IBM Systems Group
May 2004

Implementing PeopleSoft EnterpriseOne ERP 8.0 Using an Independent Auxiliary Storage Pool

By Marie Gann
Software Engineer
iSeries ERP Development – PeopleSoft Support

IBM

PeopleSoft

Table of contents

1. Introduction

High availability solutions come in various shapes and forms. One of the most straightforward to implement and to use is generically called a switched disk or switchable device topology. In the IBM® eServer™ iSeries™ world, we refer to this solution type as switchable IASPs (Independent Auxiliary Storage Pools. This paper, though technical in nature, demonstrates how straightforward it is to set up and operate a switched disk cluster topology for a PeopleSoft® EnterpriseOne 8.0 environment. Customers and consultants alike should find this paper both interesting and informative.

1.1 Hardware environment architecture

The architecture of the hardware environment tested and deployed in this paper consists of two iSeries servers and one DASD tower, all of which were connected in an HSL loop configuration. The two iSeries servers were configured in a two-node cluster with one system functioning as the primary node and the other as the backup node (see Figure 1). The DASD tower was configured to serve as a switchable IASP, to be used by either system. A cluster resource group (CRG) was configured between the two nodes and used to associate an IP address with the switchable IASP (see Figure 1). This IP address is called a takeover IP because, when control of the IASP is switched from one node to the other, the IP address is switched as well. In other words, the IP address is taken over by whichever node is controlling the IASP, hence the name takeover IP.

IMPORTANT NOTE: For this project, the numeric IP address used to configure the takeover IP interface was resolved to a specific name using the network DNS. When the takeover IP is referred to throughout this document, it is the IP address name (not the numeric IP address) that is being referenced.

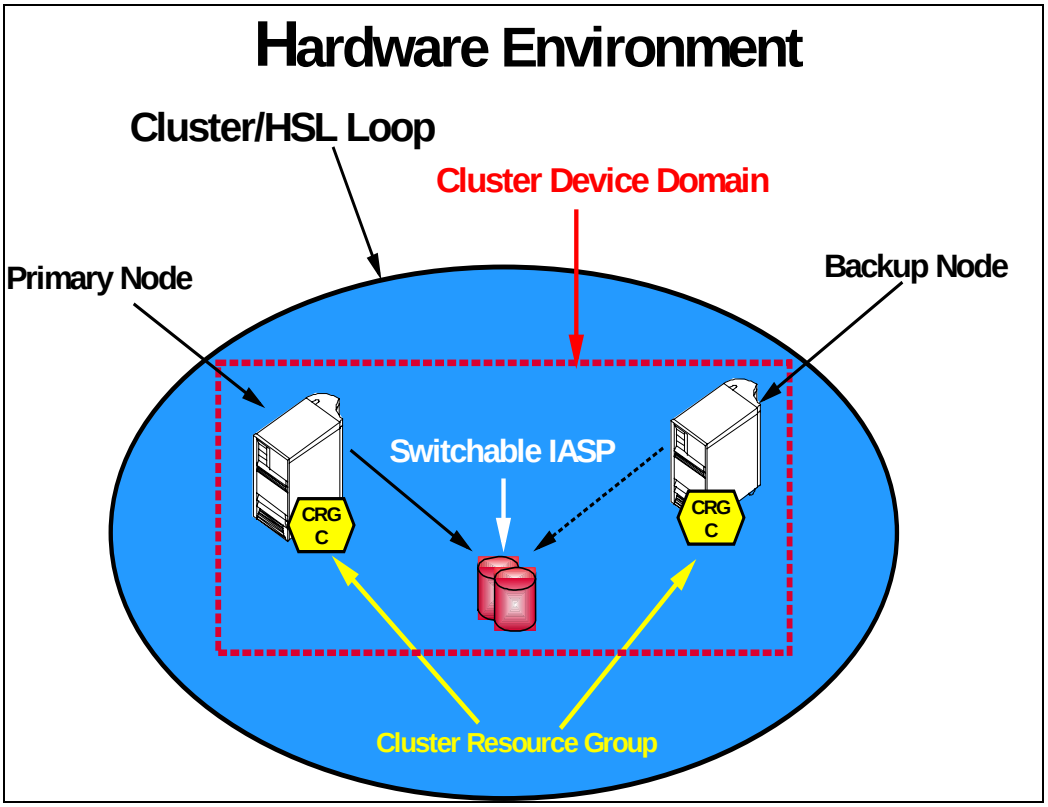

Figure 1

1.2 Software environment architecture

The PeopleSoft EnterpriseOne ERP 8.0 software environment architecture used in this project consisted of an application server, a database and a Java™ Application Server (JAS). The application server and the database were first installed on the primary node. Once the installation was completed, the following list of objects was moved from the system ASP of the primary node to the IASP.

1. All PeopleSoft EnterpriseOne ERP 8.0 application server directories in the IFS
2. All PeopleSoft EnterpriseOne ERP 8.0 application server libraries, ***except the B7334SYS library***
3. All PeopleSoft EnterpriseOne ERP 8.0 database libraries

Since the IASP was configured to be switchable, all the libraries and directories in the IASP moved from one iSeries server to the other when control of the IASP was switched from one node to the other. (See Figure 2.)

A copy of the application server library B7334SYS, which remained in the system ASP of the primary node, was placed in the system ASP of the backup node. Therefore, PeopleSoft EnterpriseOne services could be started on whichever node was currently controlling the IASP. (See Figure 2.)

NOTE: All PeopleSoft EnterpriseOne user profiles that exist on the primary node system must also be created on the backup node system. This means any user profile that uses the PeopleSoft EnterpriseOne application for any reason must exist on both the primary and backup system nodes.

The Java Application Server location will not be mentioned because this configuration will work as long as the following two conditions are true.

1. The html clients are able to establish a network connection to the Java Application Server
2. The Java Application Server is able to establish a network connection to the PeopleSoft EnterpriseOne Application Server by using the takeover IP

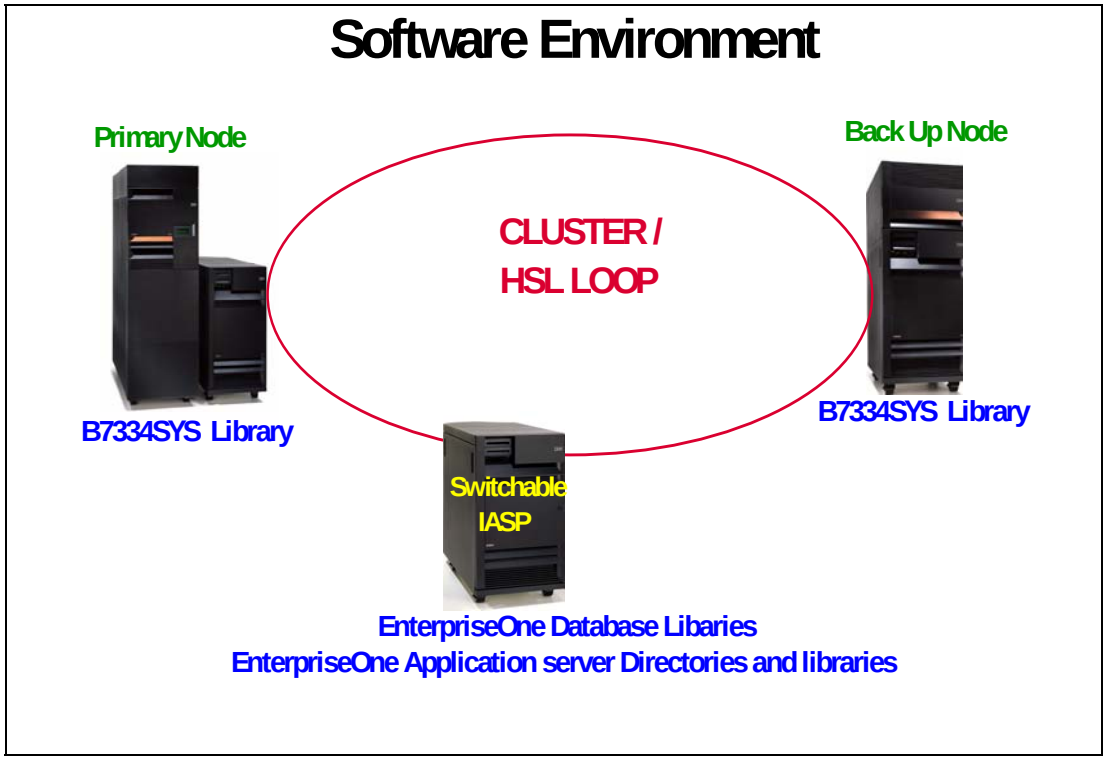

Figure 2

2. Configuring the IASP for PeopleSoft EnterpriseOne

The focus of this document is on the PeopleSoft EnterpriseOne application and the steps needed to implement it in a switchable IASP environment. Therefore, the steps necessary to create the cluster, the IASP, and the CRG are not documented here. With that said, there is an important IASP configuration issue that must be addressed to insure that the PeopleSoft EnterpriseOne application will function correctly. The issue is as follows:

The name of the Relational Database Directory Entry associated with the IASP must match the name that is used for the takeover IP.

If you are not familiar with the concept of Relational Database Directory (RDB) entries and associating them to an IASP, see Appendix A.

Creating the IASP with the same name as the takeover IP name is an easy way to insure that its RDB entry name is the same as the takeover IP name. This is because, when an IASP is created, the system adds an RDB entry for the IASP with the same name as the IASP. However, if the IASP has already been created with a name that does not match the name of the takeover IP, a configuration change will need to be made. There are two possible configuration changes that will accomplish the goal of making the name of the takeover IP match the name of the RDB entry associated with the IASP.

1. The name of the takeover IP can be changed to match the name of the RDB entry associated with the IASP. To accomplish this, contact your network administrator.
2. The name of the RDB entry associated with the IASP can be changed to match the name of the takeover IP. To accomplish this, follow the two-step directions below:
   - First, create an RDB entry with the same name as the takeover IP. To create a new RDB entry, from the green screen, enter the following command:
     ADDRDBDIRE RDB(name of takeover IP**) RMTLOCNAME(LOOPBACK *IP)**
   - Second, associate the RDB entry that was just created with the IASP. To do this, from the green screen, enter the following command:
     CHGDEVASP DEVD(name of IASP**) RDB(**name of takeover IP**)**

3. Moving PeopleSoft EnterpriseOne libraries to the IASP

A full system save to tape should be done before attempting the instructions given in this chapter.

The PeopleSoft EnterpriseOne environment used in this project consisted of the libraries listed in table 1.

LIBRARY NAMES
B7334SYS
PD7334
PD7334FA
SYS7334
SVM7334
COPD7334
PD7334DNT
OL7334
DD7334
OWJRNL
PRODDTA
PRODCTL

Table 1

All of the Libraries in table 1, ***with the exception of the B7334SYS library***, were moved to the IASP. ***The B7334SYS library was not moved to the IASP because it contains objects that are not supported in the IASP***.

NOTE: When moving database libraries that are journaled to the IASP, the journals and journal receivers for these libraries must be moved as well. This is because it is not possible to journal a library from the system ASP to a journal in the IASP and vice versa. That is, a library and its journal must both be either in the system ASP or the IASP.

