

IBM i switchable IASPs — a working configuration by example

A step-by-step guide to getting this working with IBM i clustering

Ian Marshall

June 18, 2012

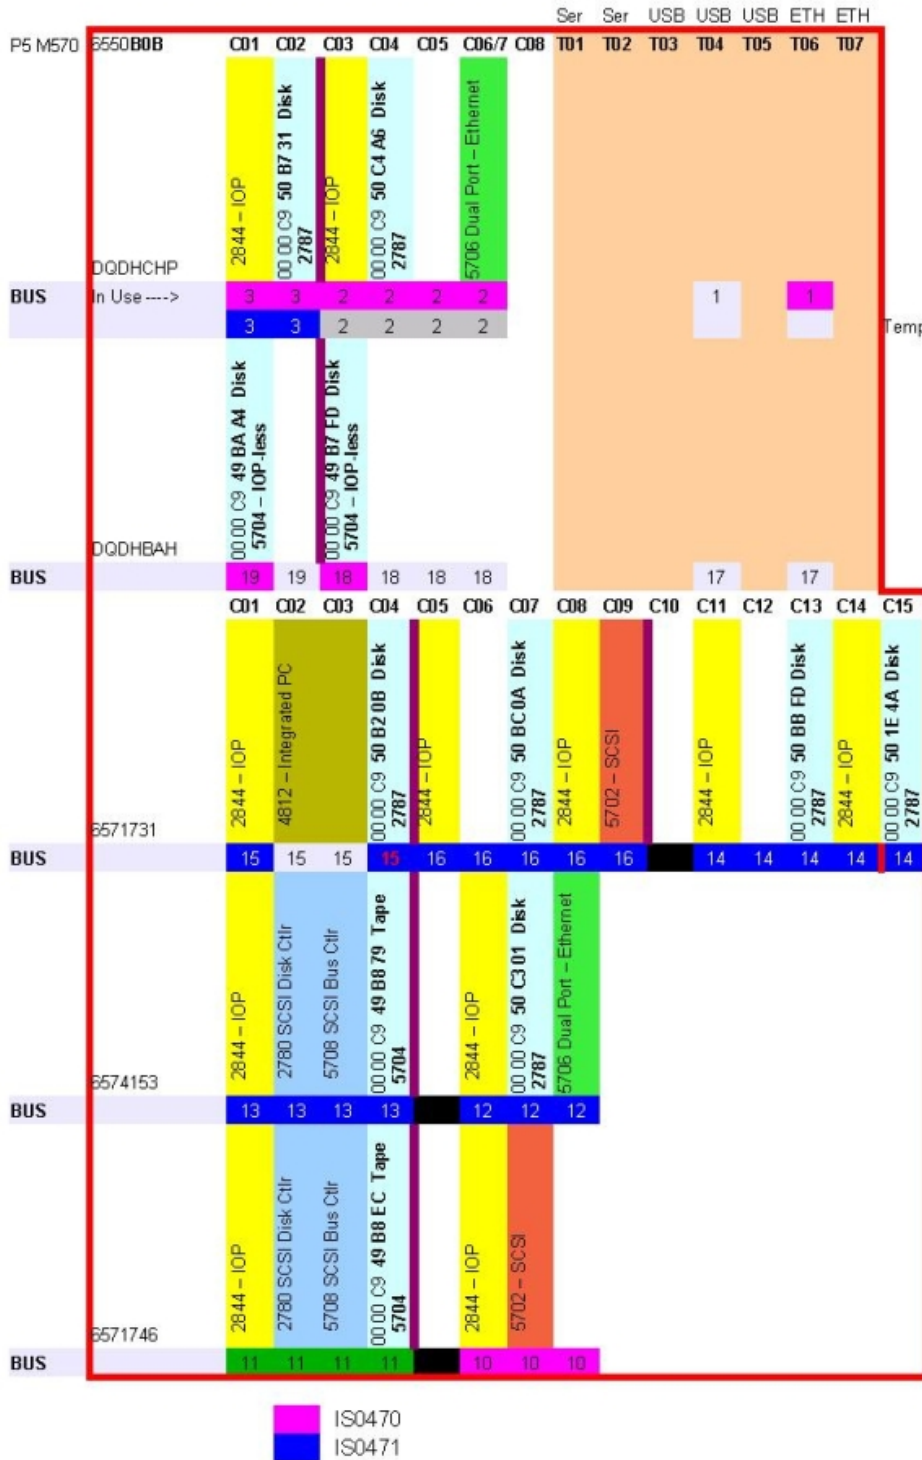
IBM i independent auxiliary storage pool (IASP) technology is the fundamental building blocks for providing switchable resource across IBM i servers. How easy is it to deploy this technology? There are several references on the subject, namely: [IBM i 6.1 Independent ASPs: A Guide to Quick Implementation of Independent ASPs](#), [IBM eServer iSeries Independent ASPs: A Guide to Moving Applications to IASPs, Clustering and IASPs for Higher Availability on the IBM eServer iSeries Server](#) and [Independent ASP Performance Study on the IBM eServer iSeries Server](#). However, the material, although comprehensive, was for either older versions of the operating system or did not provide a concise easy, step-by-step guide to getting switchable IASPs configured and working from hardware upwards. As a result, I decided to compile my own contribution comprising a 1,2,3 guide for getting IASPs configured for two IBM i partitions and made switchable, so that you can move storage from a primary to a secondary backup node in a pseudo-failover scenario.

