



# Implementing Oracle's JD Edwards World A7.3 in an IASP environment on IBM System i5



*Kathy Mack*  
*System i5 ERP Development – Oracle Support*  
*ISV Business Strategy and Enablement*  
*July 2006*



## Table of contents

<b>Abstract</b> .....	<b>1</b>
<b>Introduction</b> .....	<b>1</b>
<b>Considering the JD Edwards World environment</b> .....	<b>2</b>
Implementation overview .....	2
Software prerequisites .....	3
<b>Configuring the IASP</b> .....	<b>4</b>
Using HMC to configure hardware.....	4
Creating the cluster and CRG.....	8
Configuring the disk for the IASP.....	9
<b>Moving JD Edwards World libraries to the IASP</b> .....	<b>10</b>
Identifying libraries used in the JD Edwards World environment .....	10
Saving libraries to save files .....	12
Removing libraries in the system ASP.....	12
Restoring libraries to the IASP.....	12
<b>Configuring JD Edwards World user profiles</b> .....	<b>12</b>
Creating user profiles.....	12
Creating job descriptions .....	13
Changing user profiles to access new job descriptions.....	14
<b>Switching between nodes</b> .....	<b>14</b>
Planned IASP switch.....	14
Unplanned IASP switch .....	14
Performance considerations .....	14
<b>Summary</b> .....	<b>15</b>
<b>Appendix A: Hardware and software tested</b> .....	<b>16</b>
<b>Appendix B: Useful IASP commands</b> .....	<b>16</b>
<b>Appendix C: Resources</b> .....	<b>17</b>
<b>Appendix D: About the author</b> .....	<b>17</b>
<b>Trademarks and special notices</b> .....	<b>18</b>



## Abstract

---

*The purpose of this paper is to demonstrate how to set up, configure and move a JD Edwards World A7.3 application to a cluster-based, IBM System i environment that implements an independent auxiliary storage pool (IASP).*

## Introduction

---

High availability (HA) solutions come in various forms and levels of complexity. In the IBM® System i™ environment, the switchable independent auxiliary storage pool (IASP) is one of the easier HA choices to implement and use.

Because of loop limitations with high-speed link (HSL) cables, the systems must be within several hundred meters of each other. Therefore, IASPs are not useful in a catastrophe as a disaster recovery solution.

It is assumed that the Oracle® JD Edwards® World application is successfully installed and runs without errors in the system ASP on a System i model. This paper will step through the process of setting up and configuring a switchable IASP, then moving the JD Edwards World application from the system ASP to the IASP, and finally, the actual switching of the IASP from the primary node to the backup node.

## Considering the JD Edwards World environment

When planning the use of IASPs in your environment, several factors are important to consider: the needs of the business, performance and the existing System i configuration. The IBM Redbook entitled *IBM eServer iSeries Independent ASPs: A Guide to Moving Applications to IASPs* walks the reader through these different requirements, pointing out issues to consider during the planning phase.

### Implementation overview

The cluster environment used for this IASP implementation consisted of two partitions on a System i model; one was used as the primary node, and a second partition was used as the backup node. An expansion tower within this system was configured as the IASP. Because the primary node, backup node, and the IASP were all located physically within the same system, there was no requirement for additional HSL cabling.

It is possible to run JD Edwards World™ in an IASP within a cluster/HSL loop that is configured with multiple systems. The process to install the JD Edwards World application in this environment is the same when the cluster/HSL loop is in place. For information on configuring and using HSL connections in a clustered environment, see the additional resources listed in “Appendix C.” Figure 1 contains a diagram of the JD Edwards World environment used for this paper.