Moving the PeopleSoft EnterpriseOne libraries to the IASP is done by following the three steps below:

1. Save the libraries to save files.
2. Delete the libraries from system space.
3. Restore the libraries to the IASP.

3.1 Saving libraries into save files

To create a save file for each of the libraries in table 1, ***except the B7334SYS library***, and save each library into the appropriate save file, use the following commands from the green screen:

CRTSAVF FILE(Name of library to containing save file / Save File Name**)**
SAVLIB LIB(Name of library being saved **) DEV(*SAVF) +**
SAVF(Name of library containing save file /Save File Name**) ACCPTH(*YES) + DTACPR(*YES)**

3.2 Deleting libraries from the system ASP

Before proceeding with this section, doublecheck that all the libraries in the tables on the previous page have been fully saved into save files. Also it is recommended that if a full system save was not completed before beginning, one should be done now.

Deleting all the original libraries from the system ASP is necessary because the same library name cannot exist in the system ASP and the independent ASP at the same time. So, before any libraries can be restored to the independent ASP, they must be removed from the system ASP. Therefore, all the libraries in table 1 above, ***except the B7334SYS library***, must be deleted from the system ASP.

To delete the libraries from the system ASP, from the green screen, use the following command:

DLTLIB LIB(Name of library to be deleted**)**

3.3 Restoring DB libraries to the IASP

Once all the libraries have been removed from the system ASP, they can be restored to the IASP from the save files that were created in section 3.1 “Saving libraries into save files.” Therefore, all the libraries in table 1, ***except the B7334SYS library***, should be restored to the IASP.

To restore the libraries to the IASP, from the green screen, use the following command:

RSTLIB SAVLIB(Name of library being restored**) DEV(*SAVF) +**
SAVF(Name of library where save file is / Save File Name**) + RSTASP(**Name of the IASP**)**

4. Moving the PeopleSoft EnterpriseOne IFS directories to the IASP

When an IASP is created, an Integrated File System (IFS) directory is created in the system’s root directory with the same name as the IASP. Any IFS objects that resided in this directory are actually in the IASP’s IFS—not the system ASP’s IFS. So, moving IFS objects to the IASP is accomplished by simply copying them into the IASP’s IFS directory and then deleting them from the system ASP’s IFS. The PeopleSoft EnterpriseOne environment used in the project consisted of the IFS directories listed in table 2.

Directory Names
B7334SYS
Jdeb7334
PD7334
OneWorld

Table 2

All the IFS directories listed in table 2 were moved to the IASP directory. Moving directories to the IASP is a two-step process.

1. Copy the directories to the IASP's IFS directory and delete them from the system ASP's IFS.
2. Create symbolic links, in the system ASP's IFS on both nodes, pointing to the directories' new locations in the IASP's IFS directory.

4.1 Copying directories to the IASP

To copy the directories listed in table 2 to the IASP Directory, use iSeries Navigator—click the (+) plus sign beside **My Connections** to expand the list of systems that are available. Then click the (+) plus sign beside the name of the system where the IASP resides. Next, click the (+) plus sign beside **File Systems** to expand the directory list for that system. The iSeries Navigator screen should now look similar to the one in figure 3.

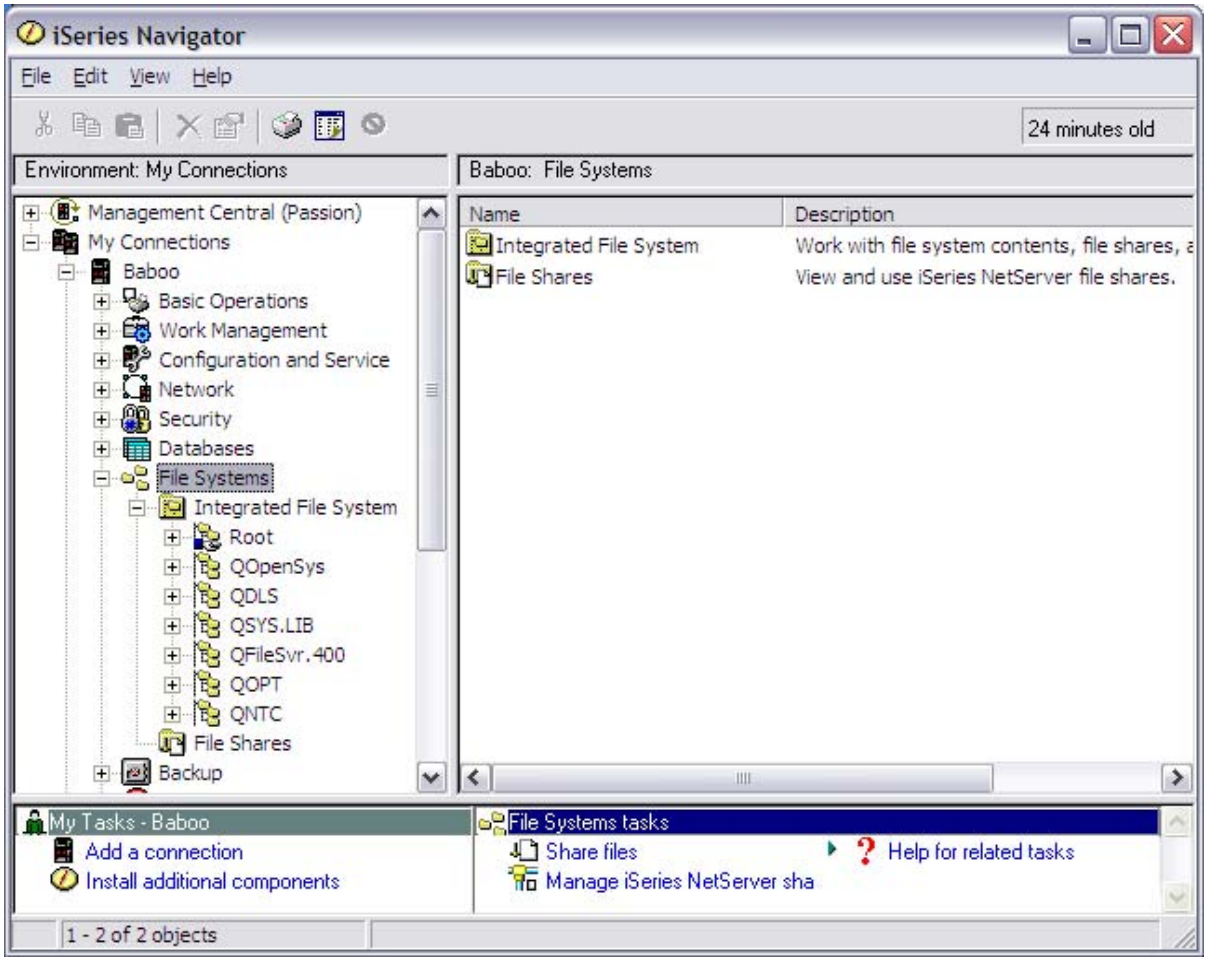

Figure 3

The **Root** directory is where the IFS directory for the IASP resides, so click the (+) plus sign beside the **Root** directory. Once all the directories under **Root** are visible, find the directory with the same name as the IASP (see figure 4). This is the directory into which all the PeopleSoft EnterpriseOne directories listed in table 2 must be copied. To do this, click on the name of the directory to be copied and use the copy and paste commands found on the edit menu. When the copying is finished, the directories listed in table 2 must be deleted from the system ASP's IFS, which is done by simply clicking on the name of the directory to be deleted and pressing the delete key.

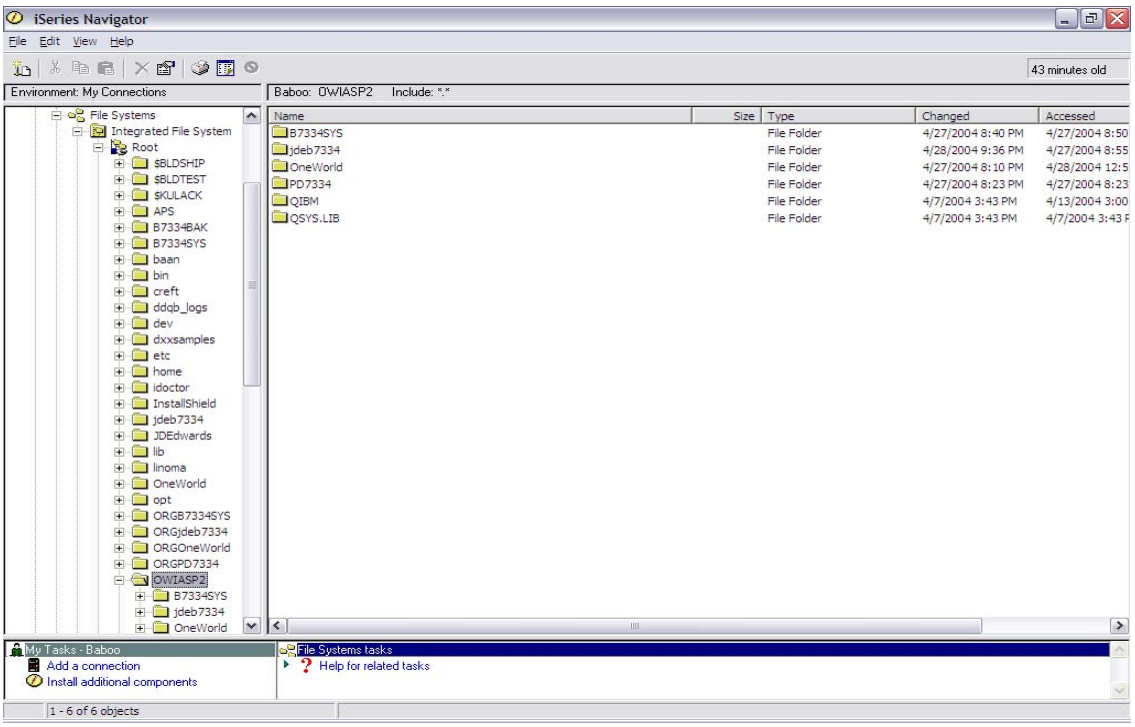

Figure 4

4.2 Creating symbolic links

A symbolic link must be created for every IFS directory that is moved from the system ASP's IFS to the IASP's IFS directory. A symbolic link with the same name as the directory is placed in the directory's original location and points to its new location in the IASP's IFS directory. This is done so that the PeopleSoft EnterpriseOne application will be able to find the directories in their new location in the IASP's IFS directory. To create a symbolic link for each of the directories moved to the IASP's IFS directory, from the green screen, use the following command:

ADDLNK OBJ('IASP Name/Directory Name'**) NEWLNK(**'Directory Name'**)**

A symbolic link must be added to the system ASP's IFS, ***on both the primary and back up node***, for each directory that was moved to the IASP's IFS directory. For a complete list of the PeopleSoft Enter-

priseOne directories that were moved to the IASP's IFS refer to table 2. To view a symbolic link that has been created, from the green screen, use the following command:

WRKLNK DETAIL(*EXTENDED)

A screen similar to figure 5 should appear.