[View more content in this series](#)

Physical hardware

First, we should assess my starting hardware environment. Figure 1 is a schematic illustration of the physical hardware of my the IBM Power Systems™ server detailing much of the hardware as viewed from the rear of the rack. You can see from the schematic that there is color coding that indicates how I allocate the various buses to each of the partitions. There are two partitions, IS0470 and IS0471, with bus and card allocations. Resources colored magenta for IS470 and resources colored blue for IS471. You notice that bus 11, color-coded green, has an 2844 IOP that is driving a SCSI bus controller and a fibre channel tape controller. These system bus resources are the resources that are shared (and thus switched) between the two partitions. The SCSI disk controller controls all 12 drives within the lower expansion unit.

Figure 1. Rack layout schematic



HMC configuration
Wanted and required allocations

We now need to configure the hardware according to the layout illustrated. On the Hardware Management Console (HMC) partition profiles, we set them up so that the allocations that are

magenta or blue are set as wanted resources for their respective partitions. The remaining bus 11, shown in green, should be marked as wanted for both of the partitions.


Figure 2. Allocations for IS0470 (Primary Partition)

Resource ID	Resource Name	Status	Bus
<input type="checkbox"/> U0595.001.6571746-CB1-C01	PCI I/O Processor	Desired	11
<input type="checkbox"/> U0595.001.6571746-CB1-C02	PCI Ultra4 SCSI Disk Controller	Desired	11
<input type="checkbox"/> U0595.001.6571746-CB1-C03	SCSI bus controller	Desired	11
<input type="checkbox"/> U0595.001.6571746-CB1-C04	Empty slot	Desired	11
<input type="checkbox"/> U0595.001.6574153-CB1-C06	PCI I/O Processor	Required	12
<input type="checkbox"/> U0595.001.6574153-CB1-C07	PCI Fibre Channel Disk Controller	Required	12
<input type="checkbox"/> U0595.001.6574153-CB1-C08	PCI 10/100/1000Mbps Ethernet UTP 2-port	Required	12
<input type="checkbox"/> U0595.001.6574153-CB1-C01	PCI I/O Processor	Required	13
<input type="checkbox"/> U0595.001.6574153-CB1-C02	PCI Ultra4 SCSI Disk Controller	Required	13
<input type="checkbox"/> U0595.001.6574153-CB1-C03	SCSI bus controller	Required	13
<input type="checkbox"/> U0595.001.6574153-CB1-C04	PCI Fibre Channel Tape Controller	Required	13

