

iSeries and Availability

IBM @server. For the next generation of e-business.

Agenda

Review save restore and clustering highlights through V4R5

V5R1 Journaling

V5R1 Backup Recovery and Media Services for iSeries, 5722-BR1

New Hardware Options

V5R1 Clustering

V5R1 Independent Auxiliary Storage Pool and Switched Disks

Availability Software

Additional Availability Solutions

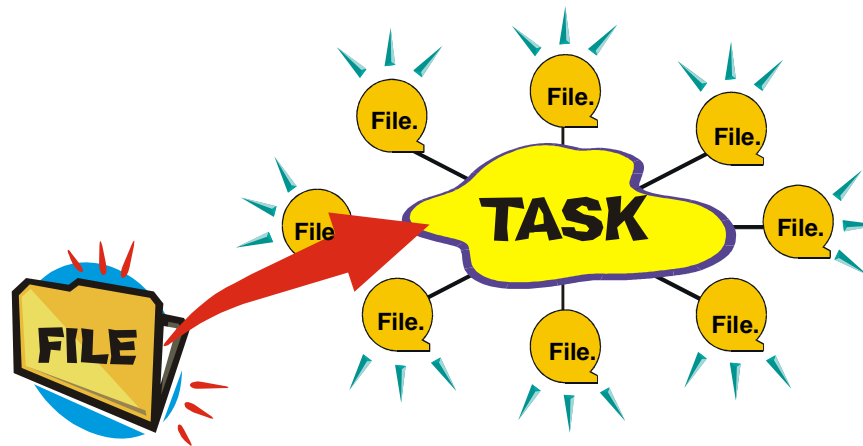
Summary

Save/Restore Performance V4R5

IBM @server. For the next generation of e-business.

Recent releases introduced

- Parallel save (Same library/object saved across multiple tape drives):




- V4R5 introduced iSeries High Speed Link connections

Parallel Save/Restore

- Up to 2.2 TB/Hr
- Controlled by User or BRMS

Concurrent Save/Restore

- Up to 2.6 TB/Hr



Notes: Save/Restore Performance V4R5

This foil is a quick review of key performance increments in saving and restoring major objects on the AS/400 and iSeries systems.

Concurrent of different objects and, in V4R4, parallel save/restore support offered performance increments over previous release capabilities.

The parallel save/restore support delivered fastest throughput, but required use of a Media Definition Object (*MEDDFN) that defined the tape drives to be used concurrently and portions of the save object would be contained on each tape. This typically required Backup and Recovery Media Services (BRMS) to manage this object and associated save or restore functions.

Even with the parallel support maximum throughput was limited by the AS/400 V4R4 SPD bus architecture

With the introduction of the High Speed Link (HSL) bus architecture on the V4R5 iSeries models, the parallel save achieved significant performance improvement with test showing up to 2.2 TB/Hr. The Concurrent save improved to up to 2.6 TB/Hr.

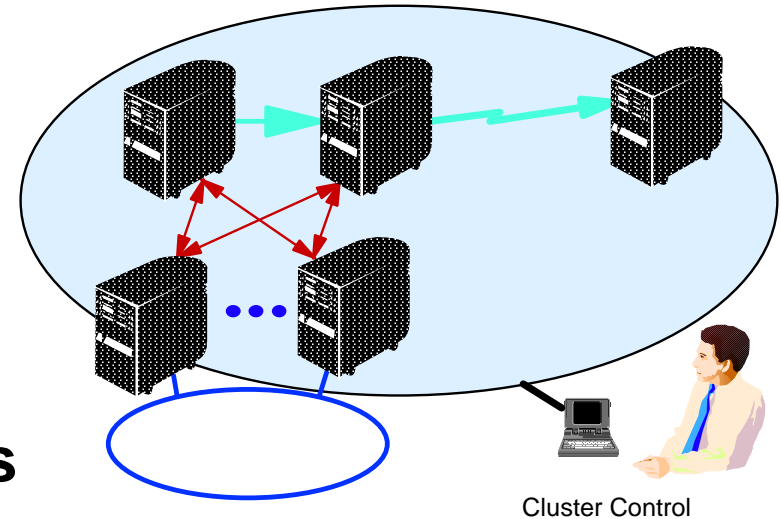
The V5R1 Performance presentation has performance test results updates for new V5R1 tape device Fibre Channel attachments using 3590 E11/E1A models and 3584 LTO models.

V4R5 Clustering Highlights

IBM @server. For the next generation of e-business.

Cluster Resources

- Single resource view for multiple nodes
- IP address takeover
- Data Resiliency and Application Resiliency resource groups



Integrated Cluster Resource Services

- Manage cluster resource groups
- Integrated services for heartbeating, reliable message delivery, switchover administration and distributed activities

Cluster Management

- 2 to 128 nodes support
- Easy to configure/manage
- High availability solutions from Business Partners

IBM  server. For the next generation of e-business.

Notes: V4R5 Clustering Overview

In V4R4 the AS/400 is enhanced with Continuous Availability Clustering. The definition of a cluster is "a group of independent systems working together as a single system."

AS/400 or iSeries clustering lets you efficiently group your systems together to set up an environment that provides availability that approaches 100% for your critical applications and your critical data. Clustering also provides simplified systems management and increased scalability to seamlessly add new components as a customer's business grows.

Cluster resources include multiple nodes, heartbeat services, configurations, IP address takeover and so on. Under Cluster Management, our continuous availability and data resilience software and solutions come from IBM's AS/400, iSeries High Availability Business Partners (HABP).

There are four aspects to continuous availability, four components to providing data resiliency and application resiliency. The foundation is Cluster Resource Services, which exists within OS/400. On top of that is cluster management provided by business partners which code to an open set of APIs. These products will control the operations within the cluster, providing a GUI interface to the customer.

Data resiliency is keeping track of copies of the data. The HABPs provide this function, based on remote journaling functions for database and IFS stream file objects. Non-database objects which are not journaled require a variety of techniques and business partners provide these solutions.

Application resiliency involves having the application written to "restart" when a "fail over" or "switch over" occurs. Application developers can write highly available applications according to a defined architected interface provided by OS/400.


AS/400 systems that makes up the cluster: Nodes

Special subset of Nodes: Recovery Domain

Resilient Objects and Resilient Applications: Cluster Resources

Fail over and Switch over: Cluster Policies

OS/400 system object: Cluster Resource Group



A **cluster node** is any AS/400 system, iSeries that is a member of a cluster. You can use any name that you want. However, it might be simplest if the node name is the same name used for either the host name or the system name. This cluster node name is then mapped to an 8-character cluster node identifier (CNI) that is associated with one or more Internet Protocol (IP) addresses that represent an AS/400, iSeries system. Cluster communications that run over IP connections provide the communications path between cluster services on each node in the cluster. A cluster node can also be referred to as a **point of access**. The set of cluster nodes that are configured as part of the cluster are referred to as the cluster membership list. A cluster node or point of access can have different roles.

Cluster Resources are sets of resilient data and resilient applications. Resilient data is data that is identified as a resource to the cluster and access to that data can be managed through the cluster resource services. The technique for maintaining that resiliency is, to a great extent, transparent to the application implementing and demanding access to that data. That data could be copied or "held in a safe place, "independent of the application's knowledge." Resilient data can be used by ClusterProven applications as a means to allow them to have the application's "data store" available to the cluster nodes where the application itself is "resilient".

In the "switched IASP" case, resilient data does not use copies or replication to provide access to itself from cluster nodes but rather provides "resiliency" to cluster nodes by using a clustered managed object (device CRG) that allows nodes in the cluster switched access to the data.

A resilient application is an application that can be restarted on a different cluster node without requiring you to reconfigure the requesters or clients. A resilient application needs the ability to recognize the temporary loss of the Internet Protocol (IP) connection between the requesters or client and the server. The requester or client application must be aware that the IP connection will be temporarily unavailable and must retry access rather than ending or initiating a fail over. Similarly, if you are performing a switch over, server applications need to be aware that the IP connection is no longer available. Eventually, an error condition is returned to the server application. Once this error condition is received, it is best if the server application recognizes the condition and ends normally. IP address takeover is a high availability function that is used to protect requesters or clients from application server outages. The concept is to use IP address aliasing to define a floating IP address that is associated with multiple application servers or hosts. When one application server in a cluster fails, another cluster node assumes the responsibilities of the application server without requiring you to reconfigure the requesters or clients.

A **cluster resource group** is an OS/400 system object that is a set or grouping of cluster resources. The group describes a recovery domain (see next page) and supplies the name of the cluster resource group exit program that manages cluster-related events for that group. One such event would be moving an access point from one node to another node.

Cluster resource group objects are either defined as data resilient (type-1) or application resilient (type-2). Data resiliency enables multiple copies of data that is maintained on more than one node in a cluster. Application resiliency enables an application (program) to be restarted on either the same node or a different node in the cluster. Every cluster resource group has a cluster resource group exit program associated with it.

A **primary node** is the cluster node that is the point of access and principle copy of a resource. If this node fails, all cluster resource group objects having this node as the primary access point will fail over to a backup node.

A **backup node** is the cluster node that will take over the role of primary access if the present primary node fails. This cluster node contains a copy of a cluster resource. In the case of a data cluster resource group, copies of the data are kept current with replication.

A **replicate node** is a cluster node that has copies of cluster resources, but is unable to assume the role of primary or backup. Use the Change Cluster Resource Group API to change a replicate node to a backup node.

A **recovery domain** is a subset of nodes in the cluster that are grouped together in a cluster resource group for a common purpose such as performing a recovery action. A domain represents those nodes of the cluster from which cluster resource can be accessed. This subset of cluster nodes that is assigned to a particular cluster resource group either supports the primary point of access, secondary (backup) point of access, or replicate. Each node in the recovery domain has a role with respect to the current operational environment of the cluster. This is called its current role in the recovery domain. As the cluster goes through operational changes such as nodes ending, nodes starting, and nodes failing, the node's current role is changed accordingly. Each node in the recovery domain also has a role with respect to the preferred or ideal cluster environment. This is called its preferred role in the recovery domain. The preferred role is a static definition that is initially set when the cluster resource group is created. As the cluster environment changes, this role is not changed. However, you can change this role through one of the available API's (see cluster API reference).

Cluster Policies; Fail over and Switch over. A fail over means that the system automatically switches over to one or more backup systems in the event of a system failure. A switch over happens if you manually switch access from one system to another. You would usually do this if you wanted to perform system maintenance such as applying program temporary fixes (PTFs), installing a new release, or upgrading your system.

Switch over and fail over order is the relationship (or order) that you have defined among the primary node and backup nodes in a recovery domain. In a recovery domain, there can be multiple backup nodes. You specify one node as first backup, another as second backup, and so on. If a primary node fails, the access point for the resilient resources switches to the first backup node.

Join means to become a new member of some entity such as a cluster. Rejoin means to become an active member of a cluster after having been a nonparticipating member. For example, when clustering is restarted on a node after the node has been inactive, the cluster node rejoins the cluster.

Integrated Cluster Services

IBM  server iSeries

Message Function

Heartbeat Function

OS/400 cluster service jobs

IBM  server. For the next generation of e-business.

Cluster resource services consists of a set of multithreaded jobs. When clustering is active on an AS/400, the jobs run in the QSYSWRK subsystem. The jobs run using the QDFTJOB job description. Should any cluster resource services job fail, no job log will be produced. In order to provide a job log, change the LOG parameter of the job description to a level that produces job logs.

- Cluster control consists of one job that is named QCSTCTL.
- Cluster resource group manager consists of one job that is named QCSTCRGM.
- Cluster resource groups consist of one job per cluster resource group object. The job name is the same as the cluster resource group name.

Using the message function and heartbeat function to monitor the status of your cluster:

The reliable **message function** of cluster resource services keeps track of each node in a cluster and ensures that all nodes have consistent information about the state of cluster resources. Reliable messaging uses retry and time-out values that are unique to clustering. These values are preset and cannot be changed. These values are used to determine how many times a message will be sent to a node before a failure or partition situation is signaled. For a local area network (LAN), the amount of time it will take to go through the number of retries before a failure or partition condition is signaled is approximately 45 seconds. For a remote network, more time is allowed to determine whether a failure or partition condition exists. You can figure approximately four minutes and 15 seconds for a remote network.

Heartbeat monitoring ensures that each node is active. When the heartbeat for a node fails, the condition is reported so the cluster can automatically fail over resilient resources to a backup node. A heartbeat message is sent every 3 seconds from every node in the cluster to its upstream neighbor. In a network, the nodes expect acknowledgment to their heartbeat from the upstream node as well as incoming heartbeats from the downstream node, thus creating a heartbeat ring. By using routers and relay nodes, the nodes on different networks can monitor each other and signal any node failures.

V5R1 Topics

IBM  server iSeries

HSL enhancements

Clustering and IASPs

Journaling

BRMS

IBM  server. For the next generation of e-business.

V5R1 Availability, Clustering, Switchable Independent Auxiliary Storage Pools

IBM @server. For the next generation of e-business.

iSeries Clusters have been significantly enhanced with switch disk tower clustering support, enabled with Independent ASP and enhanced HSL Loop capabilities. Additional HSL OptiConnect support is now available with high speed system interconnect using the HSL Loop technology. For replication clusters there is now Journaling of IFS Objects, Data Areas and Data Queues. Finally, within Operations Navigator and Management Central there is a simple Cluster Management Utility that allows for simple two node switch disk clusters to be created without the need for the full support available with the cluster middleware products from DataMirror, Lakeview Technology and Vision Solutions.

Domino for iSeries, release 5.0.7 is a ClusterProven for iSeries application, using the switch disk cluster implementation.

High Speed Link OptiConnect

Server to Server connectivity over HSL

- 1 GBytes per second connectivity

All iSeries Models with V5R1 hardware are enabled with new adapters

- Year 2000 iSeries Models 830 and 840 require new orderable features

If implementing switchable IASP:

- HSL OptiConnect hardware required*
- OS/400 option 41 (HA Switchable Resources, extra cost) required
- OS/400 option 23 (OptiConnect intersystem communications) not required

If implementing OptiConnect intersystem communications:

- Can use: HSL OptiConnect hardware, 1 Gbits/second Ethernet, SPD Optical via Migration Tower
- OS/400 option 23 (OptiConnect intersystem communications) required
- OS/400 option 41 (HA Switchable Resources) not required

*Not required for IOPs in LPAR partitions

HSL Loop technology was introduced in V4R5 as the means for attaching I/O towers to the base system unit.

You can use this HSL fabric for high-speed server to server interconnect running more than ten times faster than existing SPD OptiConnect which **HSL OptiConnect** replaces. This new support does not need the new V5R1 HSL adapter hardware. The result is greatly expanded capability for high-availability options and distributed application scenarios. In the world of e-business, continuous availability and distributed workload are minimum requirements.

HSL OptiConnect Loop capability requires the new (updated) HSL hardware, available in the new iSeries models announced April 2001. This enables both HSL OptiConnect connectivity and switched (disk) I/O tower functions.

The V4R5 iSeries models 830 and 840 (available before April 2001) can order the same level HSL OptiConnect/switchable IASP adapter. The V4R5 iSeries 270s and 820s cannot have this newer technology adapter installed on their system. A customer with such a 270 or 820 must order or upgrade to a new April 2001 270 or 820 if they wish to use this new adapter.

See the Hardware overview presentation for details of the HSL features available in combination with V5R1 for the Models 830 and 840 announced in 2000.

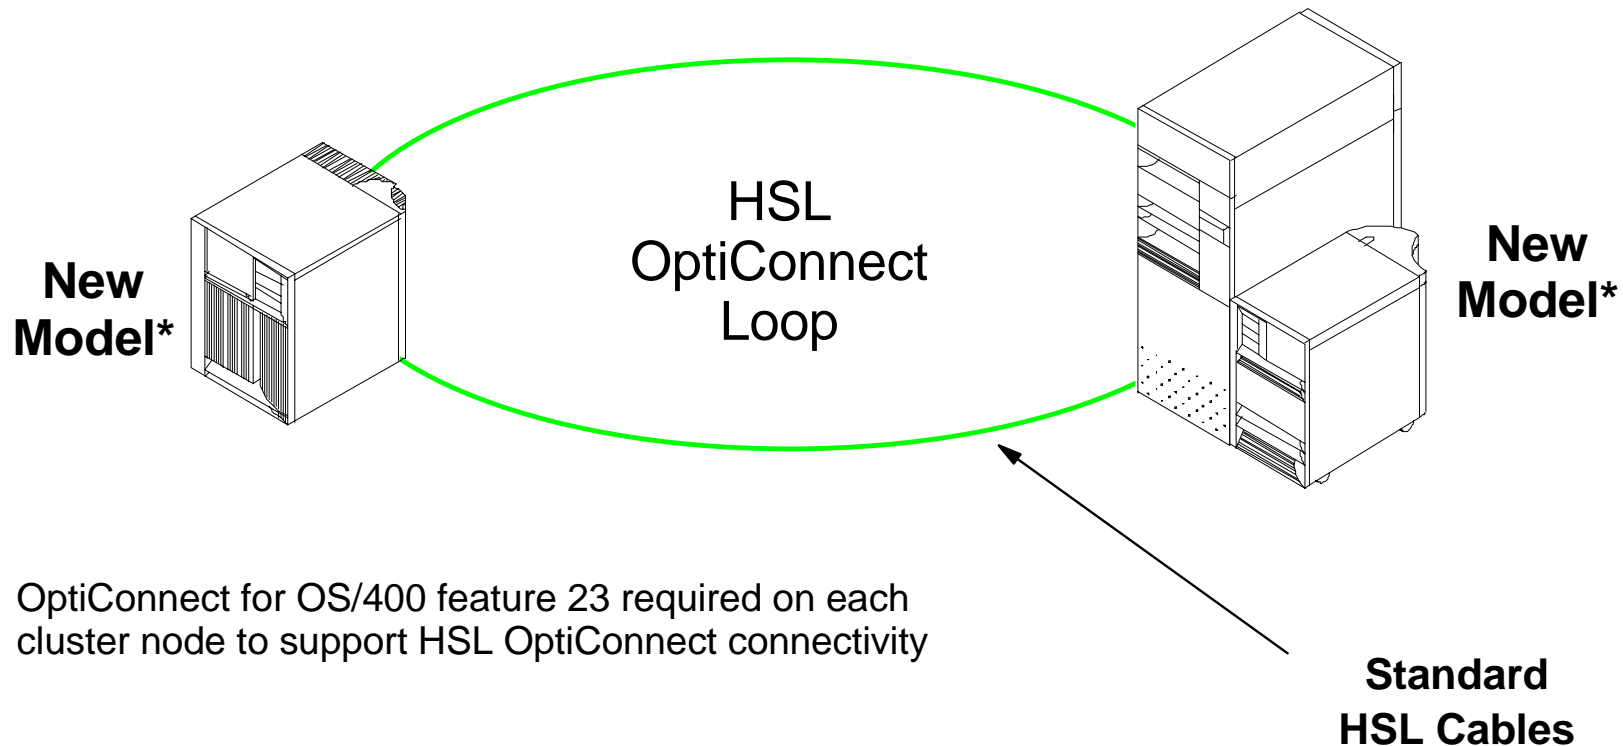
Besides the "new adapter" technology required, you must also have OptiConnect for OS/400 option 23 installed in order to support HSL OptiConnect high speed connectivity (for example fast DDM and remote journaling). This support is separate and independent of switched disk support.

Switched disk support via IASPs does require the new HSL adapter and OS/400 HA Switchable Resources, option 41.

The following chart shows an example of a simple HSL OptiConnect Loop.

Simple HSL OptiConnect

IBM  server iSeries



OptiConnect for OS/400 feature 23 required on each cluster node to support HSL OptiConnect connectivity

*830 and 840 Models available before April 2001 need to order a new HSL adapter. OptiConnect for OS/400 feature required on each cluster node.

IBM  server. For the next generation of e-business.

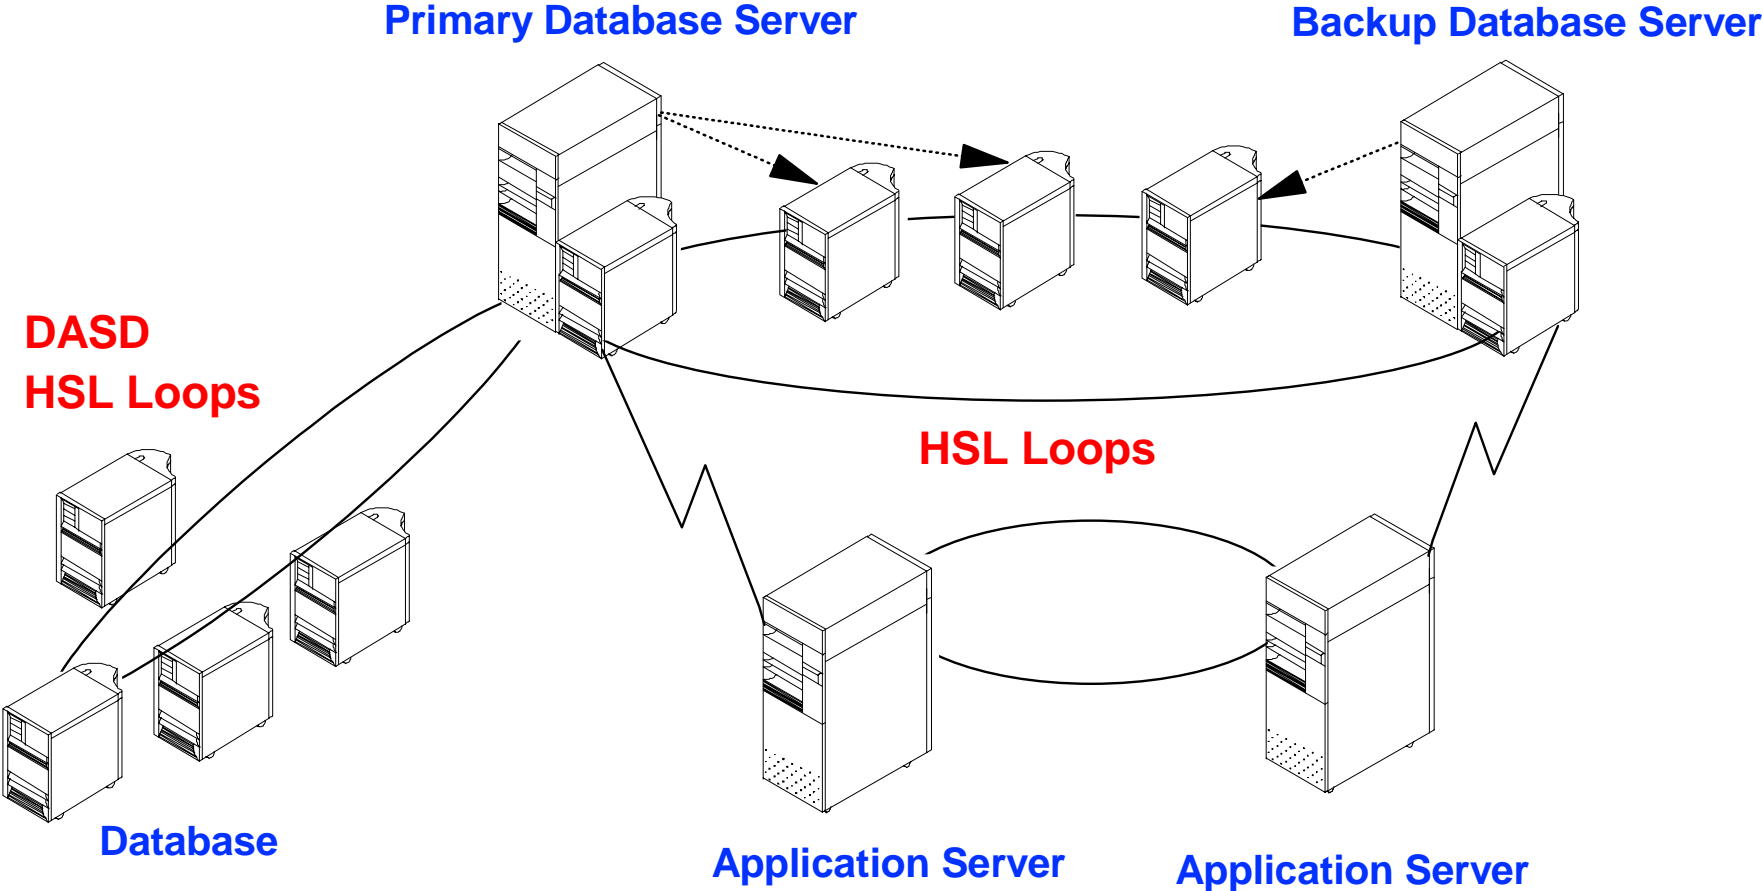
This foil shows a simple HSL loop using the new OptiConnect adapters.

All 270, 820, and 840 models announced April 2001 come with HSL adapters enabled for OptiConnect. The 830 and 840 models that were available before April 2001 can have one of the new OptiConnect capable adapters installed by ordering:

- #2754 Bus Expansion with 8 HSL ports: Enables clustering over HSL on iSeries 830 and SB2 (all processors except processor feature #2400)
- #2777 Bus Expansion with 8 HSL ports: Enables clustering over HSL on iSeries 830 processor # 2400
- #2755 Bus Expansion with 16 HSL ports: Enables clustering over HSL on iSeries 840 and SB3

The pre April 2001 270 and 820 models use the same boards to drive the processor(s) and HSL. Therefore they do not have the capacity to drive the new OptiConnect capable adapters.

Complex Switch Disk Cluster Example



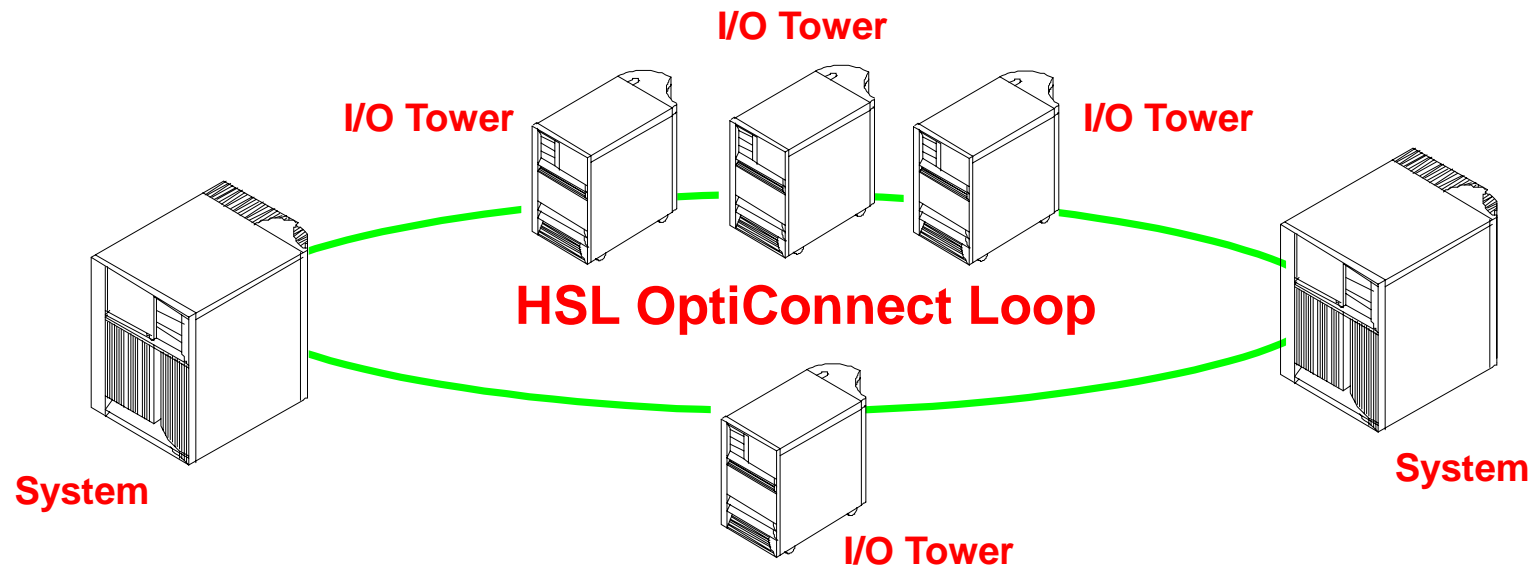
Notes: Complex Switch Disk Cluster Example

This configuration shows a network of HSL loops that could be used for full High Availability Business Partner clustering support of duplicating data and, or new Independent ASP switching. Exactly which functions would be used depends on the customer requirements and hardware cost considerations.,

When using the new OptiConnect HSL and simple clustering (switched disk) support, there are some hardware placement resource considerations that must be understood when planning to use the new support.

The next two foils graphically depict HSL OptiConnect connectivity rules and placement rules. See the Hardware presentation for more details.

V5R1 Switch Disk and HSL OptiConnect



Maximum of

- 2 cluster nodes (systems) per loop
- 4 external towers (including IXS towers)
- 3 external towers per Loop segment
- All switchable towers on one loop segment are in the same device CRG and switch together
- All switchable towers on one loop segment are in the same SPCN power domain

OS/400 HA Switchable Resources, optional feature 41 required
on each cluster node to support switchable I/O tower capability

HSL OptiConnect Loop connectivity rules

IBM  server. For the next generation of e-business.


