

IBM Power Systems Technical University
featuring IBM AIX and Linux

September 8 – 12, 2008 – Chicago, IL



Session Title: Upgrading from AIX v5.3 to AIX v6

Session ID: PAI21

Speaker Name: Jaqui Lynch



© 2008 IBM Corporation

Upgrading from AIX v5.3 to AIX v6

Jaqui Lynch
Systems Architect
Mainline Information Systems
Jaqui.lynch@mainline.com

Mainline: solutions you need
from people you trust 2

Agenda

- Prerequisites and Planning
- Options for Migration
- The Migration Process itself
- Post Migration steps
- Tunables
- Gotchas

Mainline: solutions you need
from people you trust 3

Why bother?

- All releases prior to AIX v5.3 are out of standard service
- New hardware and function support
- PowerVM requires AIX v5.3 or v6 and POWER5 or POWER6
- New functions in v6
 - WPARs
 - RBAC
 - Encrypted JFS2 filesystem (EFS)
- NOTE
 - AIX v5.3 and v6 do not support MCA, PReP or ISA hardware

Mainline: solutions you need
from people you trust 4

AIX 6.1 Features

AIX 6.1 Feature

PowerVM Workload Partitions
PowerVM Live Application Mobility
Application Storage Keys
Kernel Storage Keys
Automatic Variable Page Size
Firmware Assisted Dump
Hardware Decimal Floating-Point
Role Based Access Control
Encrypting Filesystem
Trusted AIX
probevueDynamic Tracing

Platforms Supported

POWER4, PPC970, POWER5 and POWER6
POWER4, PPC970, POWER5 and POWER6
POWER6 (also supported by AIX5.3)
POWER6
POWER6
POWER6
POWER6 (also supported by AIX5.3)
POWER4, PPC970, POWER5 and POWER6
POWER4, PPC970, POWER5 and POWER6
POWER4, PPC970, POWER5 and POWER6
POWER4, PPC970, POWER5 and POWER6

Taken from:

<ftp://ftp.software.ibm.com/common/ssi/pm/sp/n/pod03007usen/POD03007USEN.PDF>

Mainline: solutions you need
from people you trust 5

Prerequisites

- Runs on POWER4, POWER5, POWER6 and PPC970
- Version changes but it is a free upgrade if your AIX v5 SWMA is current
- Check your applications are available and supported
 - If you require certification then check that also
- Document hardware, software and network
- Ensure you have a clean rootvg
 - Contains no application data
- Plan for success and disaster
- Take backups (mksysb and data) and test them before starting the process (unmirror rootvg first)
- Read the latest README and install guide

Mainline: solutions you need
from people you trust 6

Preparation

- Bring your firmware up to date – this is the scary part of the process
 - Pre-migration script will check this but you should too
- Check required VIO and HMC versions are in place if this is an LPAR
- Check you have the supported drivers for any third party hardware
 - MPIO drivers, etc
- Make sure you have all the media and it is readable
 - AIX v6 install DVDs (2 volume set)
 - If your server only has a CD drive make sure you get the multi-CD set instead
 - AIX Toolbox for Linux Applications
 - AIX Expansion Kit
 - Any update CDs
- Download the latest v6 TLs and service packs
 - <http://www-933.ibm.com/eserver/support/fixes/fixcentral/pfixpacks/61>
 - Burn them to CD or make them NFS mountable

Mainline: solutions you need
from people you trust 7

More Preparation

- You may need to increase the size of / as there are some extra filesets you will need to install
 - AIX v5.3 boot LV is around 12MB and v6.1 is 20MB minimum
 - Migration will try to increase as needed
- If using NIM, the NIM server must be updated to the latest AIX v6 first
 - It can still support your earlier releases
- Migration is supported for AIX v5.1 and higher – 4.3 required mkysyb install support
- Make sure none of your critical filesystems in rootvg are full or close to full

