Useful AIX commands

svmon
svmon -P <pid>

Further:
use can user svmon command to monitor memory usage as follows;

(A) #svmon -P -v -t 10 | more   (will give top ten processes)
(B) #svmon -U -v -t 10 | more   (will give top ten user)

---

smit install requires "inutoc ." first. It'll autogenerate a .toc for you
I believe, but if you later add more .bff's to the same directory, then
the inutoc . becomes important. It is of course, a table of contents.

---

dump -ov /dir/xcoff-file

---

topas, -P is useful # similar to top

When creating really big filesystems, this is very helpful:
chlv -x 6552 lv08
Word on the net is that this is required for filesystems over 512M.

esmf04m-root> crfs -v jfs -g'ptmpvg' -a size='884998144' -m'/ptmp2'
-A"`locale yesstr | awk -F: '{print $1}'`" -p'rw' -t"`locale yesstr |
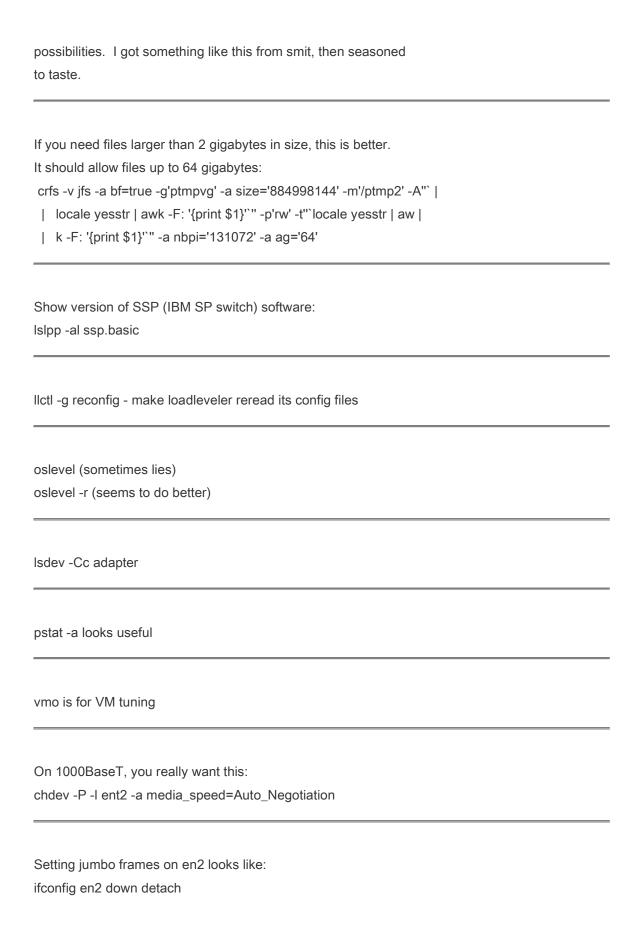awk -F: '{print $1}'`" -a frag='4096' -a nbpi='131072' -a ag='64'
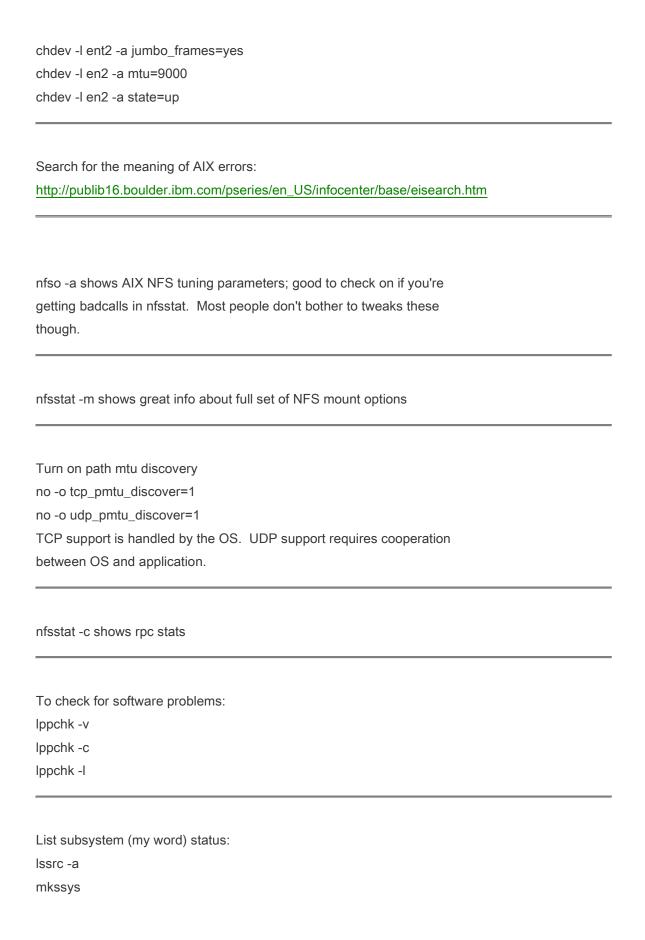Based on the parameters chosen, the new /ptmp2 JFS file system
is limited to a maximum size of 2147483648 (512 byte blocks)
New File System size is 884998144
esmf04m-root>

If you give a bad combination of parameters, the command will list

possibilities. I got something like this from smit, then seasoned
to taste.

If you need files larger than 2 gigabytes in size, this is better.
It should allow files up to 64 gigabytes:

```
 crfs -v jfs -a bf=true -g'ptmpvg' -a size='884998144' -m'/ptmp2' -A'` |
  |   locale yesstr | awk -F: '{print $1}'`'" -p'rw' -t'`locale yesstr | aw |
  |   k -F: '{print $1}'`'" -a nbpi='131072' -a ag='64'
```

Show version of SSP (IBM SP switch) software:
lslpp -al ssp.basic

---

llctl -g reconfig - make loadleveler reread its config files

---

oslevel (sometimes lies)
oslevel -r (seems to do better)

---

lsdev -Cc adapter

---

pstat -a looks useful

---

vmo is for VM tuning

On 1000BaseT, you really want this:
chdev -P -l ent2 -a media_speed=Auto_Negotiation

Setting jumbo frames on en2 looks like:
ifconfig en2 down detach

chdev -l ent2 -a jumbo_frames=yes

chdev -l en2 -a mtu=9000

chdev -l en2 -a state=up

Search for the meaning of AIX errors:

http://publib16.boulder.ibm.com/pseries/en_US/infocenter/base/eisearch.htm

---

nfso -a shows AIX NFS tuning parameters; good to check on if you're
getting badcalls in nfsstat. Most people don't bother to tweaks these
though.

---

nfsstat -m shows great info about full set of NFS mount options

---

Turn on path mtu discovery

no -o tcp_pmtu_discover=1

no -o udp_pmtu_discover=1

TCP support is handled by the OS. UDP support requires cooperation
between OS and application.

---

nfsstat -c shows rpc stats

To check for software problems:

lppchk -v

lppchk -c

lppchk -l

List subsystem (my word) status:

lssrc -a

mkssys

rmssys

chssys

auditpr

refresh

startsrc

stopsrc

traceson

tracesoff

This starts sendmail:

startsrc -s sendmail -a "-bd -q30m"

This makes inetd reread its config file. Not sure if it kills and

restarts or just HUP's or what:

refresh -s inetd

---

lsps is used to list the characteristics of paging space.

