

Upgrading from POWER5 to POWER6

Planning and preparation for a POWER6 upgrade: One customer's experience

Skill Level: Introductory

[Chris Gibson](mailto:chrisgibson@gibsonnet.net) (chrisgibson@gibsonnet.net)

AIX Specialist

Australia Post

09 Jun 2009

Read about and learn from my experiences in upgrading a POWER5® 595 to a new POWER6® 595.

Introduction

This article discusses my experiences when upgrading a POWER5 595 to a new POWER6 595. This is not intended as an official "how-to" guide, but a discussion on how I performed the upgrade and what decisions and considerations I made during the planning and execution phases. I hope that this information will help others who need to perform similar tasks within their own organizations or those of their customers.

Let me start off by stating that each environment is different. Most sites customize their AIX® operating system and POWER® hardware configuration to meet their requirements, so what I describe here may not match what you have in your environment. So please use your best judgment and apply what you need from this article, but only if it is appropriate for your site. Only you can make this call, as you know more about how your AIX and POWER infrastructure is configured (and why) than anyone else!

An important note about my AIX environment: All of my LPARs were virtualized; that is, they were all micro-partitioned and they all used Virtual I/O (VIO) for all disk and

network devices. None of my AIX LPARs had any dedicated physical hardware. All physical devices were owned by the Virtual I/O servers (VIOS).

I will outline the steps I performed before and after the POWER6 upgrade. The focus is on the VIOS, AIX, and HACMP tasks I executed after the hardware had been physically upgraded to POWER6.

Upgrade overview

I needed to upgrade my existing System p® 595 (9119-595) landscape to the new POWER6 595 (9119-FHA). Please refer to Figure 1 below. When I say upgrade, I mean this was an MES upgrade. MES stands for Miscellaneous Equipment Specification. An MES upgrade includes any server hardware change, which can be an addition, improvement, removal, or any combination of these. An important feature of an MES upgrade is that the systems serial number does not change.

Figure 1. 9119-595 and the 9119-FHA



Essentially, our upgrade from POWER5 to POWER6 involved moving the existing I/O drawers (including internal disks, FC, and Ethernet adapters) from the POWER5 frame to the POWER6 frame. Once this was completed, the POWER6 system would be powered up and the IBM® CE (Customer Engineer) would then hand back the

system to me. Then I would attempt to bring up the LPARs on the new POWER6 server.

This was the first time I had migrated to a newer POWER platform using the MES upgrade method, and I had concerns.

In the past I had migrated AIX systems to newer platforms with both the old and new systems sitting side by side. For example, several years ago, when [migrating from POWER4 to POWER5](#), we purchased a new 9119-595 and sat it next to the old p4 p690. We connected the new 595 to our SAN and network and started moving LPARs from the p690 (one at a time) by restoring a mksysb using Network Installation Manager (NIM). The advantage with this method was if we had an issue on the new 595, we could easily fallback to the p690, as the original LPAR was still available. It also allowed us time to test the 595 before we unleashed any workload onto the system. This gave us greater confidence that all our components were compatible (such as software and firmware) and functioning as expected. It essentially gave us time to shake out any bugs or issues with the new hardware.

This method was what I considered, at the time, my preferred way of performing the migration to POWER6.

With the MES upgrade method, the old p5 system would be shut down, rolled out the door, and the new p6 moved into its place. The IBM CE would then transfer the I/O drawers, configure the system, verify it was OK, hand it back to me, and walk away (so to speak!). With the *'big bang'* upgrade approach, I would not be able to rehearse or test the upgrade process and there was potential to be caught out by unknown issues.

My main concern here was that if there was a problem with the 9119-FHA, we did not have a way to easily back out to the old system. We could not simply power up the old p5 and start the LPARs. Nor could we test that the new hardware was functioning OK, well in advance, before activating the LPARs and running production workload.

Given that this was an MES upgrade and that wasn't going to change, I set about planning for the upgrade.

Planning and preparation

The most pressing decision I had to make was what migration approach I was going to use for my AIX LPARs. I had two choices here; I could either rebuild the VIOS and the LPARs from a mksysb restore or attempt to boot them from disk.

I understood that the only documented and official method to migrate LPARs to newer or different hardware was using a "mksysb clone" operation, which means

taking a mksysb of the LPAR and restoring it on the new p6 system. However, I was interested in simply booting the LPARs on the new p6 595.

This was not guaranteed to work and I could certainly understand why. In order to boot on the new system, you would need the appropriate device driver filesets to support the new platform. This meant you would need to ensure that all your systems were installed with `Enable System Backups` to install any system set to `Yes`. This enables systems to be installed on any other system (using cloning) by installing all devices and kernels. No guarantee is implied by this setting. However, when I think about how Live Partition Mobility works, and the fact that you can move a Virtual I/O client (VIOC) LPAR from one physical system to another (without a mksysb restore), I wonder if this may change in the future? This is the default setting when installing AIX and I had always ensured it was set when loading the operating system. You can refer to the `/usr/lpp/bosinst/bosinst.template.README` file for more details.

Some evidence on the AIX forums suggested that this method may work. One customer had reported using this method when they upgraded from a p5 570 to a p6 570.