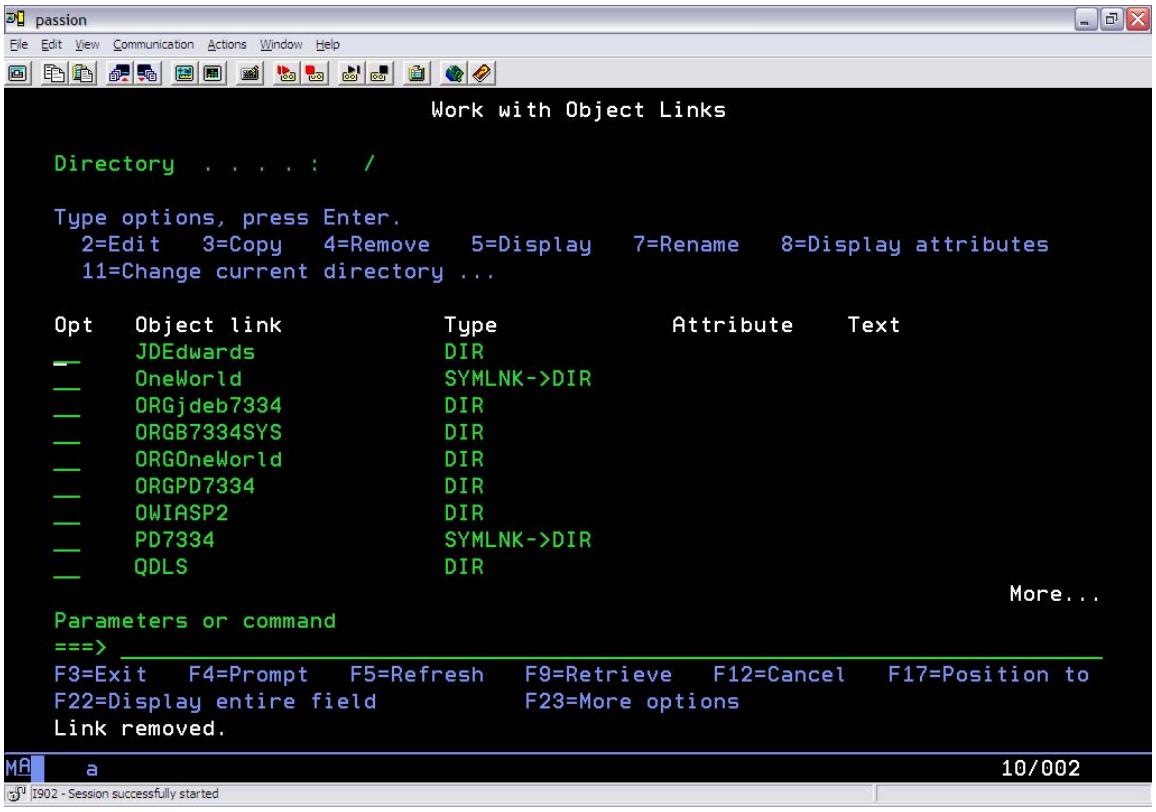

Figure 5

5. Configuring PeopleSoft EnterpriseOne data sources and OCM tables

NOTE: The following instructions were written for PeopleSoft EnterpriseOne ERP 8.0. It is possible that changes will be made to the libraries and tables mentioned here in later releases of PeopleSoft EnterpriseOne. If you are using a PeopleSoft EnterpriseOne release later than PeopleSoft EnterpriseOne ERP 8.0, consult PeopleSoft EnterpriseOne documentation before proceeding with these instructions.

The PeopleSoft EnterpriseOne data sources tables (named the F98611 tables) define how the application should access either data or servers, and are found in both the SYS7334 and SVM7334 libraries. These tables work in conjunction with the PeopleSoft EnterpriseOne Object Configuration Manager or OCM

tables. The OCM tables (named the F986101 tables) define where the database tables are located, where business functions run, and where UBEs run, and are also found in the SYS7334 and SVM7334 libraries.

Since these tables were installed on the primary node, they contain the primary node system's host name. In order for PeopleSoft EnterpriseOne to function correctly in a clustered IASP environment, these tables must be edited to reflect the takeover IP name. This is important so the mapping in these tables is not tied to either system's host name, but instead will work regardless of whichever system has control of the IASP. To achieve this, use the SQL commands given in the following four sections.

NOTE: The SQL commands given in this chapter were run using iSeries Navigator's SQL environment, the syntax must be changed for them to run correctly in other SQL environments.

5.1 Editing the F98611 table in the SVM7334 library
Takeover IP = name of takeover IP address.
PnodeName = host name of primary node system.

```
UPDATE svm7334.F98611 SET (OMDATP) = ('Takeover IP - B7334 Server Map') where OMDATP =
'PNodeName - B7334 Server Map'   ;

UPDATE  svm7334.F98611 SET (OMDATP) = ('Takeover IP') where OMDATP = 'PNodeName'   ;

UPDATE  svm7334.F98611 SET (OMSRVR) = ('Takeover IP') where OMSRVR = 'PNodeName'   ;

UPDATE  svm7334.F98611 SET (OMDATB) = ('Takeover IP - B7334 Server Map') where OMDATB =
'PNodeName - B7334 Server Map'   ;

UPDATE  svm7334.F98611 SET (OMLL) = ('Takeover IP') where OMLL = 'PNodeName'   ;
```

5.2 Editing the F986101 table in the SVM7334 library
Takeover IP = name of takeover IP address.
PnodeName = host name of primary node system.

```
UPDATE svm7334.F986101 SET (OMDATP) = ('Takeover IP') where OMDATP = 'PNodeName'   ;
```

5.3 Editing the F98611 table in the SYS7334 library
Takeover IP = name of takeover IP address.
PnodeName = host name of primary node system.

```
UPDATE sys7334.F98611 SET (OMDATP) = ('Takeover IP - B7334 Server Map') where OMDATP =
'PNodeName - B7334 Server Map'   ;
```

UPDATE sys7334.F98611 SET (OMDATP) = ('Takeover IP') where OMDATP = 'PNodeName' ;

UPDATE sys7334.F98611 SET (OMDATP) = ('Takeover IP - Logic') where OMDATP = 'PNodeName - Logic' ;

UPDATE sys7334.F98611 SET (OMSRVR) = ('Takeover IP') where OMSRVR = 'PNodeName' ;

UPDATE sys7334.F98611 SET (OMDATB) = ('Takeover IP - B7334 Server Map') where OMDATB = 'PNodeName - B7334 Server Map' ;

UPDATE sys7334.F98611 SET (OMLL) = ('Takeover IP') where OMLL = 'PNodeName' ;

5.4 Editing the F986101 table in the SYS7334 library
Takeover IP = name of takeover IP address.
PnodeName = host name of primary node system.

UPDATE sys7334.F986101 SET (OMDATP) = ('Takeover IP') where OMDATP = 'PNodeName' ;

UPDATE sys7334.F986101 SET (OMDATP) = ('Takeover IP - Logic') where OMDATP = 'PNodeName - Logic' ;

6. Configuring PeopleSoft EnterpriseOne user profiles

Since, most of the PeopleSoft EnterpriseOne objects have been moved to the IASP, we need to insure that the library name space being connected to by the PeopleSoft EnterpriseOne users corresponds to IASP. When a user logs on to an iSeries, by default they are connected to the system ASP library name space. To insure that a user, at log in, will connect to the IASP library name space, the user's job description must be edited to point to the IASP.

In the case of the PeopleSoft EnterpriseOne application all PeopleSoft EnterpriseOne user profiles on both nodes must have their job descriptions edited to point to the IASP. Meaning any profile, on either node, that uses the PeopleSoft EnterpriseOne application for any reason must have its job description edited to point to the IASP.

NOTE: Since all PeopleSoft EnterpriseOne user profiles must exist on both the primary and backup system nodes, the instructions in this chapter must be completed for both nodes.

6.1 Creating a job description
In the case of the systems used for this project, the PeopleSoft EnterpriseOne user profiles utilized a job description named ONEWORLD. A copy of the ONEWORLD job description was made and then edited.

To find the job description used by the PeopleSoft EnterpriseOne user profiles on your system, from the green screen, use the following command: **DSPUSRPRF USRPRF(**EnterpriseOne user name**)**

The **Display User Profile – Basic** screen will appear, see figure 6.

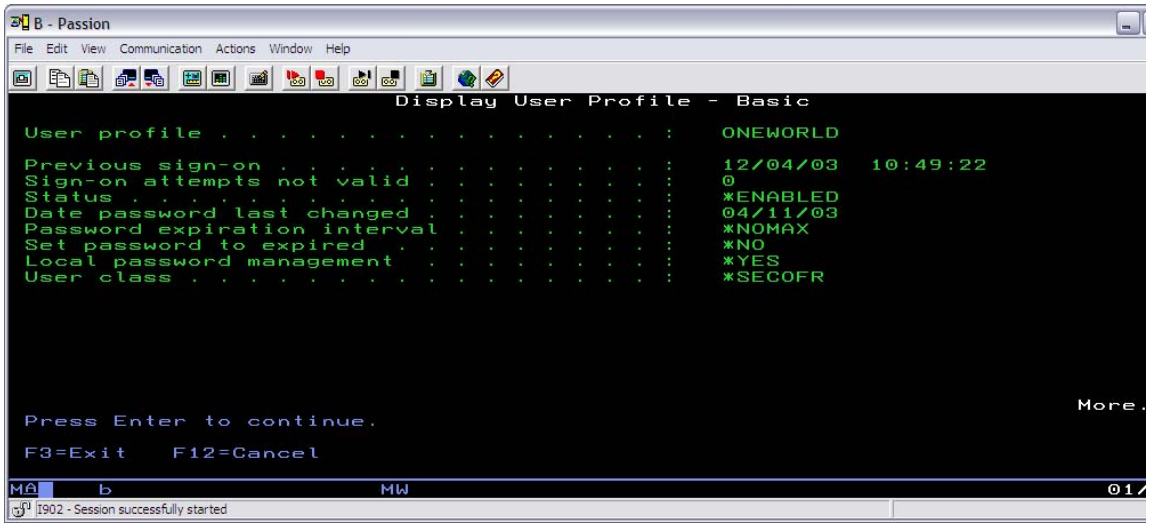

Figure 6

Press the <**Page Down**> key until you see the line labeled **Job description**, see figure 7. This is the job description that the PeopleSoft EnterpriseOne user profiles are using.