Turning off ip forwarding:

/usr/sbin/no -o ipforwarding=0

Detailed info about a specific error:

errpt -a -jE85C5C4C

BTW, Rajiv Bendale tells me that errors are stored in NVRAM on AIX,

so you don't have to put time into replicating an error as often.

Some or all of these will list more than one number. Trust the first,

not the second.

lslpp -l ppe.poe

...should list the version of poe installed on the system

Check on compiler versions:

lslpp -l vac.C

lslpp -l vacpp.cmp.core

Check on loadleveler version:

lslpp -l LoadL.full

If you want to check the bootlist do bootlist -o -m normal if you want to

update bootlist do bootlist -m normal hdisk* hdisk* cd* rmt*

---

prtconf

Run the ssadiag against the drive and the adapter and it will tell you if it

fails or not. Then if its a hot plugable it can be replaced online.

You can get patches from:

http://www-912.ibm.com/eserver/support/fixes

You'll need to click through a bit of red tape before getting to where

you actually can list corequisites and start a download.

BTW, "Add to my download list" does not work in konqueror, but it does

work in mozilla.

Backup to tape:

env - /usr/bin/mksysb '-m' '-i' '-X' /dev/rmt0

The "env -" is because some sort of environment variable can confuse

mksysb, making it error out instead of doing your backup

There's also "smitty mksysb"

You can create an image using the savevg command i.e.

savevg -v -n -9 / _rootvg.img rootvg

This can be used to build a NIM installable image to recover your systems

alternatively, the command line call for a mksysb to tape (to include map
and exclude files) is /usr/bin/mksysb '-m' '-e' '-i' /dev/rmt0

Finding which lpp contains a file:
lslpp -w /usr/sbin/umount

There exists a "diag CD" for AIX.

/usr/samples/kernel/vmtune

---

lsattr -El sys0 | grep realmem
lsattr -El mem0

See if you AIX system's hardware is CHRP (some sort of PowerPC reference
platform spec, I believe) :
bootinfo -p
chrp

Some really funky hardware-looking problems can be fixed by draining
the NVRAM battery for 5 minutes, and then reinstalling the microcode.
We used to do this on some IBM RT's in Cincinnati, and a recent poster
to AIX-L indicates that it's still useful in some situations.

From AIX-L:

AIX 4.3.2 -> AIX 4.3.3 is the most easiest Upgrade of ALL. Place the
AIX 4.3.3 Vol 1 of CD on the CDROM drive and run smitty update_all ,
this will upgrade the OS

On the subject of running out of paging space, from AIX-L:
Look into npswarn, npskill stuff in Performance Management Guide

---

Changing the boot device order:
Boot the server, and hit 1 or F1 (depending if you have an ascii console
or a graphics console) when the logos come up to get to sms mode. In
the menus select multiboot, select boot devices, select boot order.

You should start tracing for inetd subsystem with

traceson -s inetd

and then issue:

trpt -j

you will see the protocols control blocks (PID) you're tracing, and then with:

trpt -p <PID>

you should see output for telnet communications. But this is not working.

Why don't you try using iptrace and ipreport to see the behavior of your
telnet sessions ??

Purportedly orks with JFS 1 and 2:

To split off a mirrored copy of the /home/xyz file system to a new mount
point named /jfsstaticcopy, type the following:

chfs -a splitcopy=/jfsstaticcopy /home/xyz

You can control which mirrored copy is used as the backup by using the
copy attribute. The second mirrored copy is the default if a copy is
not specified by the user. For example:

 chfs -a splitcopy=/jfsstaticcopy -a copy=1 /home/xyz

At this point, a read-only copy of the file system is available in
/jfsstaticcopy. Any changes made to the original file system after the
copy is split off are not reflected in the backup copy.

To reintegrate the JFS split image as a mirrored copy at the /testcopy
mount point, use the following command:

 rmfs /testcopy

The rmfs command removes the file system copy from its split-off state
and allows it to be reintegrated as a mirrored copy.

Working around a loader domain problem:

```
esmf04m-strombrg> /usr/local/bin/gribmap
exec(): 0509-036 Cannot load program /usr/local/bin/gribmap because of
the following errors:
    0509-030 Insufficient permission to create loader domain
    /usr/lib/libiconv.a
    0509-026 System error: The file access permissions do not allow
    the specified action.

esmf04m-strombrg> LIBPATH=$TMPDIR/gribmap-ld /usr/local/bin/gribmap
gribmap v1.4 for GrADS Version 1.8SL11
```

Apparently you can also link your application with -L$TMPDIR/loaderdomain

or so, but you'd need a unique one for each set of shared libraries.

This one apparently must be the first -L in the link line.

Please see also:

http://dcs.nac.uci.edu/~strombrg/AIX-shared-libs.html

/usr/bin/uname -M

Anyway, set the timezone variable TZ in /etc/environment like this:

    TZ=MST7

...takes effect after everyone logs out and back in. This is just an

example, not something for California.

"svmon" will give u this output which give u the information regarding

ur memory.

	size	inuse	free	pin	virtual
memory	1310711	1298235	12476	103782	711466
pg space	2097152	585219			

	work	pers	clnt	lpage
pin	103782	0	0	0
in use	438570	10130	849535	0

---

acledit

Scott (of IBM, onsite hardware tech) stuff:

lsdev -Cc adapter

"defined" means at one time the piece of hardware was on system - as opposed to "available". A card which is being newly added should not temporarily pass through "defined" state. Hardware should be in the "available" state.

/////

lsslot -c pci

p1-i1 is the first slot on the back left

/////

diag

diadiagnostic routines

problem determination

sfp: phones home (to IBM) over modem

previously reported problem

/////

task selection

hot plug task

pci or scsi

identify function will blink light, so you can make sure the hardware and software are on the same page.

u1.1 drawer address, bottom left

/////

EIA numbers on right and left of rack, goes to lowest of the numbers adjacent to the equipment in question. EG, something in the rack might be 3 EIA numbers high - the software should identify the hardware by the lowest number of the 3.

/////

hotplug in os removes voltage from slot, slot light should blink yellow, same as for identify.

/////

we have older "hotswap cassettes" - which means lots of screws. Newer ones snap together. It also can take a bit of wrestling to get the card in and out of the old cassettes (shades of Sun's IPX's :)

/////

yellow llight continues blinking after card inserted, until software is told to let the slot have voltage again.

/////

advanced diagnostics, search for thing to test visually

/////

cfgmgr

takes awhile to run, checks all devices on machine

no output, but then lsdev -Cc adapter again should change afterward

/////

ps -ef | grep Worm

splstdata -a

should not say not_configured

use rc.switch to make it configured

ps -ef pipe | Worm again, should show up now

Eunfence 49 - 49 is 04m

/////

spmon -d

"d" for diagnostic

like front panel leds

"host responds" and "switch responds" should say yes for all css adapters

/////

errpt (no args)

/////

Scott says that sometimes an SP2 system will refuse to acknowledge the

new adapter, in which case you have to lie to the ODM to make the system

see the card. He suggested that maybe we did not need to do that this

time, because we have the latest pssp (ssp.*) software on the system.

/////

We also had to Eunfence the node whose card was replaced.

Rajiv tells me that it does not matter which host is Eprimary, as long as

one of the nodes is, and there aren't things fenced off that shouldn't be.

