iSeries Availability and Independent Auxiliary Storage Pools

Agenda

Independent Disk Pools

Cluster Resource Services

Journal enhancements

HTTP

Windows & xSeries Integration

Backup & Recovery

BRMS

Tivoli Storage Manager

Notes: Agenda

This foil lists the topics covered in this presentation.

In this presentation Independent Disk Pools (Independent Auxiliary Storage Pools (IASPs) are also discussed from a server consolidation viewpoint. This is because V5R2 IASP support provides functions than enable having multiple "name spaces" (commonly thought of as "multiple databases") on the same system or, in an LPAR environment, within the same partition.

In V5R1 IASPs became available as tools for increasing system availability. In V5R1 IASPs did not support objects stored in the iSeries file system QSYS.LIB. With V5R2 most QSYS.LIB objects, including database objects, such as tables, views and indexes can be supported in an IASP.

The Database presentation has additional details for using an IASP to support multiple databases on the same system or partition

This enables V5R2 IASPs to be considered both in an availability environment as well as a server consolidation environment.

For more detailed assistance in planning and implementing the use of IASPs in either a server consolidation or increased availability environment, we recommend contacting the iSeries Technology Center in Rochester, Minnesota, USA at:

http://www.ibm.com/eserver/iseries/service/itc

Independent Disk Storage Pools

Independent Disk Pools

Disk pools

- Support for library-based objects
- Switchable and non-switchable

Disk Pool Groups

Multiple databases

Object Management

Planning

Backup, Recovery of IASPs

Advantages of using IASPs

- and other considerations

Scenarios

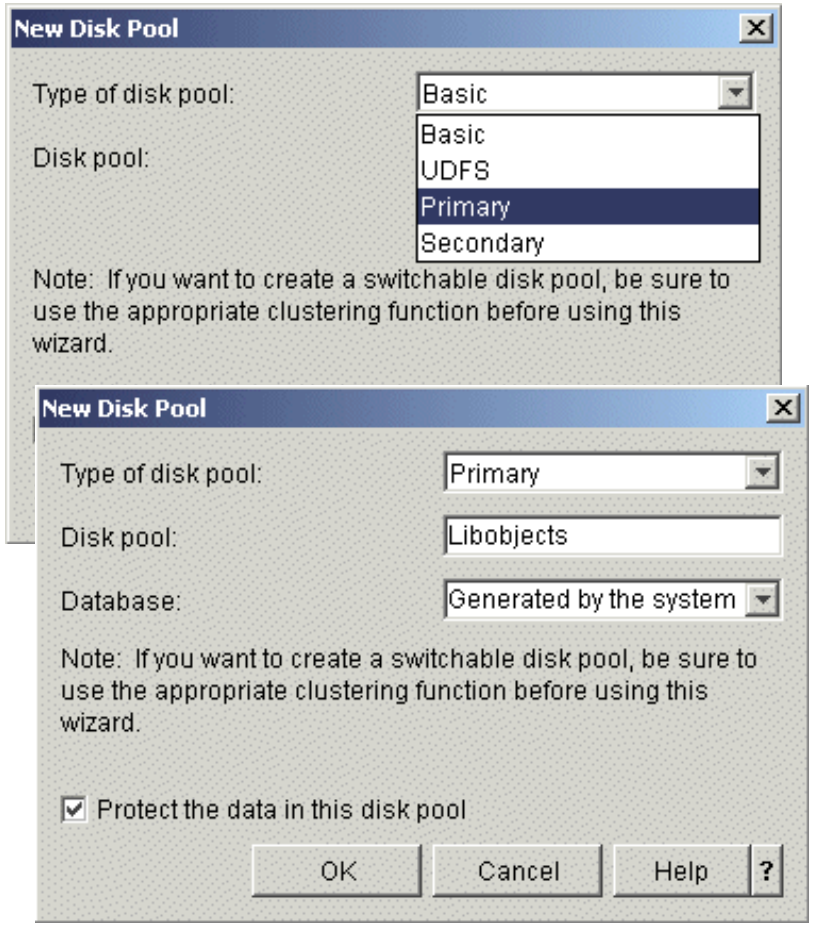

Notes: Independent Disk Pools

Independent disk pools provide the ability to group together storage that can be taken off-line or brought online independent of system data or other unrelated data. An independent disk pool can be either:

- Privately connected to a single system. Also known as stand-alone or primary IASPs
- Switched among multiple systems or partitions in a clustered environment

Clearly, this is quite a departure from the way in which auxiliary storage (disk) has been regarded prior to V5R1. Until then, all iSeries disks were considered to be owned and usable by a single system only. Enhancements made to OS/400 in V5R1 and again in V5R2 make the use of independent disk pools a very attractive option for many customers looking for higher levels of availability or server consolidation.

Note: In this presentation we interchangeably use the acronym IASP to be identical to the term "independent disk pools." As yo u will see in this presentation, iSeries Navigator uses the term "disk pools" to refer to Auxiliary Storage Pools (ASPs) and a prefix adjective that uniquely classifies the ASP to be dependent or independent (UDFS, Primary, Secondary).

Primary and Secondary can contain OS/400 libraries and library objects. UDFS corresponds to the V5R1 support which does not include OS/400 library-based objects. The next foil has more information.

Disk Pools

System ASP

- ASP# 1
- OS/400

Basic

- ASP# 2-32
- Also known as User or Dependent ASPs

Independent (IASP)

- ASP# 33-255
- User Defined File System (UDFS)
  - No QSYS.LIB objects. Can be converted to Primary or Secondary pools for library based object support
- Primary (QSYS.LIB objects)
  - Secondary

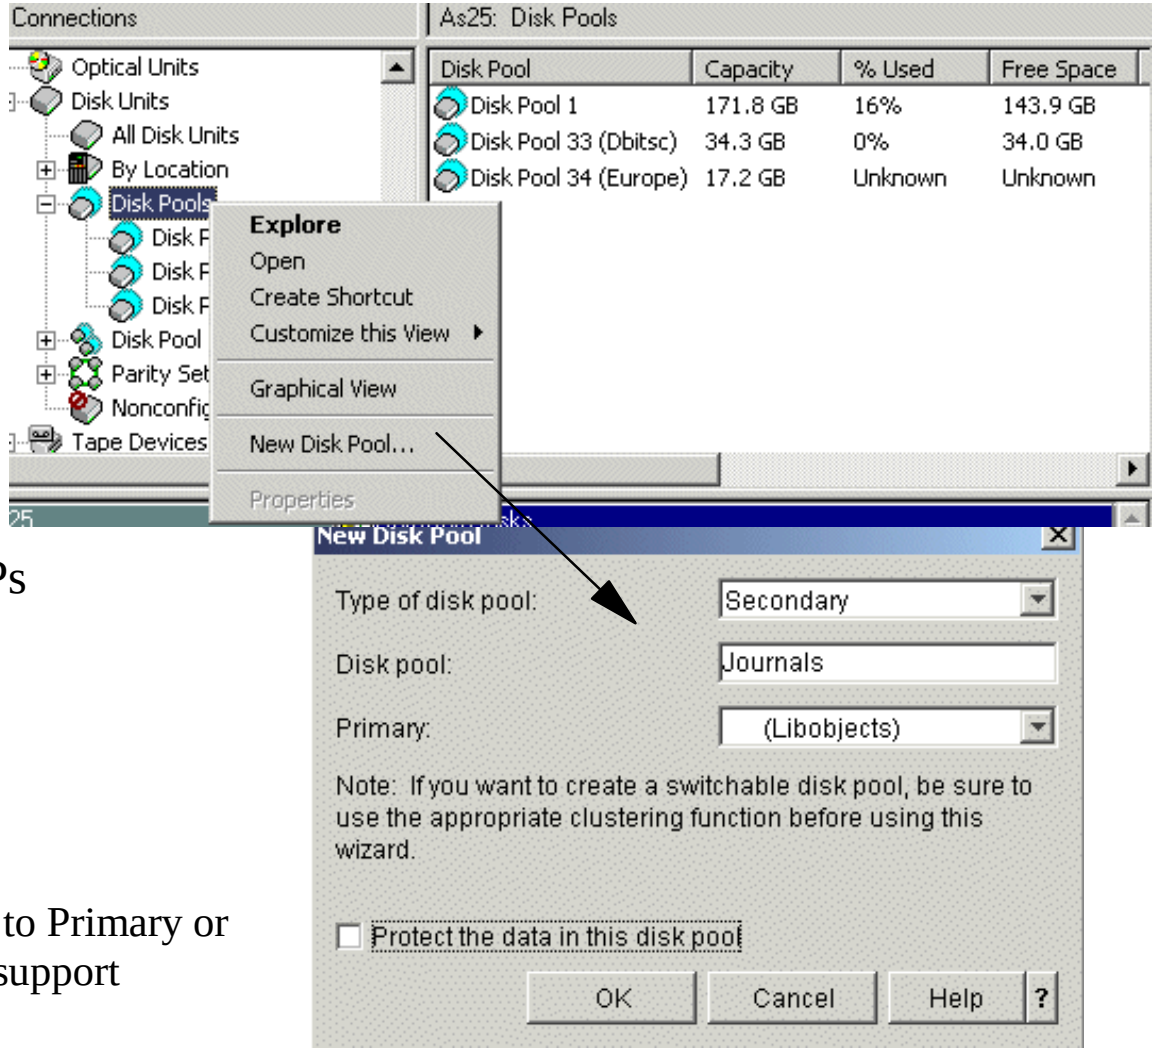

Notes: Disk Pools

A set of one or more storage units selected from all available disk units. Disk pools (also known as auxiliary storage pools) provide a means of placing certain objects on specific disk units to prevent the loss of data due to failures of disk units in other disk pools.

When independent disk pools were introduced in V5R1, they supported user-defined file systems (UDFS) only. Support for library-based (QSYS.LIB) objects has been added in V5R2. You can now create as many as 223 independent disk pools. Previous releases only supported 67 independent disk pools. In V5R1 independent disk pools were numbered from 33-99. That range has been expanded to 33-255 at V5R2. See "Supported and unsupported OS/400 object types" in the iSeries Information Center for details. There are three types of disk pools, also known as auxiliary storage pools (ASPs):

- **System**
  The system automatically creates the system disk pool (Disk Pool 1) which contains disk unit 1 and all other configured disks that are not assigned to a basic or independent disk pool. The system disk pool contains all system objects for the OS/400 licensed program and all user objects that are not assigned to a basic or independent disk pool. You can see the system ASP as Disk Pool 1 in foil.
- **Basic**
  Prior to V5R2, Auxiliary Storage Pools 2 to 32 were known as User ASPs. The function of these has not changed but they are now often referred to Basic User ASPs or Basic Disk Pools in V5R2. Prior to V5R2 they were referred to as Dependent Pools.
- **Independent**
  A disk pool that contains objects, the directories or libraries that contain the objects, and other object attributes such as authorization and ownership attributes. 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 one iSeries server and another iSeries server in a clustered environment.

  An independent disk pool that is not associated with a cluster resource group is referred to in OS/400 application programming interfaces (APIs) as a private disk pool or non-switchable IASP. Independent disk pools can also function in conjunction with other independent disk pools in a disk *pool group.*

In the upper right window pane you can see "statistics" are known for the system ASP and IASP 33 (named Dbitsc). IASP Dbitsc has been varied on to the local system so the system "knows the disk configuration is active. IASP disk pool 34 (Europe) % Used and amount of Free Space. statistics are not known because this

Notes: Disk Pools 2

Three type of IASPs can be defined in V5R2:

- **User-defined file system (UDFS)**
  An independent disk pool that contains only user-defined file systems. A UDFS disk pool cannot be a member of a disk pool group unless it is converted to a primary or secondary disk pool. An UDFS IASP is generally the equivalent to the only IASP support in V5R1.
- **Primary**
  An independent disk pool that defines a collection of directories and libraries and may have other secondary disk pools associated with it. A primary disk pool also defines a database for itself and other disk pools that may be added in its disk pool group. Primary disk pools can only be implemented on V5R2 or later of OS/400.
- **Secondary**
  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. A secondary pool is optional. Though not practical, you can , in theory, have as many as 200 or more secondary pools. Additional details are provided in the following foils.

Disk Pool Groups

Combine iASPs to one entity

Made up of:

- a primary disk pool
- zero or more secondary disk pools

Logically connects disk pools

- Vary them On/Off together
- Switch them together

Share the same database

- Similar as System ASP and basic ASPs
- For example:
  - Primary independent ASP for libraries and database files
  - Secondary independent ASP for journals and journal receivers

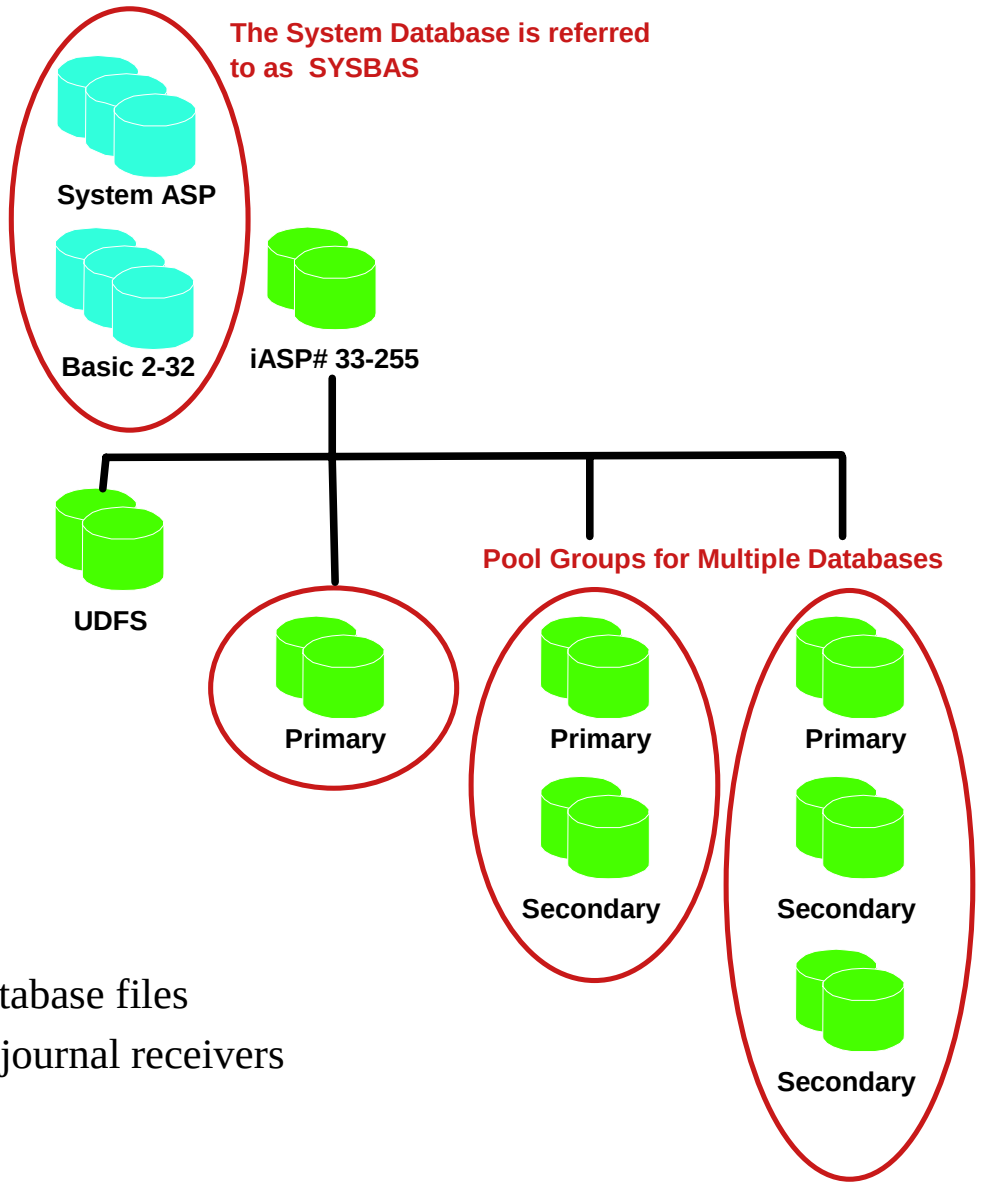

Notes: Disk Pool Groups

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. The primary and secondary disk pools also share the same database.

An example of a practical use for a secondary IASP disk pool group would be for journals and journal receivers. The primary disk pool could contain the libraries and objects to be journaled, while the secondary disk pools could contain the associated journal receivers. The journals and journal receivers would remain separate for maximum availability and performance, but they would function together in the disk pool group. Disk pool groups can only be implemented on V5R2 or later of OS/400.

This is similar to the system ASP (disk pool 1) and any base pool being treated as an entity sharing the local System Database and any associated journal objects. In V5R2 SYSBAS is the term assigned to refer to the system ASP and any Basic user ASP.

IFS Journaling

With V5R1 you could journal IFS objects but not the ones in an UDFS independent pool. This is because journal is a library based object. Now with V5R2, you can place IFS objects in a library capable pool (or convert an UDFS pool to a library based pool), journal them and place the jouranl receiver in an associated secondary independent pool. **Only library capable pools can be grouped together.**

Contrast basic and independent disk pools

When the server IPLs, all of the disk units configured in SYSBAS must be accounted for in order for the server to continue the IPL. Independent disk pools are not included in the IPL. When you vary on the independent disk pool the node then verifies that all disk units are present.

Structuring disk pool groups

The iSeries server supports up to 223 independent disk pools, any number of which can be primary, secondary, or 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 could be placed in a single disk pool group which consists of one primary disk pool and one secondary disk pool. Alternatively, you could create several disk pool groups, some with only a primary disk pool and some with one or more secondary disk pools.

There are other factors to consider when planning the placement of your data in disk pools. These are dependent on the way your iASPs are used.. Please refer to the iSeries information center for more details.

Multiple Databases

Independent disk pools

- Distinct Databases

Requires V5R2

- A single system can even be a 6xx or 7xx

SETASPGRP command

- Sets the namespace for the current thread

Reclaim Storage

- In parallel
- No restricted state required
  - except for SYSBAS

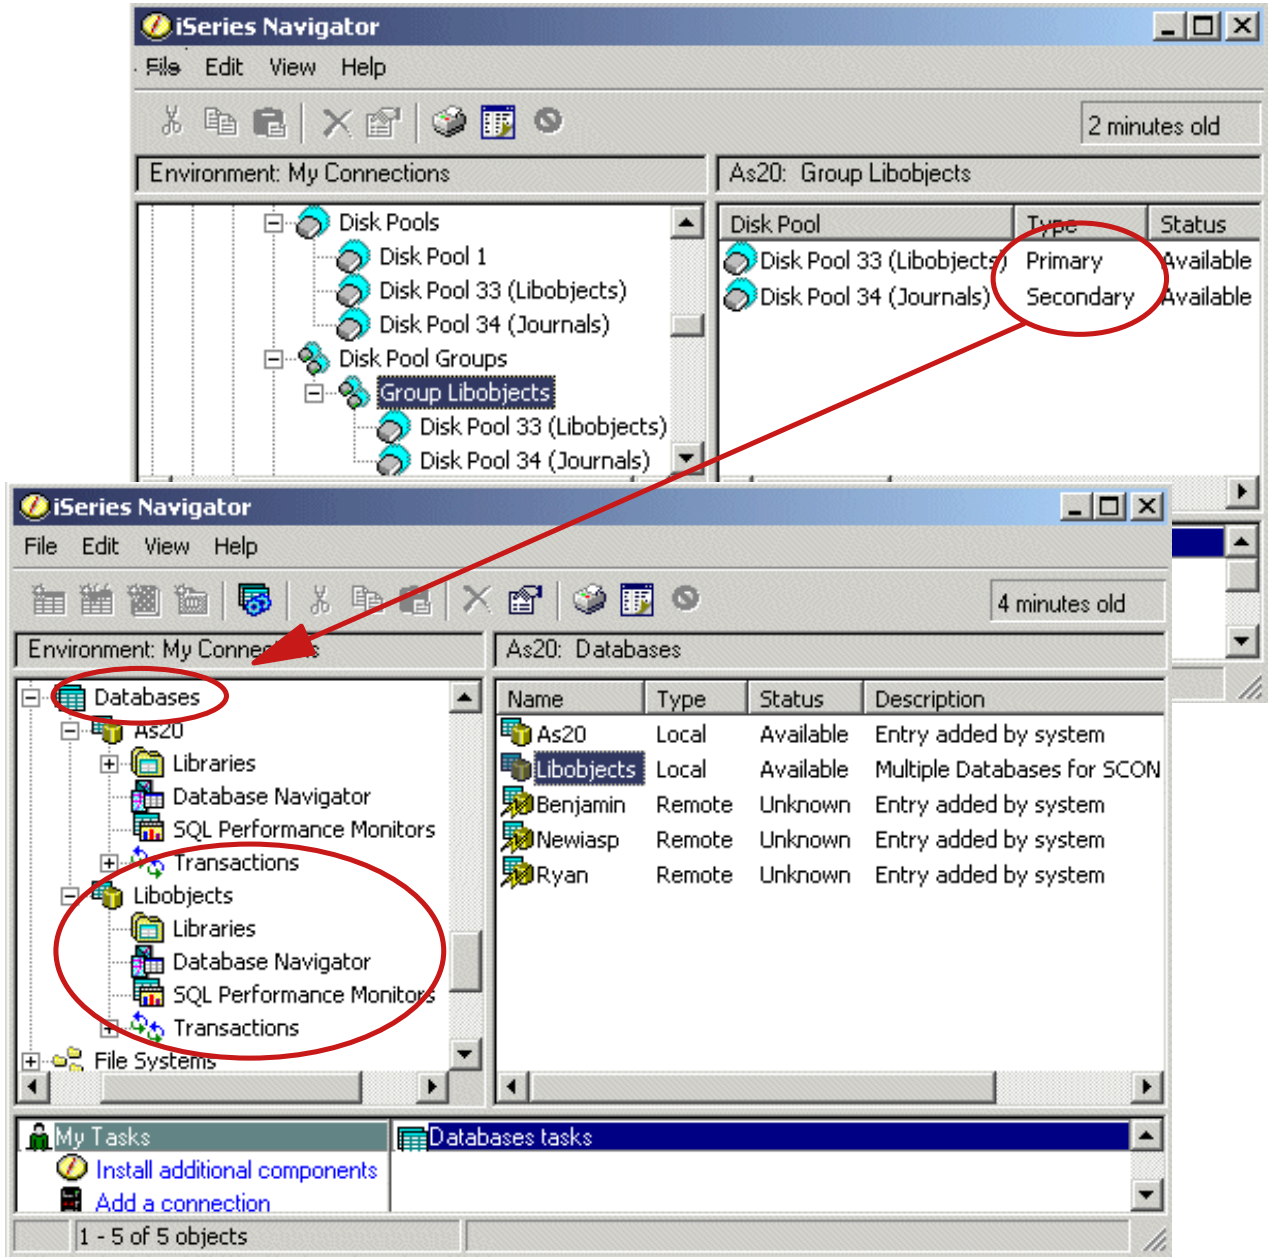

Notes: Multiple Databases

An independent disk pool can contain many kinds of objects in addition to database objects. By definition an inThis e foilWhen an independent disk pool is created, it will appear as a distinct user database on the server. This is separate from the system database, which was the only database available per system in previous releases.

Independent disk pools

Independent disk pools are used to set up user databases on the iSeries server. There are three types of independent disk pools: primary, secondary, and user-defined file system (UDFS). Databases are set up using primary independent disk pools. When a primary independent disk pool is configured, a new user database is defined that is separate from the system database. Once the independent disk pool is created, it appears as another database under the Database function of iSeries Navigator. The user database also includes any secondary disk pools that are associated with the primary disk pool. By default, the database and the independent disk pool will have the same name. You administer the user database with the same functions that you use for the system database. Multiple databases can be created on 6xx, 7xx and 8xx servers, as long as they are on V5R2.

SETASPGRP cmd

The Set Auxiliary Storage Pool Group (SETASPGRP) command sets the auxiliary storage pool (ASP) group for the current thread. Additionally, this command allows you to change the libraries in the library list for the current thread. If an ASP group had already been set, this command will remove the old ASP group from the current thread and set the specified ASP group for the current thread. Once the specified ASP group has been set for the current thread, all libraries in the independent ASPs in the ASP group are accessible and objects in those libraries can be referenced using regular library-qualified object name syntax. The libraries in the independent ASPs in the specified ASP group plus the libraries in the system ASP (ASP number 1) and basic user ASPs (ASP numbers 2-32) form the library name space for the thread. All libraries in the library list need to be in the new library name space or the library list is not changed and the new ASP group is not set.

Reclaim storage

With the introduction of IASPs comes the capability to run Reclaim Storage (RCLSTG) on an IASP while the rest of the system keeps running. This implies that multiple IASP RCLSTG processes can execute simultaneously. V5R1 functional changes to the RCLSTG command added to support IASPs are:

- *SYSBAS values If the *SYSBAS value is specified for the ASP device, the Reclaim Storage operation runs as it does on systems prior to V5R1. The reclaim operation is performed on the system and on basic user-defined ASPs. If the value specified is an ASP device name, then that ASP is reclaimed.
- Reclaim Storage for an ASP device (that is, an IASP) can be run without the system being in restricted state. Multiple ASP devices can be reclaimed in parallel.

Note: Reclaiming an auxiliary storage pool device requires that there are no active users of the ASP device that is the subject of the reclaim. The system will place the iASP in active state which is an intermediate state between unavailable and available.