Mainline: solutions you need
from people you trust 8

More Preparation

- Console gotcha - Pre AIX v5.3 console
 - Console was /dev/tty
 - On 5.3 and higher it is now /dev/vty
 - Make sure your bosinst data (if NIM) says CONSOLE=Default
 - If it says /dev/tty then you won't be able to login after migration
- PowerPath and non-IBM multi-path drivers
 - Check with EMC or your vendor
 - You may need to uninstall their software (i.e. PowerPath) before the migration and reinstall it after the migration
 - There were some system hang issues after migration early on with PowerPath unless this was done

Mainline: solutions you need
from people you trust 9

Options for migration

- Complete Overwrite
- Migrations
 - Media
 - Alt disk Install
 - Nim or standalone
 - NIM
 - Clone or migration
 - VIO optical library
 - Can use this instead of media

Mainline: solutions you need
from people you trust 10

Using VIO Optical Library

- File backed Virtual SCSI
- Comes with VIOS 1.5 FP 10.1 so your VIO needs to be current
- Can even use it to fdo backups
- Add disk into a new volume group (violibvg)
- Using IVM (have not done this yet on my HMC – uses command line I think)
 - Connect to the VIO server partition (<https://viosname>)
 - Click on View/modify virtual storage
 - At the panel select the tab for optical devices
 - Create library and give it a size (I used 40gb) and choose violibvg
 - Creates /var/vio/VMLibrary as a filesystem and mounts it
 - Upload the iso formatted DVDs or CDs to /var/vio/VMLibrary
 - You can now allocate those iso files to LPARs to boot from or use, as if they were on a direct attached CD or DVD
- Useful URLs
 - <http://www.redbooks.ibm.com/redbooks/pdfs/sg247940.pdf>
 - Section 3.5.4 covers this
 - <http://www.ibm.com/developerworks/wikis/display/virtualization/IVM>
 - http://www.ibm.com/developerworks/wikis/download/attachments/51773454/IVM_FAQ.pdf?version=1
 - Movie from IBM on this = http://www.ibm.com/developerworks/wikis/download/attachments/53871873/IVM_optical.wmv

Mainline: solutions you need
from people you trust 11

Overwrite versus Migration

- | | |
|---|---|
| <ul style="list-style-type: none">• Overwrite<ul style="list-style-type: none">– Clean install– Easy switch to using JFS2 for rootvg– Starts with 64 bit kernel– Required if want to use “hard” security– Extensive configuration needed post install | <ul style="list-style-type: none">• Migration<ul style="list-style-type: none">– Preserves current configuration and data– Only recovery is restore from backup or reboot from old rootvg if alt-disk-install– Can’t convert to JFS2 – need to use backup/restore– Can’t be used if changing to new security model |
|---|---|

Mainline: solutions you need
from people you trust 12

The Migration Process

- Check oslevel so you know where you are coming from:
oslevel -s should show: 5300-07-03-0811
- Ensure primary authentication method shows as system
lsuser -a auth1 root should show: root auth1=SYSTEM
- Ensure all users are logged off
- Run errpt -a and ensure any errors are corrected
- Clean tape drive (if using tape backup)
- If rootvg is mirrored then unmirror it before proceeding
- Make a note of the disk address for rootvg
 - For me it was hdisk0 Available 01-08-00-3,0 (output from lsdev -Cdisk)
- Run backup of unmirrored rootvg and then check to see it was done correctly
- You may want to exportvg any data VGs prior to the install
- Run the premigration script on the install CD
 - It will be on the CD as /usr/lpp/bos/pre_migration
 - Do not remove the data this creates – it will be used later

Mainline: solutions you need
from people you trust 13

Pre-migration script

Checks for known conditions that may cause migration problems

Runs from install CD: <mntpoint>/usr/lpp/bos/pre_migration

Performs the following:

- List the device filesets being removed and other filesets being removed.
- List the saved base configuration files that will not be merged.
- List configuration files that will be merged.
- Verify fileset version consistency.
- Create a list of all filesets installed, to be used by the **post_migration command**.
- Check the size and location of the boot logical volume.
- Check the major number for rootvg is 10.
- Check for the missing DB directory for the **bos.net.ipsec.keymgt fileset**.
- Determine if Kerberos is being used.
- Check disk and memory sizes.
- Check the firmware level for IBM System p5™ 7025/7026 systems.
- If migrating from an earlier version of AIX, verify that the correct updates are applied.
- Verify system.
- Print a recommendation that a system backup be made before the migration.