---

mount -v cdrfs -o ro /dev/cd0 /mnt

Mount iso9660 filesystem

More on cfgmgr, from aix-l:

you can execute cfgmgr on line without trouble

normally cfgmgr have 3 steps named phases :
phase 1 during boot
phase 2 during normal boot (after phase1)
phase 3 durinf service boot (after phase1)

if you run cfgmgr without flags (-p or -f) cfgmgr executes phase 2 only by default

in fact cfgmgr and cfgmgr -p2 are the sames commands
flag -v for verbose

AIX 5.2 has builtin CIFS client?
mount -v cifs -n winserver/myuser/mypassword /home /mnt

Can also "smitty cifs_fs"

This is supposed to be included in lpp bos.cifs_fs

Apparently this was added in AIX 5.2

---

please check if your cd device is being used by some process by running:

fuser -c /dev/cd0

you can also chack if cdromd is up and running by:

lssrc -a | grep cd

if cdromd is running, then try with the following commands:

cdumount

cdeject

here is some commands to manipulate the ODM directly (I don't suggest you
do so, at least you know exactly what you are doing).
odmget, odmshow, odmchange, odmadd, odmdelete, odmdrop

---

lsps -a

---

nmon - free, unsupported download from IBM
What's this about chmod'ing kmem to be world readable though?!

---

esmf04m-dcsew> instfix -i | grep ML
   All filesets for 5.1.0.0_AIX_ML were found.
   All filesets for 5100-01_AIX_ML were found.
   All filesets for 5100-02_AIX_ML were found.
   All filesets for 5100-03_AIX_ML were found.
   All filesets for 5100-04_AIX_ML were found.
esmf04m-dcsew>

The specifix fixes can be checked using instfix command:
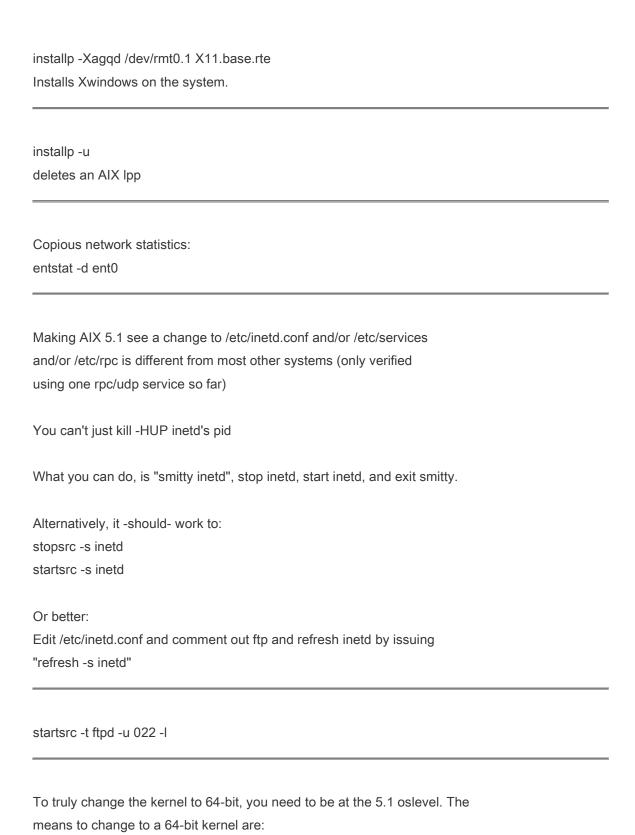#instfix -ivk <APAR no.>
e.g #instfix -ivk IY56076

---

instfix -ciqk 4330-08_AIX_ML | grep ":-:"
Lists what filesets need to be installed for instfix to show "All filesets
for 4330-08 were found".

---

instfix -k "IX#####" -d /dev/rmt0.1
Installs the APAR and its prerequisites on the system.

installp -Xagqd /dev/rmt0.1 X11.base.rte

Installs Xwindows on the system.

---

installp -u

deletes an AIX lpp

Copious network statistics:

entstat -d ent0

Making AIX 5.1 see a change to /etc/inetd.conf and/or /etc/services
and/or /etc/rpc is different from most other systems (only verified
using one rpc/udp service so far)

You can't just kill -HUP inetd's pid

What you can do, is "smitty inetd", stop inetd, start inetd, and exit smitty.

Alternatively, it -should- work to:

stopsrc -s inetd

startsrc -s inetd

Or better:

Edit /etc/inetd.conf and comment out ftp and refresh inetd by issuing
"refresh -s inetd"

---

startsrc -t ftpd -u 022 -l

To truly change the kernel to 64-bit, you need to be at the 5.1 oslevel. The
means to change to a 64-bit kernel are:

From 32-bit to 64-bit:

ln -sf /usr/lib/boot/unix_64 /unix

ln -sf /usr/lib/boot/unix_64 /usr/lib/boot/unix

lslv -m hd5

bosboot -ad /dev/ipldevice

shutdown -Fr

bootinfo -K (should now be 64)

To change the kernel back to 32-bit:

From 64-bit to 32-bit:

ln -sf /usr/lib/boot/unix_mp /unix

ln -sf /usr/lib/boot/unix_mp /usr/lib/boot/unix

lslv -m hd5

bosboot -ad /dev/ipldevice

shutdown -Fr

bootinfo -K (should now be 32)If you are running AIX 5.1

Switching From 32 to 64 Bit Mode

To switch from 32-bit mode to 64-bit mode run the following commands,
in the given order:

1.ln -sf /usr/lib/boot/unix_64 /unix

2.ln -sf /usr/lib/boot/unix_64 /usr/lib/boot/unix

3.bosboot -ad /dev/ipldevice

4.shutdown -Fr

5.bootinfo -K (should now show 64)

Switching From 64 To 32-Bit Mode

To switch from 64-bit mode to 32-bit mode run the following commands,
in the given order:

1.ln -sf /usr/lib/boot/unix_mp /unix

2.ln -sf /usr/lib/boot/unix_mp /usr/lib/boot/unix

3.bosboot -ad /dev/ipldevice

4.shutdown -Fr

5.bootinfo -K (should now show 32)

_-_-_-_-_-_-_-_-_-_-_-_-_-_-_-_-_-_-_-
Moulay Rachid BOUSSETA

To see if you're running with a 32 bit or 64 bit kernel, run:

bootinfo -K
...or...
prtconf -k

EG:
esmf04m-root> PATH=/usr/bin:/usr/sbin prtconf -k
Kernel Type: 64-bit
esmf04m-root> bootinfo -K
64

bootinfo -s hdiskxxx

lspv hdiskXX   as well is defined on a Volume Group

lsattr -El hdiskXX

lscfg -vp -l hdiskXX

These should give you the raw disk capacity

Go to: http://www-912.ibm.com/eserver/support/fixes/fcgui.jsp
Choose: 1)pSeries family
        2)AIX OS,Java, compilers
        3)Specifi fix
        4)Your OS

Press continue....

type in your requisites in the text box. thats it!