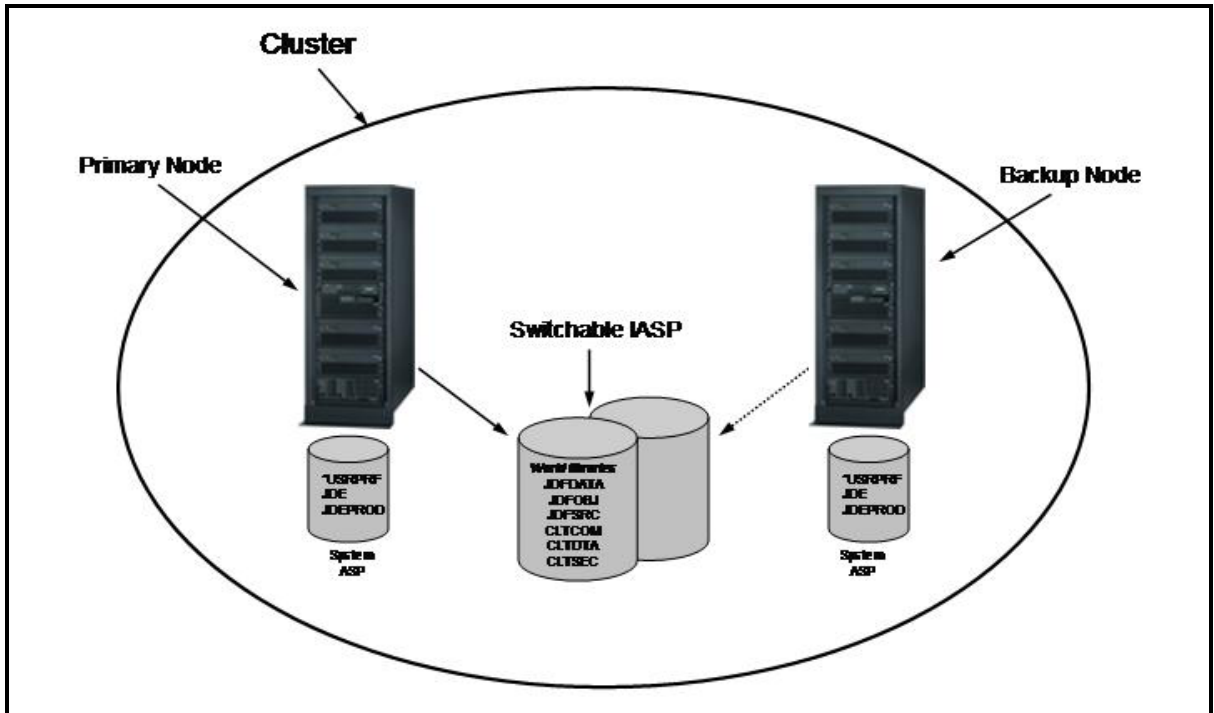


Figure 1. JD Edwards World IASP implementation



## Software prerequisites

The required IBM i5/OS® operating system software products for configuration and use of an IASP are listed below. These products need to be in place on both the primary and backup nodes.

5722SS1	41	HA Switchable Resources
5722TC1	*BASE	IBM TCP/IP Connectivity Utilities for i5/OS
5722XE1	*BASE	IBM eServer™ iSeries™ Access for Windows

A workstation running iSeries Navigator is also needed.

The following steps must be taken:

1. Ensure that the INETD TCP server is started on all nodes before beginning the cluster configuration. Figure 2 shows the INETD in an iSeries Navigator window. If this server needs to be started, right-click on the server in the iSeries Navigator window and select *Start* or run the following command from a 5250 screen: `STRTCPSVR SERVER(*INETD)`