Mainline: solutions you need
from people you trust 14

pre_migration script

```
./pre_migration
```

All saved information can be found in: /home/pre_migration.080824120559
Checking size of boot logical volume (hd5).
Listing software that will be removed from the system.
Listing configuration files that will not be merged.
Listing configuration files that will be merged.
Saving configuration files that will be merged.
Running lppchk commands. This may take awhile.

Please check /home/pre_migration.080824120559/software_link_existence_check for possible errors.

It is recommended that you create a bootable system backup of your system before migrating.

Mainline: solutions you need
from people you trust 15

pre_migration script output

```
/home/pre_migration.080824120559
```

```
# ls -l
```

```
total 184
```

-rw-r--r--	1	root	system	3253	Aug 24 12:06	configuration_files_being_merged
-rw-r--r--	1	root	system	15	Aug 24 12:06	configuration_files_missing_from_system
-rw-r--r--	1	root	system	67	Aug 24 12:06	configuration_files_not_merged
-rw-r--r--	1	root	system	1739	Aug 24 12:06	device_software_being_removed
drwxr-xr-x	5	root	system	256	Aug 24 12:06	saved_configuration_files
-rw-r--r--	1	root	system	0	Aug 24 12:06	software_checksum_verification
-rw-r--r--	1	root	system	0	Aug 24 12:06	software_consistency
-rw-r--r--	1	root	system	0	Aug 24 12:06	software_file_existence_check
-rw-r--r--	1	root	system	72244	Aug 24 12:05	software_installed_before_migration
-rw-r--r--	1	root	system	92	Aug 24 12:06	software_link_existence_check

Do not erase these as post_migration will use them later

Mainline: solutions you need
from people you trust 16

The Migration Process

- Put the DVD or CD in the drive and reboot using shutdown -r
- Standalone PCI server.
 - Insert the CD or DVD into the drive and power on the system.
 - At the first beep (when the keyboard and other prompts come up) hit F5 (or 5 if this is an ASCII terminal) several times, causing the system to boot in SMS mode.
 - NOTE – if F5 or 5 do not work then try F1 or 1
- LPAR (for P5 and P6)
 - At the HMC, check partition definitions
 - Assign the DVD/CD drive as an optional resource and make sure no one else has the DVD/CD assigned.
 - Right-click on the partition profile and select “Activate”.
 - Check the Open Terminal box (so you can use the vterm) and select “Advanced.”
 - Select “SMS” for the boot mode and select “OK” and “OK” again.
 - This ends the selection and opens the vterm, which should come up in SMS mode.

Mainline: solutions you need
from people you trust 17

Check Network

- SMS Mode Install
 - press 0 to select this as console
 - Select 2 for setup remote ipl
 - select 1 for correct ethernet adapter
 - Most likely the one that was the primary network before you rebooted
 - select 1 for IP parameters
 - Set client IP, server IP, gateway IP and subnet
 - Esc to go back
 - Select 2 adapter config –
 - 2 spanning tree - disable it (check this is supported by the switch)
 - 3 protocol - set to standard
 - 3 ping test
 - M – return to main menu

Note – your selection numbers and options may differ so please check them

Mainline: solutions you need
from people you trust 18

SMS Mode Install

- SMS Mode Install
 - From Main menu
 - select 5 boot options
 - 1 select boot device
 - Select the CD/DVD if booting from media or VIO Optical library
 - OR Select network if booting from NIM
 - 2 normal mode boot
 - 1 yes I want to exit
 - At prompt - 0 for console
 - 1 to use this as console
 - 1 for english during install
 - 2 check install settings
 - 1 Check hdisk it has chosen (rootvg is unmirrored so make sure to get the right one)
 - Select upgrade method – be sure to select MIGRATION
 - 0 to continue with those choices
 - The migration should now start and at the end the system will reboot as AIX v6.1
 - You will need to respond to various prompts possibly