I/O pool definition

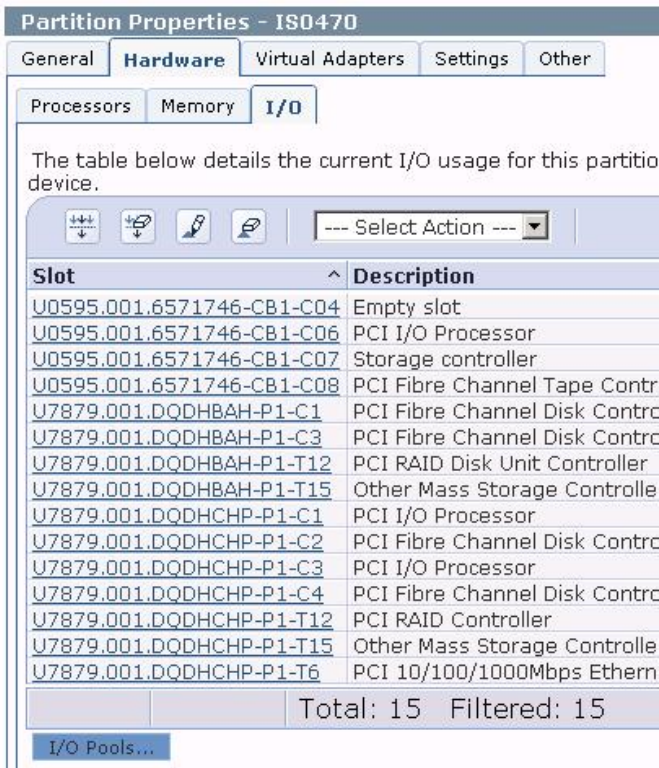
Now we need to configure an I/O pool for each of the partitions, so that we can add our shared resources to it in order to allow the server to switch them dynamically. On the Partition screen for the server, select one of the two partitions by clicking the adjacent twisty and select **Properties** from the drop-down menu.

Figure 3. Select Properties for a partition



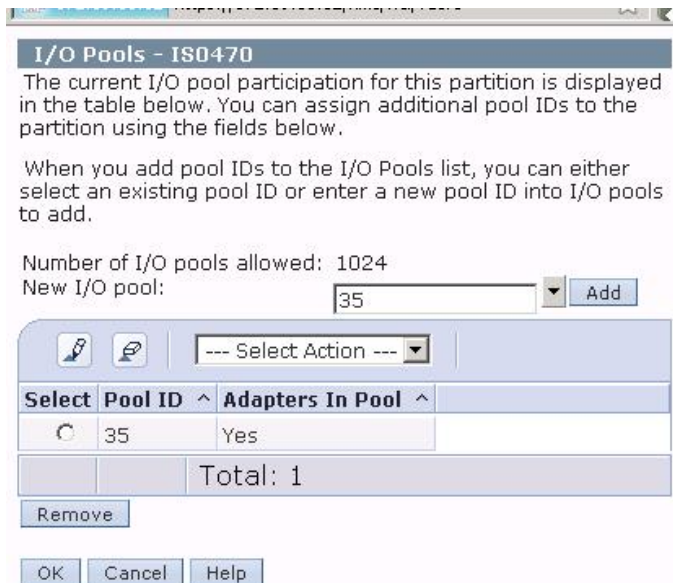
Now select **Hardware** and then **I/O**. At the bottom of the list of adapters, select **I/O Pools**.

Figure 4. Partition I/O pools



You then see the I/O Pools panel, which allows you to generate additional pool IDs for a partition. If you click the down arrow for the pool number field, you can see the next allocatable pool number. Select this number and click **Add**. You can see your new pool in the list.

Figure 5. Defining a new I/O pool



Click **OK** to accept the addition of the new pool.

You will need to repeat the pool-addition process for the other sharing partition. You use the same pool ID for the sharing partition.