LPP history:

lslpp -h

---

sar 1 10

---

bootinfo -b   reports last device the system booted from
bootinfo -k   reports keyswitch position
                       1=secure, 2=service, 3=normal

bootinfo -r   reports amount of memory (/ by 1024)
bootinfo -s (disk device)   reports size of disk drive
bootinfo -T   reports type of machine
                       ie rspc,rs6ksmp,rspc or chrp
bootinfo -y   reports your hardware arquitecture (32
bits or 64 bits)

bootinfo -K   reports if the kernel in memory is 32
bits or 64 bits

You can submit/check a pSeries PMR via the web at:

https://techsupport.services.ibm.com/ssr/ssr.slprob

Force a user to change their password on their next login:
pwdadm -f ADMCHG username
Note that this works with some sshd's and not others

Identifying hard disk issues:

svmon -G
vmstat 1 20
iostat -d hdisk0 1 20
ps avg | sort +3r -n | head -25

Maximum number of processes a user can have:
lsattr -E -l sys0 -a maxuproc

---

smitty chgsys
Also allows one to change the max number of processes per user, among
other things

AIX and SNMP:

by Host Resource you mean the AIX SNMP component that monitors system
resources ??

if so, then there's a conf file for the daemon aixmibd named
/etc/aixmibd.conf where you can configure the thresholds for many
monitors. Once you have configure this then you should activate the daemon
by issuing:

startsrc -s aixmibd

Please remember to uncomment the line that starts aixmibd in /etc/rc.tcpip
file.

On AIX patches:

1) An APAR (Authorized Program Analysis Report) is a bunch of software patches that solves many problems while a PTF (is the same as Fix and means Program Temporary Fix) is a patch that solves one specific problem. You will download Maintenance Levels (ML) as APARs from IBM Software Web Site.

2) You should install the latest Maintenance Level for the AIX version you have installed (usually a big bunch of software up to 650 MB that needs almost 1GB space to be decompressed and installed). As AIX 5L is new technology from IBM they're patching many problems and generating ML very often. You can download from
http://www-1.ibm.com/servers/eserver/support/pseries/aixfixes.html

3) First, you have to know which Fix or PTF to install, then download it from the above web link, then copy to a location in the server (usually PTF's are copied to /usr/sys/inst.images directory as well as there's enough space (what i do is to create a new FS of some 2 GB dize and mount it over /usr/sys/inst.images, after installing the APAR or PTF i just delete the FS without deleting the mount point). Then uncompress or unzip, untar, whatever, and using the fastpath   smitty update_all in AIX you can install or preview the installation of any patches. I recommend using preview option before real installation and also recommend installing patches in APPLIED status, that is, both either original or old version and newest version of the software are installed, so you can REJECT the installation of any patch.

4) You can remove any single fileset with the fastpath   smitty remove

5) A COMMITed software is installed and the only way to reject it is by uninstalling the software fileset while a APPLIED software is installed and the preview versions of filesets are installed too so if you REJECT the APPLIED software then those older versions will be active again.

Checking on known maintenance levels:

esmf04m-strombrg> oslevel -qr

Known Recommended Maintenance Levels

5100-04

5100-03

5100-02

5100-01

esmf04m-strombrg> lppchk -v

Dual booting AIX:

>Okay you install AIX 5.1 on hdisk0 as example and boot your maschine. th=

>an

>you clone your rootvg to hdisk1 :

>alt_disk_install -C hdisk1

>so you have hdisk0 with old_rootvg

>and hdisk1 with alt_*rootvg

>

>bootlist -m hdisk0 hdisk1 (means you boot from hdisk0 first and hdisk1

>second)

>

>boot with AIX5.2 CD and install with Migartion Option from prompt on

>hdisk0.

>

>now you have Aix5.2 on hdisk0 and aix5.1 on hdisk1

>

>if you want to remove the alternate disk install:

>alt_disk_install -X

Installing an IBM maintenance release upgrade:

Go to the IBM Support Fix Central site:

http://www-912.ibm.com/eserver/support/fixes/fcgui.jsp

* Server

Select "Pseries family" or the series that your server is.

* Product or fix type

Select "AIX OS, java, compilers"

* Ordering option

Select "Mainteneance packages"

* OS level

Select "AIX 5.1"

Select "continue" for next screen

Current level

Select "5100-04"

Desired Level

Select "5100-05"

Select "go"

Download "510405.tar.gz " at the bottom of the page

Follow the instructions

Locking an account:

The following procedure can be used to lock a user's account;
(1) smitty user
(2) select, change the characteristics of a user
(3) Expiration Date: input the effective date, when this account will be expiring / closing
(4) Is this user account locked: false, use tab key to choose true
(5) User can login:true, use tab key to change true to false
(6) user can login remotely:true, use tab key to change true to false

(7) Press enter key and account will be locked

(8) for further security also change the password

to permit the user to login after 30 days / specfied time revert the above

fields to original values.

If an ESMF node mostly falls off the net (strobe shows only about 5
ports open), then:

1) Go down to the ESMF HMC

2) Log in

3) Locate the right window to use

4) Log in to the trouble machine

5) kill and restart srcmstr

6) startsrc -s inetd

7) startsrc -s sshd

8) startsrc -s automountd

9) /etc/nfs.clean

10) /etc/rc.nfs

There may be other things that need to be started up as well, but this
has been sufficient so far.

Following the documentation if you issue the following command you will
activate HMT or Hardware MultiThreading

```
# bosdebug -H on
Memory debugger          off
Memory sizes             0
Network memory sizes      0
Kernel debugger          off
Real Time Kernel          off
HMT                      on
```

...but only if your hardware -supports- HMT!

Definiing a virtual network interface:

ifconfig en# alias xxx.xxx.xxx.xxx

Checking if NFS is active:

lssrc -a | egrep nfs

```
 biod          nfs          20752  active
 nfsd          nfs          21426  active
 rpc.mountd     nfs            27888  active
 rpc.statd     nfs           22730  active
 rpc.lockd      nfs            24280  active
```

---

nfso -o nfs_use_reserved_ports=1

Find where gzip lives, package-wise:

which_fileset gzip

Get the machine model:

esmf04m-strombrg> /usr/bin/uname -M

IBM,7039-651

esmf04m-strombrg>

---

lsconf

Looks a lot like prtconf?

You can check microcode version by issuing the following command

lsmcode

if this does not work, then

lscfg -vp | grep -i alterable

You can download Fixes and microcodes not only for your Server nut for any

peripheral devices from

techsupport.services.ibm.com

---

1. Type no -o tcp_keepinit=3750 The initial timeout for TCP/IP will change
   from 75 seconds to 31.25 minutes. The time (3750) is in 1/2 seconds.

2. Type no -o tcp_keepidle=86400 The connection will be kept alive
for 12 hours.

The above two items will not be active once a reboot is done. If this
solves your problem you can add the statements to your /etc/rc.tcpip file.

---

filemon Command

Monitors the performance of the file system, and reports the I/O activity on
behalf of logical files, virtual memory segments, logical volumes, and physical
volumes.

---

lsfs
...can be used to check what kind of filesystem a filesystem is

---

portmir
Apparently can be used to snoop on a tty/pty on AIX? A bit like screen
or VNC, but without the forethought requirement.

Restoring from a mksysb tape:

You can either boot from your mksysb Medium (band Streamer or cdrom) and restore.

change your bootlist:

if you have a Band Streamer, so you can boot from AIX Installation Medium and choose point 3 (Maintenance mode) and restore from media.

Determing what needs to be upgraded to advance to a higher os level:
you can do an "instfix -i | grep ML" to list which maintenance level is incomplete and then show what filesets are required i.e. if AIX 5.2 ML02 is incomplete do "instfix -ivk 5200-01_AIX_ML | grep ":" | grep not"

Nice page with AIX OpenSSH bff's, a script for creating bff's, a script for setting up LBX for use with ssh, and more.
http://www.zip.com.au/~dtucker/openssh/

An example mksysb backup:

mksysb /dev/rmt0

Creating tape boot image ...