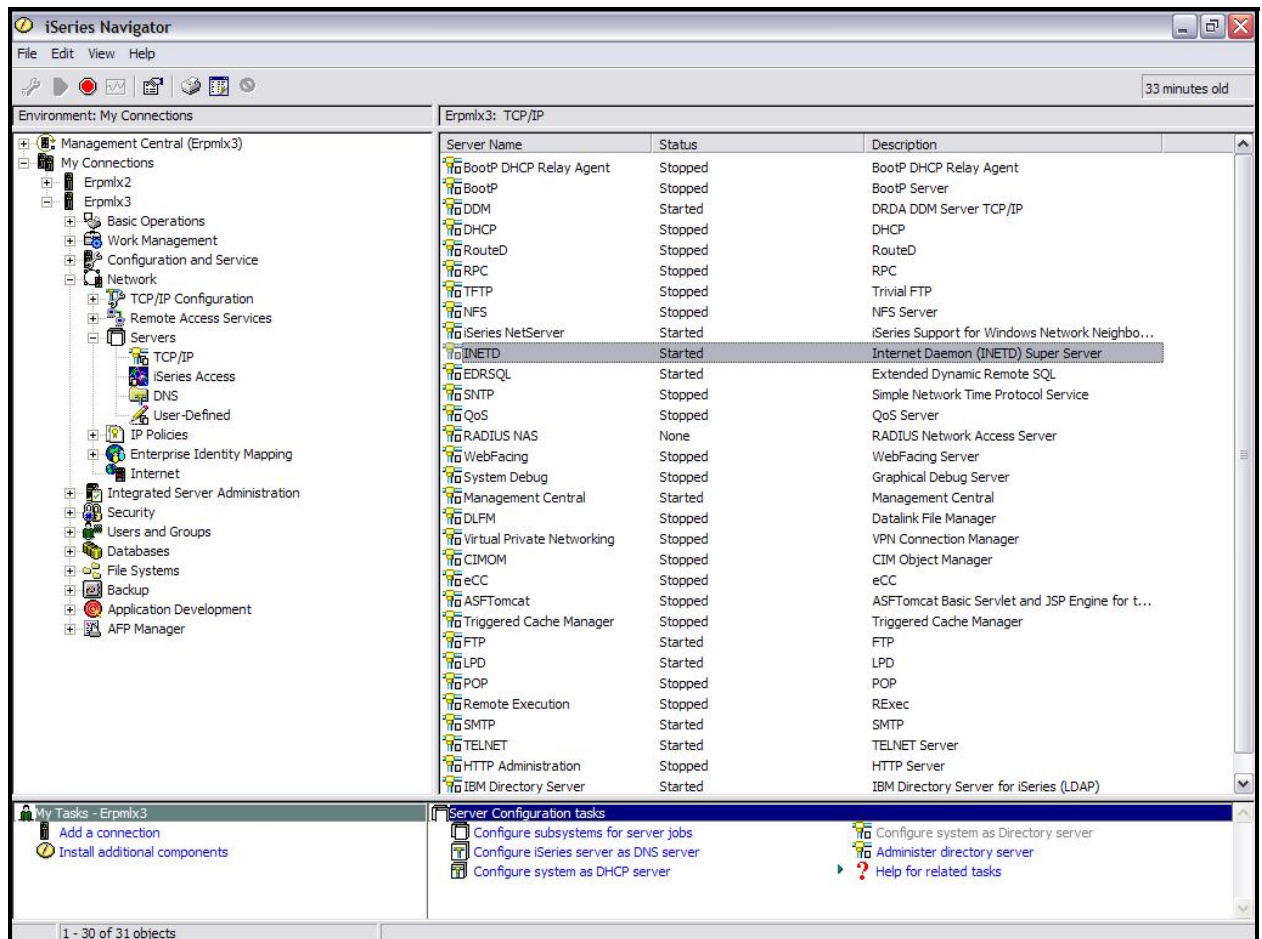


Figure 2. iSeries Navigator, TCP/IP Servers



2. Check the network attribute ALWADDCLU is set to \*ANY for each node in the cluster. Enter **DSPNETA** and page down to the line labeled *Allow add to cluster*.

```
Display Network Attributes
System:  ERPMLX3
Maximum hop count . . . . . : 16
DDM request access . . . . . : *OBJAUT
Client request access . . . . . : *OBJAUT
Default ISDN network type . . . . . :
Default ISDN connection list . . . . . : QDCCNNLANY
Allow AnyNet support . . . . . : *NO
Network server domain . . . . . : D10421CF
Allow APPN virtual support . . . . . : *NO
Allow HPR transport tower support . . . . . : *NO
Virtual controller autocreate APPC device limit : 100
HPR path switch timers:
  Network priority . . . . . : 1
  High priority . . . . . : 2
  Medium priority . . . . . : 4
  Low priority . . . . . : 8
Allow add to cluster . . . . . : *ANY
Modem country or region ID . . . . . :
Press Enter to continue.
F3=Exit  F12=Cancel
Bottom
```

Figure 3. DSPNETA

3. If this attribute is not set to \*ANY, use the following command to change it:  
CHGNETA ALWADDCLU(\*ANY)

## Configuring the IASP

This section of the paper provides an overview of the steps to create the cluster, the cluster resource group (CRG), and the IASP. For a detailed description of these steps and the functions of each of the elements, refer to the Redbook entitled *IBM eServer iSeries Independent ASPs: A Guide to Moving Applications to IASPs*.

### Using HMC to configure hardware

The first task in configuring an IASP is to ensure that the desired disk drives are shareable across partitions. This is done through the IBM Hardware Management Console (HMC) for System i.

1. Sign on to the HMC.
2. Expand *Server* and *Partition*.
3. Select *Server Management*.
4. In the primary node, take the following steps (see Figures 4 and 5):
  - a. In the right pane, expand *Partitions <primary node>*.
  - b. Right-click the partition profile and select *Properties*.
  - c. Select the *Physical I/O* tab.
  - d. Select the switchable bus in the *Managed system I/O devices* pane and click *Add* as desired.
  - e. In the *profile I/O device* pane, click the I/O pool column for the switchable bus and enter a number for the I/O pool. (The acceptable numbers for an IASP range from 33 to 255).

- f. Expand the switchable bus and ensure that everything associated with the bus now belongs to this I/O pool.
- g. Click the *Advanced* button and ensure that the new I/O pool is listed under I/O pools.
- h. Select the *OptiConnect* tab.
- i. Check the boxes for *Use virtual OptiConnect* and *Use High Speed Link (HSL) OptiConnect*.
- j. Power down the partition from the green screen by using the following command:  

```
PWRDWN SYS OPTION(*IMMED) RESTART(*NO)
```
- k. When the system is down, restart the partition from the HMC by selecting the partition profile and selecting *Activate*.