Some other considerations when choosing the migration approach were around the I/O bus numbering and LPAR profiles. According to the 9119-595 to 9119-FHA MES upgrade instructions, the I/O bus numbering did not change after the upgrade. Were my LPAR profiles going to be recovered, and intact, on the new p6, or would I need to rebuild? The MES upgrade instructions stated the IBM CE should perform a `Recover Partition Data` operation, using the HMC, after the upgrade. This meant I would not have to recreate all of my LPAR profiles from scratch (either using the System Planning tool (SPT) or a scripting method). I also knew that the system serial number was guaranteed not to change, so I wasn't going to have application licensing problems because the system ID had changed.

I finally settled on my approach to the upgrade. I would boot my LPARs (including my VIOS) and use mksysb restore only if I had serious issues bringing up the systems (in a clean state). My procedure would be:

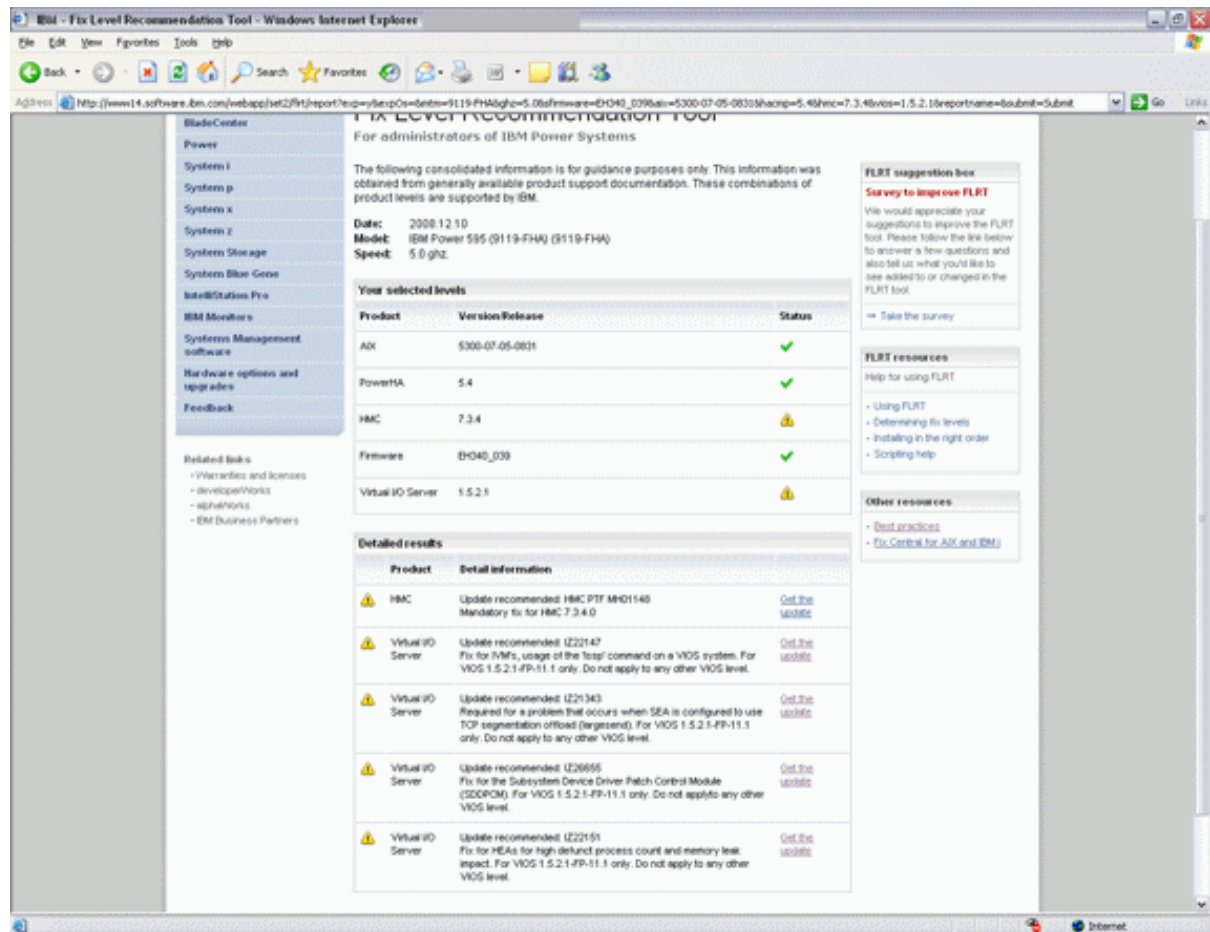
- Document the current virtual device mappings on each VIOS. I used `lsmap -all` and `lsmap -all -net` for this task.
- Collect PVID and LUNID information for all the Virtual target devices backed by physical disks.
- Verify that all Virtual Slot IDs were greater than 10. Starting with HMC v7, the HMC reserves the first ten virtual adapter slots on each VIOS for internal HMC use.
- Take a mksysb of all LPARs and VIOS. Use these mksysb images to restore from, if required.

