

UPGRADING AIX TO V6 OR V7

Presented at Common Virtual Conference 2/26/2013

Handout at:

<http://www.circle4.com/papers/aixupgrade-common.pdf>



Jaqui Lynch
lynchj@forsythe.com



AGENDA

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



WHY BOTHER?

All releases prior to AIX v6.1 are out of standard service

New hardware and function support

PowerVM requires AIX v5.3 or higher and POWER5 or POWER6 or POWER7

New functions in v6 or v7

- WPARs
- RBAC
- Trusted AIX Install option
- Encrypted JFS2 filesystem (EFS)
 - Can now use new POWER7+ crypto accelerator!

NOTE

- From AIX v5.3 on there is no support for MCA, PReP or ISA hardware

3



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 (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>

4



AIX 6 FEATURES....

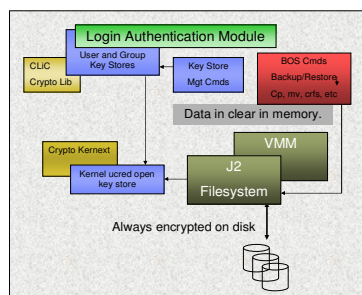
POWER6 Exploitation
 Software Reliability Availability Serviceability
 Enhancements to existing Virtualization Technologies
 Workload Partitions (Software based Virtualization)
 Application Mobility (Cross system Workload Mobility)
 64-bit Kernel only
 Integrated Multilevel Security
 Role Base Access Control (Partial Root base)
 Encrypted File system
 CAPP EAL4+ and LSPP Security Certification
 Solution Performance Tuning
 AIX Kernel Hot-Patching
 Dynamic Tracing for AIX
 Ease of Use

- Portal base SMT, LPAR Simplification



5

AIX ENCRYPTED FILE SYSTEM



- Enables improved security by reducing unauthorized access to data, even by privileged users
- Secure backups reduces the exposure of data compromised when backup media is taken outside of secure facilities
- Automatic management of protection keys can reduce the administrative effort of using encrypted data

- The capability to automatically encrypt data in a JFS2 filesystem
- Data can be protected from access by privileged users
- Backup in encrypted or clear formats
- Automated key management - key store open on login, integrated into AIX security authentication
- Each file encrypted with a unique key
- No keys stored in clear in kernel memory
- A variety of AES, and RSA cryptography keys supported

NEW accelerator in P7+



6

LPM – LIVE PARTITION MOBILITY

Uses

- Server Consolidation
- Workload Balancing
- Can use to allow power off of unused servers to assist in Green IT efforts
- Planned maintenance and upgrades

Inactive partition migration moves a powered-off partition

Partitions cannot be migrated from failed machines

PowerVM Enterprise

Requires minimum of 5300-07-01, 6100-00-01 or 7.1

Specific Firmware, HMC and VIOS requirements

POWER6/6+ and POWER7/7+

IT IS NOT A REPLACEMENT FOR HACMP OR OTHER HA

7



WPAR INTRO

Aim

- Reduce the number of images to maintain
- Reduce install and updating of AIX, backup, recovery, etc
- Encapsulate applications as units of control
- Rapid creation of new application environments
- Reduces memory needs as only one copy of AIX
- Mobility – for performance balancing as well as planned maintenance

Requires

- AIX v6 or higher on Power4, Power5 or Power6 hardware

Common operating system running a group of WPARs

- Each WPAR gets a regulated share of processor and memory resources
- Each WPAR has separate network and file systems
- Each WPAR is a separate administrative and security domain
- Shared resources are I/O devices, Processor, operating system and shared library and text
- Allows for automatic, policy based relocation of workloads

8



AIX 7.1 ADDS:

Fully exploit POWER7

- New hardware support

Versioned WPARs (5.2 and 5.3) for AIX v7

Vertical Scalability up to 256 cores / 1024 threads

- 1000 LPARs on a 795, 640 on 770 and 780 and 320 on the 750

Cluster Aware AIX infrastructure to simplify high availability and management of scale out workloads