Adding shared resource to the I/O pool

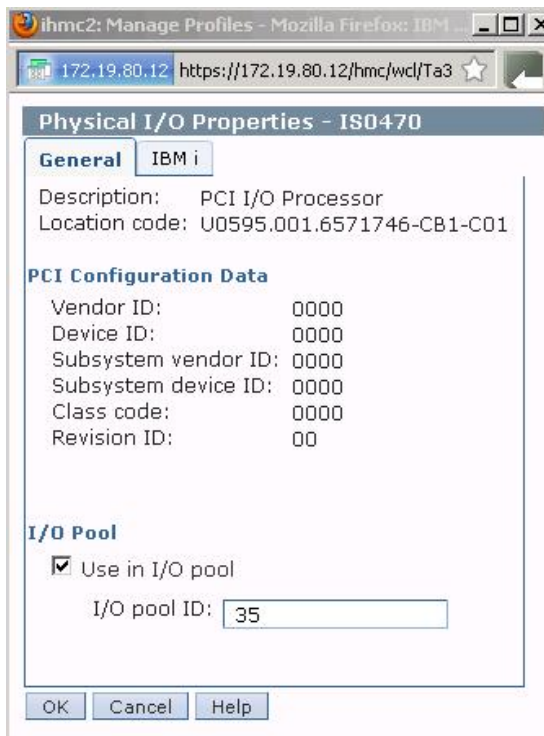
Now we need to add the shared (required) adapters to the pool we have just created. Select your primary partition in the partition list on the HMC. And select the twisty as before, but this time, select **Manage profiles**.

Figure 6. Select the configuration profile for a partition



To edit the profile, select your partition profile, then **Actions -> Edit**. I'm assuming this is familiar to you. Select the I/O tab and scroll down to the adapters that we have marked as desired. Click on the link for one of the shared adapters. You will see the following panel. Select **Use in I/O Pool** and input in the pool ID for the pool that you generated earlier. Then click **OK**.

Figure 7. Associate card resource with your I/O pool



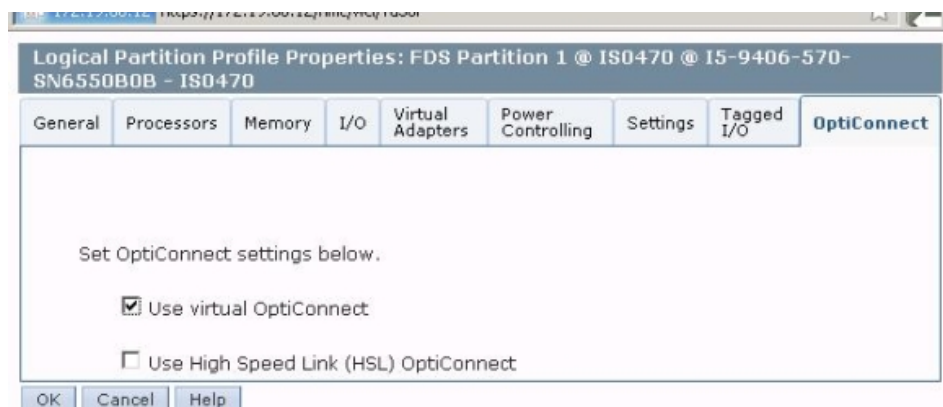
You have to repeat this for each adapter that you want to share, in my case all those on bus 11.

Enabling virtual OptiConnect

Now you have completed the process of adding all shared adapters to the I/O pool and you are back to the I/O panel within the profile configuration.

Select **OptiConnect** and select **Use virtual OptiConnect** for this partition profile.

Figure 8. Enable virtual Opticonnect



Click **OK** to save the profile.

You need to then select the profile for the backup partition and repeat the steps above to add the adapters in the profile for that partition to the same I/O pool and enable virtual OptiConnect.

HMC configuration summary

For both partitions, we have added a shared I/O pool with the same identification number for both to enable each to participate in the sharing of the adapter resources and any peripheral devices attached, in our case a number of disk drives. In addition, within both partition profiles, for each shared adapter, we added them to the I/O pool that has been defined and we enabled virtual OptiConnect within each of the partition profiles.

Activate primary and backup partitions

You now need to activate each of the partitions using the profiles configured, restore the licensed internal code (LIC) and operating system as normal, being careful on each partition to only allocate and format drives to ASP1 that are not attached to the adapters on the shared bus, in my case bus 11. Allocation and formatting of the shared drives will be done later.


after you have two operational and fully patched servers running, we can proceed to set up the object definitions to support a clustered environment. Clustering is the support mechanism to allow you to manage a switched IASP environment.

TCP/IP support for a clustered environment

Before you begin, you need to make sure that you have an operational TCP/IP environment. In my case, I took the recommended advice and allocated separate IP addresses for the IASP control

to those used for regular Telnet, FTP and other services. In Figure 9 you will see that I have two interfaces on my primary partition 172.19.78.98 for regular TCPIP activity and 172.16.78.99 that will be used for Clustering support (with addresses 172.16.78.102 and 172.16.78.103 respectively used on my backup server). I have however had this working with a single interface on each partition.