- Back up the managed systems partition profile data.
- IBM CE to perform the hardware upgrade to POWER6.
- IBM CE to restore the managed systems profile data from the previous backup.
- Verify the partition profile data for each AIX LPAR and VIOS is correct on the HMC.
- Upon successful verification of the LPAR and VIOS profiles, boot each VIOS. Enter the SMS menu, confirm the boot list, and boot the VIOS.
- Verify the virtual device configuration and status on each VIOS. Perform a health check on each VIOS. If the health check is not successful, then restore the VIOS from mksysb using NIM.
- Upon successful verification of each VIOS, boot each LPAR. Enter the SMS menu, confirm the boot list, and boot the LPAR.
- If booting an LPAR failed, then restore the LPAR from a mksysb image.
- Correct the boot list on each LPAR. Start functional verification of the environment, such as VIOS failover and application startup and test.

I had to ensure that the process would be executed with a great deal of care and attention. If I had any unforeseen issues that I could not resolve in a timely manner, I would revert to mksysb restore immediately.

Another area of consideration was around the appropriate software and hardware levels to support the new platform. I needed to make sure that I had the correct levels installed before the p6 upgrade. I used the IBM Fix Level Recommendation Tool (FLRT) to determine what software and firmware levels were compatible with the 9119-FHA. The FLRT provides minimum recommended fix level information on key components of IBM Power Systems. I highly recommend this tool when planning any type of AIX or POWER upgrade activity. The tool generates reports you can use when planning upgrades. Refer to Figure 2.

Figure 2. FLRT report



(click to enlarge)

In the months leading up to the upgrade, we updated the following components to the following levels, in the following order:

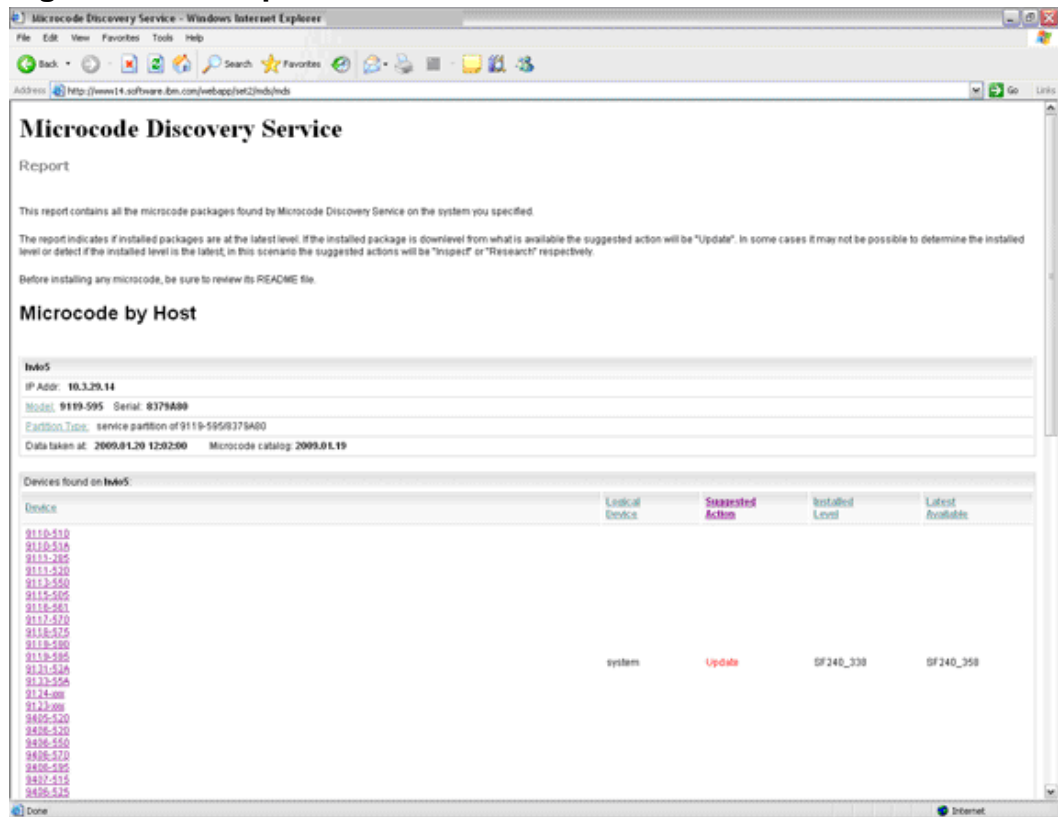
1. HMC V7R3.4.0 + MH01152
2. Firmware Updated various H/W (for example, FC, SCSI and Ethernet adapters)
3. VIOS 1.5.2.1-FP-11.1 + SDDPCM 2.2.0.4
4. AIX 5300-07-05-0831
5. HACMP 5.4.1.3 + RSCT fixes

Prior to the upgrade, I captured a plethora of configuration information relating to the 595, AIX, VIOS, HACMP, and the HMC in my environment. If I needed to recover any or all of the systems from scratch, for whatever reason, I wanted to be well

prepared and have a wealth of information at hand should I need it. The following highlights just some of the data I collected using scripts and other methods:

- Ran my *AIXinfo* script to collect vast amounts of information relating to the AIX configuration of each and every LPAR. The script ran several system commands, such as `oslevel`, `lppchk`, `instfix`, `lsconf`, `lscfg`, `lsdev`, `lsattr` and many more. The information was stored in a text file on another system.
- Created a Microcode Discovery Service (MDS) Report for each VIOS and any LPAR that contained physical hardware like FC adapters or SCSI disks. This involved downloading that latest microcode catalog file from the IBM support website, running the `invscout` command on each VIOS/LPAR and then uploading the resultant MUP file into the online MDS reporting tool, which is shown in Figure 3. The MDS tool determines if microcode installed on your systems is at the latest level.