Creating list of files to back up .
Backing up 68614 files.............................
17379 of 68614 files backed up (25%).............................
25331 of 68614 files backed up (36%).............................
25341 of 68614 files backed up (36%).............................
55359 of 68614 files backed up (80%)................
68614 of 68614 files backed up (100%)
0512-038 mksysb: Backup Completed Successfully.

```
# echo $PATH
/usr/ucb:/bin:/usr/bin:/etc:/usr/lpp/ssp/bin:/usr/lib/instl:/usr/sbin:/usr/local/bin
#
```

Note the PATH! The backup failed when I had a larger PATH.

IBM's document describing AIX to Solaris admins:

http://www.redbooks.ibm.com/abstracts/sg246584.html?Open

Changing the boot device:

Boot from aix cd's into maint shell and run the bosboot -ad /dev/hdisk0
command.

Or if the hd5 boot device is mirrored on hdisk0 and hdisk1 all you need
to do is boot into sms menu and ensure both disks are selected in the
boot order.

To access sms hit 1 before it does a speaker test.
/////
You can boot it up into what used to be called SMS mode .. i.e. hit F1 at
the 'keyboard' prompt ... You can change the boot device from there and
then make sure that you rerun your bosboot once you have booted up.
Couldn't be simpler

```
manctsr/ >lsvg rootvg -p
rootvg:
PV=5FNAME          PV STATE    TOTAL PPs   FREE PPs   FREE DISTRIBUTION
hdisk0          active    542      245       28..00..00..108..109
hdisk1          active    542      245       28..00..00..108..109
manctsr/ >lsvg rootvg -l
rootvg:
LV NAME           TYPE      LPs  PPs  PVs  LV STATE     MOUNT POINT
hd5             boot    1   2   2   closed/syncd  N/A
```

```
hd6          paging   64   128  2   open/syncd   N/A
hd8          jfslog   1    2    2   open/syncd   N/A
hd4          jfs      1    2    2   open/syncd   /
hd2          jfs      27   54   2   open/syncd   /usr
hd9var       jfs      3    6    2   open/syncd   /var
hd3          jfs      5    10   2   open/syncd   /tmp
hd1          jfs      1    2    2   open/syncd   /home
apachelv     jfs      5    10   2   open/syncd   /apache
cv4=5Fhome   jfs      172  344  2   open/syncd
/export/cv4=5Fhome
cv4=5Fdec    jfs      15   30   2   open/syncd /export/cv4=5Fd=
ec
lv00         jfs      2    4    2   open/syncd   /mn/script
```

```
# lsvg rootvg -l
rootvg:
LV NAME        TYPE      LPs  PPs  PVs  LV STATE      MOUNT POINT
hd5            boot      1    2    2    closed/syncd N/A
hd6            paging    40   80   2    open/syncd    N/A
hd8            jfslog    1    2    2    open/syncd    N/A
hd4            jfs       1    2    2    open/syncd    /
hd2            jfs       71   142  2    open/syncd    /usr
hd9var         jfs       1    2    2    open/syncd    /var
hd3            jfs       2    4    2    open/syncd    /tmp
hd1            jfs       4    8    2    open/syncd    /home
hd10opt        jfs       2    4    2    open/syncd    /opt
log1           jfslog    1    2    2    closed/syncd N/A
paging01       paging    9    18   2    open/syncd    N/A
hd14           jfs       4    8    2    closed/syncd N/A
```
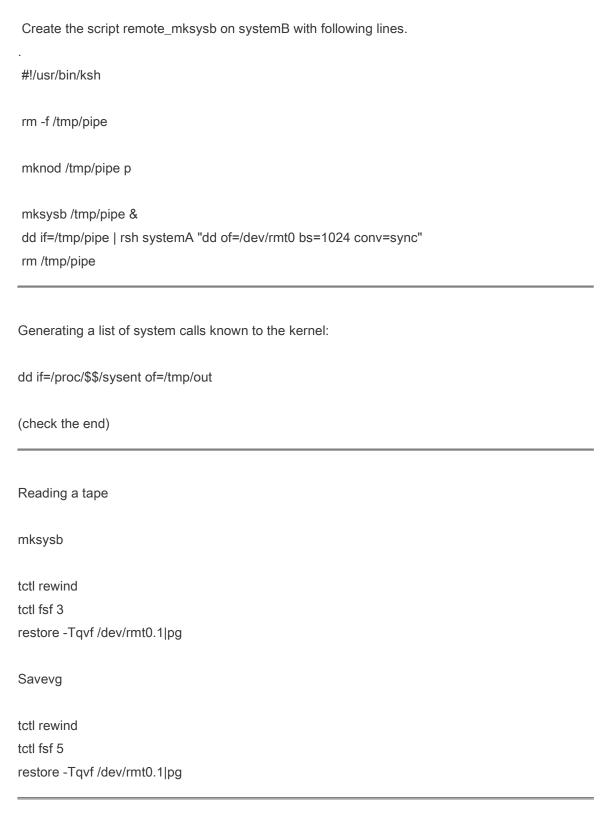
PSSP has it's own 5 CD set (PSSP-3.5) and has to be ordered. Thanks

Don't panic! DISK_ERR4 (in errpt) is just a bad block relocation and
is a somewhat

"normal" occurrence. You only need to be concerned about these errors
if you notice them increasing in number on the same disk. So - you need
to track it but not necessarily replace it.

What kernel level (lslpp -l 'bos.[um]p*')

The hardware must be CHRP (Common Hardware Reference Platform) in order for
5.2 or greater to be supported.

You can determine that by issuing "bootinfo -p".

Nice article on AIX backups:

http://www.ahinc.com/aix/backup.htm

A fix for some kinds of tape backup problems:

please post the output of the following command

lsattr -El rmt0

we are looking for the value "ret error", if this is set to true then i'd
recommend changing it to false by issuing a smitty devices->Tape
devices->Change Tape Devices

How to create mksysb to a remote tape drive.
.
 **** Note mksysb will not be bootable ***
.
 Lets say tape drive is on systemA and you need to create
 mksysb of system
 You should be able to do rsh from systemB to systemA

Create the script remote_mksysb on systemB with following lines.

.

#!/usr/bin/ksh

rm -f /tmp/pipe

mknod /tmp/pipe p

mksysb /tmp/pipe &
dd if=/tmp/pipe | rsh systemA "dd of=/dev/rmt0 bs=1024 conv=sync"
rm /tmp/pipe

Generating a list of system calls known to the kernel:

dd if=/proc/$$/sysent of=/tmp/out

(check the end)

Reading a tape

mksysb

tctl rewind
tctl fsf 3
restore -Tqvf /dev/rmt0.1|pg

Savevg

tctl rewind
tctl fsf 5
restore -Tqvf /dev/rmt0.1|pg

I think that all these following commands mean the same thing :

```
# bootinfo -y
32
# prtconf -c
CPU Type: 32-bit
# bootinfo -K
32
```

On alt_disk_install:

We use it mainly to reduce downtime while upgrading the systems and also
to have a quick back out path. You can have the new built image install
on the alt disks. Switch boot device to the new partition and your newly
upgraded system up and running. If your system has any problems you
cannot fix with adjustments, you can switch back to the old partition
and bring out the old software.

Outage time is little over a reboot worth of time.