Note – your selection numbers and options may differ so please check them

**Mainline: solutions you need
from people you trust** 19

Post Migration Steps

- Install Assistant will come up
 - Accept licenses, Check network, etc settings
 - Exit – select finish now and do not start configuration assistant when restarting AIX
- Run the `/usr/lpp/bos/post_migration` script
- You may need to increase the size of `/` as there are some extra filesets you will need to install
- In `/etc/inittab` remove the following line:
 - `/usr/sni/aix53/rc.ml`
- Install missing filesets from AIX v6 CDs
 - `bos.compat`, `bos.mls`, message sets, `sysmgmt.pconsole` (needed for WPARs)
- Install the latest TL and service packs
- The following were completely replaced and will need re-editing:
 - `/etc/inetd.conf`, `/etc/inittab`, `/etc/motd`, `/usr/dt/config/Xservers`
 - `/usr/sbin/skulker`
 - Check your environment and profile settings as well
 - Clean out or comment out unused services from `inetd.conf` and `rc.tcpip`

**Mainline: solutions you need
from people you trust** 20

Post Migration Steps

- Check files
 - Check all files you tailor normally after an install
 - Check integrity of installed packages
 - `lppchk -v` and `lppchk -c`
- Tunables
 - Tunables will migrate as they were
 - If you are coming from 5.1 or were using vmtune on 5.2 it will not work so remove that line from inittab or wherever you have it
 - Check `/etc/tunables/nextboot` and `lastboot`
 - Best method is to use `vmo`, `ioo`, and `no` to reset all tunables to defaults (write them down first)
 - i.e. `vmo -D`
 - Now go to the starter set of tunables and see if any changes need to be made

Mainline: solutions you need
from people you trust 21

Post Migration

- Increase the size of `/` as there are some extra filesets you will need to install
 - `/` went from 262144 to 52428
 - `/opt` may need to be increased
 - I always create a separate JFS2 filesystem for `/usr/local` and I set my logging up to use `/usr/local/logs`
 - `/usr/local` is normally (for me) around 1 to 2gb because I put all 3rd party software here
- Third Party Tools
 - You will need to reinstall TCP Wrappers if you use it
 - Use the IBM provided versions where possible (next slide)

Mainline: solutions you need
from people you trust 22

Third Party Apps

- Make sure you have /usr/local as a separate filesystem
- Add /usr/local/bin to your path
- Install nmon12e or higher into /usr/local/bin
- Create a group and id called freeware and another called sshd (if they do not already exist)
- Any SSH etc versions that you already have will have been migrated across so you may need to uninstall them to install the following

Mainline: solutions you need
from people you trust 23

Linux Toolkit

- I typically install the following – several are only needed if you are installing gcc though
- Ensure these are not against your corporate policies
- | | |
|-------------------|----------------------------|
| • bison-1.875 | flex 2.5.4a |
| • freetype-1.3.1 | gcc-4.0.0 and gcc-4.0.0C++ |
| • gettext-0.10.40 | less-382 |
| • libgcc-4.0.0 | libpcap 0.8.3 |
| • libstdC++ | libstdC++-devel |
| • sudo 1.6.7p5 | zlib 1.2.3 |
- Note: versions may differ on newer CDs
- If installing gcc you may want to create a softlink to it as cc if you do not have CC installed

Mainline: solutions you need
from people you trust 24

Expansion Pack

- Ensure these are not against your corporate policies
- To take advantage of the new security features and Java6
 - java6 bos.crypto*
 - clic.* des
 - modcrypt rsct.crypt
 - sysmgt.websm-security-us
- OpenSSH and openssl
 - I had problems getting these to compile so I use the versions on the expansion pack
 - They do need some of the security software above I think
 - Openssh openssl
- lsof
 - I usually install this as I find it useful
 - It can be used to list open files, open network connections and so on

