

# IBM PowerHA SystemMirror cluster migration to IBM POWER7

Chris Gibson

April 04, 2013

This article provides the tips for migrating an IBM® PowerHA® SystemMirror cluster from IBM POWER6® to IBM POWER7® processor-based servers. This step-by-step guide describes how to migrate a high-availability (HA) cluster to POWER7, upgrade the cluster from IBM HACMP™ to PowerHA, and migrate shared cluster volume groups for fast disk takeover.

## Purpose

The purpose of this article is to provide a step-by-step guide for migrating an existing IBM HACMP (PowerHA) cluster from a POWER6 processor-based server to a new POWER7 processor-based server. This article is based on a real-world customer scenario. Though your environment and requirements might differ from those presented here, similar methodology can be applied in most other cases.

The customer had purchased two new IBM Power Systems™ 795 (9119-FHB) server based on POWER7 processor-based technology. They needed to move their existing HACMP clusters from their old POWER6 hardware to the new systems. Along with the server migration, they also required an upgrade of HACMP (as their current version was due to go out of support very soon). Upgrading the cluster as part of the POWER7 migration was deemed appropriate given that the cluster would need to be offline during the migration anyway.

The customers' workload was distributed across two existing IBM Power 595 (9119-FHA) systems based on POWER6 processor-based technology. The existing cluster was installed with HACMP version 5.4. It was a two-node cluster, with a cluster node on each Power 595 server. The LPARs and the cluster were built many years ago. Therefore, as part of the migration process, the following components would be changed or upgraded:

- New POWER7 logical partitions (LPARs), with new physical 8 GB Fibre Channel (FC) and 1 GB Ethernet adapters. The old POWER6 LPARs would be decommissioned, along with the Power 595 servers.
- New disk storage had been introduced; we would need to move the cluster-shared volume group disks from IBM System Storage® DS8300 disk to Hitachi Data Systems (HDS) virtual service processor (VSP) disk devices.

- The Subsystem Device Drivers (SDDs) and associated vpath devices would need to be removed as the IBM storage devices would no longer exist on the new LPARs.
- To support the new HDS disk devices we would need to install the HDS Object Data Manager (ODM) file set (and configure new HDS VSP hdisk devices).
- The shared volume groups in the cluster would be converted from standard to enhanced-concurrent mode.
- New network switches were introduced. The 1GB Ethernet adapters in the POWER7 LPARs would be connected to the new switches.
- IBM AIX® 5.3\* updated to TL12 (from TL8).
- HACMP 5.4 migrated to PowerHA 6.1.

**\*Note:** AIX 5.3 is no longer supported. At the time of planning for the POWER7 migration (in September 2011) the customer preferred to remain on AIX 5.3 due to the fact that their application software had not yet been verified to run on AIX 7.1 (or 6.1). They planned on upgrading the cluster to AIX 7.1 this year (2013).

## Migration overview

Before starting the migration, we produced a (very) high-level list of the steps required to achieve our goal. The high-level migration steps are as follows:

1. Restore the LPARs from mksysb through Network Installation Management (NIM) (recover devices = yes) on new POWER7 LPARs.
2. Migrate the shared and non-shared disk configuration from IBM to HDS disk.
3. Import the non-shared volume groups.
4. Import the shared volume groups.
5. Perform HACMP Discovery to find new (HDS) hdisk devices instead of vpaths (SDD).
6. Synchronize/Verify the cluster.
7. Start cluster services.
8. Reconfigure the disk heartbeat to use hdisks instead of vpaths.
9. Synchronize/Verify the cluster.
10. Enable shared volume groups for fast disk takeover (enhanced-concurrent mode).
11. Synchronize/Verify the cluster.
12. Stop cluster services on both nodes.
13. Start cluster services on both nodes.
14. Verify shared volume groups are varied on in enhanced-concurrent mode.
15. Stop cluster services on both nodes.
16. Upgrade (migrate) HACMP 5.4 to PowerHA 6.1 SP6 on both nodes.
17. Reboot both nodes.
18. Start cluster services.
19. Synchronize/Verify the cluster.
20. Perform cluster failover tests.
21. Ensure that migration is complete.