Recommended structure for iASPs

For best possible single system protection and performance

- Place the majority of your application **database objects** into IASPs
- Place a minimal number of **database objects** in SYSBAS
  - Primarily operating system objects, licensed program product libraries, few user libraries
  - Depending on customer requirement, place all application programs in SYSBAS or in IASP along with database objects
- Recommended: Use disks in a separate RAID group for each IASP

Application data

- Is isolated from unrelated faults
- Can also be processed independently of other system activity

Vary on and switchover times are optimized with this structure

- Additional processing is required for database cross reference files
  - Information from the SYSBAS cross reference files is merged into the cross reference files in that independent ASP.
- Shorter IPL times
  - Remember that the database recovery phase is a major part of an IPL

Notes: Recommended structure for iASPs

The recommended usage structure for independent disk pools is to place the majority of your application database objects into independent disk pools and a minimal number of database 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 product libraries, and very 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:

- Since a database network cannot span an independent disk pool boundary, entire database networks (SYSBAS and an IASP) are contained within disk pool groups.
- Library names can be duplicated across disk pool groups, but not between a disk pool group and the libraries in SYSBAS. Applications can access different databases will little or no changes. As discussed later in Multiple Database Support Setup, you can see the ways impact to applications programs can be zero or minimal.

Although the above is the recommended structure, this does not exclude other configurations.

For example, you may start by migrating only a small portion of your data to a disk pool group and keeping the bulk of your data in SYSBAS. However, you should expect longer vary on and switchover times with this configuration since additional processing is required to merge database cross reference information ifrom SYSBAS with the database reference information in the independent disk pool group.

Merging database cross reference information

At the time an independent pool is made available, the information of the SYSBAS database cross reference information is merged into the database cross reference information of the independent pool. The information in the database cross reference information of the independent pool is not merged into the database cross reference information of the system ASP. Keeping the system ASP cross reference finformation small enables faster vary-on, faster switchover of the independent pools and shorter IPL times.

Object Management

Library naming rules

- Multiple occurrences in multiple iASPs
- Libraries in an IASP cannot be in SYSBAS

Object management

- Most commands that move objects are enabled (save, restore, ...)
- New Job Description parameter - set initial ASP Group
- Most commands default to the current name space (no ASPGRP parameters)
- SETASPGRP needs to be done
  - Or use the JOBD linked to the user profile

IASPs are not available after an IPL

- Must be explicitly varied on/made available by an operator

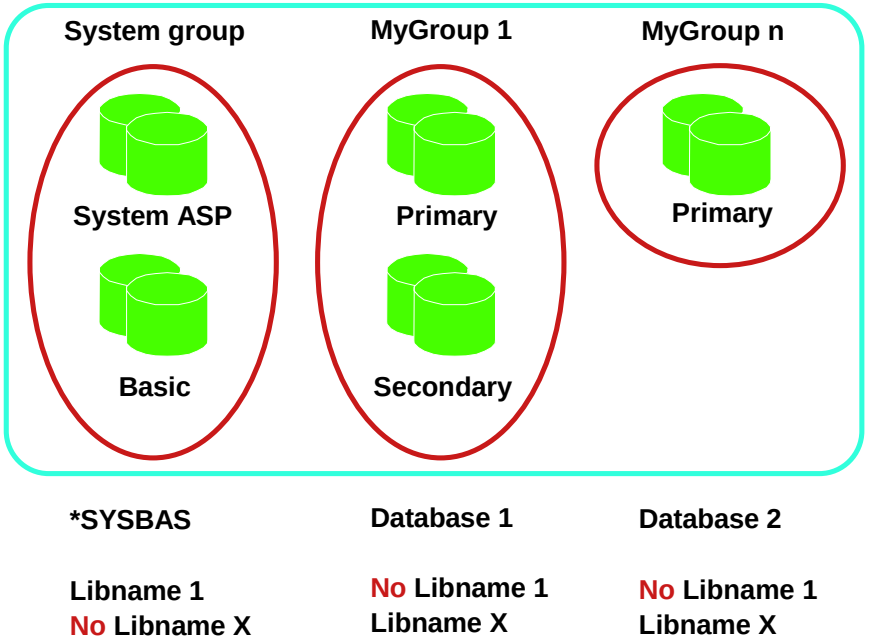

Notes: Object Management

Library names

With the independent ASPs it is now possible to have multiple databases. In a pool group there can reside only one primary independent ASP and zero or more secondary independent ASPs. The primary and secondary disk pools share the same database since they share the same database name space. So you could say that a pool group is equivalent to a database.

The same library name (or SQL collection name) can exist in two different independent disk pool groups though the same name cannot exist in the system disk pool. (ASP1).

Object management

Because the existence of an independent disk pool on a server implies that multiple databases will exist on a single server, identifying an object requires more complex planning than 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 don't necessarily uniquely identify an object. There will be times when you'll 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 don't necessarily have to match. A database namein DB2 UDB for iSeries can be up to 18 characters long, while an independent disk pool name can be up to 10 characters long. The SETASPGRP command points the current job/thread to any pool group.

iASPs are NOT automatically available after an IPL

Independent disk pools are available only after you have varied them on / made them available (iSeries Navigator term). They are not made available during a normal restart of your server, unless you include code in your Start Up Program to make them available. When you select to make a disk pool available, the disk pool goes through a process similar to that of a server restart (IPL). While this processing takes place, the disk pool cannot be used with the application.

Multiple Database Support - Setup Example: New IASP

5250 workstation

CRTDEVASP DEVD(MYCPY1)
RSRCNAME(MYCPY1) RDB(*GEN)

- Go to SST to add unconfigured disks to this IASP*

- **ADDRDBDIRE entry added by the system.**
- **In iSeries Navigator "make available" or in OS/400 "vary on;" populate with library with objects; the library content is then available to applications.**
- **Repeat for each IASP<-> Database required.**

iSeries Navigator Configuration and Service -> Hardware -> Disk pools - > New Pool *

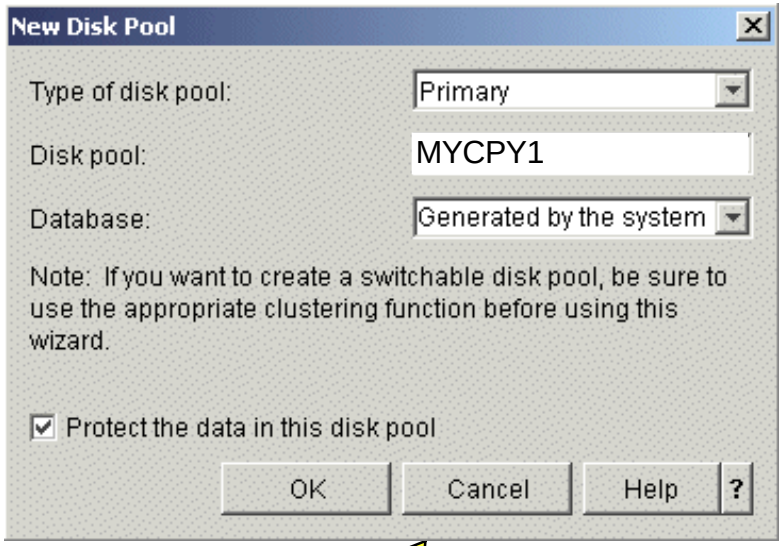

- Continue in iSeries Navigator to add unconfigured disks to this IASP

*** System Service Tools signon required**

Notes: Multiple Database Support - Setup: New IASP

This is an overview of the process to create an Independent ASP being used to support multiple databases on the same system. In our example, as shown on the following foils, we use 3 IASPs that are named Mycpy1, Mycpy2, and Mycpy3.

We show examples of creating a new IASP from either a 5250 command interface or an iSeries Navigator interface. We do not show the steps to add unconfigured disks to an IASP but do note authority to use System Service Tools (SST) is required to add the disks to the IASP. V5R1 introduced this Start Service Tools (SST) sign-on requirement for disk management. For a user not familiar with the SST functions, we recommend using the iSeries Navigator interface to create the new pool and assign disks to that pool.

In the example shown we are defaulting to the database and the IASP having the same name. When you define a new IASP and associated database name, the system implicitly invokes the Add Relational DataBase Entry (ADDRBDIRE) function.

Multiple Database Support - Setup Example

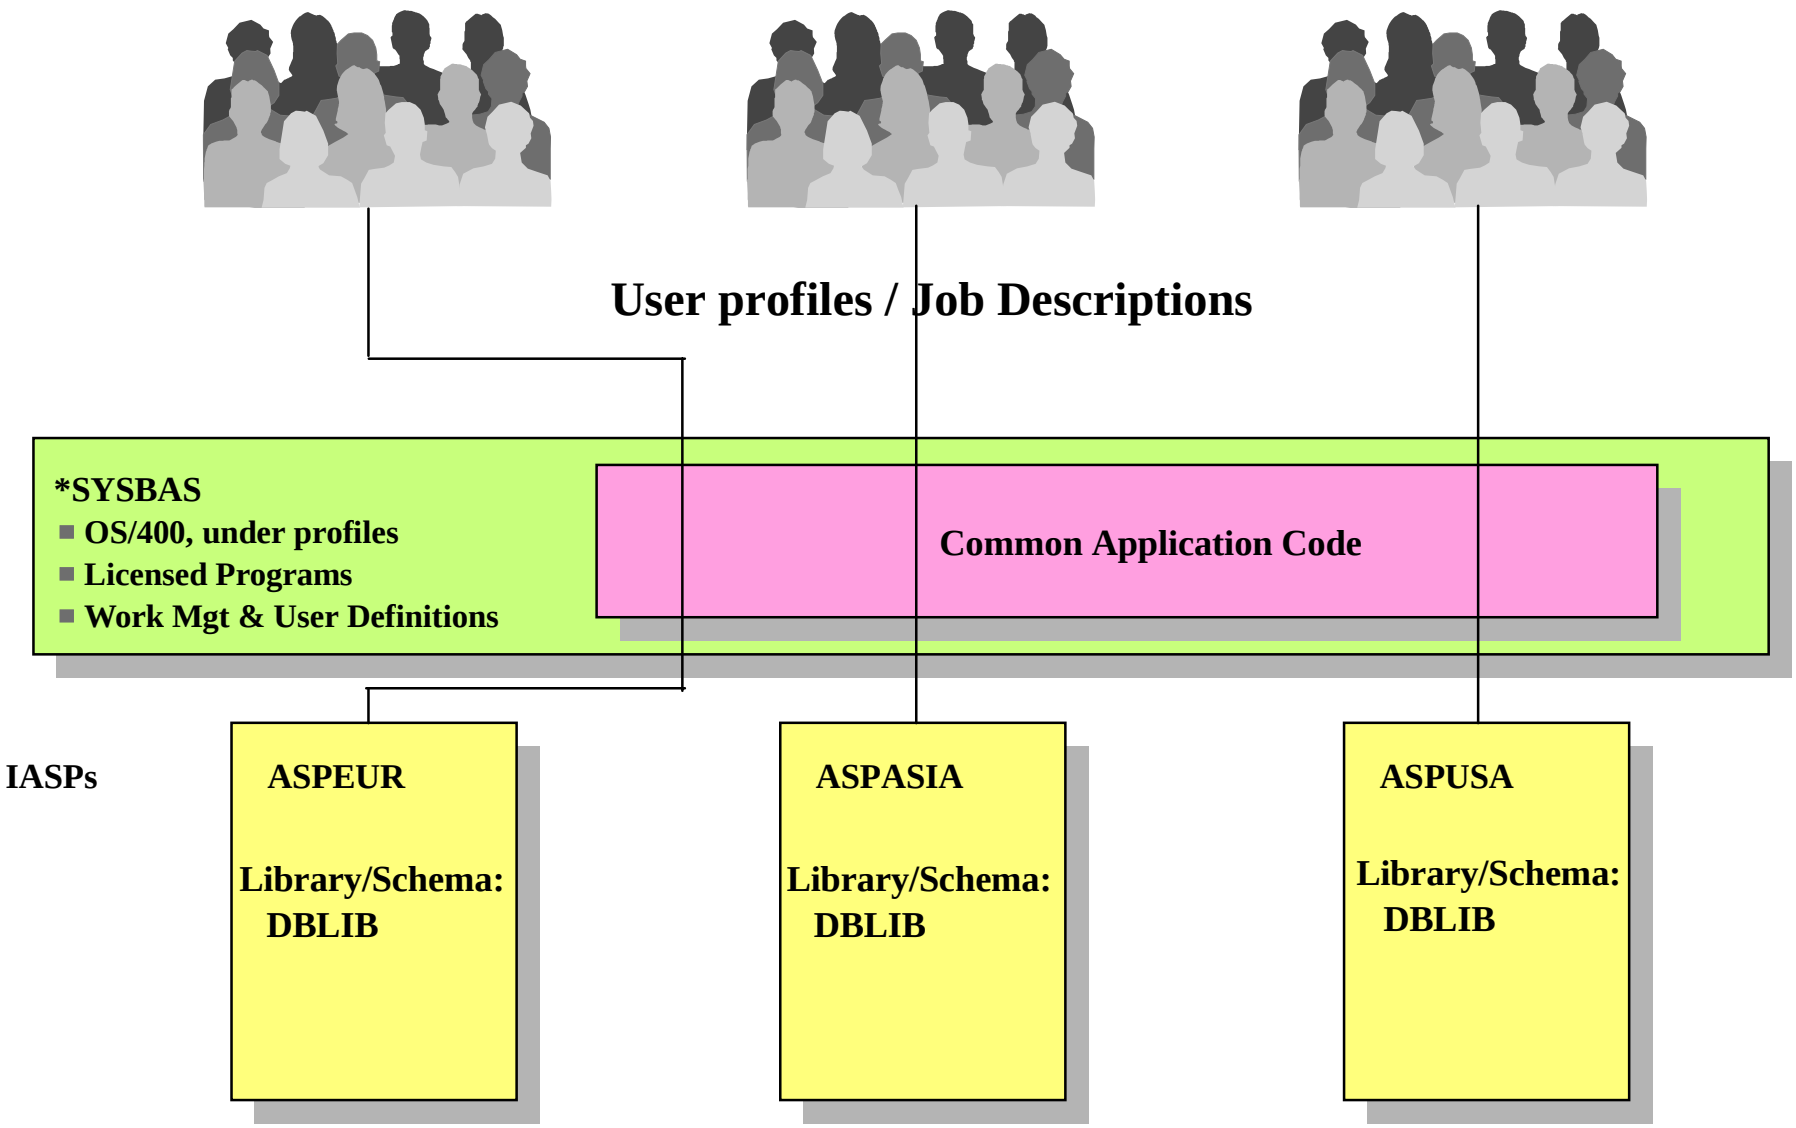

Notes: Multiple Database Support - Setup Example

This chart explains how to set up a consolidated environment on a single, not partitioned LPAR, where a common set of application code is used to access individual files with the same names, but of which the data has to be kept separately for business reasons.

Three IASPs are defined on this system, of which each will contain the data libraries of each entity in the company. Each IASP will therefore contain the same libraries with the same tables, views and indexes, but with different data - Europe, Asia, and USA. For purposes of simplifying change management, the application programs are stored only once in the system ASP.

The user profiles and their attributes are migrated from each system, and are associated with a job description that directs them to the IASP device where their data resides. This is done by specifying the name of the IASP device on the Initial ASP Group (INLASPGRP) parameter of the Job Description associated with each user profile. This way, a user that signs on to the system will have access to the common application code, residing in the system ASP, and to his data libraries, residing in the IASP.

If this system is running an application that needs to be enabled for support of different cultural values, it can still be achieved in this way, as long as the individual set up of the cultural objects follows the design principles as detailed in the International Application Development (SC41-5603) manual. In this case, multiple instances of the different cultural elements can be stored in the IASPs in their libraries, much as is the case on a system without multiple database support.

The advantage of such a consolidation, which can maintain its naming schema, is simple:

- No need to change the naming of all data containers, data objects or data attributes to fit in a consolidated IT service delivery environment; no secondary changes to the security setup either.
- Transparency for change and problem management including help desk support organizations.

Multiple Database Support - Setup Details

User Profile

- Initial Program
- Job Description

Initial Program:

- Identifies DB libraries to be worked with, uses SETASPGRP

Job Description, per IASP user(s):

- INLASPGRP
- Library List = *SYSVAL

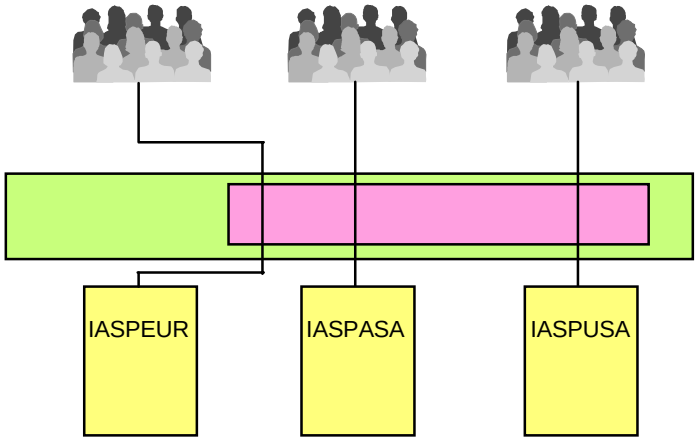

System Value QUSRLIBL:

- PGMGLAPAR
- PGMOE
- PGMPROD
- PGMMISC
- ...

Code Libraries in SYSBAS:

- PGMGLAPAR
- PGMOE
- PGMPROD
- PGMMISC

...DB Libraries in SYSBAS:

- MASTER

IASP Definition

Libraries per IASP DB:

- DBGLIB
- ...

Notes: Multiple Database Support - Setup Details

This figure shows more details of using multiple IASPs on a single system. This example is intended only as what we believe to be an implementation adhering to the principle of separate placing of processes and data.

The base assumption is to keep as much as possible the code libraries (programs, modules, environment descriptions) separated in the system ASP, while the data libraries are contained in separate IASPs, as required by the business entities. This is as discussed earlier in the notes for the foil: Multiple Database Support - Details.

The system value QUSRLIBL contains the names of the common code libraries and is used in the job descriptions associated with the different user profiles; each user profile has an initial program that sets up the data libraries the user will work with.

In a multilingual environment, you can add the libraries, containing the MRI or any other culturally dependent objects, in the system ASP also. These might also reside in the IASP.

In this example, we are showing the use of the SETASPGRP command and the job description Initial ASP Group parameter. The SQL CONNECT or a DDM file may also be used.

We show the SYSBAS library MASTER as a reminder the application can access objects in either the SYSASP or the "currently active" ASP Group.

Managing Libraries in an IASP examples

OS/400 commands:

- CRTLIB LIB(JIMCDLB1) ASP(*ASPDEV) ASPDEV(DLB1)
- WRKLIB/DSPLIB LIB(*ALL) ASPDEV(*ALLAVL , ASPname, SYSBAS, *ASPGRP, ...)
- SAVLIB, RSTLIB, ...

DSPLIB *ALL *ALLAVL example

Opt	Library	Type	ASP Device
	DLBLIB	PROD	DLB1
	JIMCDLB1	PROD	DLB1
	QRCL00033	PROD	DLB1
	QRCY00033	PROD	DLB1
	QRPL00033	PROD	DLB1
	QSYS00033	PROD	DLB1
	QSYS200033	PROD	DLB1
	SYSIB00033	PROD	DLB1
	#LIBRARY	PROD	
	#RPGLIB	PROD	
	ISALIB	PROD	
	MRK	PROD	

Operations Navigator - IFS

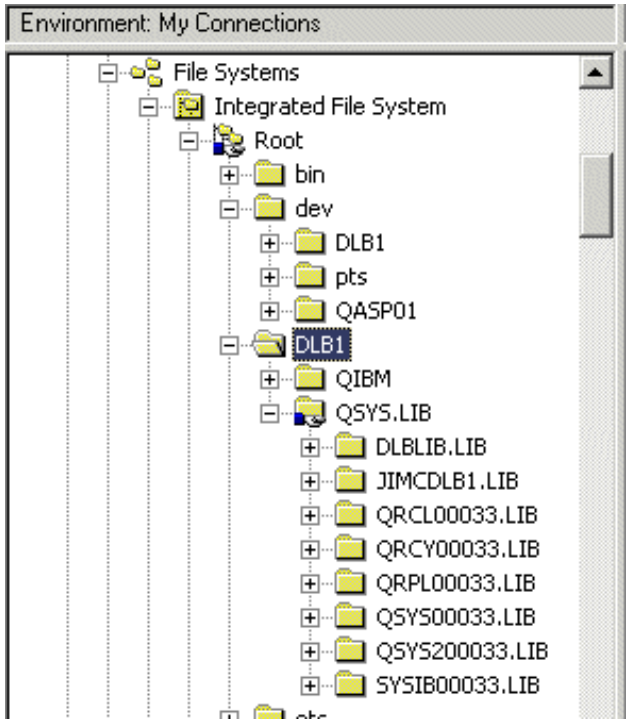

Planning is critical !!!

Business Needs

Performance Considerations

Software Requirements

Application Integration

Security Considerations

Capacity Planning

Hardware Configuration

Physical Planning Requirements

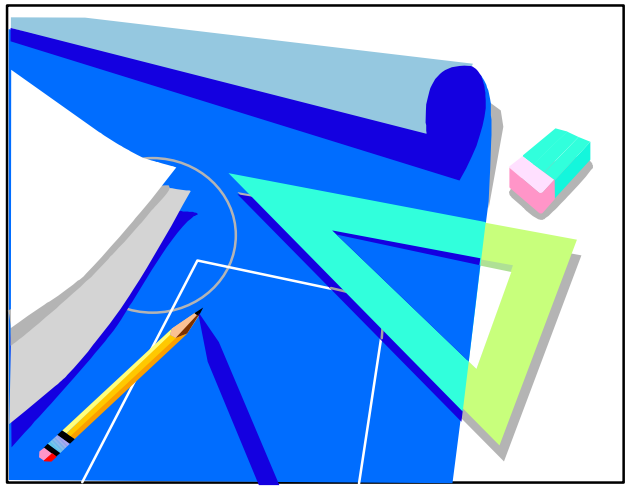

For further assistance please see
http://www-1.ibm.com/servers/eserver/iseries/ha/haplanning.htm

Notes: Planning is critical !!!

To implement an independent ASP solution is NOT something that simple. To be more explicit, it can be very complex. Planning is critical !!!
The planning considerations below are not intended to be complete since every customer situation vary, but they should give you some starting points.
For further assistance please see: http://www-1.ibm.com/servers/eserver/iseries/ha/haplanning.htm

Business Needs
Business needs are quantified, in terms of volumes of data, throughput, response time, etc. When doing the planning for IASP's, keep the business needs close by. In each step of the planning process, the business need should be satisfied before approval and adoption of the plan.

Performance Considerations
Performance is usually one of the metrics of the business needs. Even if it is not, maximizing throughput of the configuration generally requires a small effort with a potentially high return. The following areas should be considered:

- Disk Drives - number of disk arms versus total disk storage capacity. For example a few 35 GB disks may have sufficient capacity but be two few disk arms in an environment that has thousands of disk write operations.
- Placement for data and journals for better performance: The general recommendation to place journal receivers in a separate ASP applies.
- Current name space: iSeries Database support is more efficient if the current name space is the one containing the library and database object referenced by the currently active application program. The use of a job description to specify the initial IASP or the SETASPGRP command can be used to establish the current name space as shown earlier in this presentation. Otherwise, for example, an SQL CONNECT statement will use the Distributed Relational Database support within the same system. This is faster than actually accessing a remote system database, but is a longer code path through the system. This is also discussed in other foils on Multiple Database Support - Setup Details

Software Requirements
System ASP’s and User ASP’s are supported under all releases. UDFS ASP’s are supported under V5R1. Primary and Secondary ASP’s are supported under V5R2.

Notes: Planning is critical !!! -2

Application Integration

On V5R1, only IFS supported objects are supported in an IASP. On V5R2, the list of supported and unsupported objects can be found on the iSeries Information Center. Traditionally, data objects pertaining to an application area are stored in a data library; program objects pertaining to an application area are stored in a program library, and other objects common to the application area are stored in system libraries, or libraries designated as common to that application. User profiles, authorization lists, job queues, and spool output queues are some of the object types that cannot exist in an IASP. These are in system libraries at application installation time in most cases. However, consideration should be given to how these are replicated, saved, restored, and/or used in light of Independent ASP's.

Authority Considerations

User profile information is stored in the system ASP. Each user profile object is an object type of *USRPRF. Copies of *USRPRF objects are not in any independent pool. However, some user profile information must be maintained on the IASP itself. Additional storage (above that consumed by objects) is required for these system security structures. This is necessary to make the independent ASP self-contained. When creating independent disk pools in a clustered system environment, it is assumed the content of any user profiles is synchronized across the cluster in which the user profile exists.

Capacity Planning

Consider the amount of disk required, the performance requirements, then the amount of rack space required, then the switchability of the config, then the amount of floor space required. Sizing the disk pool All IASP's or Disk Pools within the rack will be switched when the rack is switched. When the application in Pool33 is switched, and Pool34 exists on disk in the same rack, Pool34 (and its resident application) will be switched as well.

Hardware Configuration

The first step after any capacity planning session is to begin the sizing and configuration of a system to fulfill the requirements. When building the configuration, keep in mind and factor into the planning the following items:

Physical vs. Logical Switching

Remember that when you create an IASP with disks that are part of the same parity protection set they might have parity stripes on them, so plan to select the correct drives to meet your capacity requirements.

Notes: Planning is critical !!! -3

Physical Planning Requirements

There is a lot to be considered in respect to physical planning. Except for the "normal" physical requirements like floor layouts and power requirements, things like HSL cable placement and SPCN cable placement can be complex.

For placement and cabling assistance, see iSeries Information Center at:

- http://www.ibm.com/eserver/iseries/infocenter
- Select Planning for Hardware and Software

Building a drawing of the desired configuration, especially the towers and racks involved, can often save costly errors when assembling the configuration upon delivery, especially in the area of cables. Design it one way, assemble it another way, and, the end result may be the need for additional cables.