You can have a switch disk cluster without using the system-to-system high speed communications provided by HSL OptiConnect. In this case you do not need feature 23. In a multi-system environment, you still need the HSL OptiConnect 23 to send and receive communications data..

You can have HSL OptiConnect connectivity between two systems without having a switchable tower. In this case you do not need feature 41, but you do need feature 23.

Nomenclature review:

- **HSL Loop:** The technology for connecting I/O Towers to the system
- **HSL OptiConnect Loop:** The same HSL Loop technology but in this case two systems are on the same loop. We distinguish between the two loops in the following manner:
 - Model 840: 8 HSL Loops, 4 of these loops can be HSL OptiConnect Loops
 - Model 830: 4 HSL Loops, 2 of these loops can be HSL OptiConnect Loops
 - Model 820: 1 HSL Loop. This loop can also be an HSL OptiConnect Loop
 - Model 270: 1 HSL Loop. This loop can also be an HSL OptiConnect Loop

V5R1 Switch Disk Placement Rules



Adjacency rule

- ★ switched tower must be physically adjacent to the alternate system or tower owned by the alternate system

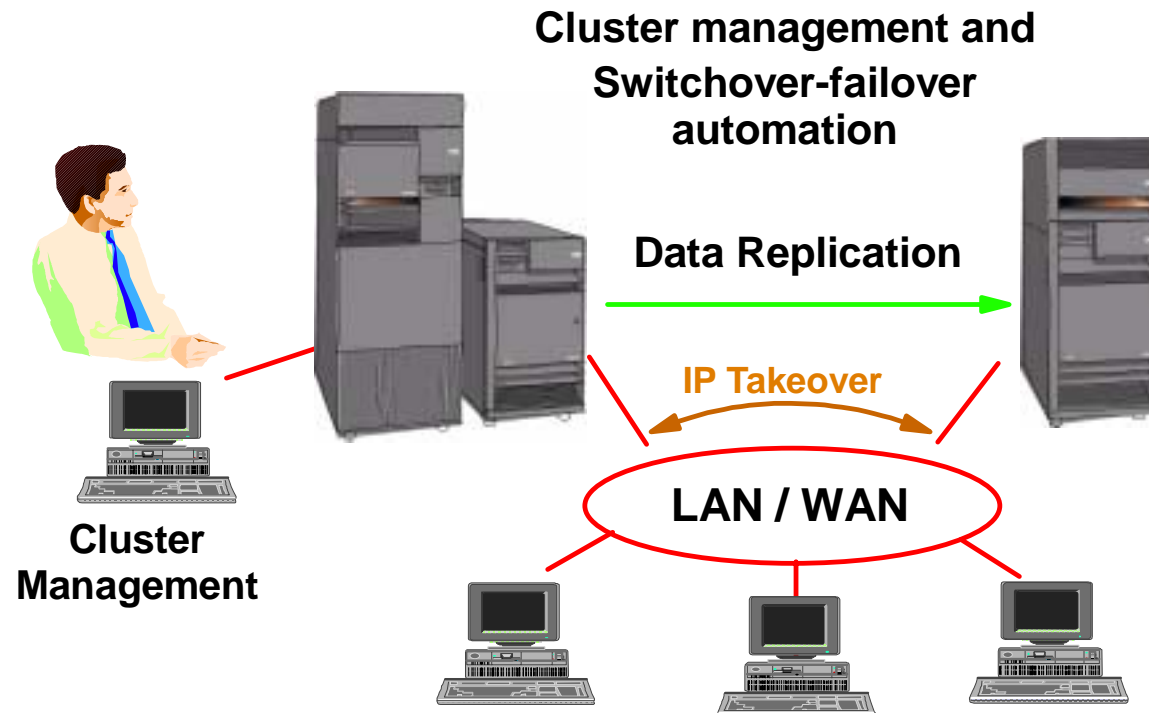
HSL Segment Rule

- ★ switch tower must reside on the HSL segment connecting home and adjacent systems for that tower

V5R1 Clustering Enhancements

IBM @server. For the next generation of e-business.

Cluster Solution for iSeries



- Data & Application Resiliency
 - Addresses Disaster Recovery
 - Addresses Planned, Unplanned outages
 - Addresses Resource Management
 - All resources concurrently usable
 - Save Window Elimination
- Continuous Availability Solution

OS/400
Functions

IBM Cluster
Middleware
Business Partners



IBM  server. For the next generation of e-business.

This foil summarizes the full range of capabilities available through clustering support, before we go over some V5R1 enhancements and then into simple clustering or "switched disks ." support:

- Tools to create and manage clusters, the ability to detect a failure within a cluster, and switch over and fail over mechanisms to move work between cluster nodes for planned or unplanned outages
- A common method for setting up object replication for nodes within a cluster. This includes the data and program objects necessary to run applications that are cluster enabled
- Mechanisms to automatically switch applications and users from a primary to a backup node within a cluster for planned or unplanned outages.

This clustering framework is built around a set of system APIs, system services, and exit programs. This clustering architecture requires teamwork between IBM and business partners to provide the total solution. Data replication services and the cluster management interface is provided by IBM's HABPs. iSeries software solution providers deliver ClusterProven applications that utilize OS/400 cluster support, or you can implement OS/400 cluster support within your own applications.

A well managed iSeries cluster can provide the highest levels of availability of any individual server in the industry. Small outages, tolerated just a few years ago, can now mean significant losses of revenue and future opportunities for your business today. Clusters are the best solution for continuous availability requirements on an iSeries, providing rapid recovery for the widest range of outages possible, with minimal cost.

Combined with the benefits of continuous availability, the second server can be treated as a production server by performing tasks such as save operations, database queries, batch reporting, and act as a web server for inquiries only thereby freeing up some of the resources from the production server.

OS/400 Clustering support includes the connectivity ("cluster fabric") of SPD OptiConnect software and hardware, communications lines (including ATM or 1 Gbps Ethernet) and now, with V5R1 the new OptiConnect over HSL support. OptiConnect is the most elaborate cluster fabric, supporting special application services and exceptional performance.

Switchable Independent Auxiliary Storage Pools*

- Resilient Cluster Device
- Device Domain

IBM Cluster Management Utility within Management Central central server

Tuning cluster performance

Distribute Information

Cluster versioning

Improved handling and recovery for cluster partitions

Example commands and exit program

*Note: The V5R1 Independent ASP - switched disks support includes IFS objects but not library-based objects such as OS/400 Database and Journal objects. This means objects within a system or user ASP may be journaled, but objects within an IASP cannot be journaled in V5R1. Journaling IASP data is planned for next release.

Switchable Independent Auxiliary Storage Pools (additional foils follow)

- Resilient Cluster Device: a hardware resource (IASP in V5R1) represented by a contiguous object that can be switched between system in the event of a system outage (failed or planned)
- Device Domain: a subset of cluster nodes across which a set of resilient devices can be "shared". A Device Domain prevents conflicts that could cause resilient device switching to fail. The Resilient Cluster Device IASP can be active only on one system at a time.

IBM Cluster Management Utility within Management Central: You can use the IBM Simple Cluster Management utility to create and manage a two-node, switched disk cluster. IBM provides a Simple Cluster Management interface that is available through Operations Navigator and accessible through Option 41 of OS/400. The utility allows you to create and manage a cluster that uses switchable IASPs to ensure data availability. Simple Cluster Management features a wizard which steps you through the creation of a simple, two-node cluster. Additional cluster management can be accomplished using this interface including tasks such as:

- Adding a node to an existing one-node cluster
- Adding a switchable hardware group to a cluster
- Adding a switchable software product to a cluster
- Changing the cluster description
- Changing the exit program name
- Changing the takeover IP address for a switchable software product
- Deleting a cluster
- Starting clustering
- Stopping clustering
- Switching cluster resources from the primary node to the backup node
- Viewing messages about cluster activity

Note: When using the IBM Simple Cluster Management utility, you should avoid using the commands in QUSRTOOL.

Notes: V5R1 OS/400 Clustering Enhancements-2

IBM  server iSeries

Tuning cluster performance: APIs are available for basic tuning, such as allowing you to set the tuning parameters to a predefined set of values identified for high, low, and normal time-out and messaging interval values.

Distribute Information: The Distribute Information (QcstDistributeInformation) API provides you a mechanism to send

information from one node in the CRG's recovery domain to other nodes in the recovery domain. This can be a useful mechanism to communicate application activity or to send small amounts of information related to the application to affected nodes.

Cluster versioning: A cluster version represents the level of function available on the cluster. Versioning is a technique that allows the cluster to contain systems at multiple release levels and fully interoperate by determining the communications protocol level to be used.

Improved handling and recovery for cluster partitions: In addition to better detection of some failover conditions, cluster resource services provides an easier way to change partition nodes to failed.

Example commands and exit program: A set of example commands are provided in QUSRTOOL that can be used to create and manage a cluster in some environments. See the member TCSTINFO in the file QUSRTOOL/QATTINFO for more information on these example commands. An example application CRG exit program is also included in the QUSRTOOL library. The sample source code can be used as the basis for writing an exit program. See the TCSTAPPEXT member in the QATTSYSC file for an example written in ILE C.

See Clusters in V5R1 Information Center for more information:

- <http://www.ibm.com/eserver/series/infocenter>

IBM  server. For the next generation of e-business.

Resilient Cluster Device is: a hardware resource represented by a configuration object that can be switched between systems in the event of a system outage.

Possible examples:

- **IASPs (CRTDEVASP)**
- Removable media (not V5R1)
- Communications devices (not V5R1)

Requires several new cluster constructs:

- Device domain
- Resilient device CRG

Enforced relationship between CRG and switchable entities

- Enabled by OS/400 optional feature 41 HA Switchable Resources

A Resilient Cluster Device represents the hardware-based object that can be switched between systems in the event of a planned or unplanned fail over. There is a plan for object types to be included as a resilient cluster device that includes removable media and communications hardware, but for V5R1 the only supported object is an Independent Auxiliary Storage Pool (IASP).

The new for V5R1 switchable IASP (switched disks) support requires the new cluster constructs. These are discussed on following foils:

- Device Domain
- Resilient Device CRG (Cluster Resource Group)

There is an enforced relationship between the Resilient Device CRG and the switchable entities that is enabled through OS/400 Option 41 - HA Switchable Resources, which is an additional cost, licensed option.

Device domain:

- Subset of cluster nodes across which a set of resilient devices can possibly be "shared"
- Prevents conflicts and that would cause resilient device switching to fail

Cluster resources negotiated across a device domain include:

- IASP number assignments
- DASD Unit number assignments
- Virtual addresses

New cluster interfaces:

- Add / remove device domain entry
- List device domain information

The device domain is a subset of cluster nodes across which a set of resilient devices can be possible "shared." The sharing is not concurrent from each node. During configuration on the "primary node" the secondary node is made aware of the individual hardware within the CRG and is "ready to receive the CRG" should it be switched.

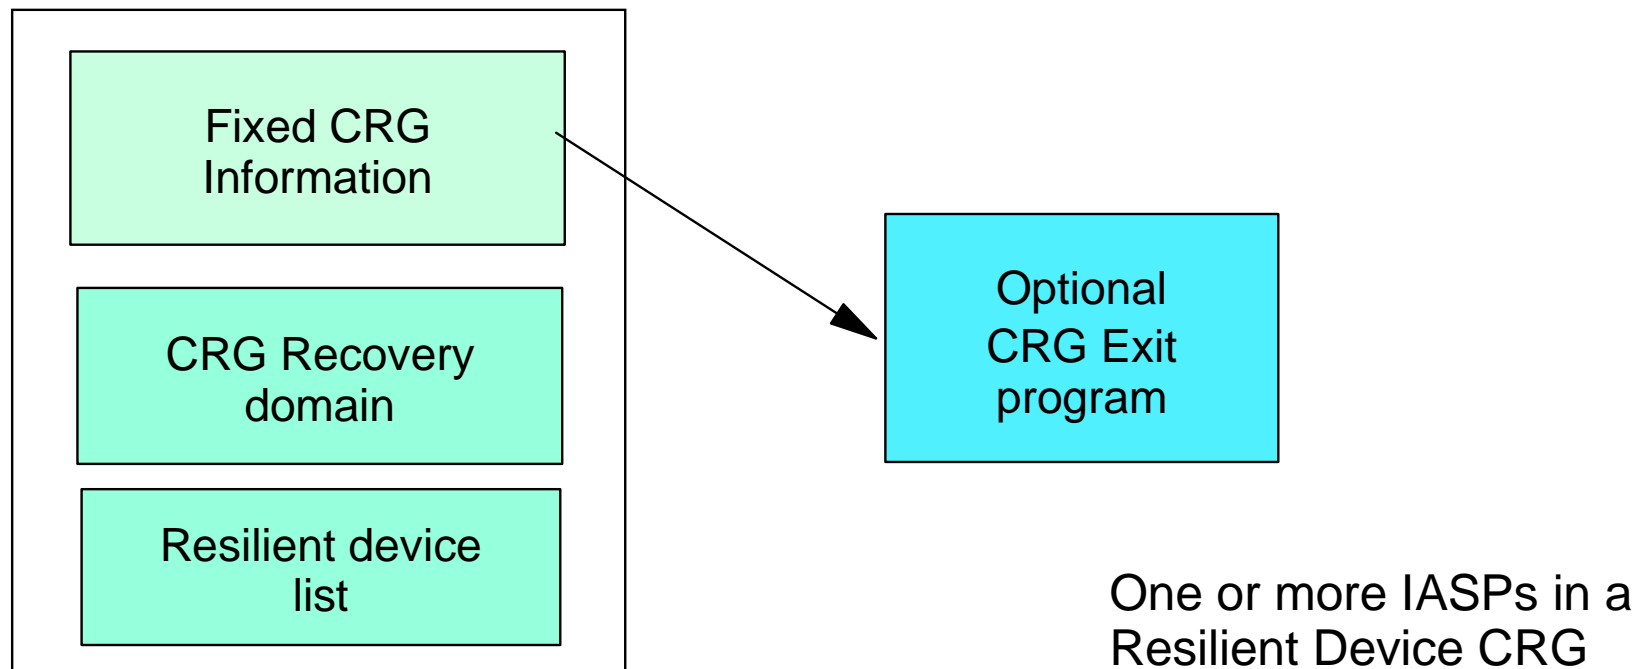
The following cluster resources are negotiated across a device domain:

- IASP number assignments: The IASP numbers are 33-99 and automatically assigned to the named (by the user) resource. The system manages the assigned IASP numbers, which may not always be in numerical order, based upon creation date.
- DASD unit number assignments: So as to not conflict with "permanently attached disk units to each node" the IASP disk unit numbers start with 4001.
- Virtual address assignments: The cluster configuration determines the virtual address space required on the primary node and communicates that to the secondary node.

New cluster interfaces to support Device Domains include:

- Add or remove Device Domain entry
- List Device Domain information

Note that this presentation will focus on the Operations Navigator interface to this new support.



New type of CRG in V5R1 -- Resilient device

- Like Resilient Data and Resilient Application, but ...

CRG switchover / failover order is device, data, then application

New cluster interfaces:

- Manage device CRG
- Add / remove resilient device list entry

In V5R1 the Resilient Device CRG is comprised of one or more IASPs.

A Resilient Device CRG can be manual switched (switchover) automatically switched (failover) to a secondary node.

New Cluster interfaces are required to support this new type of CRG:


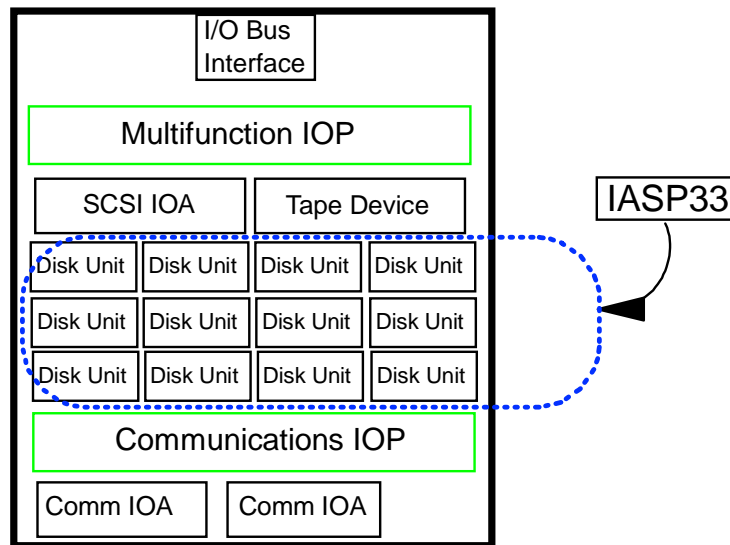
- Manage Device CRG
- Add or Remove resilient device list entry

A CRG exit program can be optionally written. In the simple clustering support highlighted in V5R1 the exit program is not required. This support is available for High Availability Business Partners solutions.


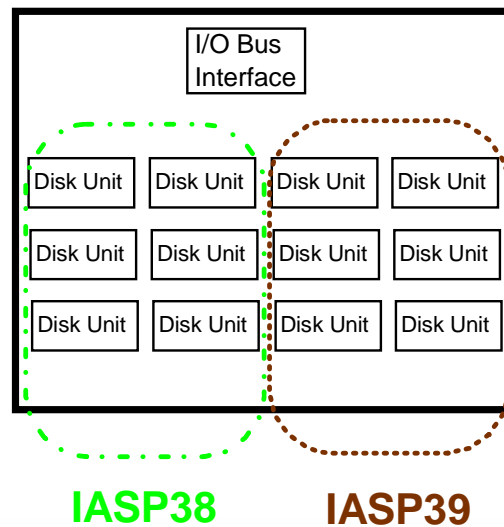
There can be one or more IASPs in a device list for a specific Device CRG. Some of the possible combinations are shown in the following foils.

I/O Tower, IASP, Device CRG Examples

1 tower, 1 IASP,
1 RD CRG



1 tower, multiple
ASPs, 1 RD CRG



In this foil we show two examples of the new IASP with Resilient Device (RD) Cluster Resource Group.

In the simplest case (top) you see a single set of disk units configured into a single IASP (IASP33) and assigned to that RD CRG on the upper right. The disk units are in a single tower, which is required for switching between separate systems. If switching between partitions on a single system, you need the disk units on a separated IOP. The device list entry is the IASP name.

In a slightly more complex example (bottom) you see two IASPs assigned to 1 RD CRG. The device list has two entries (IASPs).

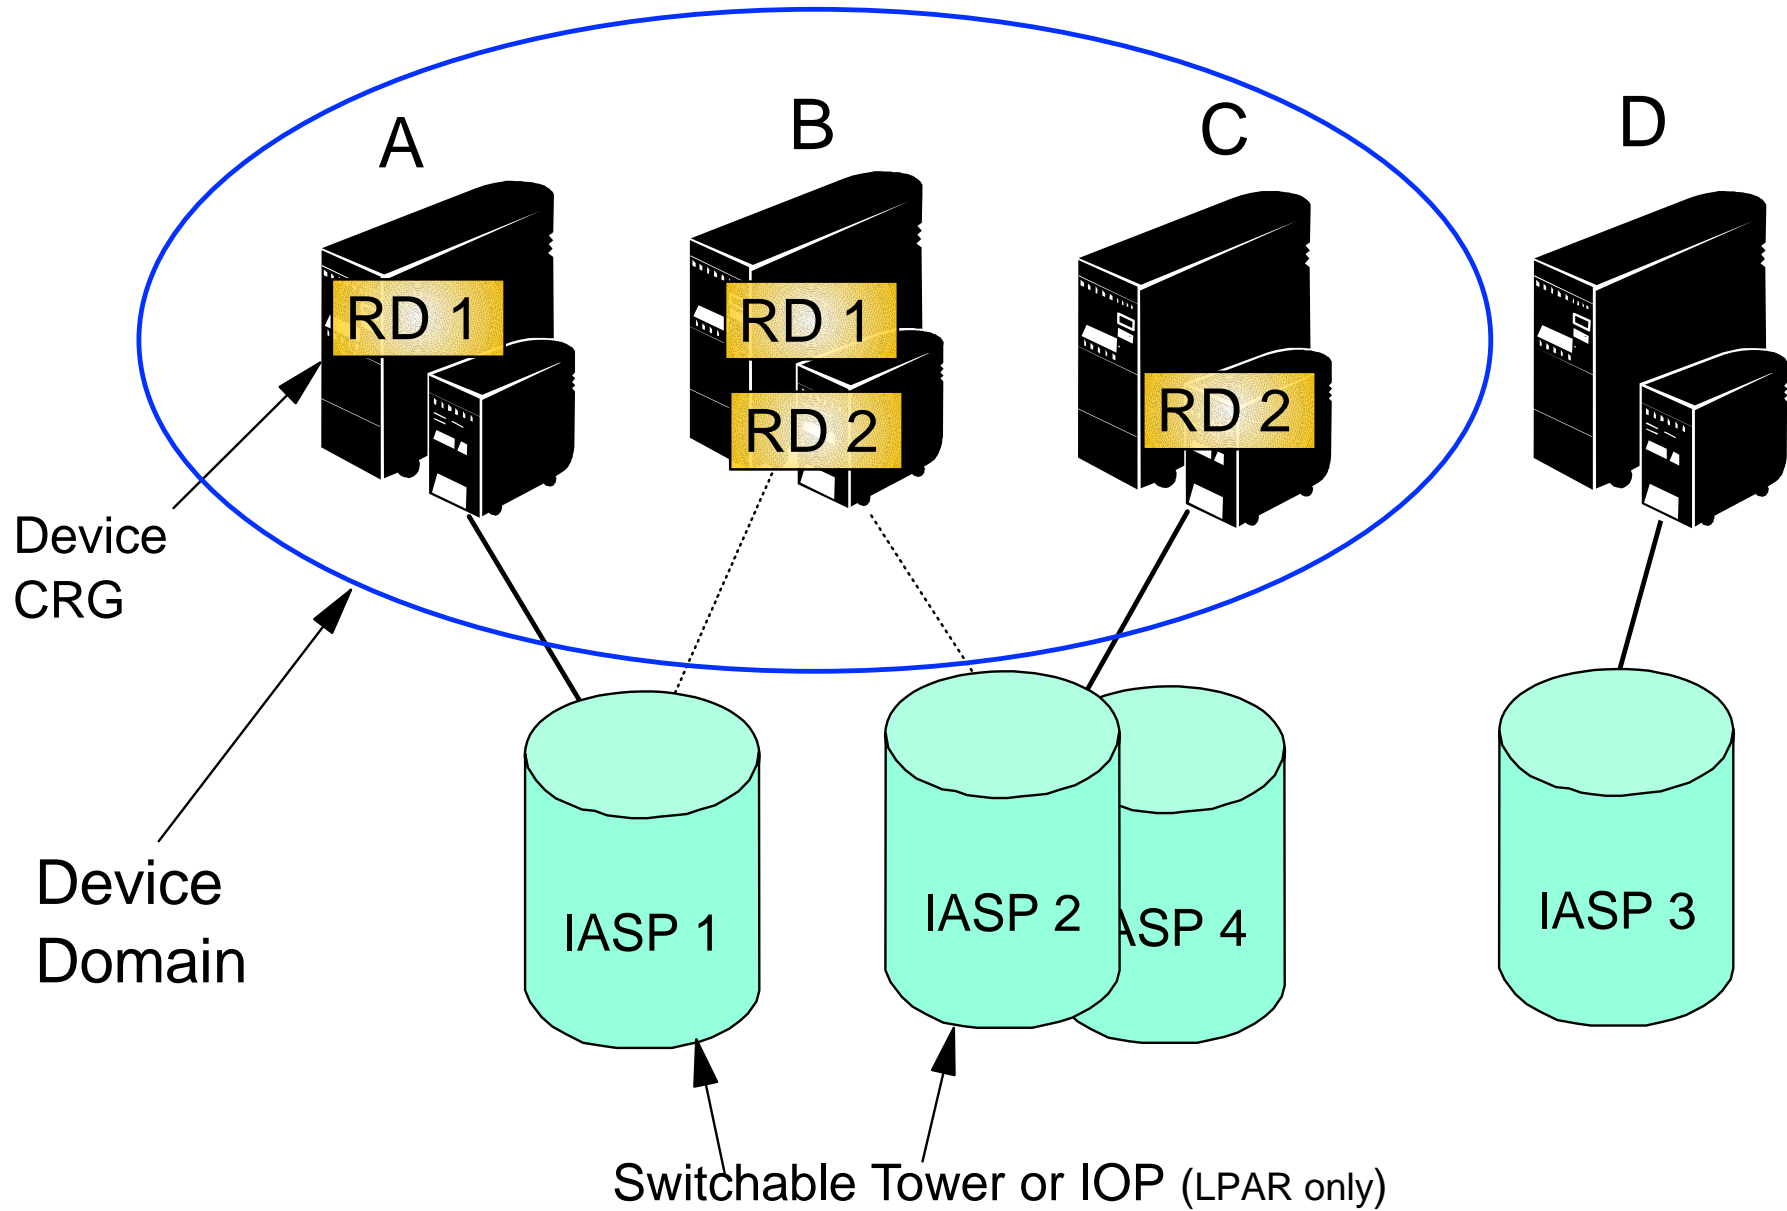
There are other more complex configurations supported, such as:

- Multiple towers, 1 IASP, 1 RD CRG
- Multiple towers, Multiple IASPs, 1 RD CRG

Complex examples and scenarios are beyond the scope of this presentation. However, some Operations Navigator Cluster and IASP (disk pool) screen captures are shown. A simple switch over example sequence is also shown.

Note: When an I/O tower is switched to another node in the cluster, all its resources, such as an internal tape or CD-ROM are also switched to the second system. After a few minutes these "other non-disk" devices are recognized by the "switched to" system and become ready for use on that system.

Switchable IASPs, Device CRG, Device Domain



IBM e server. For the next generation of e-business.

Notes: Switchable IASPs, Device CRG, Domain

This foils depicts a moderately complex clustering setup for switchable IASPs.

You see a Device Domain represented by the oval line and Resilient Device CRGs with the solid lines between the nodes and the IASPs - I/O tower or IOP if the nodes are actually partitions within the same system.

IBM @server. For the next generation of e-business.

Device list can contain 1 or more IASP device entries

A device can belong to at most 1 device CRG

Configurations that hinder switching are not allowed

Changes are prevented when correctness of the change cannot be verified (for example, there is an inactive node or cluster partition)

Device CRG must contain at least 1 device entry before it can be started

The IOP (LPAR configuration) or I/O Tower must be accessible to all nodes in the recovery domain


Configuration object can be varied on (activated) independently of the CRG, but it only on the primary system

This foil lists rules and requirements for configuring and switching resilient devices.

The statement "the IOP or I/O tower must be accessible to all nodes in the recovery domain" means there are specific HSL loop position locations that are required. See the Hardware presentation for more information on this.

Simple Clustering Example - Steps

0. Complete environment plan
1. Create the cluster
2. Identify which nodes are to be in the device domain
3. IASP device descriptions created on appropriate nodes
4. Create the device CRG
5. Configure disk units into IASPs (done w/ DASD mgmt utility)
6. Populate the IASP with user defined file system



To set up a cluster environment that includes resilient devices, care must be taken so that conflicts are avoided across the cluster. The following set of steps help you achieve a successful configuration for collections of switchable resources.

- **Complete the environment plan.** Plan for both the hardware and software configurations. For example, determine what I/O towers will be switching between what systems. Also, identify any systems that might potentially be added to the recovery domain for a device CRG. If you do this now, you can avoid conflicting resource assignments that might prevent you from doing this later. Determine what resources will be included in the switchable tower -- remember, when a switchover or failover occurs, the entire switchable tower is moved to the backup system. Also determine other physical requirements, such as floor space and power domains.
- **Create the cluster.** If it does not already exist, create the cluster. Add all of the nodes to the cluster. To take advantage of resilient device capabilities, all nodes must be at potential cluster version 2 and the current cluster version must be set to 2. Start all nodes in the cluster or at least those that will be in device domains.
- **Identify which nodes are to be in device domains.** Add nodes to their respective device domains, including those nodes that might eventually participate in the switching action for a device CRG. When this is done, the resources associated with the device domain (such as disk unit numbers and virtual address ranges) are assigned across all nodes in the device domain to avoid resource conflicts.
- **Create the IASP device description.** On every node that is to be in the recovery domain for an IASP, create a device description using the same parameter values. Under Operation Navigator this is implicitly done for you when configuring on the primary node.
- **Create the device CRG.**

- **Identify which nodes will initially be in the recovery domain.** In the CRG device list, specify which IASPs will be switchable under the CRG and if the system is to vary them on when a switchover or failover occurs. Identify an exit program for the CRG if one is desired.
- **Configure the disk units into the IASP.** The DASD management utility (used by Operations Navigator) provides a useful tool to do this. When this step is done, it is best to have every node in the CRG's recovery domain active. That allows the nodes to communicate and determine if all disk units can be accessed by all nodes and thus avoid trouble later.
This is required when using the Operations Navigator interface.
- **Make Available (vary on) the IASP.** When you create the IASP device the system automatically generates a User Defined File System (UDFS) named the same as the device (IASP) name. You must place files meaningful to your application environment into this default UDFS. While the IASP is "varied off/unavailable" any files in the UDFS cannot be accessed. In fact, through the Operations Navigator-Integrated File System neither the root/dev/JIMC ASP name nor the specific UDFS name JIMC can be seen.
When that IASP is "made available/varied on" the default file system is automatically mounted in the root directory of the system's IFS. The root/dev/JIMC ASP and UDFS JIMC can now be seen through Operations Navigator-IFS view.
The UDFS is accessible only when the IASP device is varied on and then only when the device CRG ("Switch" in the examples that follow) is active and "primary" within the cluster. When the device DCR is varied off or switched to a secondary node, the UDFS is "unmounted" and the objects within that file system are no longer accessible under IFS on original node.
- **Populate the IASP with data.** Create or move the appropriate user-defined file systems into the IASP.