At the end of the migration, the cluster nodes reside on the POWER7 processor-based systems, that is, one node on each Power 795 server. Each node will be running PowerHA 6.1 and will be utilizing HDS disk devices for all AIX volume groups.

## The cookbook

What follows is essentially our 'cookbook' for migrating each cluster to the new POWER7 processor-based servers. There were six two-node clusters that needed to be migrated to POWER7. Our 'cookbook' provided simple steps that can be followed by any member of the customers UNIX administration team during the POWER7 migration project. These steps were developed and tested using a lab/test cluster environment on the POWER6 and POWER7 processor-based systems. The team tested and refined the migration procedures several times before implementing in the production environment. This testing was crucial to the success of the migration project.

### A. Preparatory steps

1. The new LPARs on the POWER7 system were pre-provisioned with new HDS disks for the AIX OS (rootvg). The disk was assigned and tested before starting any migration activities. Tests would usually include performing a dummy mksysb restore to the new LPARs to ensure that the disk was operational and that the mksysb could be recovered successfully.
2. The Ethernet and Fibre Channel adapters assigned to the new LPARs were precabled to the correct network and SAN switches. It is important to ensure that the network configuration is correct in order for PowerHA to continue functioning after migration. We ensured that the network interfaces on the HA adapters (en0 and en2) were configured in the appropriate virtual local area network (VLAN). This configuration was identical to the configuration for the nodes on the POWER6 system. Tests were performed to confirm that interfaces were in the correct VLAN and that each node could communicate with its partner before migration.  
**Note:** If the network adapters/interfaces are not assigned to the same VLAN, then the cluster may behave unexpectedly. For example, during testing, we found that our test cluster was unable to communicate through one of the boot interfaces after a resource group move. The issue was traced to the fact that the HA (boot) interfaces were in different VLANs on the network switch. Moving both interfaces into the same VLAN on the switch resolved the problem.
3. Before we start, first we ensure that we have disabled monitoring on our two-node cluster. We place the cluster nodes into maintenance mode (in our case, we disabled the customers Nagios monitoring on the cluster nodes).  
The nodes and the cluster will essentially be unavailable during the migration. This step prevents any unwanted alerts during the migration. We also document the current configuration of each AIX system before implementing the change.
4. Before migration, we need to ensure that the cluster is healthy and stable. We perform a *HACMP Synchronize and verify* operation to verify whether this is the case. If the cluster is not stable or does not synchronize, we will correct this situation before starting the migration. There is little point in moving the cluster when it is broken; doing so would most likely result in a failed migration. We also use the tools, **clstat** and **cldump** to view the current status of the cluster and the nodes.

```
# smit hacmp
Initialization and Standard Configuration
Verify and Synchronize HACMP Configuration

# clstat < On both nodes. Is everything UP?

# cldump < On both nodes. Is the cluster currently STABLE?

# vi /tmp/hacmp.out < On both nodes, checking for any events, errors or failures.

# errpt
```

5. It is always a good idea to take a snapshot of the cluster (from the home node) before you make any change to a HA cluster. This snapshot can be used to recover the cluster to a known state when issues arise. A script was run to perform a snapshot.

```
# /usr/local/bin/cluster_snap.sh
```

6. We take a mksysb backup of the LPARs to our NIM master. This backup image can be used to restore an AIX node to a known state when we encounter issues with the migration. We also perform a file data backup to IBM Tivoli® Storage Manager.

```
# /usr/local/bin/saveskelvg >> /var/log/saveskelvg.log 2>&1
# /usr/local/bin/mksysb2nim -s rootvg >> /var/log/mksysb2nim.log 2>&1
# dsmc i
```