Figure 3. MDS report



(click to enlarge)


- Captured information from the HMC, such as LPAR profile information (CPU/Memory allocation), managed system properties, Physical I/O adapter assignment, Virtual adapter definitions. This could be captured

using the SPT or simple screen captures from the HMC.

- Virtual I/O mapping and configuration (`lsmap -all` and `lsmap -all -net`) output. I wrote a script to capture a whole bunch of data relating to the VIOS configuration and devices, such as Shared Ethernet Adapter (SEA) settings, VTD mapping, `pcmpath` output, vhost slot numbers, and more. I also captured the location codes for the VIOS rootvg disks, which proved to be an important step, as shown later.
- Cable locations and labels. The IBM CE was going to disconnect all of my SAN and network connections from all of my adapters in order to move the I/O drawers to the new frame. I created a map of each cable and which adapter it was plugged into. I also made sure that each cable was labeled so that I could check it had been plugged back into the correct adapter.
- Build documentation. I had my original system build documentation on hand in case I needed to refer to it. This outlined how the systems had been built and configured originally.
- HACMP information. I had several HACMP nodes on the frame so I captured cluster information from such commands as `clstat`, `cltopinfo`, `clsnapshot`, `cldump`, `clRGinfo`, and `cldisp`. I also exported the cluster configuration using the HACMP Online Planning Worksheets, such as `# smit cl_export_def_olpw.dialog`. The HACMP configuration for each cluster was also verified and synchronized prior to the upgrade.
- Review the AIX and VIOS error report to check for any serious errors with the `errpt` and `errlog` commands. Catching (and resolving) these sorts of issues before a major upgrade can save you from headaches later on.
- Check the HMC for any open hardware events, like `$ lssvcevents -t hardware`.
- I ran the HMC readiness checker to identify any 595 hardware issues that may impact the upgrade. You'll find this task on the HMC under "Updates." Once you select a managed system, you can click on **Check system readiness**.

Of course, I also ensured that I had a good backup of all the components involved in the upgrade, such as a `mksysb` of all AIX LPARs, a `savevg` of all volume group structures, a data backup of all applications and databases, a backup of each VIOS, a HMC backup to DVD-RAM, and similar. Most importantly, I performed a backup of the managed systems partition profile data, using the HMC. See Figure 4 for the backup partition data. This would be a critical step, as the IBM CE would use this backup to recover my partition data after the upgrade. Without it, I would have to rebuild all of my LPAR profiles again.

Figure 4. Backup partition data



POWER6 upgrade

On the day of the upgrade, I shut down all the LPARs and handed the system over to the IBM CE. He spent the next six hours performing the hardware upgrade. We had HACMP clustered systems, so production workload was processed on another 595 while this one was down for the upgrade.

When the CE gave the system back to me, the first thing I did was check that all of my cables were plugged in correctly. They were. Next I verified that all my LPAR profiles had been recovered successfully on the new system. They had, as shown in [Figure 5](#). I double checked each profile and found that the partition id and the bus numbers had not changed – refer to [Figure 6](#). Also, the location codes for all adapters (e.g. TY-P1-C02) had not changed. The serial number, as promised, had not changed. This was all good news!

Figure 5. HMC view of LPARs on the 9119-FHA

Select	Name	ID	Status	Processing Units	Memory (GB)	Active Profile	Environment
<input type="checkbox"/>	bvio1	2	Running	0.5	2	default	Virtual IO Server
<input type="checkbox"/>	bvio2	3	Running	0.5	2	default	Virtual IO Server
<input type="checkbox"/>	bxb01adm	4	Running	2	100	default	ADK or Linux
<input type="checkbox"/>	bxb02adm	5	Running	2	13	default	ADK or Linux
<input type="checkbox"/>	bxb03	6	Running	0.5	2	default	ADK or Linux
<input type="checkbox"/>	bxb04adm	7	Running	0.1	1	default	ADK or Linux
<input type="checkbox"/>	bxb05adm	10	Running	0.1	1	default	ADK or Linux
<input type="checkbox"/>	bxb06adm	8	Running	0.2	8	templememory	ADK or Linux
<input type="checkbox"/>	bxb07adm	9	Running	1	1	default	ADK or Linux
<input type="checkbox"/>	bxb08adm	1	Running	0.2	1	default	ADK or Linux

(click to enlarge)