The next set of foils address disk and other I/O device considerations within an IASP with examples.

IASP disk and other I/O device selection considerations

Non-Switchable Independent Auxiliary Storage Pool

Single server implementation of IASP

- Select any Disk Unit(s) except the Load Source Unit
- Implications with random (unplanned) selections
  - Protection
  - Recovery
- Number of disk units consideration

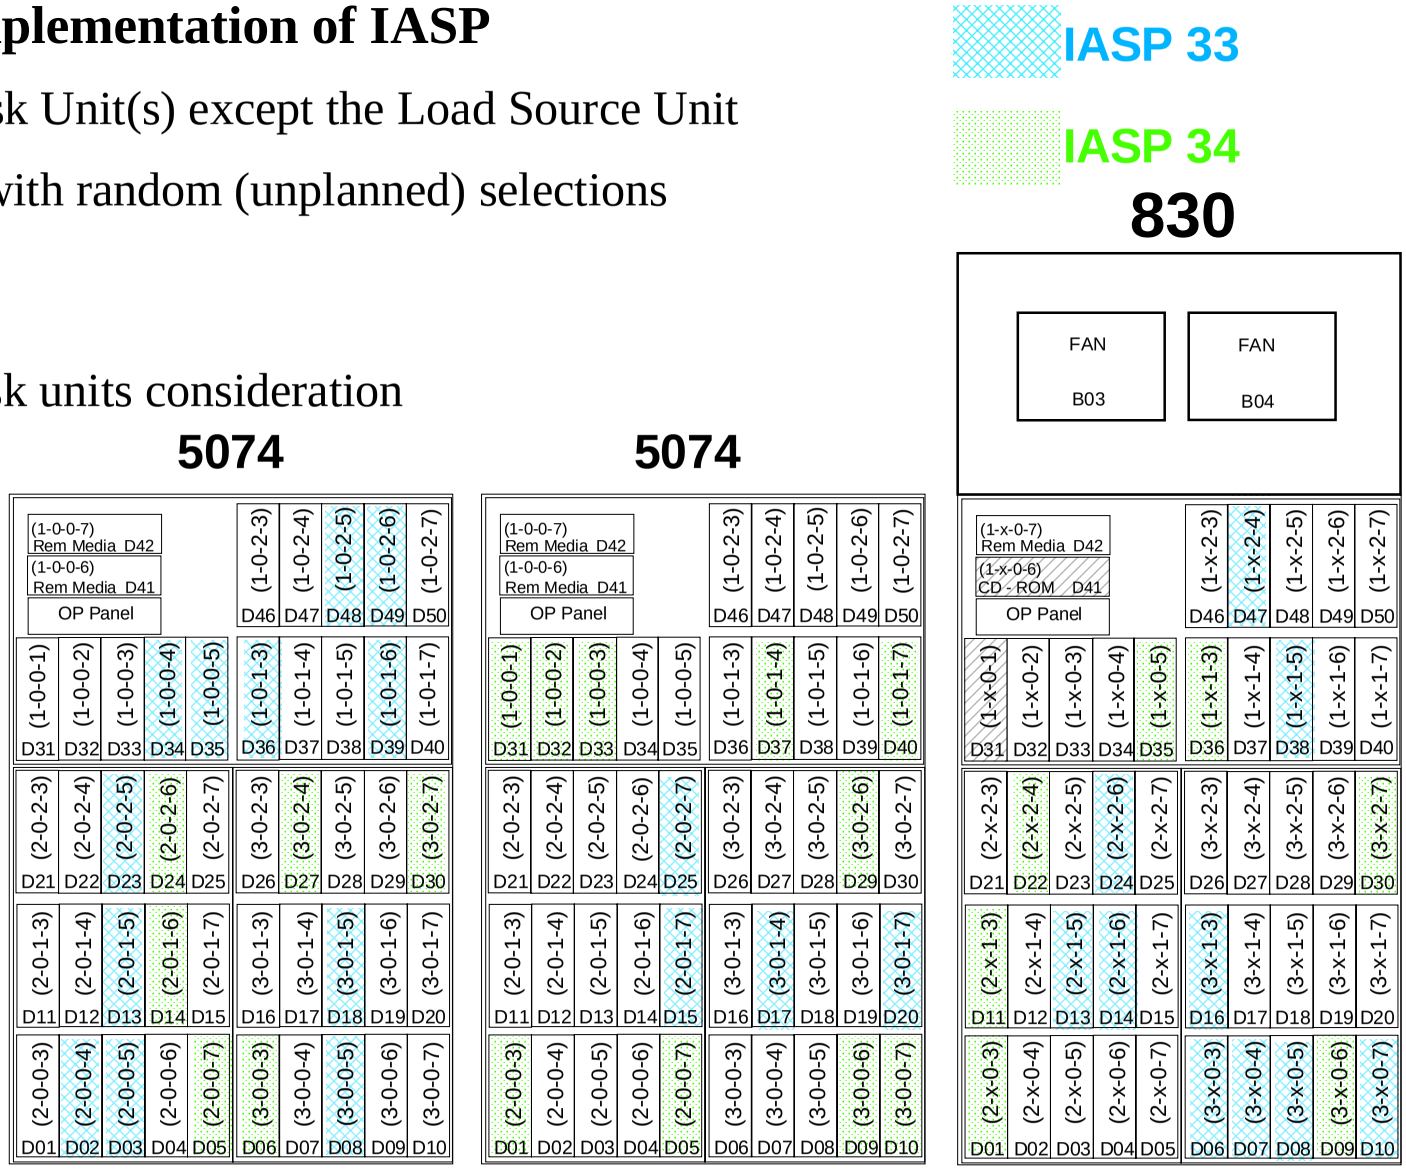

Notes: Non-Switchable Independent Auxiliary Storage Pool

Creating a non-switchable Independent Auxiliary Storage Pool (IASP) for use on a single system is intended for, but not limited to, support of multiple databases on the same system, as described in this presentation and the Database presentation. Selecting disks for this IASP can be done without taking any rules into account ("random or unplanned selection") - except for the fact that the load source of the system is required to be in the system ASP.

Although, not required, we recommend some logic be applied in determining which disks are placed into an independent ASP. When one understands the mechanics and rules of RAID5 protection and the physical I/O behavior using this protection method, it makes sense to select one or more Raid5 sets to form a group of disk units to be configured into the IASP. The high reliability of Raid5 protection with the capability of concurrent replacement and rebuild of a failing unit in a RAID5 set , makes the reload of data due to a hardware failure a very rare event. However, in the exceptional case that two disks in a single RAID5 set would fail, the alignment of an ASP (whatever type) with one (or more) RAID5 set(s) will prevent data loss in the system ASP or any other ASP.

Starting mirroring in an ASP (any type) is easy. However, the system licensed internal code will automatically obtain the highest protection level possible with the disk units selected for a particular ASP when mirroring for that ASP is started. Disk units must be selected depending on the level of mirroring protection that needs to be achieved: Disk level, IOP level, Bus level or Tower level.

For acceptable performance, the appropriate number of disk arms should be selected to configure the ASP. The number of disk arms in an IASP can influence the performance behavior of an application using the data in that IASP. Mixing disk sizes in a single ASP can also effect the disk I/O behavior. Therefore, it is recommended to use the same type and size of disk units when configuring disk units in an ASP. Correct sizing must be done based on an accurate interpretation of performance measurements.

For example, you could review disk activity over time with the WRKDSKSTS command or viewing the Disk Resource section of the System Report available with Performance Tools for iSeries, 5722PT1, to understand which disks are the most busy in your application environment.

Switchable Independent Auxiliary Storage Pool

Single server implementation of switchable IASP with LPAR

- Cluster requirement between partitions
- Select Disk Units on IOP level
- Respect shared bus rules
- Number of disk units consideration

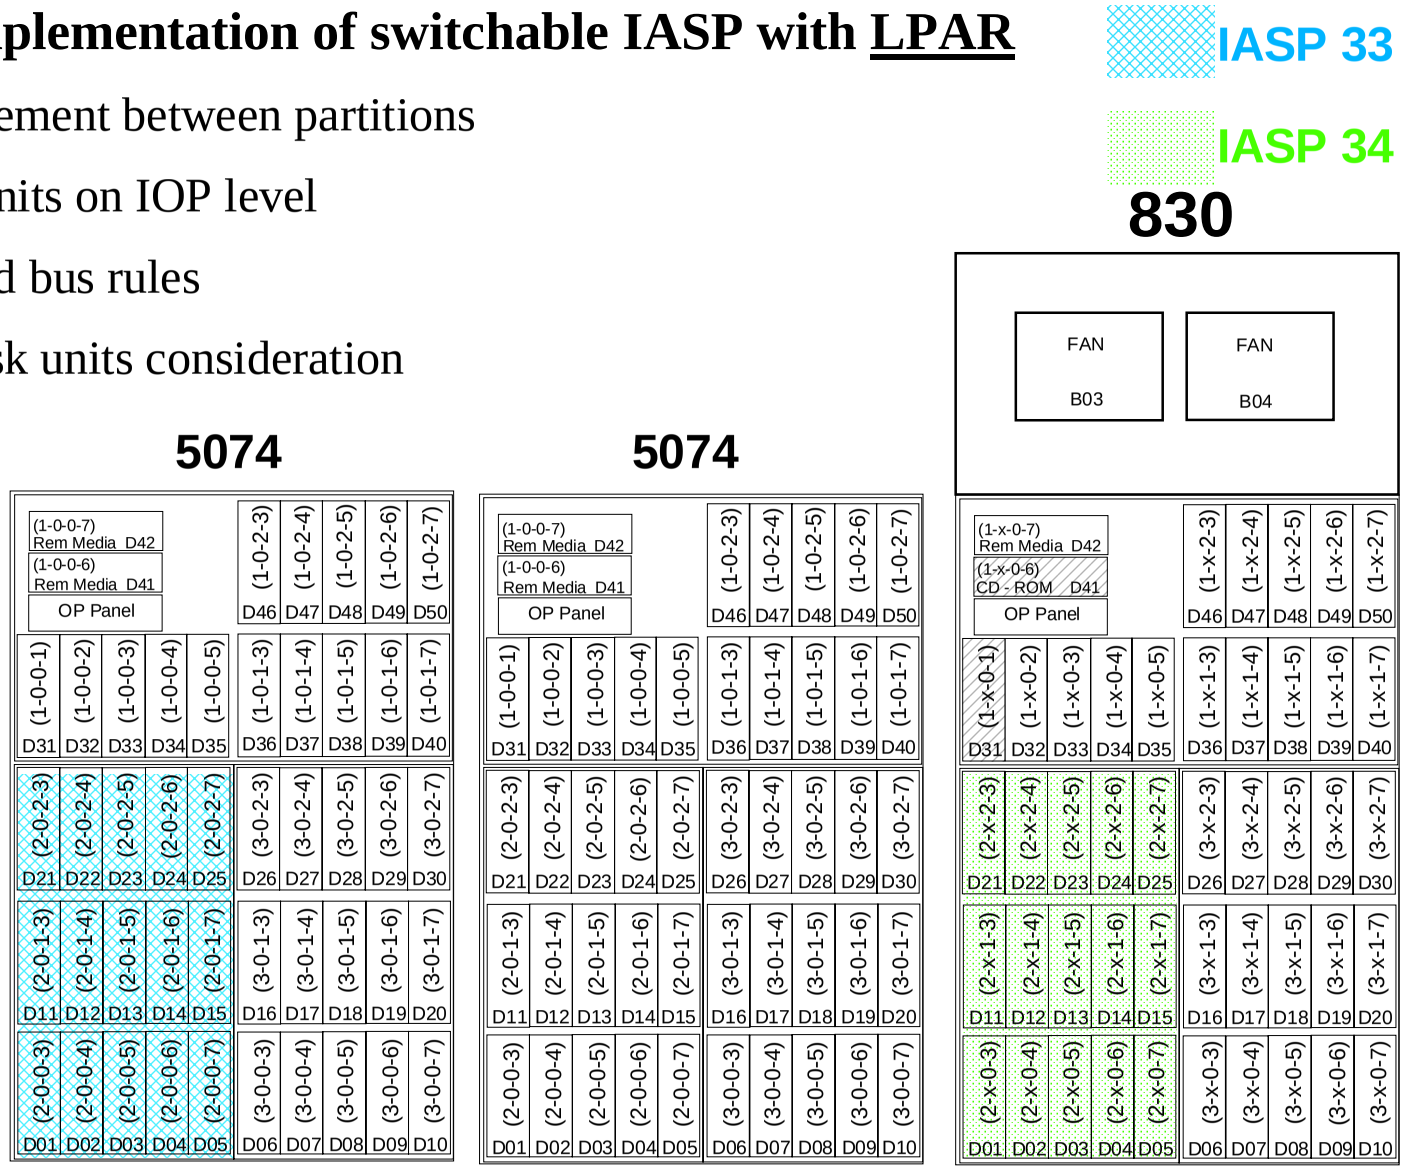

Notes: Switchable Independent Auxiliary Storage Pool

Switchable Independent Auxiliary Storage Pools can be used with a cluster of partitions within a single system. The configuration of that type of switchable IASP can only be done at the IOP level. This is simply because the IOP is the lowest possible hardware level that can be removed from or moved (switched) to partitions that share the bus the IOP resides in. This rule implies that if a number of adapters or disk controllers are assigned to a certain IOP and that IOP is removed from partition A and moved into partition B, all associated assigned adapters and controllers move together with that IOP.

Placement or assignment of disk controllers that drive the disk units configured in a switchable IASP must be carefully planned. There are a number of other considerations such as bus ownership in the logical partition setup, cluster rules, disk controller capabilities and more that need proper investigation before configuring a switchable IASP.

For performance, the same rules and considerations as for the non- switchable IASP apply . In additional to these considerations, the way the switchable IASP will be used on the clustered nodes has to be considered.

As described elsewhere in this presentation, the vary-on time for the IASP device can be considerably reduced by minimizing the number of database objects in the *SYSBAS environment. Recall that SYSBAS means the system ASP and any user (dependent) ASPs (ASP 1 to 32). This is because the cross reference of the *SYSBAS environment (primarily database objects and object authorities) must be propagated into the IASP that is brought online. The necessary checking must be performed and ownership and authorization information, called user profile extensions, must be written into the IASP.

In a clustered environment, after the IASP has been switched, the time needed to make the data in a Switchable Independent Auxiliary Storage Pool available to a cluster node is directly impacted by the way the applications and data are distributed over the *SYSBAS ASPs and the "newly arrived" IASP. That is, if there are a large number of database objects in SYSBAS on the switched to system and several database objects in the "newly arrived" IASP, then the vary on (make available) process of this IASP on the target system could be longer than expected.

Switchable Independent Auxiliary Storage Pool

Multiple server implementation of switchable IASP

- Cluster configuration requirement
- Select disk units on Tower level
- Respect switchable tower rules
- Number of disk units consideration

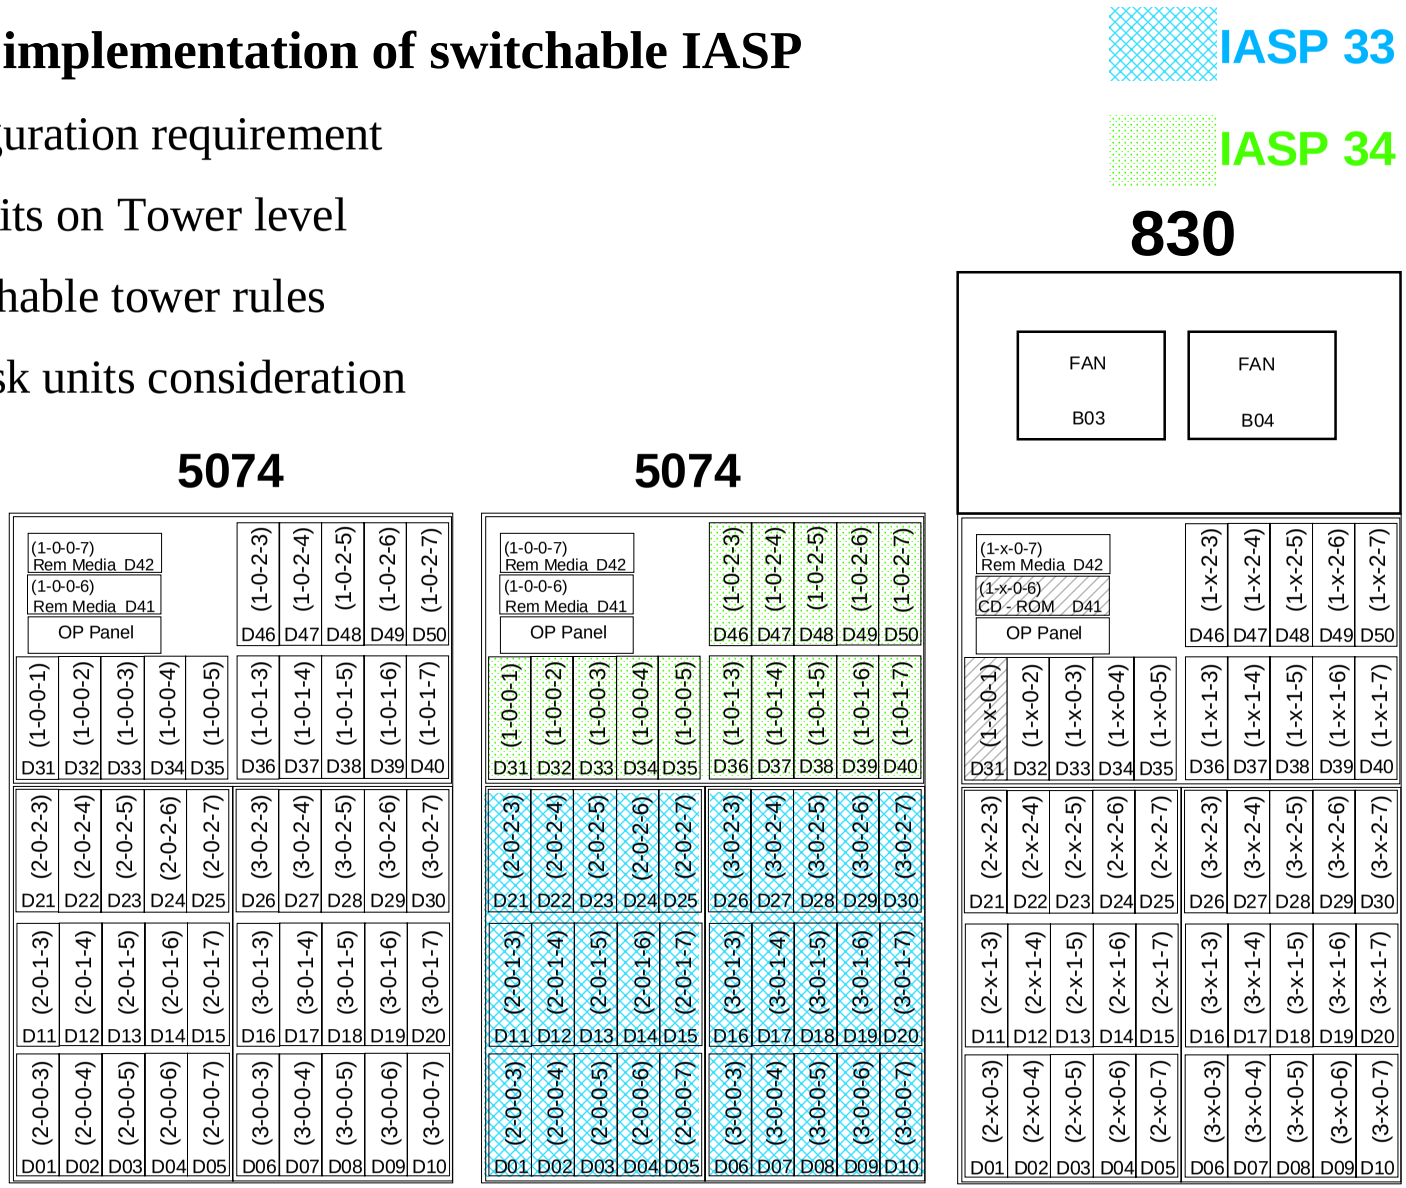

Notes: Switchable Independent Auxiliary Storage Pool

Configuring a Switchable Independent Auxiliary Storage Pool for use on two clustered iSeries servers is essentially the same as configuring the switchable IASP for use on two clustered iSeries partitions on a single iSeries server.

The only additional consideration here is that the IASP needs to physically reside in one or more switchable HSL towers. The HSL tower is considered as the lowest switchable hardware level resource that can be switched between two iSeries clustered servers. The complete tower(s) is switched from one clustered iSeries server to the other. Therefore, in this example configuration both IASP33 and IASP34 would be switched between systems.

Note also, all hardware resources contained in the tower are brought over as well as the IASP resources. To use the hardware resources other than the IASP hardware, a configuration object for each of these hardware resources must exist.

The HSL rules for clustered iSeries servers must be respected. Only four I/O towers can connect to a clustered loop with two iSeries servers and the maximum of towers is three per HSL segment. All other aspects mentioned in the two previous foils should be considered also.

We did not cover IASP configurations on external DASD (ESS) attached to the #2766. The rules are essentially the same. Since a #2766 requires a dedicated IOP, all LUN's attached to the #2766 will be moved together with the IOP the #2766 is assigned to. Moving a tower containing multiple #2766 adapters will of course move all associated hardware attached to and contained in that tower to the other iSeries cluster node.

It is also possible to "randomly" configure IASPs with a selected number of LUN's that are made available for use on a single iSeries server.

Backup and Recovery of IASPs

Carefully plan your backup strategy

- Disk Migrate While Active when data is already on disks to be placed into an ASP

Saving IASPs

- And saving your entire system

Restoring IASPs

Saving and Restoring Linux & Windows NWSSTG on IASP

Using BRMS for IASPs

Notes: Backup and Recovery of IASPs

This next set of foils describe techniques and strategies for backup and recovery of independent disk pools. When you add an independent disk pool to your system configuration, you will need to plan for the backup and recovery of the user data on these devices, **because these devices operate differently than the system or basic user auxiliary storage pools**.

These differences will mean that you will have to carefully plan your backup strategy to assure you have a complete system backup.

Disk Migration While Active

Enables disk migration task

- Reduce downtime
- Helpful in migration scenarios

Start ASP Balance (STRASPBAL) TYPE(CAPACITY, *USAGE, *HSM ...)

- *ENDALC V5R2
- *MOVDTA V5R2

User may specify a run time limit on the function

V5R1 Support

- Controlled availability: V5R1 PTF SI03310:
  - Not via the STRASPBAL command
  - Refer to the Special instructions that are in the PTF cover letter

Notes: Disk Migration While Active

To enable disk migration set of tasks, used for example during system upgrades can also be applied to setting up and existing configuration to use IASPs. Essentially this set of tasks are used to reduce the down time associated with removing existing data from a set of disks and remove the disk units so they be added to an IASP.

The Start ASP Balance (STRASPBAL) command enables you to (re)move data from disk unit(s) onto other disk units , while the system is active.

Move data from units

A unit that is scheduled for removal can be marked to end allocations by specifying UNIT(unit-number) and TYPE(*ENDALC). This will keep new allocations away from this unit. For all units marked *ENDALC, specifying TYPE(*MOVDTA) will move data from the marked units to other units in the same ASP. To resume allocations for units marked *ENDALC, specify UNIT(unit-number) and TYPE(*RSMALC). New allocations will once again be allowed to this unit. The Check ASP Balance (CHKASPBAL) command can be used to determine which units are currently marked *ENDALC.

As with previous OS/400 releases the administrator may specify a time limit that the function is to run for each ASP being balanced or the balance can be set to run to completion. If the balance function needs to be ended, use the End ASP Balance (ENDASPBAL) command. A message will be sent to the system history (QHST) log when the balancing function is started for each ASP. A message will also be sent to the QHST log when the balancing function completes or is ended.

If the balance function is run for a few hours and then stopped, it will continue from where it left off when the balance function restarts. This allows the balancing to be run during off hours over a several day period.

Saving iASPs

OS/400 SAVxxx and RSTxxx commands

- Use the SETASPGRP command or
- SAVxxx ASPDEVparameter
  - *, *SYSBAS, *CURASPGRP, ASP device name
  - Special considerations when saving to a save file

On a full-system save (Option 21) or a save of all user data (Option 23)

- Independent disk pools are saved alphabetically
- Secondary ASPs are saved along with their primary

Example:

Save Order	Independant ASP Name	Independant ASP Type	What is Saved	Commands
1	Apples	Primary	Libraries	SAVLIB LIB (*NONSYS or *ALLUSR)
	Cantaloupe	Secondary		
2	Apples	Primary	User-defined file systems	SAV OBJ(('/dev/*'))
	Cantaloupe	Secondary		
3	Bananas	UDFS	User-defined file systems	SAV OBJ(('/dev/*'))

Notes: Saving iASPs

The native OS/400 SAVxxx and RSTxxx commands have been enhanced to provide support for IASPs. Using these commands in your own CL programs to backup the system is relatively straightforward, as you are in control of the environment when these are running. In general, these native commands must have access to the namespace where the objects to be saved reside. This can be achieved by using the SETASPGRP command prior to issuing the Save command or by using the new with V5R2 ASPDEVparameter on the SAVxxx and RSTxxx commands.

If you have grasped the concept of using an IASP, you should be able to save or restore specific libraries or objects in that IASP. However, if you are more familiar with using the Save and Restore menus to save or restore your system or its components (for example, *NONSYS, *ALLUSR or *IBM saves), you will need to understand the way in which these are affected by addressability to the IASPs. This is particularly important if you use Option 21 = Entire system.