Profile-based AIX configuration management from IBM Systems Director

Guaranteed application binary compatibility with previous AIX versions

Supports IBM hardware based on POWER7, POWER6, POWER5, POWER4 and PPC970 processors

AIX Strength to Strength

- <http://public.dhe.ibm.com/common/ssi/ecm/en/poo03022usen/POO03022USEN.PDF>

9



AIX 7.1 TL2 AND AIX 6.1 TL8 ADD:

Support for new D model 770 and 780

Exploitation of POWER7+ Cryptographic offload engine

Exploitation of POWER7+ Random number generator

LPAR to WPAR migration tool (AIX 7 only)

Remove need for rsh when using NIM

IPV6 support for NFSV3

IPV6 support for WPAR specific routing

Active System Optimizer now included in AIX 6 TL8

New schedo option for aggressive SMT4 use

Specific HMC, VIOS and firmware pre-reqs

AIX Strength to Strength

- <http://public.dhe.ibm.com/common/ssi/ecm/en/poo03022usen/POO03022USEN.PDF>

10



UPGRADING

11



PREREQUISITES

Make sure SWMA is current and check entitlement on the web

- <https://www-05.ibm.com/servers/eserver/ess/ProtectedServlet.wss>

Runs on POWER4, POWER5, POWER6 and PPC970 and POWER7

Version changes but it is a free upgrade if your AIX v5 SWMA is current

If you are licensed for (with SWMA) v5, v6 or v7 you can run any or all of them on that server

Check your applications are available and supported

- If you require certification then check that also
- Make sure the vendor has the versions of all applications so they can check for issues

Take a performance baseline

Bring your firmware up to date – this is the scary part of the process

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

12



AIX v6 OR v7 UPGRADE PREPARATION

You may need to increase the size of / as there are some extra filesets you will need to install – also possibly /opt

- AIX v5.3 boot LV is around 12MB and v6.1/v7.1 are 20MB minimum
- Migration will try to increase as needed

If using NIM, the NIM server must be updated to the latest AIX v6/v7 first

- It can still support your earlier releases

Migration is supported for AIX v5.1 and higher – 4.3 requires mkysb install support

Make sure none of your critical filesystems in rootvg are full or close to full

Create a sysplan if you use HMC or IVM – this is great documentation

Or use HMC Scanner:

- <http://www.ibm.com/developerworks/wikis/display/WikiPtype/HMC+Scanner>

Grab an lscfg and an lscfg -v

13



PREPARATION

Take a performance baseline

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 old and new media, and it is readable

- AIX v6 or v7 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
- May need the old media if you mess up the bootlist

Download the latest v6/v7 TLs and service packs

- <http://www-933.ibm.com/eserver/support/fixes/fixcentral/pfixpacks/61>
- <http://www-933.ibm.com/eserver/support/fixes/fixcentral/pfixpacks/71>
- Burn them to CD or make them NFS mountable
- Or create ISOs and use FBO

14



PREPARATION

Save contents of /etc/tunables to your PC

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

64 bit kernel

- AIX v6.1 & 7.1 only support a 64 bit kernel
- 32 and 64 bit applications will run fine
- 32 bit device drivers or kernel extensions NOT supported

Take a backup!! mksysb

15



OPTIONS FOR MIGRATION

Complete Overwrite and Migration

- Media
- Alt disk Install
 - Nim or standalone
- NIM
 - Clone or migration
- VIO optical library (file backed optical)
 - Can use this instead of media

16



USING VIO OPTICAL LIBRARY

File backed Virtual SCSI (file backed optical)

Comes with VIOS 1.5 FP 10.1 so your VIO needs to be current (latest release is 2.1)

Can even use it to backups

Add disk into a new volume group (violibvg)

Using IVM (HMC uses command line)

- 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

17



FBO LINE COMMANDS ON VIOS

Virtual Media Commands

```
chrep          chvopt
loadopt        lsrep
lsvopt         mkrep
mkvopt         rmrep
rmvopt         unloadopt
help <command>
loadopt -disk aix61tl06sp2 -vtd vtopt0
```

```
mkvdev -fbo -vadapter vhost0
        vtopt0 Available
```

```
lsrep
Size(mb) Free(mb) Parent Pool    Parent Size    Parent
Free
  69352   47003 datavg            69888          256

Name          File Size Optical Access
aix61tl06sp2.iso  3561      None      rw
```

18



OVERWRITE VERSUS MIGRATION

Overwrite

- 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

Migration

- 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

19



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 -Ccdisk)

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

20



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.