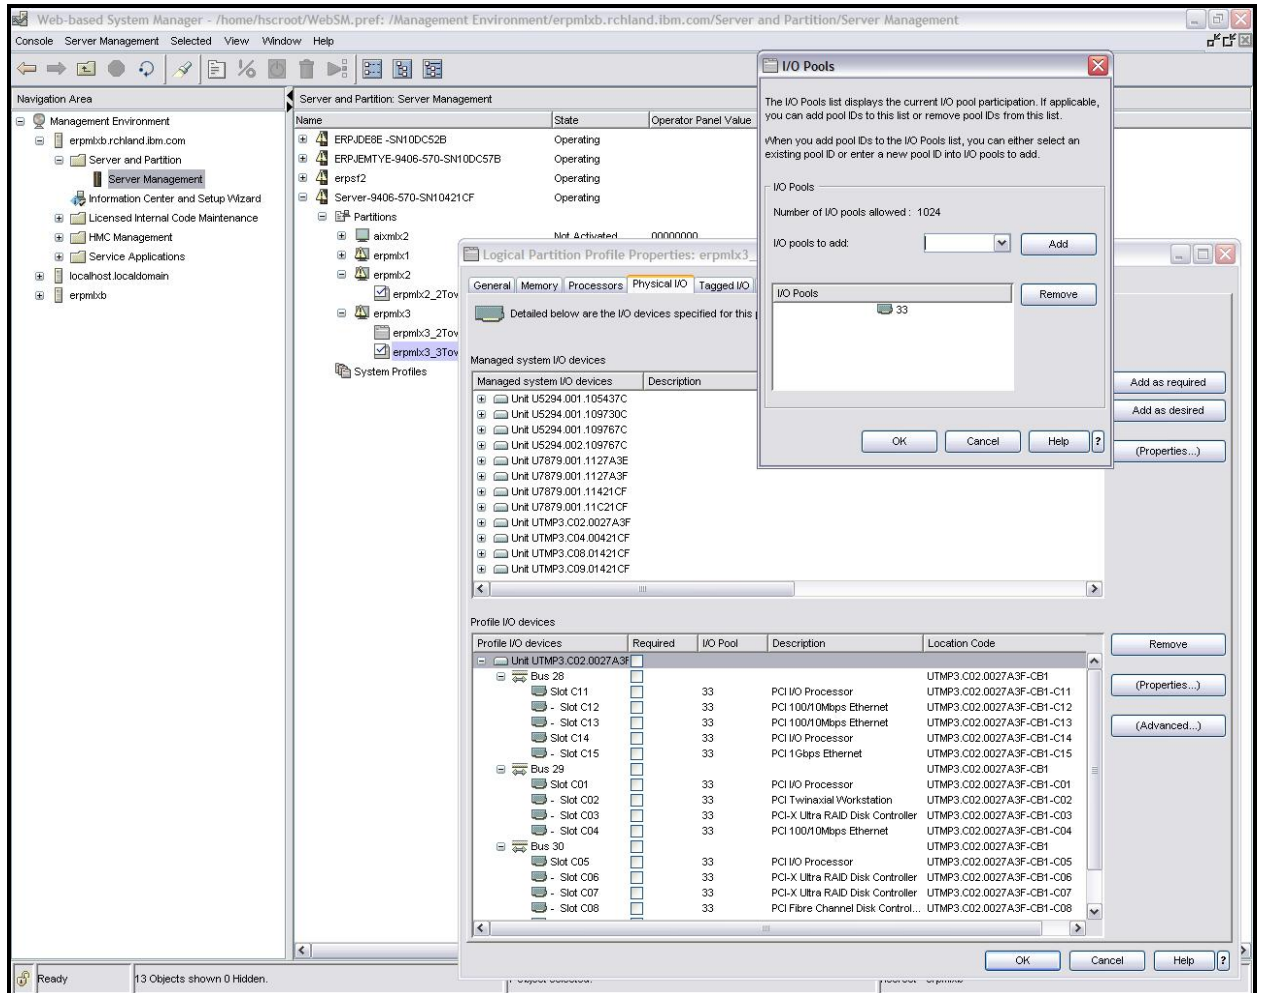


Figure 4. HMC configuration for primary node

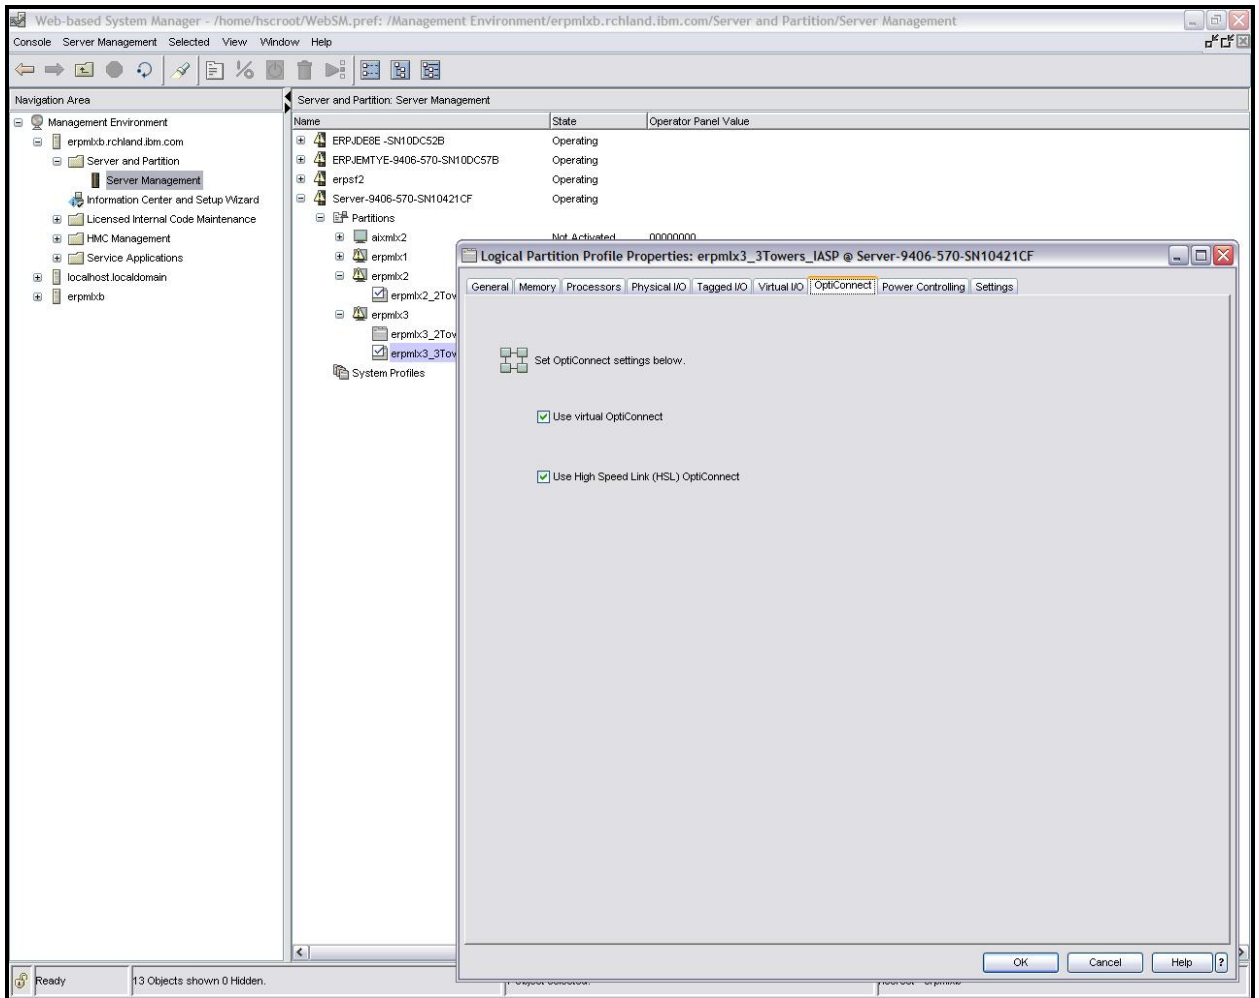


Figure 5. HMC configuration for primary node

5. In the backup node, take the following steps (see Figure 6):
  - a. In the right pane, expand *Partitions <backup node>*.
  - b. Right-click the partition profile and select *Properties*.
  - c. Select the *Physical I/O* tab.
  - d. Ensure that the switchable bus is not present in the *profile I/O device* pane.
  - e. Click the *Advanced* button.
  - f. Enter the newly created I/O pool number in the *I/O pools to add* field and click *Add*.
  - g. Ensure that the new I/O pool is listed under I/O pools.
  - h. Select the *OptiConnect* tab.
  - i. Check the boxes for *Use virtual OptiConnect* and *Use High Speed Link (HSL) OptiConnect*.
  - j. Power down the partition from the green screen by using the following command:
 

```
PWRDWN SYS OPTION (*IMMED) RESTART (*NO)
```
  - k. When the system is down, restart the partition from HMC by selecting the partition profile and selecting *Activate*.