Special considerations on save commands

The ASPDEV parameter will allow you to save the independent ASP or object in SYSBAS and, if active in the job/thread, an IASP - without changing your job thread, depending on the value you use for this parameter. If you are saving to save files, this parameter does not affect the DEVICE parameter of the save commands. You must use the SETASPGRP command if you are saving to a save file that exists in an independent ASP. This also allows you to save to a save file that exists in a different independent ASP than the one you are saving. Rather, the ASPDEV parameter acts as a filter on the SAV command. The following is an example:

There is a file called MOLN and a save file called CORZ in a library called BER in an IASP called ROCHESTER.
There also is a save file called ADMS in library QGPL.

The following command **will not work unless** SETASPGRP has been previously run:
SAVOBJ OBJ(MOLN) LIB(BER) DEV(*SAVF) SAVF(BER/CORZ)

Even had we added the ASPDEV such as in the following command, it still will not find the save file.
The following command **will not work:**
SAVOBJ OBJ(MOLN) LIB(BER) DEV(*SAVF) SAVF(BER/CORZ) ASPDEV(ROCHESTER)

To save the file MOLN to the save file CORZ, the SETASPGRP command has to be used.
The following is the correct sequence to use the save file:
SETASPGRP ASPGRP(ROCHESTER)
SAVOBJ OBJ(MOLN) LIB(BER)DEV(*SAVF) SAVF(BER/CORZ)

Notes: Saving iASPs -2

Here is a quick summary of the supported values for ASPDEV parameter using Save Library command as an example:

- *****: Saves a library located in the system ASP (ASP number 1) and in any basic user ASPs (ASP numbers 2-32). If the current thread also has an ASP group, all independent ASPs in the ASP group will also be saved.
- ***SYSBAS**: Saves a library located in the system ASP (ASP number 1) and in any basic user ASPs (ASP numbers 2-32).
- ***CURASPGRP:** If the current thread has an ASP group, saves a library in all independent ASPs in the IASP group are included in the save operation.
- **auxiliary-storage-pool-device-name**: saves a library in the specified IASP group.

Notes: Saving iASPs -3

You can save your independent ASPs as part of a full system save (GO SAVE: Option 21), or when you save all user data (GO SAVE: Option 23). You can also use BRMS to save your independent ASPs.

In either case, you must make the independent ASPs available before you perform the save. You can also just save the current ASP group. Saving the ASP group saves both he primary ASP and any associated secondary ASPs. If you make independent ASPs available, they will be included in an Option 21/option 23 save. Before you end subsystems and restrict your server, make sure that your current job does not use integrated file system objects in the independent ASP and do not perform a SETASPGRP cmd.

The following scenarios show examples of using different options to save a library or object within an independent ASP.

Save the current ASP group (Following are the commands to save the current independent ASP group)

- SETASPGRP ASPGRP(primary-ASP-name)
- SAVSECDTA ASPDEV(*CURASPGRP)
- SAVLIB LIB(*ALLUSR) ASPDEV(*CURASPGRP)
- Unmount any QDEFAULT user-defined file systems in the current Independent ASP group
- SAV OBJ(('/dev/*')) UPDHST(*YES) ASPDEV(*CURASPGRP)
- Mount any QDEFAULT user-defined file systems that were unmounted in an earlier step

Save UDFS ASP (Following are the commands to save an available UDFS ASP)

- SAVSECDTA ASPDEV(ASP-name)
- Unmount any QDEFAULT user-defined file systems in the current Independent ASP group
- SAV OBJ(('/dev/*')) UPDHST(*YES) ASPDEV(*CURASPGRP)
- Mount any QDEFAULT user-defined file systems that were unmounted in an earlier step

Save libraries and objects within an Independent ASPs as part of a full system save (Option 21) or when you save all user data (Option 23)

In addition to the commands listed in the GO SAVE Options, the server performs the following commands for each available ASP:

- SETASPGRP ASPGRP(asp-group-name)
- SAVLIB LIB(*ALLUSR) ASPDEV(*CURASPGRP)
- SAV OBJ(('/dev/*')) UPDHST(*YES) ASPDEV(*CURASPGRP)
  The server then performs the following command for each available user-defined file system (UDFS) ASP.
- SAV OBJ(('/dev/*')) UPDHST(*YES) ASPDEV(udfs-asp-name)

The system will also perform a CHKTAP ENDOPT(*UNLOAD) command after the last SAV command it processes.

Restoring iASPs

The system ASP must be restored first

- Then restore any independent disk pool
- Or pools must be manually created through iSeries Navigator

Special consideration for the RSTAUT command

- Use after you have recovered all of your independent ASPs

Restoring iASPs

Restoring an entire system that is using independent disk pools becomes a more complicated matter. The system ASP must be restored first, then the any independent disk pool. Or an IASP must be manually created through iSeries Navigator or the Create Device ASP (CRTDEVASP) command must be issued and a restore into the IASP be performed.

This requires knowledge of the original disk pool sizes and names. Also, if you are using the restore menu to recover user data, you should consider not using the RSTAUT command until you have recovered all of your independent ASPs. An example recovery procedure of your system would be as follows (do not forget to review the V5R2 *Backup and Recovery Guide*, SC41-5304):

- Restore your SAVSYS
- If you are recovering your system from an option 21 save, you may perform option 21 restore, which does the following steps: (You may want to prompt for commands)
  - RSTUSRPRF
  - RSTCFG OBJ(*ALL)
  - RSTLIB SAVLIB(*NOSYS)
  - RSTDLO DLO(*ALL)FLR(*ANY)
  - RST ((‘/*’)(‘/QSYS.LIB ’ *OMIT)(‘/QDLS ’ *OMIT))

Perform the following steps for each IASP you need:

- In OS/400 create your independent storage pools from iSeries Navigator or use the CRTDEVASP command.
- RSTLIB SAVLIB(*NONSYS) ASPDEV(youriasp)
- RST OBJ(‘/dev/*’)

Then perform RSTAUT USRPRF(*ALL). Note, If RSTAUT has been performed prior to recovering your independent ASPs, you must RSTUSRPRF first, then use RSTAUT.

Saving and Restoring Linux & Windows NWSSTG on iASP

You have additional save and restore considerations compared to an OS/400 partition:

- Information under the QFPNSSTG in the system ASP
- The network storage space ((NWSSTG)) named under the /dev/IASPname directory

Restore the network server description as last recovery step

- Automatically link the network server description (NWSD) to the network server storage space (NWSSTG)

If the NWSD is on a switchable IASP

- You may save the pointers and the network server description
- Restore NWSD to the other system in the cluster
- By switching the IASP, Linux & Windows becomes usable on the switched to system

Notes: Saving and Restoring Linux & Windows NWSSTG on iASP

When a network server storage space is created on an IASP, it still creates its pointers in the QFPNSSTG directory which is stored in the system ASP. This means to save and restore a network storage space that was created in an IASP, you must save the information under the QFPNSSTG in the system ASP, and also the network storage space named under the /dev/IASPname directory. The following is an example creating a Linux storage space:

```
CRTNWSSTG NWSSTG(LINUXSTG)NWSSIZE(3000)FORMAT(*OPEN)ASP(40)
```

This creates a network storage space called LINUXSTG in ASP40, which, in our example we have previously named ROCHESTER. The actual storage space resides in User Defined File System - /dev/rochester/linuxstg.udfs. It also creates an entry in /qfpnsstg/linuxstg/qfpcontrol along with a /mount directory under /qfpnsstg/linuxstg.

To use a network storage space, you must also create a network storage description to link to the storage space. The following is an example creating a Linux network storage description and adding the link to the network server storage space:

```
CRTNWSD NWSD(LINUXSVR)RSRCNAME(*NONE)TYPE(*GUEST)PARTITION(LINUX) +
ADDNWSSTGL NWSSTG(LINUXSTG)NWSD(LINUXSVR)
```

If the network storage space was created on a switchable IASP, you may save the pointers and the network server description and restore it to the other system in the cluster, and by switching the IASP, Linux would be usable on the switched to system.

The following commands (shown on next page) will save the objects needed for Linux from System A, put them in a save file in the switchable IASP, restore the objects to System B, then allow Linux to be active in the partition on system B. Note that System B would have to already have a Linux partition configured, but the network server space and the network server description will be restored from System A.

Notes: Saving and Restoring Linux & Windows NWSSTG on iASP -2

On system A:
CRTLIB LIB(MYLIB)ASP(*ASPDEV)ASPDEV(ROCHESTER)
SETASPGRP ASPGRP(ROCHESTER)
CRTSAVF FILE(MYLIB/SAVEFILE1)
CRTSAVF FILE(MYLIB/SAVEFILE2)
SAV DEV('/ROCHESTER/QSYS.LIB/MYLIB.LIB/SAVEFILE1.FILE')OBJ(('/QFPNWSSTG/LINUXSTG'))
SAVCFG DEV(*SAVF)SAVF(MYLIB/SAVEFILE)
SETASPGRP ASPGRP(*NONE)

On system B:
After the IASP has been switched to System B, the following commands will make Linux usable on System B (assuming that the partition already exits):

Vary on or Make Available (iSeries Navigator
SETASPGRP ASPGRP(ROCHESTER)
RST DEV('/ROCHESTER/QSYS.LIB/MYLIB.LIB/SAVEFILE1.FILE')OBJ(('/QFPNWSSTG/LINUXSTG'))
RSTCFG OBJ(LINUXSVR)DEV(*SAVF)OBJTYPE(*NWSD)SAVF(MYLIB/SAVFILE2)
SETASPGRP ASPGRP(*NONE)

It is very important to restore the network server description last, as this will automatically link the network server description to the network server storage space. At this point, varying on the **network server storage** description should bring up the Linux partition.

Using BRMS for iASPs

Updated Edit Backup Control Group Entries screen

- Default values should not affect your current backup strategy

IASP recovery steps are included on the System Recovery Report

```
                    Edit Backup Control Group Entries               RCHASM20
Group . . . . . . . . . . . . : TEST
Default activity  . . . . . . *BKUPCY
Text  . . . . . . . . . . . . . test
Type information, press Enter.

                          Auxiliary   Weekly    Retain  Save      SWA
      Backup      List    Storage     Activity  Object  While     Message
Seq   Items       Type    Pool Device SMTWTFS   Detail  Active    Queue

  10  *EXIT
  20  *SAVSYS
  30  *IBM
  40  *ALLUSR             *SYSBAS
  50  *ALLDLO
  60  *LINK               *ALLAVL
  70  *EXIT

F3=Exit                     F5=Refresh
F11=Display exits           F12=Cancel
```

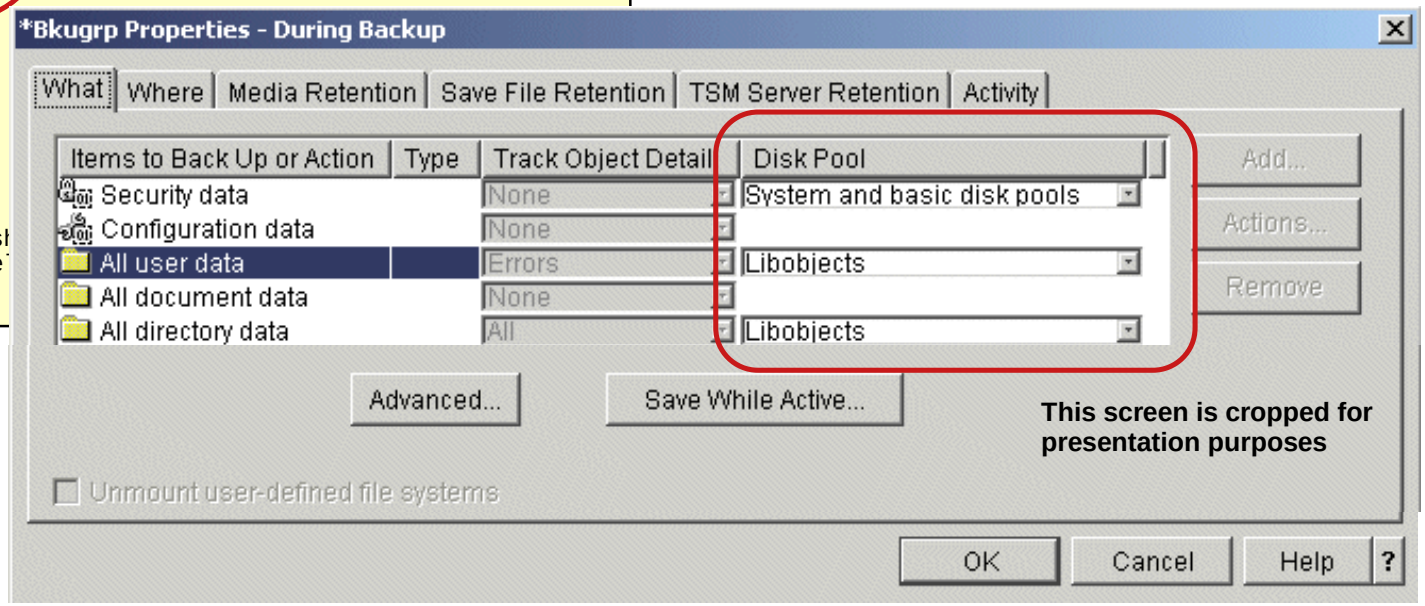

Notes: Using BRMS for iASPs

The V5R2 Backup and Recovery Media Services (BRMS) licensed program (5722-BR1) Edit Backup Control Group Entries command screen has been updated as shown to include a new Auxiliary Storage Pool Device field. In the screen example, the IASP name is Libobjects.

This field will not appear on some backup items entries. Typically this occurs for backup items which cannot reside on auxiliary storage pool devices. The Auxiliary Storage Pool Device prompt will be automatically filled in for entries of your existing backup control groups to reflect the scope of the save across auxiliary storage pool devices.

These default values should not affect your current backup strategy and should be consistent with what was being saved by the control group in V5R1.

- The *SYSBAS value on the *ALLUSR backup item saves all user libraries on the system (1) and any basic user (2-32) auxiliary storage pools.
- The *ALLAVL value for the *LINK backup items saves the links on the system (1) and any basic user (2-32) auxiliary storage pools as well as the links on all available auxiliary storage pool devices.

Note: When saving directory and files, you should unmount any mounted user-defined file systems (UDFSs) prior to the save to assure the objects in the mounted over directories are saved. UDFSs are automatically unmounted on auxiliary storage pool devices when the system is in restricted state. UDFSs on the system or basic user auxiliary storage pools need to be explicitly unmounted. Any unmounted UDFSes need to be remounted after the save.

Advantages of using IASPs - single system

Single-system environment

- Isolate low-use data with ability to bring online only when needed
- Reduce system start time (IASP not automatically varied on during IPL)
- 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
  - Server consolidation with multiple databases:
    - Enables multiple database support
    - if 2 RAID disk failures occur in IASP only the data in that failing IASP becomes unavailable; the application can continue to run against data in a different IASP
- Perform application maintenance that does not affect entire system

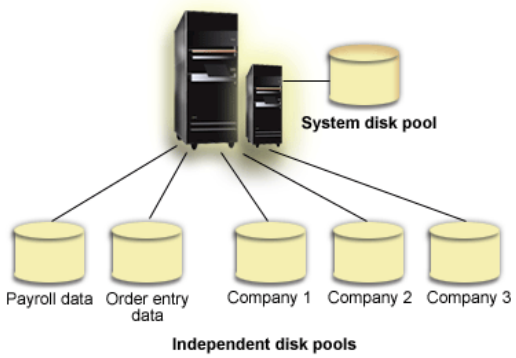

Advantages of using iASPs - multiple systems

Multi-system clustered environment

- Availability
  - Keep data available even in the event of a single system outage, either scheduled or unscheduled.
- Server Consolidation of “Branch Office” type systems
- Isolate data associated with specific applications
- Workload switching across multiple servers

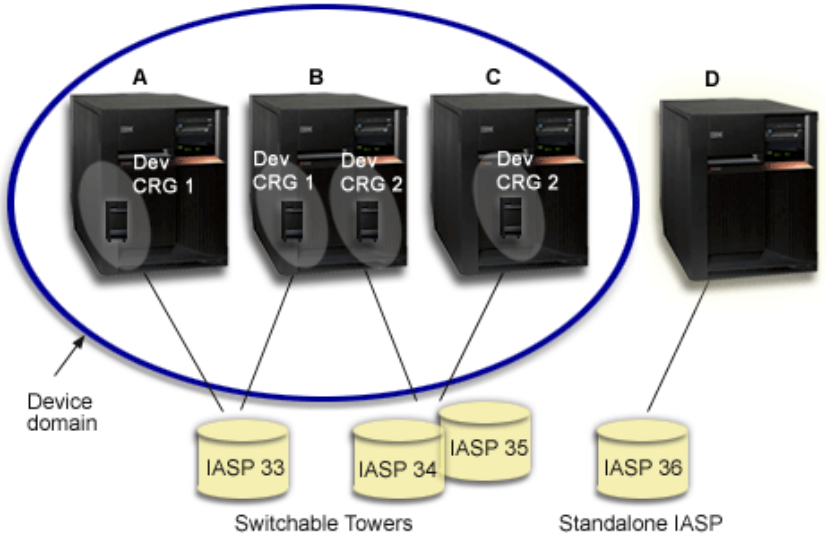

Notes: Advantages of using iASPs

The advantages of using iASPs mentioned on this foil are self explanatory.

Beyond the capabilities of replication solutions

Manage Disk protection options

- RAID5 and Mirroring protection can be stopped or started from within OS/400

Higher level of availability without the need to buy a duplicate set of disks

- Need not maintain multiple copies of data, programs, and other objects

There is minimal additional system overhead with IASP

- Replication requires more CPU cycles and involves network traffic
- Less resources necessary for system functions
- Objects are not “in flight” in the event of a failure

Notes: Beyond the capabilities of replication solutions

iSeries availability is enhanced through the use of IASPs, different than the capabilities of replication solutions.

Manage Disk protection options

When you make restricted changes to disk configuration in a basic disk pool you must have your server restarted to Dedicated Service Tools (DST). In an off-line independent disk pool you do not have to have your server in DST mode to start or stop mirroring, start device parity protection, start compression, remove a disk unit, etc.

Higher level of availability without the need to buy a duplicate set of disks

It is not necessary to maintain multiple copies of data, programs, and other objects. Multiple copies of objects is a function of replication. In a sense, IASPs are the poor man's option for higher availability.

There is minimal additional system overhead with IASP

- Replication requires more CPU cycles when replicating to a backup system. There is no network traffic associated with IASP. Replication across a LAN or WAN involves network traffic.
- There is less work for system functions such as IPL, reclaim storage, and some save operations.
  In a single system environment, an independent ASP can be used to store certain data off-line except for the periods when it is actually needed. The isolation provided by storing data off-line means that there is less work necessary for system functions.
- Objects are not “in flight” in the event of a failure.
  With replication, it is possible that journal entries become “trapped” on the source system at the time of failure and do not arrive at the target machine.

IASP Availability considerations

IASPs represent a single point of failure in the system

Not a disaster recovery solution

Pre-V5R1 HSL adapters do not work with IASPs

IASP and balancing workloads

IASPs can coexist with HABP solutions

- HABP solutions provide geographical dispersal of the data
- The IASP cannot be used for balancing workload

Notes: IASP Availability considerations

Using only IASPs, there are some deficiencies in total system availability remain: for example:

IASPs represent a single point of failure in the system: If the disks in the IASP are permanently damaged, then the data is unrecoverable so special care is need to ensure that the correct level of disk protection is chosen for the IASP. For critical data, tower level mirroring may be appropriate. For less critical data, RAID protected disks may be the best solution. Additionally, a secondary IASP used for journaling should exist in a separate RAID set to ensure that the latest updates are available if the data in a primary IASP becomes damaged or lost.

Not a disaster recovery solution: Because of loop limitations with HSL, the systems must be within 250 meters of each other (15 meters with some configurations). The production and backup systems can be several thousand kilometers apart when replication is used. IASPs are, therefore, not useful as a disaster recovery solution.

Pre-V5R1 HSL adapters do not work with IASP: If the IASP configuration involves an HSL loop, a V5R1 supported HSL adapter (such as feature #2739 Optical Bus Adapter) is required. HSL adapters available prior to V5R1 (for example the #2695 Optical Bus Adapter) do not work with IASPs. Models 830 and 840 systems with original HSL cabling can be upgraded to newer HSL features. Models 270 and 820 will require a model upgrade to 830 or higher to support the required HSL technology.

IASP and balancing workload: Careful planning should be done to balance workloads across systems with IASPs when ever possible. The way to achieve the is to use multiple IASP groups with separate applications or application instances. For example, if used with Domino, split the Domino servers across at lest 2 IASP groups, with each system in the loop acting as the primary for only 1 of the groups.

IASPs coexists with HABP solutions: While it is true that IASPs provide additional availability tools on the iSeries server, IASP functions do not replace High Availability Business Partner (HABP) solutions. Customers who simply want to increase availability at a lower cost can use independent ASPs without adding disks for the backup system. Independent ASPs coexists with HABP solutions. Consider these characteristics of IASP and HABP solutions:

- **HABP solutions provide geographical dispersal of the data:** The production and backup systems can be several thousand kilometers apart. This is an important factor for effective disaster recovery. With an IASP solution, the systems must be within 250 meters of each other because of the limitations of the HSL Loop. With some V5R1 configurations, the distance is limited 15 meters.
- **An HABP solution provides switchover capability between two systems:** The HABP (High Availability Business Partner) approach for switchover and failover can be complex. Monitoring is often performed at a high level, although newer versions of the HABP products now use IBM's cluster monitor. In comparison, using switchable IASP with clustering provides a means to handle a complex requirement in a relatively simple way. The heartbeat monitoring that is implemented with IBM clustering is very sophisticated. Once properly setup, the switchover or failover to the backup system can be nearly seamless.

IASP and LPAR server consolidations

Choose one or combine

Planning is required for either environment or using separate IASPs within a single system or partition

Need for copies of programs and, or data must be analyzed

Recommendation: Contact the Rochester iSeries Technology Center

- http://www.ibm.com/eserver/iseries/service/itc/

Notes: iASP and/or LPAR consolidation

Both LPAR and iASPs can be used for degrees of consolidation. Since each business has it's own compelling reasons for consolidation and will require specific associated service level agreements, choices for certain technologies should be made on the value these technologies will have to benefit the business.

The detailed Database presentation contains a list of performance considerations when supporting multiple databases (IASPS) within the same partition or system.

Scenarios

Non-switchable simple independent disk pool

Switchable independent disk pool

- Switchable tower
- Switchable IOP with logical partitions
- Switchable tower with logical partitions

Integrated xSeries Servers and independent disk pools

Linux and independent disk pools

Notes: Scenarios

The following foils provide several implementation and usage examples of independent disk pools. They range from the ultra-simple to extremely complex. These are meant to demonstrate the flexibility and possibilities of independent disk pools in the day-to-day iSeries world.
The scenarios we will have a brief look at are:

- Non-switchable simple independent disk pool
- Switchable independent disk pool
  - Switchable tower
  - Switchable IOP with logical partitions
  - Switchable tower with logical partitions
- Integrated xSeries Servers and independent disk pools
- Linux and independent disk pools

Non-switchable simple independent disk pool

Failure in one Database does NOT disrupt other Databases

While the system is active

- Can be independently taken off-line
- Can be independently brought online
  - no initial program load (IPL) required

Data can be off-line until it is needed

- Shorten processing time for operations such as IPL and reclaim storage

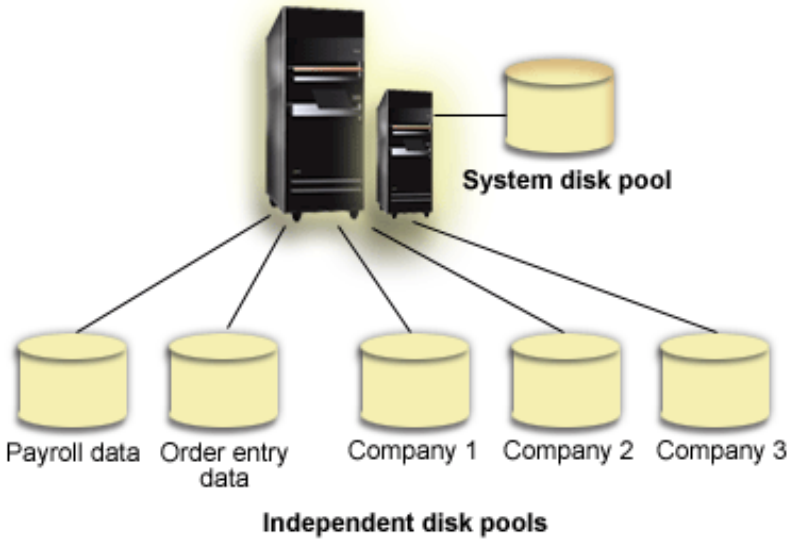

Notes: Non-switchable simple independent disk pool

In a single-system environment, a standalone, or dedicated, independent disk pool can be taken off-line independent of other disk pools because the data in the independent disk pool is self-contained. In this scenario, the user has five independent disk pools. They could represent three different applications where the third application may 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.

Failure in one Database does NOT disrupt other Databases

The most important advantage of this isolation is that should an IASP fail, that database will not be available. All other databases will not be impacted by this failure.

While the system is active

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 brought online while the system is active; no initial program load (IPL) required. Using independent disk pools this way can be very useful, for example, if you have large amounts of data that are not needed for normal day-to-day business processing. The independent disk pool containing this data can be left off-line until it is needed. When large amounts of storage are normally kept off-line, you can shorten processing time for operations such as IPL and reclaim storage.

Data can be off-line until it is needed

Data that does not need to be on-line can be brought off-line until it is needed. This way the system does not have to maintain this data and the internal structures that are associated with it. For example the system will have smaller database cross reference files for the data that is online. Besides operational advantages, this also results in a shorter processing time for operations such as IPL and reclaim storage.