7. The cluster is migrating from IBM storage to HDS disk. We must install the latest version of the HDS ODM driver **on each node** to support this new device. This is installed using an NFS mount to the NIM master software repository.

```
# mount nim1:/software/HDS/HDS_odm_driver/HTCMP103 /mnt
# cd /mnt
# smit install
devices.fcp.disk.Hitachi.array.mpio.rte
```

8. The migration process can now start. We stop cluster services on **both nodes** first. At this point, all cluster resources are taken offline. All clustered applications are no longer available.

```
nodeA# smit clstop
```

9. All data volume groups must be varied off and exported before migration. The associated devices (vpaths) must also be removed (with **rmdev**) from the AIX ODM.

```
# varyoffvg vgroupname
# exportvg vgroupname
# rmdev -dl vpathX
# rmdev -dl hdiskX
```

10. Both nodes and LPARs are now shut down on each 595 server.

```
node1 # shutdown -F
node2 # shutdown -F
```

## B. Migrate LPAR to POWER7 using NIM

1. The cluster nodes were previously configured with mirrored rootvg disks. This is no longer required on the new storage; the new LPARs are configured with a single disk for rootvg. To ensure that the mksysb restore to the new POWER7 LPAR (with a single rootvg disk) is successful, we must create a custom **image.data** file for each node. This **image.data** file will ensure that the mksysb restore process does not attempt to mirror the root volume group. If we leave the **image.data** file untouched, the restore process fails, stating that there is insufficient disk space to cater for the mksysb image.

We extract the **image.data** file from the node's mksysb image on NIM. We change the **LV\_SOURCE\_DISK\_LIST**, **PP=**, and **COPIES=** values to reflect a non-mirrored disk configuration for rootvg. Then, we create a new **image\_data** NIM resource using the custom **image.data** file.

```

root@nim1 : /home/cgibson # restore -xvqf /export/images/nodeB-mksysb.530802-.Thu
./image.data

root@nim1 : /home/cgibson # cp image.data node2.image.data

root@nim1 : /home/cgibson # vi node2.image.data

root@nim1 : /home/cgibson # grep LV_SOU node1.image.data
LV_SOURCE_DISK_LIST= hdisk0
...etc...

root@nim1 : /home/cgibson # grep PP= node2.image.data
PP= 1
...etc..

root@nim1 : /home/cgibson # grep COPIES node1.image.data
COPIES= 1
...etc...

```

The process (above) was manual. You can automate this procedure using this [script](#) (courtesy of Kristijan Milos).

A new NIM **image\_data** resource was defined using our custom **image.data** file.

```

root@nim1 : / # smit nim_mkres
image_data

                                Define a Resource

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                [Entry Fields]
* Resource Name                  [cg-node2-image-data]
* Resource Type                  image_data
* Server of Resource             [master]
* Location of Resource           [/home/cgibson/node2.image.data]

```

2. Now, we can restore the mksysb of each LPAR to each Power 795 server using NIM. We must specify the custom image data file that we created in the previous step.

- Use the AIX 531204 lpp\_source and SPOT.
- Select recover devices and import user data.

**IMAGE\_DATA** to use during installation [cg-node2-image-data] +

Refer to the following articles for detailed information on migrating AIX LPARs to new Power Systems hardware.

- [Avoiding the gotchas of AIX LPAR migrations](#)

3. After the mksysb restore is complete, check to make sure that the system is now running at the required AIX level.

```

# oslevel -s
5300-12-04-1119
# instfix -i |grep AIX
# instfix -icqk 5300-12_AIX_ML| grep ":-:"
# lppchk -m3 -v

```

## C. Verify LPAR network configuration

1. Ensure each network is configured correctly for HA on each node. The networks, 10.1.2 (en0) and 10.1.3 (en2) were used for the customers' HACMP network configuration.