Figure 6. 595 I/O bus numbering

Unit	Bus	Slot	Description	I/O Pool Id
U5791.001.9980XTY-P1	10	T6	SCSI bus controller	Unassigned
U5791.001.9980XTY-P1	10	C08	4 Gb Dual-Port Fibre Channel PCI-X 2.0 DDR Adapter	Unassigned
U5791.001.9980XTY-P1	10	C09	Ethernet controller	Unassigned
U5791.001.9980XTY-P1	10	C10	SCSI bus controller	Unassigned
U5791.001.9980XTY-P1	11	C01	Ethernet controller	Unassigned
U5791.001.9980XTY-P1	11	C02	Ethernet controller	Unassigned
U5791.001.9980XTY-P1	11	C03	4 Gb Dual-Port Fibre Channel PCI-X 2.0 DDR Adapter	Unassigned
U5791.001.9980XTY-P1	11	C04	RAID Controller	Unassigned
U5791.001.9980XTY-P1	12	T5	SCSI bus controller	Unassigned
U5791.001.9980XTY-P1	12	C05	4 Gb Dual-Port Fibre Channel PCI-X 2.0 DDR Adapter	Unassigned
U5791.001.9980XTY-P1	12	C06	Ethernet controller	Unassigned
U5791.001.9980XTY-P1	12	C07	4 Gb Dual-Port Fibre Channel PCI-X 2.0 DDR Adapter	Unassigned
U5791.001.9980XTY-P2	13	T6	SCSI bus controller	Unassigned
U5791.001.9980XTY-P2	13	C08	4 Gb Dual-Port Fibre Channel PCI-X 2.0 DDR Adapter	Unassigned
U5791.001.9980XTY-P2	13	C09	Ethernet controller	Unassigned

Total: 72 Filtered: 72

As expected, the properties of the 595 showed the "Type/Model" had changed, from 9119-595 to 9119-FHA. See Figures 7 and 8 below for illustrations of this.

Figure 7. POWER5 Type/Model

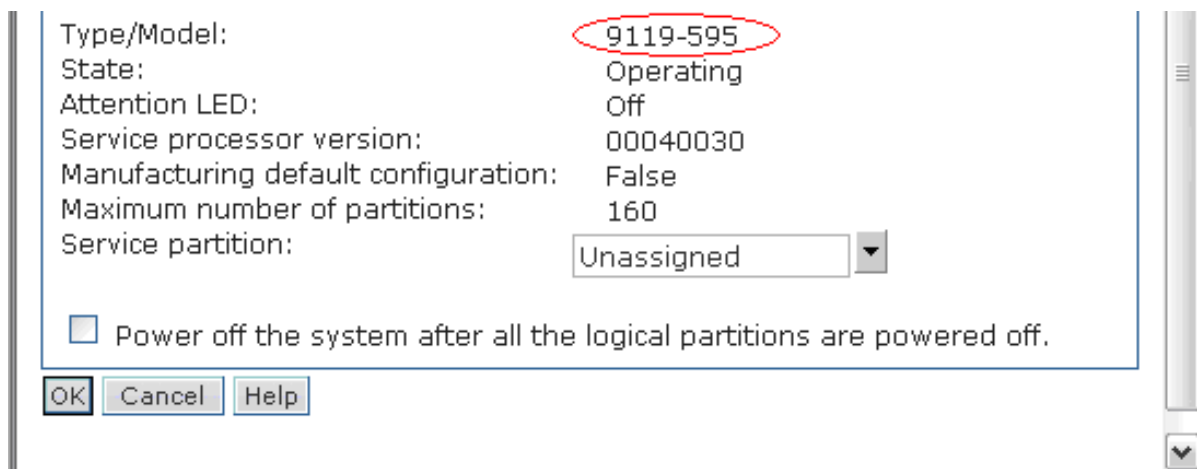
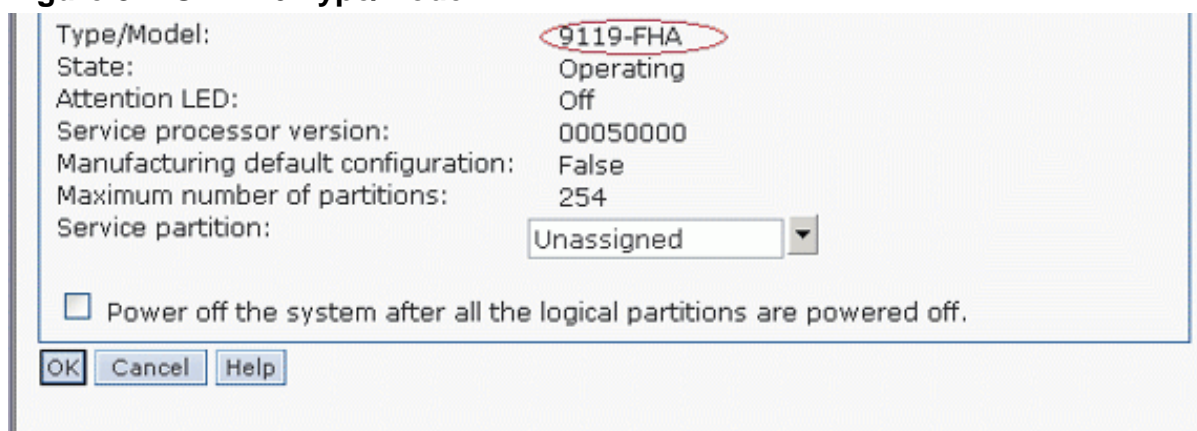


Figure 8. POWER6 Type/Model



Booting each VIOS

The next step was to boot the VIOS. I activated one of the VIOS and waited while it started to boot. This took a long time, as I had several disks (100+) assigned to my VIOS, and four paths to the SAN, all of which had to be scanned. After several minutes, I finally entered the SMS menu and confirmed that the boot list contained the correct rootvg disk for this VIOS (which is where my documentation came in very handy). I needed to be careful here.