Switchable independent disk pool

Switchable towers

- First tower can be switched between nodes A and B
- Second tower can be switched between nodes B and C
- Node D can only access IASP36, a private independent disk pool

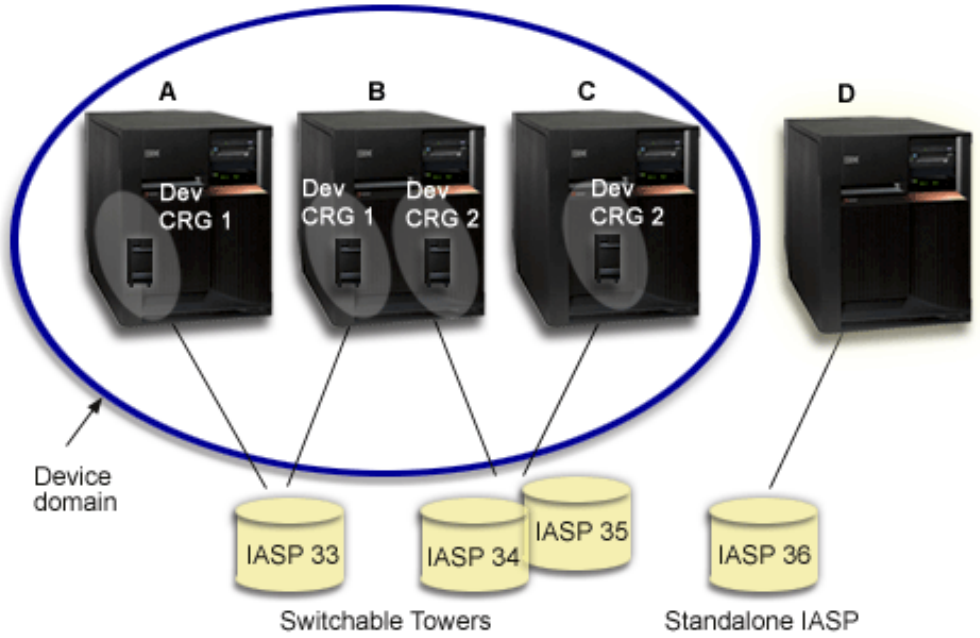

Notes: Switchable independent disk pool

In this example, the 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 towers - one contains IASP33 and the other contains IASP34 and IASP35. The tower containing IASP33 is on an HSL loop that also contains nodes A and B. This first tower can be switched between nodes A and B. The tower containing IASP34 and IASP35 could be on another HSL loop that also contains nodes B and C. This second tower 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 standalone, or dedicated, independent disk pool.

For details on cluster domains please refer to the iSeries Information Center.

Switchable independent disk pool

Switchable IOP with logical partitions

- Four logical partitions on a single iSeries server
- IOP Y is on the shared bus
  - It can be switched between all of the nodes in the cluster.
  - When the IOP is switched
  - Everything that is physically connected to that IOP is also moved to the new primary node.

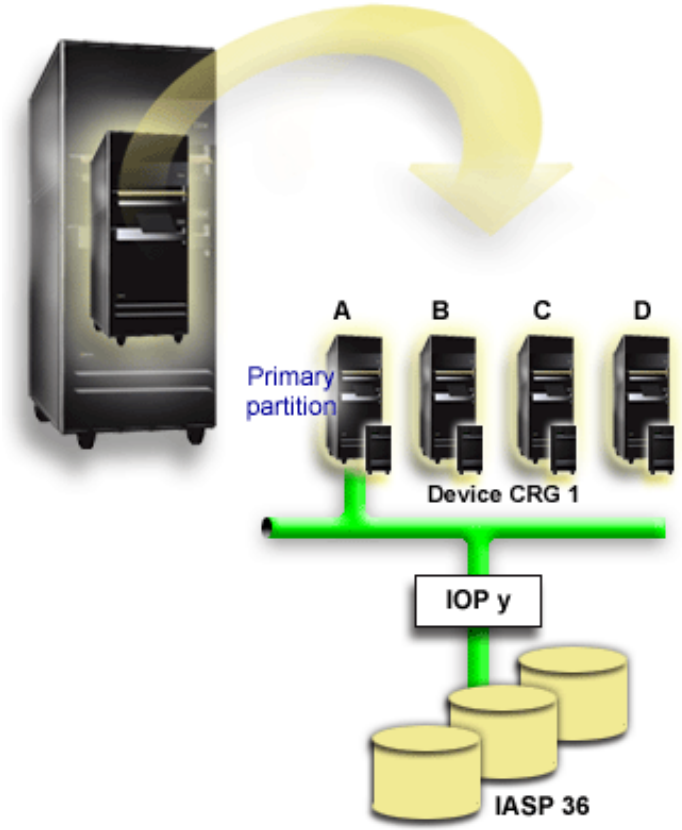

Notes: Switchable independent disk pool

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.

Remember when the IOP is switched, everything that is physically connected to that IOP is also moved to the new primary node.

Switchable independent disk pool

Switchable tower with logical partitions

- A combination of the previous two examples

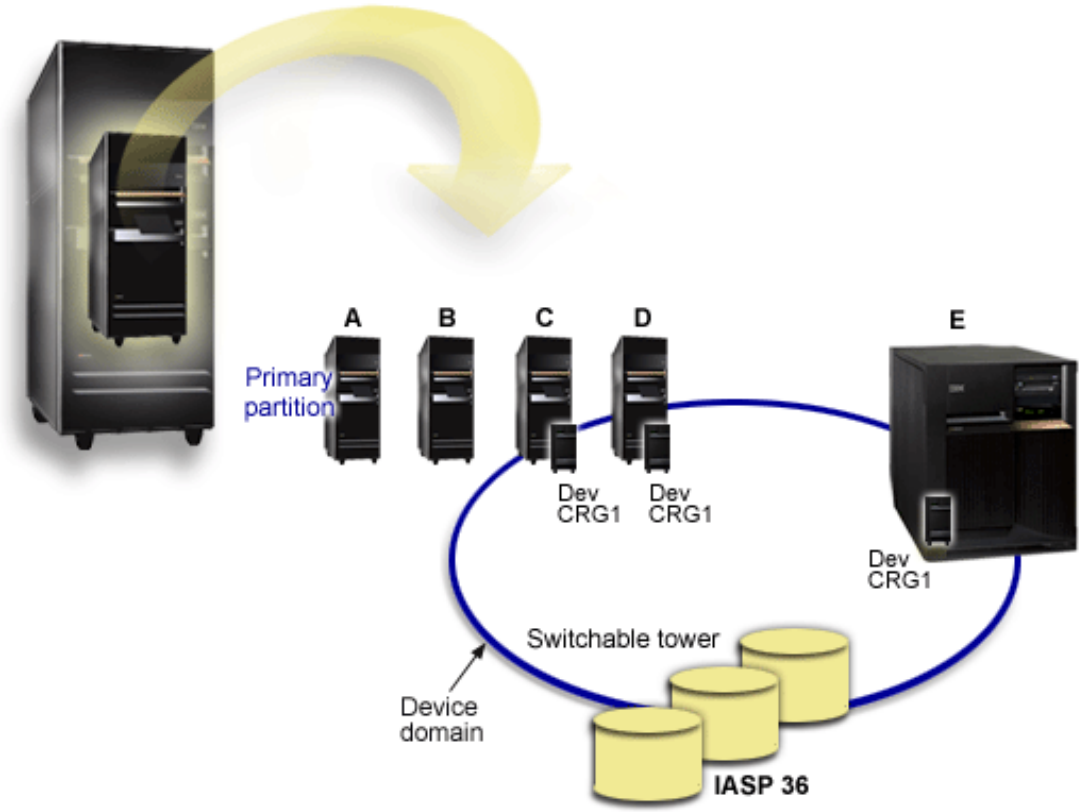

Notes: Switchable independent disk pool

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

Integrated xSeries Servers and independent disk pools

A real world example

Notes: Integrated xSeries Servers and independent disk pools

This is a real world example of using Integrated xSeries Adapters and independent disk pools. In this environment the disks containing the server storage for the IXAs and the IXAs themselves are in switchable towers and switchable independent disk pools. In the event of a failure or during routine maintenance on the primary system the independent disk pool and the IXAs can be switched to the backup system. For planned outages it allows the applications running on the IXAs to continue to be available. For unplanned outages you will have a determinable time for recovery.

Linux and independent disk pools

A higher available Linux system

Primary system with a production Linux partitions

Secondary system with backup Linux partitions

Disk pool with Linux server storage objects

Notes: Linux and independent disk pools

Linux on iSeries was introduced in V5R1. With the introduction of switchable independent disk pools you now can have a more highly available Linux "system." To use Linux with independent disk pools you need to use hosted or virtual DASD. The disk space for the Linux partition is created using network server storage spaces. One of the selling points for putting Linux on an iSeries partition was the ease of backing up the network server storage spaces and the ability to duplicate a Linux installation by copying the storage space. If that server space happens to be on an independent disk pool then you can switch that disk pool to another system and bring up that Linux system on another partition. This sheet shows a possible configuration with Linux and independent disk pools.

If your Linux disk requirements are significantly large then it is not practical to save the storage spaces and move them to a backup system. In this case a switchable independent disk pool provides for a convenient and effective way to move a Linux system from one system to another. If the network storage space is the only object on the switchable disk pool in the event of a hard failure there should be little recovery time when switching the disk pool to the backup system and bringing it up.

Remember - Planning is Critical !!!

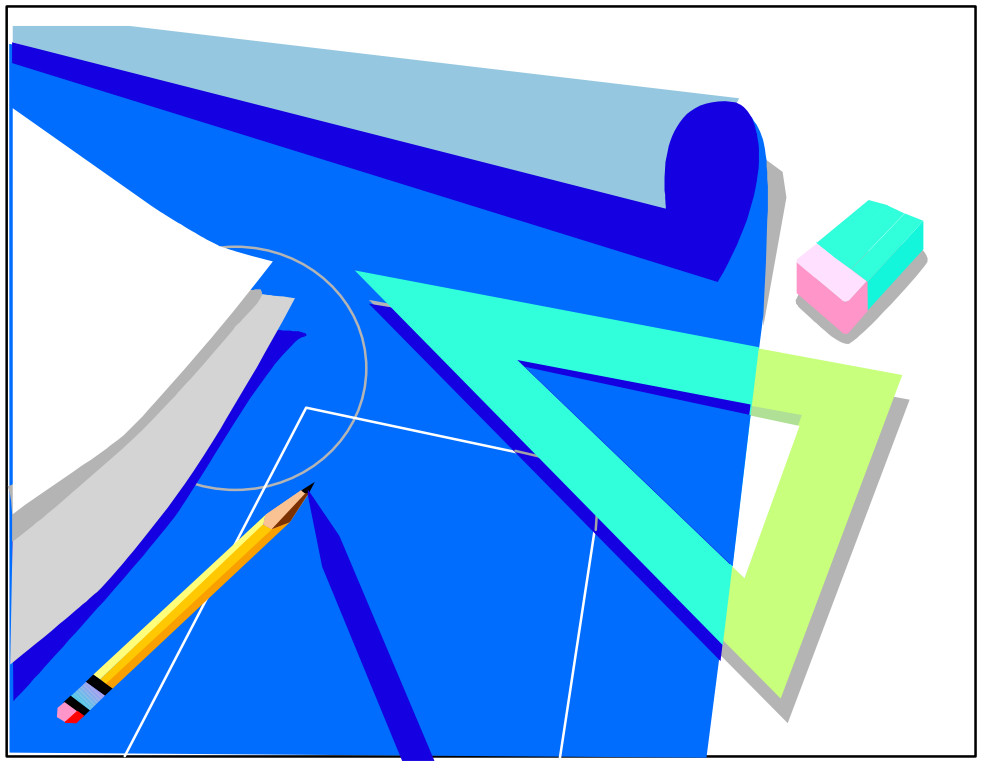

For further assistance please see:

http://www.ibm.com/servers/eserver/iseries/ha/haplanning.htm **OR**
http://www.ibm.com/servers/eserver/iseries/service/itc

Cluster Resource Services

Cluster Resource Services

Primary and secondary independent disk pools

Clustered hash table

Cluster Management Enhancements

Multiple-Release Clusters

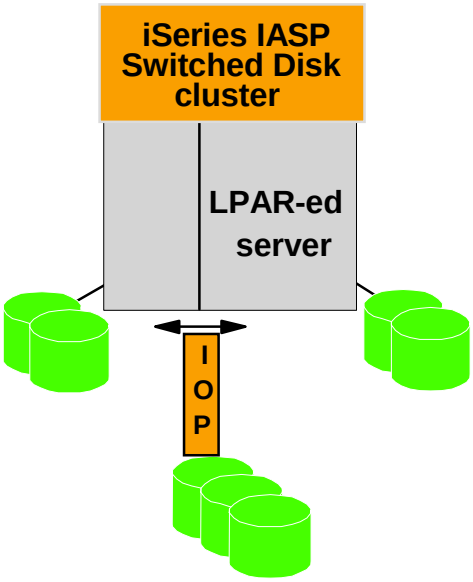

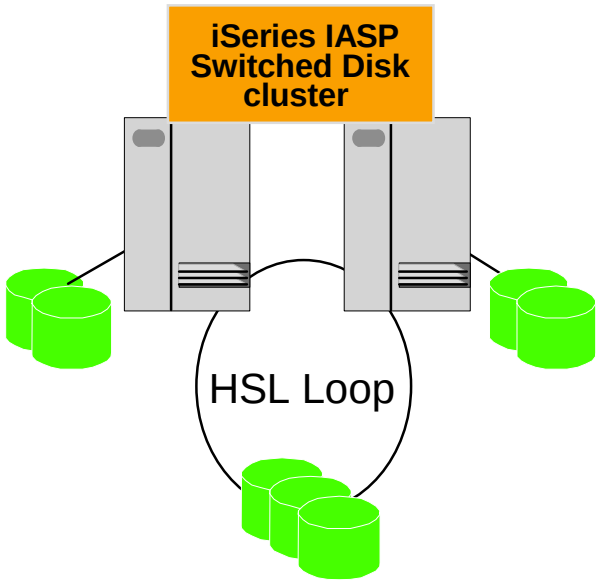

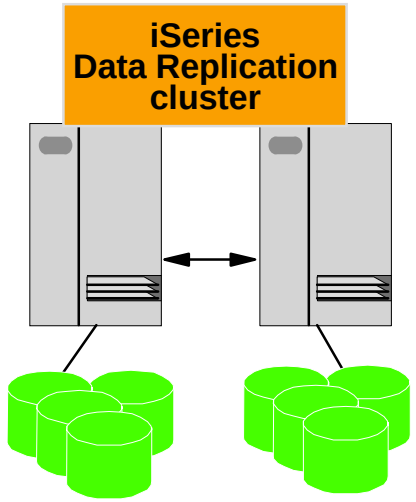

Notes: Cluster Resource Services

Primary and secondary independent disk pools

Support for library-based objects through the use of primary and secondary disk pools is present at V5R2. When independent disk pools were introduced in V5R1, they supported user-defined file systems (UDFS) only. Support for library-based objects allows independent disk pools residing on switchable devices to be composed of library-based objects.

The V5R2 level enhanced Clustered Hash Table, Cluster Management Enhancements and Multiple-Release Clusters are explained on the following foils.

Clustered Hash Table

Used by IBM ClusterProven™ Highly Available HTTP Server (Powered by Apache)

Share non-persistent data

- Available to store a CGI program written for high availability transaction state ("persistent CGI state") information
  - Requires CGI program to write to the hash table API's
- Appropriately written CGI program can maintain its state
  - Including after the application switches to a new node in the cluster
- Part of the "state replication" mechanism - part of iSeries Cluster Resource Services

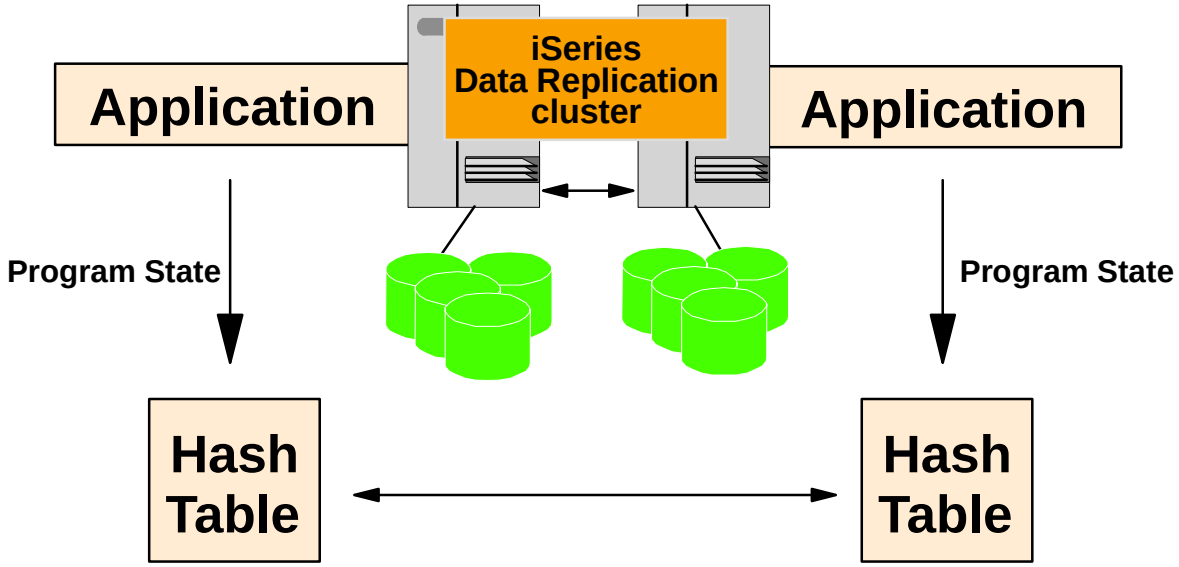

Notes: Clustered Hash Table

By using highly available Cluster hash table APIs, a CGI program can become a highly available CGI program. Highly available CGI program's state information can be saved in a clustered hash table. After a switch over, the other system participating in the highly available HTTP server recovery domain retrieves that highly available CGI program's state information from the clustered hash table. The clustered hash table is part of the state replication mechanism.

Clustered hash table

The clustered hash table APIs enable sharing and replicating of data between cluster nodes. The storage for the clustered hash table is not persistent. Not persistent means the storage for the clustered hash table is only known to the server on the local node and only available until the clustered hash table server is ended. All requests to store entries are replicated to other nodes in the clustered hash table domain. When an entry is stored, a time to live value is specified. The entry can become expired, when the time to live value has expired. Expired entries will be removed when processing various functions. For example, when adding another cluster node to the domain of an existing clustered hash table server. The existing clustered hash table entries, if any, are replicated to the cluster hash table domain node being added. Expired entries are removed from the clustered hash table during this process.

For an example on the usage of a clustered hash table, please refer to the Highly Available HTTP Server (Powered by Apache) section in this presentation.

Conflicting entries

Should you experience entry conflicts after a cluster was partitioned and then merged message CPI BD02 "Entry mismatch detected in clustered hash table &1" will be posted to the QSYSOPR message queue. You would then list the keys in the table to determine which keys contain a conflict, determine the correct information for the key and request to store it in the clustered hash table server. To list the keys in the table use the List Clustered Hash Table Keys (QcstListCHTKeys) API.

Clustered Hash Table

Two levels of Security

- Clustered hash table server
- Entry stored in a clustered hash table

No Security on the replication

- VPN recommended

See the iSeries Information Center section on clustering for the API details

Notes: Clustered Hash Table

Two levels of Security

There are two levels of security supported by a clustered hash table. One security level is associated with a clustered hash table server. This security is provided through the authorization list parameter on the STRCHTSVR command. This provides the ability to specify users that are allowed to start, end and connect to a clustered hash table server. For more details on the authorization list see the AUTL parameter on the STRCHTSVR command. The second security level is provided on an entry stored in a clustered hash table. The authority access level is specified when an entry is stored in a clustered hash table. This provides the ability to restrict access to retrieving and updating an entry. For more details on the authority access level for an entry see the Store Clustered Hash Table Entry (QcstStoreCHTEntry) API.

No Security on the replication

There is no encrypting of the information that is replicated and stored in the clustered hash table. Like other cluster messaging, the replication of the Clustered Hash Table is not secured and since the table can hold application and user data it is recommended to do this over a virtual private network.

Cluster Management

iSeries Navigator

- For Simple Clusters only

Cluster CL commands & APIs

- Examples in QUSRTOOL

Auto-Failover message queue

- Defined when the cluster is created

Self-starting a cluster node

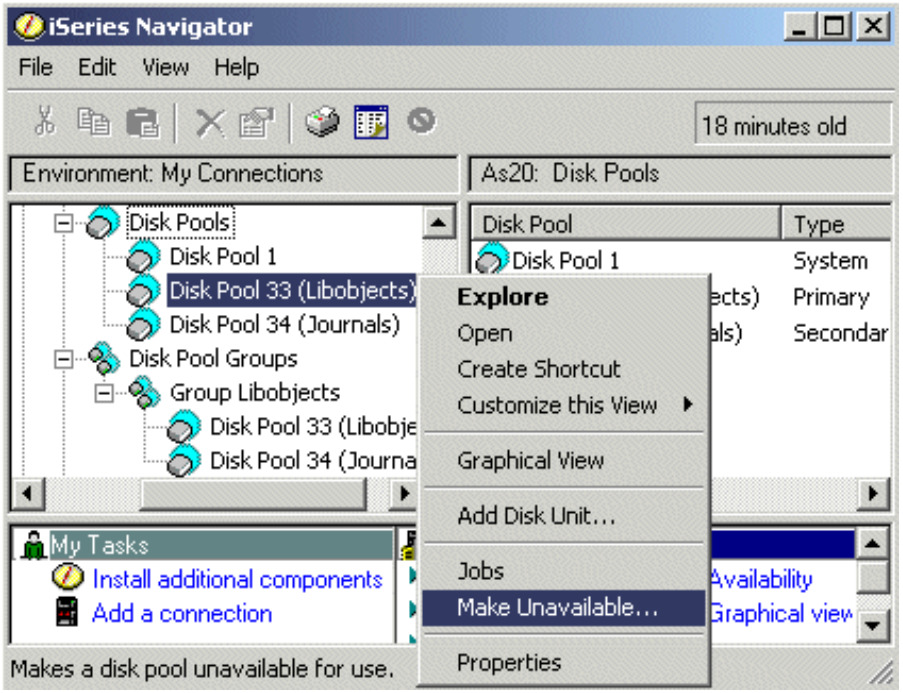

Notes: Cluster Management

IBM offers a Simple Cluster Management interface that is available through iSeries Navigator and accessible through Option 41 (OS/400 - HA Switchable Resources). This interface allows you to create and manage a cluster that uses switchable independent disk pools (switchable independent ASPs - IASPs) to ensure data availability. See iSeries Navigator in the iSeries Access presentation for more general information on iSeries Navigator.

iSeries Navigator

The iSeries Navigator and Simple Cluster Management interface does not contain all of the capabilities provided by iSeries cluster resource services. While iSeries Navigator provides many functions necessary to configure and manage a cluster, be aware that there are some capabilities that are only available through the cluster commands and APIs, or perhaps through a cluster middleware business partner application, depending upon the particular application. For example, the iSeries clustering architecture supports up to 128 nodes in a cluster, however the iSeries Navigator interface only supports up to four nodes in a cluster. With iSeries Navigator, you can create a simple cluster consisting of one or two nodes. Once you have established a cluster in iSeries Navigator, you can then add a node to an existing cluster, up to as many as four total nodes. If your clustering needs exceed this, you should consider using the full set of OS/400 cluster commands and APIs or cluster middleware business partner products.

Cluster CL commands

Cluster control language (CL) commands and APIs have been added to allow system programmers and system administrators easier access to cluster capabilities. They are design for configuring, activating, and managing a cluster and also nodes and cluster resource groups in a cluster.
Cluster resource services also provides a set of example commands in the QUSRTOOL library that map to the CL commands and APIs mentioned above. The QUSRTOOL commands might be useful in some environments. For example, one can easily set up a cluster for testing cluster-enabled applications. See the member TCSTINFO in the file QUSRTOOL/QATTINFO for more information on these example commands.

Auto-Failover message queue

The auto-failover message queue receives messages regarding failover activity. Using the failover message queue allows an administrator to be notified before a failover occurs. This gives the administrator the ability to cancel the failover if the desired behavior is to prevent the failover at this time. The failover message queue is defined when creating a cluster resource group using the Create Cluster (CRTCLU) command or the Create Cluster (QcstCreateCluster) API. It can also be modified using the CL command and API for changing a cluster resource group. The failover message queue cannot be used with the iSeries Navigator Simple Cluster Management interface.

Self-starting a cluster node

A node can start itself and can rejoin the current active cluster, provided it can find an active node in the cluster. Prior to V5R2 when a node needed to be started and rejoined this had to be initiated from an active node in the cluster. With V5R2 this can also be done from the node that needs to be started and rejoined by the use of the STRCLUNOD cmd or the Start Cluster Node API (QcstStartClusterNode).

