

Here are some notes on LPARs/DLPARs under AIX 5.2 and HMC code Release 3, Version 1.0 and later (the minimum required for Dynamic LPARs).

Memory Assignment Restrictions

For AIX 5.2 DLPARs, it is highly recommended that you utilize the "Small Real Mode Address Region" option when defining the partition. Please keep in mind that by default this option is NOT selected. When this option is used the memory allocation follows this rule: the maximum memory allocation is 64 times the minimum memory allocation. The Small Real Mode Address Region option provides significant flexibility in assigning LPAR memory sizes.

Example:

An AIX 5.2 LPAR in "Small Real Mode" with a minimum memory size of 256MB could have a maximum size of 16 GB. A similar LPAR with a minimum 512MB memory allocation could have a maximum value of 32GB.

The memory size restrictions outlined below apply in those partitions where the "Small Real Mode Address Region" option has NOT been selected. Once again, this mode is not recommended for AIX 5.2.

The values for required, minimum and maximum memory allocation must all be within the same range. The allowable ranges are:

256MB - 4GB
1GB - 16GB
16GB - 256GB
256GB - 4TB

The following text was in the original messages on this topic but no longer applies: ~~So, we still have to be careful when we talk about memory allocations and changes for the LPARs. For example, it is not possible to create an LPAR with 8 GB and subsequently increase it to 32 GB. If you want an LPAR that can grow to 32 GB the minimum size is 16 GB.~~

DLPAR

Dynamic LPAR relies on Service Focal Point and RSCT. This means that TCP/IP (including the hostname) on the HMC must be configured. The /etc/hosts files on the HMC and all LPARs must have entries for all entities with the hostname appearing first in any list of aliases.

I ran into a problem with DLPAR after using the altinst_rootvg cloning procedure to create the second LPAR. The file /etc/ct_node_id had the same contents on both LPARs so the RSCT thought it had only one LPAR to manage. Development is still deciding whether or not the RSCT should have recognized and corrected this, or if a procedural change needs to be made to the altinst_rootvg write-up in the [AIX LPAR Guide \(SC23-4382\)](#). In the meantime, a workaround is to run these two commands after the first boot of the new LPAR:

```
/usr/sbin/rsct/install/bin/uncfgct -n  
/usr/sbin/rsct/install/bin/cfgct
```

A subsequent reboot should get everyone synchronized from an RSCT perspective.

The duplicate /etc/ct_node_id condition are likely to cause negative behaviors in a HACMP scenario between cloned LPARs. Cloning an LPAR with a mksysb tape/DVD will likely cause the same problem. I verified that NIM installs avoid this problem as a unique /etc/ct_node_id file is created.

Adding/Removing Adapters with DLPAR

Adding an adapter is just as you would expect. Select the active partition on the HMC and choose the DLPAR menu and adapters from the sub-menu. Then add the adapter. Run "cfgmgr" on the LPAR to make the new device available. NOTE that an added adapter is RETAINED across a reboot (shutdown -Fr). I suspect if you shutdown and start with a different profile the adapter is not retained but did not have time to confirm this.

Removing is a little more complicated:

- Move the child devices from "available" to "defined" (rmdev -l device)
- Move the parent device from "available" to "defined" (rmdev -l device)
- Move the associated PCI bus device from "available" to "defined" (rmdev -l pciXX)
- Use the DRSLLOT command to "remove" the card from the slot.
- Use the DLPAR menu on the HMC to remove the adapter.

If you are uncertain about child devices, use something like "odmget CuDv | grep -p parentdevice" to get the stanzas for the children.

NOTE: This script is useful when you are making the PCI bus associated with an adapter "defined" before the adapter can be moved to a different partition. It gathers key information about an adapter (including the PCI bus) into a single file.

```
#!/bin/ksh
#
# As-is script to display adapter information on an AIX system.
# Please forward changes to Bill Moraca at wmoraca@us.ibm.com
#
OUTFILE=/tmp/adapterinfo.txt
TITLE="DEVICE\tHW ADDRESS\tLOCN\tPCIBus\t DESCRIPTION"
#
print $TITLE >$OUTFILE
for i in `lsdev -Cc adapter | grep -i available | grep -v Standard |
grep -v Keyboard | grep -v Mouse | grep -vi etherchannel | cut -f1 -d"
 "`
do
  line=`lscfg -vl $i | head -1`
  devc=`print $line | awk '{ print $1 }'`
  hwslot=`print $line | awk '{ print $2}'`
  desc=`print $line | cut -f3- -d" "`
  locn=`lsdev -Cc adapter | grep ^$i" " | awk '{ print $3 }'`
  pcib=`odmget -q name=$i CuDv | grep parent | cut -f2 -d' '`
  print "$devc\t$hwslot\t$locn\t$pcib\t $desc" >>$OUTFILE
done
```

As an example on the above sequence, if you were removing ent4 you would do the following:

```
ifconfig en4 down
ifconfig en4 detach
rmdev -l en4
ifconfig et4 down
ifconfig et4 detach
rmdev -l et4
rmdev -l ent4
rmdev -l pci10      # assumes that ent4 is in pci10
DLPAR menu sequence
```