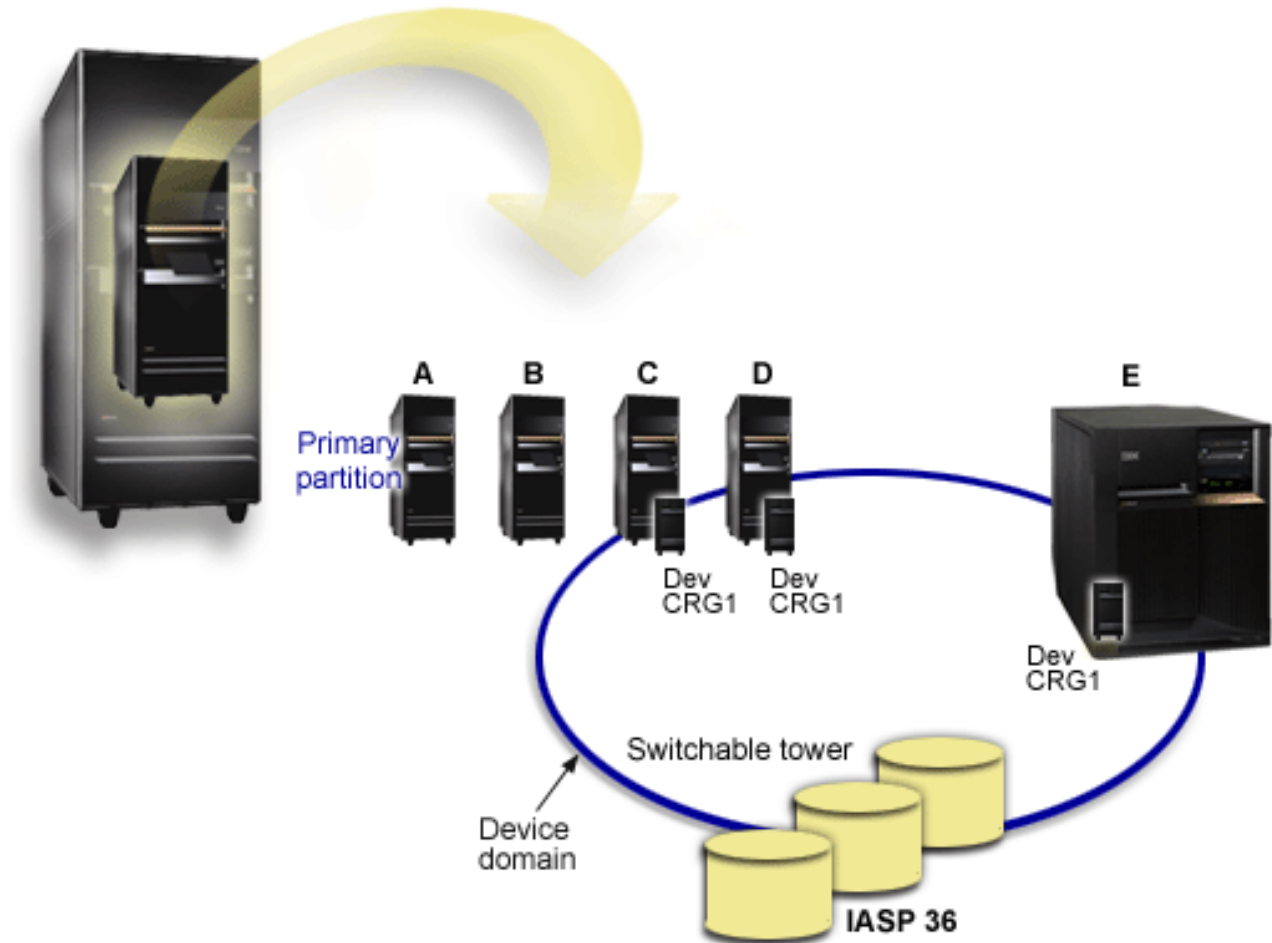
Example: Switchable IOP with logical partitions

In this logical partition example, the following figure shows a cluster consisting of four logical partitions on a single iSeries server. All four nodes belong to the same device domain. IASP36 is composed of disk units accessible through IOP Y. IOP Y is on the shared bus so it can be switched between all of the nodes in the cluster: A, B, C, and D. When the IOP is switched, everything that is physically connected to that IOP is also moved to the new primary node.