Multiple-Release Clusters

From V5R1 to V5R2

- UDFS disk pools only

From V5R2 to V5R1

- Not possible due to internal changes

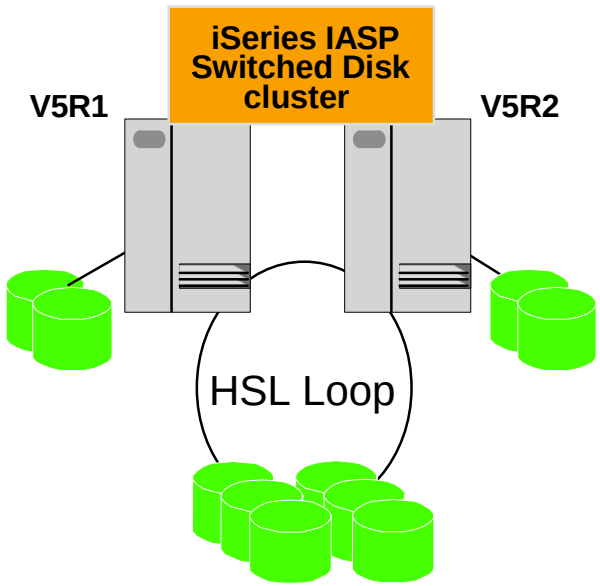

Notes: Multiple-Release Clusters

A cluster version represents the level of function available on the cluster. Versioning is a technique that allows the cluster to contain servers at multiple release levels and fully interoperate by determining the communications protocol level to be used. If you are implementing a cluster that will contain servers of varying release levels, see Multiple-release clusters. There are actually two cluster versions:

1. Potential cluster version
   Represents the most advanced level of cluster function available for a given node. This is the version at which the node is capable of communicating with the other cluster nodes.
2. Current cluster version
   Represents the version currently being used for all cluster operations. This is the version of communications between the nodes in the cluster.

The potential cluster version is incremented on every OS/400 release which has significant new clustering functionality not available in earlier cluster versions. If the current cluster version is less than the potential cluster version, then that function cannot be used since some nodes would not be able to recognize or process the request. To take advantage of such new function, every server in the cluster would need to be at the same potential cluster version and the current cluster version must also be set to that level.

If creating a cluster that will include nodes at multiple cluster versions, then certain steps are required when you create the cluster. By default, the current cluster version will be set to the potential cluster version of the first node added to the cluster. This approach is appropriate if this node is at the lowest version level to be in the cluster. However, if this node is at a later version level, then you will subsequently be unable to add nodes with a lower version level. The alternative is to use the target cluster version value on create cluster to set the current cluster version to one less than the potential cluster version of the first node added to the cluster. For example, consider the case where a two-node cluster is to be created. The nodes for this cluster are:

Node Identifier	Release	Potential Cluster Version
Node A	V5R1	2
Node B	V5R2	3

If the cluster is to be created from Node B, care must be taken to indicate that this will be a mixed release cluster. The target cluster version must be set to indicate that the nodes of the cluster will communicate at one less than the requesting node's potential node version.

Important

Once an independent disk pool is made available on a V5R2 system, it cannot be made available on a V5R1 system. It is possible to switch a V5R1 independent disk pool to a V5R2 system and make it available. However, once it has been made available on the V5R2 system, its internal contents are changed, and it cannot be made available to the V5R1 system again.

Highly Available HTTP Server

(Powered by Apache)

Highly Available HTTP Server (Powered by Apache)

Takes advantage of iSeries Cluster technology

- Clustered Hash Table, and other iSeries cluster capabilities

Supported models

- Primary/backup with takeover IP model
- Primary/backup with a network dispatcher model
- Peer model

Greatly assists developing Highly Available websites

- ClusterProven status

Notes: Highly Available HTTP Server (Powered by Apache)

IBM HTTP Server for iSeries includes the latest version of Apache Software Foundation's Apache Web server V2.

Highly Available HTTP Server (Powered by Apache) takes advantage of iSeries Clustering technology and makes it possible to build a highly available Web site, improving the availability of business-critical Web applications built with Common Gateway Interface (CGI) programs.

Highly available HTTP servers provide function that monitors a URL that is part of your Web site (for example, your home page) and will take recovery action if the Web server is no longer serving your Web content. For example, the monitor function will try to end and start your Web server or may initiate a switchover to move your HTTP server function to the backup node in the cluster. In addition, you may choose to write your CGI programs using highly available CGI APIs to save the state of each CGI program instance into the Clustered Hash Table within the iSeries cluster. In the event of a failed node in the cluster, an appropriately -written CGI program can maintain its state, even after the application switches to a new node in the cluster.

The following foils give some simple examples of networks using the V5R1 high availability features with the HTTP Server for iSeries Powered by Apache.

It is the responsibility of the CGI programmer to take advantage of this clustered hash table support.

Takeover IP model

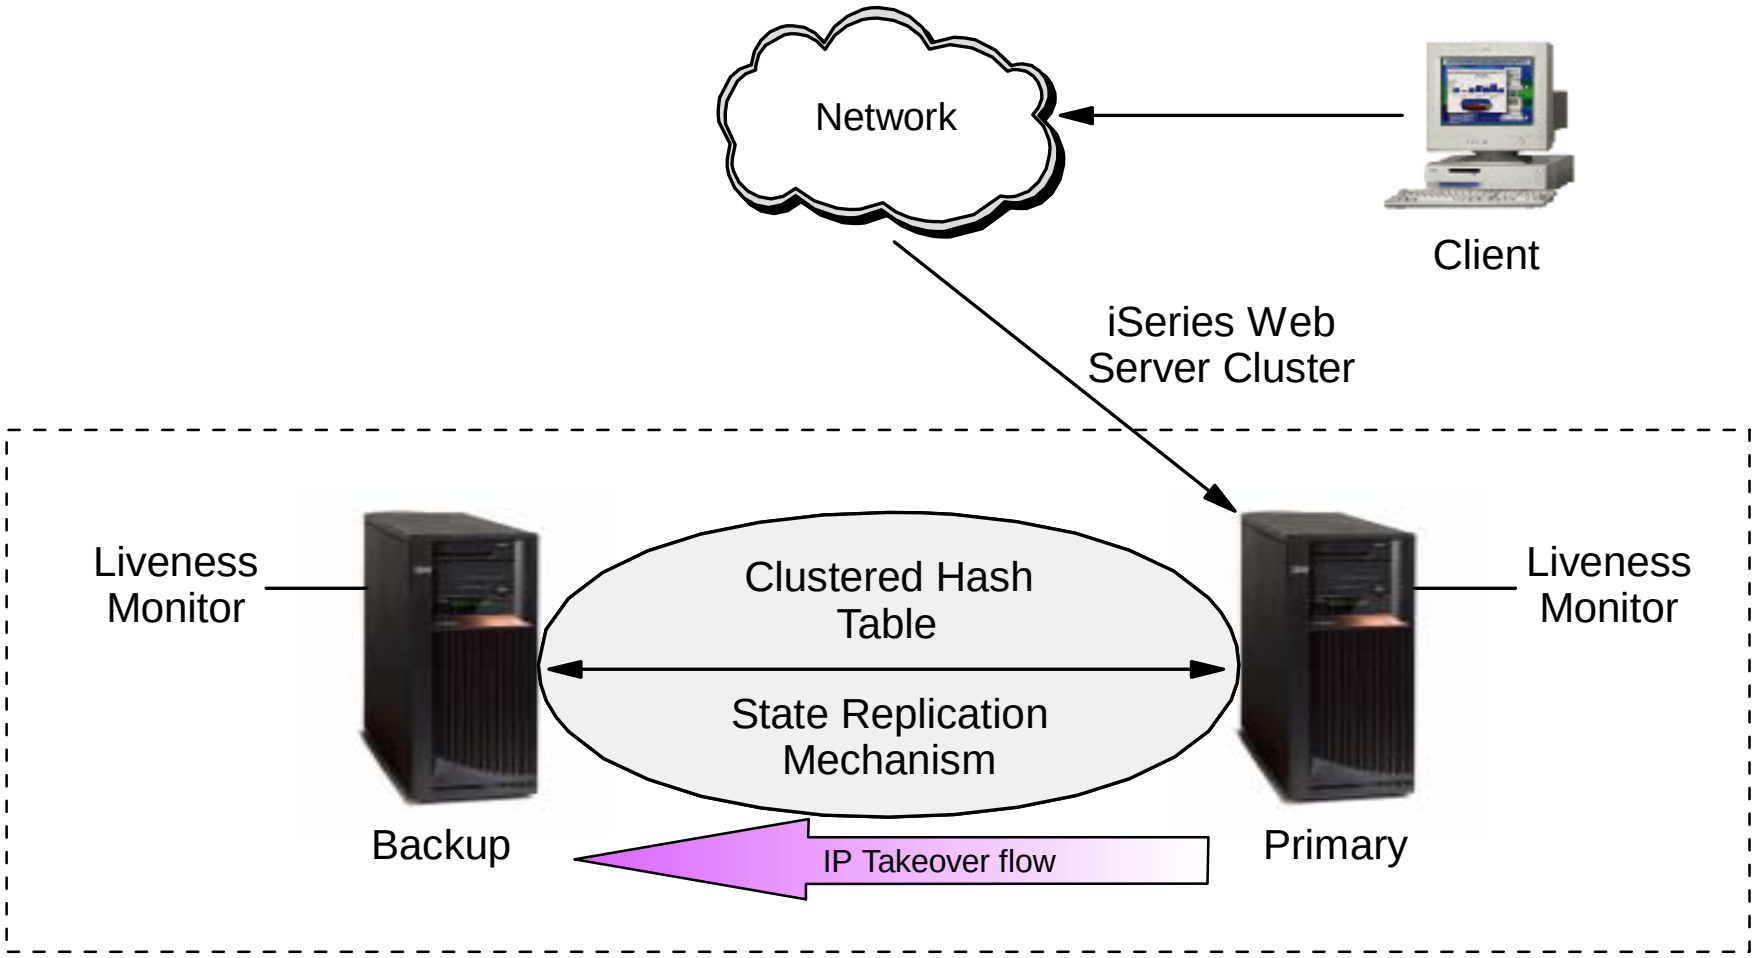

Notes: Takeover IP model

In this model, the Web server runs on the primary and all backup nodes. The backup node or nodes are in a idle state, ready to become the primary Web server should the primary Web server fail (failover), or a switchover takes place. **All client requests are always served by the primary node.**

When the primary node fails (failover), or is brought down by the administrator, the failover/switchover process begins. The following steps are performed during failover/switchover:

1. The "Liveness Monitor" using iSeries Cluster support is monitoring the active status of each node.
2. One of the backup servers becomes the primary (the first backup in the switchover order).
3. The client requests are redirected to the new primary node.
4. If the new primary receives a user request that belongs to a long-running-session (a CGI program updated to be a highly available CGI program), the server will restore the request's state. The new primary retrieves that highly available CGI program's state information from the clustered hash table. The clustered hash table is part of the state replication mechanism.
5. After the failed node recovers, the highly available Web server instance can be restarted and it will become the backup system. If the system administrator wants the failed node to become primary again, a manual switchover must be performed (this can be accomplished with the IBM Simple Cluster Management interface available through Operations Navigator or a business partner tool).

Network Dispatcher Model

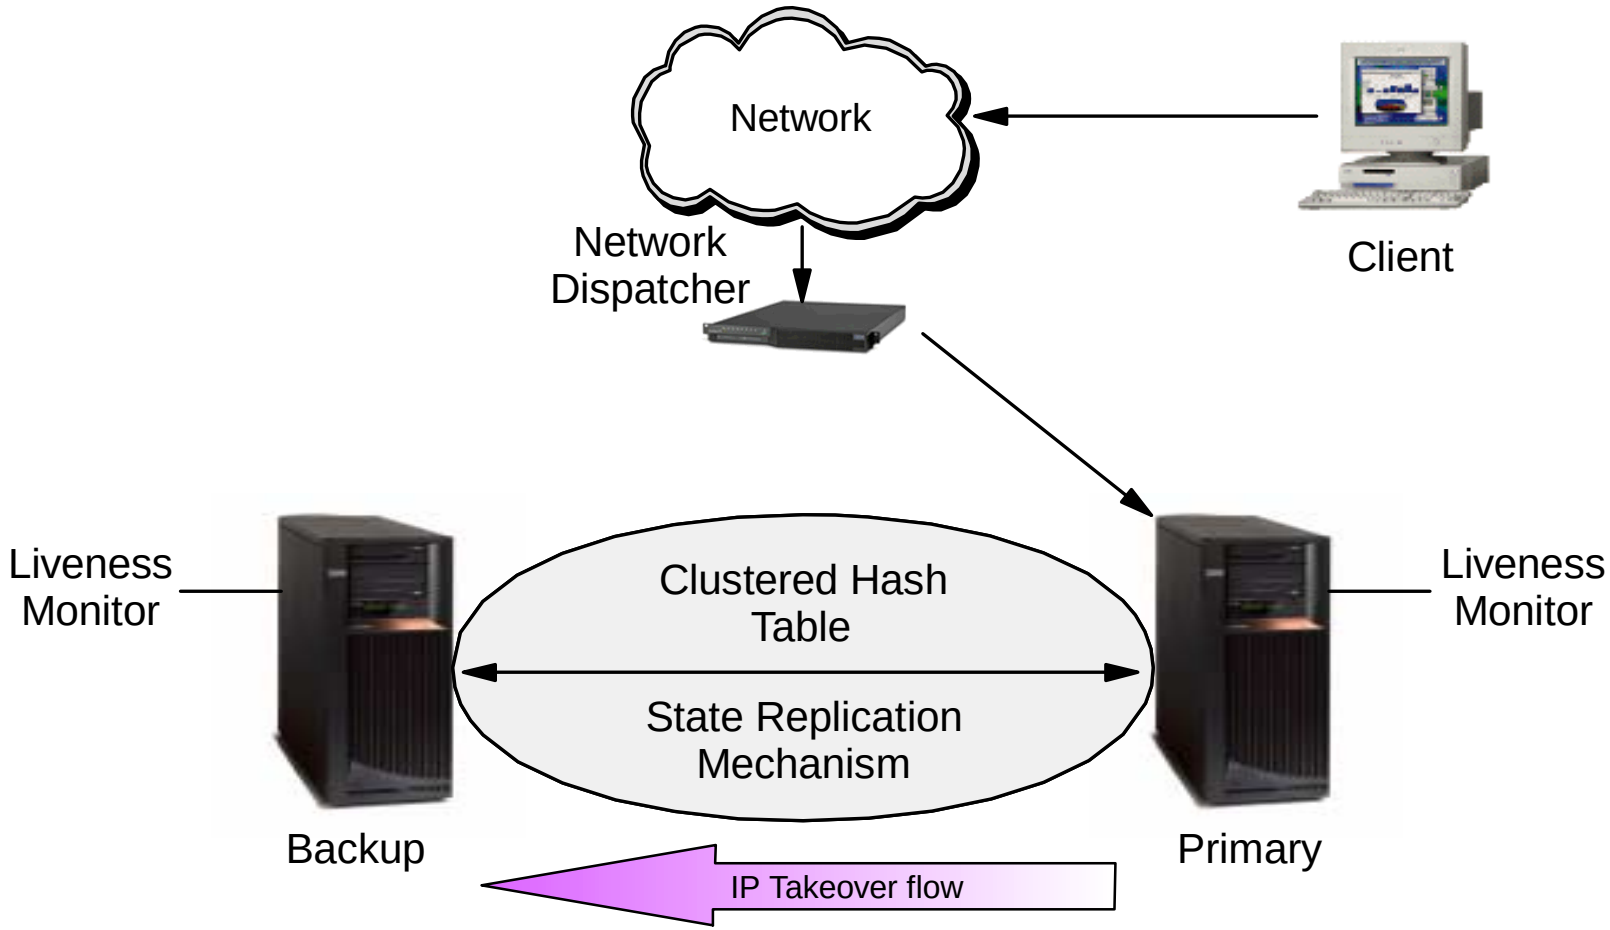

Notes: Network Dispatcher Model

In this model, just like the primary/backup with takeover IP model, the Web server runs on the primary and all backup nodes. The backup nodes are in an idle state and all client requests are served by the primary node. A network dispatcher (for example the IBM WebSphere Edge Server) sends client requests to the primary Web server.

When the primary node fails (failover), or a switchover takes place, the failover/switchover process begins. The following steps are performed during failover/switchover:

1. One of the backup servers becomes the primary (the first backup in the switchover order).
2. The client requests are sent to the new primary node by the network dispatcher.
3. If the new primary receives a user request that belongs to a long-running-session, the server needs to restore the request's state. The new primary searches for the state either locally or in the clustered hash table. The clustered hash table is part of the state replication mechanism.
4. After the failed node recovers, the system administrator can restart the Web server instance and it will become a backup Web server. If the system administrator wants the failed node to become primary again, a manual switchover must be performed.

Note: A node can join a recovery domain as primary only if the cluster resource group is in inactive mode.

Peer model

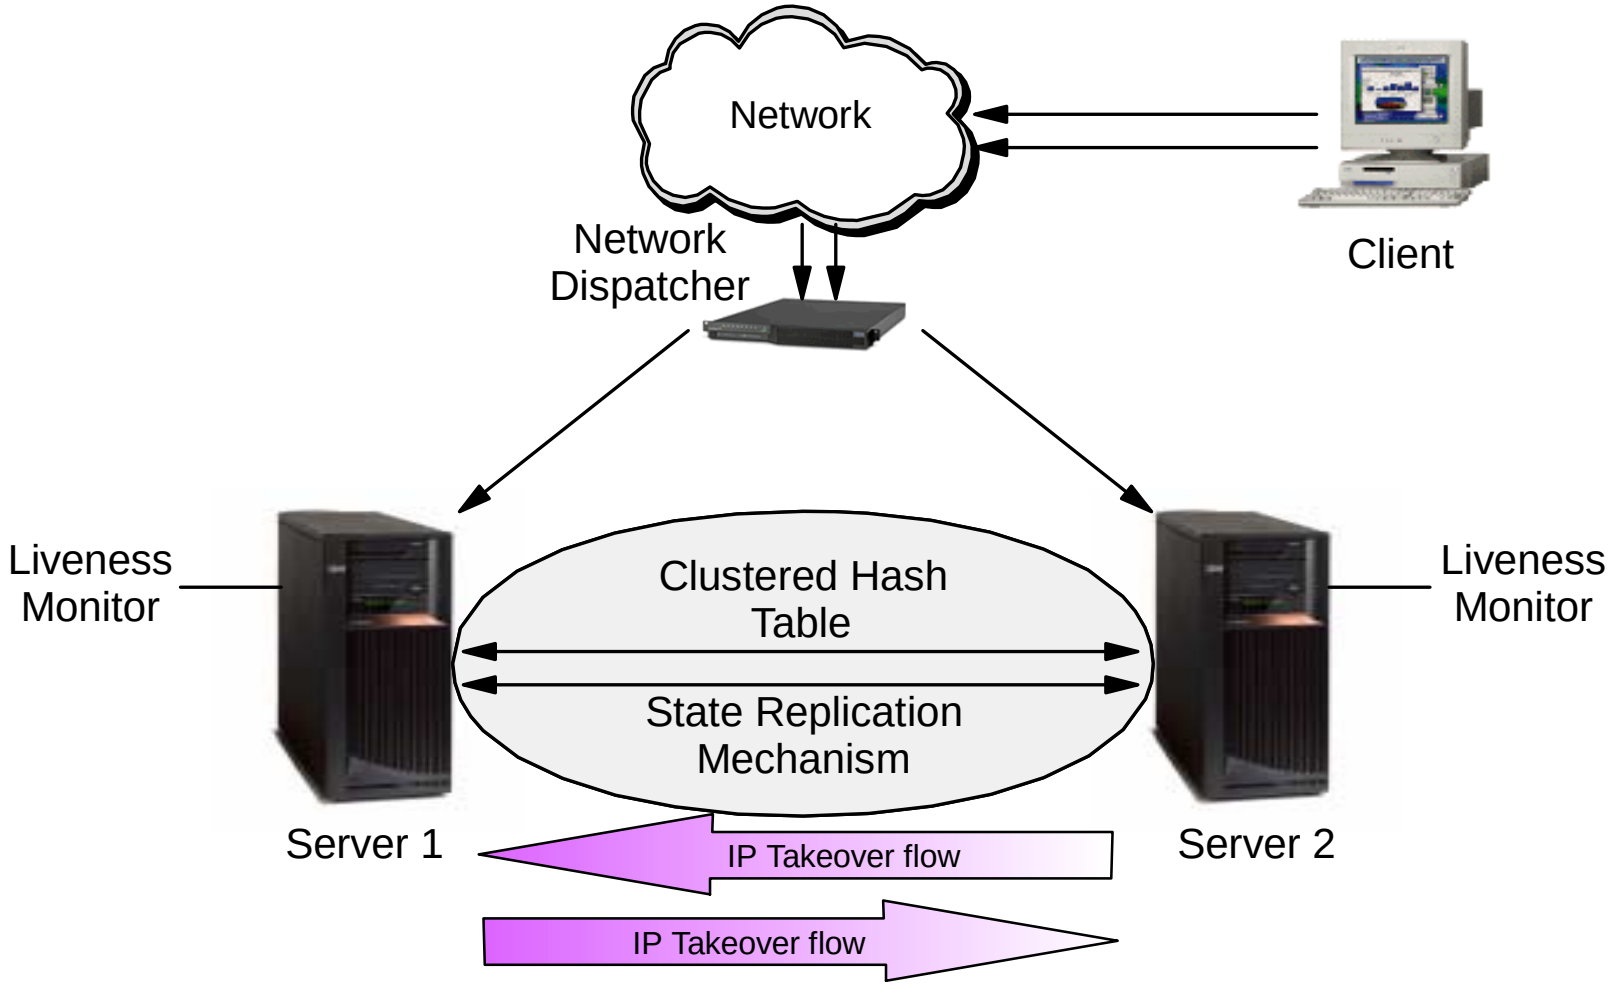
Network
Client
Network
Dispatcher
Liveness
Monitor
Clustered Hash
Table
State Replication
Mechanism
Liveness
Monitor
Server 1
IP Takeover flow
Server 2
IP Takeover flow

Notes: Peer model

In this model, there is no declared primary node. All nodes are in an active state and serve client requests. A network dispatcher (for example the IBM WebSphere Edge Server) evenly distributes requests to different cluster nodes. This guarantees distribution of resources in case of heavy load. Linear scalability is not guaranteed beyond a small number of nodes. After some number of nodes are added, scalability can disappear, and the cluster performance can deteriorate.

In the event that one node fails (failover), the failed Web server traffic is routed to one of the other operational Web servers according to the configuration of the network dispatcher.

Proxy ARP Routing for Virtual IP interfaces

Fault tolerance and high availability

- *VIRTUALIP address on the same network as defined on other physical interfaces

iSeries will respond to ARP requests for virtual interfaces

Eliminating

- The need to define static routes, or
- Run a dynamic routing protocol such as RIP

Example scenarios

- Adapter failover without local clients
- Adapter failover with local clients
- Adapter failover with multiple iSeries and local clients

Notes: Proxy ARP Routing for Virtual IP interfaces

In support for fault tolerance and high availability, the iSeries in V5R2 allows you to define a *VIRTUALIP address on the same network as defined on other physical interfaces. In this configuration, the iSeries will respond to ARP (Address Resolution Protocol) requests for virtual interfaces, eliminating the need to define static routes or run a dynamic routing protocol such as RIP.

This enables a new way to approach load balancing. In the event of an adapter failover you can now use proxy Address Resolution Protocol (ARP), along with virtual IP, to provide seamless availability to clients. This enables new scenarios where you provide network availability to the iSeries for its clients in the event of an adapter failure.

Adapter failover using Virtual IP and Proxy ARP

- Situation
  Your production iSeries handles data entry from both remote and LAN clients. It has the company's critical application on it. As the company has grown, so has its demand on the iSeries and the network. Because of the growth, it has become imperative that this iSeries be available on the network *without an unscheduled down time*. If, for any reason, a network adapter becomes unavailable, other network adapters on the iSeries should take over and the network clients should be unaware of any failures.
- Objectives
  The concept of availability has many different aspects of redundancy and backup for failing components. In this scenario, the goal is to provide network availability to the iSeries for its clients in the event of an adapter failure.

The following foils show the following example scenarios:

- Adapter failover without local clients
- Adapter failover with local clients
- Adapter failover with multiple iSeries and local clients

Adapter failover without local clients

Each of these physical connections would have a different IP address

- Assign a virtual IP address to the system

All remote clients will communicate through a external load balancing server

- Network dispatcher will route the virtual IP addresses to one of the network adapters

Local clients would unnecessarily overload the network dispatcher

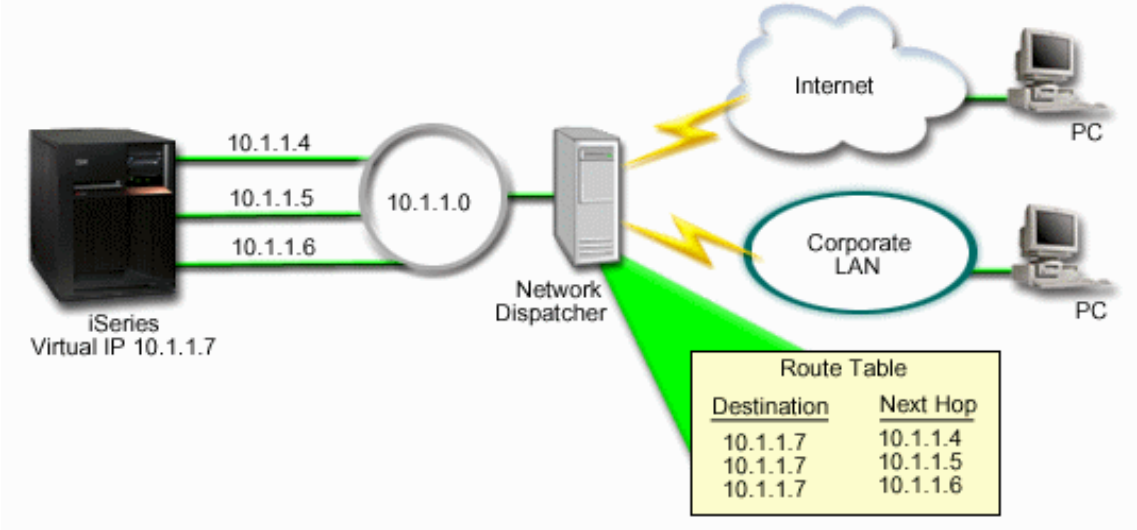

Notes: Adapter failover without local clients

In this example, each of the 3 physical connections have a different IP address (10.1.1.4 - 10.1.1.6). Then assign virtual a IP address (10.1.1.7) to the system. **This virtual IP address is the IP address that all of its clients will recognize it by**.