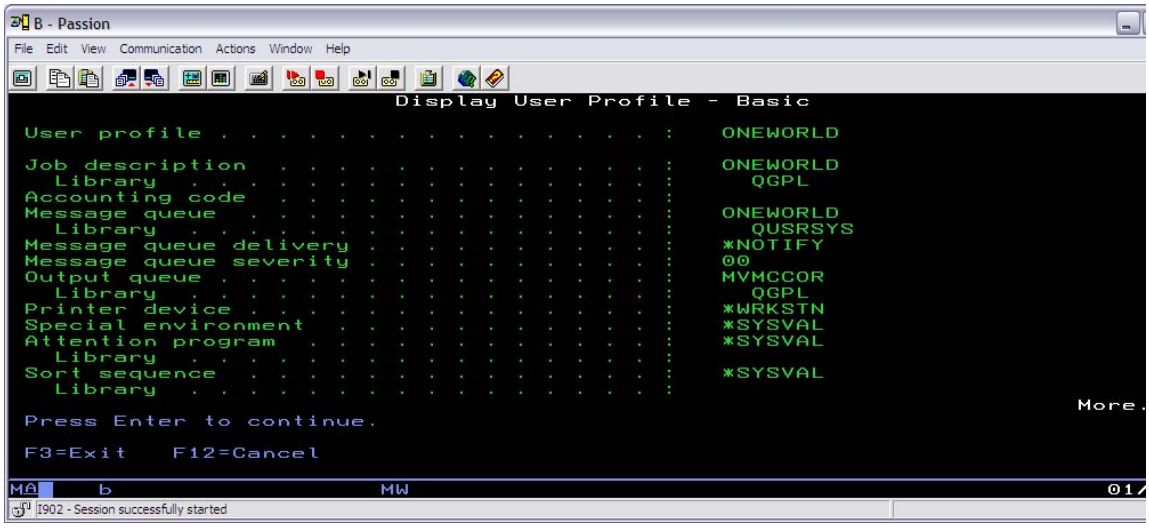

Figure 7

It is recommended that you make a copy of the job description that is currently used by the PeopleSoft EnterpriseOne users on your system. To copy the job description, from the green screen, use the following command:

CRTDUPOBJ OBJ(name of EnterpriseOne User's jobd**) FROMLIB(**library where jobd resides) **OBJTYPE(*JOBD) NEWOBJ(**new jobd name**)**

NOTE: Make sure the new job description is created, on both nodes, in a library in the system ASP, not the IASP. The Create Duplicate Object command used above will, by default, place the new object in the same library as the object being copied.

Once the new job description is created, it must be edited to default to the IASP. To do this, from the green screen, use the this command:

CHGJOBD JOBD(library where jobd resides / new jobd name**) INLASPGRP(**IASP name**)**

6.2 Adding the new job description to user profiles

Now that the job description has been created, ***all*** PeopleSoft EnterpriseOne user profiles, on both nodes, must be edited to use the new job description. In this way, the user profiles will default to the IASP when a user logs on. To edit the PeopleSoft EnterpriseOne user profiles, from the green screen use the following command.

CHGUSRPRF USRPRF(EnterpriseOne user name**) JOBD(**library where jobd resides / new jobd **)**

NOTE: Make sure that all PeopleSoft EnterpriseOne user profiles, on both nodes, are changed to use the new job description, which points to the IASP.

7. Editing the PeopleSoft EnterpriseOne enterprise server's JDE INI file

The PeopleSoft EnterpriseOne Application Server uses the settings in its JDE INI file to establish its connectivity to various database sources and servers. In order for PeopleSoft EnterpriseOne to function correctly in a clustered IASP environment the JDE INI file will need to be edited. This is to insure that PeopleSoft EnterpriseOne will establish its connections based on the takeover IP and not any system host names.

The following instruction in this chapter will need to be executed on the primary and backup node systems. This is because JDE INI resides in the B7334SYS library and a copy of this library is in the system ASP on both the primary and back up node systems.

To edit the Enterprise Server's JDE INI file, sign on to the system as ONEWORLD and enter the **SAW** command when the **OneWorld Server Administration Workbench** screen appears (figure 8), choose option 5 (**Work with Server INI file**).

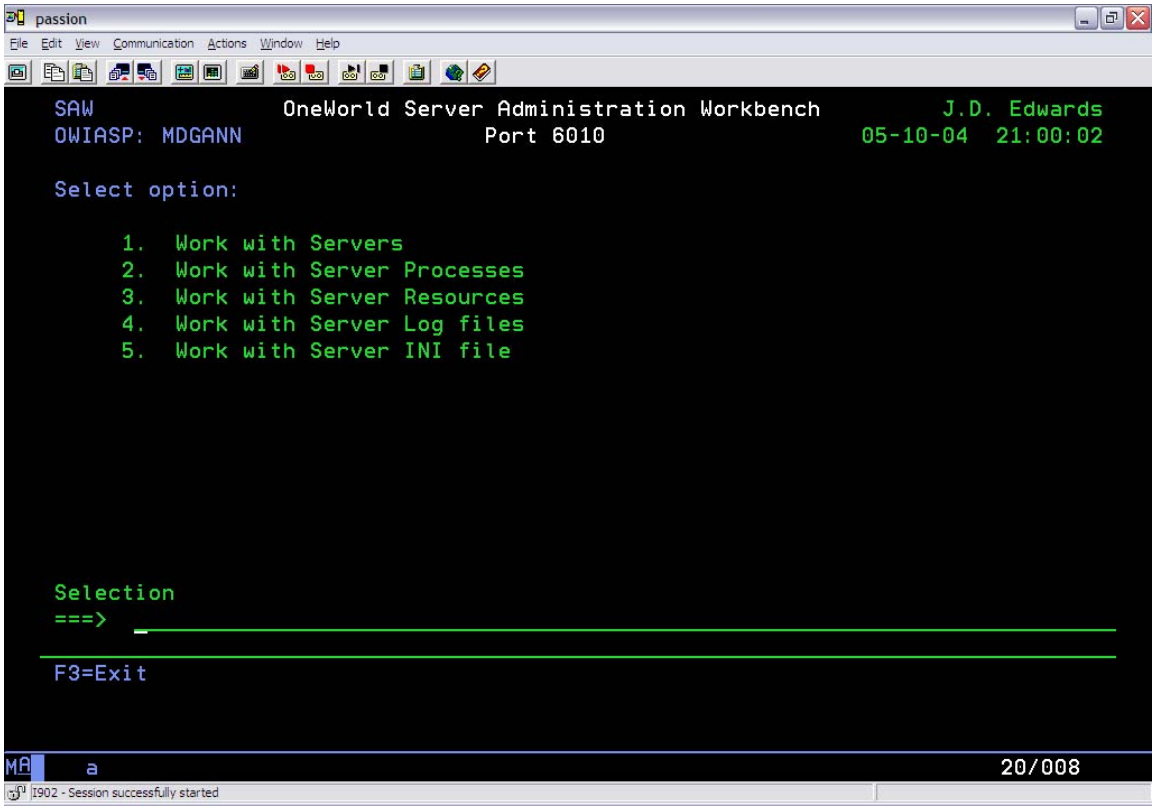

Figure 8

The following section labeled **[CLUSTER]** will need to be added to the JDE INI file. This section is only used for clustered environments and is not in the JDE INI file by default.

[CLUSTER]
PrimaryNode=Takeover IP

The following section labeled **[DB SYSTEM SETTINGS]** will already be in the JDE INI file but will need to have the settings that appear under it edited.

[DB SYSTEM SETTINGS]
Base Datasource=Takeover IP - B7334 Server Map
Server=Takeover IP
Database=Takeover IP - B7334 Server Map

The following section labeled **[SECURITY]** will also already be in the JDE INI file but will need to have the settings that appear under it edited.

[SECURITY]
SecurityServer=Takeover IP

After you have made all the above changes, be certain to save the JDE INI file.

8. Adding a host table entry for the takeover IP

The PeopleSoft EnterpriseOne Application Server uses TCPIP host table entries to establish connectivity when running in a clustered IASP environment. In order for PeopleSoft EnterpriseOne to function properly in a clustered IASP environment, ***host table entries for the takeover IP must be added to both the primary and backup node systems.***

To add a host table entry to the system, from the green screen, use the following command:

ADDTCPHTE INTNETADR('numeric take over IP address'**) HOSTNAME((**'takeover IP Name'**)**
('takeover IP Name.domain name.com'**))**

Remember to add a host table entry for the takeover IP on both the primary and backup systems.

9. Configuring PeopleSoft EnterpriseOne fat clients

In order for the fat client machines to work correctly in an IASP environment there are two configuration connections issues that need to be addressed. The fat client ODBC connections and the fat client network connections need to be configured to connect to the server using the takeover IP address, rather than a system host name. This is to insure that the fat clients will work regardless of whichever node is currently controlling IASP.

The instructions given in this chapter will need to be repeated for all PeopleSoft EnterpriseOne data sources that exist on all fat client machines.

9.1 Changing the ODBC connections

To change the ODBC settings on a fat client machine you need to launch the **ODBC Data Source Administrator**. This can be found from the **start menu/control panel/ administrative tools/data sources (odbc)**. Double-click on **data sources (odbc)** and the **ODBC Data Source Administrator** dialog box will be launched, see figure 9.

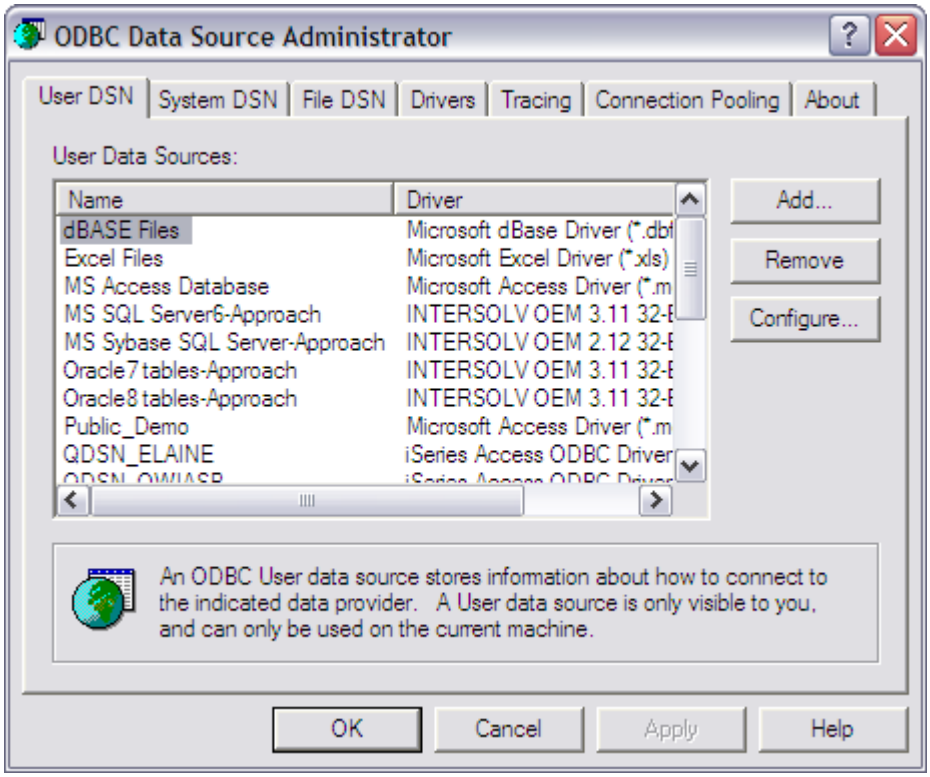

Figure 9

Once the **ODBC Data Source Administrator** dialog box is launched click on the **System DSN** tab, see figure 10.