First, try to start the switch adapter daemon (worm) with rc.switch.

Good luck - these SP switch problems are notoriously hard to fix.

Enabling quotas on a JFS filesystem (and perhaps others) :

Edit /etc/filesystems and edit in quota=userquota on the relevant filesystem.

```
esmf04m-root> chfs -a "quota = userquota" /home
esmf04m-root> quotaon /home
esmf04m-root> quotacheck /home
```

If a program proves too large to compile with the default options due
to a toc overflow, please try adding:

-Wl,-b -Wl,bigtoc

...to your $CC or $LDFLAGS

---

bash-2.05b$ lsattr -El ent0

alt_addr   0x000000000000   Alternate ethernet address
True
busintr   553   Bus interrupt level
False
busmem   0xf8080000   Bus memory address
False
chksum_offload yes   Enable hardware transmit and
receive checksum True
compat_mode   no   Gigabit Backward compatability
True
copy_bytes   2048   Copy packet if this many or less
bytes   True
flow_ctrl   yes   Enable Transmit and Receive
Flow Control   True
intr_priority 3   Interrupt priority
False
intr_rate   10000   Interrupt events processed per
interrupt   True
jumbo_frames   no   Transmit jumbo frames
True
large_send   yes   Enable hardware TX TCP
resegmentation   True
media_speed   Auto_Negotiation Media speed
True
rom_mem   0xf8040000   ROM memory address
False
rx_hog   1000   RX buffers processed per RX
interrupt   True
rxbuf_pool_sz 2048   Rcv buffer pool, make 2X rxdesc_que_sz

True

rxdesc_que_sz 1024   RX descriptor queue size
True

slih_hog   10   Max Interrupt events processed
per interrupt True

tx_que_sz   8192   Software transmit queue size
True

txdesc_que_sz 1024   TX descriptor queue size
True

use_alt_addr   no   Enable alternate ethernet address
True

---

lsslot -c pci
Slot   Description   Device(s)
U0.1-P1-I1 PCI-X capable, 64 bit, 133MHz slot Empty
U0.1-P1-I2 PCI-X capable, 32 bit, 66MHz slot   Empty
U0.1-P1-I3 PCI-X capable, 32 bit, 66MHz slot   pci9 lai0
U0.1-P1-I4 PCI-X capable, 64 bit, 133MHz slot Empty
U0.1-P1-I5 PCI-X capable, 64 bit, 133MHz slot Empty
U0.1-P1-I6 PCI-X capable, 64 bit, 133MHz slot Empty

Operating System and Devices

Split a Mirrored Disk from a Volume Group

Beginning with AIX 5.2, snapshot support helps you protect the consistency of your mirrored volume groups from potential disk failure. Using the snapshot feature, you can split off a mirrored disk or disks to use as a reliable (from the standpoint of the LVM metadata) point-in-time backup of a volume group, and, when needed, reliably reintegrate the split disks into the volume group. In the following procedure, you first split off a mirrored disk from a volume group and then you merge the split-off disk into the original volume group. To further ensure the reliability of your snapshot, file systems must be unmounted and applications that use raw logical volumes must be in a known state (a state from which the application can recover if you need to use the backup).

A volume group cannot be split if any one of the following is true:

A disk is already missing.

The last non-stale partition would be on the split-off volume group.

Any stale partitions exist in the volume group, unless you use the force
flag (-f) with the splitvg command.

Furthermore, the snapshot feature (specifically, the splitvg command)
cannot be used in enhanced or classic concurrent mode. The split-off
volume group cannot be made concurrent or enhanced concurrent and there
are limitations to the changes allowed for both the split-off and the
original volume group. For details, read the chvg command description in
AIX 5L Version 5.2 Commands Reference.

Ensure that the volume group is fully mirrored and that the mirror
exists on a disk or set of disks that contains only this set of mirrors.

To enable snapshot support, split off the original volume group (origVG)
to another disk or set of disks, using the following command:
splitvg origVG
At this point, you now have a reliable point-in-time backup of the
original volume group. Be aware, however, that you cannot change the
allocation on the split-off volume group.
Reactivate the split-off disk and merge it into the original volume
group using the following command:
joinvg origVG
At this point, the split-off volume group is now reintegrated with the
original volume group.

Configuring ntp

1) Stop the xntpd daemon

The xntpd daemon is managed by the System Resource Controller (SRC).
To verify that the xntpd daemon is active : lssrc -s xntpd : status
should be "active"

To stop the xntpd subsystem : stopsrc -s xntpd

Note : xntpd is automatically started in /etc/rc.tcpip. To verify this :
cat /etc/rc.tcpip | grep xntpd.

2) Modify the /etc/ntp.conf file

Put the following lines in the /etc/ntp.conf file :

server <IP address or FQDN of timeserver> prefer
driftfile /etc/ntp.drift
tracefile /etc/ntp.trace

3) Restart the xntp daemon

To restart the xntpd daemon :
startsrc -s xntpd

4) Check status of time synchronization

To check the status of the time synchronisation, use the ntpq utility.
ntpq -i : start ntpq interactively

ntpq> peer

```
    remote           refid        st t when poll reach
delay   offset   disp
=======================================================================
======
* .PPS.         1 u  863 1024  377    0.92   0.160   0.47
```

The "offset" field displays the difference (in milliseconds) between the
system time and the reference time.
Type "quit" to exit the ntpq utility.

Kind of like ldd:

dump -X32 -Tv /bin/ls

Getting security notices from IBM:

https://techsupport.services.ibm.com/server/pseries.subscriptionSvcs?mode=2

Changing prngd to listen on a socket, using chsys:

esmfcws-root> chssys -s prngd -a '-f /dev/egd-pool -m 666 tcp/localhost:708'
0513-077 Subsystem has been changed.
esmfcws-root> ps -ef | grep prng
   root 303186 1015878   0 19:19:43 pts/2 0:00 grep prng
   root 1007836 262212   0 19:04:42   - 0:01
   /opt/freeware/sbin/prngd -f /dev/egd-pool -m 666
esmfcws-root> stopsrc -s prngd
0513-044 The prngd Subsystem was requested to stop.
esmfcws-root> startsrc -s prngd
0513-059 The prngd Subsystem has been started. Subsystem PID is 852062.
esmfcws-root> ps -ef | grep prng
   root 852062 262212   0 19:20:42   - 0:01
   /opt/freeware/sbin/prngd -f /dev/egd-pool -m 666 tcp/localhost:708
   root 1007846 1015878   0 19:20:54 pts/2 0:00 grep prng
esmfcws-root> /usr/lo
local   lost+found
esmfcws-root> /usr/local/sbin/gen-pas
Not bad, using prngd for entropy
cf4b01142c33d9bd06f1e50d6968f4da
esmfcws-root>

Or if prngd isn't already partially set up:

esmf04m-root> mkssys -s prngd -p /opt/freeware/sbin/prngd -u root -a
'-f /dev/egd-pool -m 666 tcp/localhost:708'

0513-071 The prngd Subsystem has been added.

esmf04m-root> lssrc -s prngd

Subsystem   Group   PID   Status

 prngd   inoperative

esmf04m-root>

esmf04m-root> for i in 1 2 3 4 5 6 7 8; do ssh esmf0${i}m "mkssys -s

prngd -p /opt/freeware/sbin/prngd -u root -a '-f /dev/egd-pool -m 666

tcp/localhost:708'"; done

0513-071 The prngd Subsystem has been added.