All remote clients (clients that are not physically attached to the same LAN as the iSeries) will communicate with the iSeries through a external load balancing server such as a network dispatcher. When the IP requests from the remote clients go through the network dispatcher, the network dispatcher will route the virtual IP addresses to one of network adapters on the iSeries.

If the LAN that the iSeries is connected to has clients, this configuration enables these clients to not use the network dispatcher to direct their locally bound traffic. By eliminating the need for local clients to go through the network dispatcher, you help ensure the network dispatcher is not unnecessarily overloaded.

You could create route entries on each client that were similar to the route tables in the network dispatcher, but given a large number of clients, this would be a very impractical to carry out. This situation is addressed in the following scenario foil.

Adapter failover with local clients

Route entries on each client

- Given a large number of clients, this is very impractical to put into effect

As of OS/400 V5R2:

- Local clients can now connect to the iSeries virtual IP address through ARP
- Enables local clients to have an adapter failover solution as well

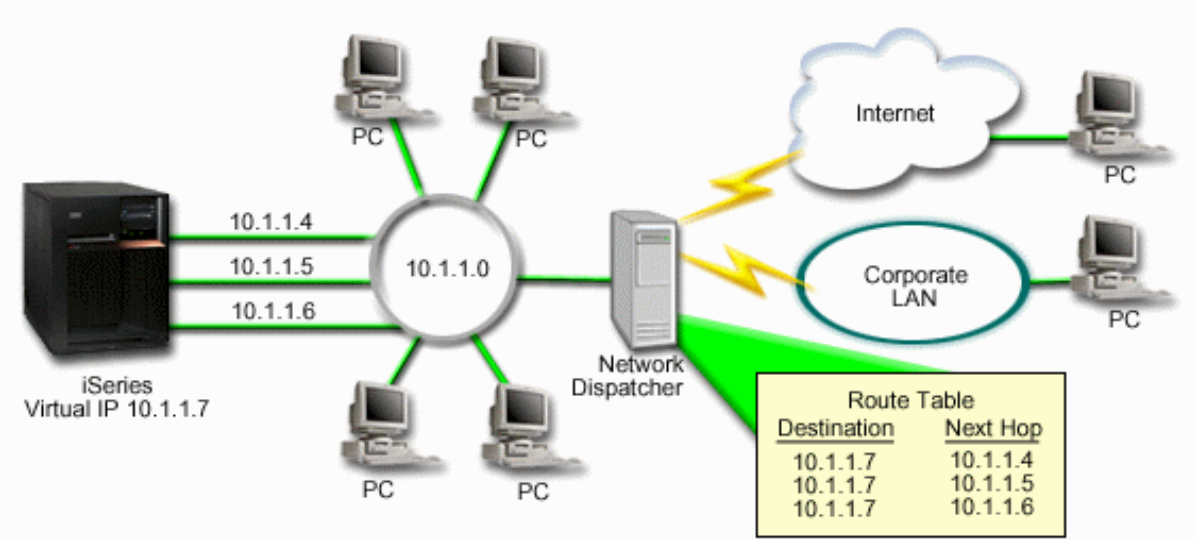

Notes: Adapter failover with local clients

As of OS/400 V5R2, local clients (clients that are attached to the same LAN as the iSeries) can now connect to the iSeries virtual IP address through ARP. This allows local clients to have an adapter failover solution as well.

The solution can also involve using two or more iSeries servers to support each other. If one of the iSeries systems become unavailable, then the second system can serve as a failover. The following scenario shows the same setup using two iSeries servers.

Adapter failover with multiple iSeries and local clients

The packet routing is the same as routing for a single iSeries and its remote clients

Distinct difference for the local clients

- With multiple iSeries using the same virtual IP address
  - Can only proxy for one of the iSeries
  - In this example: iSeries with the two LAN connections serves as the proxy.

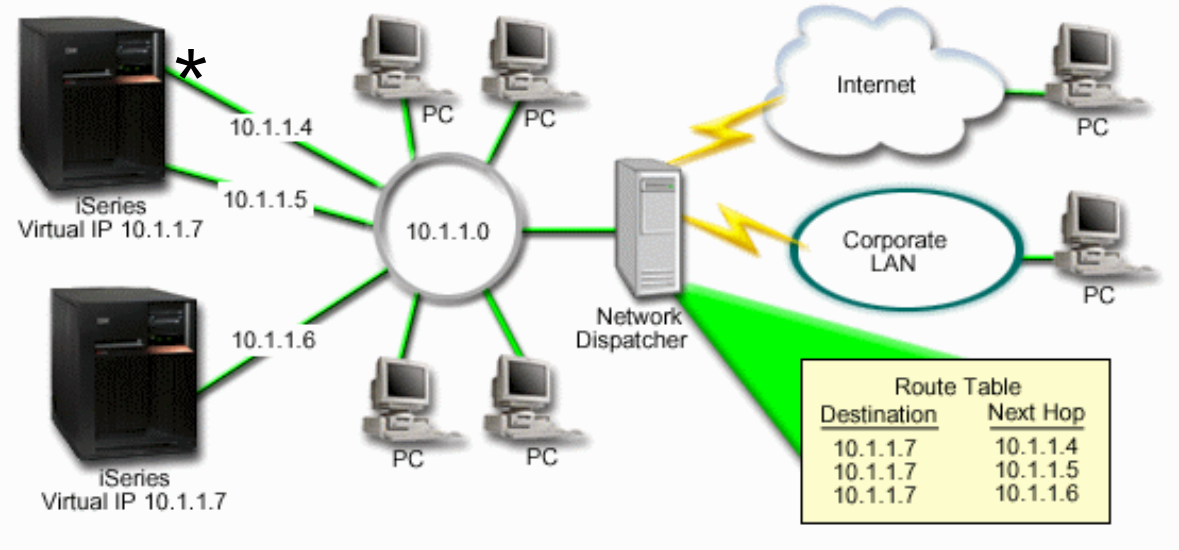

* Proxy server system

Notes: Adapter failover with multiple iSeries and local clients

The packet routing is the same as routing for a single iSeries and its remote clients; however, there is a distinct difference for the local clients. If you have multiple iSeries using the same virtual IP address, you can only proxy for one of the iSeries. In this case, you would have the iSeries with the two LAN connections serve as the proxy.

Journal Enhancements

Journal Enhancements

HA Journal Performance (OS/400 option 42)

Support for library-capable independent disk pools

Fixed-length options for journal entries

Delay automatic journal change

Delay the next attempt to delete a journal receiver

HA Journal Performance (OS/400 option 42)

Journal Caching feature for performance

- PRPQ 84486 now part of OS/400 option 42
- Significantly reduces journaling synchronous disk write operations
- Considerations
  - Not recommended when most recent change must be "maintained"
  - May not be suitable for interactive applications where single system recovery is the primary reason for journaling

Journal Standby feature for Faster Switchover

- Used on the Backup system
- Skip starting journaling for all objects
- Minimizes journaling overhead
  - Most journal entries are not deposited when the journal is in standby state
  - No errors indicating that the entry was not deposited and no error messages are sent to the application

Notes: HA Journal Performance (OS/400 option 42)

For V5R2, there are a number of improvements and additions to journal management. For the most demanding high-availability clustering environments supported by our high availability Business Partners, Journal Standby Mode and Asynchronous Journaling capabilities enable faster failover and reduce performance bottlenecks. Both the Journal Caching feature and the Journal Standby feature are provided by installing OS/400 option 42.

The Journal Caching feature (PRPQ 84486 before V5R2)

The journal caching feature allows batch applications to substantially reduce the number of synchronous disk write operations performed thereby reducing overall elapsed batch run time. You can specify journal caching with the JRNCACHE parameter on the Create Journal (CRTJRN) or Change Journal (CHGJRN) commands. Journal caching provides significant performance improvement for batch applications which perform large numbers of add, update, or delete operations against journaled objects. Applications using commitment control will see less improvement (commitment control already performs some journal caching). Journal caching modifies the behavior of traditional non-cached journaling in batch. Without journal caching, a batch job waits for each new journal entry to be written to disk. Journal caching allows most operations to no longer be held up waiting for synchronous disk writes to the journal receiver. Journal caching is especially useful for situations where journaling is being used to enable replication to a second system.

- **Important**
  Journal caching is not recommended if it is unacceptable for you to lose even one recent change in the event of a system failure - where the contents of main memory are not preserved. This type of journaling is directed primarily toward batch jobs and may not be suitable for interactive applications where single system recovery is the primary reason for using journaling.

The Journal Standby feature

You can use the CHGJRN command to put the journal in standby state. You might want to put a journal in standby state if the journal is on a backup system. By having the journal in standby state, a switchover to the target system can be accomplished more quickly because all objects on the backup system can be journaled thus allowing the switchover processing to skip the costly step of starting journaling for all objects. At the same time though, the backup system is not incurring the overhead of journaling because most journal entries are not deposited when the journal is in standby state.

With the journal state in standby, journal entries are not deposited into the journal. If an attempt is made to deposit an entry into the journal, there will be no errors indicating that the entry was not deposited and no error messages are sent to the application.. While in *STANDBY state, journaling can be started or stopped. However, using commitment control is not allowed while in *STANDBY state.

Notes: See the Backup and Recovery information in the iSeries Information Center at http://www.iseries.ibm.com/infocenter for the journal entry types that can be deposited in this state.

More Journal Enhancements...

Support for library-capable independent disk pools

- Use a secondary disk pool in a disk pool group for isolation

Fixed-length options for journal entries

- audit security related activity for journaled objects

Delay automatic journal change

- Wait before automatically attach a new journal receiver

Delay the next attempt to delete a journal receiver

Notes: More Journal Enhancements...

Support for library-capable independent disk pools

Starting with V5R2 you can journal objects on library-capable independent disk pools.

Fixed-length options for journal entries

You can use the Fixed Length Data (FIXLENDTA) parameter of the Create Journal (CRTJRN) and Change Journal (CHGJRN) commands to audit security related activity for journaled objects on your system.

With the FIXLENDTA parameter, you can specify that the following data is included in the journal entries that are deposited into the attached journal receiver: Job name, User profile name, Program name, Program library name, System sequence number, Remote address, Thread identifier, Logical unit of work identifier, Transaction identifier

Delay automatic journal change

You can use the Manage Receiver Delay Time (MNGRCVDLY) parameter of the CHGJRN or CRTJRN commands to cause the system to wait the length of time that you specify before its next attempt to automatically attach a new journal receiver. See Manual versus system journal-receiver management for details in the iSeries Information Center.

Delay the next attempt to delete a journal receiver

Use the Delete Receiver Delay Time (DLTRCVDLY) parameter of the CHGJRN or CRTJRN commands to cause the system to wait the length of time that you specify before its next attempt to automatically delete a journal receiver. See Automatic deletion of journal receivers for details in the iSeries Information Center.

Windows and xSeries Integration

xSeries & Windows Integration

Windows Clustering Service

- Must reside within a single iSeries partition
- IXS or IXA resource type of 2890, 2892, or 2689
- Running either Windows 2000 Advanced Server or Windows .NET Enterprise Server
  - IBM plans support for Windows .NET Standard and Enterprise Server when available from Microsoft
- 16 New Shared Storage Spaces

Virtual LAN support

- TCPPORTCFG parameter on the CRTNWSD command

Support for Auto Cartridge Loader (ACL) tape devices

- 3570, 3580 and 3590

iSeries Navigator has install/uninstall assistance for new release, service pack, and hot fixes.

Notes: xSeries & Windows Integration

Support for Windows Clustering Service

Support for installing and configuring the Windows Cluster Service on the iSeries Integrated Windows Servers has been added in V5R2. Support is available only for Windows servers that have an Integrated xSeries Server or Adapter with a resource type of 2890, 2892, or 2689; running either Windows 2000 Advanced Server or Windows .NET Enterprise Server. Official support for Windows .NET Standard and Enterprise Server will be provided within 30 days of their general availability.

Note:

- Windows clustered network server nodes must reside within a single iSeries partition in order to be clustered.
- IXS code for Windows NT 4.0 remains in V5.2 but on an 'as-is' basis only, no support is planned.
- Windows 2000 Advanced Server supports a two-node cluster while Windows .NET Enterprise Server supports four-node clusters.
- Datacenter versions of Windows are not supported.

Although the traditional Windows clustered server solution requires a shared physical SCSI or Fibre Channel device, the Integrated Windows Server solution uses a virtual Fibre Channel bus to share the virtual disk devices between the nodes of a cluster.

Support for linking up to 16 shared disk drives for clustered Windows servers which increases the total number of disk drives on clustered servers to 48. Windows system drive size has increased to 64000 MB for supported Integrated xSeries Servers and Integrated xSeries

In addition, the new support for Virtual Ethernet enables high-performance; secure communication for the internal node-to-node communication between clustered nodes.

Virtual LAN support

The iSeries virtual LAN allows IXSs to securely "talk" to each other internally through OS/400 rather than go out on an external network and then back into the iSeries. The performance of this virtual LAN for IXSs (internal and external) is improved. Each Windows server will participate as a peer on the virtual ethernet network that interconnects on the iSeries (interconnecting amongst LPARs also). This results in a powerful and flexible internal network structure that can interconnect instances of OS/400, Linux, and Windows.

You can configure the Windows server to use virtual LAN on the TCPPORTCFG parameter in the CRTNWSD command.

Support for Auto Cartridge Loader (ACL) tape devices

Customers who have systems with large amounts of data often have Auto Cartridge Loader (ACL) tape devices (3570, 3580 and 3590) which have the ability to load another tape cartridge when done with the current one. Support is now added in the IXS tape device driver to handle commands for ACLs so multiple tape cartridges can be accessed during backup/restore operations initiated from the Windows Server.

Backup & Recovery

New 30/60 GB Internal QIC Tape Drive

SLR60 technology

- Up to 10x speed of 4GB tape drive, up to 2x speed of 25GB tape drive
- 30GB (uncompressed) or 60 GB (with 2x compression) per tape
- Planned availability June 14, 2002
- Supported by OS/400 V4 and V5
  - Supported on 270, 820, 830, 840, 890, 250, 170, 150, 6xx, Sxx 7xx

SLR60 is planned to be the last IBM QIC drive announced which will read the older low capacity media 2GB, 4GB, 2DC and 4DC

New 30/60 GB QIC Tape Drive - Details

FC # of SLR60 Drive	iSeries or AS/400 Model system unit into which the 30/60 QIC is installed	I/O Tower or expansion unit into which the 30/60 QIC is installed
#4584	270, 820	
#4684	830, 840, 890	5065, 5066, 5074, 5079
#6384	150, 170, 250,640, 650, S30, S40, 730, 740	5077, 9077, 5072, 5073, 5082, 5083
#6484	600, S10, 620, S20, 720	5033, 5034, 5035

Notes: 30/60 GB QIC Tape Drive

This chart summarizes the feature numbers and media compatibility of the new 30/60 GB taped drive.

The following table contains an enlarged table showing media compatibility with much older Quarter Inch Cartridge media.

30/60 GB QIC Tape: Media Compatibility

Media Type	Capacity	#4x82	#4x83	#4x86	#4x84	#4x87
SLR100-50GB	50GB	N	N	N	N	R/W
SLR100-5GB	5GB	N	N	N	#4x84	R/W
SLR60-30GB	30GB	N	N	N	R/W	R/W
MLR3-25GB	25GB	N	N	R/W	R/W	R/W
MLR1-16GB	16GB	N	R/W	R/W	R/W	R
MLR1-2GB	2GB	N	R/W	R/W	R/W	R
SLR5-4GB	4GB	R/W	R	R	R	R
DC9250 *	2.5GB	R/W	R	R	R *	N
DC9120	1.2GB	R/W	N	N	N	N
DC6525	525MB	R/W	N	N	N	N
DC6150	120MB	R/W	N	N	N	N

* DC9250 format QIC2DC NOT allowed , format QIC2GB read only with #4x84

Backup & Recovery Performance

Save-While-Active performance improvements

- Reduced checkpoint processing time
  - Test case example: An SAP library showed that checkpoint processing that previously took 19:36 minutes now took only 1:36 minutes

Enhanced save and restore times to save files

- Certain internal benchmark save/restore times have improved with V5R2. See section 15.17 of the Performance Capability Reference Manual for details

BRMS

BRMS

GUI Enhancements

Improved disaster recovery

Online incremental Lotus server backup and recovery

Additional Tivoli Storage Manager (TSM) client functions

GUI Enhancements

- Save System
  - New delay function for restricted save in System Policy
  - Function of iSeries Navigator only
- Native save/restore
- Movement
  - Move Policy Wizard
  - Perform Movement
  - Verify Movement
- Lotus Server Point in Time Recovery
  - Available through iSeries Navigator only
- Backup Policy
  - Tape Library Support
  - Subsystem and Job Queues added
  - Parallel, TSM servers, Savefiles, Save-While-Active (SWA), and Independent ASPs added
- List Management
  - Create and update of your backup lists

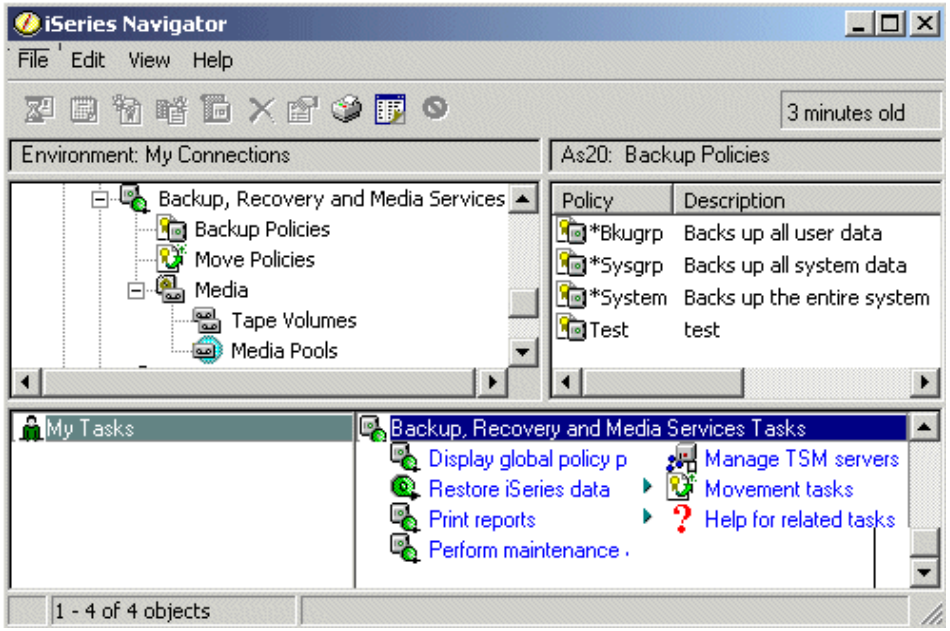

Notes: GUI Enhancements

Operation Navigator Plugin Enhancements and Wizard enhancements for V5R2 mentioned on this foil are self explanatory.

Improved Disaster Recovery

System Recovery Report (QP1ARCY) Enhancements

- Setting up standalone or media library devices for recovery
- Initialize BRMS functional authority information
- Update system name in BRMS media information
- Configure auxiliary storage pool devices (IASP support)
- Verify auxiliary storage pool device names
- Recovering Linux servers in a logical partition

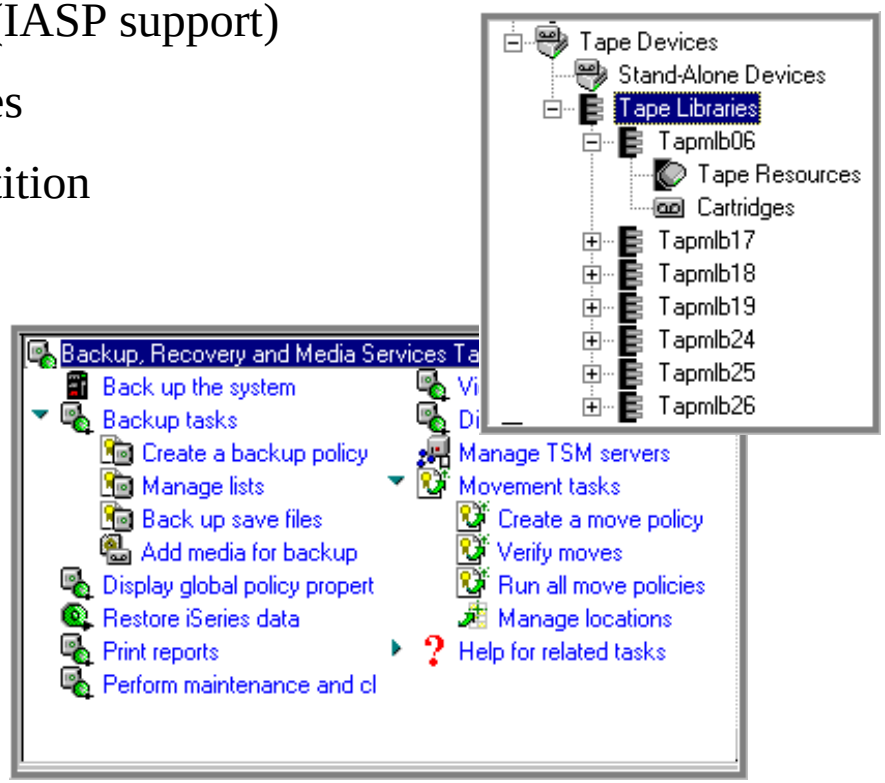

Notes: Improved Disaster Recovery

The BRMS System Recovery Report (QP1ARCY) has been enhanced with to following new steps.

- Setting up standalone or media library devices for recovery
- Initialize BRMS functional authority information
- Update system name in BRMS media information
- Configure auxiliary storage pool devices
- Verify auxiliary storage pool device names
- Recovering Linux servers in a logical partition

More BRMS Enhancements...

Enable shutdown of hosted logical partitions (Linux)

Native TCP/IP support for BRMS network

Performance enhancements for BRMS maintenance

Lotus Server Point in Time Recovery

- Online incremental Lotus server backup and recovery
- http://www-1.ibm.com/servers/eserver/iseries/service/brms/domIncremental.htm
- V4R5 PTF SF66782
- V5R1 PTF SI02780

Additional Tivoli Storage Manager (TSM) client functions

- Stronger Password Encryption
- Support the PASSWORDACCESS GENERATE client option
  - when used with the Tivoli Storage Manager OS/400 Client API, Version 4 Release 2 Level 1

Notes: More BRMS Enhancements...

Enable shutdown of hosted logical partitions (Linux)

Support was added in V5R2 so BRMS can control hosted logical partitions. This results in a higher level of backup automation.

There is no BRMS automated "power on" for hosted logical partitions in this release. Typically customers would manual start up their Linux partitions again or use startup programs.

Native TCP/IP support for BRMS network

BRMS added supported for TCP/IP in the BRMS network. So when we go look at another system, TCP/IP is used, if set up.

Performance enhancements for BRMS maintenance

In DUPMEDBRM and STRMNTBRM performance enhancements have been incorporated.

Lotus Server Point in Time Recovery

- Online incremental Lotus server backup and recovery
- http://www-1.ibm.com/servers/eserver/iseries/service/brms/domIncremental.htm
- V4R5 PTF SF66782
- V5R1 PTF SI02780

Additional Tivoli Storage Manager (TSM) client functions

- Media policies used for TSM operations have been updated to allow for a 64-character password, the maximum allowed by TSM servers. In addition, these passwords are stored by BRMS using a stronger encryption algorithm than was provided in previous releases.
- The BRMS TSM client has been enhanced to support the PASSWORDACCESS GENERATE client option when used with the Tivoli Storage Manager OS/400 Client API, Version 4 Release 2 Level 1 .
- See the following foils for more information on Tivoli Storage Manager V5.1 support on iSeries

Tivoli Storage Manager V5.1

Tivoli (TM) Storage Manager

Version 5.1 on iSeries generally corresponds to V 5.1 on other operating systems

- Centralized backup processes for iSeries and other platforms
  - Servers: OS/400, Windows NT, Sun Solaris, HP-UX, IBM AIX, MVS
  - Clients: OS/400 using BRMS, Window 95/98/NT/2000, AIX, Apple Macintosh. HP-UX, NewWare, OS/300 UNIX System Services, OS/2, SCO UnixWare, Sequent PTX, SGI IRIX, Sun Solaris, Tru64, Windows 32-bit DEC Alpha
- Significant improvement in function and performance over previously available V3.1 on iSeries

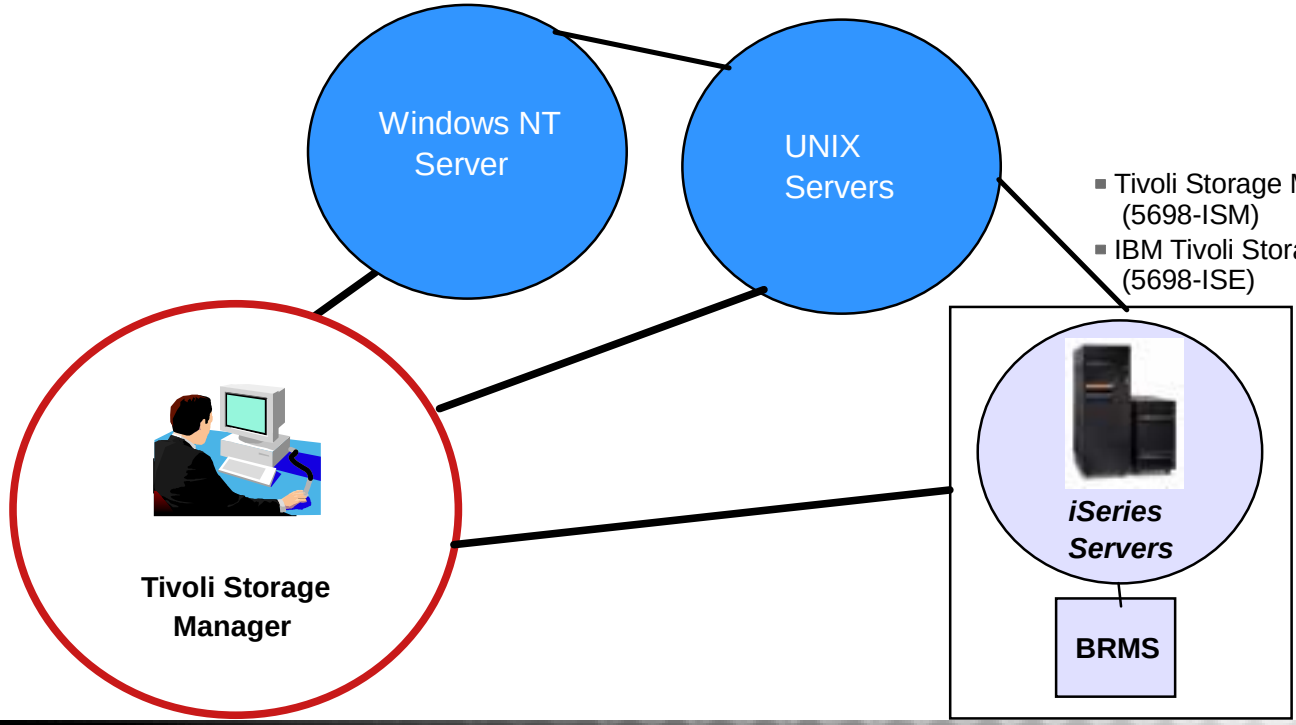

iSeries support available October, 2002

Notes: Tivoli (TM) Storage Manager

This foil places the new Tivoli Storage Manager V5.1 for iSeries along with TSM V5.1 on other platforms.