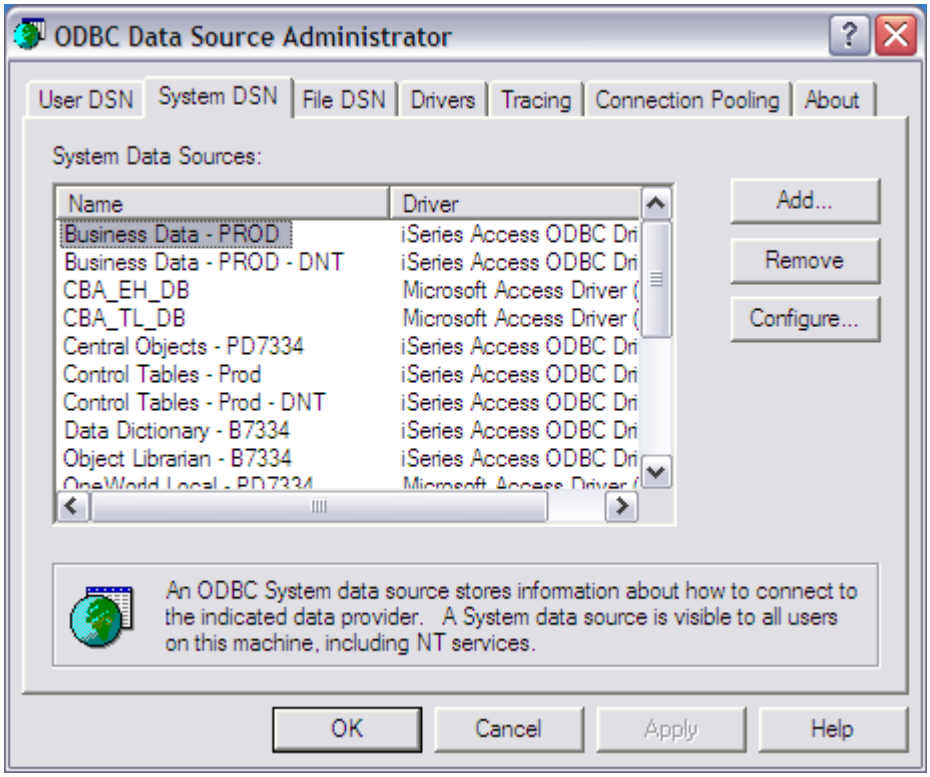

Figure 10

On the **System DSN** tab, click on the **Data Source Name** to highlight it, see figure 10. Then, click the **Configure** button so that the **iSeries Access for Windows ODBC Setup** dialog appears, see figure 11.

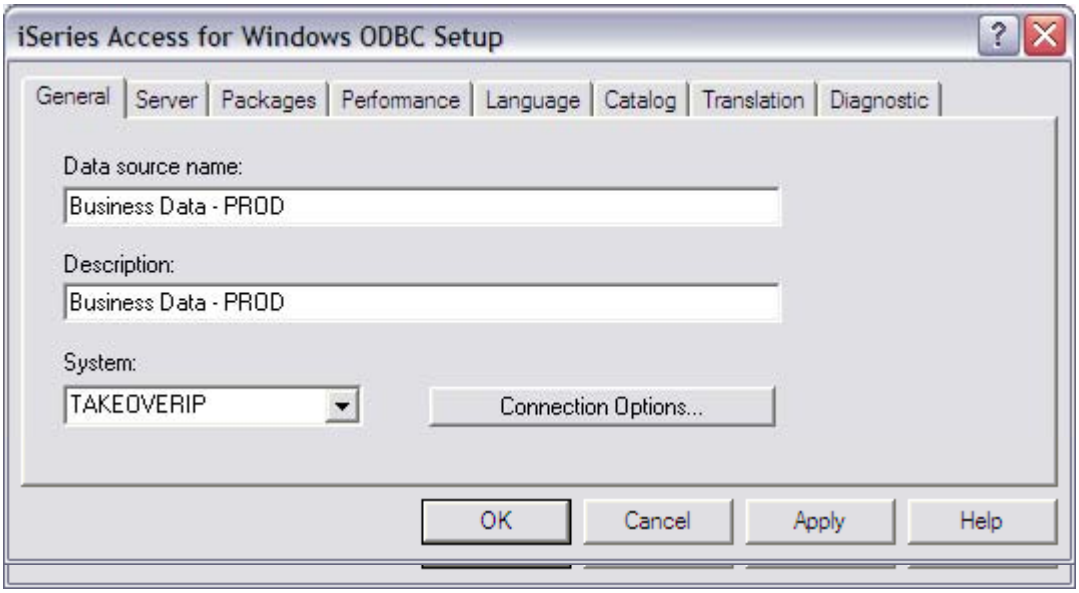

Figure 11

Once the **iSeries Access for Windows ODBC Setup** dialog appears, type the **takeover IP name** in the **System** pull down text box. Once the **takeover IP name** is typed, click the **Apply** button and then click the **OK** button to close the **iSeries Access for Windows ODBC Setup** dialog.

9.2 Network connections

To configure the fat client to make its network connections use the takeover IP address, edit the JDE.ini file on the fat client machine. To do this, locate the JDE.ini file on the fat client machine. It is usually found in the following directory path: **C:\WINDOWS\JDE.INI** .

Once the file is located, open it, using any text editor you prefer, and make the following changes.

In the **[DB SYSTEM SETTINGS]** section, change the following:
Server=TAKEOVERIP
Library List= TAKEOVERIP

In the **[SECURITY]** section, change the following:
SecurityServer=TAKEOVERIP

After you have made the above changes be sure to save the file.

10. Configuring the PeopleSoft EnterpriseOne Java Application Server

To insure that the PeopleSoft EnterpriseOne Java Application Server will function correctly in an IASP environment, only one configuration issue needs to be addressed. The JAS's network connections must be configured to connect to the server using the takeover IP rather than a system host name. This is to insure that JAS will function correctly regardless of whichever node is currently controlling the IASP.

To configure the Java Application Server to make network connections using the takeover IP, it is necessary to edit the **JAS.ini** file on the system where the Java Application Server resides.

To edit the JAS.INI file, from the green screen, use the following command:

wrklnk "/JDEdwards/JAS/EA_JDEdwards_1.ear/webclient.war/ WEB-INF/jas.ini"

Once the **Work With Object Link** screen appears (figure 12), type a **2** next to file name and press the <**Enter**> key. This will open the file in edit mode.

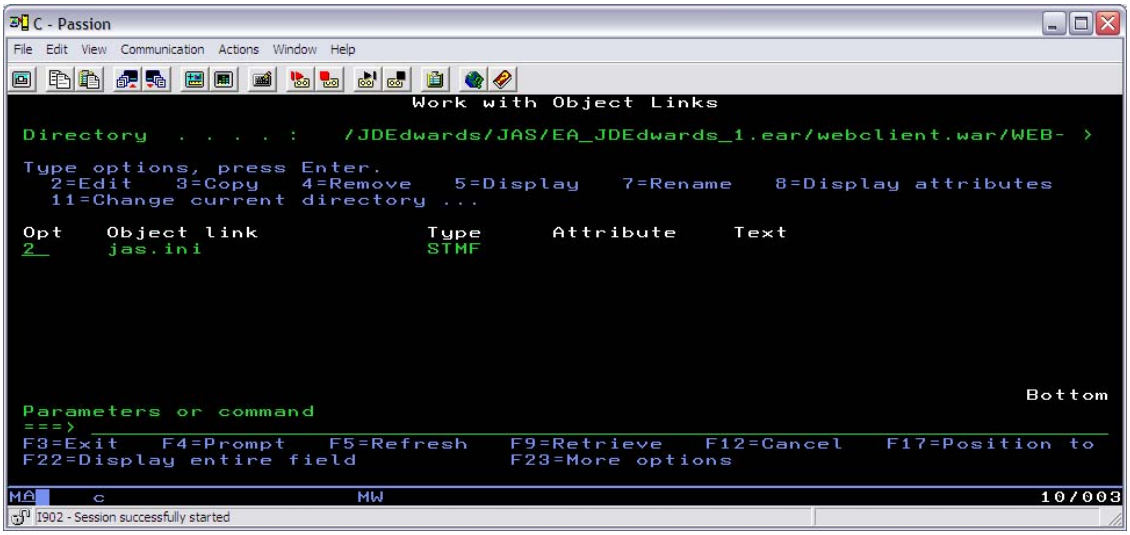

Figure 12

Once the jas.ini file is opened in edit mode, edit the following sections.

In the **[DB SYSTEM SETTINGS]** section, change the following:
Server=Takeover IP

In the **[SECURITY]** section change the following:
SecurityServer=Takeover IP

After you have made the above changes be sure and save the file by pressing the <F2> key.

11. Switching between nodes

NOTE: If running a PeopleSoft EnterpriseOne Java Application Server, you must stop the JAS instance and restart it after you have completed the steps outlined in this chapter. It is also recommended that you stop and restart the PeopleSoft EnterpriseOne HTTP server, as well.

This chapter outlines three basic steps necessary to switch from the primary node system to the backup node system.

1. Stop the PeopleSoft EnterpriseOne Application Sever on the primary node system.

   Log on to the primary node system as ONEWORLD and type the following command:
   endnet and press the <**Enter**> key

The PeopleSoft EnterpriseOne Application Server on this system will now shut down.

2. Switch control of the IASP from the primary system node to the backup system node. The steps to achieve this are documented in the following redbook: IBM eServer iSeries Independent ASPs: A Guide to Moving Applications to IASPs

3. Start the PeopleSoft EnterpriseOne Application Sever on the backup node system.

   Log on to backup node system as ONEWORLD and type the following command:
   strnet and press the <**Enter**> key

   The PeopleSoft EnterpriseOne Application Server on this system will now start.

Once you have completed the previous three steps, your PeopleSoft EnterpriseOne environment is fully operational. The fat clients and HTML clients will function normally, unaware that the PeopleSoft EnterpriseOne Application Server is running on the backup system.

11.1 Important switch over performance considerations

There are several factors that can affect how long it takes the IASP to switch from one node to the other. However, it was discovered during this project that there were two main issues that seemed to affect switchover time the most. The first was keeping the number of database objects in the systems ASP as low as possible. In fact it is recommended that, if possible, all database objects be placed in the IASP. The second was making sure that user profiles that own objects in the IASP have the same user id numbers (UIDs) and group id numbers (GIDs) on both the primary and backup node systems. The easiest way to synchronize UIDs and GIDs on both nodes is to use iSeries Navigator. The steps to synchronize across both nodes using iSeries Navigator are documented in Appendix B.

12. Conclusion and further resources

It should be apparent that developing and operating a PeopleSoft EnterpriseOne IASP solution is a relatively straightforward procedure. For educational opportunities and /or implementation assistance, visit the iSeries Technology Center Web site at:
ibm.com/servers/eserver/iseries/service/itc/index.html

Additional information about iSeries High Availability solutions can also be found on the High Availability Web site at:
ibm.com/servers/eserver/iseries/ha/

13. Appendix A

Relational Database Directory Entry or RDB entry might be a new concept for many iSeries users. On an iSeries system without an IASP there is only one database and therefore only one RDB entry. Once an IASP is created the system automatically adds an additional RDB entry with the same name as the IASP. To look at the RDB entries on a system use the following command: **WRKRDBDIRE *all**

The **Work With Relational Database Directory Entries** screen will appear. See figure 13.