0513-071 The prngd Subsystem has been added.

0513-071 The prngd Subsystem has been added.

0513-075 The new subsystem name is already on file.

0513-071 The prngd Subsystem has been added.

0513-071 The prngd Subsystem has been added.

0513-071 The prngd Subsystem has been added.

0513-071 The prngd Subsystem has been added.

esmf04m-root> for i in 1 2 3 4 5 6 7 8;

do ssh esmf0${i}m "stopsrc -s prngd"; done

0513-004 The Subsystem or Group, prngd, is currently inoperative.

0513-004 The Subsystem or Group, prngd, is currently inoperative.

0513-004 The Subsystem or Group, prngd, is currently inoperative.

0513-044 The prngd Subsystem was requested to stop.

0513-004 The Subsystem or Group, prngd, is currently inoperative.

0513-004 The Subsystem or Group, prngd, is currently inoperative.

0513-004 The Subsystem or Group, prngd, is currently inoperative.

0513-004 The Subsystem or Group, prngd, is currently inoperative.

esmf04m-root> for i in 1 2 3 4 5 6 7 8; do ssh esmf0${i}m "startsrc

-s prngd"; done

0513-059 The prngd Subsystem has been started. Subsystem PID is 25880.

0513-059 The prngd Subsystem has been started. Subsystem PID is 34508.

0513-059 The prngd Subsystem has been started. Subsystem PID is 30670.

0513-029 The prngd Subsystem is already active.

Multiple instances are not supported.

0513-059 The prngd Subsystem has been started. Subsystem PID is 37450.

0513-059 The prngd Subsystem has been started. Subsystem PID is 21266.

0513-059 The prngd Subsystem has been started. Subsystem PID is 27662.

0513-059 The prngd Subsystem has been started. Subsystem PID is 42666.
esmf04m-root>

Don't forget /etc/prngd.conf

---

please issue the following commands:

fuser -c /mnt

and check for any PID that maybe locking your CD device. If you have any,
you can kill them all bye issuing

fuser -ck /mnt

and try to eject the CDROM. If this doesn't work at all, then check for
the cdromd daemon (new feature ported for Solaris into AIX) with the
following command:

lssrc -a | grep cdrom

if cdromd is running, then you should umount the cdrom device:

cdmount
cdumount /cdrom/cdXX
cdeject

Linux, by default, requires any NFS mount to use a reserved port below 1024.
AIX, by default, uses ports above 1024. Use the following command to
restrict AIX to the reserved port range:

/usr/sbin/nfso -o nfs_use_reserved_ports=1

Creating a subsystem:
mkssys -s smbd -p /opt/freeware/sbin/smbd -u 0 -a "-D" -d -q -S -n 15
-f 9 -G tcpip

But it's useless since smbd make fork.

Sincerely,

Lev

AIX system firmware upgrade (pSeries?) :

Sysplanar is something like motherboard in Intel domain, i.e. it is hardware.
It is possible to upgrade firmware when in maintenance mode - when there
is E1F1 on the LCD display right on the machine press key 1 (not on the
numeric keyboard) if you have ASCII terminal.
If you have graphical console press functional key 'F1'
you will be directed to standalone diagnostics menu

the firmware you can find here together with description:
http://techsupport.services.ibm.com/server/mdownload2/download.html

if you cannot boot and have the shell prompt you can do it according
the paragraph 'Updating with the Diagnostic Service Aid Method' - see
the description from the link mentioned above.

in the diagnostics menu you can find 'current firmware as well (there
is something like 'Display config' there')

Diagnostics can be ran against a single device while online

use the
diag -d devicename

---

bindprocessor -q (will give you the number of proc.)
lscfg -v (will give your system info.)
lsmcode -A (will give you the proc. firmware + others)

---

chuser maxage=0 username

Some good stuff on OpenMP and AIX (among other things):

http://www.rz.rwth-aachen.de/ewomp03/OMPtools.html

Someone on AIX-L indicated that this was a good vmtune for a database system:

/usr/samples/kernel/vmtune -p 5 -P 20   (to set the max perm and min

perm values)

Getting an AIX machine's serial number:

esmf04m-root> uname -m

0020D3FA4C00

LoadLeveler upgrade PMR# 70374-227 - website only showing linux downloads

of loadleveler, no AIX downloads

From a post on AIX-L:

IBM recommends the following formula to calculate the amount of paging

space you need...

For memories larger than 256 MB, the following is recommended:

total paging space = 512 MB + (memory size - 256 MB) * 1.25

For 1024MB RAM = 1600MB Paging Space

1 LP = 64 MB = add 17 LP's to = 1600MB

This is what we use while running AIX 5L.

Changing a forgotten root password on AIX:

1.   Insert the product media for the same version and level as the current installation into the appropriate drive.
2.   Power on the machine.
3.   When the screen of icons appears, or when you hear a double beep, press the F1 key repeatedly until the System Management Services menu appears.
4.   Select Multiboot.
5.   Select Install From.
6.   Select the device that holds the product media and then select Install.
7.   Select the AIX version icon.
8.   Define your current system as the system console by pressing the F1 key and then press Enter.
9.   Select the number of your preferred language and press Enter.
10.   Choose Start Maintenance Mode for System Recovery by typing 3 and press Enter.
11.   Select Access a Root Volume Group. A message displays explaining that you will not be able to return to the Installation menus without rebooting if you change the root volume group at this point.
12.   Type 0 and press Enter.
13.   Type the number of the appropriate volume group from the list and press Enter.
14.   Select Access this Volume Group and start a shell by typing 1 and press Enter.
15.   At the # (number sign) prompt, type the passwd command at the command line prompt to reset the root password. For example:

16.   # passwd
17.   Changing password for "root"
18.   root's New password:
Enter the new password again:

19.   To write everything from the buffer to the hard disk and reboot the system, type the following:

sync;sync;sync;reboot

---

turning off diagnostic lights:

/usr/lpp/diagnostics/bin/usysfault -s normal

AIX filesystems and quotas:

http://unix.derkeiler.com/Newsgroups/comp.unix.aix/2003-11/0744.html

---

2005-06-23
IBM informs me that PSSP is never going to be ported to AIX 5.3. There is
a followon product like PSSP called "CSM", and it runs on recent AIX and
Linux, but it is not going to support an SP2 switch, like the ESMF has.

Redirect console messages to a specific file of your choosing:
swcons /tmp/console.messages

Checking if an AIX machine is still marketed and/or supported by IBM:

http://www-306.ibm.com/common/ssi/OIX.wss

Like tcpdump/ethereal?

iptrace -e -i lo0 /tmp/iptrace.out, (let it run for 5 minutes, kill it)
ipreport /tmp/iptrace.out

---

lscfg -vp | grep -e "Memory DIMM" -e "Size"

Memory DIMM:

  Size........................256

Memory DIMM:

  Size........................256

Memory DIMM:

  Size........................256

Memory DIMM:

  Size........................256

Clipped from a message on AIX-L - outlines the procedure for replacing
a bad disk in a logical volume:

u must procee in tyhos order:

1- unmirror the rootvg (unmirror rootvg hdisk1)

2- extrcat hdisk1 from rootvg (reducevg rootvg hdisk1) hidsk1 should not
have any other data, if yes, move them first

3- rmdev -dl hdisk1

4- put the new pv

5- cvrmgrl

6- extendvg rootvg "the new pv"

7-mirrorvg rootvg hdsikxxx

/////
And another:

Use this redbook, page 182, section 6.5.1.

http://www.redbooks.ibm.com/abstracts/SG245496.html?Open

On -some- IBM (PowerPC) machines, you boot to singleuser by hitting F5
during the boot

Where to get firmware for pSeries machines:

"I believe the p in p-Series stands for Performance.