Figure 9. Separate IP interface for Clustering Support




In order to support the communications for clustering, I started the TCPIP INET Daemon, on both servers, which is the fundamental support for cluster monitoring.

STRTCPSVR SERVER(*INETD)

As shown in Figure 10, I then went into operations navigator and changed the settings to autostart this daemon for both servers.

Figure 10. AutoStart INET Daemon



Configuration of the clustered environment

Now we are in a position to create the clustered environment. These are the commands that I used to generate the objects required. It should be noted that I only needed to run the CRTCLU, ADDDEVDMNE and CRTCRG commands on the primary partition because the INETD was started and replicated the configuration to the backup partition.

Creating the cluster

Add both nodes, IS0470 and IS0471 to the cluster. For each node, we nominate a dedicated TCP/IP address to use for cluster communications and signaling.

```
CRTCLU CLUSTER(CLUSTER) NODE((IS0470 ('172.16.78.99')) (IS0471 ('172.16.78.103')))
```

Create the device domain

We need to add each of the Nodes that appear in the cluster to the device domain entity.

```
ADDDEVDMNE CLUSTER(CLUSTER) DEVDMN(DEVDMN) NODE(IS0470)
```

```
ADDDEVDMNE CLUSTER(CLUSTER) DEVDMN(DEVDMN) NODE(IS0471)
```

Create a device description for the IASP

Create a device description, on each partitions by running the following command on both partition (INETD does not replicate this command). I kept the name and the resource name the same and made the DEVD name on each partition the same.

```
CRTDEVASP DEVD(IASPDEVICE) RSRNAME(IASPDEVICE) TEXT('IASP device to be switched')
```

Create the cluster resource group

Finally we create the cluster resource group, which links the the operating system switchable object (in my case IASPDEVICE) to the cluster and defines what role each of the nodes in that cluster plays (primary, backup, etc)

```
CRTCRCG CLUSTER(CLUSTER) CRG(CRG) CRGTYPE(*DEV) EXITPGM(*NONE) USRPRF(*NONE)
RCYDMN((IS0470 *PRIMARY *LAST *NONE *NONE) (IS0471 *BACKUP *LAST *NONE *NONE))
CFGOBJ((*NONE *DEVD *ONLINE *NONE))
```

All of the operating system objects have now been generated

Configure the disk pool

All that remains is to create the disk pool itself, format the drives and link them to the IASPDEVICE device description. You can only do this using the operations navigator Disk Pool Wizard.

Start up IBM Navigator for i

If you do not already have a connection to your server, create one so that it is shown in **My Connections**.

Click **iSeries Server** to expand it (you might need to sign on). Click **Configuration and Service**, then **Hardware**, then **Disk Units** (you will be prompted for the Dedicated Service Tools user ID and password).

Figure 11. SST signon



Supply this to enable disk configuration

Right-click **Disk Pools** and select **New Disk Pool**.