Simple Clustering Example - Steps ...

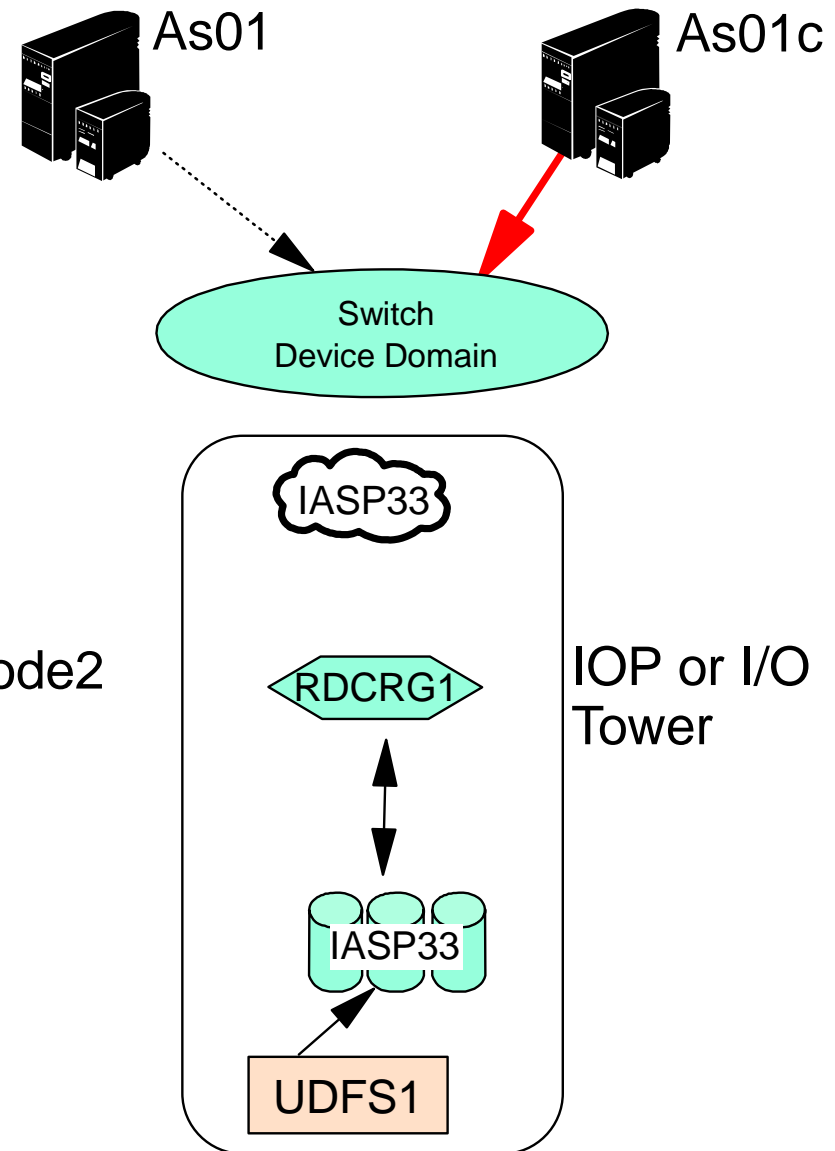
7. To initiate a switch over, quiesce the affected application(s) on node1

8. Make the IASP unavailable

9. Switch the device Cluster Resource Group

10. Make the newly received IASP available on node2

11. Activate the application on node 2



Quiesce the affected application(s). You need to shut down the application(s) using the objects within the IASP User Defined File System. V5R1 Operations Navigator does Opearg nod oOusRemember that after the Device Cluster Resource Group that contains this IASP is switched to a secondary system, you must "Make Available" that IASP after that switching function has completed. You will be able to see the disk devices and associated IASP on the secondary system once the switch has completed. But you will not be able to see the UDFS until it is mounted with "Make Available/Vary On."

Switch disk support integrated into cluster architecture

Cluster

collection of iSeries servers

Device Domain

Collection of cluster nodes that share resources (switchable DASD towers)
Manages assignment of common IASP id, disk unit & virtual addresses across domain

Device CRG

Cluster Control object for a set of I-ASPs (switchable towers)



Device Description

Logical control name for varying on/off an I-ASP

I-ASP

Independent Disk Pool (I-ASP) defines a physical set of switchable drives

prereq: cluster



prereq: cluster, device description and Optional Feature 41 license

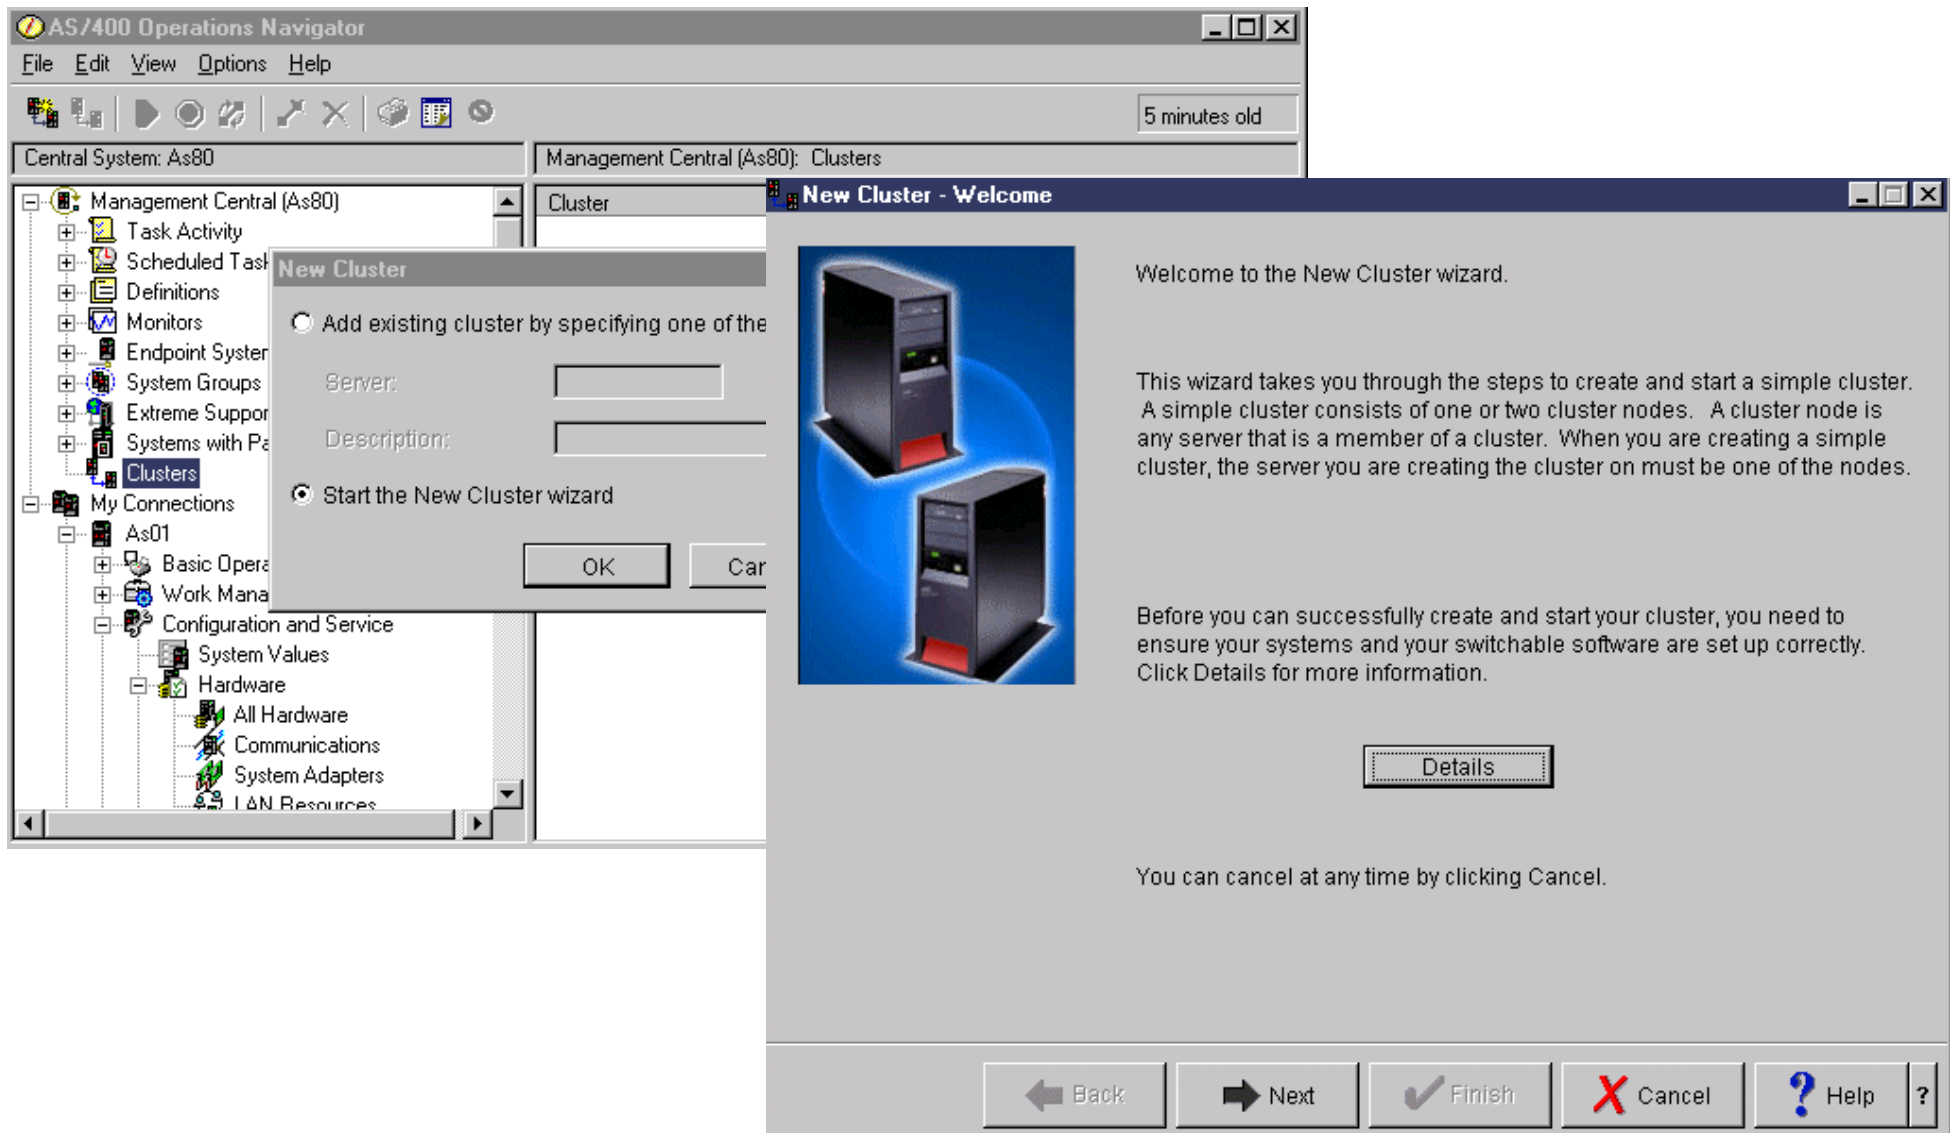
prereq: cluster and Optional Feature 41 license

prereq: TCP/IP connectivity to all nodes, V5R1 HSL port hardware, I/O Tower or DASD on LPAR shared bus

Sample Operations Navigator Interface to Simple Clustering - Switched Disks

IBM @server. For the next generation of e-business.

New Cluster



AS/400 Operations Navigator

File Edit View Options Help

5 minutes old

Central System: As80 Management Central (As80): Clusters

Management Central (As80)

- Task Activity
- Scheduled Task
- Definitions
- Monitors
- Endpoint System
- System Groups
- Extreme Support
- Systems with Pa
- Clusters

My Connections

- As01
 - Basic Operat
 - Work Mana
 - Configuration and Service
 - System Values
 - Hardware
 - All Hardware
 - Communications
 - System Adapters
 - LAN Resources

New Cluster

Add existing cluster by specifying one of the

Server:

Description:

Start the New Cluster wizard

OK Cancel

New Cluster - Welcome

Welcome to the New Cluster wizard.

This wizard takes you through the steps to create and start a simple cluster. A simple cluster consists of one or two cluster nodes. A cluster node is any server that is a member of a cluster. When you are creating a simple cluster, the server you are creating the cluster on must be one of the nodes.

Before you can successfully create and start your cluster, you need to ensure your systems and your switchable software are set up correctly. Click Details for more information.

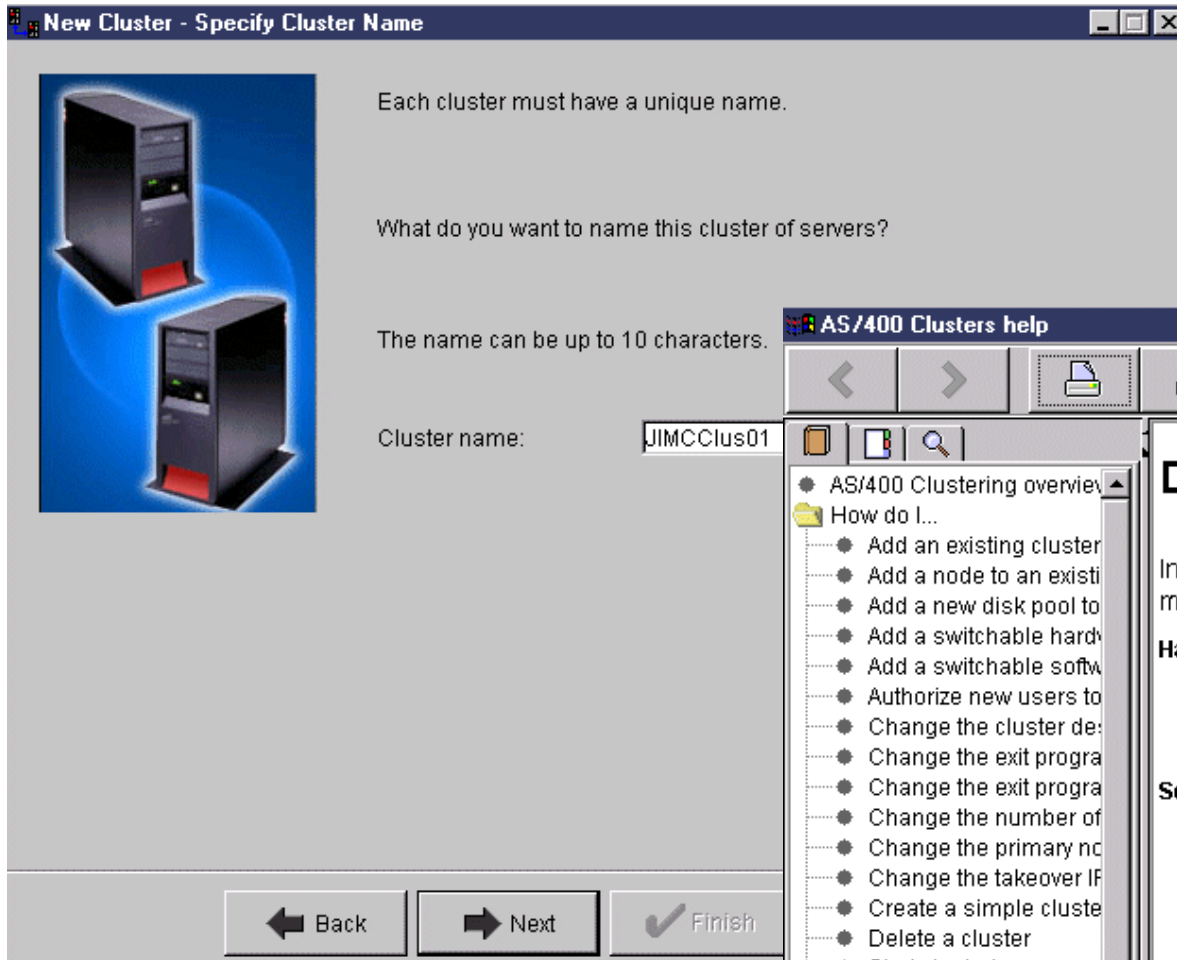
Details

You can cancel at any time by clicking Cancel.

Back Next Finish Cancel Help ?

IBM  server. For the next generation of e-business.

New Cluster name, and software requirements IBM server iSeries




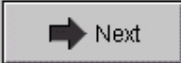

New Cluster - Specify Cluster Name

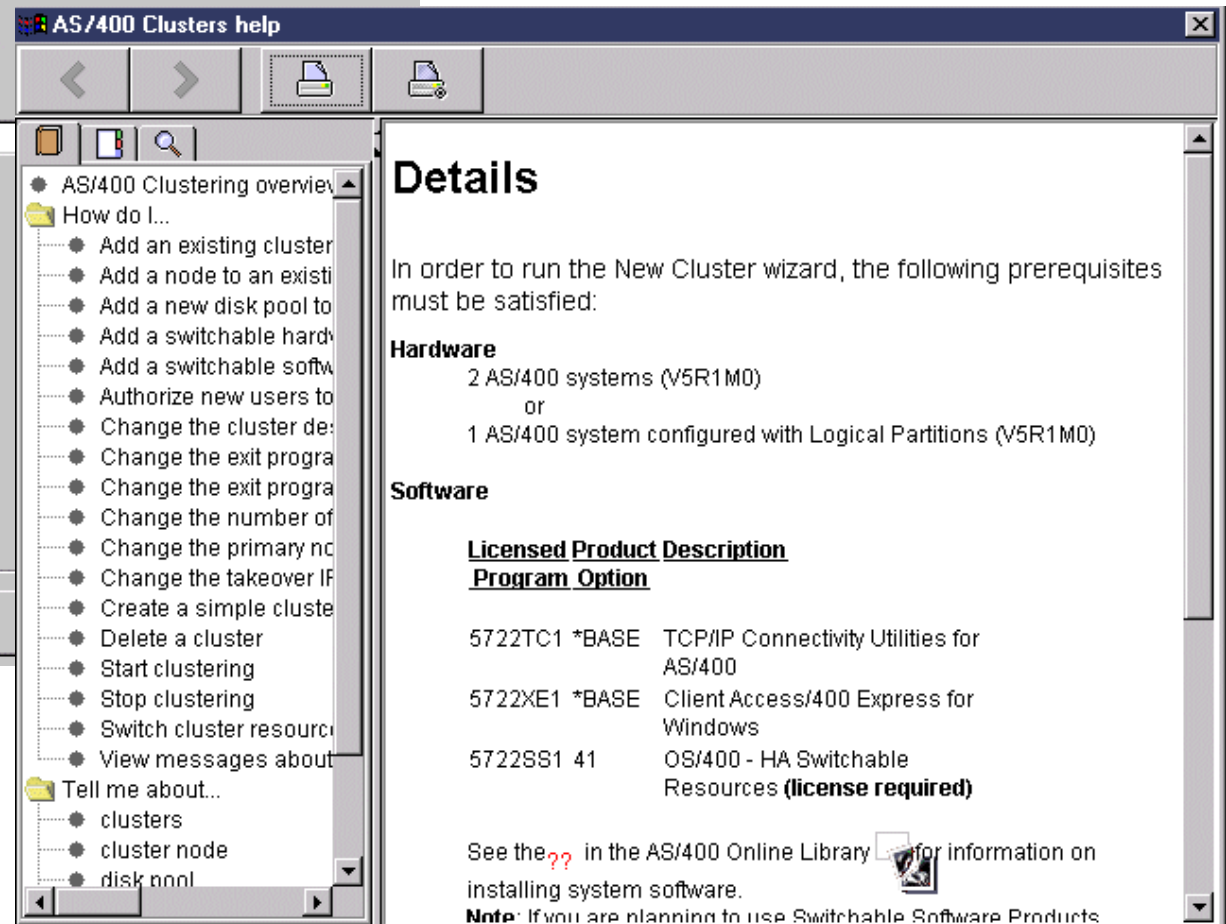
Each cluster must have a unique name.

What do you want to name this cluster of servers?

The name can be up to 10 characters.

Cluster name:



AS/400 Clusters help

Details


In order to run the New Cluster wizard, the following prerequisites must be satisfied:

Hardware

- 2 AS/400 systems (V5R1M0)
- or
- 1 AS/400 system configured with Logical Partitions (V5R1M0)

Software

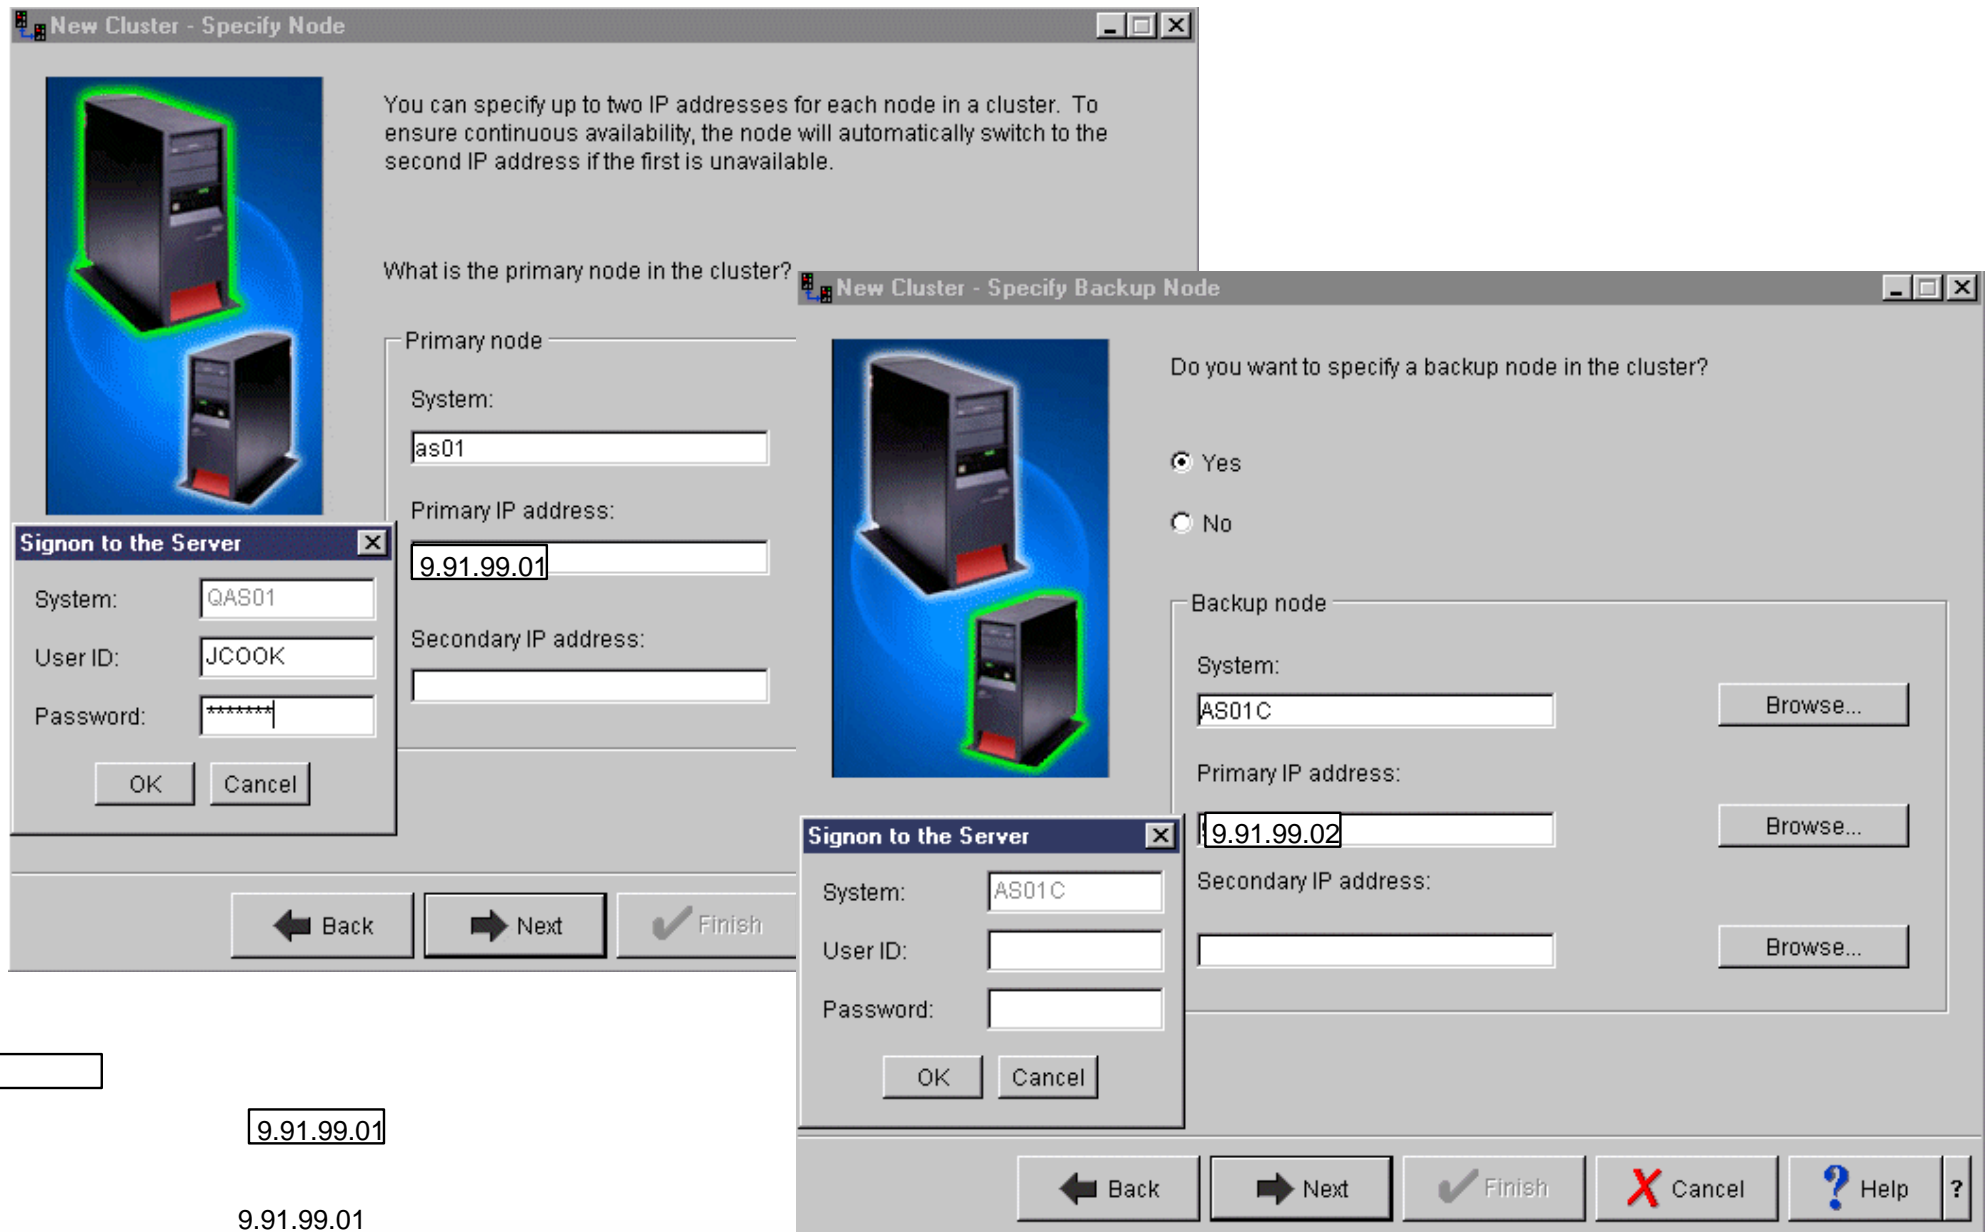
<u>Licensed Product Description</u>	<u>Program Option</u>
5722TC1 *BASE TCP/IP Connectivity Utilities for AS/400	
5722XE1 *BASE Client Access/400 Express for Windows	
5722SS1 41 OS/400 - HA Switchable Resources (license required)	

See the ?? in the AS/400 Online Library  for information on installing system software.

Note: If you are planning to use Switchable Software Products

IBM  For the next generation of e-business.

New Cluster-Node Definition



New Cluster - Specify Node

You can specify up to two IP addresses for each node in a cluster. To ensure continuous availability, the node will automatically switch to the second IP address if the first is unavailable.

What is the primary node in the cluster?