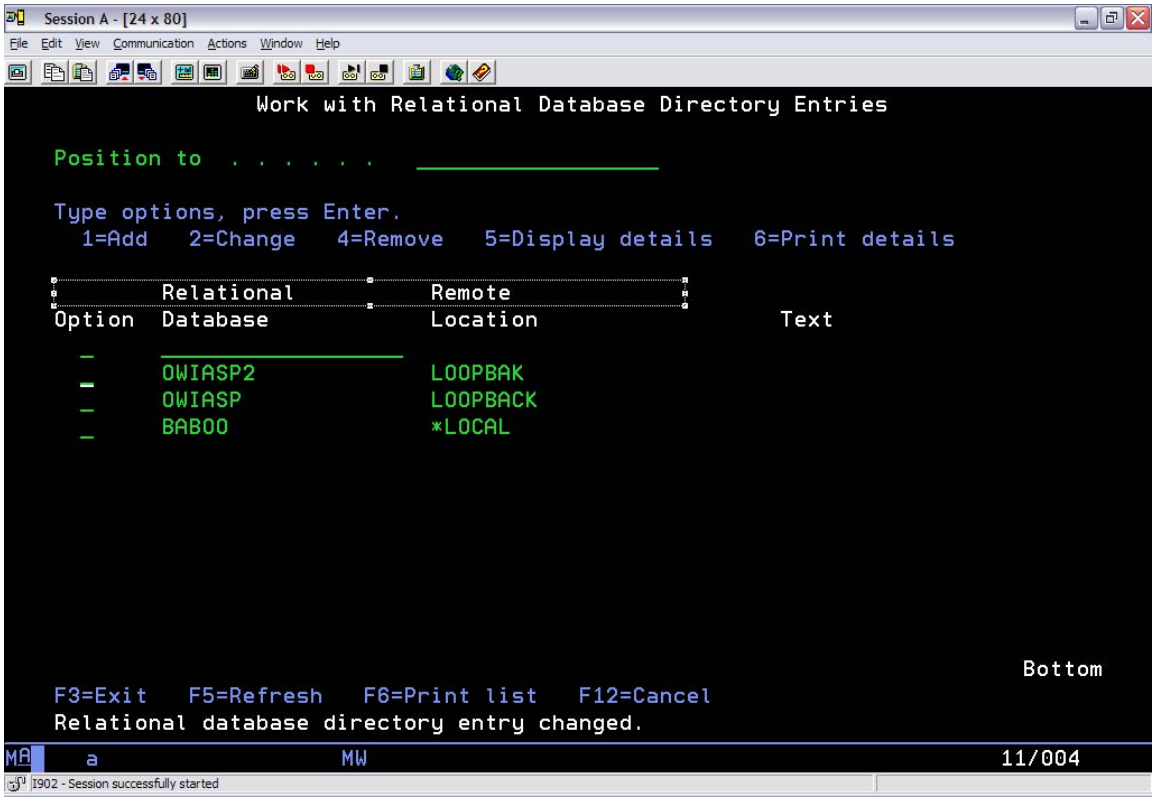

Figure 13

You will notice there are two different kinds RDB entries one with the remote location as *LOCAL and the others with the remote locations as LOOPBACK. The RDB entry with the remote location as *LOCAL is associated with system space and the system will only allow one *LOCAL RDB entry. However it is possible to have several LOOPBACK RDB entries, this is because it is possible to have more than one IASP on a given system.

To see which LOOPBACK RDB entry is associated with an IASP use the following command:

dspdevd name of IASP

The **Display Device Description** screen will appear. See figure 14.

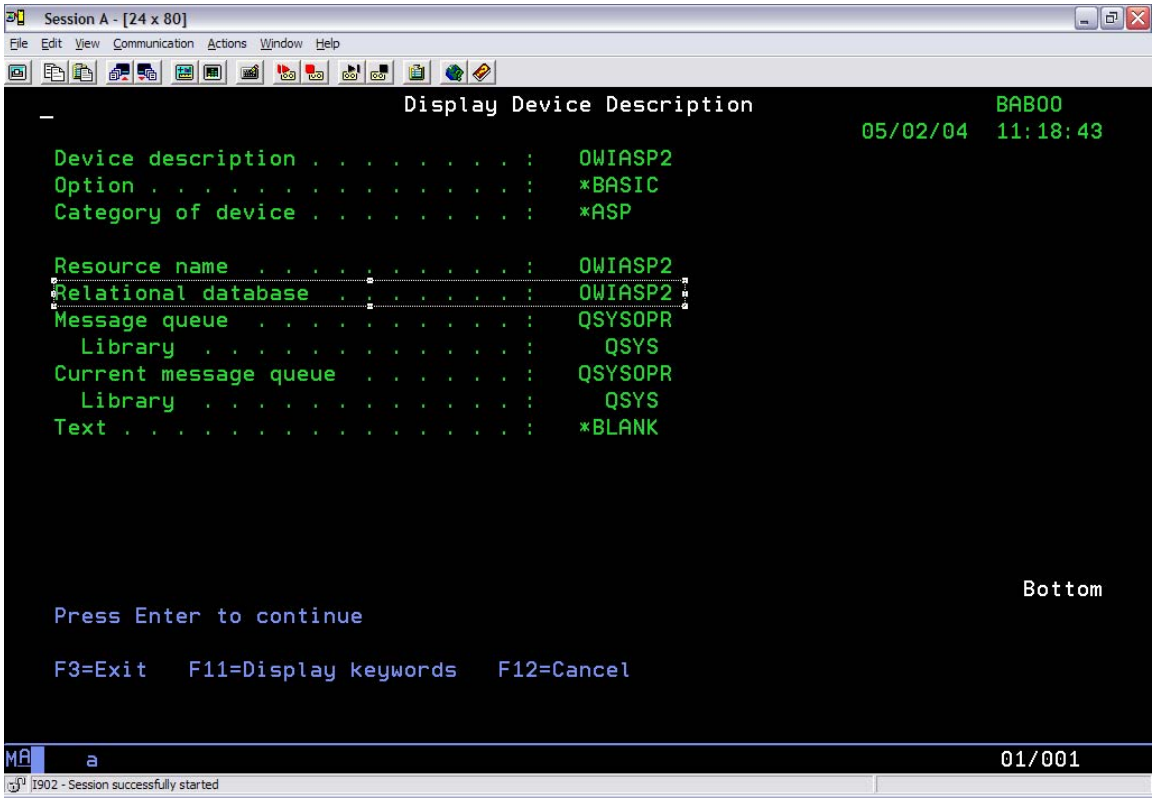

Figure 14

The IASP in the figure above is named OWIASP2 and it is associated with RDB entry OWIASP2.

14. Appendix B

This section provides step-by-step instructions on how to synchronize user profiles across nodes, using iSeries Navigator.

Begin by collecting an inventory of groups and users in the cluster.

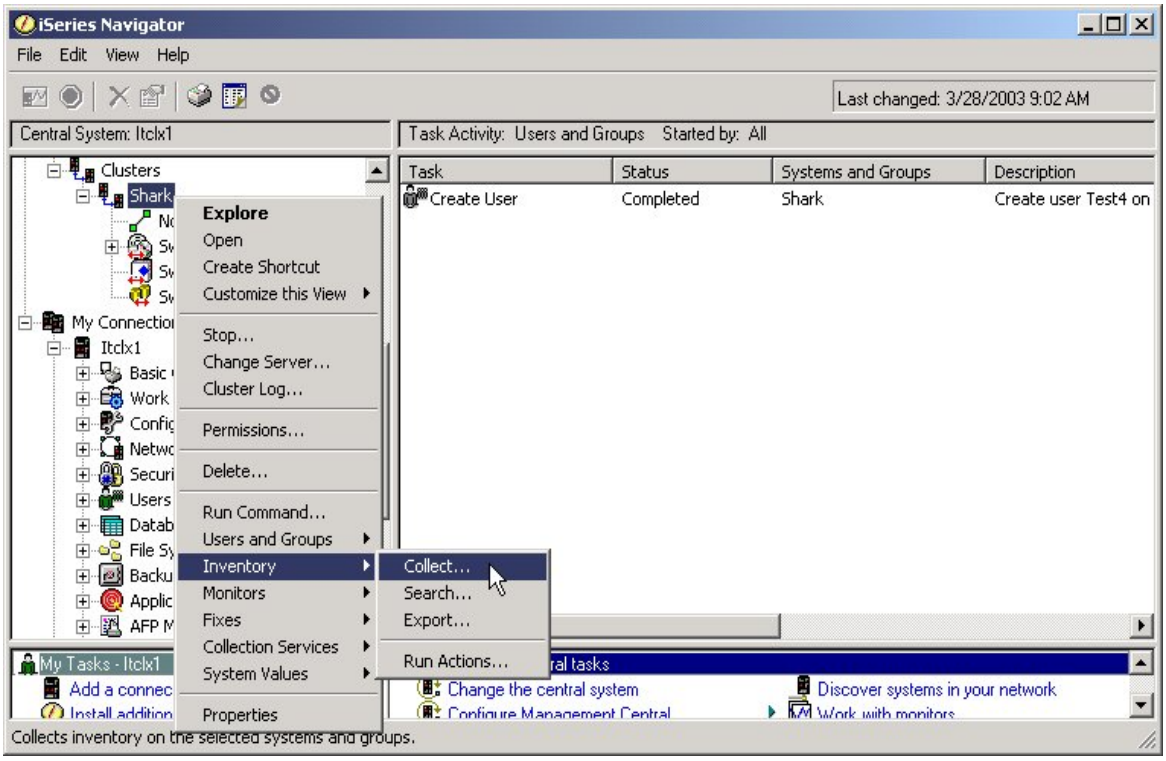

Figure 15

Select **"Users and Groups"** for the type of inventory to collect, then click on **"OK."** Verify the task completed successfully by clicking on **"Inventory"** under **"Task Activity."**

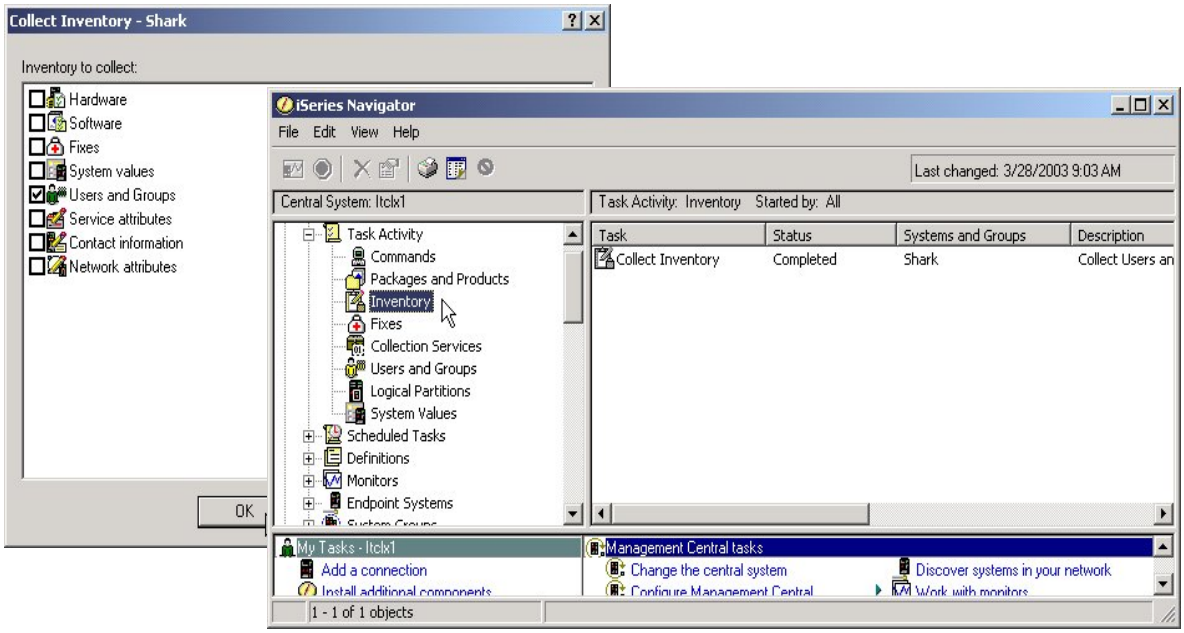

Figure 16

Next, select **"Edit Users"** under **"Users and Groups"** for the cluster.