Example: Switchable expansion unit with logical partitions

The example, shown in the figure below, depicts a combination of the previous two examples. IASP36 is composed of disk units contained in a switchable expansion unit. The expansion unit is on the same HSL loop as two systems, one of which is made up of four logical partitions. Assume that nodes C and D and the second server, node E, are defined to be in the same device domain, and that the independent disk pool can be switched between those three nodes.



Scenario: Consolidate servers using switchable independent disk pools

Situation

Your company's network currently uses 30 small servers distributed within a single country, all in the same time zone, using the same language, and running the same release of the operating system and programming code. The amount of time and effort you spend maintaining the small systems and keeping them at the same operating system and application release levels is significant.

Objectives

To reduce the resource required to maintain and administer your servers, you want to consolidate by reducing the number of servers in your network.

The objectives of this scenario are as follows:

- To consolidate from 30 small servers to one larger server at a central location
- To maintain data independence for each geographic region

Details

None of the 30 small servers in your network require more than four disk units.

Prerequisites and assumptions

A potential consolidation answer for your network is logical partitioning (LPAR). However, in your scenario, consolidating the 30 locations with logical partitioning is not ideal because:

- The effort required to manage the partitions is approximately the same as managing 30 distributed systems.
- Each partition requires an IOP in order to support a load source for the partition. As a result, 30 IOPs are required for the consolidated system.
- Additional expansion units are required to hold the IOPs needed for the 30 partitions. Since each location uses only a few disk units, the expansion units might be nearly empty.

As a result, the LPAR solution is not justifiable from an economic point of view for your scenario.

A better solution for your particular scenario is to use switchable independent disk pools to provide server consolidation. By creating one switchable independent disk pool for each of the 30 branch offices, you will be able to reduce the number of IOPs from 30 to 7, while requiring just two expansion units. This is an economically attractive alternative.

Design

To understand how to use switchable independent disk pools, see “Create a switchable independent disk pool” on page 28. In addition to the planning and configuration steps for implementing switchable independent disk pools, the following strategies can be used to ensure that your users at the respective branch offices can seamlessly access data:

- To ensure that users receive access to the correct set of data, your run-time environment can be changed to make sure that users from different branch offices connect to their data in the corresponding independent disk pool. This can be accomplished through a simple adjustment to user profiles and to the job descriptions that are specified by user profiles.

All user profiles from a particular branch office will use one job description. The job description will specify the independent disk pool that contains the user’s data, and create the library list that each job will use. With these simple changes, the task of getting each user to the correct set of data is completed.

- Another run-time problem to be addressed is the resolution of duplicate subsystems and job queues. Each branch office uses a cloned subsystem description to run batch jobs. Each of the subsystems uses job queues that have the same name on each of the branch office subsystems. If a single subsystem and a single set of job queues are used in the consolidated environment, jobs submitted by users from different branch offices will all be placed on the same set of queues and initiated by a single subsystem. This results in work flow that is inconsistent with the run-time environment of the distributed systems.

To resolve this problem, the subsystems will be given unique names. Then, a command to start all of the subsystems will be added to the startup program. Finally, each of the job queues used by the subsystem will be moved into a library that is unique to each of the job descriptions that are used by the branch offices. As a result, any application that submits a job will require no changes in order to submit batch jobs to its unique queue.

Example: Independent disk pools with geographic mirroring

The following example shows one way that geographic mirroring can be configured. Node A and Node B are located in New York City. Node C and Node D are located in Boston. All four nodes are configured in the same recovery domain. The production copy can be switched between nodes A and B. The mirror copy can be switched between nodes C and D. Because all of the nodes are in the same recovery domain, the source system in New York can also exchange roles with the target system in Boston, allowing Boston to host the production copy.