e.g. nodeA.

```

en0: flags=5e080863, c0<UP, BROADCAST, NOTRAILERS, RUNNING, SIMPLEX, MULTICAST, GROUPRT,
    64BIT, CHECKSUM_OFFLOAD(ACTIVE), PSEG, LARGESEND, CHAIN>
    inet 10.1.2.19 netmask 0xffffffff00 broadcast 10.1.2.255

en2: flags=5e080863, c0<UP, BROADCAST, NOTRAILERS, RUNNING, SIMPLEX, MULTICAST, GROUPRT,
    64BIT, CHECKSUM_OFFLOAD(ACTIVE), PSEG, LARGESEND, CHAIN>
    inet 10.1.3.14 netmask 0xffffffff00 broadcast 10.1.3.255

en4: flags=5e080863, c0<UP, BROADCAST, NOTRAILERS, RUNNING, SIMPLEX, MULTICAST, GROUPRT,
    64BIT, CHECKSUM_OFFLOAD(ACTIVE), PSEG, LARGESEND, CHAIN>
    inet 10.1.4.4 netmask 0xffffffff00 broadcast 10.1.4.255

en6: flags=5e080863, c0<UP, BROADCAST, NOTRAILERS, RUNNING, SIMPLEX, MULTICAST, GROUPRT,
    64BIT, CHECKSUM_OFFLOAD(ACTIVE), PSEG, LARGESEND, CHAIN>
    inet 10.1.5.7 netmask 0xfffffc00 broadcast 10.1.5.255

lo0: flags=e08084b<UP, BROADCAST, LOOPBACK, RUNNING, SIMPLEX, MULTICAST, GROUPRT, 64BIT>
    inet 127.0.0.1 netmask 0xff000000 broadcast 127.255.255.255
    inet6 ::1/0

```

2. Ensure that each boot interface can ping it's partner interface on the alternate node. Verify whether the host names resolve to the correct IP addresses for the boot and service labels.

For example:

On nodeA:

```

# ping nodeBb1
# ping nodeBb2

# host nodeAb1
# host nodeAb2
# host nodeBb1
# host nodeBb2
# host nodeAsvc
# host nodeBsvc

```

On nodeB:

```

# ping nodeAb1
# ping nodeAb2

# host nodeAb1
# host nodeAb2
# host nodeBb1
# host nodeBb2
# host nodeAsvc
# host nodeBsvc

```

3. Comment out the following line in **/etc/rc.net**. This errant entry was changing the cluster node host name to an unexpected value during system startup. We found it necessary to disable this on the customers' cluster.

```

# vi /etc/rc.net

# Below are examples of how you could bring up each interface using
# ifconfig. Since you can specify either a hostname or a dotted
# decimal address to set the interface address, it is convenient to
# set the hostname at this point and use it for the address of
# an interface, as shown below:
#

```

## D. Prepare for storage migration

1. Remove the old SDD filesets from each LPAR. As we are no longer accessing IBM storage, we no longer need the IBM Subsystem Device Drivers installed.

```
# stopsrc -s sddsrv
# installp -u devices.sdd.53.rte
# installp -u ibm2105.rte
# installp -u devices.fcp.disk.ibm.rte
```

- We also took this opportunity to remove any old `multibos` instance in `rootvg` (where applicable).

```
# multibos -R
```

- Reboot both LPARs.

## E. Storage migration

- **Handover to the Storage team for data logical unit number (LUN) migrations**

At this stage, the Storage team would take over. They perform the actual data migration from the IBM disk to the HDS disk. After completing data migration, the data LUNs would be rezoned to the new worldwide port names (WWPNs) associated with the new 8GB FC adapters in the new POWER7 LPARs. Essentially, when we restarted the LPARs, we would expect to see new HDS disk (not IBM). However, all the data would be intact. The data migration was achieved using the HDS storage replication technology. Using this method, we did not need to employ traditional procedures that involved AIX Logical Volume Manager (LVM) migration techniques using utilities such as **migratepv**, **mirrorvg**, and so on. If you are interested in learning more about using AIX LVM to perform storage migrations, I would refer you to my [2010 article](#) on the subject.