Primary node

System: as01

Primary IP address: 9.91.99.01

Secondary IP address:

Signon to the Server

System: QAS01

User ID: JCOOK

Password: *****

OK Cancel

← Back → Next ✓ Finish

New Cluster - Specify Backup Node

Do you want to specify a backup node in the cluster?

Yes

No

Backup node

System: AS01C

Primary IP address: 9.91.99.02

Secondary IP address:

Signon to the Server

System: AS01C

User ID:

Password:

OK Cancel

← Back → Next ✓ Finish ✗ Cancel ? Help ?

9.91.99.01

9.91.99.01

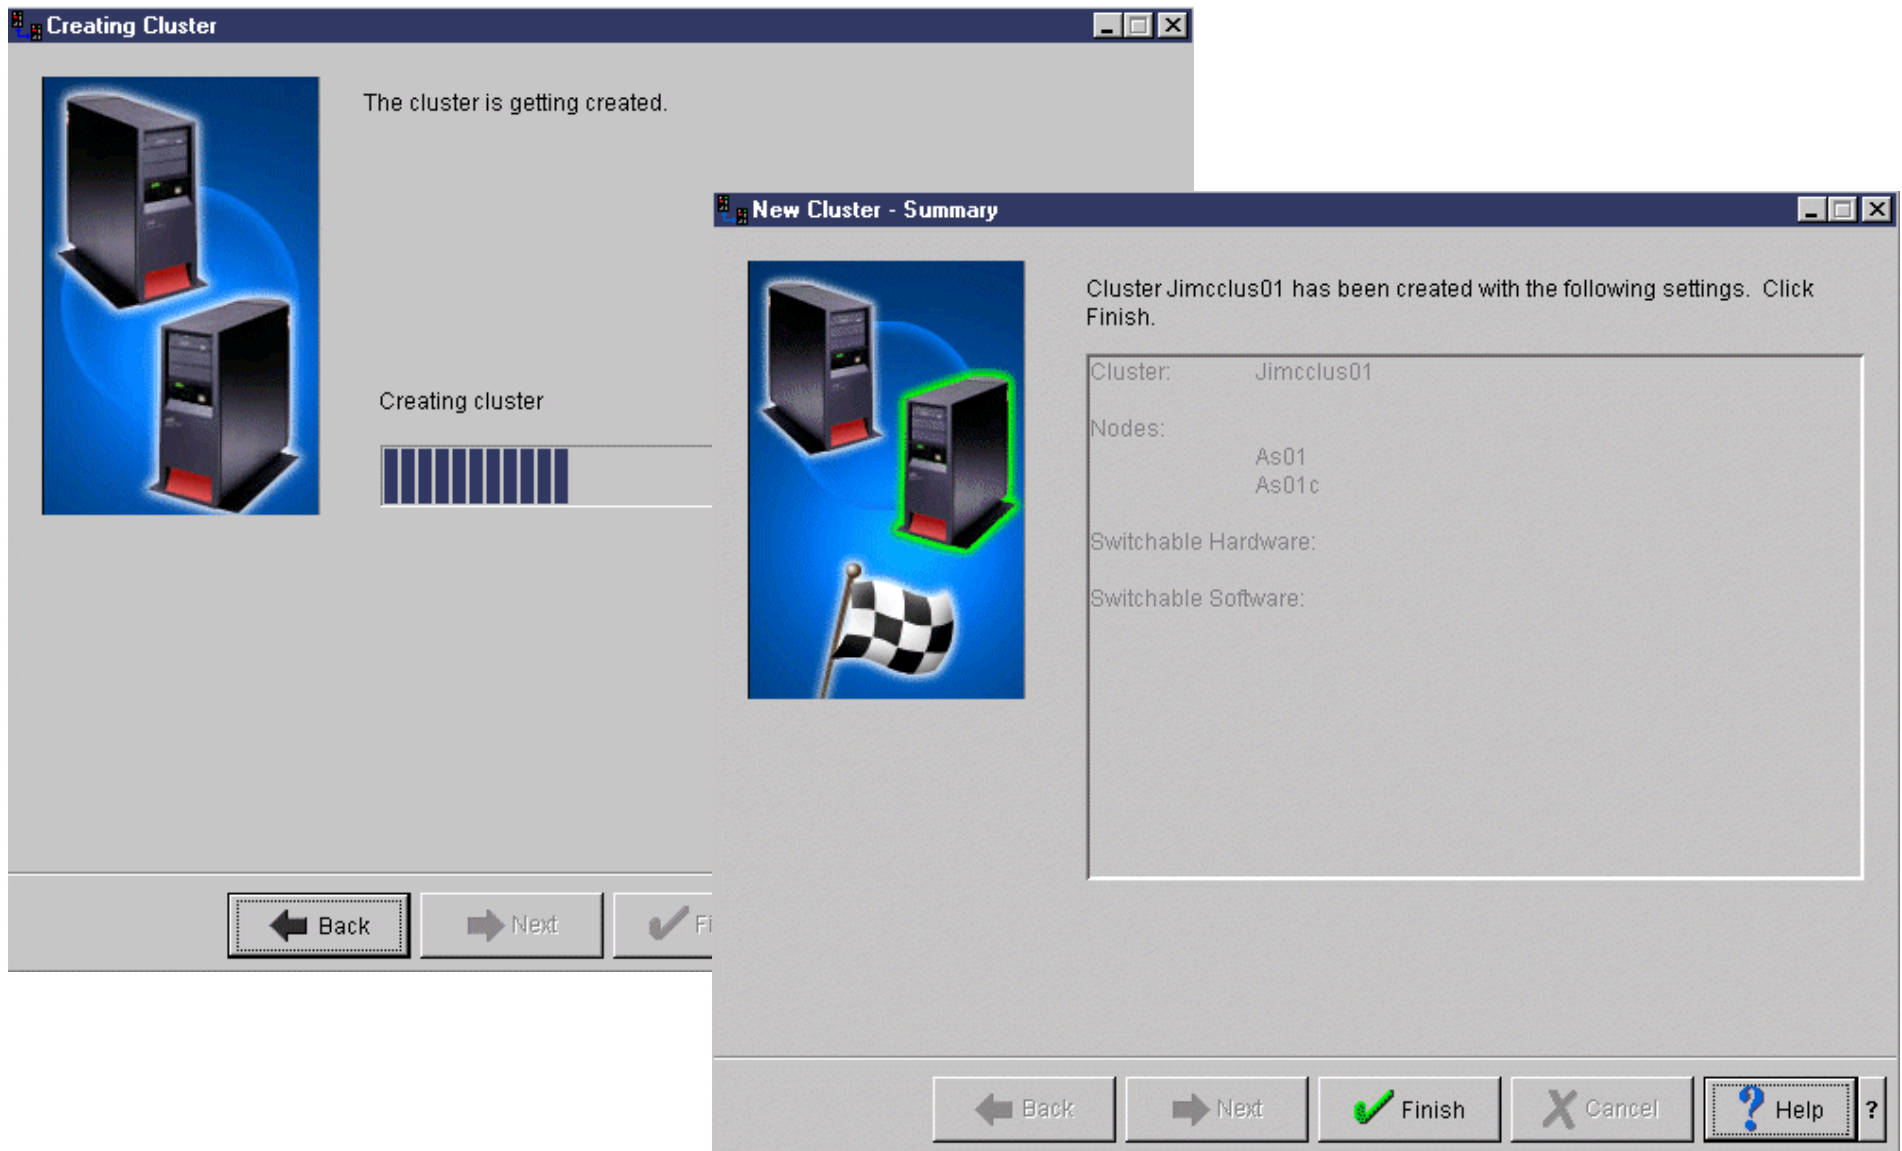
IBM  server. For the next generation of e-business.

Notes: New Cluster-Node Definition

In this example we are defining simple clustering with new nodes. We actually used a 3 system - AS80 as our Management Central central server for the two endpoint systems - AS01 and AS01C that are the cluster nodes. AS01 is the primary node where the IASP is configured and originally has the I/O devices attached.

You must specify at least one IP address for each node and optionally a secondary IP address for each node.

New Cluster Created



The screenshot displays the 'Creating Cluster' wizard interface. The main window, titled 'Creating Cluster', shows a progress bar and the text 'The cluster is getting created.' and 'Creating cluster'. An inset window titled 'New Cluster - Summary' provides details about the newly created cluster 'Jimcclus01', including its nodes 'As01' and 'As01c', and options for switchable hardware and software. Navigation buttons for 'Back', 'Next', and 'Finish' are visible at the bottom of both windows.

Creating Cluster

The cluster is getting created.

Creating cluster

New Cluster - Summary

Cluster Jimcclus01 has been created with the following settings. Click Finish.

Cluster: Jimcclus01

Nodes: As01
As01c

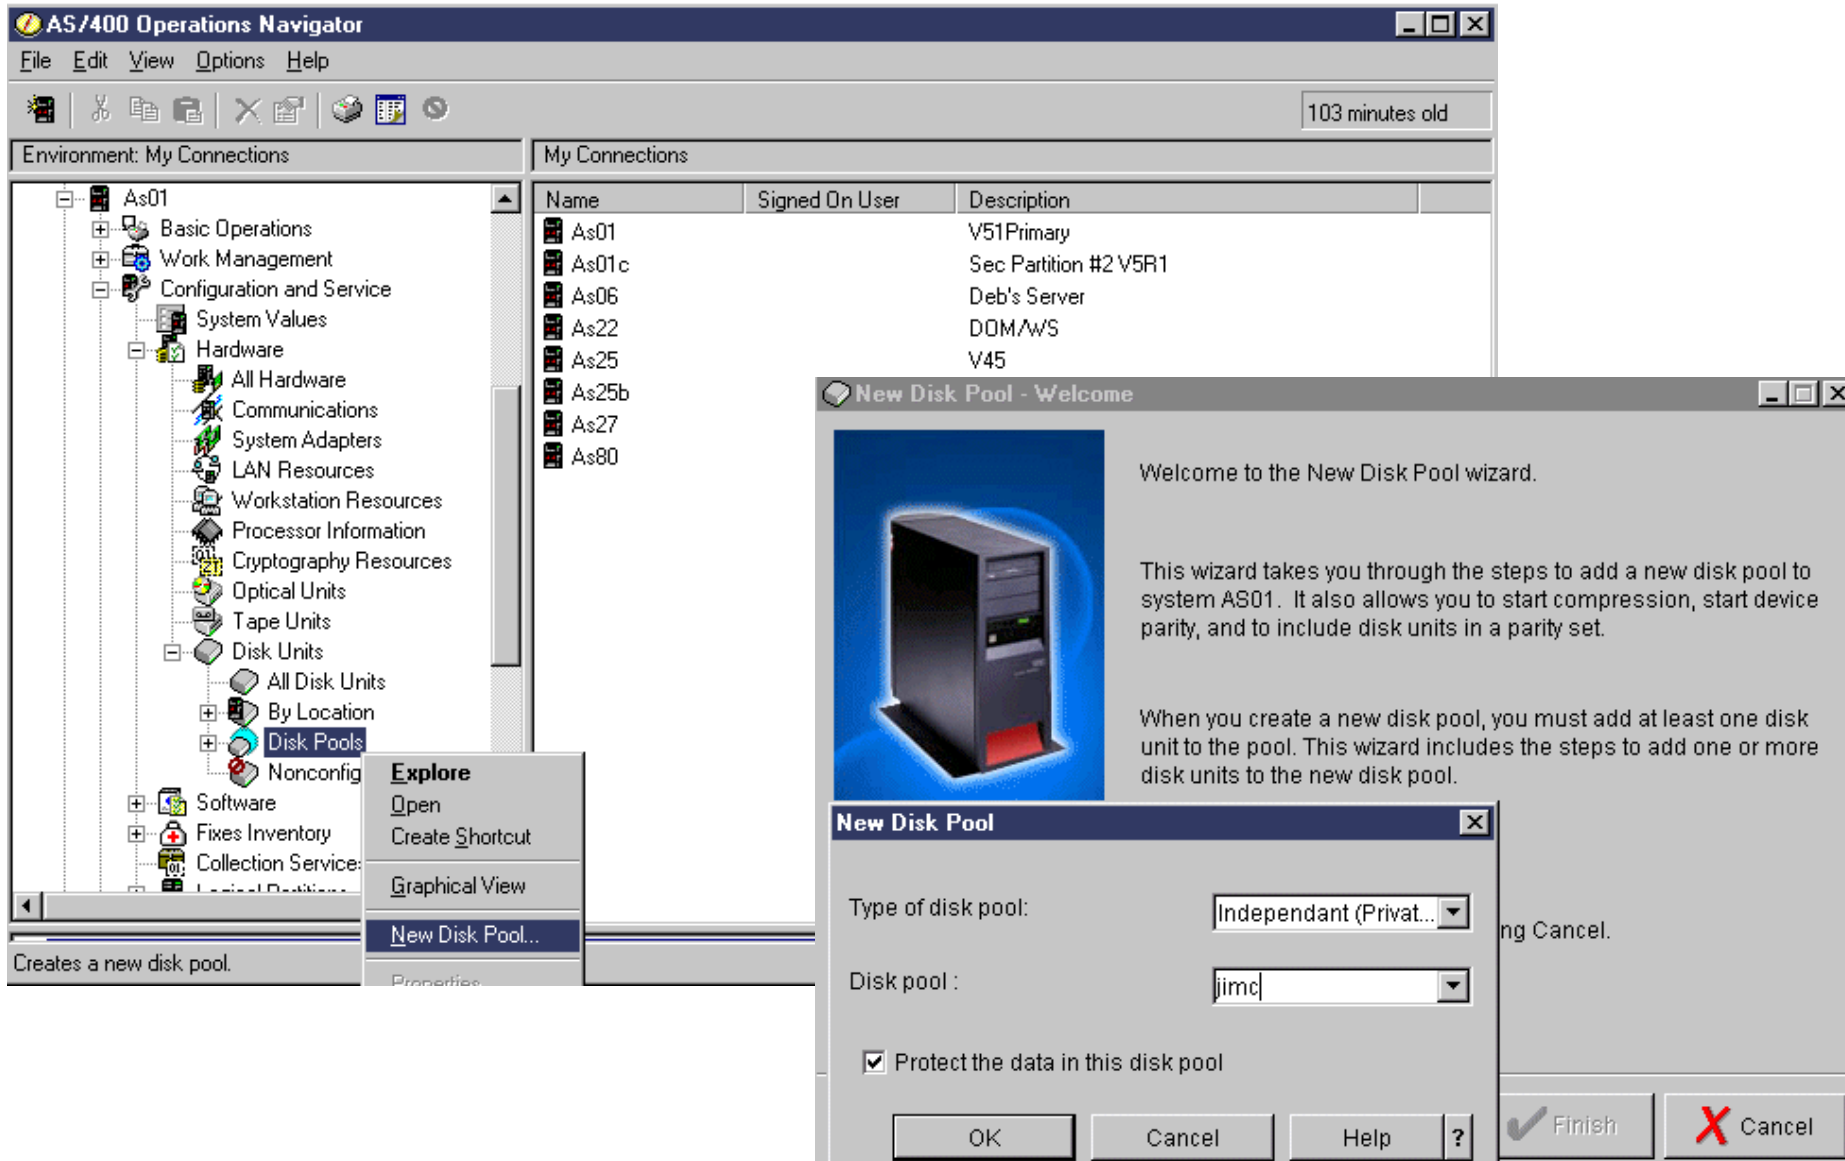
Switchable Hardware:

Switchable Software:

Back Next Finish

Back Next Finish Cancel Help ?

Configure an IASP



The screenshot displays the AS/400 Operations Navigator interface. The main window shows a tree view on the left with 'Disk Pools' selected. A 'My Connections' table is visible in the background, listing various system components. Overlaid on the interface are two windows: 'New Disk Pool - Welcome' and 'New Disk Pool'.

Name	Signed On User	Description
As01		V51Primary
As01c		Sec Partition #2 V5R1
As06		Deb's Server
As22		DOM/WS
As25		V45
As25b		
As27		
As80		

New Disk Pool - Welcome

Welcome to the New Disk Pool wizard.

This wizard takes you through the steps to add a new disk pool to system AS01. It also allows you to start compression, start device parity, and to include disk units in a parity set.

When you create a new disk pool, you must add at least one disk unit to the pool. This wizard includes the steps to add one or more disk units to the new disk pool.

New Disk Pool

Type of disk pool:

Disk pool:

Protect the data in this disk pool

Buttons: OK, Cancel, Help, ?

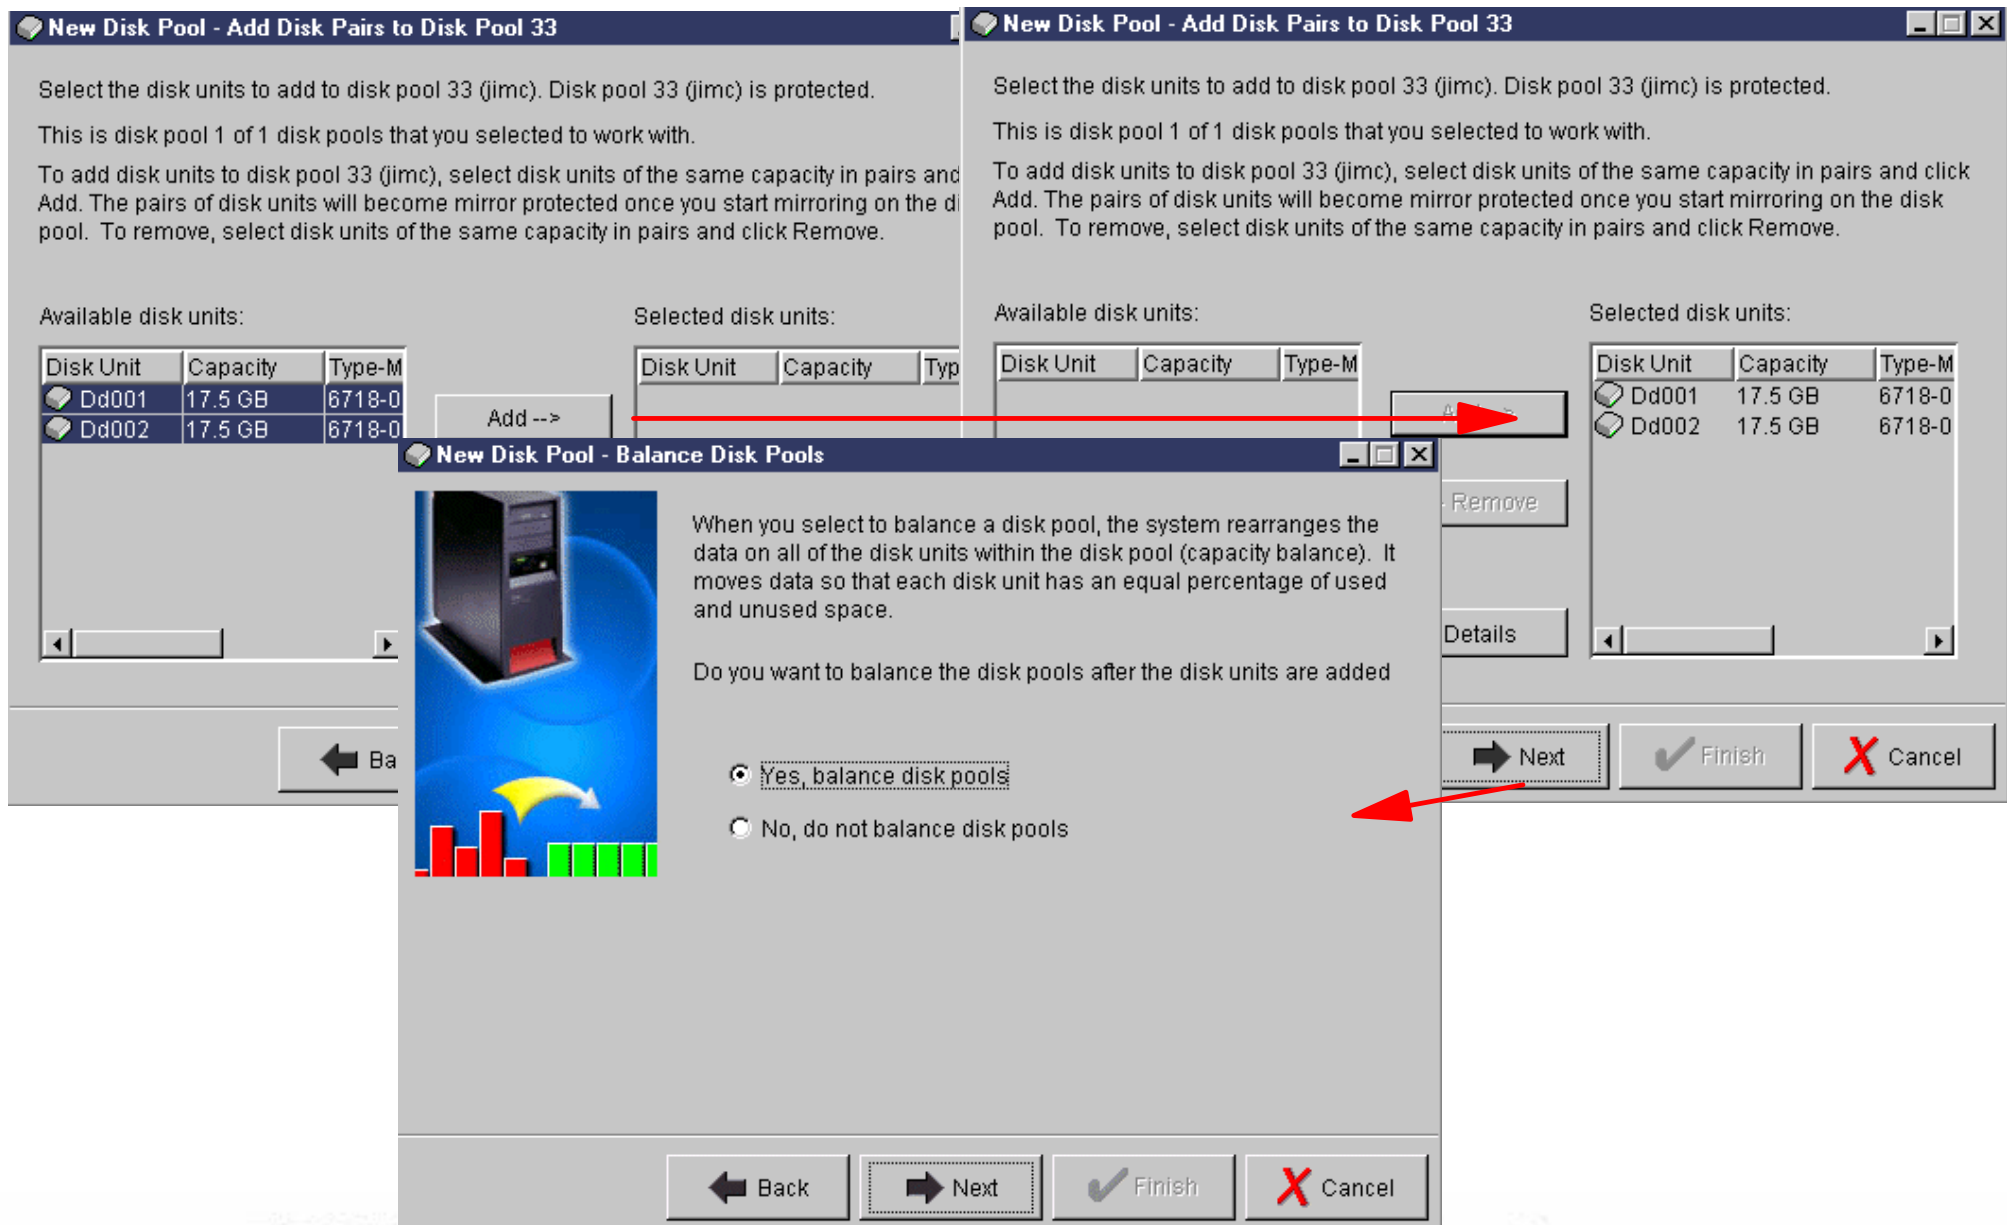
Bottom right buttons: Finish, Cancel

You can use OS/400 clustering commands to create cluster objects and the new Independent ASP objects and associate them as required. This is tedious and subject to errors, so we recommend using the V5R1 Operations Navigator interfaces as this support does "some steps and in the proper sequence " for you to minimize improper configurations.

Before you create an IASP, you need to have created the device description and added it to the CRG, otherwise OS/400 may not be able to determine which disk units are eligible to be added to the specific IASP, and the user may add invalid ones.

In this example we did part of the Cluster configuration defining the underlying resilient devices - Cluster Resource Group and then went to create the IASP (here) and then back to **Operations Navigator -> Clusters** path to complete the configuration.

New IASP - Adding Disks



New Disk Pool - Add Disk Pairs to Disk Pool 33

Select the disk units to add to disk pool 33 (j1mc). Disk pool 33 (j1mc) is protected.

This is disk pool 1 of 1 disk pools that you selected to work with.

To add disk units to disk pool 33 (j1mc), select disk units of the same capacity in pairs and click Add. The pairs of disk units will become mirror protected once you start mirroring on the disk pool. To remove, select disk units of the same capacity in pairs and click Remove.

Available disk units:

Disk Unit	Capacity	Type-M
Dd001	17.5 GB	6718-0
Dd002	17.5 GB	6718-0

Selected disk units:

Disk Unit	Capacity	Typ
-----------	----------	-----

New Disk Pool - Balance Disk Pools

When you select to balance a disk pool, the system rearranges the data on all of the disk units within the disk pool (capacity balance). It moves data so that each disk unit has an equal percentage of used and unused space.

Do you want to balance the disk pools after the disk units are added

Yes, balance disk pools

No, do not balance disk pools

New Disk Pool - Add Disk Pairs to Disk Pool 33

Select the disk units to add to disk pool 33 (j1mc). Disk pool 33 (j1mc) is protected.

This is disk pool 1 of 1 disk pools that you selected to work with.

To add disk units to disk pool 33 (j1mc), select disk units of the same capacity in pairs and click Add. The pairs of disk units will become mirror protected once you start mirroring on the disk pool. To remove, select disk units of the same capacity in pairs and click Remove.

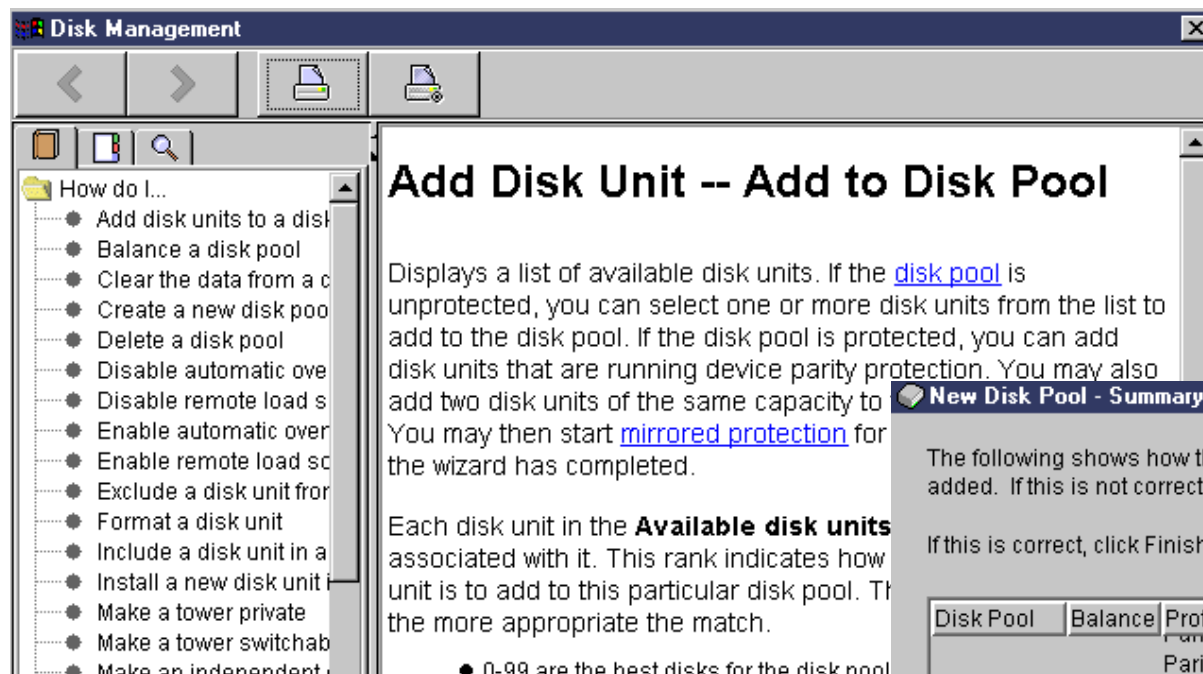
Available disk units:

Disk Unit	Capacity	Type-M
-----------	----------	--------

Selected disk units:

Disk Unit	Capacity	Type-M
Dd001	17.5 GB	6718-0
Dd002	17.5 GB	6718-0

New IASP - IASP Configured with New Disks

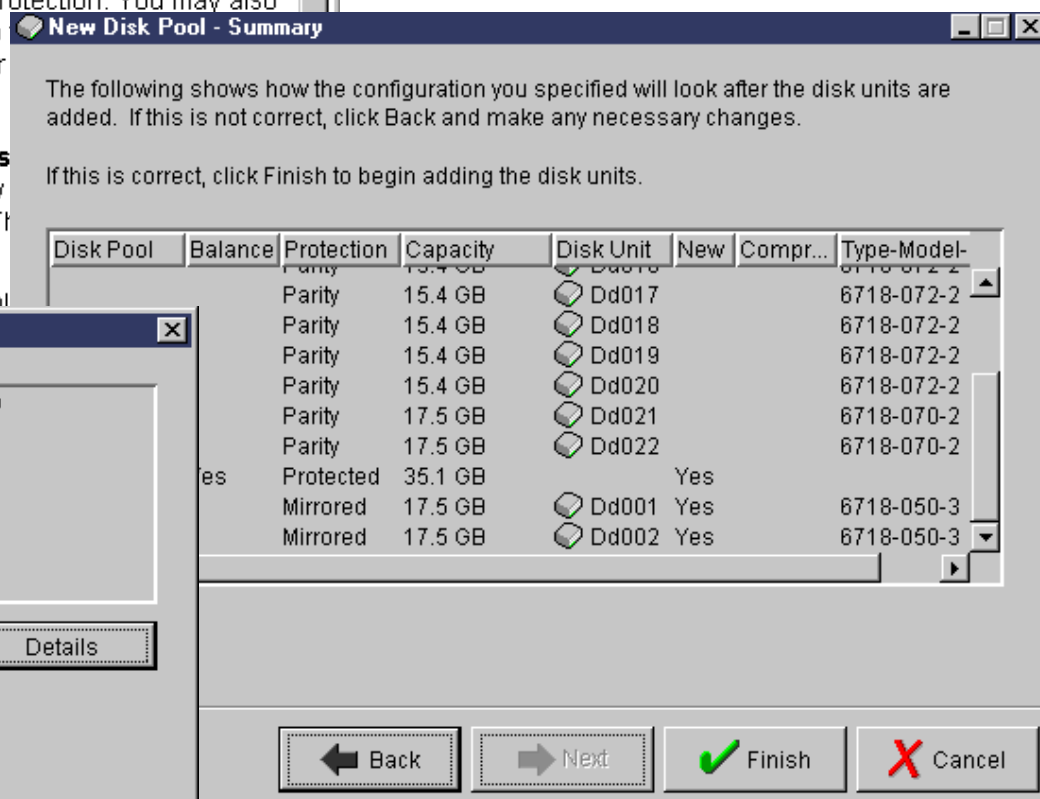


Add Disk Unit -- Add to Disk Pool

Displays a list of available disk units. If the [disk pool](#) is unprotected, you can select one or more disk units from the list to add to the disk pool. If the disk pool is protected, you can add disk units that are running device parity protection. You may also add two disk units of the same capacity to the disk pool. You may then start [mirrored protection](#) for the wizard has completed.