Mainline: solutions you need
from people you trust 25

TCP Wrappers

- I always compile this one as I am finicky about the compile options chosen
- ftp://ftp.porcupine.org/pub/security/tcp_wrappers_7.6-ipv6.4.tar.gz
- 7.6-ipv6.4 is the current version as of today
- After unzipping it you vi the Makefile and set the following:
 - /AIX – uncomment REAL_DAEMON_DIR line
 - Uncomment STYLE=-DPROCESS_OPTIONS
 - Change FACILITY=LOG_MAIL to LOG_DAEMON
 - Check SEVERITY is LOG_INFO
 - Uncomment IPV6=-DHAVE-IPV6
- Then run “make aix” and copy the tcpd binary to /usr/local/bin
- Copy the tcpd.h file to /usr/local/include and the man pages to /usr/local/man
- You will need to code a /etc/hosts.allow and a /etc/hosts.deny file as these provide the rules for the wrappers

Mainline: solutions you need
from people you trust 26

TCP Wrappers

- /etc/hosts.allow
sshd : LOCAL,all
rexecd,ftpd,telnetd : LOCAL,10.128.,192.168.
- /etc/hosts.deny
all : all
- The last step is to update /etc/inetd.conf as follows:
ftp stream tcp6 nowait root /usr/local/bin/tcpd /usr/sbin/ftpd -l ftpd
telnet stream tcp6 nowait root /usr/local/bin/tcpd /usr/sbin/telnetd telnetd -a
- Make sure you stay logged on now till you do a successful test:
refresh -s inetd
- Now try to logon from a legitimate address
- Messages should show up in the log pointed to by daemon.info in /etc/syslog.conf

Aug 24 12:20:55 p5aix5a auth|security:info sshd[94400]: Received disconnect from 192.168.118.43: 11: Disconnect requested by Windows SSH Client.

Aug 24 12:30:41 p5aix5a auth|security:info sshd[249860]: Accepted password for root from 192.168.118.43 port 2058 ssh2

Mainline: solutions you need
from people you trust 27

Post Migration

- Asynchronous I/O
 - No longer setup using minservers and maxservers
 - aioo command no longer around
 - These are now tunables set using ioo
 - Legacy
 - aio_minreqs, aio_minservers, aio_maxservers, and aio_server_inactivity
 - Posix
 - posix_aio_minreqs, posix_aio_minservers, posix_aio_maxservers, posix_aio_server_inactivity
- I/O Pacing
 - This is now turned on – not sure if it changes the settings if you were already using it
 - New: minpout = 4096 and maxpout=8193
 - Check using lsattr -El sys0

Mainline: solutions you need
from people you trust 28

Final Steps

- Double check MPIO drivers
 - Will need to reinstall Powerpath or other drivers if you had to uninstall them
- Bring back data
 - importvg any volume groups you had exported
- Backup
 - Take a mksysb again of the new system
- Remirror rootvg and reboot
- Run your tests

Mainline: solutions you need
from people you trust 29

Starter set of tunables 1/2

Typically I set the following on AIX:

NETWORK

```
no -p -o rfc1323=1
no -p -o sb_max=1310720
no -p -o tcp_sendspace=262144
no -p -o tcp_recvspace=262144
no -p -o udp_sendspace=65536
no -p -o udp_recvspace=655360
```

Also check the actual NIC interfaces and make sure they are set to at least these values

Mainline: solutions you need
from people you trust 30

Network

Interface	Speed	MTU	tcp_sendspace	tcp_recvspace	rfc1323
lo0	N/A	16896	131072	131072	1
Ethernet	10/100 mb				
Ethernet	1000 (Gb)	1500	131072	165536	1
Ethernet	1000 (Gb)	9000	262144	131072	1
Ethernet	1000 (Gb)	1500	262144	262144	1
Ethernet	1000 (Gb)	9000	262144	262144	1
Virtual Ethernet	N/A	any	262144	262144	1
InfiniBand	N/A	2044	131072	131072	1

Above taken from Page 247 SC23-4905-04 November 2007 edition

Check up to date information at:

<http://publib.boulder.ibm.com/infocenter/pseries/v5r3/topic/com.ibm.aix.prfungd/doc/prftungd/prftungd.pdf>

Mainline: solutions you need
from people you trust 31

Starter set of tunables 2/2

Memory defaults are already correctly set and should not be changed

If you upgrade from a previous version of AIX using migration then you need to check the settings though (/etc/tunables/nextboot and lastboot)

Maxfree may need to be changed depending on j2_maxPageReadAhead and whether you are using JFS2 or jfs (see following slide)

The parameters below should be reviewed and changed

PBUFS

pv_min_pbuf is now a restricted tunable so tune for the individual volume group

Run lvmo -a to determine what needs changing

lvmo -v datavg -o pv_pbuf_count=1024

This changes datavg

JFS2

ioo -p -o j2_maxPageReadAhead=128

(default above may need to be changed for sequential – affects maxfree)

j2_dynamicBufferPreallocation=16

Default that may need tuning – I use 32

Replaces tuning j2_nBufferPerPagerDevice

JFS

ioo -p -o numfsbufs=1024

(now restricted)

ioo -p -o maxpgahead=16

(now restricted)

Mainline: solutions you need
from people you trust 32

minfree and maxfree

```
vmo -a | grep mempools  
mempools = 3
```

```
vmo -a | grep free  
maxfree = 1088  
minfree = 960
```

Calculation is:
 $\text{minfree} = (\max(960, (120 * \text{lcpus}) / \text{memory pools}))$
 $\text{maxfree} = \text{minfree} + (\text{Max}(\text{maxpgahead}, \text{j2_maxPageReadahead}) * \text{lcpus}) / \text{memory pools}$

So if I have the following:
Memory pools = 3 (from vmo -a)
J2_maxPageReadahead = 128
CPUS = 6 and SMT on so lcpu = 12
So minfree = $(\max(960, (120 * 12) / 3)) = 1440 / 3 = 480$ or 960 whichever is larger
And maxfree = $480 + (128 * 12) / 3 = 480 + 512 = 992$
But if we actually set minfree to 960 then maxfree should be $960 + 512 = 1472$

If you overallocate these values it is possible that you will see high values in the "fre" column of a vmstat and yet you will be paging.

Mainline: solutions you need
from people you trust 33

Gotchas

- Binary compatibility
 - Means it will run
 - Does not mean performance – recompile or reinstall applications if at all possible!
 - <http://www-03.ibm.com/systems/p/os/aix/compatibility/index.html>
- IBM C compiler outperforms GCC significantly on Power
- Multipath drivers
- TTY versus VTY
- Not testing backups

Mainline: solutions you need
from people you trust 34

Useful Links

- AIX Migration from Media Webpage
 - http://publib.boulder.ibm.com/infocenter/systems/scope/aix/topic/com.ibm.aix.howtos/doc/howto/HT_insgdrf_migrationinstall.htm?tocNode=int_191216
- Video on migrating
 - <http://publib.boulder.ibm.com/infocenter/pseries/v5r3/topic/com.ibm.aix.elearning/doc/aix.elearning/migration/migration.html>
 - It is based on upgrading to v5.3 but is useful
- AIX v6
 - AIX General
 - <http://www-03.ibm.com/systems/p/os/aix/index.html>
 - <http://publib.boulder.ibm.com/infocenter/systems/scope/aix/index.jsp>
 - Maintenance
 - <http://www-933.ibm.com/eserver/support/fixes/fixcentral/pfixpacks/61>
 - <http://publib.boulder.ibm.com/infocenter/pseries/v6r1/topic/com.ibm.aix.install/doc/insgdrf/insgdrf.pdf> Install Guide
 - <http://www.redbooks.ibm.com/Redbooks.nsf/RedpieceAbstracts/sg247559.html> Differences
- PowerVM
 - <http://www-03.ibm.com/systems/power/software/virtualization/index.html>
- Movies - <http://www.ibm.com/developerworks/wikis/display/WikiPtype/Movies>

Mainline: solutions you need
from people you trust 35



Questions???

Mainline: solutions you need
from people you trust 36