## F. Migrate cluster storage devices

- Run the **cfgmgr** command and configure the HDS disk and FC adapters on both nodes.

```
# cfgmgr
# lsdev -Cc disk
# chdev -l fscsi{y} -a fc_err_recov=fast_fail -a dyntrk=yes -P
# chdev -l hdisk{x} -a reserve_policy=no_reserve -P
# chdev -l hdisk{x} -a algorithm=round_robin -P
# lsattr -El fscsi{y}
# lsattr -El hdisk{x}
```

where `fscsi{y}` is each FC adapter in LPAR  
where `hdisk{x}` is ALL disks.

- Check whether the Physical Volume Identifier (PVIDs) are intact (that is, have not changed after the LUN rezone/migration) and are identical on both nodes.

```
nodeA# lspv
nodeB# lspv
```

- Import data volume groups, both shared and non-shared.

- **Non-shared volume group:**

```
# importvg -y appvg hdisk4
```

- **Shared volume group:**

- Perform **importvg** on both nodes. For example:

On **nodeA**:

```
# importvg -y sharedvgA hdisk2
# lsvg -l sharedvgA
# varyoffvg sharedvgA
```

One **nodeB**:

```
# importvg -y sharedvgA hdisk2
# lsvg -l sharedvgA
# varyoffvg sharedvgA
```

- Mount non-shared file systems. Check whether the file system mount order is correct, as this may have changed after the volume group (VG) was imported. Ensure that there are no file system overmounts.

```
# for i in `lsvgfs appvg`
do
  mount $i
done
# lsvg -l appvg
# df
# lspath
```

4. Perform a HACMP discovery. This will find the new HDS hdisk devices for the shared volume groups.

```
# smit hacmp
Extended Configuration
Discover HACMP-related Information from Configured Nodes
```

5. Perform a "Verify and Synchronize HACMP Configuration" operation now.
6. Start cluster services on both nodes. Confirm that the cluster is stable with `clstat`.

## G. Configure new cluster heartbeat devices

1. Reconfigure disk heartbeat devices to use hdisks instead of vpaths.

```
# smit hacmp
Extended ConfigurationConfigure HACMP Communication Interfaces/DevicesChange/Show Communication
Interfaces/Devices
  Select one or more Communication Interfaces/Devices to Remove
Move cursor to desired item and press F7. Use arrow keys to scroll.
ONE OR MORE items can be selected.
Press Enter AFTER making all selections.

# Node / Network
#   Interface/Device  IP Label/Device Path          IP Address

# nodeA / net_diskhb_01
vpath1                nodeA_vpath1_01          /dev/vpath

# nodeA / net_ether_01
en0                    node1b1                   10.1.2.
en2                    node1b2                   10.1.3.

# nodeA / net_ether_02
en4                    nodeA                      10.1.4.

# nodeB / net_diskhb_01
vpath0                nodeB_vpath0_01          /dev/vpath

# nodeB / net_ether_01
en0                    node2b1                   10.1.2.
en2                    node2b2                   10.1.3.

# nodeB / net_ether_02
en4                    nodeB                      10.1.4.

Remove Communication Interfaces/Devices

# nodeA / net_diskhb_01
vpath1                nodeA_vpath1_01          /dev/vpath
# nodeB / net_diskhb_01
vpath0                nodeB_vpath0_01          /dev/vpath

Interface and Devices

Select Point-to-Point Pair of Discovered Communication Devices to Add
```

```

Move cursor to desired item and press F7.
ONE OR MORE items can be selected.
Press Enter AFTER making all selections.

# Node                Device  Pvid
>nodeA                hdisk1  00c69b9edd04057b
>nodeB                hdisk1  00c69b9edd04057b

Change/Show Communication Interfaces/Devices

      Select a Communication Interface/Device to Change/Show

Move cursor to desired item and press Enter. Use arrow keys to scroll.

# Node / Network
#   Interface/Device  IP Label/Device Path          IP Address

# nodeA / net_diskhb_01
   hdisk1             nodeA_hdisk1_01              /dev/hdisk

# nodeA / net_ether_01
   en0                node1b1                      10.1.2.
   en2                node1b2                      10.1.3.

# nodeA / net_ether_02
   en4                nodeA                        10.1.4.

# nodeB / net_diskhb_01
   hdisk1             nodeB_hdisk1_01              /dev/hdisk

# nodeB / net_ether_01
   en0                node2b1                      10.1.2.
   en2                node2b2                      10.1.3.

# nodeB / net_ether_02
   en4                nodeB                        10.139.64.