Each disk unit in the **Available disk units** associated with it. This rank indicates how unit is to add to this particular disk pool. The more appropriate the match.

- 0-99 are the best disks for the disk pool

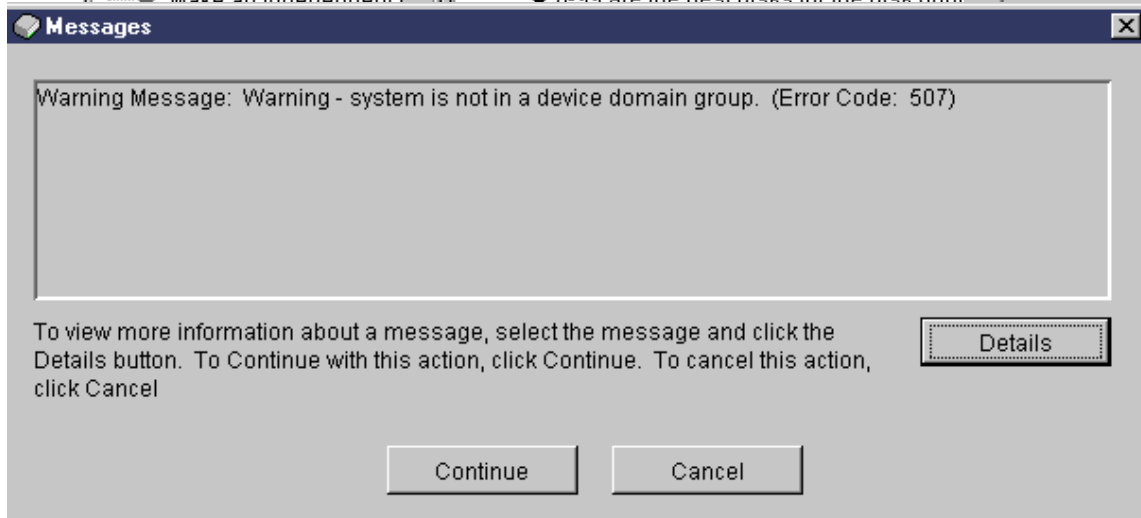


The following shows how the configuration you specified will look after the disk units are added. If this is not correct, click Back and make any necessary changes.

If this is correct, click Finish to begin adding the disk units.

Disk Pool	Balance	Protection	Capacity	Disk Unit	New	Compr...	Type-Model
	Parity	15.4 GB		Dd016			6718-072-2
	Parity	15.4 GB		Dd017			6718-072-2
	Parity	15.4 GB		Dd018			6718-072-2
	Parity	15.4 GB		Dd019			6718-072-2
	Parity	15.4 GB		Dd020			6718-072-2
	Parity	17.5 GB		Dd021			6718-070-2
	Parity	17.5 GB		Dd022			6718-070-2
es	Protected	35.1 GB			Yes		
	Mirrored	17.5 GB		Dd001	Yes		6718-050-3
	Mirrored	17.5 GB		Dd002	Yes		6718-050-3

Buttons: Back, Next, Finish, Cancel



Messages

Warning Message: Warning - system is not in a device domain group. (Error Code: 507)

To view more information about a message, select the message and click the Details button. To Continue with this action, click Continue. To cancel this action, click Cancel

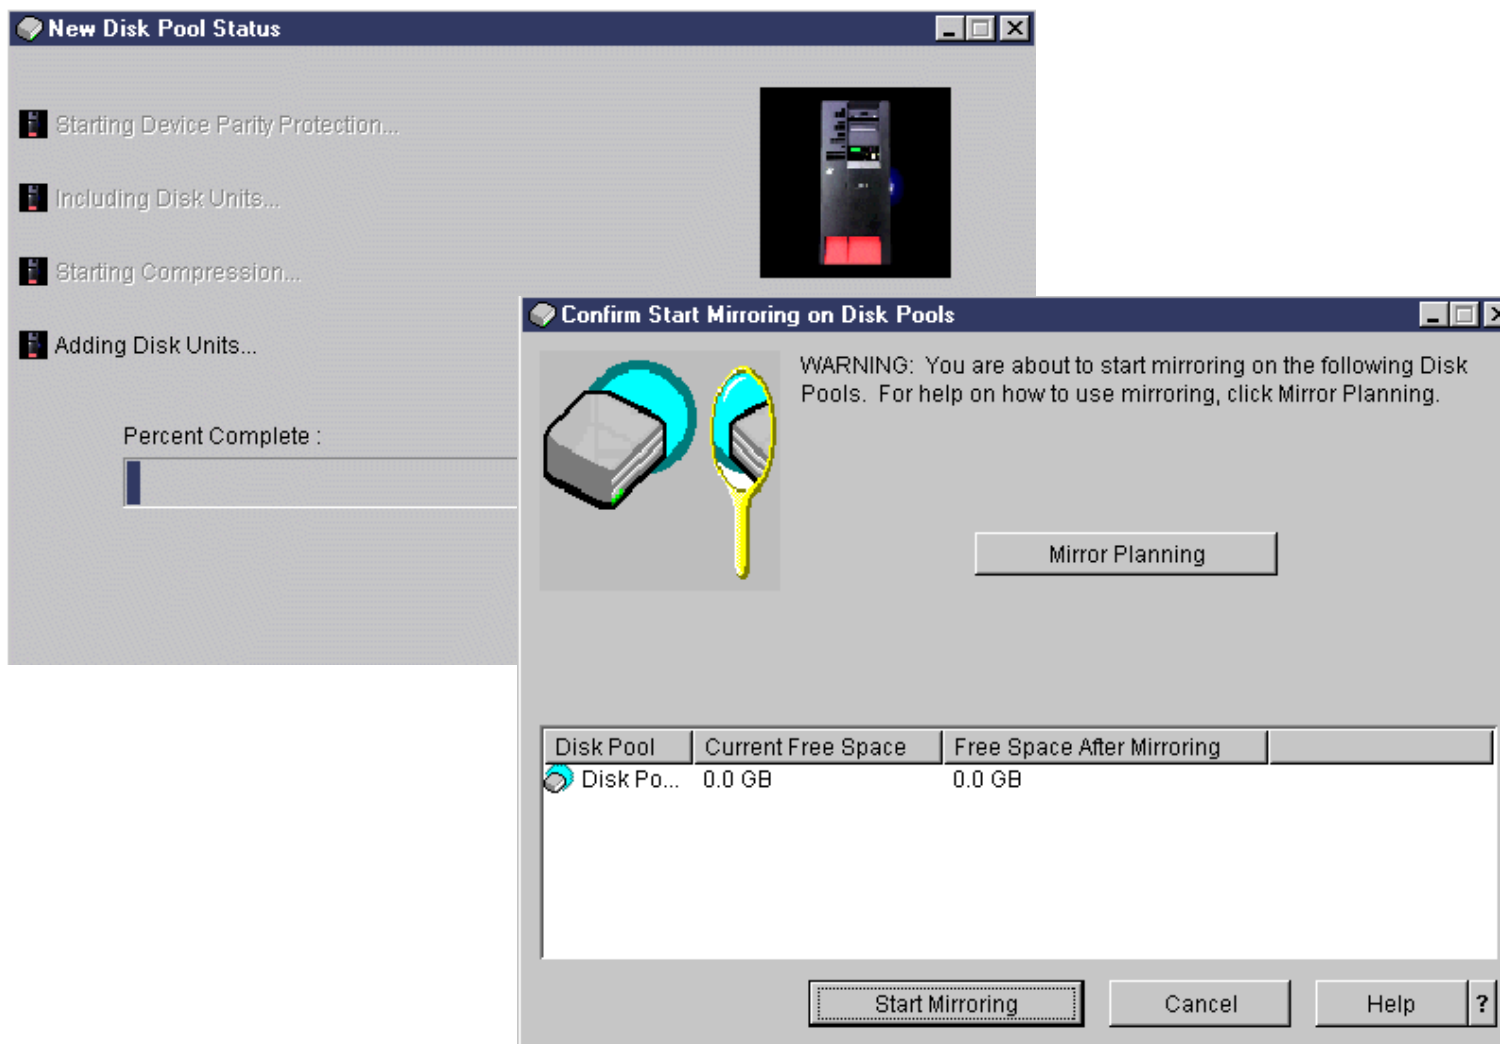
Buttons: Continue, Cancel, Details

This foil shows part of the process of adding disks to a new IASP.


You may get an error window as shown in the lower left window.

You may, in your configuration of clustering get a message similar to the one shown here. This message is issued when the system you are defining the ASP is not part of a cluster (device domain group). If you keep going there is a possibility that if the user decides later to make the system part of a cluster the, the user may not be able to do that without delete exist IASPs.

Completing the new IASP

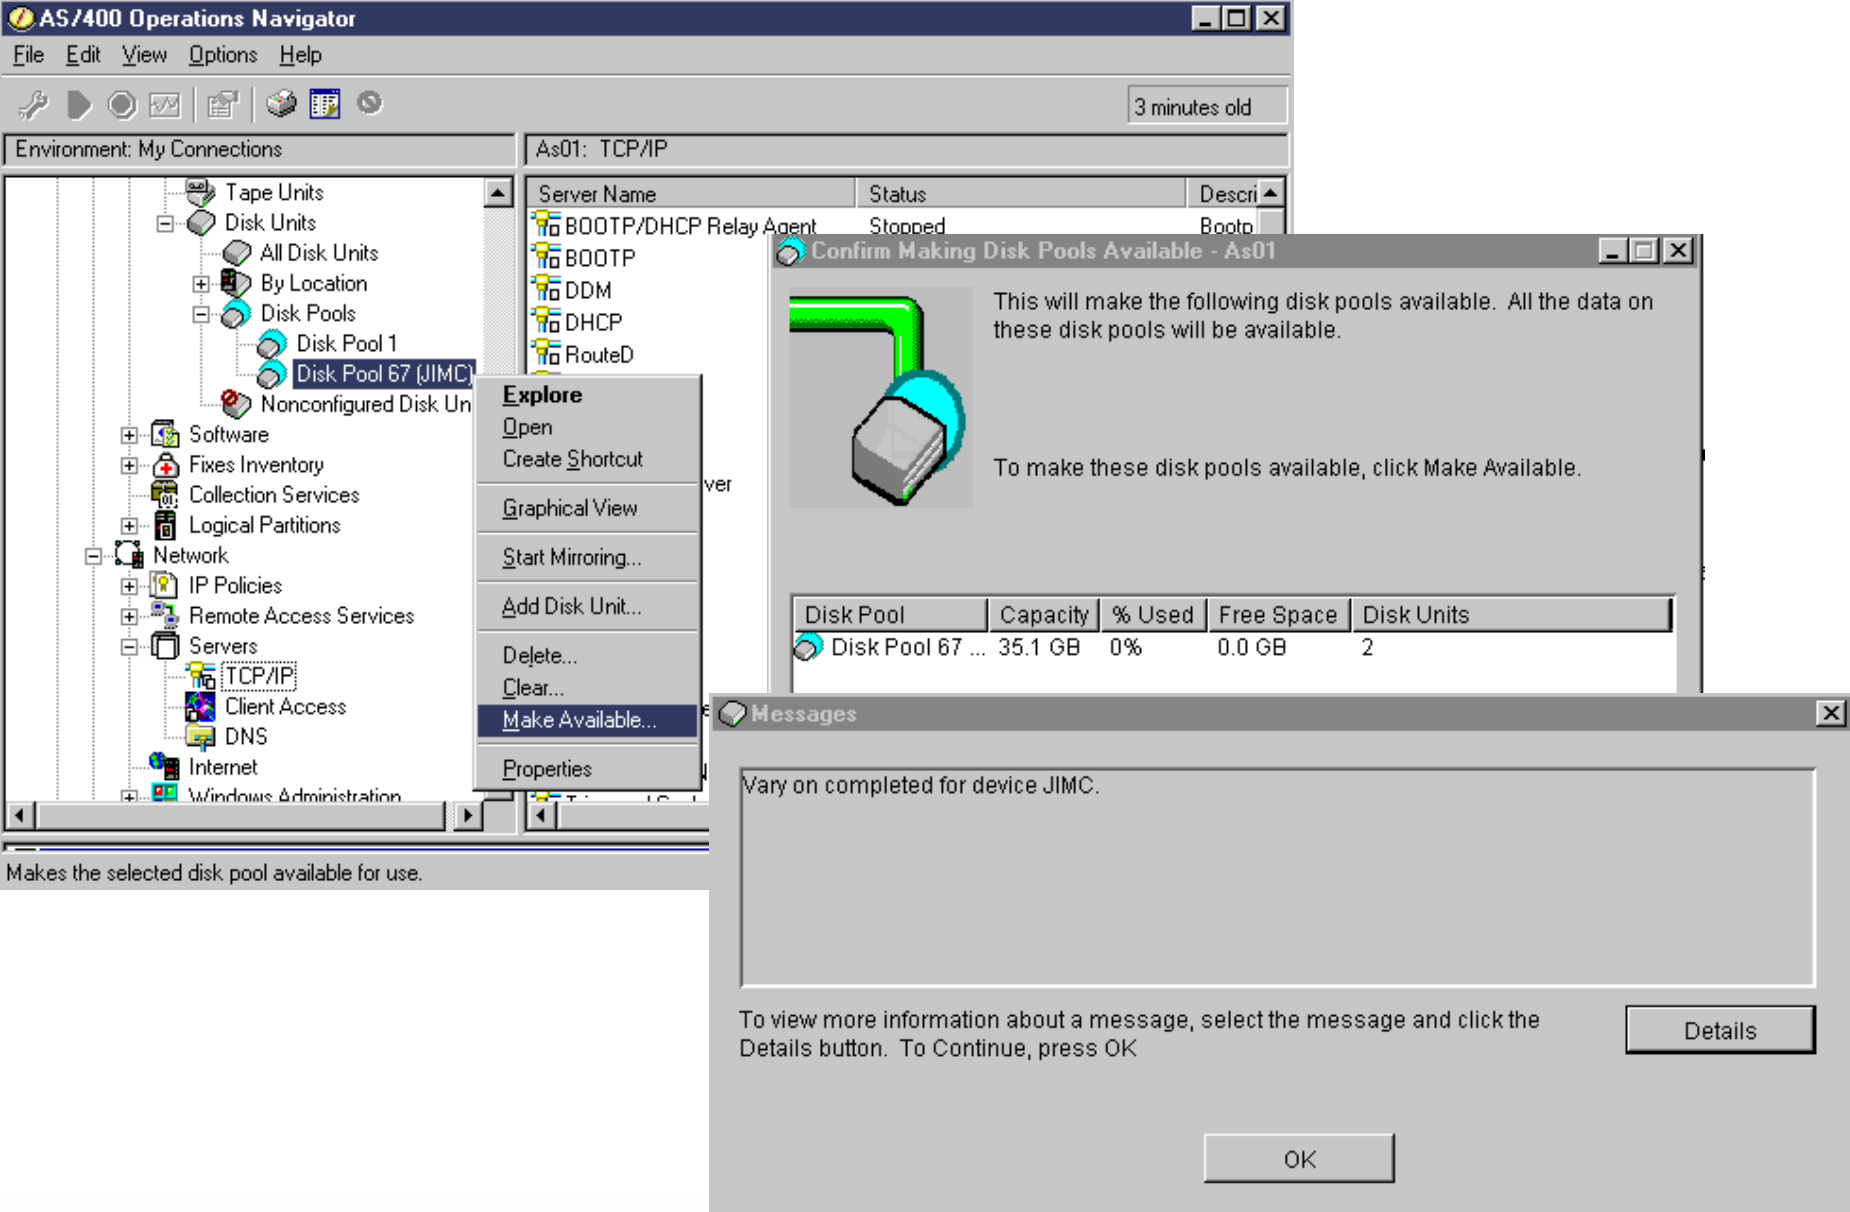


The screenshot shows two overlapping windows from the IASP (IBM Advanced System Programming) interface. The background window, titled "New Disk Pool Status", displays a progress list on the left: "Starting Device Parity Protection...", "Including Disk Units...", "Starting Compression...", and "Adding Disk Units...". Below this list is a "Percent Complete" progress bar. To the right of the list is a small image of a server rack. The foreground window, titled "Confirm Start Mirroring on Disk Pools", contains a warning message: "WARNING: You are about to start mirroring on the following Disk Pools. For help on how to use mirroring, click Mirror Planning." Below the warning is a "Mirror Planning" button. At the bottom of this window is a table with the following data:

Disk Pool	Current Free Space	Free Space After Mirroring
 Disk Po...	0.0 GB	0.0 GB

At the bottom of the foreground window are three buttons: "Start Mirroring", "Cancel", and "Help".

Making the Pool Available (Vary On) for use



AS/400 Operations Navigator

Environment: My Connections As01: TCP/IP

3 minutes old

Server Name Status Descri

BOOTP/DHCP Relay Agent	Stopped	Boo
BOOTP		
DDM		
DHCP		
Routed		

Confirm Making Disk Pools Available - As01

This will make the following disk pools available. All the data on these disk pools will be available.

To make these disk pools available, click Make Available.

Disk Pool	Capacity	% Used	Free Space	Disk Units
Disk Pool 67 ...	35.1 GB	0%	0.0 GB	2

Messages

Vary on completed for device JIMC.

To view more information about a message, select the message and click the Details button. To Continue, press OK

Details

OK

Makes the selected disk pool available for use.

Notes: Making the Pool Available for use

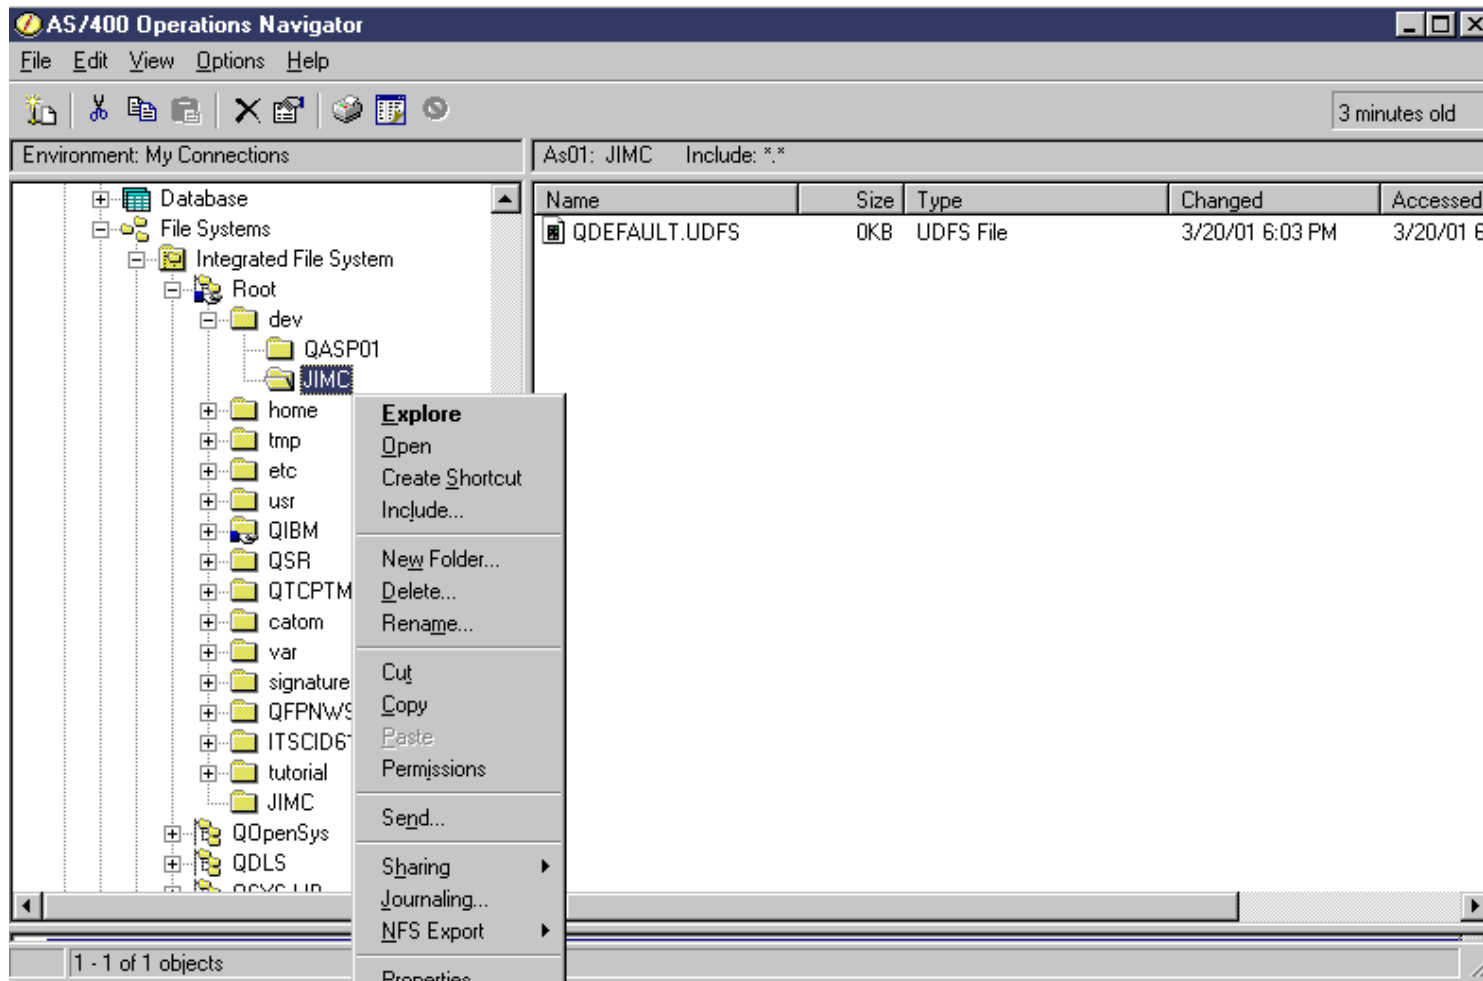
IBM  server iSeries

Before the resources within a new IASP (Disk Pool) can be used by the system (node) the IASP must be varied on. Operations Navigator uses the term "Make Available" for "vary on" and "Make Unavailable" for "vary off."

IBM  server. For the next generation of e-business.

Mount User Defined File System

IBM  server iSeries



The screenshot shows the AS/400 Operations Navigator interface. The left pane displays a file system tree under 'Environment: My Connections'. The tree structure is as follows:

- Database
- File Systems
 - Integrated File System
 - Root
 - dev
 - QASP01
 - JIMC
 - home
 - tmp
 - etc
 - usr
 - QIBM
 - QSR
 - QTCPTM
 - catom
 - var
 - signature
 - QFPNWS
 - ITSCID6
 - tutorial
 - JIMC
 - QOpenSys
 - QDLS
 - QCVCLP

The right pane shows a table of objects for 'As01: JIMC Include: *.*'. The table has the following data:

Name	Size	Type	Changed	Accessed
QDEFAULT.UDFS	0KB	UDFS File	3/20/01 6:03 PM	3/20/01 6:

A context menu is open over the 'JIMC' folder in the tree, listing options such as Explore, Open, Create Shortcut, Include..., New Folder..., Delete..., Rename..., Cut, Copy, Paste, Permissions, Send..., Sharing, Journaling..., NFS Export, and Properties.

IBM  server. For the next generation of e-business.

Notes: Mount User Defined File System

When you create the IASP device the system automatically generates a User Defined File System (UDFS) named the same as the device (IASP) name. In the examples used here that name is JIMC.

You must place files meaningful to your application environment into this default UDFS. While the IASP is "varied off/unavailable" any files in the UDFS cannot be accessed. In fact, through the Operations Navigator-Integrated File System neither the root/dev/JIMC ASP name nor the specific UDFS name JIMC can be seen

When that IASP is "made available/varied on" the default file system is automatically mounted in the root directory of the system's IFS. The root/dev/JIMC ASP and UDFS JIMC can now be seen through Operations Navigator-IFS view

The UDFS is accessible only when the IASP device is varied on and then only when the device CRG ("Switch" in the examples that follow) is active and "primary" within the cluster. When the device DCR is varied off or switched to a secondary node, the UDFS is "unmounted" and the objects within that file system are no longer accessible under IFS on original node.