21



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.

TEST IT

22



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

23



THE MIGRATION PROCESS

Put the DVD or CD in the drive and reboot using shutdown -r

- This could also be NIM or File Backed Optical

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 F1 (or 1 if this is an ASCII terminal) several times, causing the system to boot in SMS mode.
- NOTE – if F1 or 1 do not work then try F5 or 5

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.
- *Instead of open terminal you can use vtmenu after SSHing to the HMC*

24



IF INSTALLING OVER NETWORK

SMS Mode Install

- press 0 to select this as console
- Select 2 for setup remote ip
- 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

25



SMS MODE DVD 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

26



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

- Check it first and see what it is going to do

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, `sysmgt.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`

27



POST MIGRATION STEPS

Check files

- Check all files you tailor normally after an install
- Check integrity of installed packages
 - `lppchk -v` and `lppchk -c`
 - Look for missing filesets – use `oslevel` and `instfix` to check there are no surprises

Tunables

- Tunables will migrate as they were
- If you are coming from 5.1 or were using `vm tune` 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
- Compare to the copy you saved earlier

28



POST MIGRATION

- Increase the size of / as there are some extra filesets you will need to install
 - / went from 262144 to 524288
 - /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)

29



THIRD PARTY APPS

Make sure you have /usr/local as a separate filesystem

Add /usr/local/bin to your path

nmon

- Previous nmon versions to nmon12 do not pick up all the information
- Nmon is now built into AIX with AIX v6.1 tl02, AIX v5.3 tl09 and VIO Server v2.1 and is automatically installed
- For AIX v5.3 I still install nmon myself but for 6 and 7 I use the one provided. You may need to change your scripts to point to the new nmon

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

30



LINUX TOOLKIT

I typically install the following – several are only needed if you are installing gcc though

These are rpm installs

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 xlc installed

31



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
- Openssh
- openssl

lsuf

- I usually install this as I find it useful
- It can be used to list open files, open network connections and so on

32



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

33



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

```
ftp  stream tcp6  nowait root  /usr/local/bin/tcpd /bin/false
```

- Add /bin/false to /etc/security/login.cfg in the shells section

```
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
```

34



POST MIGRATION

Asynchronous I/O

- No longer setup using minservers and maxservers in sys0
- No longer need to enable it – it is done for you
 - Check the AIX differences guide for v6 and the one for v7
- aioo command no longer around
- These tunables are now set using ioo – defaults – min=3, max=30
- 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

35



FINAL STEPS

Double check MPIO drivers

- Will need to reinstall Powerpath or other drivers/HAKs if you had to uninstall them

Bring back data

- importvg any volume groups you had exported

Recompile performance critical code with latest xLC

Backup

- Take a mksysb again of the new system

Remirror rootvg and reboot

Run your tests

36



MIGRATIONS AND UPDATES WITH NIM

Migration

- Changes version or release

Update

- Preserves version and release
- Basically changes a TL or SP

Both migrations and updates can use alternate copies of rootvg if there is an unused disk available

After changes boot from altinst_rootvg and test

Migrate or update NIM Master first

Then update LPP_SOURCE and SPOT

Use nimadm for migrations, to install a down level mksysb and then migrate it or to install a new golden image

Use nim_alt_clone with update_all to update a TL or SP

With multibos the standby copy of AIX can be on the same physical disk as the current rootvg

37



NIMADM

nimadm (Network Install Manager Alternate Disk Migration) command Utility that allows the system administrator to:

1. Create a copy of rootvg to a free disk (or disks) and simultaneously migrate it to a new version or release level of AIX.
2. Using a copy of rootvg, create a new NIM mksysb resource that has been migrated to a new version or release level of AIX.
3. Using a NIM mksysb resource, create a new NIM mksysb resource that has been migrated to a new version or release level of AIX.
4. Using a NIM mksysb resource, restore to a free disk (or disks) and simultaneously migrate to a new version or release level of AIX.