```

2. Perform a "Verify and Synchronize HACMP Configuration" operation now. Wait for several minutes and verify that the new **diskhb** device appears in the **c1stat** output.

```

c1stat - HACMP Cluster Status Monitor
-----

Cluster: HA1      (1110178227)
Tue Oct 11 15:07:17 EETDT 2011
State: UP          Nodes: 2
SubState: STABLE

Node: nodeA      State: UP
Interface: nodeA (2)      Address: 10.1.4.4
                          State: UP
Interface: node1b1 (1)    Address: 10.1.2.19
                          State: UP
Interface: node1b2 (1)    Address: 10.1.3.14
                          State: UP
Interface: nodeA_hdisk1_01 (0)  Address: 0.0.0.0
                          State: UP
Interface: node1 (1)      Address: 10.1.1.9
                          State: UP
Resource Group: ResGrpA      State: On line

Node: nodeB      State: UP
Interface: nodeB (2)      Address: 10.1.4.5
                          State: UP
Interface: node2b1 (1)    Address: 10.1.2.12
                          State: UP

```

```

Interface: node2b2 (1)           Address: 10.1.3.15
                                State: UP
Interface: nodeB_hdisk1_01 (0)  Address: 0.0.0.0
                                State: UP
Interface: node2 (1)           Address: 10.1.1.3
                                State: UP
Resource Group: ResGrpA         State: On line

```

3. Verify whether the nodes can communicate through the new disk heartbeat.

#### On nodeA:

```

# /usr/sbin/rsct/bin/dhb_read -p hdisk1 -t
DHB CLASSIC MODE
  First node byte offset: 61440
  Second node byte offset: 62976
  Handshaking byte offset: 65024
    Test byte offset: 64512

Receive Mode:
Waiting for response . . .
Magic number = 0x87654321
Magic number = 0x87654321
Magic number = 0x87654321
Magic number = 0x87654321
Link operating normally

```

#### On nodeB:

```

# /usr/sbin/rsct/bin/dhb_read -p hdisk1 -r
DHB CLASSIC MODE
  First node byte offset: 61440
  Second node byte offset: 62976
  Handshaking byte offset: 65024
    Test byte offset: 64512

Transmit Mode:
Magic number = 0x87654321
Detected remote utility in receive mode.  Waiting for response . . .
Magic number = 0x87654321
Magic number = 0x87654321
Link operating normally

```

## H. Convert shared volume groups to enhanced-concurrent mode

1. Enable the shared volume groups for fast disk takeover, that is, enhanced concurrent mode. First, ensure that the **bos.c1vm.enh** file set is installed on both nodes. Perform the following task for each shared volume group. Afterwards, synchronize the shared volume group definition immediately.

```

# smit hacmp
System Management (C-SPOC)
HACMP Logical Volume Management
Shared Volume Groups
Enable a Shared Volume Group for Fast Disk Takeover
ResGrpA  sharedvgA
* Resource Group Name          ResGrpA
* SHARED VOLUME GROUP name    sharedvgA

# smit cl_lvm
Synchronize a Shared Volume Group Definition
sharedvgA