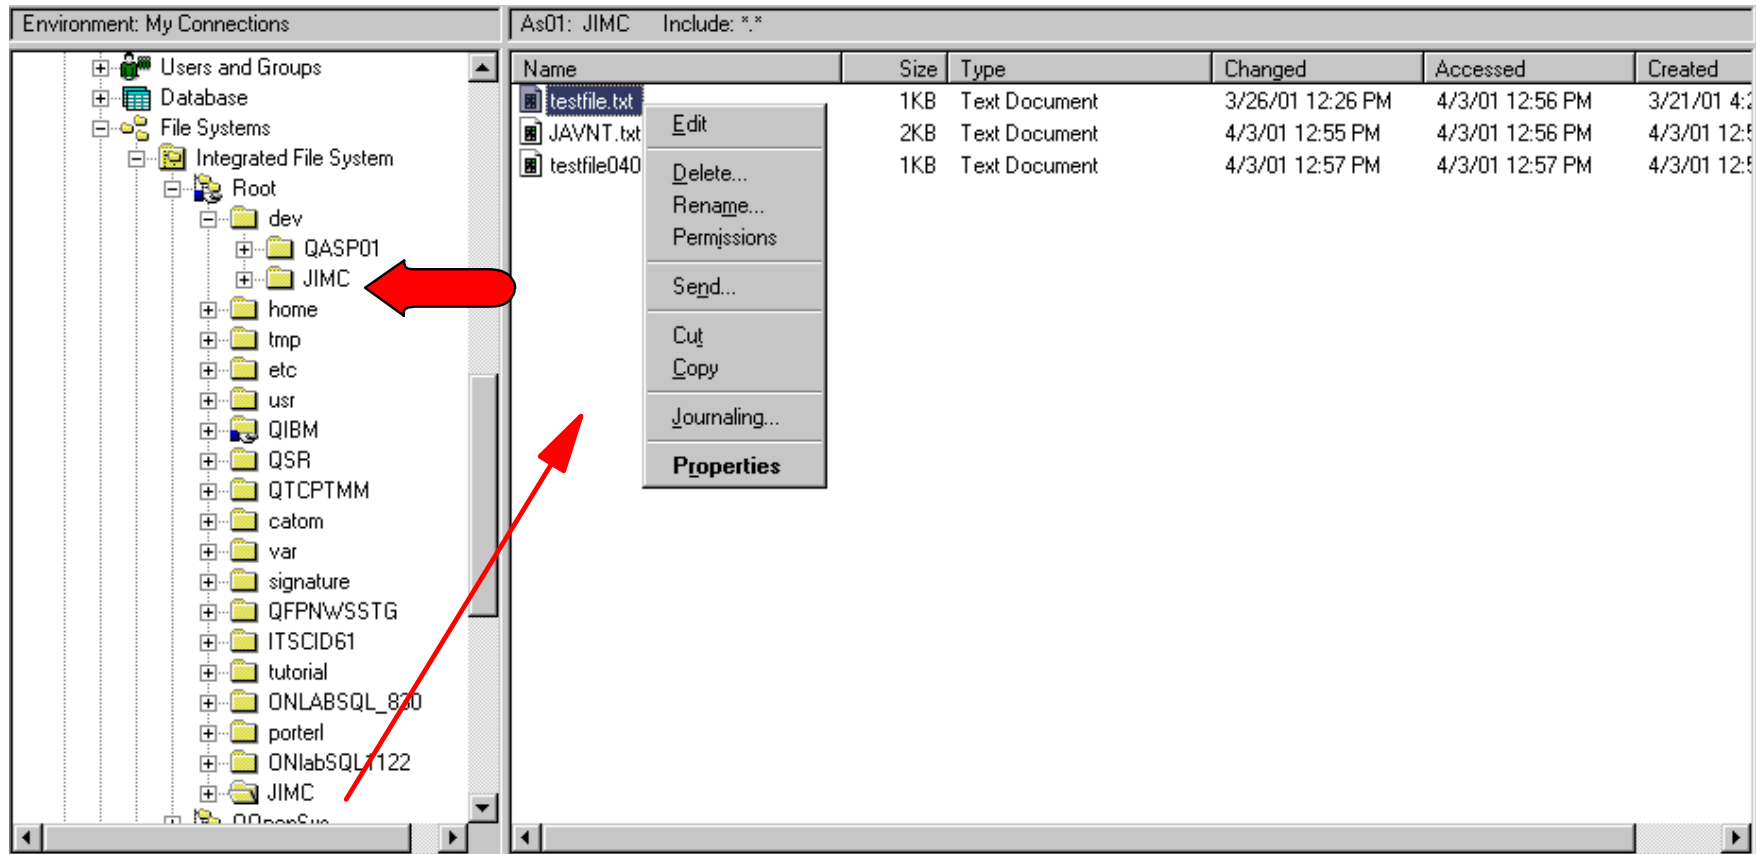
IFS file objects in the UDFS are accessible via the IFS file interfaces using a path name that includes the UDFS name.

In the Domino ClusterProven™ discussion later in this presentation, where the Domino Server's data directory is placed in an IASP, an example path for the data directory could be:

- /JIMC/DOMINO/domino-server-name

Showing Contents of Mounted UDFS

IASP Device and associated UDFS files



Environment: My Connections

As01: JIMC Include: *.*

Name	Size	Type	Changed	Accessed	Created
testfile.txt	1KB	Text Document	3/26/01 12:26 PM	4/3/01 12:56 PM	3/21/01 4:2
JAVNT.txt	2KB	Text Document	4/3/01 12:55 PM	4/3/01 12:56 PM	4/3/01 12:5
testfile040	1KB	Text Document	4/3/01 12:57 PM	4/3/01 12:57 PM	4/3/01 12:5

File System Tree:

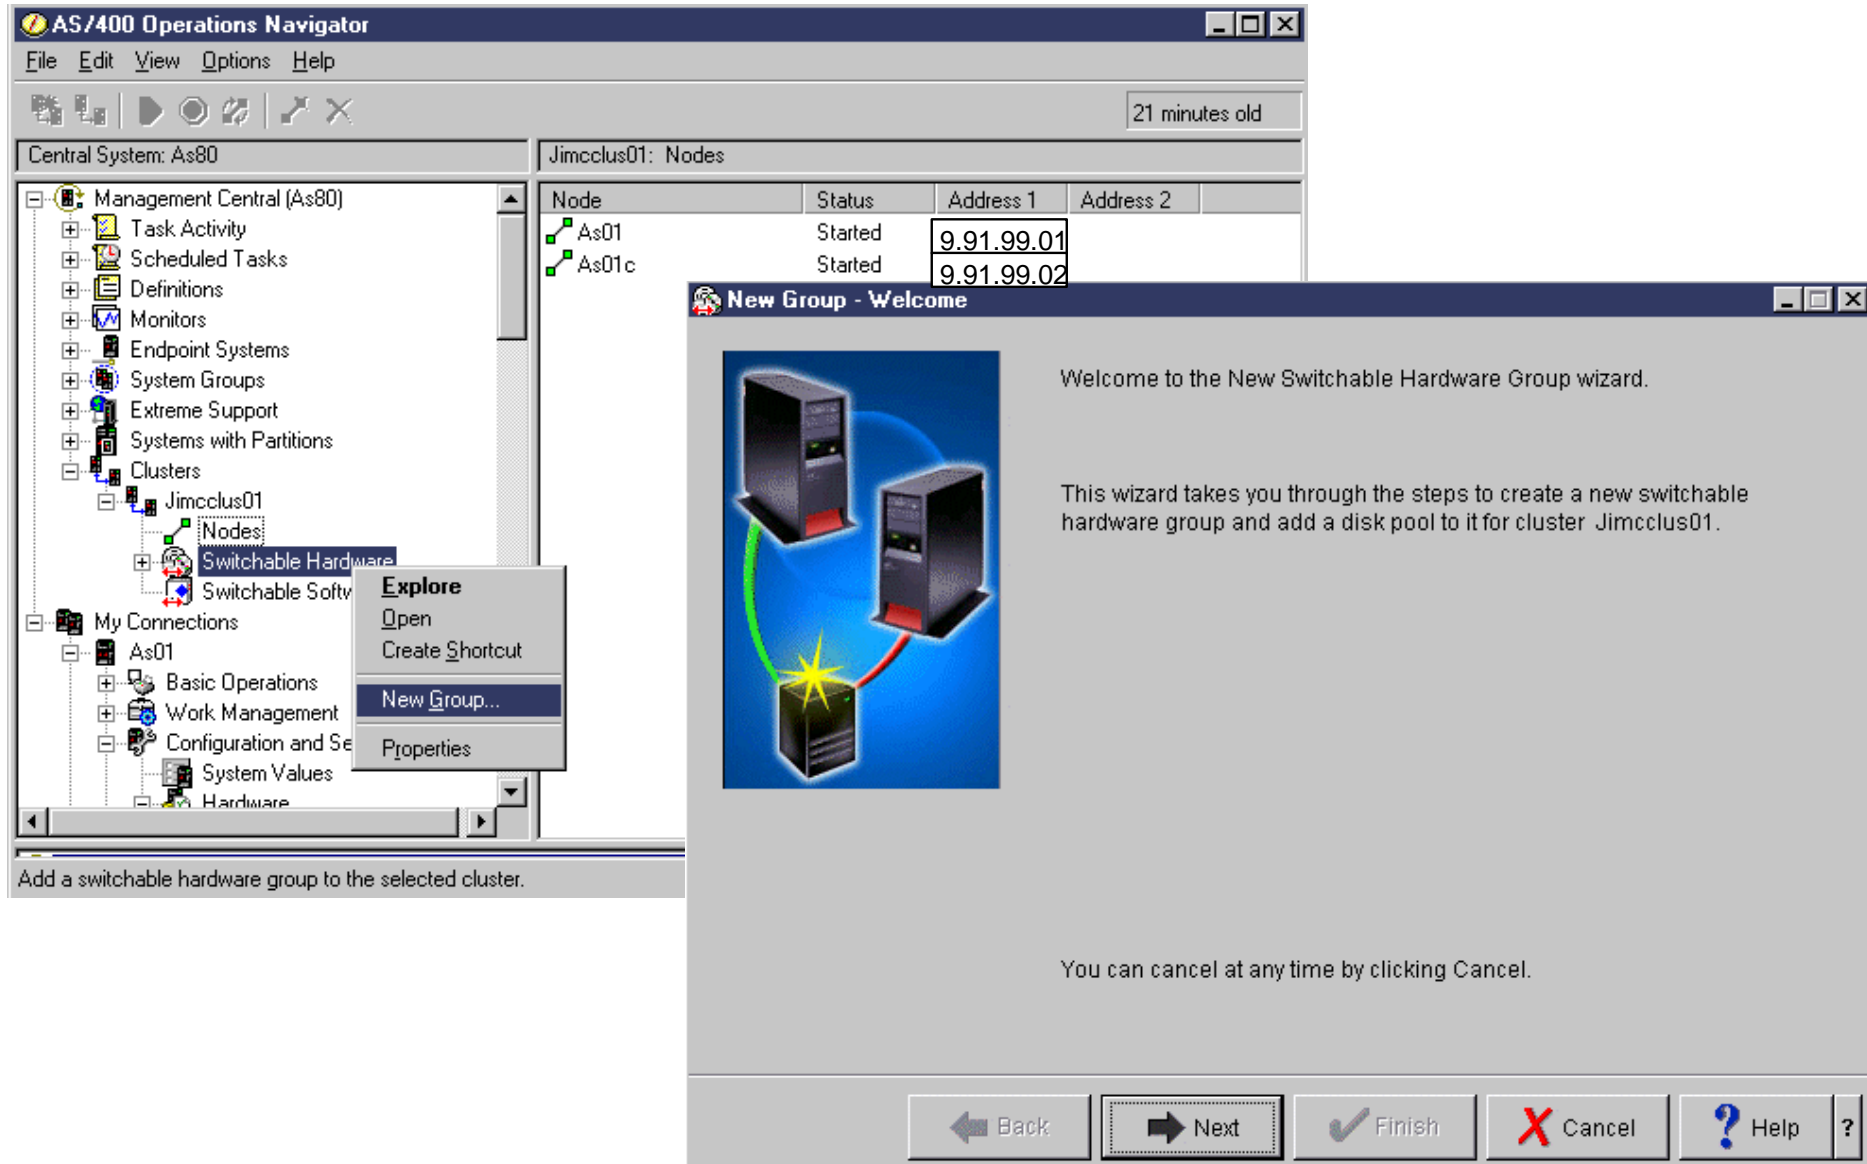
- Users and Groups
- Database
- File Systems
 - Integrated File System
 - Root
 - dev
 - QASP01
 - JIMC
 - home
 - tmp
 - etc
 - usr
 - QIBM
 - QSR
 - QTCPTMM
 - catom
 - var
 - signature
 - QFPNWSSTG
 - ITSCID61
 - tutorial
 - ONLABSQL_800
 - portel
 - ONlabSQL_1122
 - JIMC
 - QDexSus

Notes: Showing Contents of Mounted UDFS

This foils shows the IASP device name JIMC and the associated default UDFS JIMC under AS01 system's Operations Navigator-Integrated File System directory structure - after "Make Available" completed successfully.

Remember that after the Device Cluster Resource Group that contains this IASP is switched to a secondary system, you must "Make Available" that IASP after that switching function has completed. You will be able to see the disk devices and associated IASP on the secondary system once the switch has completed. But you will not be able to see the UDFS until it is mounted with "Make Available/Vary On."

Creating Resilient Device CRG



AS/400 Operations Navigator

File Edit View Options Help

21 minutes old

Central System: As80

Jimclus01: Nodes

Node	Status	Address 1	Address 2
As01	Started	9.91.99.01	
As01c	Started	9.91.99.02	

New Group - Welcome

Welcome to the New Switchable Hardware Group wizard.

This wizard takes you through the steps to create a new switchable hardware group and add a disk pool to it for cluster Jimclus01.

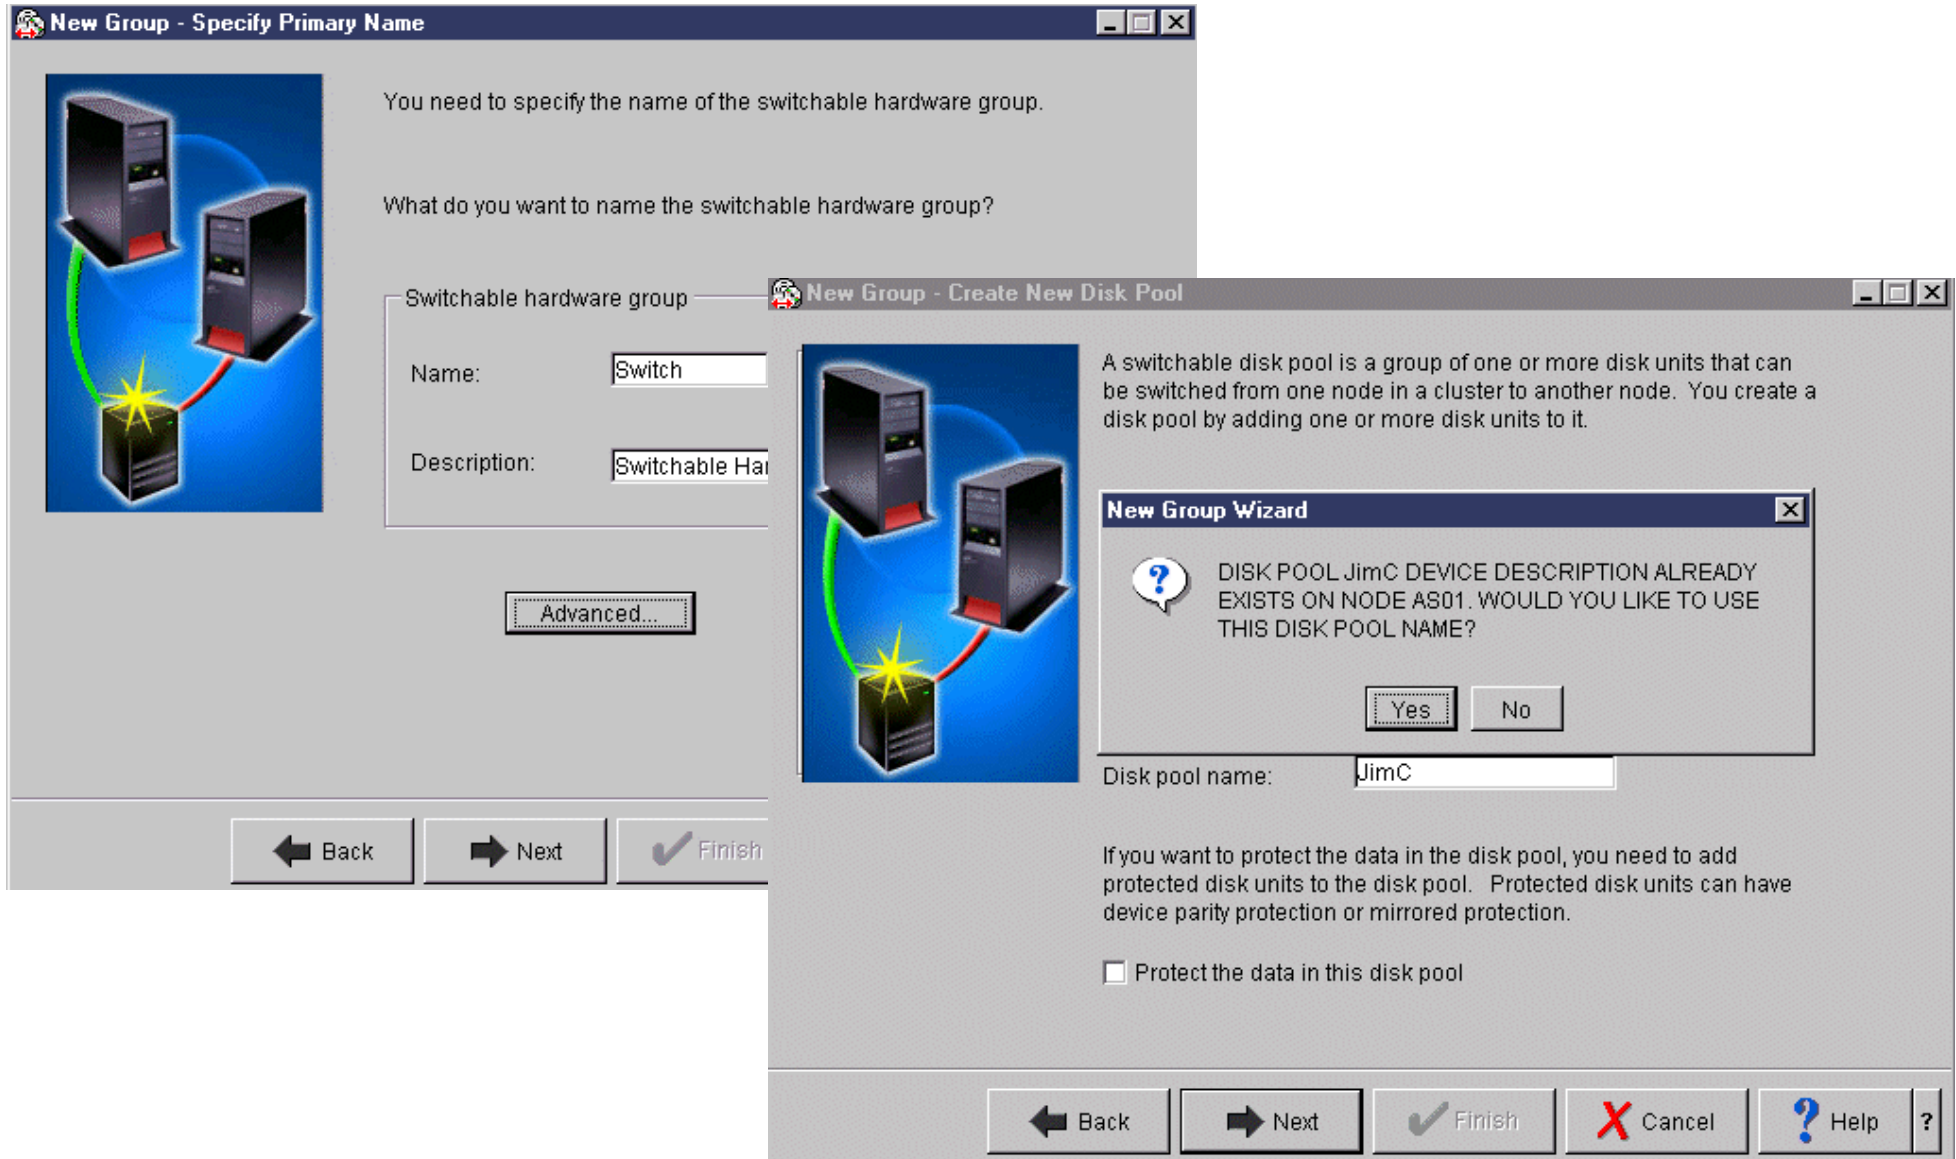
You can cancel at any time by clicking Cancel.

Back Next Finish Cancel Help ?

Management Central (As80)

- Task Activity
- Scheduled Tasks
- Definitions
- Monitors
- Endpoint Systems
- System Groups
- Extreme Support
- Systems with Partitions
- Clusters
 - Jimclus01
 - Nodes
 - Switchable Hardware
 - Switchable Softw
- My Connections
 - As01
 - Basic Operations
 - Work Management
 - Configuration and Se
 - System Values
 - Hardware

New RD Cluster Resource Group Name



The screenshot displays the 'New Group - Specify Primary Name' wizard window. It contains an illustration of a cluster with three nodes and a central disk unit. The text reads: 'You need to specify the name of the switchable hardware group. What do you want to name the switchable hardware group?'. Below this, there are input fields for 'Name' (containing 'Switch') and 'Description' (containing 'Switchable Ha'). An 'Advanced...' button is visible below the description field. At the bottom of this window are 'Back', 'Next', and 'Finish' buttons.

Overlaid on this is the 'New Group - Create New Disk Pool' wizard window. It features the same cluster illustration and explains: 'A switchable disk pool is a group of one or more disk units that can be switched from one node in a cluster to another node. You create a disk pool by adding one or more disk units to it.' Below this is a 'New Group Wizard' dialog box with a question mark icon and the text: 'DISK POOL JimC DEVICE DESCRIPTION ALREADY EXISTS ON NODE AS01. WOULD YOU LIKE TO USE THIS DISK POOL NAME?'. It has 'Yes' and 'No' buttons. Below the dialog is a 'Disk pool name:' field containing 'JimC'. Further down, it says: 'If you want to protect the data in the disk pool, you need to add protected disk units to the disk pool. Protected disk units can have device parity protection or mirrored protection.' and includes a checkbox labeled 'Protect the data in this disk pool' which is currently unchecked. At the bottom of this window are 'Back', 'Next', 'Finish', 'Cancel', and 'Help' buttons.

Cluster Group Start

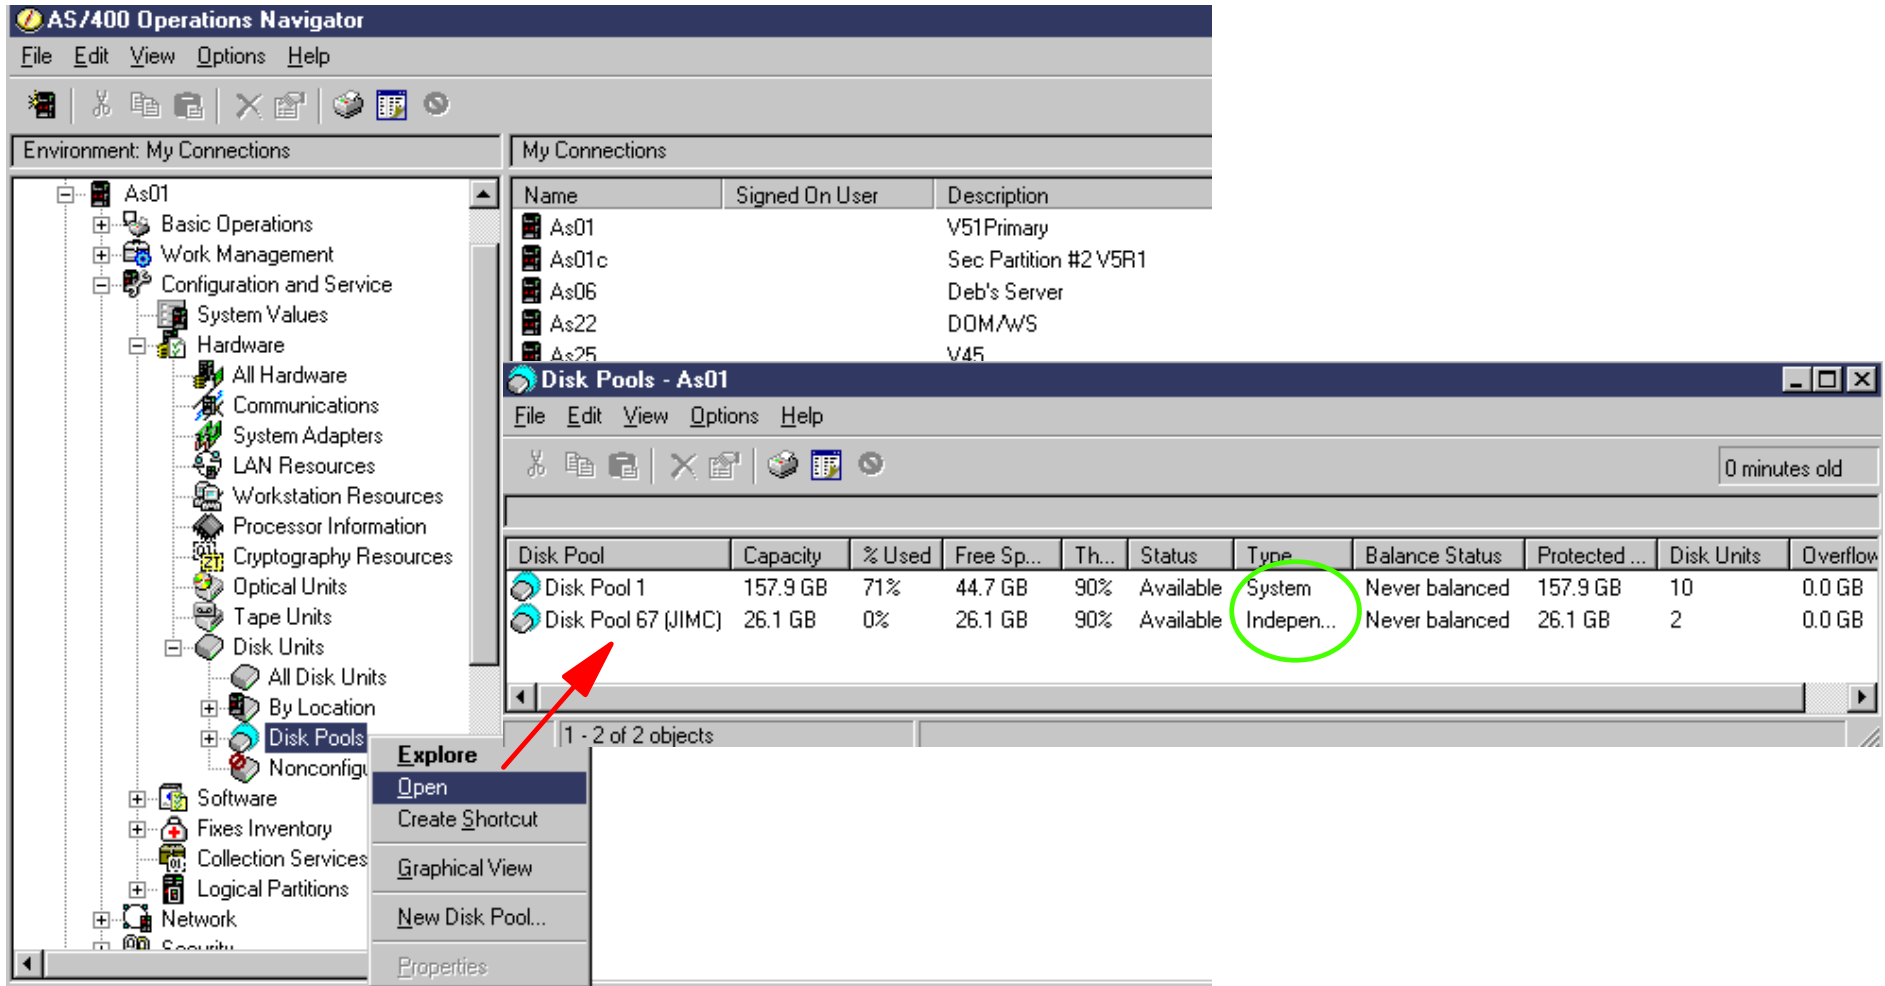
The image displays two screenshots of the AS/400 Operations Navigator interface, illustrating the process of starting a cluster group.

Top Screenshot: Shows the 'Jimclus01: Nodes' table with the following data:

Node	Status	Address 1	Address 2
As01	Stopped	9.91.99.01	
As01c	Stopped	9.91.99.02	

Bottom Screenshot: Shows the same interface with a context menu open over the 'As01' node. The menu options are: Cluster, Start, Stop..., Switch..., Properties, and Remove... The 'Start' option is highlighted.

View Independent ASP before switch



AS/400 Operations Navigator

Environment: My Connections

My Connections

Name	Signed On User	Description
As01		V51Primary
As01c		Sec Partition #2 V5R1
As06		Deb's Server
As22		DOM/WS
As25		V45

Disk Pools - As01

0 minutes old

Disk Pool	Capacity	% Used	Free Sp...	Th...	Status	Type	Balance Status	Protected ...	Disk Units	Overflow
Disk Pool 1	157.9 GB	71%	44.7 GB	90%	Available	System	Never balanced	157.9 GB	10	0.0 GB
Disk Pool 67 (JIMC)	26.1 GB	0%	26.1 GB	90%	Available	Indepen...	Never balanced	26.1 GB	2	0.0 GB

1 - 2 of 2 objects

- Explore
- Open
- Create Shortcut
- Graphical View
- New Disk Pool...
- Properties

Switch to secondary node example

The screenshot shows the AS/400 Operations Navigator interface. On the left, a tree view shows the hierarchy: Management Central (As80) > Clusters > Jimclus01 > Nodes. A context menu is open over the 'Nodes' folder, with 'Switch...' selected. The main pane displays a table of nodes:

Node	Status	Address 1	Address 2
As01	Started	9.91.99.01	
As01c	Started	9.91.99.02	

A red arrow labeled "AS01 to AS01C" points from the 'Nodes' folder in the tree to the 'Switch...' menu option.

Below the main interface, a 'Work with Disk Status' window is open, showing a table of disk units. The table has columns: Unit, Type, Size (M), % Used, I/O Rqs, Request Size (K), Read Rqs, Write Rqs, Read (K), Write (K), and % Busy. The row for unit 4001 is highlighted with a red box.