This company has defined the following roles for the nodes in the recovery domain:

Node	Role
Node A	Primary
Node B	Backup 1
Node C	Backup 2
Node D	Backup 3

In the event of a natural disaster in New York, Node C in Boston becomes the primary node by upgrading its mirror copy to a production copy. Node C becomes the source system for geographic mirroring, although geographic mirroring will be suspended because there is no target node because of the natural disaster in New York. When the New York site recovers, Node A becomes a backup node and its previous production copy becomes the mirror copy.

Code disclaimer information

This document contains programming examples.

SUBJECT TO ANY STATUTORY WARRANTIES WHICH CANNOT BE EXCLUDED, IBM, ITS PROGRAM DEVELOPERS AND SUPPLIERS MAKE NO WARRANTIES OR CONDITIONS EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OR CONDITIONS OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT, REGARDING THE PROGRAM OR TECHNICAL SUPPORT, IF ANY.

UNDER NO CIRCUMSTANCES IS IBM, ITS PROGRAM DEVELOPERS OR SUPPLIERS LIABLE FOR ANY OF THE FOLLOWING, EVEN IF INFORMED OF THEIR POSSIBILITY:

1. LOSS OF, OR DAMAGE TO, DATA;
2. SPECIAL, INCIDENTAL, OR INDIRECT DAMAGES, OR FOR ANY ECONOMIC CONSEQUENTIAL DAMAGES; OR
3. LOST PROFITS, BUSINESS, REVENUE, GOODWILL, OR ANTICIPATED SAVINGS.

SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO SOME OR ALL OF THE ABOVE LIMITATIONS OR EXCLUSIONS MAY NOT APPLY TO YOU.

Frequently asked questions

Here is a list of independent disk pool questions and answers. If you have a question that is not on this page, please contact us.

General

1. How do independent disk pools work?
2. How can independent disk pools be implemented in my environment?
3. How should I structure my independent disk pools?
4. What is a disk pool group?
5. What is geographic mirroring?

iSeries Navigator graphical user interface

1. How do I access the iSeries Navigator disk management function?
2. What is the difference between the disk management functions in iSeries Navigator and the character-based command interface?
3. How do I access the disk management function when the system is in dedicated service tools (DST) mode?
4. What is the service tools server?
5. Why does the data I see in iSeries Navigator appear to be out of date?
6. Why can't I connect to the service tools server after I add the service table entry?

Configuring

1. How do I create a new disk pool or independent disk pool?
2. How do I create a disk pool group?
3. How do I configure geographic mirroring?

Performance

1. Why is performance slow?
2. Why is performance slow for geographic mirroring?

Troubleshooting

1. Why do no disk units appear as eligible to be added to my disk pool?
2. Why doesn't the device description get deleted when I delete the disk pool?

3. Why do I get a message saying the device description is already created?
4. Why does the primary or secondary disk pool I try to create appear to be a UDFS disk pool?
5. Why do I get a message that says my disk pool is not the right type when I try to create a library in it?
6. What is a CPF9898 message, and how can it be fixed?

General

How do independent disk pools work?

The key characteristic of an independent disk pool is its ability to be, of course, independent of the rest of the storage on a server. It is independent because the data in the independent disk pool is self-contained. This means that all of the necessary system information associated with the data resides within the independent disk pool. See “How independent disk pools work” on page 6 for details.

[Back to questions](#)

How can independent disk pools be implemented in my environment?

There are two basic environments in which you can take advantage of independent disk pools: a multi-system environment managed by an iSeries cluster, and a single-system environment with a single iSeries server. See “Switchable and stand-alone independent disk pools” on page 13 for details.

[Back to questions](#)

How should I structure my independent disk pools?

IBM provides some recommendations for structuring and populating your independent disk pools. See “Recommended structure for independent disk pools” on page 12 for details.

[Back to questions](#)

What is a disk pool group?

A disk pool group is made up of a primary disk pool and zero or more secondary disk pools. Each disk pool is independent in regard to data storage, but in the disk pool group they combine to act as one entity. See “Disk pool groups” on page 14 for details.