```

2. Perform a "Verify and Synchronize HACMP Configuration" operation now.
3. Stop cluster services on both nodes.

4. Start cluster services on both nodes. Verify that all shared volume groups are now **varied on in the enhanced-concurrent mode, that is, active-read/write** on the primary node and **active-passive-only** on the standby node.

```

nodeA# lspv
hdisk0          00c342c6c73137a9          rootvg          active
hdisk1          00c69b9edd04057b          diskhbvg       concurrent
hdisk2          00c40a8e97f5cca1          sharedvgA      concurrent
hdisk3          00c40a8e97f5dc8a          sharedvgB      concurrent
hdisk4          00c40a8e04e87819          appvg          active
hdisk5          00c40a8e97f0dda8          appvg          active

nodeB# lspv
hdisk0          00c334b6c7cca4b1          rootvg          active
hdisk1          00c69b9edd04057b          diskhbvg       concurrent
hdisk2          00c40a8e97f5cca1          sharedvgA      concurrent
hdisk3          00c40a8e97f5dc8a          sharedvgB      concurrent
hdisk4          00c69b9e97fe3331          appvg          active
hdisk5          00c69b9e04e9bc20          appvg          active

nodeA# lsvg sharedvgA
VOLUME GROUP:  sharedvgA          VG IDENTIFIER:  00c40a8e00004c000000010280d7161f
VG STATE:      active              PP SIZE:        64 megabyte(s)
VG PERMISSION: read/write          TOTAL PPs:      527 (33728 megabytes)
MAX LVs:       512                FREE PPs:       357 (22848 megabytes)
LVs:           4                  USED PPs:       170 (10880 megabytes)
OPEN LVs:      4                  QUORUM:         2 (Enabled)
TOTAL PVs:     1                  VG DESCRIPTORS: 2
STALE PVs:     0                  STALE PPs:      0
ACTIVE PVs:    1                  AUTO ON:        no
Concurrent:    Enhanced-Capable    Auto-Concurrent: Disabled
VG Mode:       Concurrent
Node ID:       1                  Active Nodes:    2
MAX PPs per VG: 130048
MAX PPs per PV: 1016
LTG size:      128 kilobyte(s)     MAX PVs:        128
HOT SPARE:     no                  AUTO SYNC:      no
BB POLICY:     relocatable         BB POLICY:      relocatable

nodeB# lsvg sharedvgB
VOLUME GROUP:  sharedvgB          VG IDENTIFIER:  00c40a8e00004c000000010280d7161f
VG STATE:      active              PP SIZE:        64 megabyte(s)
VG PERMISSION: passive-only        TOTAL PPs:      527 (33728 megabytes)
MAX LVs:       512                FREE PPs:       357 (22848 megabytes)
LVs:           4                  USED PPs:       170 (10880 megabytes)
OPEN LVs:      0                  QUORUM:         2 (Enabled)
TOTAL PVs:     1                  VG DESCRIPTORS: 2
STALE PVs:     0                  STALE PPs:      0
ACTIVE PVs:    1                  AUTO ON:        no
Concurrent:    Enhanced-Capable    Auto-Concurrent: Disabled
VG Mode:       Concurrent
Node ID:       2                  Active Nodes:    1
MAX PPs per VG: 130048
MAX PPs per PV: 1016
LTG size:      128 kilobyte(s)     MAX PVs:        128
HOT SPARE:     no                  AUTO SYNC:      no
BB POLICY:     relocatable         BB POLICY:      relocatable