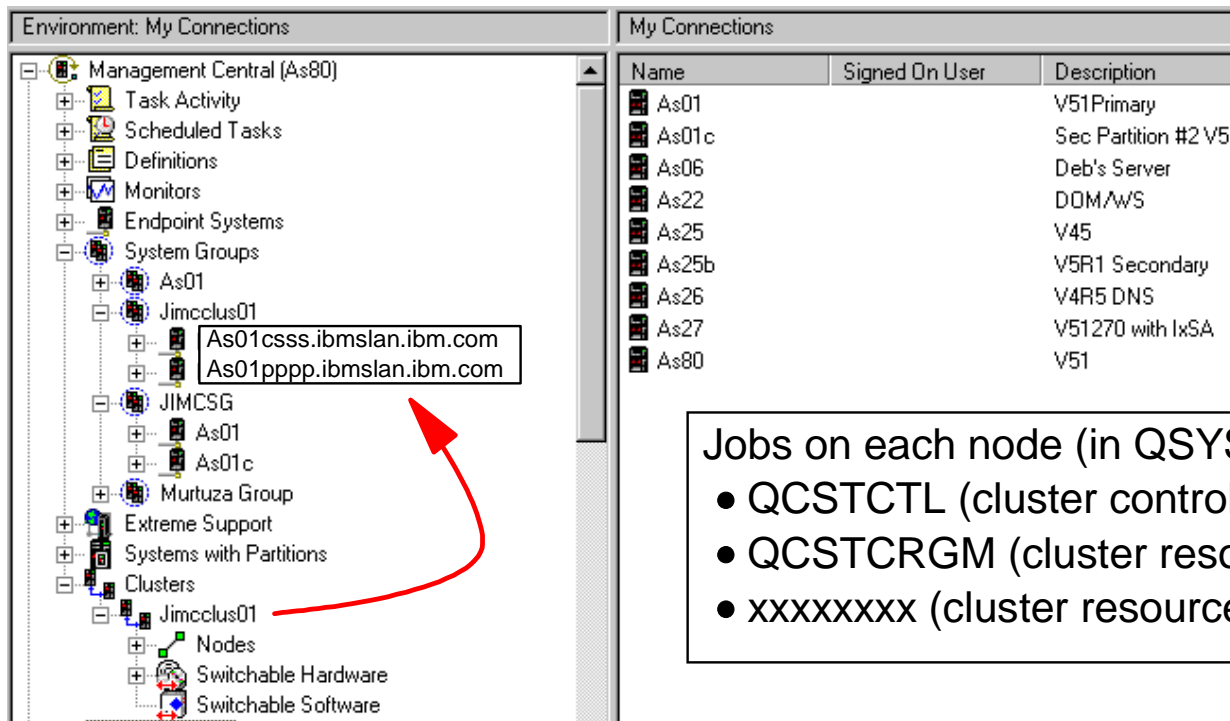
Unit	Type	Size (M)	% Used	I/O Rqs	Request Size (K)	Read Rqs	Write Rqs	Read (K)	Write (K)	% Busy
14	6607	4194	11.5	.0	.0	.0	.0	.0	.0	0
15	6713	7516	11.5	.0	.0	.0	.0	.0	.0	0
16	6713	7516	11.6	.0	.0	.0	.0	.0	.0	0
4001	6718	17548	.0	.0	.0	.0	.0	.0	.0	0
4001	6718	17548	.0	.0	.0	.0	.0	.0	.0	0

Notes: Switch to secondary node example

After the switch has completed, we used the 5250 command WRKDSKSTS to determine that the IASP disks have been successfully switched. Remember to do "Make Available/Vary On" on AS01C to actually access the data within the IASP.

Central Server, Cluster Node View

- Cluster known and managed from Management Central central server system
- System group created by clustering support
- CHGNETA ALWADDCLU(*ANY | RQSAUT)
- INETD servers started



Name	Signed On User	Description
As01		V51Primary
As01c		Sec Partition #2 V51
As06		Deb's Server
As22		DOMAWS
As25		V45
As25b		V5R1 Secondary
As26		V4R5 DNS
As27		V51270 with IxSA
As80		V51

- Jobs on each node (in QSYSWRK)
- QCSTCTL (cluster control manager job)
 - QCSTCRGM (cluster resource group manager job)
 - xxxxxxxx (cluster resource group job)

Notes: Central Server, Cluster Node View

As you may have noted, configuring a cluster, resilient device CRG, and a switch are done under the central server. When configuring the cluster under Operations Navigator a system group is created that is "tightly coupled" with the cluster. In this example the cluster and system group have then same name - Jimclus01.

The central server is used to activate the cluster objects, but the IASP (disk pool) is managed from the specific system (for example, AS01 under My Connections).

Although the central system can be one of the cluster nodes it is not a requirement. In the example, we used, AS80 is the central server and the cluster nodes are AS01, and AS01C. If the cluster is deleted, all cluster name-specific branches" under the central server are removed.

Throughout the cluster configuration and switching functions, both nodes must be active and that requires the Network Attribute to "Allow add to cluster" (ALWADDCLU) be set to either *ANY or *RQSAUT.

ALWASSCLU specifies whether this system will allow another system to add it as a node in a cluster. *RQSAUT specifies that any other system can add this system as a node in a cluster only after the cluster add request has been authenticated.

Also the TCP/IP INETD server on each node must be active.

On each node in a cluster (which may or may not include the Management Central central server system) the following cluster-related jobs run by default in IBM-supplied subsystem QSYSWRK:

- QCSTCTL (cluster control manager job) processes cluster control messages between the nodes
- QCSTCRGM (cluster resource group manager job) processes cluster resource group messages between the nodes
- xxxxxxxx (cluster resource group job) processes the specific cluster resource group. The job is named the same as the cluster resource group name. In the example Operations Navigator screens that follow later in this presentation the name we used is "Switch."

Availability Software

IBM @server. For the next generation of e-business.

Application Environments

IBM  server iSeries

Domino for AS/400

HTTP Server

Windows 2000 on Integrated xSeries Server

High Availability Business Partner solutions

IBM  server. For the next generation of e-business.

- ▶ Adds changes to Domino for iSeries (in release 5.0.7) that allows Domino servers to be managed through iSeries Cluster Management APIs as a ClusterProven™ application
- ▶ A ClusterProven application can be restarted on multiple nodes (systems) in an iSeries cluster using Cluster Management APIs or automatically by Cluster Management due to system failover
- ▶ ClusterProven application resources are automatically configured on cluster nodes when it is configured to the cluster
- ▶ iSeries only (requires OS/400 V5R1)

<http://www.ibm.com/eserver/series/domino>

Notes: Making Domino for iSeries ClusterProven™

IBM  server iSeries

Domino for AS/400, release 5.0.7 has been certified by the Rochester laboratory as ClusterProven™. This means the Domino application, when configured according to the "white paper" located at the URL shown on this foil, can have its configuration and files switched over to a second system (cluster node) with minimal or no interruption in the running Domino environment.,

The next foils give an idea of the Domino ClusterProven™ considerations.

IBM  server. For the next generation of e-business.

Multiple Domino server partitions allowed on any one system

- Independently created and managed without impacting other partitions
- Some Domino servers provide specific services (Quickplace and Sametime)

Requires iSeries Cluster Management APIs to independently manage a specific Domino server

- Create, Start, End, Switch, Failover,...
- Does not affect other servers on same cluster node

Domino servers can be configured for a specific IP address

- Cluster Management provides automatic IP address takeover

Domino servers can support Domino applications

- Domino applications not managed directly by Cluster Management

- ▶ **Clustering for high availability**
 - Requires only one copy of server's Data Directory and its Domino databases
 - Reduced or eliminated Domino replication overhead
 - Savings on DASD usage

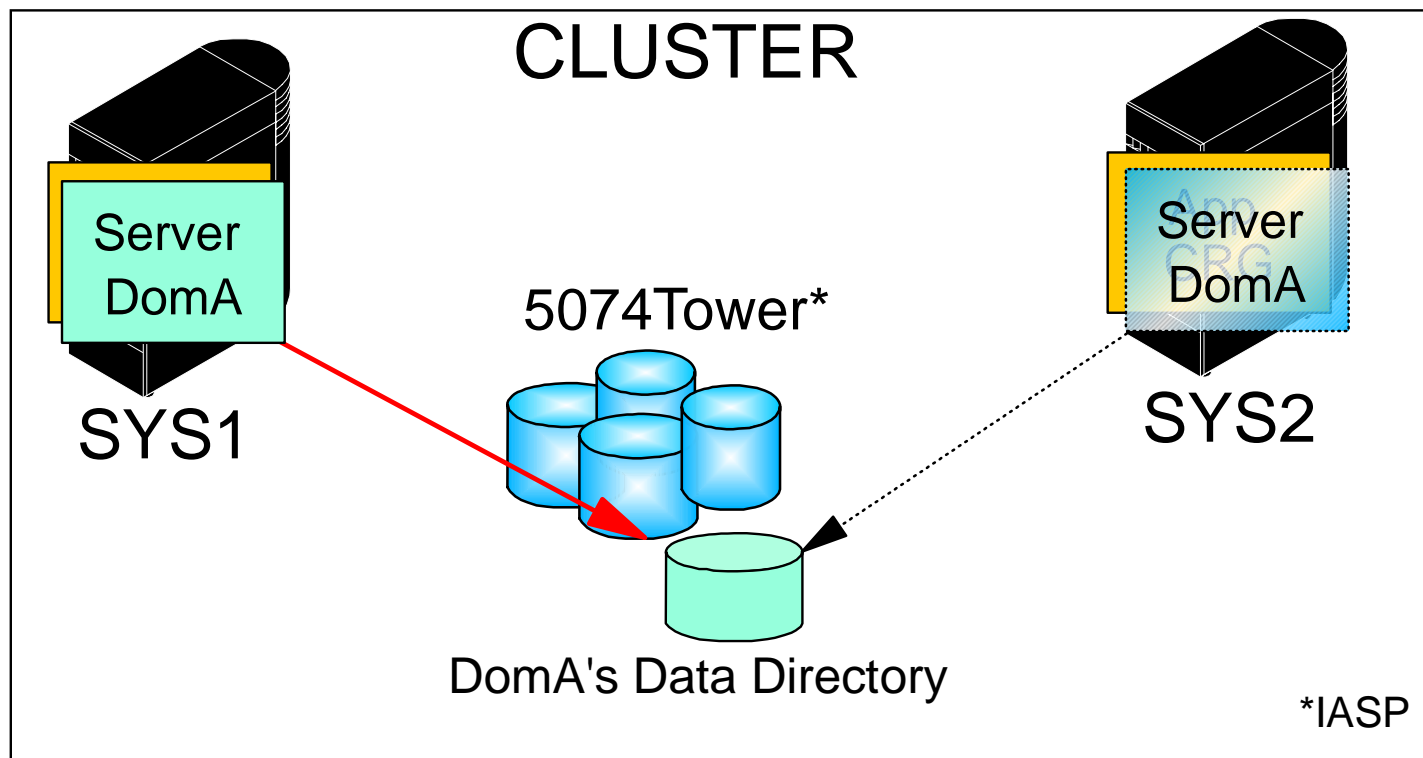
- ▶ **Same server can be "moved" between systems or logical partitions**
 - Scheduled logical partition or system maintenance

- ▶ **Cluster management Failover can automatically start Server on secondary system**

- ▶ **Domino servers can be managed by BP provided Cluster Management tools or IBM's Cluster Management Utility (Operations Navigator)**
 - (or OS/400 APIs and CL commands)

How does Domino use Switched IASPs?

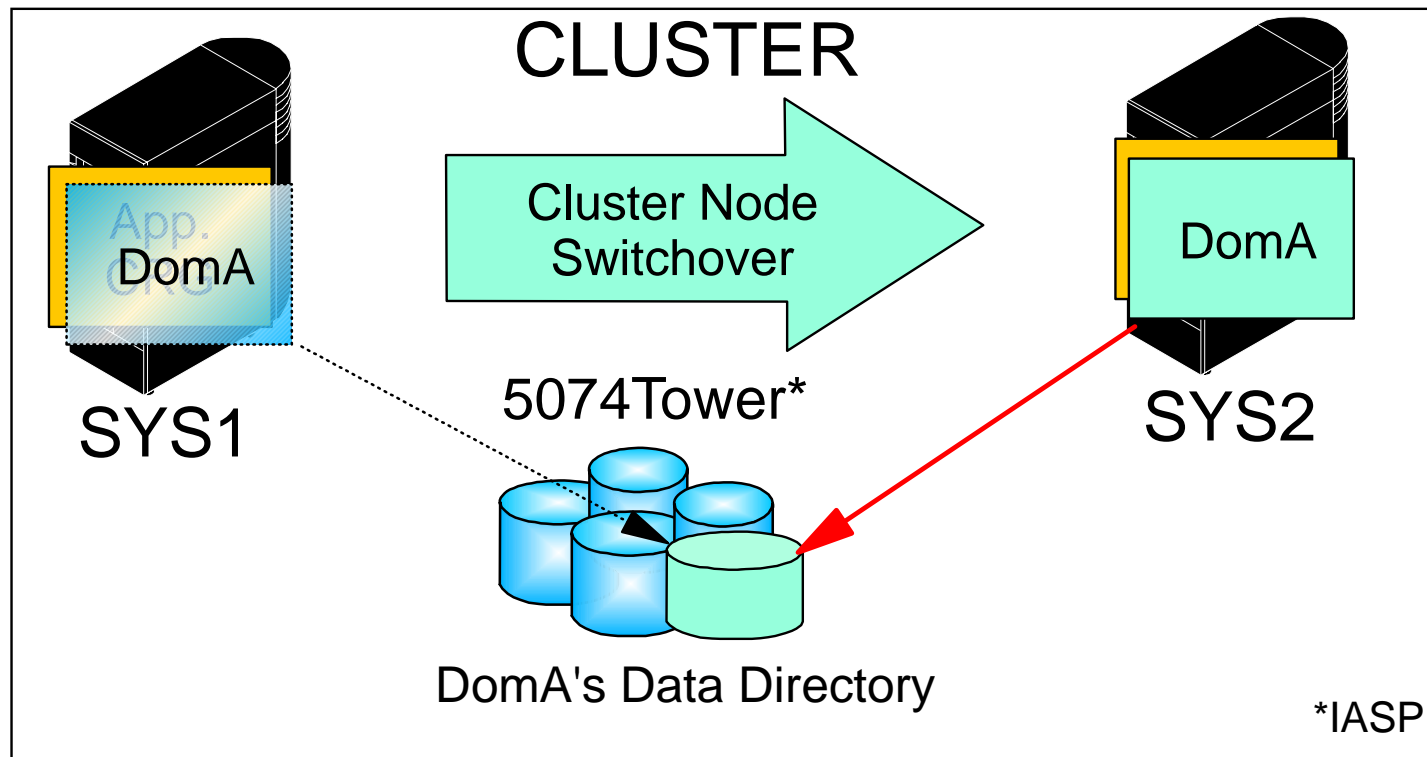
Domino server can be defined on more than one system. Its data directory can reside on switchable storage.



Allows a Domino server's data to be switched from one iSeries system to another

IBM  server. For the next generation of e-business.

Domino Switchover / Failover



Same server can be started on another iSeries system and access same Data Directory

This shows the "flow" of switching the IASP containing the Domino data directory and making it available on the second system using the same Domino Server.

For enhanced availability and recovery, the Domino Server on each node in the cluster can be setup to be part of an Application Cluster Resource Group. Once it is defined to OS/400 Cluster Management (CM) as part of an application CRG, the application can be started or stopped using CM interface and switched to start on another cluster node (system) where it is also configured.

For a detailed discussion on complete Domino for iSeries clustering capabilities refer to:
<http://www.ibm.com/eserver/series/domino>

Webserver can be defined to be highly available

- Web server state maintained across recovery domain
- Session state retrieved in event of switchover or failover
- Defined through the STRTCPSVR and STRCHTSVR interfaces

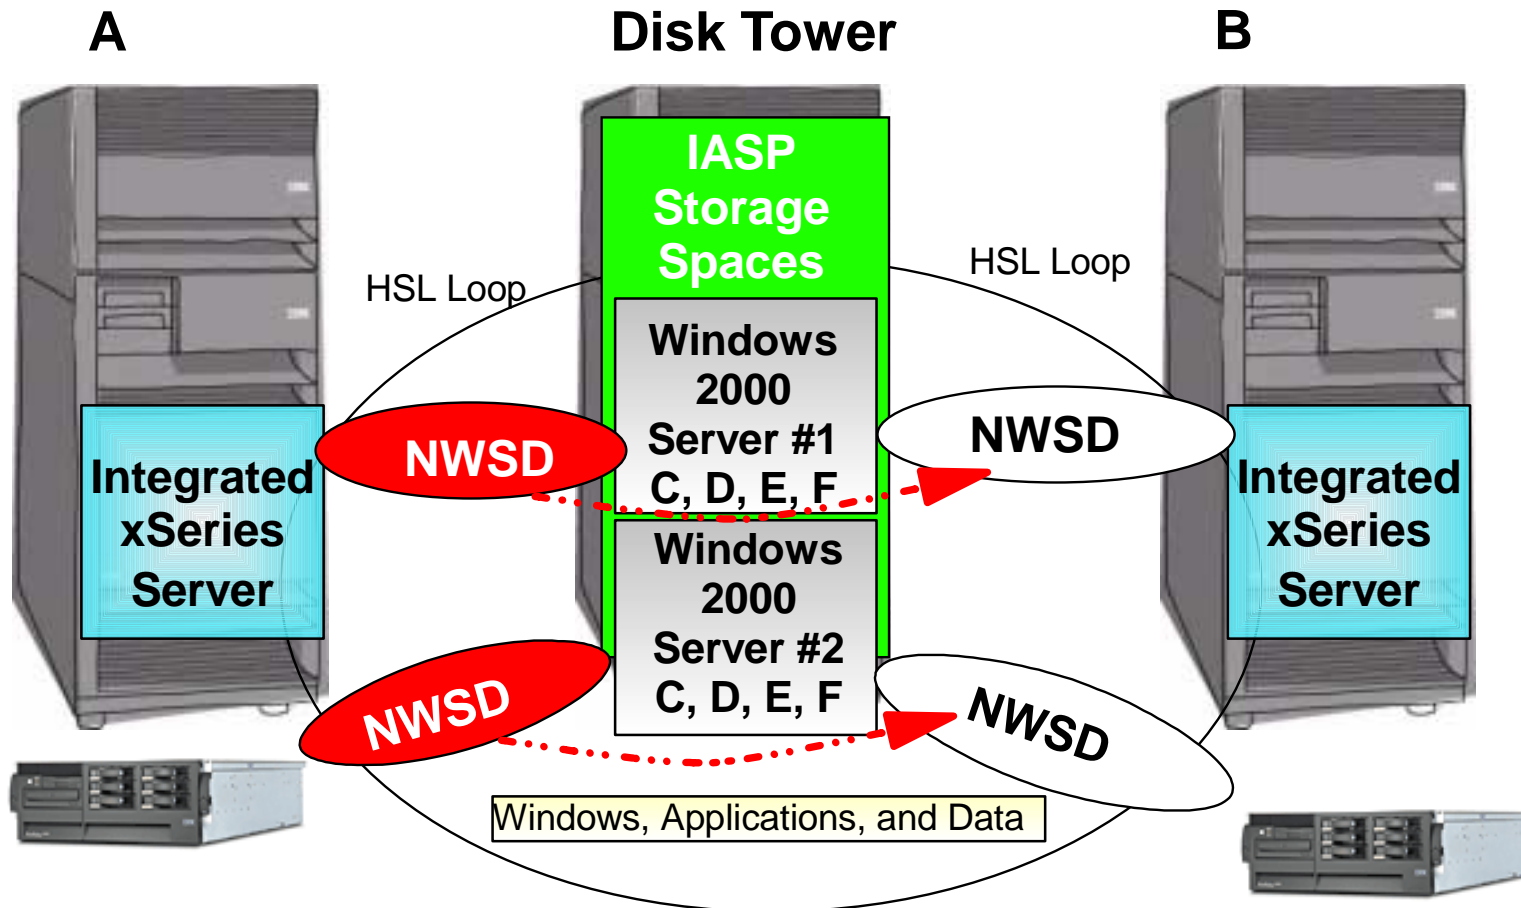
HA Models:

- Primary / backup
- Peer

Handling of application data is separate

Integrated xSeries Server and IASPs

Solution for planned and unplanned iSeries server outages



Support for Windows disks in Independent ASP

- ▶ Server A is running with IXS A and / or direct attached xSeries server A
- ▶ Take iSeries A offline. Disk Tower switches to iSeries B. Manually link NWSDs to B resource names, reboot Windows servers. Windows servers back online on B.
- ▶ xSeries servers need to have the same configuration

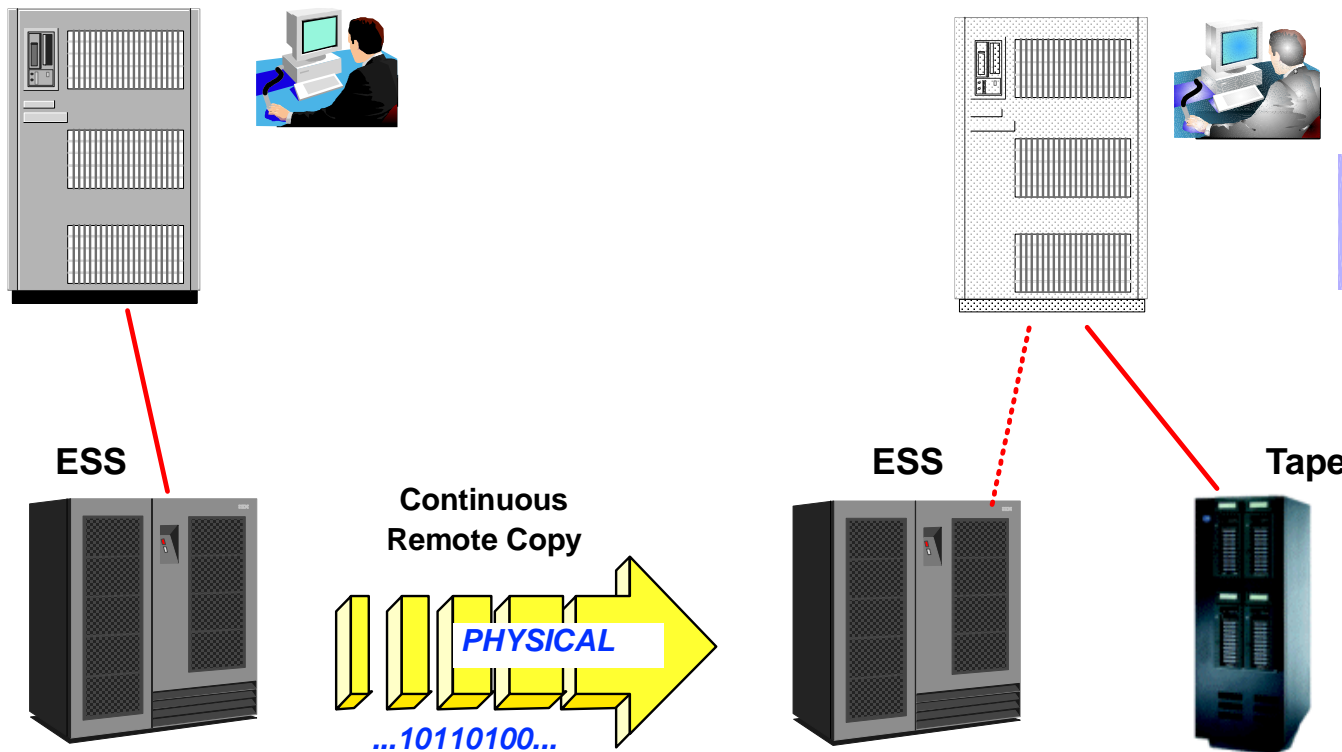
IBM  server. For the next generation of e-business.

Other Availability Solutions to Consider

IBM @server. For the next generation of e-business.

Peer to Peer Remote Copy*: Data Centric Approach

IBM  server iSeries



- Solution for data copy
- Disk-level approach only Disaster Recovery
- Must be synchronous connection
- Not designed for continuous availability
- Primary must be brought to restricted state or power down

- Second copy static at IPL
- Second copy not useable until IPL
- Full resynchronization required
- Understand the limitations before using

IBM  server. For the next generation of e-business.

This is a data resiliency strategy that can be used for disaster recovery. This is purely a disk based solution, where none of the hosts (the primary system or the stand by secondary system) are aware of a second set of disk storage that can theoretically be put quite a distance away from the primary machine.

You need to review the issue of price/performance. When attempting to utilize a storage mirroring configuration for disaster/recovery you effectively end up with twice the processing capacity installed, but are effectively utilizing only half of it. The backup copy of data that is there for disaster protection; it cannot be accessed for real time processing such as off-line batch or queries as can be done with an iSeries data replication cluster. Therefore, based on your availability requirements - you may want to fully understand the restrictions this solution has.

Further, the two systems must be connected synchronously to preserve order in the database on the second system. Synchronous means distance is limited to at best, 10 Km. Asynchronous database I/O does not maintain data integrity. Since the disk subsystems do not possess the understanding of when an application I/O is complete or not, it is recommended that the synchronization between the two disk storage subsystems is in synchronous mode. This ensures that the application I/O is always written to the disk and not sitting in the main memory for the usual FIFO rules to page it out to the disk units. If you allow asynchronous replication between the two storage servers, then there will be issues with data integrity since I/O in main memory of the source machine may not have been paged out to the disk units.

Are your primary concerns planned outages? For example, to perform a save operation using this storage feature, your system will need to be brought to a restricted state (an outage) prior to detaching the objects from the host environment. This is necessary to insure that you have a "clean" consistent copy, a snapshot of your system data at the time of the operation. And depending on the storage configuration, during this period the disk units may be left in an unprotected state, creating an exposure for a very long, unplanned outage should a single disk unit fail.

Further, there are application design considerations too. A suggestion would be to switch off journaling in exchange of using this approach for disaster recovery. As discussed already, the disk subsystems do not have knowledge of the application and its I/O requirements so the notion of switching off an audit trail of all updates is not valid. The fact is, journaling, sometimes called event logging on other platforms is about data integrity within a single footprint. Journaling is about moving data from mainstore to disk in a manner that locks the progress of the application to the removal of data from memory. (To make the discussion over simplified). The ultimate design for applications is to deploy commitment control - which requires application changes, if not already implemented.

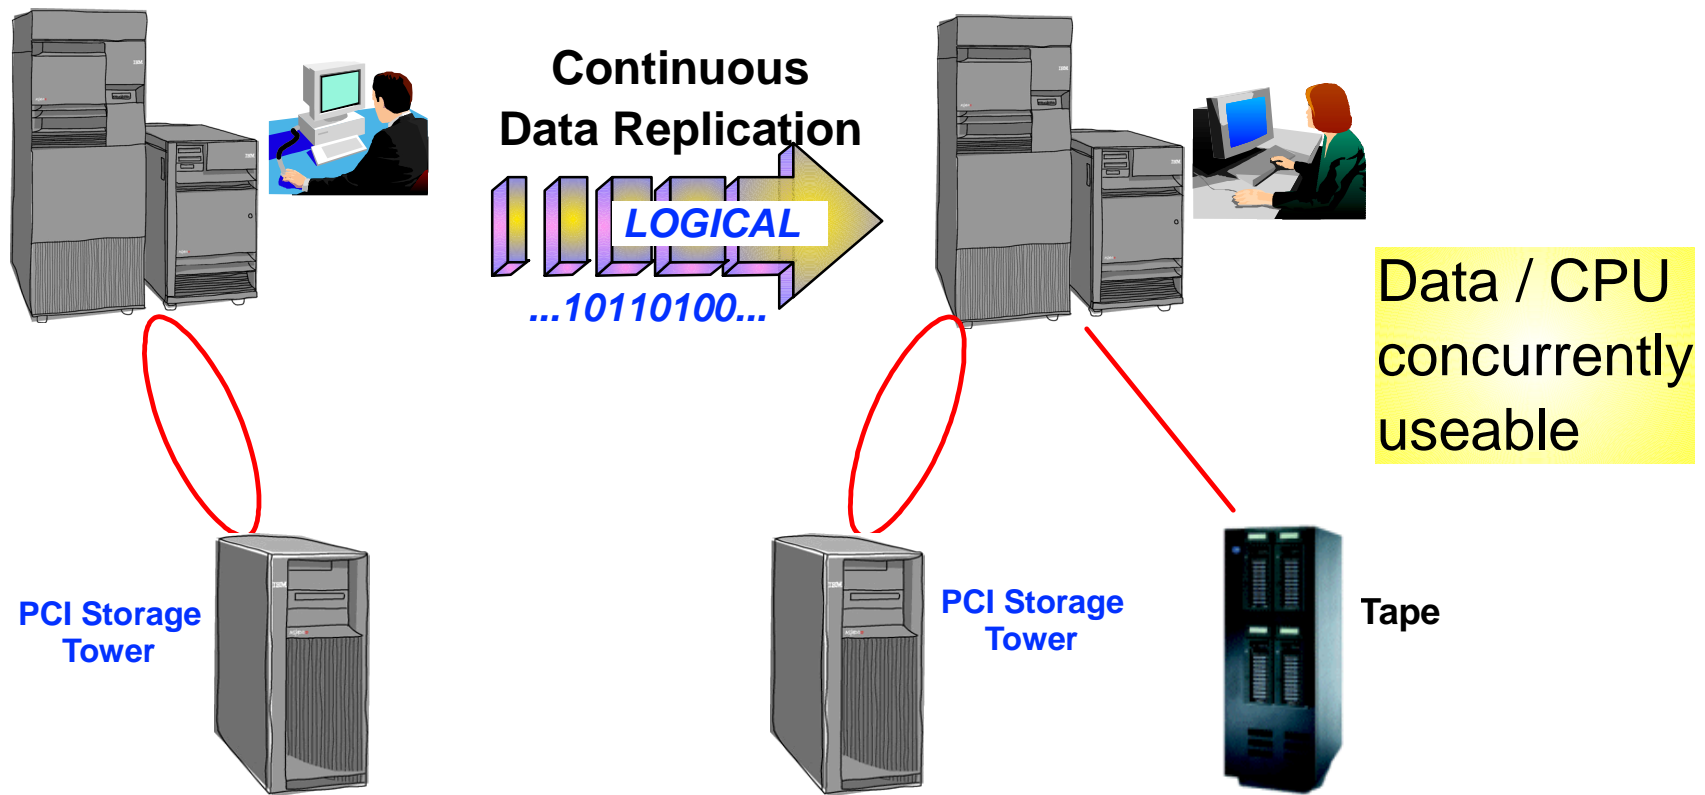
The ESS Shark product has several availability and recovery offerings, including Peer to Peer Remote Copy (PPRC) and Flash Copy. These offerings are being considered for iSeries support, but as of May 2001 formal support has not been announced. Look for availability information at the web site:

- <http://www.storage.ibm.com/>

For further iSeries information see "Roadmap to planning for Continuous Availability" whitepaper on iSeries website: <http://www.iseries.ibm.com/ha>

Data Resiliency Services

IBM  server iSeries



- Data Resiliency - Part of Cluster Architecture
 - System level approach to High Availability
 - Offloads read only workload to secondary server
 - Building block for Clustering
 - Addresses Disaster Recovery, Planned and Unplanned outages
 - Eliminates Save Window
 - Recommend for iSeries Environments

IBM  server. For the next generation of e-business.