Figure 12. Launch the Disk Pool wizard

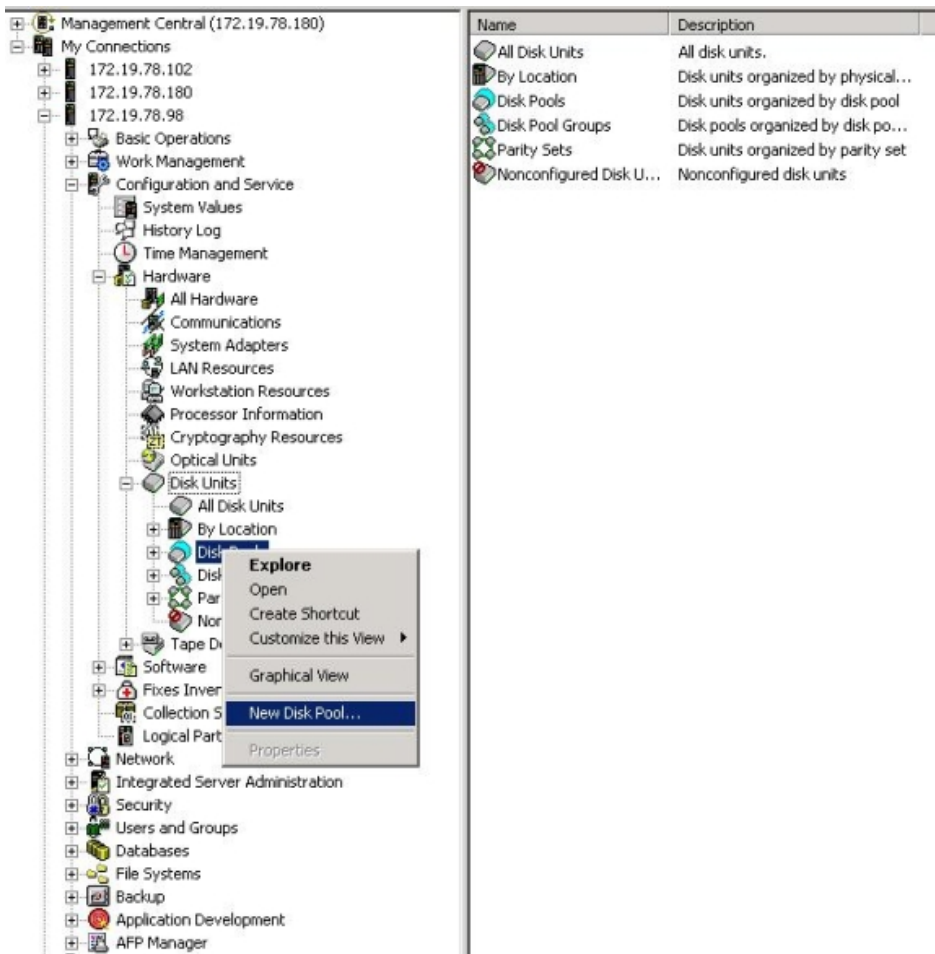


Figure 13. New Disk Pool wizard welcome



Click **Next**. Then click **New Disk Pool**.

Figure 14. New disk pool definition



For the Type of Disk Pool, select **Primary**.

Enter the name of the IASP device that you created earlier. (In my case, IASPDEVICE). Select or clear **Protect the data in the disk pool** as appropriate (I chose unprotected). Click **OK**.

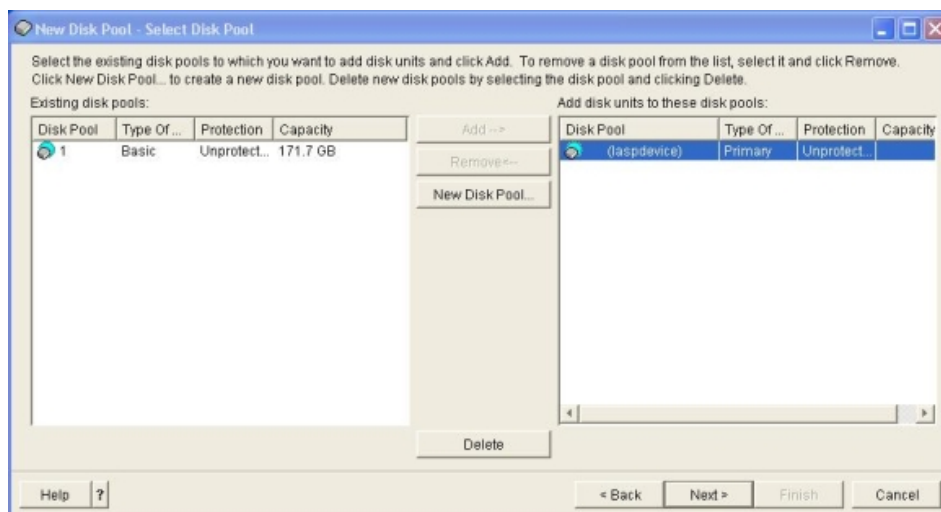
Figure 15. Device description already found



All being well, you should receive this message, that the device description already exists (I created it earlier on both servers). Click **OK** to continue.