[Back to questions](#)

What is geographic mirroring?

Geographic mirroring is a function that generates a mirror copy of an independent disk pool on a system that is (optionally) geographically distant from the originating site for availability or protection purposes. See “Geographic mirroring” on page 15 for details.

[Back to questions](#)

iSeries Navigator graphical user interface

How do I access the iSeries Navigator disk management function?

Before you can access disk management functions in iSeries Navigator, you must complete some setup tasks. See [Access disk management functions](#) for details.

Back to questions

What is the difference between the disk management functions in iSeries Navigator and in the character-based (command) interface?

Support for many independent disk pool tasks are only available through iSeries Navigator. Almost all disk management functions that are available from the system service tools (SST) mode are available through iSeries Navigator. A number of disk management functions that are only available from the dedicated service tools (DST) mode are also available.

Back to questions

How do I access the disk management function when the system is in dedicated service tools (DST) mode?

Beginning with V5R1, the Disk Units folder in iSeries Navigator is available when the system is in dedicated service tools (DST) mode.

Back to questions

What is the service tools server?

The service tools server allows you to use your PC to perform service tools functions through TCP/IP. Before you attempt to use any disk management functions, you must configure the service tools server. See "Set up communication" on page 22 for details.

Back to questions

Why does the data I see in the iSeries Navigator window appear to be out of date?

The disk management function in iSeries Navigator caches information, and therefore needs to be refreshed to have the most current data visible. After you make a configuration change, iSeries Navigator should refresh itself. If it does not, however, you can manually refresh it by clicking the Refresh button on the iSeries Navigator toolbar. You can also set iSeries Navigator to refresh periodically. Depending on the size of your server, however, you might not want to do this. Disk unit configuration data tends to be fairly static and does not need to be refreshed often. If your system is large, it can take a significant amount of time to download all information.

Back to questions

Why can't I connect to the service tools server after I add the service table entry?

The Add Service Table Entry (ADDSRVTBLE) command is case sensitive. In particular, it is important to ensure that the Protocol = 'tcp', and not 'TCP'. To ensure this is the case, use the Work with Service Table Entry (WRKSRVTBLE) command, and check the as-sts server field. Make sure that TCP is lowercase. If it is not, then remove the entry, and re-create it by issuing the following command exactly as shown:

```
ADDSRVTBLE SERVICE('as-sts') PORT(3000) PROTOCOL('tcp') TEXT('Service Tools Server')  
ALIAS('AS-STs')
```

Back to questions

Configuring

How do I create a new independent disk pool?

You can create an independent disk pool in a clustered, multisystem environment or on a single system. See the following topics for details:

- “Create a switchable independent disk pool” on page 28
- “Create a dedicated independent disk pool” on page 27

[Back to questions](#)

How do I create a disk pool group?

See “Create a new disk pool group” on page 32 for details.

[Back to questions](#)

How do I configure geographic mirroring?

You can configure geographic mirroring for independent disk pools that are dedicated or switchable between systems. See the following topics for details:

- “Configure geographic mirroring with dedicated independent disk pools” on page 33
- “Configure geographic mirroring with dedicated independent disk pools” on page 33

[Back to questions](#)

Performance

Why is performance slow?

There are several factors that can influence performance. Make sure your PC’s TCP/IP settings are configured correctly. Specifically, make sure that you do not have an incorrect secondary gateway. If you do have a secondary gateway, remove it. This should provide a significant increase in performance. See Requirements for geographic mirroring for more detailed information.

[Back to questions](#)

Why is performance slow for geographic mirroring?

For geographic mirroring you should consider the distance that the independent disk pool is being mirrored. The type and number of communication lines as well as their bandwidth have an effect on performance. You can configure up to four TCP/IP communication interfaces on multiple adapters. You should consider configuring multiple communication lines to allow for the highest performance. The volume of disk unit writes that your applications require also plays a role in the performance of your geographically mirrored independent disk pool.

[Back to questions](#)

Troubleshooting

Why do no disk units appear as eligible to be added to my disk pool?