TSM is used by 80 Fortune 500 companies to protect approximately 1 million systems worldwide on the platform operating systems shown in this foil.

Tivoli Storage Manager has two unique features - its incremental backups that only backup the files that have changed since the last backup, and its ability to effectively utilize the disk and tapes in its storage hierarchy. The combination of these two mean less data being sent over your network, less tape drives needed to write the data and fewer tapes needed to store that data.

Tivoli Storage Manager does more than just backup your clients data - it has numerous products that integrate directly with the centralized server. These products allow you to backup your application databases, help you generate a plan for recovering from a disaster, move seldom used files off of clients local disk to near line storage, and generate in-depth reports to help you analyze and forecast your backup procedures.

Notes: Tivoli (TM) Storage Manager -2

On iSeries TSM can be optionally combined with Backup and Recovery Media Services (5722BR1) to provide a single centralized backup and recovery solution to your entire storage server environment. TSM helps reduces management costs by enabling users to perform backups and recoveries from a graphical interface on Web clients and set up automated backups according to the time considerations for each network environment.

There are two separately packaged Tivoli Storage Manager products:

- IBM Tivoli Storage Manager Version 5.1 (5698-ISM) — for basic backup-archive over a LAN
- IBM Tivoli Storage Manager Enterprise Edition V5.1 (5698-ISE) — to exploit advanced functions including LAN-free backup-restore to tape and disk, library sharing, space management, disaster recovery management, and NDMP support . IBM Tivoli Storage Manager Enterprise Edition is required if you want to use a library that has more than two drives or more than 40 tape slots. The functions that were in Tivoli Disaster Recovery Manager, Tivoli Space Manager, and Tivoli Data Protection (TDP) for NDMP, are included in IBM Tivoli Storage Manager Enterprise Edition.
  More about these "sub functions" functions later in this presentation.

Tivoli Storage Manager V5.1 for iSeries provides a significant set of enhanced functions and improved performance over the previous version available on the iSeries - at the V3.1 function level. TSM has made many changes since then and most, except for SAN-based support functions, apply to the new iSeries TSM V5.1.

Support for iSeries V5R1 and V5R2 planned availability is October, 2002 .TSM 5.1 support on other operating systems ("platforms") became available during April 2002. iSeries support will be integrated into the already existing TSM V5.1 support functions as follows:

- Customers ordering 5698-ISM after the October 5.1.5 GA will get the iSeries package with all the other servers; existing customers will get a refreshed autoship of the package which will include the iSeries server.
- Customers will get the new iSeries entitled publications (Quick Start) as part of the new or autoshipped package described above. After GA a user can download nonEntitled documentation from the Tivoli Information Center or the IBM publication center or purchase from the IBM Publications Center.

The next few foils describe Tivoli Storage Manager (TSM) V5.1 capabilities on iSeries. There are some additional TSM V5.1 functions that are uniquely available in each platform (operating system) that would not be on iSeries.

Tivoli Storage Manager V5.1 for iSeries: General Information

Significant enhancements over previous (V3.1) level TSM support on iSeries

TSM V5.1 runs in OS/400 PASE

- OS/400 V5R1, V5R2 under PASE
  - Starting with V5R2 OS/400 PASE (OS/400 option 33) is now a no-charge, no order required option
- Install Tivoli Storage Manager V5.1 for iSeries, 5698-ISM or Tivoli Storage Manager V5.1 for iSeries Enterprise Edition, 5698-ISE, with RSTLICPGM

Improved integration with Backup and Recovery Media Services (BRMS), 5722-BR1

TSM V5.1 for iSeries Improved Availability

- Application availability is improved during backups by providing on-line, nondisruptive image backups
- High availability support is provided with HACMP failover for the AIX backup-archive and HSM clients

Notes: TSM V5.1 for iSeries: General Information

The current (prior to October 2000 availability of TSM V5.1 for iSeries) TSM support on iSeries is at the older iSeries TSM level V3.1 function level.

The IBM Tivoli Storage Manager V5.1 products provide a server that runs on OS/400 PASE (Portable Application Solutions Environment). OS/400 PASE is an integrated runtime environment for AIX (or other UNIX®-like) applications running on the IBM iSeries system. With V5R1 OS/400 you must explicitly order PASE, option 33. Starting with OS/400 V5R2, option 33 comes with OS/400. In both cases Option 33 must be installed.

The IBM Tivoli Storage Manager V5.1 OS/400 PASE server has a UNIX look and feel, but contains the necessary modifications to permit it to function in the PASE environment. Tivoli Storage Manager V5.1 OS/400 PASE server, which generally supports the non-SAN environment enhancements that are part of V5.1, as discussed in this set of foils.

BRMS for iSeries, 5722BR1 contains many save and restore functions on iSeries, which can optionally be used with TSM V5.1 as some TSM V5.1 functions provided on other platforms (operating systems) can be performed by BRMS on iSeries.

TSM V5.1 Availability enhancements for iSeries:

- Online Image Backups: Application availability is improved during backups by providing online, nondisruptive image backups. Applications remain available while online image backups create a point-in-time image backup of the file system. Image backups complement progressive incremental backups and can be used for faster recovery of large amounts of data in file systems with large numbers of files. Online image backups are available for Windows 2000 for FAT, FAT32, NTFS, and RAW volumes and for Linux x86 for Ext2, ReisterFS, and RAW LVM volumes.
- HACMP (High Availability Cluster Multiprocessing) support: Support for high availability is provided with HACMP failover for the AIX backup-archive, HSM (Hierarchical Storage Management) and API clients. HACMP failover automatically resumes normal operation on another node in the cluster after failover. Scheduled processes which have not completed are automatically restarted within the schedule start window. Otherwise, normal operation resumes with the next scheduled event. Support for automatic fallback when failed node rejoins the cluster is also provided as an option.

Limitations

SAN attached tape and disk (ESS) devices are supported. However, the Tivoli Storage Manager running under OS/400 PASE does not support specific Storage Area Network (SAN) capabilities, including support for LAN-free clients, tape library sharing, and server-free backup.
Also SNMP support and optical device support are not available with TSM V5.1 running under OS/400 PASE.

TSM V5.1 for iSeries: Improved Recovery

TSM V5.1 for iSeries Improved Recovery and Disaster Recovery

- Performance:
  - Recovery performance on Windows 2000 and Linux enhanced by utilizing fast image restores
  - Recovery performance further enhanced with parallel restores provided by multi-session restore
  - Time to create recovery tapes for local and off-site vaulting is reduced with simultaneous writes to multiple copy storage pools during backup and archive with LAN clients
  - Self-tuning
  - Restore time can be improved by staging data to disk or by consolidating data on tape
  - Client backup set restore times are also improved
- Adaptive Differencing minimizes the amount of data transmitted over a network
- Improved backup performance available with journal-based backup supported on cluster configurations
- Disaster Recovery Manager functions included in Enterprise Edition

Notes: TSM V5.1 for iSeries Improved Recovery

This foil lists some of the key recovery and disaster recovery important performance-related enhancements:

- Fast Image restores: Recovery performance on Windows 2000 and Linux x86 is enhanced by utilizing fast image restores. Image backups complement progressive incremental backups to provide full file system restores for disaster recovery or when a large percentage of a file system is to be restored.
- Multi-session Restores: Restore time is reduced when multi-session restore sessions are used with the backup-archive clients. This enable restores from tape to run in parallel when data is stored on multiple tapes and the tape drives needed are available at restore time. The maximum number of tape drives that can be used concurrently can be configured by the administrator.
- Simultaneous Writes To a Copy Storage Pool: The time to create recovery media for local and off-site vaulting is reduced with simultaneous writes to multiple copy storage pools. Duplicate copies of data for disaster recovery can now be created concurrently when backing up or archiving data with the LAN clients, provided that there are sufficient devices for each of the primary and copy storage pools. This can reduce the need for the sequential procedure of first completing the backup or archive and then creating the copies for vaulting.
- Self Tuning: Tivoli Storage Manager uses an intelligent, adaptive algorithm to optimize performance for individual customer environments, while minimizing administrative intervention.
  - The batch size and the threshold for server migration and storage pool backup will be adjusted to obtain better performance without running out of log space
  - Server expiration processing will reset the database buffer pool, based on the cache hit ratio
- Backup-Archive Client Multiple Sessions: Exploitation of backup-archive client multi-threading allows automatic multiple sessions for backup and archive tasks. Both the GUI and the command line are enhanced to display data from multiple backup streams.
- Backup Sets: Backup sets can be used to create media for rapid recovery or instant archive. Backup sets can be used to restore data to backup-archive clients locally without a network or Tivoli Storage Manager server.
- Instant Archive: Instant archive enables client archive sets to be created from backup versions of files already stored on the Tivoli Storage Manager server. In a networked storage management environment, this enables LAN-free (non-networked) records retention (archive) capabilities. The archive sets are created when the BACKUPSET command is issued to any supported sequential media, such as tape volumes. At creation, the sets can optionally be given a description and retention period to simplify tracking and expiration. The sets are tracked in the volume history file.

Notes: TSM V5.1 for iSeries Improved Recovery -2

Performance-related enhancements continued:

- Fault-Tolerance: Enhanced fault-tolerant server support permits the skipping over of read-errors on tape if no readable copy of the file can be found. Before skipping over the error, the server will attempt to restore the file from a copy stored in an online storage pool, if any. The error will be reported to the server and the client. It is also possible that a good copy may reside in an off-site copy storage pool volume. If so, an informational message will inform the administrator that a volume can be brought on-site to help with the restore.
- Full File System and Raw Logical Volume Backups: Full file system or raw logical volume images can be backed up and are managed as a single object by Tivoli Storage Manager policy like any other object on the server. **The function is implemented on the supported versions of the AIX, HP-UX, and Solaris backup-archive client platforms**. The Tivoli Storage Manager backup-archive client has a new option in the command line and GUI programs to create a file system image to backup which the client will send to the server as a single object. When using image backup at the same time with progressive incremental, the user will be able to restore images and incremental backups based on the image. A new option allows the user to perform an incremental by image date backup. Restores of single files from within an image will not be possible.
- Out-of-Band Database Backups to Media: You need to prepare for disaster recovery by taking database backup media off-site. Out-of-band (snap shot) backups allow you to create Tivoli Storage Manager server database backups for movement off site, while maintaining your full and incremental database backup series on site, for availability purposes. Off-site management of the Tivoli Storage Manager server database snap shot backups can be performed by the Tivoli Disaster Recovery Manager.

Notes: TSM V5.1 for iSeries Improved Recovery -3

Tivoli Storage Manager for iSeries V5.1 functions continued - client-based functions:

- Adaptive Differencing: Adaptive Differencing technology fundamentally changes the way data may be transferred throughout the enterprise. Available for the AIX, Sun Solaris, HP-UX, Windows NT, Windows 2000, OS/400 PASE, and MVS™ **server platforms**, Adaptive Differencing technology transfers data by byte, block, or file level based on data size. This technology supports a variety of connectivity strategies, including LANs, WANs, Internet, and dialup connections. This application is designed for mobile computer users and other users with a need to minimize the amount of data transmitted over a network.

  Mobile user data protection is enhanced with support for adaptive subfile level backup and encryption of data before it is transmitted over the network. With Tivoli Storage Manager Progressive Backup Methodology, only files that have changed are candidates to be backed up, therefore, eliminating unnecessary data transfers that rob your network and CPUs of vital power and productivity. With Adaptive Differencing, the backup-archive client dynamically determines the most efficient approach for creating backup copies changed bytes, changed blocks or changed files, delivering improved backup performance over dialup connections. With encryption, the backups being sent over the public phone lines are more secure, as are the files being stored on the Tivoli Storage Manager server.

  Adaptive Differencing and encryption technology are supported on Tivoli Storage Manager backup-archive clients for Windows NT, Windows 2000, Windows ME, and Windows XP.

  A file backed up using Adaptive Differencing initially sends a reference file to the Tivoli Storage Manager. Subsequent backups for that file send a delta file, which consists of the changed data from the reference copy. A restore operation reconstructs the file on the Tivoli Storage Manager client using the reference file and the delta file that represents the point in time for the restore request.
- Windows 2000 Support: The Tivoli Storage Manager client for Windows 2000 offers comprehensive protection of the Windows 2000 System State, including Active Directory, COM+ Class Registration Database, System Volume, and System Files. Other new Windows 2000 changes include disk quotas, encrypted files, distributed file systems, certificate server databases, MS cluster databases, reparse points utilized by volume mount points, directory junctions and removable storage manager databases, which are also fully protected through Tivoli Storage Manager backups and archives.
- Journal-Based Backups: Tivoli Storage Manager client journaling improves overall incremental backup performance for Windows NT and Windows 2000 clients. Performance is improved because the files to be backed up are tracked in a journal. Journal-based backups eliminate the need for the client to scan the local file system or query the server to determine which files to process. It also reduces network traffic between the client and server.

Notes: TSM V5.1 for iSeries Improved Recovery -4

This The Tivoli Disaster Recovery Manager functions (included in the Enterprise Edition) helps you maintain business continuance by:

- Establishing and helping to automate a thorough disaster recovery plan
- Automating vital recovery steps to bring your business back to normal
- Managing and identifying off-site media needed for recovery
- Tracking and reporting systems destroyed, in event of disaster
- Performing restores in order of priority

The Tivoli Space Manager function uses hierarchical storage management (HSM) to automatically and transparently migrate rarely accessed files to Tivoli Storage Manager storage while the files most frequently used remain in the local file systems. By migrating rarely accessed files to the server storage, Tivoli Space Manager frees administrators and users from manual file system pruning tasks by enabling you to have sufficient free storage at your workstation or file server, deferring the need to purchase additional disk storage. **The Tivoli Space Manager function is supported on the AIX and Sun Solaris platforms.** For the iSeries environment you can use TSM along with Backup and Recovery Media Services (BRMS) product, 5722-BR1, to accomplish corresponding HSM capabilities.

The Tivoli Data Protection for NDMP function provides backup and recovery support on Tivoli Storage Manager servers for network-attached storage (NAS) file servers from Network Appliance. NAS file servers often require a unique approach to providing backup and recovery services, because these file servers typically will not run third-party software. Tivoli Data Protection for NDMP utilizes the Network Data Management Protocol (NDMP) to communicate with and provide backup and recovery services for NAS file servers. NDMP is an industry-standard protocol that allows a network storage-management application to control the backup and recovery of an NDMP-compliant file server without installing third-party software on that server. The implementation of the NDMP server protocol enables the NAS file servers to be backup-ready and enables higher-performance backup to tape devices without moving the data over the LAN.

TSM for iSeries V5.1: Details

Browser-based interface now supports JRE 1.3.1 Swing-enabled browser

OS/400 IFS file system QOpenSys contains TSM disk volumes

Standard OS/400 tape and disk devices supported, including:

- ESS LUNs
- 3590, 3570, LTO (358x)
- 8mm (not on i890), QIC

All library and tape definitions are performed with DEFINE LIBRARY, DEFINE DRIVE commands

Notes: TSM for iSeries V5.1: Details

The Web client support is enhanced to support Java Runtime Environment (JRE) 1.3.1 Swing-enabled browsers. The first time a Web browser accesses a TSM V5.1 server you are prompted to download and install a Java (TM) runtime plug-in.

The IFS QOpenSys file system is similar to the AIX file system hierarchy, which "fits with the TSM V5.1 implementation on iSeries under PASE. All disk volumes defined with the *define volume*, *define dbvol*, *define logvol* commands are placed into /QOpenSys/usr/tivoli/tsm/server/bin directory.

PASE tape support drivers were added to support all tape devices supported by OS/400. The ESS (Shark) device definitions are transparent to TSM.

All library and tape definitions are supported as they were on the previous TSM versions running on iSeries - via the DEFINE DRIVE and DEFINE LIBRARY commands.

TSM for iSeries V5.1: More details

TSM V5.1 and BRMS

- TSM V5.1 only or
- TSM V5.1 and BRMS on same server

TSM V3.7 clients and later are supported

TSM export function used to migrate from older versions of TSM to V5.1

Performance updates planned for August 2002 availability:

- Server side: http://www.tivoli.com
- Client side: http://www.ibm.com/servers/eserver/iseries/service/brms

Backing up Domino files options:

- TSM V5.1 TDP, or ...
- BRMS base for single system, additional cost for multiple systems

Notes: TSM for iSeries V5.1: More details

TSM V5.1 server functions can run without BRMS.

The BRMS Application Client to TSM allows BRMS users to perform save and restore operations using Tivoli Storage Manager (TSM) servers as the storage medium.

If BRMS and TSM V5.1 server are installed on the iSeries, TSM can share the tape library and media using the same techniques used with the older TSM V3.1 - BRMS interfaces- the same exits for mount, dismount, or delete . TSM V5.1 still uses the BRMS media classes to separate tapes from BRMS backups.

During OS/400 V4R5 BRMS was enabled to back up Domino files. BRMS provides true online backup support for Lotus Servers such as Domino and QuickPlace. These servers can be saved while they are in use with no save while active synchronization points. Incremental backup and point in time restore of Lotus Server databases are supported.

BRMS is available with some packaging/priced options. The base license manages a single system. If you are already using BRMS to backup your Domino files you can continue to do so. If you do not have BRMS installed, you can use the Tivoli Data Protection (TDP) support to backup Domino files.

TSM for iSeries V5.1: Restrictions

TSM V5.1 functions not supported on iSeries:

- LAN Free (Storage Agent)
- Server Free
- Tape Library Sharing
- Raw logical volumes
- Optical devices
- NAS (NDMP)
- SNMP Subagent

OS/400 Save System (SAVSYS), Save Security Data (SAVSECDTA), Save Configuration (SAVCFG) command functions are not supported

iSeries TSM V5.1: Restrictions

As supplemental information, the following describes some of the TSM V5.1 functions not supported on iSeries servers.

LAN Free

Introduced with TSM V 3.7, enough of Tivoli Storage Manager server code (called Storage Agent) is installed on the TSM client which has a SAN attachment. When LAN Free support is invoked it can send managed TSM client data over the SAN directly to the tape device - typically a TSM Server managed tape library. The LAN is used to exchange control information between TSM server and TSM client, but only the SAN network is used to transfer data from the TSM client device directly to the SAN-attached tape device.

Server Free

The impact on application servers during backups is minimized with server-free data movement. Server-free data movement reduces CPU utilization by removing backup and restore processing from your production machines. Similar to LAN-free (but without the client workstation Storage Agent), server-free also lowers traffic on your LAN by off loading your backups and restores to the SAN. Scalability on the IBM Tivoli Storage Manager server is improved and it can handle more concurrent client connections and server operations because the server will not be copying data.

NAS (NDMP) support

The Tivoli Data Protection for NDMP function provides backup and recovery support on Tivoli Storage Manager servers for network attached storage (NAS) file servers from Network Appliance. NAS file servers often require a unique approach to providing backup and recovery services, because these file servers typically will not run third-party software. Tivoli Data Protection for NDMP utilizes the Network Data Management Protocol (NDMP) to communicate with and provide backup and recovery services for NAS file servers.

NDMP is an industry-standard protocol that allows a network storage-management application to control the backup and recovery of an NDMP-compliant file server without installing third-party software on that server.

Raw logical volumes

A raw logical volume means "all the data on a disk" (regardless of file system)

Notes: iSeries TSM V5.1: Restrictions -2

iSeries OS/400 Save Security Data (SAVSECDTA), Save Configuration (SAVCFG), Save System (SAVSYS) command functions not supported

The Save System function is not supported. SAVSYS requires the system to be in restricted state and TSM cannot run in this restricted state. OS/400 supports running SAVSECDTA, and SAVCFG in non-restricted state. However, TSM V5.1 and BRMS, when working together, do not support saving and then restoring these sets of security and hardware configuration information.

The save function for SAVSECDTA and SAVCFG could be supported, but in a "disaster recovery" mode, this saved information could not be restored. This is because the iSeries restricted state required for the disaster recovery scenario does not support remote LAN communications or an additional OS/400 subsystem to be active, which is required to run TSM.

Planned Tivoli Storage Manager for iSeries support

Currency with iSeries TSM on all supported platforms:

- iSeries PASE TSM will be kept current along with other TSM servers

Exceptions:

- No Storage Area Network (SAN)-based support planned (LAN Free, ...)
- Domino support changes
  - TSM Domino support will not depend on nor require BRMS
  - June 2002: BRMS-based solution to backup Domino (data does not go to TSM)
  - Future: TSM for Mail/Domino solution will backup to TSM

IBM plans and directions are subject to change without notice

Notes: Planned Tivoli Storage Manager iSeries support

TSM V5.1 planned enhancements are grouped into the topics listed on this foil.

Limitations

As previously stated The Tivoli Storage Manager V5.1 OS/400 PASE server does not support and has no plans to add the SAN-based functions, including support for LAN-free clients, tape library sharing server-free backup, and Server free, There are no plans to add SNMP support and optical device support.

Additional information:

See the following for additional information:

- The Tivoli Storage Manager V5.1 announcement letters, at http://www.ibm.com/support -> Announcement Letters for:
  - US: 202-078, April 09, 2002
  - EMEA: ZP02-0172, April 09, 2002
  - AP: AP01-1106, April 09, 2002
- http://www.ibm.com/tivoli
- http://www.tivoli.com/products/solutions/storage/news.html
- Redbook Tivoli Storage Manager Version 5.1 Technical Guide, SG24-6554, available at http://www.ibm.com/redbooks

Trademarks and Disclaimers

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.
The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

400	BRMS	Host Integration Series	JustMail	Payment Manager	Stylized e
ADSTAR	Client Series	Host on Demand	MQSeries	Payment Server	SystemView
Advanced Function Printing	ClusterProven	Host Publisher	MQSeries Integrator	PCOM	VisualAge for Java
AFP	CODE/400	HTTP Server for AS/400	Net.Commerce	PowerPC	VisualAge for RPG
AIX	DataGuide	IBM	Net.Data	PowerPC AS	WebSphere
AnyNet	DB2	IBM Logo	Netfinity	Print Service Facility	WebSphere Advanced Edition
Application Development	DB2 Extenders	IBM Network Station	NetView	pSeries	WebSphere Commerce Suite
APPN	DB2 UDB for AS/400	Information Warehouse	NUMA-Q	PSF	WebSphere Development Tools for AS/400
AS/400	DB2 Universal	Integrated Language Environment	OfficeVision	S/390	WebSphere Standard Edition
AS/400e	e-business logo	Intelligent Printer Data Stream	OS/2	SanFrancisco	Workpad
AT	e(logo) Server	IPDS	Operating System/400	Screen Publisher	xSeries
BrioQuery	Enterprise Storage Server	iSeries	OS/400	SmoothStart	

cc:Mail, Domino.Doc, Freelance, LearningSpace, Lotus, Lotus Domino, Lotus Notes, iNotes, QuickPlace, Sametime, and Word Pro are trademarks of Lotus Development Corporation in the United States, other countries, or both.
Tivoli and NetView are trademarks of Tivoli Systems Inc. in the United States, other countries, or both.
C-bus is a trademark of Corollary, Inc. in the United States, other countries, or both.
Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
PC Direct is a trademark of Ziff Communications Company in the United States, other countries, or both and is used by IBM Corporation under license.
ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States, other countries, or both.
IBM's VisualAge products and services are not associated with or sponsored by Visual Edge Software, Ltd.
Linux is a registered trademark of Linus Torvalds.
UNIX is a registered trademark of The Open Group in the United States and other countries.
SET and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.
Other company, product and service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information in this presentation concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of the specific Statement of Direction.

Some information in this presentation addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.