While the i in i-Series stands for Integrated."

"I believe the p in pSeries stands for Power as in the power 5 chip

architecture the hardware uses."

OK, from the (0)> prompt enter either ? or h - these subcommands list

all the available subcommands you can key into the kdb at the (0)>

prompt. Unfortunately, unless you know what you are looking for its

hard to understand the output.

The common commands to use are stat and staus - which will show the

status of the system and dump, vmlog and vmstat will show any memory

errors that may have caused the dump.

You really need an indepth knowledge of how the system works to

decipher most of the output and Im afraid theres no easy way to do it.

This link has a list of all the kdb subcommands

---

bindprocessor is for binding a process to a specific CPU

---

esmf04m-root> sysdumpdev -l

primary   /dev/lv00

secondary   /dev/sysdumpnull

copy directory   /var/adm/ras

forced copy flag   TRUE

always allow dump   TRUE

dump compression   OFF
Wed Oct 26 13:43:31

From a IBM AIX partner:

GIL is a kernel process, which does TCP/IP timing. It handles
transmission errors, ACKs, etc. Normally it shouldn't consume too much
CPU, but it can take quite a lot of CPU when the system is using the
network a lot (like with NFS filesystems which are heavily used).
.
The kproc gil runs the TCP/IP timer driven operations. Every 200ms, and
every 500ms the GIL thread is kicked to go run protocol timers. With TCP
up (which is ALWAYS the case), TCP timers are called which end up
looking at every connection on the system (to do retransmission, delayed
acks,etc). In version 4 this work is all done on a multi-threaded kproc
to promote concurrency and SMP scalability.gil.

GIL is one of the kprocs (kernel processes) in AIX 4.3.3, 5.1 and 5.2.
Since the advent of topas in AIX 4.3.3 and changes made to the ps
command in AIX 5.1, system administrators have become aware of this
class of processes, which are not new to AIX. These kprocs have no
user interfaces and have been largely undocumented in base
documentation. Once a kproc is started, typically it stays in the
process table until the next reboot. The system resources used by any
one kproc are accounted as kernel resources, so no separate account is
kept of resources used by an individual kproc.
.
Most of these kprocs are NOT described in base AIX documentation and
the descriptions below may be the most complete that can be found.
.
GIL term is an acronym for "Generalized Interrupt Level" and was
created by the Open Software Foundation (OSF), This is the networking
daemon responsible for processing all the network interrupts, including
incoming packets, tcp timers, etc.
Exactly how these kprocs function and much of their expected behavior
is considered IBM proprietary information.

In the event of a power failure, from "jessie" on the AIX-L mailing list:

check you error report for an entry that states
EPOW_SUS_CHRP
if there is an entry post it in detail to have a look at the Power status
registers, and the sense data.

If it is not a true failure such as a fan, or power supply then you would
notice in the logs that the problem started after a shutdown, or power
failure...

"pstat -S will associate processor to process but not
process to processor. It is a matter of opinion if
this is what you want. "

Superb page on AIX:

http://www.douzhe.com/docs/jh/9/97757.html

...but I think there may be a bit of a mistake on how to do backups to
a remote tape drive... dd -should- work for that, but IME, it doesn't.

AIX supports large pages with 32-bit and 64-bit kernels. Applications,
either 32-bit or 64-bit,
can take advantage of large pages. The extended common object file format
(XCOFF or
XCOFF64), the object file format for AIX, provides a flag to identify
binaries if they are set (or
cleared) to use large pages (or turn the large pages flag) through ldedit10.
The flag can also
be turned on at load time (ld)10 with the following commands:

ld command: ld -blpdata -o a.out

ldedit command: ldedit -blpdata a.out (or -bnolpdata a.out)

An AIX upgrade procedure:

I just went through this with my company, and wrote some directions as
to what we should do; I will share this document with you.

You should do a complete configuration management scheme/snapshot of
your system:
1) execute df -lk
2) execute lsvg, lsvg -p for each vg, and lsvg -l for each vg
3) execute lspv
4) execute bootlist -m normal -o and bootlist -m service -o
5) execute bootinfo -y and bootinfo -k
6) execute lspv -a
7) execute lsvg -M rootvg
8) execute lsconf

You want to document everything from above so that you can have this to
re-create your system should there be any mistakes or unfortunate
events.

This just helps you to know exactly what your system looks like, before
you make any changes.

Go to this site and you will get exactly what you need:
http://www-03.ibm.com/servers/eserver/support/unixservers/aixfixes.html

Choose the -> AIX 5.3 link and choose follow the prompts to get you the
correct maintenance level(s).

Please let me know if this is of any help.

Thanks.

LeRoy S. Phillips 'Phil'
UNIX System Administrator (AIX/SAP)

From a message on IBM-AIX-L:

I get these stupid messages all the time and I just filter them and send
them to junk.
I've tried making the sysdumpdev bigger, but it comes back and wants it
to be just a little bigger than I made it.
IBM does recommend that you use a second sysdumpdev.

/////////////////////////////////
                   SYSTEM DUMP
/////////////////////////////////
IBM recommends:
 Don't mirror the system dump device
 Don't use compression on the dump device
 Don't use a secondary dump device unless it is on a separate device,
separate cable and separate i/o card.

sysdumpdev -l   Lists current dump destination.
sysdumpdev -e   Estimates dumpsize of the current system in bytes.
sysdumpdev -L   Displays information about the previous dump.
sysdumpdev -c   <-- the system dump device will not be compressed
when the next dump is taken
sysdumpdev -p (dump device) -P   Sets the default dump device, permanently
sysdumpdev -P -s /dev/sysdumpnull   <--   makes the secondary
dump device a bit bucket (recommended)

sysdumpstart -p   Starts a dump and writes to the primary dump device.
sysdumpstart -s   Starts a dump and writes to the secondary dump device.

(MCA machine can also dump if key is in service position and the reset
 button is pressed)

Analyze dump file :-

echo "stat status t -m" | crash /var/adm/ras/vmcore.0

$ errpt

IDENTIFIER TIMESTAMP T C RESOURCE_NAME DESCRIPTION

F89FB899   0822150005 P O dumpcheck   The copy directory is too small

This message is the result of a dump device check. You can fix this by increasing the size of your dump device. If you are using the default dump device (/dev/hd6) then increase your paging size or go to smit dump and "select System Dump Compression". Myself, I don't like to use the default dump device so I create a sysdumplv and make sure I have enough space. To check space needed go to smit dump and select "Show Estimated Dump Size" this will give you an idea about the size needed.

The copy directory is whatever sysdumpdev says it is.

Run sysdumpdev and you will get something like

#sysdumpdev

primary   /dev/hd6

secondary   /dev/sysdumpnull

copy directory   /var/adm/ras

forced copy flag   TRUE

always allow dump   FALSE

dump compression   ON

sysdumpdev -e

0453-041 Estimated dump size in bytes: 57881395

Divide this number by 1024. This is the free space that is needed in your copy directory. Compare it to a df -k or divide this number by 512. This is the free space that is needed in your copy directory. Compare it to a df