There are a number of possible reasons for this. First, you must have an unconfigured disk unit to add. If the disk pool is protected, you can only add parity disks, or disks in pairs, so that they can be mirrored.

If your system is in a clustered environment, disk unit eligibility is more complex. Each disk unit is assigned a rank, which indicates its eligibility to be added to a particular disk pool. If the rank of the disk unit is above 300, then the disk is ineligible. A complete list of the ranks, and what they mean, is available in the disk management online help.

[Back to questions](#)

Why doesn't the device description get deleted when I delete the disk pool?

Because the device description does not always get created by the disk management function, it might not be deleted when the disk pool gets deleted. You need to manually delete it using the Delete Device Description (DLTDEVVD) command.

[Back to questions](#)

Why do I get a message saying the device description is already created?

When you create a new independent disk pool, an attempt is made to create an associated device description. If a device description of the same name as the disk pool already exists, you will receive a warning message, and the existing device description will not be changed. Most of the time, this is not a problem. However, if the device description's name and associated resource do not match, this becomes a problem, and this is why you receive the warning message.

[Back to questions](#)

Why does the primary or secondary disk pool that I try to create appear to be a UDFS disk pool?

If iSeries Navigator failed or was closed while the disk pool was being created, you might need to "Convert UDFS disk pools" on page 32 to a primary or secondary disk pool.

[Back to questions](#)

Why do I get a message that says my disk pool is not the right type when I try to create a library in it?

Make sure that the disk pool you are trying to create a library in is a primary or secondary disk pool, not a UDFS disk pool. If the disk pool is a UDFS disk pool and you want to create a library in it, you need to "Convert UDFS disk pools" on page 32 to a primary or secondary disk pool.

| What is a CPF9898 message, and how can it be fixed?

- | When a copy from an original ASP is made and there is an attempt to vary on the copy on the system
- | that had used the original ASP, message CPF9898 will appear. Before the system will accept the copy, an
- | IPL must be performed.

[Back to questions](#)

Appendix. Notices

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

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

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

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

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

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

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

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

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

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

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

IBM Corporation

Software Interoperability Coordinator, Department 49XA
3605 Highway 52 N
Rochester, MN 55901
U.S.A.

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

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

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

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

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

COPYRIGHT LICENSE:

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

SUBJECT TO ANY STATUTORY WARRANTIES WHICH CANNOT BE EXCLUDED, IBM, ITS PROGRAM DEVELOPERS AND SUPPLIERS MAKE NO WARRANTIES OR CONDITIONS EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OR CONDITIONS OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT, REGARDING THE PROGRAM OR TECHNICAL SUPPORT, IF ANY.

UNDER NO CIRCUMSTANCES IS IBM, ITS PROGRAM DEVELOPERS OR SUPPLIERS LIABLE FOR ANY OF THE FOLLOWING, EVEN IF INFORMED OF THEIR POSSIBILITY:

1. LOSS OF, OR DAMAGE TO, DATA;
2. SPECIAL, INCIDENTAL, OR INDIRECT DAMAGES, OR FOR ANY ECONOMIC CONSEQUENTIAL DAMAGES; OR
3. LOST PROFITS, BUSINESS, REVENUE, GOODWILL, OR ANTICIPATED SAVINGS.

SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO SOME OR ALL OF THE ABOVE LIMITATIONS OR EXCLUSIONS MAY NOT APPLY TO YOU.

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

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

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

Programming Interface Information

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

Trademarks

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

AS/400
Enterprise Storage Server
eServer
FlashCopy
i5/OS
IBM
iSeries
JavaScript
Operating System/400
OS/400
Redbooks
TotalStorage

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

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

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

Terms and conditions for downloading and printing information

Permissions for the use of the information you have selected for download are granted subject to the following terms and conditions and your indication of acceptance thereof.

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

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

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

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

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations. IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THIS INFORMATION. THE INFORMATION IS PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

All material copyrighted by IBM Corporation.

By downloading or printing information from this site, you have indicated your agreement with these terms and conditions.



Printed in USA