38



NIM ALT CLONE

Update within the same version/release (i.e update a TL or SP)

Build the lpp_source at the desired level with simages=yes

Update NIM master from this lpp_source and reboot master

On client

- Unmirror rootvg and cleanup (chpv -c and reducevg)
- bosboot -a and bootlist -m

Master

- smitty nim_alt_clone
- Specify target client and target disks
- Setup FIXES to install to "update_all"
- Point to lpp_source from above and accept licenses

The clone and update will take place on altinst_rootvg while still running on rootvg

Boot from altinst_rootvg and test

Either reboot from old rootvg or make this the production one

39



MULTIBOS

Creates, updates, and manages multiple versions of the Base Operating System (BOS) on a rootvg.

The multibos command allows the root level administrator to create multiple instances of AIX(R) on the same rootvg.

The multibos setup operation creates a standby Base Operating System (BOS) that boots from a distinct boot logical volume (BLV).

This creates two bootable sets of BOS on a given rootvg and the administrator can boot from either instance of BOS by specifying the respective BLV as an argument to the bootlist command or using system firmware boot operations.

Two bootable instances of BOS can be simultaneously maintained.

40



ALTERNATE DISK INSTALL

http://pic.dhe.ibm.com/infocenter/aix/v7r1/topic/com.ibm.aix.install/doc/insgdrf/HT_insgdrf_altdiskinstall_clone.htm

Make sure you have a spare hard disk or LUN

Install bos.alt_disk_install.rte and update it

Create a bundle to install to the new disk and any custom scripts

- Or take a mksysb of the system to a file

Clone rootvg using smitty alt_clone

lspv now shows a disk as rootvg and one as altinst_rootvg

Check your bootlist as the alt disk install process changes it

- bootlist -m normal -o

Correct the bootlist back to normal until you are ready

Now you can use smitty alt_disk_install to do the upgrades to the new hard drive

When happy with the upgrade you update the bootlist and reboot on the new image

41



ALT DISK FROM MKSYB

Commands to look at:

alt_disk_copy
alt_disk_install
alt_disk_mksysb
alt_rootvg_op

```
#
# lspv
hdisk0      00f6934c642af030      rootvg      active
hdisk1      none                  None
# lsdev -Ccdisk
hdisk0 Available Virtual SCSI Disk Drive
hdisk1 Available Virtual SCSI Disk Drive
#
# lspp -l | grep bos.alt
bos.alt_disk_install.boot_images
bos.alt_disk_install.rte 7.1.0.15 COMMITTED Alternate Disk Installation
bos.alt_disk_install.rte 7.1.0.15 COMMITTED Alternate Disk Installation
```

```
Alternate Disk Installation
Move cursor to desired item and press Enter.
Install mksysb on an Alternate Disk
Clone the rootvg to an Alternate Disk
NIM Alternate Disk Migration

F1=Help      F2=Refresh   F3=Cancel    F8=Image
F9=Shell     F10=Exit     Enter=Do
```

42



ALT DISK FROM MKSYSB

```

Install mksysb on an Alternate Disk

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

* Target Disk(s) to install      [Entry Fields]
* Device or image name          [hdisk1] +
Phase to execute                 [/usr/local/backups/b7] +
image.data file                 all +
Customization script            [] /
Set bootlist to boot from this disk
on next reboot?                 no +
Reboot when complete?          yes +
Verbose output?                 no +
Debug output?                   [] /
resolv.conf file

F1=Help      F2=Refresh      F3=Cancel      F4=List
Esc+S=Reset  F6=Command      F7=Edit       F8=Image
F9=Shell     F10=Exit       Enter=Do