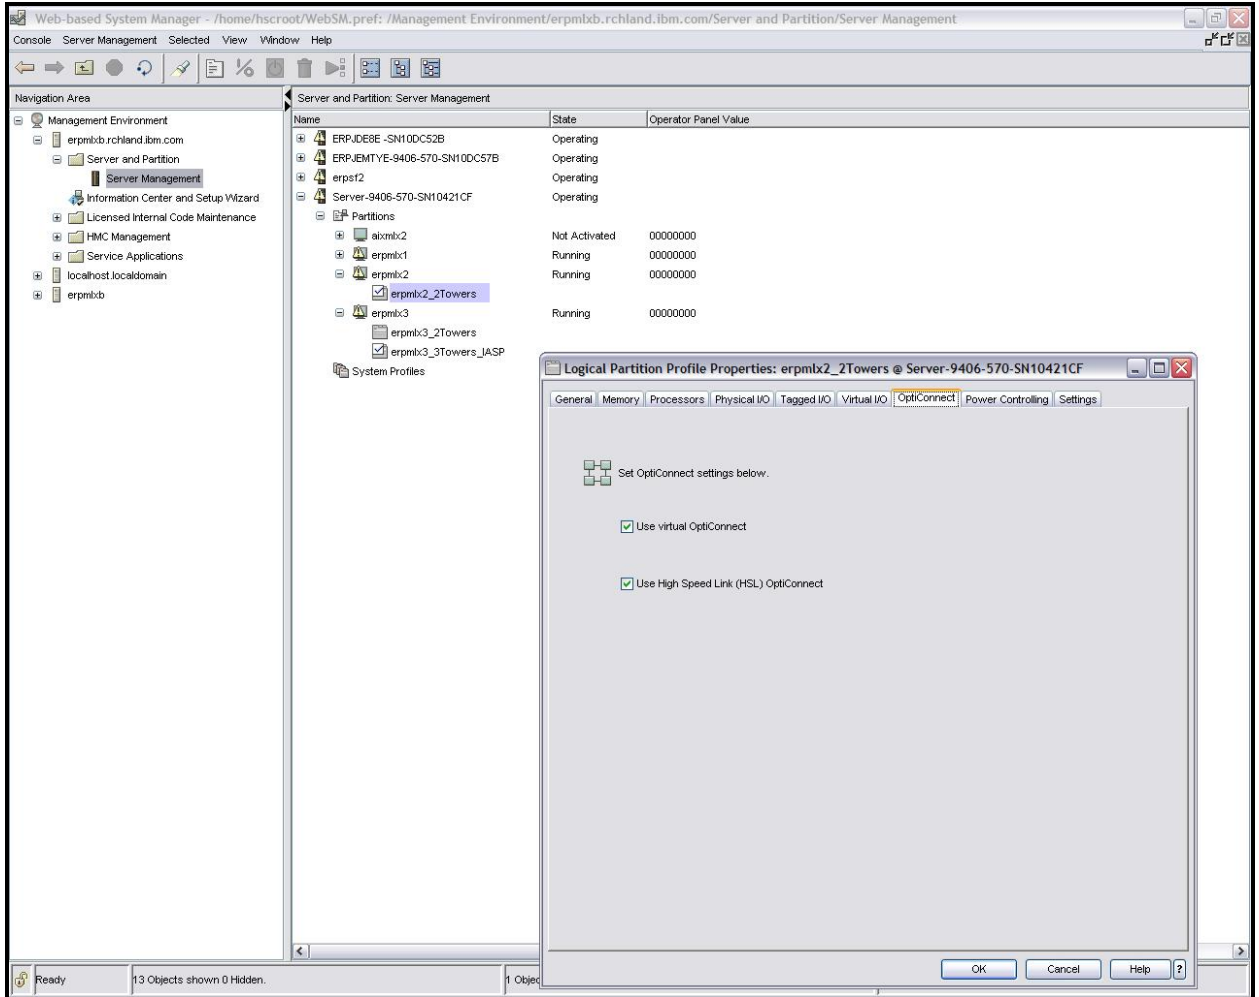


Figure 6. HMC configuration for backup node



## Creating the cluster and CRG

After the switchable hardware is identified and configured as shareable, the next step is to create the cluster. This can be done either with CL commands from a 5250 screen or through iSeries Navigator. The CL commands are provided below. For the iSeries Navigator procedure, refer to the Redbook entitled *IBM eServer iSeries Independent ASPs: A Guide to Moving Applications to IASPs*.