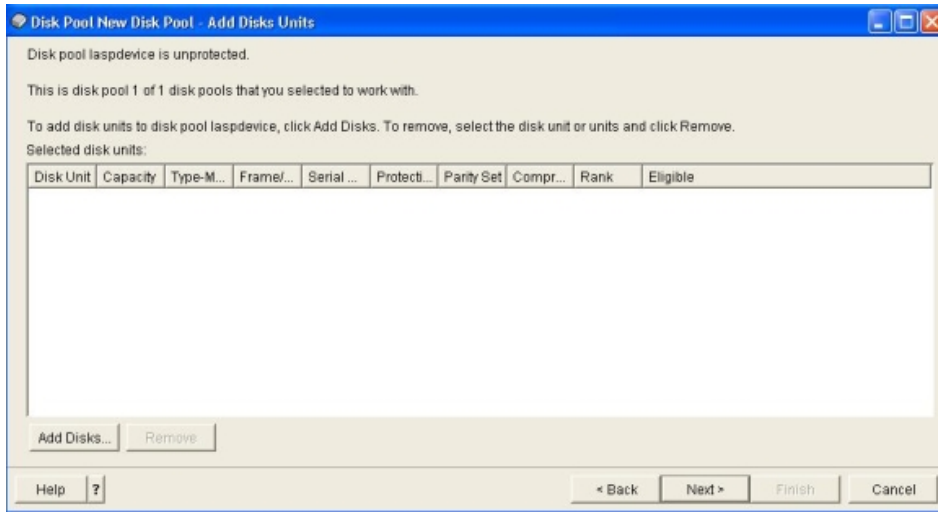
The new disk pool is shown in the right-hand pane. Highlight it and select **Next**.

Figure 16. Select disk pool



You see the Add Disk Units window. Select **Add Disks**

Figure 17. Add disk units



Highlight the disks that you want to add to the IASP and select **Add** (I selected all 12 devices).

Figure 18. Device description already found

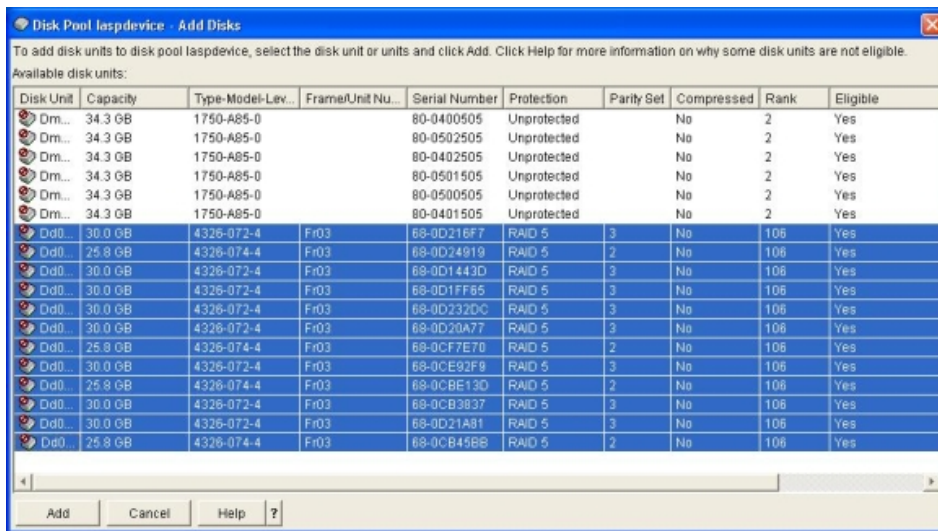


Figure 18 will continue to be shown while the disks are formatted and added to the IASP (ASP number 144 if this is your first IASP).

Notes to the reader

Note that moving IASP content from one server to another is not an instantaneous switchover and does not maintain application availability. IASPs on their own do not meet the high-availability requirements of most customers, for this IBM recommends a combination of IASP with a high-availability replication technology, such as IBM PowerHA SystemMirror for i, which exploits clustering technology to provide high availability, whether this be on native local DASD or external fibre channel-attached storage.

© Copyright IBM Corporation 2012

(www.ibm.com/legal/copytrade.shtml)

[Trademarks](#)

(www.ibm.com/developerworks/ibm/trademarks/)