```

43



ALT DISK FROM MKSYSB

```

COMMAND STATUS

Command: running      stdout: yes      stderr: no

Before command completion, additional instructions may appear below.

[MORE...6]
Restoring /image.data from mksysb image.
Checking disk sizes.
Creating cloned rootvg volume group and associated logical volumes.
Creating logical volume alt_hd5.
Creating logical volume alt_hd6.
Creating logical volume alt_hd8.
Creating logical volume alt_hd4.
Creating logical volume alt_hd2.
Creating logical volume alt_hd9var.
Creating logical volume alt_hd3.
Creating logical volume alt_hd1.
Creating logical volume alt_hd10opt.
Creating logical volume alt_hdl1admin.
Creating logical volume alt_lg_dump1v.
[BOTTOM]

```

Plus many more as it restores from the mksysb image

44



ALT DISK FROM MKSYB

Now we see:

```
# lspv
hdisk0      00f6934c642af030      rootvg      active
hdisk1      00f6934c75816830      altinst_rootvg active
```

```
# bootlist -m normal -o
hdisk0 blv=hd5 pathid=0
```

- Phase 1** Creates the **altinst_rootvg** volume group, the **alt_** "logical volumes", the **/alt_inst** file systems, and restores the mksysb or rootvg data.
- Phase 2** Runs any specified customization script, installs updates, new filesets, fixes or bundles (cloning only), copies a **resolv.conf** file if specified, and copies files over to remain a NIM client if specified.
- Phase 3** Unmounts the **/alt_inst** file systems, renames the file systems and logical volumes, removes the **alt_** logical volumes, names ODM and varies off the altinst_rootvg. It sets the bootlist and reboots if specified.

45



ALT DISK FROM MKSYB

```
# lspv -l hdisk0
hdisk0:
LV NAME      LPs  PPs  DISTRIBUTION  MOUNT POINT
hd8           1    1    00..00..01..00..00  N/A
hd6          64   64    00..64..00..00..00  N/A
hd2          80   80    00..00..80..00..00  /usr
hd4          10   10    00..00..10..00..00  /
hd3          48   48    00..00..48..00..00  /tmp
hd9var       12   12    00..00..12..00..00  /var
hd10opt      32   32    00..32..00..00..00  /opt
hd1           4    4    00..00..04..00..00  /home
hd5           1    1    01..00..00..00..00  N/A
lg_dumplv2   16   16    00..16..00..00..00  N/A
fslv00       80   80    80..00..00..00..00  /usr/local
fslv01       16   16    00..16..00..00..00  /usr/local/logs
lg_dumplv    16   16    00..16..00..00..00  N/A
livedump      4    4    00..04..00..00..00  /var/adm/ras/livedump
hd11admin     4    4    00..00..04..00..00  /admin

# lspv -l hdisk1
hdisk1:
LV NAME      LPs  PPs  DISTRIBUTION  MOUNT POINT
alt_hd10opt  32   32    00..32..00..00..00  /alt_inst/opt
alt_hd1       4    4    00..00..04..00..00  /alt_inst/home
alt_hd3       48   48    00..00..48..00..00  /alt_inst/tmp
alt_hd9var    12   12    00..00..12..00..00  /alt_inst/var
alt_hd2       80   80    00..00..80..00..00  /alt_inst/usr
alt_hd4       10   10    00..00..10..00..00  /alt_inst
alt_hd8        1    1    00..00..01..00..00  N/A
alt_hd6       64   64    00..64..00..00..00  N/A
alt_hd5        1    1    01..00..00..00..00  N/A
alt_lg_dumplv 16   16    00..16..00..00..00  N/A
alt_hd11admin  4    4    00..00..04..00..00  /alt_inst/admin
alt_lg_dumplv2 16   16    00..16..00..00..00  N/A
alt_fslv01    16   16    00..16..00..00..00  /alt_inst/usr/local/logs
alt_fslv00    80   80    80..00..00..00..00  /alt_inst/usr/local
alt_livedump  4    4    00..04..00..00..00  /alt_inst/var/adm/ras/livedump
```

You can display the above during the mksysb clone
At the end altinst_rootvg is varied offline and these (the alt ones) are all unmounted