1. From the primary node, create the cluster.  

```
CRTCLU CLUSTER(<cluname>) NODE((<node1> ('<IP addr1>')) +  
(<node2> ('<IP addr2>')) START(*NO)
```
2. From the primary node, start both the primary and backup cluster nodes.  

```
STRCLUNOD CLUSTER(<cluname>) NODE(<node1>)  
STRCLUNOD CLUSTER(<cluname>) NODE(<node2>)
```
3. Ensure that each cluster node is active.  

```
DSPCLUINF CLUSTER(<cluname>)
```
4. From the primary node, add each node to the device domain. This command automatically creates a device domain if one does not exist.  

```
ADDDEVDMNE CLUSTER(<cluname>) DEVDMN(<cludevname>) + NODE(<node1>)  
ADDDEVDMNE CLUSTER(<cluname>) DEVDMN(<cludevname>) + NODE(<node2>)
```
5. Ensure that each node has been added to the newly created device domain.  

```
DSPCLUINF CLUSTER(<cluname>)
```
6. On both nodes, create the IP address to be associated with the IASP.
  - a. For CFGTCP option 1, add the address with the line description set to \*VIRTUALIP.
  - b. For CFGTCP option 10, add the new address to the host table entries.
7. On both nodes, create a device description for the IASP.  

```
CRTDEVASP DEVD(<devdesc>) RSRNAME(<devdesc>) TEXT('World iasp')
```
8. From the primary node, create the device cluster resource group (CRG).  

```
CRTCRCG CLUSTER(<cluname>) CRG(<crgrname>) CRGTYPE(*DEV) +  
EXITPGM(*NONE) USRPRF(*NONE) RCYDMN((<node1> *PRIMARY) + (<node2>  
*BACKUP 1)) CFGOBJ((<devdesc> *DEV *ONLINE '<IASP ip + addr>'))  
TEXT('Device CRG')
```
9. From the primary node, start the CRG and ensure that the CRG status is active.  

```
STRCRG CLUSTER(<cluname>) CRG(<crgrname>)
```
10. Check the CRG status.  

```
DSPCRGINF CLUSTER(<cluname>)
```



## Configuring the disk for the IASP

After the cluster and cluster resource group are configured, it is necessary to allocate disk units to the IASP. The configuration of the hardware that uses HMC must have already taken place. Additionally, the disk units must be logically attached to the host that is acting as the primary cluster node. This step can only be done through iSeries Navigator.

1. From an iSeries Navigator window, expand **My Connections > <primary node> > Configuration and Service > Hardware > Disk Units**.
2. Right-click *Disk Pools*.
3. Select *New Disk Pool...* and click *Next*.
4. Enter the following information.

Type of disk pool:	primary
Name of disk pool:	<devdesc>
Database:	generate by system
5. Click *OK*.
6. In the *Add disk units to these disk pools* pane, the newly created disk pool is now listed. (If it is not there, add the new disk pool to this pane and click *Next*.)
7. Click *Add Parity-Protected Disk...* (This ensures that only RAID and mirrored disks are displayed.)
8. Highlight all disks on the switchable bus and click *Add*.
9. Click *Next*.
10. Click *Finish*.
11. Ignore the message that disk units have been configured. Click *Continue*.
12. Ignore the warning that the first independent disk pool is being created in a device domain group. Click *Continue*.
13. Click *OK*.
14. Click *OK*.



## Moving JD Edwards World libraries to the IASP

After configuring the IASP, the next step is to move the JD Edwards World application libraries from the System ASP to the IASP. (**Note:** A full system save is highly recommended before moving objects to the IASP.)

### Identifying libraries used in the JD Edwards World environment

One way to move the libraries to the IASP is to check the library lists of the JD Edwards World user IDs and note those that are specific to the JD Edwards World application. For this project, the user IDs were JDE and JDEPROD. Note their library lists in Figure 6 and 7.

```
Display Library List
System: ERPMLX3
Job: QPADEV0007 User: JDE Number: 049089

Type options, press Enter.
5=Display objects in library

Opt Library Type ASP Device Text
QSYS SYS System Library
QSYS2 SYS System Library for CPI's
QHLPSYS SYS
QUSRSYS SYS System Library for Users
QTEMP USR
JDFOBJ USR JDE Object Library
JDFDATA USR JDE Data Library
JDFSRC USR JDE Source Library
QGPL USR General Purpose Library

Bottom
F3=Exit F12=Cancel F16=Job menu F17=Top F18=Bottom
(C) COPYRIGHT IBM CORP. 1980, 2005.
```

Figure 7. JDE library list

```

                                Display Library List
                                System:      ERPMLX3
Job:   QPADEV0007      User:   JDEPROD      Number:   049090

Type options, press Enter.
  5=Display objects in library

Opt  Library      Type      ASP
      Device      Text
-----
QSYS      SYS      System Library
QSYS2     SYS      System Library for CPI's
QHLPSYS   SYS
QUSRSYS   SYS      System Library for Users
QTEMP     USR
JDFOBJ    USR      JDE Object Library
CLTSEC    USR      Client's Security Library
CLTCOM    USR      Client's Common Library
CLTDTA    USR      Client's Data Library
JDFSRC    USR      JDE Source Library
QGPL      USR      General Purpose Library

Bottom
F3=Exit   F12=Cancel   F16=Job menu   F17=Top   F18=Bottom
(C) COPYRIGHT IBM CORP. 1980, 2005.

```

Figure 8. JDEPROD library list

The JD Edwards World libraries that were moved to the IASP are listed in the following table.

JD Edwards World libraries
JDFDATA
JDFOBJ
JDFSRC
CLTCOM
CLTDTA
CLTSEC

**Note:** When moving database libraries that are journaled to the IASP, the journals and journal receivers for these libraries must also be moved. This requirement exists because it is not possible to journal a library in the IASP to (or from) a journal in the system ASP. Both a library and its journal must be in either the system ASP or the IASP.

Be aware that the JD Edwards World JDFINS and JDEINSTAL install libraries were not moved to the IASP. These two libraries are only used during the JD Edwards World installation process. In addition, the JDFINS library contains several objects that are not supported in an IASP (\*JOBQ and \*OUTQ). For more information on objects that are supported in an IASP, reference the Redbook entitled *IBM eServer iSeries Independent ASPs: A Guide to Moving Applications to IASPs*.