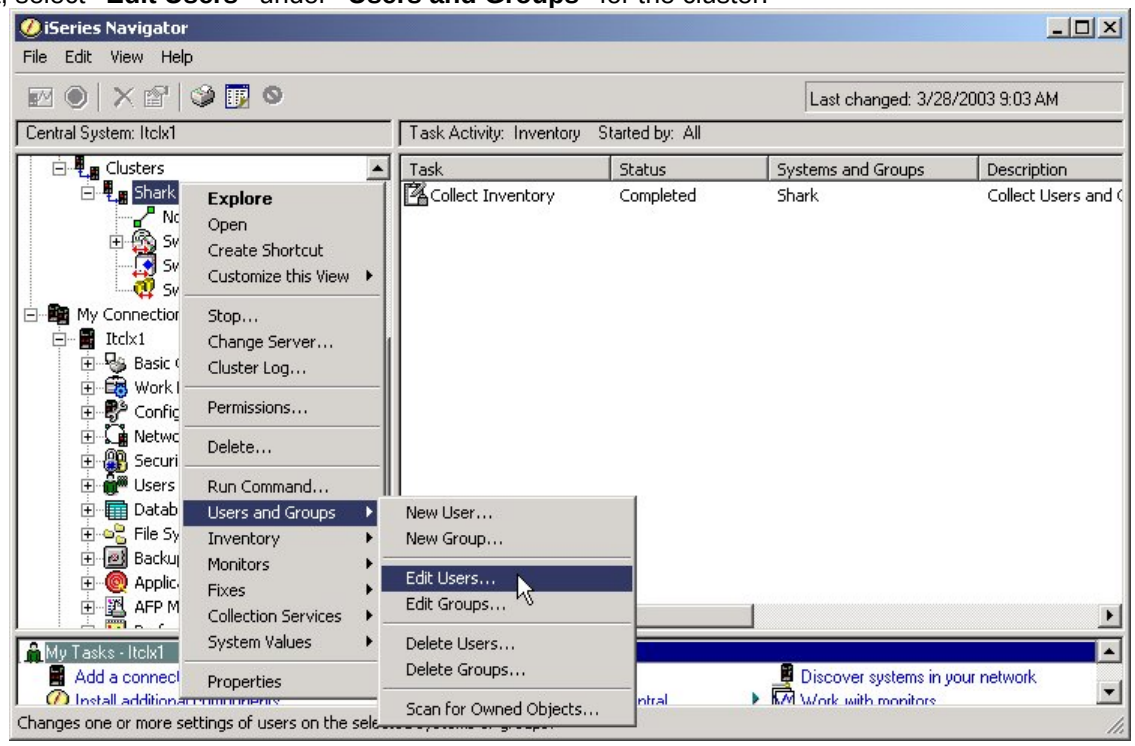

Figure 17

Click on **"Browse",** select the user profiles to synchronize, and then click on **"OK."**

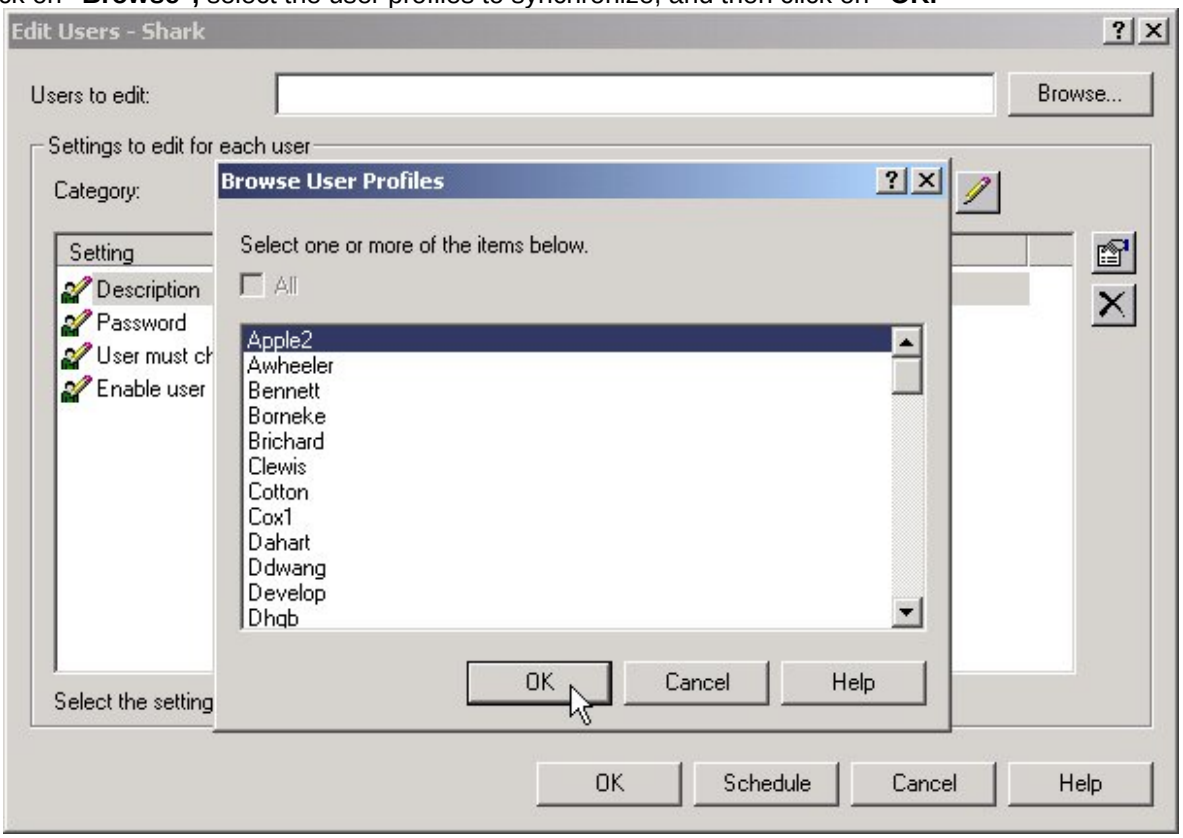

Figure 18

Select the category to edit **(Unique Identifier)**, and then click on **"Properties."**

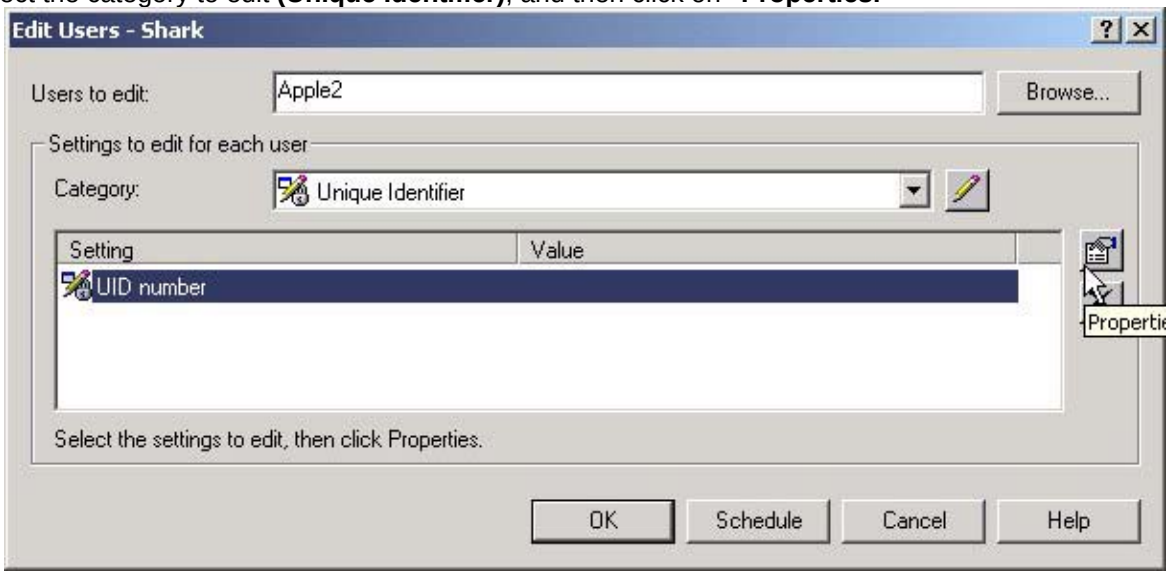

Figure 19

Select **"Find a unique identifier across all selected systems,"** then click on **"OK."**

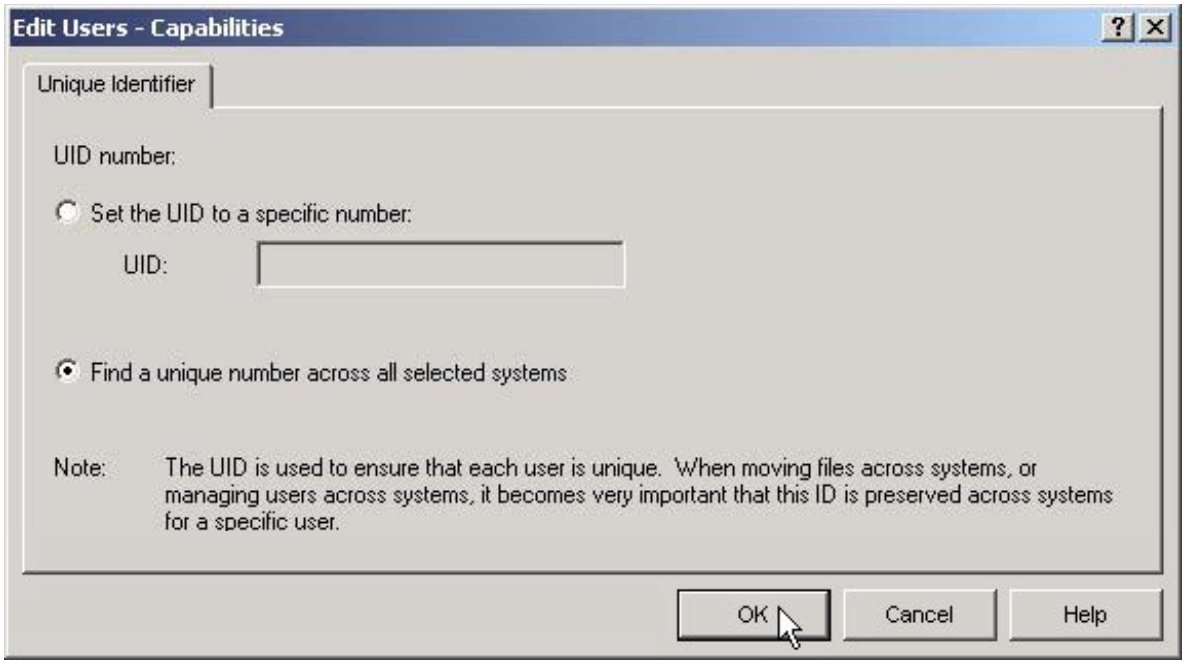

Figure 20

Click on **"OK"** to synchronize the user profile on both nodes.