Replication services is one of the iSeries methods for providing data resiliency. This service is provided in cooperation with LakeView Technology, Vision Solutions and DataMirror - IBM High Availability Business Partners (HABPs)


To accomplish data replication, IBM's HABPs make use of OS/400 journaling (called a log on other servers). When journaling is active, the system will add entries to a journal receiver when selected events occur, such as changes to database files, changes to other journaled objects, or security-relevant events. The HABP application coordinates and replicates these changes between a primary server and a backup server providing an extremely efficient method of continuous replication. OS/400 journaling is also a prerequisite to commitment control. Commitment control is an extension of journal management which allows applications to keep database files synchronized and ensures transaction integrity.

This form of replication still requires some level of operator intervention in the event of an unplanned outage. The solution to get around manual intervention is to implement a clustering solution with replication services which we will discuss next.

Complex transactions,
Continuous availability

What matters

- Highly Reliable Hardware
- Highly Reliable Operating System
- Concurrent Operations/Maintenance
- Data Resiliency
- Application Resiliency
- Transaction Monitoring
- Clustering Technology



OS/400
Functions



Data
Propagator

IBM Cluster
Middleware
Business Partners

IBM @server. For the next generation of e-business.

This foil summarizes the requirements of a full, automated as much as possible, high availability environment that uses clustering and can take advantage of new V5R1 support.

The past decade has focused on data resiliency. While mirroring and Raid5 increase the availability of the data source, Data Replication tools such as DataPropagator and the IBM Cluster Middleware Business Partners (High Availability Business Partners (HABPs), such as LakeView, Vision and DataMirror solutions are primarily about data resiliency. That is, a copy of the data together with information about the currency of the data is available to provide data availability.

The problem with solutions which only focus on data is that they cannot be 24 hours by 365 days (25 x 365). Switching between systems requires **application resiliency and transaction signaling as well as data availability**. That is why clustering technology was introduced in V4R4 and why the focus now is to include the application and the data together in a comprehensive solution called the cluster. The external disk vendors provide at best a data copy function and therefore cannot integrate their disk replicating technologies into a clustering solutions as they have no knowledge of the data currency, the transaction status and of the application architecture.

Having highlighted the importance to application resiliency, it is equally important to recognize that in order to get continuous availability, the applications have to be designed in a way which allows them to return to their previous known failure state. In other words, the job state and the application state has to be maintained. This cannot be controlled through only a data copy or through any disk storage subsystems.

For further information about designing your applications for high availability, please visit the iSeries website:
<http://www.iseries.ibm.com/ha> to review whitepapers and information on how you can enable your applications to achieve ClusterProven status.

Technology Comparison

	Replication	Switched disk
Flexibility	10's of systems	2 systems
Single Point Of Failure	none	disk subsystem
Cost factors	Additional disk capacity. Replication software.	Switchable I/O expansion tower
Performance factors	Replication overhead	(none)
Typical failover time	15 minutes*	15 min*
Typical switchover time	~ 5 minutes	~ 5 minutes
Real time coverage	Anything journaled	Anything on IASP
Geographic dispersion	Unlimited	Limited attach distance
D/R protection	Yes	No
Concurrent backup	Yes	No
Setup	Replication environment. What to replicate.	Device domain environment. Populate IASP.

* Does not include interrupted transaction, if any, recovery

iSeries Cluster Solution

Cluster Middleware Products



DataMirror[®]


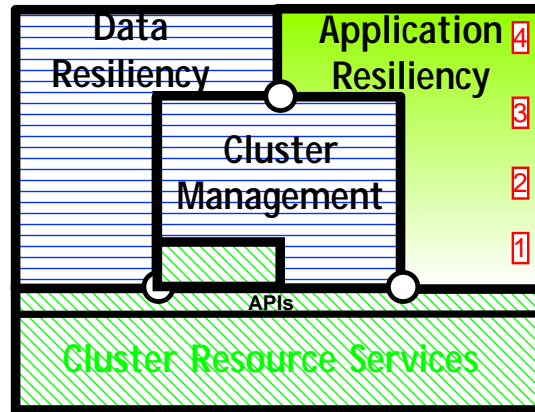
Data from where it is
to where it needs to be



LAKEVIEW
TECHNOLOGY



IBM Management Central
Cluster Management Utility
IFS two node switch disk only



Cluster
Connectivity

Cluster Enabled Applications



Baan



jack henry
& Associates Inc.



ClusterProven



Introduction of Simple Cluster Management in support of Switch Disk technology

IBM  server. For the next generation of e-business.

IBM Web Sites

- <http://www.ibm.com/eserver/iserries/ha>
- <http://www.ibm.com/servers/eserver/iserries/ha> ←
- <http://www.storage.ibm.com/>

High Availability Business Partner Web Sites

- <http://www.datamirror.com/>
- <http://www.lakeviewtech.com/>
- <http://www.visionsolutions.com/>

ITSO Redbooks

- <http://www.redbooks.ibm.com/>
- redbook - AS/400 Clusters: A Guide to Achieving Higher Availability, SG24-5194

Other Availability Enhancements

IBM @server. For the next generation of e-business.

Journaling:

- IFS Directories*
- IFS Symbolic Links
- Data Areas
 - Forward and Backout recovery
- Data Queue Journaling

Journaling Minimal Data

New Operations Navigator interfaces to OS/400 functions available with the following commands

- STRJRNOBJ
- ENDJRNOBJ
- STRJRN
- ENDJRN

Note: The V5R1 Independent ASP - switched disks support includes IFS objects but not library-based objects such as OS/400 Database and Journal objects. This means objects within a user ASP may be journaled, but objects within an IASP cannot be journaled in V5R1. Journaling IASP data is planned for the next release.

High Availability and clustering support are becoming "table stakes" for servers. Particularly important is the ability to maintain copies of data on multiple systems. Sometimes this is used for high availability (i.e. failover) support and sometimes for load balancing. Providing the ability to journal *DIR (directory) and *SYMLNK (symbolic link) objects is the second stage of providing Journaling support for IFS objects. Replication by means of a journal for a *STMF (stream file) object was made available in V4R4.

V5R1 also adds journaling data areas and data queues.

Completely new in V5R1 is the ability to recover all the above mentioned journaled objects (*STMF, *DIR, *SYMLNK, *DTAARA and *DTAQ) to a good status after an abnormal IPL or system crash even when the object was in the middle of a change action at this event. The next foil shows a 5250 command help text of the Start Journal command describing some this new "object type" support.

For object recovery from a previously saved version when the object on the server is lost, there are some restrictions.

For *STMF, *DIR and *SYMLNK type objects there is support for forward recovery only. Data areas however do support forward (APYJRNCHG) as well as back out (RMVJRNCHG) recovery if *BOTH images are preserved with the journal. For data queues, there is no forward or back-out recovery since the current implementation does not preserve the data in a data queue when it is saved.

In V5R1 a new parameter for the CRTJRN and CHGJRN commands, you can specify to make minimized journal entries. This will decrease the size of your journal entries. Entries will only be minimized if the minimized entry is smaller in size than a complete journal entry deposit would be. **Journal receivers with object types allowing minimized entry specific data cannot be saved and restored to any release prior to V5R1M0 nor can they be replicated to any remote journal on a system at a release prior to V5R1M0.**

After the next foil we show some V5R1 Operations Navigator windows.

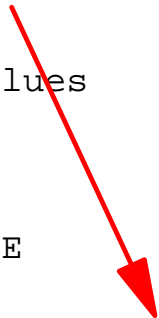
IBM  server. For the next generation of e-business.

Start Journaling (STRJRN)

Start Journal (STRJRN)

Type choices, press Enter.

```
Objects:                                OBJ
Name . . . . .
-----
Include or omit . . . . . *INCLUDE
                        + for more values
File identifier . . . . . OBJFID
                        + for more values
Journal . . . . . JRN
Directory subtree . . . . . SUBTREE *NONE
```



File identifier (OBJFID) - Help

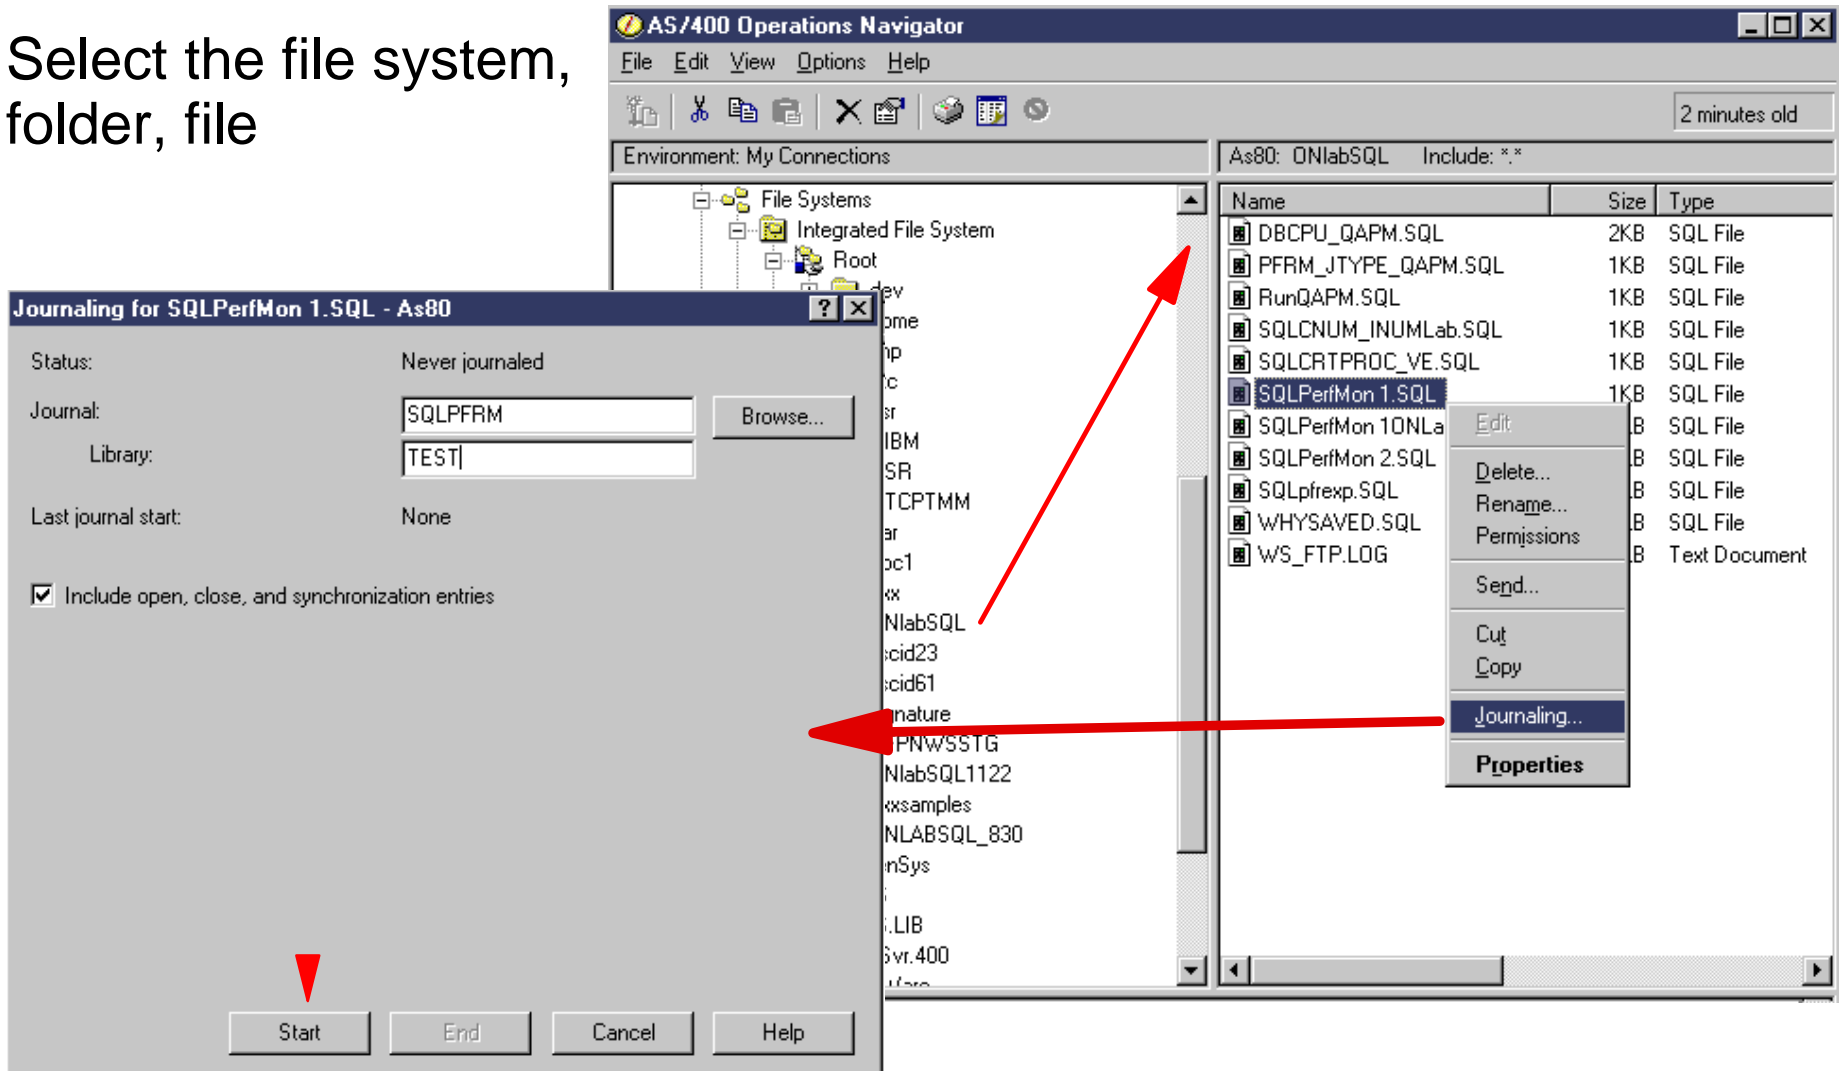
Specifies a maximum of 300 file identifiers (FID) for which changes are to be journaled. FIDs are a unique identifier associated with integrated file system related objects. This field is input in hexadecimal format. Only objects whose FID identifies an object of type *STMF, *DIR, *SYMLNK, *DTAARA or *DTAQ are supported.

Notes: Start Journaling (STRJRN)

This foil gives us examples of the new subdirectory and new object type journaling support.

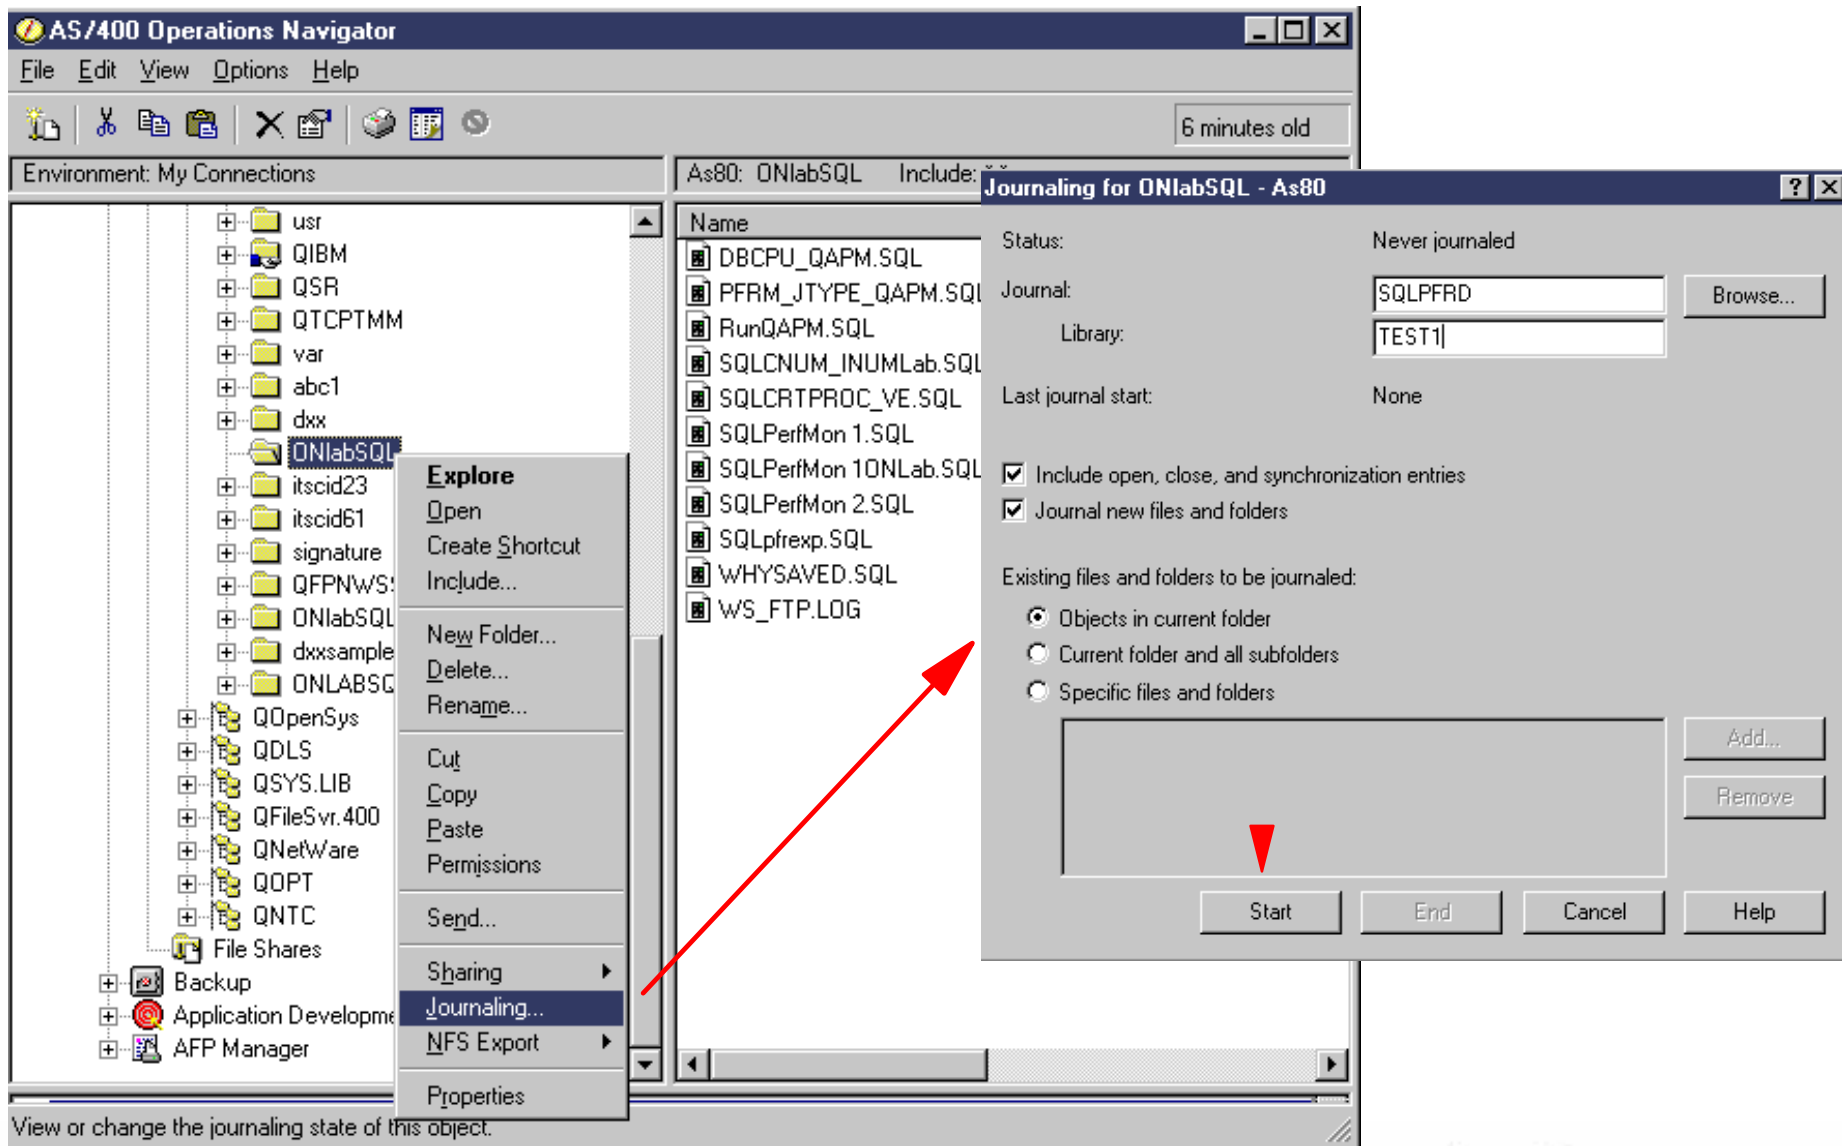
Operations Navigator-Journal stream file

- Select the file system, folder, file



Operations Navigator-Journal directory

- Select the file system, folder



BRMS Enhancements

BRMS Graphical Interface as optional Operations Navigator plug-in -First Stage

- Wizards for creating backup policies, adding media to the BRMS inventory, and restore saved objects
- Management interfaces for Backup and Restore, including scheduling

Save library using parallel devices now includes *ALLUSR, *IBM, *ALLPROD, *ALLTEST, *ASPnn, and generic library names

- Via SAVLIBBRM or Backup Control Groups

Backup Control Group *SYSTEM shipped with BRMS

STRBKUBRM parameters minimize backup control groups needed --> less media policies

Target release now supports VxRxMx syntax up to N-2

Passwords ("pass phrases") up to 128 characters and up to 32 ASPs supported


Improved report contents

Improved Domino for AS/400 incremental saves

Includes IASP support

Operations Navigator and BRMS Stage 1

Backup Recovery and Media Services for iSeries



Backup Recovery and Media Services for iSeries, 5722-BR1, has many new ease of use functions for V5R1.

Backup Recovery and Media Services (BRMS) is the IBM strategic solution for planning and managing the backup of your iSeries server. The BRMS product is available on the keyed stamped media shipped with every server. This product can be installed and used for 70 days, without charge. After 70 days a license key is required. BRMS has been refreshed and updated for V5R1 to include the following enhancements:

- BRMS now supports graphical operations by providing an optional plug-into Operations Navigator. You can install the BRMS plug-in on any workstation that has been upgraded with IBM Client Access Express for Windows, 5722-XW1 connected to any AS/400 or iSeries running BRMS 5722-BR1. When installed, a Backup Recovery and Media Services folder is added to the Operations Navigator hierarchy. The BRMS Operations Navigator plug-in simplifies backup planning by providing wizards for creating backup policies and adding media to the BRMS managed inventory. Context menu functions on the backup policies allow you to easily run and schedule backups. The BRMS restore wizard guides you through the steps to locate and restore saved objects. In addition, the BRMS Operations Navigator plug-in integrates a Backup... and Restore... function into many of the object context menus in the hierarchy allowing you to easily backup and restore these objects directly to tape media.
- The V5R1 BRMS Operations Navigator functions, though extensive, are a subset of all the functions available through other interfaces such as BRMS commands. V5R1 support should be considered "stage 1" level of support with "stage 2" planned for delivery in 2002.
- The save library support, either through the SAVLIBBRM command or Backup Control Groups, using parallel devices, has been enhanced to now include: *ALLUSR, *IBM, *ALLPROD, *ALLTEST, *ASPnn, and generic library names.
- A new Backup Control Group named *SYSTEM is shipped with the BRMS product. This control group can be used to backup the entire system including all user data. Prior to V5R1, you had to either use both the *SYSGRP and *BKUGRP control groups to complete this save or create your own customize policy. This *SYSTEM control group uses a new media policy named SYSTEM, which has a default retention of 90 days.

- The STRBKUBRM command is enhanced with two new parameters that helps you minimize the number of backup control groups you need to create and use. The ACTIVITY parameter allows you to override the weekly activity of the control group entries and can be used to force either a full or incremental backup. The RETENTION parameter allows you to override the retention settings of the media policies used by the control group. If your media polices have a retention of 30 days, you can use the RETENTION parameter to override this for a single backup to a new retention of *PER or a number of days. This might be useful if you wanted to keep a copy of your normal saves at year end for 365 days, or longer. These new control group attribute overrides are resolved at the time the command is run and do not change the stored attributes of the backup control group or media policy.
- The target release parameters on the Save Library using BRMS(SAVLIBBRM) and Save Object using BRMS (SAVOBJBRM) commands, and the target release attribute of Backup Control Groups have been updated to support the VxRxMx format for specifying a target release, where Vx is the version, Rx is the release and Mx is the modification level. This allows you to save objects that you intend to restore on previous release systems. Previously, BRMS restricted you to N-1 from the release of the save. Now you can save objects through BRMS and restore on systems that are N-2 from the release of the save.
- The BRMS System Recovery Report (QP1ARCY) has been improved. Some recovery actions that had previously included multiple tasks were moved into separate steps to minimize the likelihood of the actions being missed during recovery. Potential problems that might affect system recovery are highlighted better to minimize recovery exposures. More steps were added to the report reducing some of cross references to the iSeries Backup and Recovery book.
- The BRMS Console Monitor has been updated to support pass phrases of up to 128 characters.
- BRMS has increased the support for traditional user ASPs from 16 to 32 and supports Independent ASPs as ASP numbers 33-99. See next foil.
- BRMS is being enhanced to provide incremental Domino saves. A BRMS PTF will be needed to use this function. For PTF information refer to: <http://www-1.ibm.com/servers/eserver/series/service/brms.htm>

BRMS for iSeries and New ASP support

Save/Restore of User Defined File Systems on IASPs via SAVBRM/RSTBRM or link lists

- Required: the path of the IASP **/dev/iasp-name** is fully defined

BRMS V5R1 parameters (1-99)*

- 17-32
- 33-99 (*Link)

*Easier to use interface planned for next release

Notes: BRMS for iSeries and New ASP support

IBM  server iSeries

For V5R1 BRMs supports ASP values of 16-99, where OS/400 starts numbering IASPs at 33.

The following commands support these additional ASP numeric values:

- RSTAUTBRM, STRRCYBRM, DSPASPBRM, ADDMEDIBRM, MGRBRM
- MOVSPFBRM, WRKASPBRM, WRKMEDIBRM, WRKMGRIBRM
- SAVxxxxBRM/RSTxxxxBRM commands with SAVFASP parm.

An easier to use interface is planned for the release following V5R1. For saving and restoring a User Defined File System (UDFS), such as via the SAVBRM/RSTBRM or link lists, you must use the path of the IASP value:

- **/dev/iasp-name**

IBM  server. For the next generation of e-business.

Trademarks & Disclaimers

© Copyright International Business Machines Corporation 2001

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

AIX	Application Development	AS/400
AS/400e	DB2	Domino
IBM	OfficeVision	OS/400
Integrated Language Environment	Net.Commerce	Net.Data
PowerPC	PowerPC AS	SanFrancisco
Host on Demand	Screen Publisher	Host Publisher
PCOM	WebSphere Commerce Suite	Payment Manager
WebSphere	WebSphere Standard Edition	WebSphere Advanced Edition
MQSeries	MQSeries Integrator	Host Integration Series
WebSphere Development Tools for AS/400	VisualAge for Java	VisualAge for RPG
CODE/400	DB2 UDB for AS/400	HTTP Server for AS/400
iSeries		

Lotus, Freelance, 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.

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.

IBM  server. For the next generation of e-business.