## Saving libraries to save files

Create a save file for each of the identified JD Edwards World libraries that will be moved to the IASP. After the save files are created, save the libraries into them. Use the following CL commands to do this:

```
CRTSAVF FILE(<libname>/<filename>)  
SAVLIB LIB(<lib to be saved>) DEV(*SAVF) SAVF(<libname>/<filename>) +  
ACCPH(*YES)
```

## Removing libraries in the system ASP

Before continuing, ensure that all the JD Edwards World libraries to be moved to the IASP have been successfully put into save files. It is also recommended that a full system save be completed prior to beginning the IASP work. If that has not been done, perform a save before any libraries are deleted.

Deleting or renaming all the original JD Edwards World libraries in the system ASP is required because the same library name cannot exist in both the system ASP and the independent ASP concurrently. Before restoring any libraries into the independent ASP, they must first be removed from the system ASP.

## Restoring libraries to the IASP

After all the libraries have been removed from the system ASP, they can be restored to the IASP from the save files previously created. To restore the libraries to the IASP, use the following command:

```
RSTLIB SAVLIB(<lib to be restored>) DEV(*SAVF) +  
SAVF(<libname>/<filename>) RSTASPDEV(<IASP name>)
```

## Configuring JD Edwards World user profiles

Now that the JD Edwards World libraries have been moved to the IASP, the JD Edwards World user profiles need to be updated so that they connect to the library name space corresponding to the IASP. To ensure that a user will connect to the IASP library name space at login, the job description associated with the user profile needs to point to the IASP.

**Note:** All JD Edwards World user profiles must exist on both the primary and backup system nodes. The instructions in this section of this paper must be completed for user profiles on both nodes. Additionally, it is important not to change the QDFTJOB job description (located in QGPL).

## Creating user profiles

The JD Edwards World JDE and JDEPROD user profiles were created on the primary system during the JD Edwards World application install. They need to be created manually on the backup node.

When the user profiles exist, changes made to these objects on one node can be propagated to the other nodes automatically if the changed resource is set up correctly in the cluster administration domain. One example of where this can be used, is in keeping passwords for user profiles synchronized across nodes. For more information on using a customer administration domain, reference the link available in "Appendix C."

## Creating job descriptions

In the case of the system used for this project, the JD Edwards World JDE and JDEPROD user profiles were created where they both utilize the QDFTJOBDB job description (located in QGPL).

1. To find the job description used by the JD Edwards World user profiles on your system, use the following command:

```
DSPUSRPRF USRPRF(<user profile>)
```

The *Display User Profile – Basic* screen appears. Page down until the line labeled *Job description* appears (see Figure 8).

```

                                Display User Profile - Basic

User profile . . . . . : JDE

Storage information:
  Maximum storage allowed . . . . . : *NOMAX
  Storage used . . . . . : 1835128
  Storage used on independent ASP . . . . . : *NO
  Highest scheduling priority . . . . . : 3
  Job description . . . . . : QDFTJOBDB
  Library . . . . . : QGPL
  Accounting code . . . . . :
  Message queue . . . . . : JDE
  Library . . . . . : QUSRSYS
  Message queue delivery . . . . . : *NOTIFY
  Message queue severity . . . . . : 00
  Output queue . . . . . : *DEV
  Library . . . . . :
  Printer device . . . . . : *WRKSTN

More...
Press Enter to continue.

F3=Exit  F9=Display independent ASP information  F12=Cancel

```

Figure 9. JDE user profile

2. It is recommended that a copy be made of all the current job descriptions used by the JD Edwards World users on the system. To copy a job description, use the following command:

```
CRTDUPOBJ OBJ(<jobname>) FROMLIB(<libname>) OBJTYPE(*JOBDB) +
NEWOBJ(<new jobname>)
```

**Note:** The new job descriptions need to be created on both the primary and the backup nodes. These job descriptions must be in a library in the system ASP, not the IASP.

3. After the new job description has been created, it must be edited to default to the IASP. Use the following command to update the new job description:

```
CHGJOBDB JOBDB(<libname>/<new jobname>) INLASPGRP(<IASPname>)
```

## Changing user profiles to access new job descriptions

It is necessary to change all JD Edwards World user profiles to make use of the new job descriptions. With this update to the user profiles, they will have access to the IASP name space at login time.

1. To edit JD Edwards World user profiles, use the following command. (**Note:** Edit existing JD Edwards World user profiles to pick up the new job descriptions, or use the JD Edwards World interface to create new user profiles.)

```
CHGUSRPRF USRPRF(<World usrprf>) JOBD(<libname>/<new jobdname>)
```

## Switching between nodes

After setting up, configuring and moving the JD Edwards World application from the system ASP to an IASP, the actual switch from a primary node system to a backup node system was tested.

### Planned IASP switch

The basic steps involved in making a planned or manual switch are as follows:

1. All JD Edwards World users need to sign off the primary node system.
2. Ensure that the cluster and CRG are both active by typing the following two commands:  

```
DSPCLUINF
DSPCRGINF CLUSTER(<cluname>)
```
3. Vary off the IASP on the primary host by typing the following command.  

```
WRKCFGSTS CFGTYPE(*DEV) CFGD(*ASP)
```
4. Change the CRG from the primary node to the secondary node by typing the following command:  

```
CHGCRGPRI CLUSTER(<cluname>) CRG(<crgrname>)
```
5. Ensure that the secondary node is listed as the primary node by typing the following command:  

```
DSPCRGINF CLUSTER(<cluname>)
```
6. Vary on the IASP on the secondary node.
7. Activate the switchable IP address on the secondary node.  

```
CFGTCP option 1. Start the IP address.
```
8. JD Edwards World users can then sign on to the backup node system.