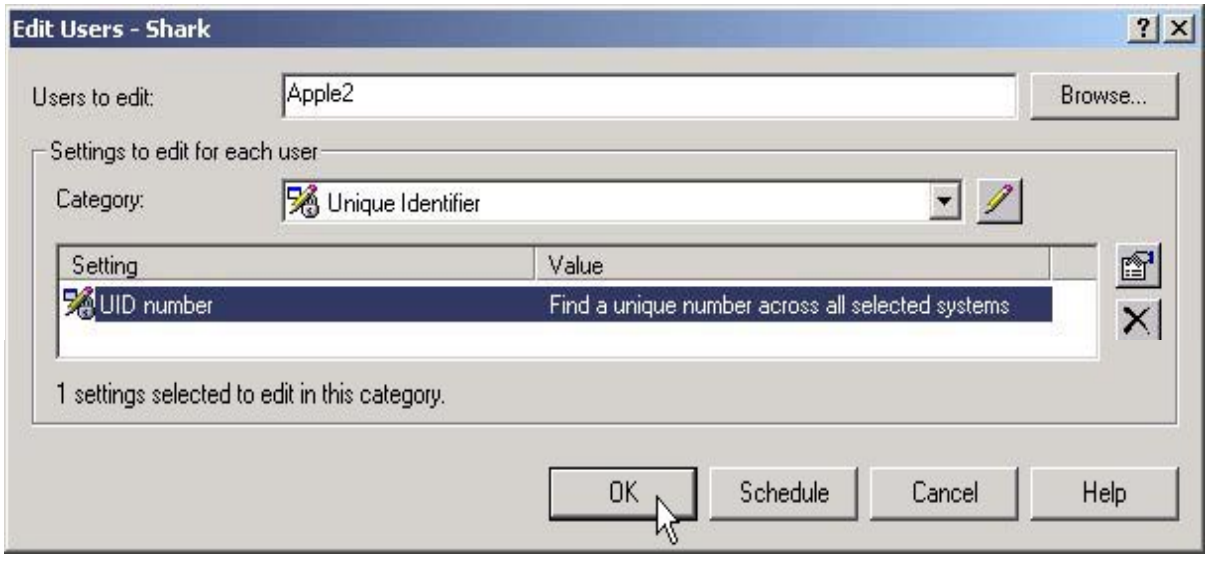

Figure 21

Verify that the task completed successfully. If it did not, double-click it to investigate.

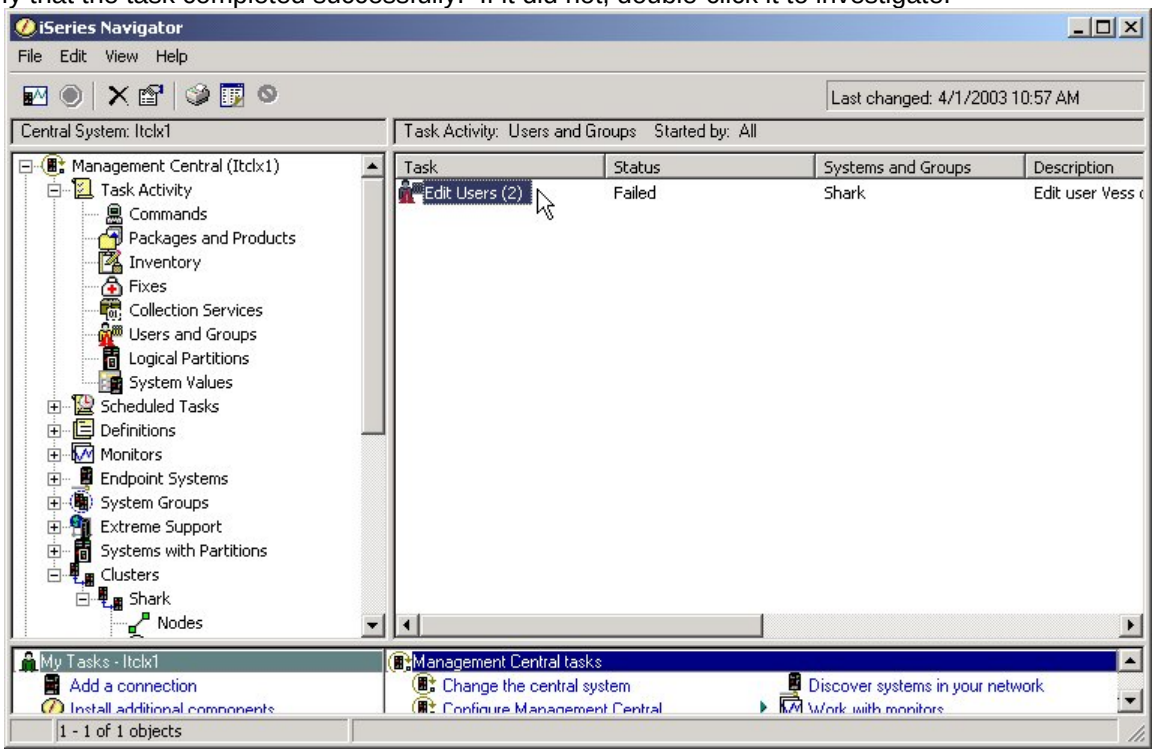

Figure 22

Often, the cause of the error will be apparent on the first status screen **(in this case, the user profile was in use when synchronization was attempted)**. If the cause of the error is not clear, right-click on each node to examine the job log on that system.

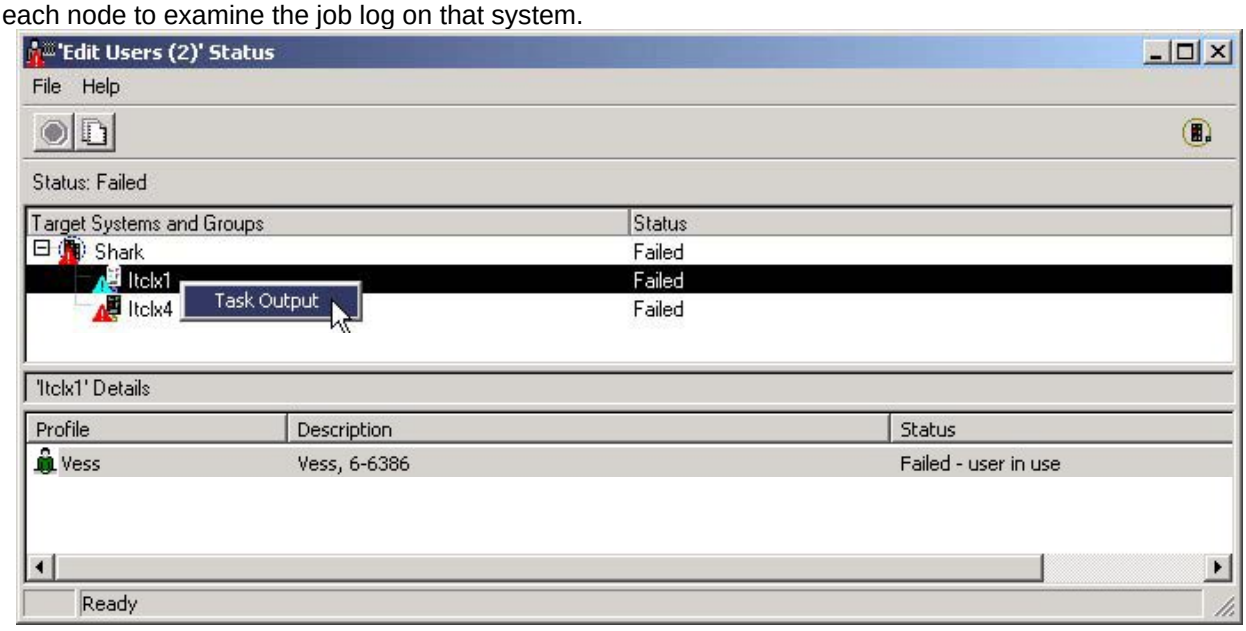

Figure 23

15. Appendix C – about the author

Marie Gann is a software engineer with IBM Rochester and is part of the ERP development team. Her primary focus is PeopleSoft and the EnterpriseOne product. She is also part of the development team for the iSeries file system functions. Marie has most recently worked with PeopleSoft on their support of independent ASPs.

IBM Corporation
Integrated Marketing Communications
Server Group
Route 100
Somers, NY 10589
U.S.A.
ibm.com

Visit *www.ibm.com/pc/safecomputing* periodically for the latest information on safe and effective computing. Warranty Information: For a copy of applicable product warranties, write to: Warranty Information, P.O. Box 12195, RTP, NC 27709, Attn: Dept. JDJA/B203. IBM makes no representation or warranty regarding third-party products or services including those designated as ServerProven or ClusterProven.

This publication was developed for products and services offered in the United States. IBM may not offer the products, services or features discussed in this document in other countries. Information is subject to change without notice. Consult your local IBM representative for information on offerings available in your area.

The examples given in this paper are hypothetical examples of how a customer can use the products described herein and examples of potential cost or efficiency savings are not based on any actual case study. There is no guarantee of comparable results. Many factors determine the sizing requirements and performance of a systems architecture. IBM assumes no liability for the methodology used for determining the configurations recommended in this document nor for the results it provides. Any performance data contained in this presentation was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements quoted in this presentation may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this presentation may have been estimated through extrapolation. Actual results may vary. Users of this presentation should verify the applicable data for their specific environment.

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

IBM, the IBM logo, the e-business logo, and iSeries are trademarks or registered trademarks of IBM Corporation in the United States, other countries, or both.

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

PeopleSoft Inc.
Corporate Headquarters
4460 Hacienda Drive
Pleasanton, California 94588 USA
Toll-free 1 888 773 8277
Tel 925 694 3000
www.peoplesoft.com

PeopleSoft, the PeopleSoft logo and Pure Internet Architecture are registered trademarks or trademarks of PeopleSoft, Inc.

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

PeopleSoft.