46



ALT DISK FROM MKSYSB

You can wake up the altinst_rootvg to mount the filesystems and put it back to sleep:

```
# alt_disk_install -W hdisk1
```

```
-----+
ATTENTION: calling new module /usr/sbin/alt_rootvg_op. Please see
the
alt_rootvg_op man page and documentation for more details.
Executing command: /usr/sbin/alt_rootvg_op -W -d hdisk1
-----+
```

```
Waking up altinst_rootvg volume group ...
```

```
# lsvg
hdisk0      00f6934c642af030      rootvg      active
hdisk1      00f6934c75816830      altinst_rootvg active
```

Back to sleep

```
# alt_disk_install -S hdisk1
```

```
-----+
ATTENTION: calling new module /usr/sbin/alt_rootvg_op. Please see
the
alt_rootvg_op man page and documentation for more details.
Executing command: /usr/sbin/alt_rootvg_op -S hdisk1
-----+
```

```
Putting volume group altinst_rootvg to sleep ...
forced unmount of /alt_inst/var/adm/ras/livedump
forced unmount of /alt_inst/var/adm/ras/livedump
forced unmount of /alt_inst/var
forced unmount of /alt_inst/var
forced unmount of /alt_inst/usr/local/logs
forced unmount of /alt_inst/usr/local/logs
forced unmount of /alt_inst/usr/local
forced unmount of /alt_inst/usr/local
forced unmount of /alt_inst/usr
forced unmount of /alt_inst/usr
forced unmount of /alt_inst/tmp
forced unmount of /alt_inst/tmp
forced unmount of /alt_inst/opt
forced unmount of /alt_inst/opt
forced unmount of /alt_inst/home
forced unmount of /alt_inst/home
forced unmount of /alt_inst/admin
forced unmount of /alt_inst/admin
forced unmount of /alt_inst
forced unmount of /alt_inst
Fixing LV control blocks...
Fixing file system superblocks...
```

```
# lsvg
hdisk0      00f6934c642af030      rootvg      active
hdisk1      00f6934c75816830      altinst_rootvg
```

47



ALT DISK FROM MKSYSB

```
# df -g
```

Filesystem	GB	blocks	Free	%Used	Used	%Used	Mounted on
/dev/hd4	0.62		0.41	35%	10330	10%	/
/dev/hd2	5.00		2.37	53%	62100	11%	/usr
/dev/hd9var	0.75		0.47	38%	5785	5%	/var
/dev/hd3	3.00		3.00	1%	72	1%	/tmp
/dev/hd1	0.25		0.25	1%	11	1%	/home
/dev/hd11admin	0.25		0.25	1%	5	1%	/admin
/proc	-	-	-	-	-	-	/proc
/dev/hd10opt	2.00		1.41	30%	13901	5%	/opt
/dev/livedump	0.25		0.25	1%	4	1%	/var/adm/ras/livedump
/dev/fslv00	5.00		4.93	2%	231	1%	/usr/local
/dev/fslv01	1.00		0.98	2%	39	1%	/usr/local/logs
/dev/alt_hd4	0.62		0.53	1%	13	1%	/alt_inst
/dev/alt_hd11admin	0.25		0.25	1%	5	1%	/alt_inst/admin
/dev/alt_hd1	0.25		0.25	1%	11	1%	/alt_inst/home
/dev/alt_hd10opt	2.00		1.41	30%	13900	5%	/alt_inst/opt
/dev/alt_hd3	3.00		3.00	1%	61	1%	/alt_inst/tmp
/dev/alt_hd2	5.00		2.37	53%	62100	11%	/alt_inst/usr
/dev/alt_fslv00	5.00		4.93	2%	230	1%	/alt_inst/usr/local
/dev/alt_fslv01	1.00		0.98	2%	39	1%	/alt_inst/usr/local/logs
/dev/alt_hd9var	0.75		0.47	38%	5761	5%	/alt_inst/var
/dev/alt_livedump	0.25		0.25	1%	4	1%	/alt_inst/var/adm/ras/livedump

48



QUICK ALT_DISK_COPY EXAMPLE

```
alt_disk_copy -d hdisk1 -F 5300-01_AIX_ML -l /updates
```

The above copies the current rootvg (happens to be at 5300-00) to hdisk1

It then applies the updates from /updates to bring the cloned rootvg on hdisk1 to 5300-01