I noticed several disks were identified as having AIX installed on them, but they did not belong to my VIOS rootvg! These were the rootvg disks that belonged to my client LPARs. If I picked the wrong disk, I'd boot the wrong system. This is why it was important to collect the location codes for my VIOS rootvg disks prior to the upgrade.

I exited SMS and let the VIOS boot as normal. The system came up without any errors. I logged in using the console as padmin and ran `lsmapi -all` and `lsmapi -all -net`. All of my virtual adapter mappings and SEAs were available. The only

difference I observed was that the vhost location codes had changed slightly, from 595 to FHA (as shown below); however, the slot numbers and serial number were identical.

```
< vhost0          U9119.595.8369B40-V2-C20          0x00000000
---
> vhost0          U9119.FHA.8369B40-V2-C20          0x00000000
```

I noticed that the boot list on each VIOS had changed. Only hdisk0 and the first network adapter (ent0) were in the list. The root volume group on the VIOS was mirrored, so I needed to change the boot list to include both bootable disks in rootvg: `# bootlist -m normal hdisk0 hdisk8`. This was expected, as the NVRAM (which contains the boot list) is not carried over from the old p5 to the new p6.

VIOS health checks

Before starting my LPARs, I first performed several health checks on each VIOS. The health checks are shown in Table 1. These steps were to check the general health of the VIOS. I was looking for anything abnormal, such as devices in a *Defined* state or permanent hardware errors in the error log.

Table 1. VIOS health checklist

VIOS checks	Description
<code>lsconf</code>	Ensure processor shows POWER6
<code>ioslevel</code>	Confirm VIOS level
<code>cfgdev</code>	Check for any missing device filesets
<code>lsmap -all</code>	Check VTD mapping
<code>lsmap -all -net</code>	Check SEA configuration
<code>lsdev -virtual</code>	Check Virtual devices are Available
<code>lsdev -type disk</code>	Check disk devices are Available
<code>lsdev -type adapter</code>	Check adapters are Available
<code>entstat -all</code>	Check SEA status and priority
<code>errlog</code>	Review errlog for any critical errors
<code>lsdev -dev fscsiX -attr</code>	Review device attributes
<code>lsdev -dev hdiskX -attr</code>	Review disk device attributes
<code>lspv</code>	Check disk pvids
<code>pcmpath query adapter</code>	Check SDDPCM adapter status
<code>pcmpath query device</code>	Check SDDPCM disk device status and

	LUNIDs
<code>pcmpath query wwpn</code>	Check FC adapter WWPNS known to SDDPCM
<code>pcmpath query version</code>	Check SDDPCM level
<code>lspath</code>	Check all disk paths are enabled
<code>netstat -nr</code>	Check default gateway

Booting the LPARs

Once both my VIOS were activated and I'd verified the virtual adapter mappings and health checks, I started an AIX LPAR. I booted the LPAR into SMS and reviewed the boot list and verified that the correct rootvg disk was in the list. It was. The LPAR booted as normal.

AIX health checks

I performed several health checks on each AIX LPAR. Again, I verified that the operating system was operating as expected. The health checks are shown in Table 2. The only change I had to make as a result of the upgrade was to reset the boot list, for LPARs with a mirrored rootvg (as I did on the VIOS): `# bootlist -m normal hdisk0 blv=bos_hd5 hdisk1 blv=bos_hd5`. Pay particular attention to this if you use multibos on your AIX systems, where you may have two boot logical volumes (BLVs). Choose the correct disk partition to boot from (for example, `part=2` or `part=4`, shown below); otherwise, you could boot from an older image of your AIX system.

```
PowerPC Firmware
Version EH340_039
SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-----
Select Device
Device   Current   Device
Number  Position  Name
1.      -        SCSI 136 GB Harddisk, part=2 (AIX 5.3.0)
        ( loc=U9119.FHA.8369B40-V22-C46-T1-L8100000000000000 )
2.      -        SCSI 136 GB Harddisk, part=4 (AIX 5.3.0)
        ( loc=U9119.FHA.8369B40-V22-C46-T1-L8100000000000000 )
```

I recommend that you remove multibos instances (`multibos -R`) prior to the upgrade to avoid confusion.