```

## I. Migrate HACMP 5.4 to PowerHA 6.1

1. Stop cluster services on both nodes.

```
# smit clstop
```

2. Upgrade HACMP 5.4 to PowerHA 6.1 SP6. Ensure '*ACCEPT new license agreements?*' is changed to **Yes**. Confirm that the HA migration from 5.4 to 6.1 completed successfully by reviewing the **clconvert.log** log file. Verify that the correct HA level is displayed.

```
# mount nim1:/software/PowerHA6.1 /mnt
```

```

# cd /mnt
# smitty update_all
# lslpp -L cluster*server*rte
Fileset                Level  State  Type  Description (Uninstaller)
-----
cluster.es.server.rte  6.1.0.0  C    F    ES Base Server Runtime

# cd /mnt/SP6/
# smitty update_all
# lslpp -L cluster*server*rte
Fileset                Level  State  Type  Description (Uninstaller)
-----
cluster.es.server.rte  6.1.0.6  A    F    ES Base Server Runtime
# cd /
# umount /mnt

# view /tmp/clconvert.log

Command line is:
/usr/es/sbin/cluster/conversion/cl_convert -F -v 5.4.1

No source product specified.
Assume source and target are same product.
Parameters read in from command line are:
Source Product is HAES.
Source Version is 5.4.1.
Target Product is HAES.
Target Version is 6.1.0.
Force Flag is set.
...etc..
odmdelete -o HACMPtopsvcs: 0518-307 odmdelete: 1 objects deleted.
odmadding HACMPtopsvcs
odmdelete -o HACMPcluster: 0518-307 odmdelete: 1 objects deleted.
odmadding HACMPcluster

*****
*** End of ODM Manager ***
*****

Done execution of ODM manipulator scripts.

Cleanup:
  Writing resulting odms to /etc/es/objrepos.
  Restoring original ODMDIR to /etc/es/objrepos.
  Removing temporary directory /tmp/tmpodmdir.

Exiting cl_convert.

Exiting with error code 0. Completed successfully.

----- end of log file for cl_convert: Sat Oct 15 16:32:33 EETDT 2011

# hlevel -s
6.1.0 SP6

```

Now update the second node before proceeding to the next step.

3. Reboot both nodes now.
4. Start cluster services on both nodes. Verify that the cluster is stable with **clstat**.
5. Perform a "Verify and Synchronize HACMP Configuration" operation now.
6. Upgrade the Tivoli Storage Manager client for PowerHA.
  - a. Upgrade the Tivoli Storage Manager client on both cluster nodes. If the Tivoli Storage Manager client is installed on your cluster nodes, refer to [my blog](#) for the procedures for upgrading the Tivoli Storage Manager client in a PowerHA environment.

b. Perform a "Verify and Synchronize HACMP Configuration" now.

## J. Test cluster resource group move

1. Move a resource group to another node, that is, move resource group A on nodeA to **nodeB**. Verify that the move is complete.

```
# smit hacmp
System Management (C-SPOC)
Resource Group and Applications
Move a Resource Group to Another Node / Site
Move Resource Groups to Another Node
ResGrpA                ONLINE                nodeA /
nodeB
# clGRinfo
# clstat
# df
# lsvg -o
# ping serviceipaddressssh to serviceipaddress (should connect to second node in cluster)
```

2. Move the resource group back to the home node, for example, **nodeA**. Verify that the move is complete.

```
# clGRinfo
# clstat
# df
# lsvg -o
# ping serviceipaddressssh to serviceipaddress (should connect to home node in cluster)
```

## K. Test cluster failover

1. Perform standard cluster failover tests (as defined by the customer) - including application testing for failover scenarios.
2. Change the old POWER6 LPAR profile definition to boot in the system management service (SMS) mode and rename to the LPAR profile to *OLD\_LPARname*. This helps to prevent the old LPAR from being started by mistake (until the Power 595 server is decommissioned).
3. Hand over to application team to test their applications.
4. Ensure that the migration is complete.

## Summary

In this article you have learnt how to migrate an existing two-node PowerHA cluster from POWER6 to POWER7, how to upgrade PowerHA, and how to convert your shared volume groups to enhanced-concurrent mode.

## Resources

The following resources were referenced during the planning phase of the PowerHA to POWER7 migration project.

- [PowerHA Upgrade Do's and Dont's](#)
- [PowerHA Upgrading on an offline cluster](#)
- [PowerHA Cluster Manager \(HACMP\) Service Packs](#)
- [HACMP Version Compatibility Matrix](#)

- [PowerHA for AIX v6.1 Release Notes](#)
- [IBM Software support lifecycle: High Availability Cluster Multi-Processing \(HACMP\) for AIX 5.4.x](#)
- [Reviewing cluster message log files](#)

© Copyright IBM Corporation 2013

([www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml))

[Trademarks](#)

([www.ibm.com/developerworks/ibm/trademarks/](http://www.ibm.com/developerworks/ibm/trademarks/))