It also sets the bootlist to boot from hdisk1

Allows you to copy the running system and apply maintenance in one step

After rebooting, the old rootvg will be named old_rootvg

Ensure you use alt_rootvg_op to remove it later

Use bootlist to go back if needs be

49



STARTER SET OF TUNABLES 1/2

Typically I set the following on AIX:

NETWORK

```
no -p -o rfc1323=1
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

ifconfig -a will show the values

Stop inetd and use chdev to change

50



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>

51



STARTER SET OF TUNABLES 2/2

On a fresh install, 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 – no need to do it globally

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)

52



MINFREE AND MAXFREE

```
vmo -a | grep mempool
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
 $\text{So minfree} = (\max(960, (120 * 12) / 3)) = 1440 / 3 = 480$ or 960 whichever is larger
 In this case we use 960 since 480 is too low for the min
 And $\text{maxfree} = 960 + (128 * 12) / 3 = 480 + 512 = 1472$
 I would bump this to 1536 (I use 1024, 1536, 1920, 2048)

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. If you underallocate these you will be paging.

53



GOTCHAS

Binary compatibility

- Means it will run
- Does not mean performance – recompile or reinstall applications if at all possible!
- `emstat -a` is your friend! Look for emulations and alignments
- <http://www-03.ibm.com/systems/p/os/aix/compatibility/index.html>

IBM C compiler outperforms GCC significantly on Power

- Recompile performance critical code with xLC and check out the new options such as `qarch`

Multipath drivers

TTY versus VTY

Not testing backups or cleaning tape drives

54



FLRT – FIX LEVEL RECOMMENDATION TOOL!

<http://www.ibm.com/support/customerarcare/flrt/>

Now in v6.5:

Support for multiple LPARs

- Ability to save an inventory file and then reupload it later when you return, editing values as needed. Can then rerun the full report again.
- Ability to supply input for multiple LPARs in a single CEC or system. Can provide inventory configuration data for each partition on your Power, PureSystems or Power blade system. Provides an easy way to generate a complete report on recommendations and software compatibility for your entire system.

Also now supports firmware recommendations for PureSystems Power compute nodes.

Other usability improvements:

- Removal of the "Processor speed" requirement.
- Opening of a separate tab when a link to Fix Central is clicked.
- Global type ahead product search.
- Session data saved via cookies.

Please send them feedback using the feedback link

55



USEFUL WEB LINKS

Sign up for Storage and System Notifications

- <https://www14.software.ibm.com/webapp/set2/subscriptions/onvdc>

Article on this topic

- <http://www.ibmsystemsmag.com/aix/trends/aix/Installing-and-Upgrading-to-AIX-6-1/>

POWER Firmware Code Matrix

- <http://www-304.ibm.com/webapp/set2/sas/f/power5cm/power7.html>

Fix Level Recommendation Tool (FLRT)

- <http://www-304.ibm.com/support/customerarcare/flrt/home>

AIX Wiki

- <https://www.ibm.com/developerworks/wikis/display/WikiPtype/AIX>

HMC Scanner

- <http://www.ibm.com/developerworks/wikis/display/WikiPtype/HMC+Scanner>

Service and support best practices

- <http://www14.software.ibm.com/webapp/set2/sas/f/best/home.html>

Fix Central - HMC, SDMC, Firmware, AIX Updates

- <http://www-933.ibm.com/support/fixcentral/>

IBM Prerequisite Tool

- https://www-912.ibm.com/e_dir/eserverprereq.nsf

IBM System Planning Tool

- <http://www-947.ibm.com/systems/support/tools/systemplanningtool/>

Migrating AIX

- <http://pic.dhe.ibm.com/infocenter/aix/v7r1/topic/com.ibm.aix.install/doc/insqdrf/insqdrf-kickoff.htm>

Forsythe Talks

- <http://www.circle4.com/forsythetalks.html>

56



THANK YOU FOR YOUR TIME



If you have questions please email me at:
lynchj@forsythe.com

This presentation at:
<http://www.circle4.com/papers/aixupgrade-common.pdf>