Table 2. AIX health checklist

AIX Checks	Description
<code>lsconf</code>	Ensure processor shows POWER6

cfgmgr	Check for any missing device filesets
lsdev -C grep Defined lsdev -Cc adapter lsdev -Cc disk	Check devices are in the correct state
lparstat -i	Verify LPAR config
oslevel -s	Verify AIX SP level
oslevel -r	Verify AIX TL.
bootlist -m normal -o	<p>Check bootlist settings.</p> <p>Expect output similar to:</p> <ul style="list-style-type: none"> - For mirrored rootvg (including the VIOS): # bootlist -m normal -o hdisk0 blv=bos_hd5 hdisk4 blv=bos_hd5 - For SAN boot rootvg: # bootlist -m normal -o hdisk0 blv=bos_hd5 hdisk0 blv=bos_hd5
instfix -i grep AIX instfix -i grep SP instfix -icqk 53-07-050831_SP grep ":-:" instfix -icqk 5300-07_AIX_ML grep ":-:"	Detects any missing AIX TL or SP filesets
lsvg -l rootvg grep stale	Check for 'stale' partitions in rootvg
df mount lsvg lsvg -il grep close	Check all necessary file systems are mounted.
df /var	Check for full /var
df /tmp	Check for full /tmp
sysdumpdev -l	Check for proper system dump config
lsattr -El mem0	Ensure proper memory configuration
emgr -l	Check 'efix' inventory if any
lppchk -v -m3	Verify installed filesets
lppchk -c -m3	Sum check installed filesets
lsps -a	Verify paging space
smtctl	Verify smt on/off all the nodes

<code>lssrc -ls xntpd grep "Reference Id"</code>	Check ntp on the nodes
<code>vmo -a</code>	Check all vmo settings
<code>no -a</code>	Check all network options settings
<code>netstat -nr</code>	Check route table
<code>lspath</code>	Check all paths are enabled
<code>errpt</code>	Check for any permanent hardware errors or other critical errors
<code>tail -100 /var/log/syslog</code>	Check for any critical errors
<code>alog -of /var/adm/ras/conslog</code>	Review errors in the console log.
<code>cat /etc/qconfig lpstat</code>	Check print queues are defined and enabled.
<code>lssrc -a</code>	Review <i>active</i> and <i>inoperative</i> subsystems.
<code>pstat -a grep aio lsattr -El aio0</code>	Verify AIO is configured.
<code>lsattr -El sys0 -a minpout -a maxpout</code>	Check I/O pacing settings.
<code>date</code>	Verify the date and time are correct.
<code>echo \$TZ grep TZ /etc/environment</code>	Check the TZ variable is set appropriately.
<code>lsattr -El hdiskX -a queue_depth</code>	Check queue depths on all hdisks.
<code>topas -C</code>	Run 'topas -C' from an LPAR. Verify all LPARs report performance util data.

Using the `lsconf` command, I was able to quickly confirm that the LPARs were now running on the new POWER6 platform. Refer to the output below, from one of the LPARs before and after the upgrade. The "System Model" had changed from 9119-595 to 9119-FHA, along with the processor type and speed.

Prior to p6 upgrade:

```
System Model: IBM,9119-595
Machine Serial Number: XXXXXXXX
```

```
Processor Type: PowerPC_POWER5
Number Of Processors: 4
Processor Clock Speed: 2302 MHz
CPU Type: 64-bit
Kernel Type: 64-bit
LPAR Info: 6 bxaix03
Memory Size: 2048 MB
Good Memory Size: 2048 MB
Platform Firmware level: SF240_338
Firmware Version: IBM,SF240_338
```

Post-p6 upgrade:

```
System Model: IBM,9119-FHA
Machine Serial Number: XXXXXXXX
Processor Type: PowerPC_POWER6
Number Of Processors: 4
Processor Clock Speed: 5000 MHz
CPU Type: 64-bit
Kernel Type: 64-bit
LPAR Info: 6 bxaix03
Memory Size: 2048 MB
Good Memory Size: 2048 MB
Platform Firmware level: EH340_039
Firmware Version: IBM,EH340_039
```

VIOS failover verification

Satisfied that each VIOS was in good shape and my VIO client (VIOC) LPARs were running fine, I performed several VIOS failover tests after the upgrade. This was to ensure that redundancy of the dual VIOS setup had not been compromised as a result of the upgrade. Some of the tests included:

- Shutdown one VIOS and ensure that all client LPARs were not impacted. For example, SEA failover, IP connectivity OK, loss of path (and/or mirror), disk traffic OK.
- Restart the VIO and ensure fallback OK. For example, SEA fallback, IP connectivity OK, path recovery (and/or re-sync mirror), disk traffic OK.
- Ensure that any LPARs with a mirrored rootvg are (re)synced before and after each VIOS shutdown/restart.
- Repeat the same verification procedures for the second VIOS.

I also unplugged each FC and network cable (one at a time) to ensure that disk and network I/O on the VIOCs was not impacted.

Plan B

If booting from disk had been unsuccessful, for whatever reason, I would have instigated Plan B. This would have involved restoring the VIOS and the LPARs from a mksysb using NIM, on the new POWER6 platform. The NIM master was located on a different 595. In fact, I tested this on a subsequent POWER6 upgrade and it worked equally well as just booting from disk. No reconfiguration (apart from the boot list) was required, even for the VIOS. If you do need to restore a VIOS from mksysb, it's a good idea to create a new SPOT on your NIM master for each VIOS. Create the SPOT from a mksysb of the VIOS (as shown below).

```

Define a Resource

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

* Resource Name          [Entry Fields]
* Resource Type          [hvio3-spot]
* Server of Resource      spot
* Source of Install Images [master]          +
* Location of Resource    [hvio3-mksysb]       +
*                         [/export/nim/spot] /
...