**Note:** Depending on the CRG configuration, Steps 3, 6 and 7 might be unnecessary.

For more explanation on the steps to make the switch over, reference the Redbook entitled *IBM eServer iSeries Independent ASPs: A Guide to Moving Applications to IASPs* (section 5.2 “Disk pool operation”).

### Unplanned IASP switch

A system failure or other major outage forces an unplanned switch of the IASP. The system and cluster code handle this code in the same manner as for a planned switch, although added delays do result.

### Performance considerations

Several factors affect the time to switch the IASP from one node to the other. The biggest consideration for the JD Edwards World environment is ensuring user profiles that own objects in the IASP have the same user and group ID numbers (UIDs and GIDs) on both the primary and backup node systems. The easiest way to synchronize UIDs and GIDs is to use iSeries Navigator. For more information, refer to the





IBM Information Center topic (*Maintain user profiles on all nodes*). See “Appendix C” for this Web site listing.

## Summary

---

The process of configuring and moving the JD Edwards World application into an IASP environment is relatively straightforward. As outlined in this paper, there are three main steps:

1. Configure the cluster and IASP.
2. Move the JD Edwards World libraries to the IASP.
3. Configure the JD Edwards World user profiles to access the IASP.

After the JD Edwards World application is up and running in the IASP, the JD Edwards World users should be unaware of the change.

To switch the IASP from one node to another is also a simple process.

1. All JD Edwards World users sign off the primary partition.
2. The IASP is varied off on the primary node.
3. The CRG is changed from the primary to the secondary node.
4. The IASP is varied on for the backup node.

After these steps are done, the JD Edwards World users will be able to sign on to the backup system and continue working.



## Appendix A: Hardware and software tested

---

The hardware and software levels used for this white paper did not play a major role in the testing. When possible, the most current levels were used. However, the System i model and the amount of memory and disk drives used were based on availability of resources. There is every reason to believe the same process of moving JD Edwards World to other supported levels of hardware and software works as described here.

For proper planning and sizing of your IASP, use the resources listed in “Appendix C.”

### Hardware:

IBM eServer™ iSeries Model 570:	2-way partition, 6 gigabytes of memory
System ASP:	90 #6718 17-GB disk units, RAID 5 protection
IASP:	45 #4326 35-GB disk units, RAID 5 protection

### Software:

JD Edwards A7.3 Service Pack 15  
i5/OS V5R4  
iSeries Navigator V5R4

## Appendix B: Useful IASP commands

---

Below is a list of i5/OS CL commands that are commonly used when working with IASP and clustered environments.

<b>GO CMDCLU</b>	Cluster commands
<b>GO CMDCRG</b>	Cluster resource group commands
<b>DSPCLUINF</b>	Display cluster information
<b>DSPCRGINF</b>	Display cluster resource group information
<b>SETASGRP</b>	Make IASP group available to current job
<b>CHGCRGPRI</b>	Switch cluster resource group to secondary system
<b>WRKCFGSTS *DEV *ASP</b>	View the current status of the IASP
<b>WRKDSKSTS</b>	View IASP disk activity and location of IASP disk storage



## Appendix C: Resources

---

These Web sites provide useful references to supplement the information contained in this document:

- IBM System i Information Center  
<http://publib.boulder.ibm.com/series/>
  - Maintaining user profiles within a clustered environment  
<http://publib.boulder.ibm.com/infocenter/series/v5r4/index.jsp?topic=/rzaig/rzaigmanage/maintainuserprofiles.htm>
  - System resources managed by a cluster administrative domain  
<http://publib.boulder.ibm.com/infocenter/series/v5r4/index.jsp?topic=/rzaig/rzaigmanage/maintainuserprofiles.htm>
  - Making hardware switchable  
<http://publib.boulder.ibm.com/infocenter/series/v5r4/index.jsp?topic=/rzaly/rzalymakehardwareswitchable.htm>
  
- IBM Publications Center  
[www.elink.ibm.com/public/applications/publications/cgi-bin/pbi.cgi?CTY=US](http://www.elink.ibm.com/public/applications/publications/cgi-bin/pbi.cgi?CTY=US)
  
- IBM Redbooks™  
[www.redbooks.ibm.com/](http://www.redbooks.ibm.com/)
  - IBM eServer iSeries Independent ASPs: A Guide to Moving Applications to IASPs  
[www.redbooks.ibm.com/abstracts/sg246802.html](http://www.redbooks.ibm.com/abstracts/sg246802.html)
  
- System i high availability and clusters  
[ibm.com/servers/eserver/series/ha/](http://ibm.com/servers/eserver/series/ha/)
  
- IBM white paper: "Implementing PeopleSoft EnterpriseOne ERP 8.0 Using an Independent Auxiliary Storage Pool"  
[ibm.com/support/docview.wss?uid=tss1wp100437](http://ibm.com/support/docview.wss?uid=tss1wp100437)

## Appendix D: About the author

---

### **Kathy Mack**

System i ERP Development – Oracle Support  
IBM System and Technology Group, Development

Kathy Mack works for IBM in the Rochester development lab providing support for Oracle JD Edwards World and JD Edwards EnterpriseOne products. She has more than 17 years of development experience.



## Trademarks and special notices

---

© Copyright. IBM Corporation 1994-2006. All rights reserved.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

eServer, i5/OS, IBM, the IBM logo, iSeries, Redbooks, and System i are trademarks or registered trademarks of International Business Machines Corporation 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.

Other company, product, or 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 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.

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.