```

Also ensure that you change "Remain NIM client after install?" to no when configuring the VIOS NIM client for a BOS installation (as shown below). This will prevent an IP address being configured on the physical network adapter used for the install. If an IP was accidentally configured on this physical interface and it was part of a SEA configuration, the SEA may fail to configure as the physical device is already in use.

```

Install the Base Operating System on Standalone Clients

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

* Installation Target    [Entry Fields]
* Installation TYPE      hvio3
* SPOT                   mksysb
LPP_SOURCE               hvio3-spot
MKSYSB                   []          +
BOSINST_DATA to use during installation [ ]          +
IMAGE_DATA to use during installation  [ ]          +
RESOLV_CONF to use for network configuration [ ]          +
Customization SCRIPT to run after installation [ ]          +
Customization FB Script to run at first reboot [ ]          +
  ACCEPT new license agreements? [yes]          +
  Remain NIM client after install? [no]          +
...

```

As I mentioned earlier, under [Bootimg each VIOS](#), I'd need to pay careful attention to which rootvg disks to restore the VIOS mksysb. There were many, many disks attached to my VIOS, some of which had an AIX image on them. If I picked the wrong disk, I'd overwrite my client LPARs AIX rootvg. Again, it was very important to document the location codes of my VIOS rootvg disks prior to restoring the VIOS

mksysb.

For both the VIOS and the LPARs, also make sure that `Recover Devices` is set to `Yes` in the BOS installation menu. This will ensure that all devices are recovered during the mksysb restore, so for the VIOS, this will ensure that your virtual adapter mappings are restored. Additionally, for the AIX LPARs, check that `Import User Volume Groupsc` is also set to `Yes`. This will import your non-rootvg volume groups during the restore. Both `Recover Devices` and `Import User Volume Groups` will be set to `Yes`, if you are restoring to the same system (if the serial number is the same, which was true in my case).

Note that after a mksysb restore, on AIX 5.3, you will need to reconfigure the Asynchronous I/O (aio0) device (if you use it). This behavior has changed under AIX 6.1

Post upgrade

Once the POWER6 upgrade was completed and all verification and health checks had been completed successfully, I re-integrated the HACMP nodes back into the cluster and performed failover and fallback tests. No issues were discovered.

With the upgrade complete, there were only a few post-upgrade tasks to perform. Some of these included:

- **Backing the up LPAR profile data, again!**
- Performing a backup of the HMC.
- Taking a mksysb of each VIOS and all the AIX LPARs.
- Reviewing any 'open' hardware events on the HMC.

Summary

I was initially concerned with the approach for this upgrade. However, having come through it unscathed, I can now say that in some cases an MES upgrade and booting the LPARs from disk is certainly an option worth considering, particularly if your LPARs are all VIO clients, use shared processors, and do not have any dedicated physical devices of any kind. There are no guarantees, so you should choose carefully and test thoroughly in your environment. Of course, mksysb clone is the supported way to migrate.

Based on my experience, both methods achieved the same satisfactory result. One downside is that mksysb restore will require more time. If you have lots of LPARs on a frame, the down-time required for the restore may not be palatable in some case,

such as if you have systems that are not clustered for High Availability.

Ultimately, how you upgrade from POWER5 to POWER6 is up to you. If you are unsure as to what method to use or just need help or guidance, talk to your friendly IBM support folks or your IBM Business Partner. I hope this article will help others who need to tackle similar upgrades.

Resources

Learn

- [Upgrade terminology for System p](#) provide definitions of IBM upgrade terms for System p.
- The [IBM Fix Level Recommendation Tool \(FLRT\)](#) can be useful for those who are planning to upgrade key components of a POWER system.
- [IBM Microcode Discovery Service \(MDS\)](#) generates a comparison report listing microcode subsystems that may need to be updated.
- Read the [PowerHA \(HACMP\) supports IBM Power 595 \(9119-FHA\)](#)
- [Cloning a system](#)
- [New to AIX and UNIX?](#): Visit the "New to AIX and UNIX" page to learn more about AIX and UNIX.
- [AIX Wiki](#): A collaborative environment for technical information related to AIX.
- [The AIX and UNIX developerWorks zone](#) provides a wealth of information relating to all aspects of AIX systems administration.
- [developerWorks technical events and webcasts](#): Stay current with developerWorks technical events and webcasts.
- [Podcasts](#): Tune in and catch up with IBM technical experts.

Discuss

- Find more information on the [IBM developerWorks PowerVM Forum](#).
- Participate in the AIX and UNIX forums:
 - [AIX Forum](#)
 - [AIX Forum for developers](#)
 - [Cluster Systems Management](#)
 - [IBM Support Assistant Forum](#)
 - [Performance Tools Forum](#)
 - [Virtualization Forum](#)
 - More [AIX and UNIX Forums](#)

About the author

Chris Gibson

Chris Gibson is a PowerVM/AIX specialist located in Melbourne, Australia. He is an IBM CATE, System p platform and AIX 5L, and a co-author of the IBM Redbooks publication, "NIM from A to Z in AIX 5L".