

Notes for migrating to EMC or Shark with Fibre Channel boot

R. Ballough, 1/15/2004

Planning Notes

1. Reference document SANBootQA0103.pdf for a complete list of requirements when booting from SAN attached disk http://w3-1.ibm.com/sales/systems/portal/_s.155/254?navID=f220s240&geolD=AM&prodID=pSeries&docID=wsfcupdate1107
2. PowerPath is allowed with boot devices; SDD is not at this time. With SDD, rootvg should be mirrored at the host level across two separate HBAs for redundancy.
3. PowerPath or SDD are preferred options for multipathing at this time. MPIO, native to AIX, is also available for multipathing; however, support for features such as loadbalancing vary by vendor, and some implementations may only support path failover.
4. Check Maintenance Level on AIX using 'oslevel' command. With AIX 5.2, make sure that ML2 is installed.
5. Make sure FC microcode is at current levels. Check with the command 'lscfg -vl fcsX', where fcsX is the fibre channel adapter to be connected. Field (Z9) indicates the microcode level. Verify this with current microcode at <https://techsupport.services.ibm.com/server/mdownload>
6. If the adapters are being added to an existing server, the following filesets may need to be installed:
devices.fcp.disk.rte -> Required for any FC attachment to AIX
devices.common.IBM.fc.rte -> Required for any FC attachment to AIX
devices.pci.df1000f7.com -> Required for any FC attachment to AIX
devices.pci.df1000f7.rte -> Required for 6227 adapter
devices.pci.df1000f7.diag -> Required for 6227 adapter
devices.pci.df1000f9.rte -> Required for 6228 adapter
devices.pci.df1000f9.diag -> Required for 6228 adapter
devices.pci.df1080f9.rte -> Required for 6239 adapter
devices.pci.df1080f9.diag -> Required for 6239 adapter
7. Identify WWPN of HBA to give to SAN administrator for LUN configuration. In 'lscfg -vl fcsX' output, this is the Network Address field.

Implementation Notes

EMC:

The following are instructions to install & configure EMC Symmetrix devices, with all devices (including rootvg) using PowerPath .

1. Install EMC ODM Support Package:
 - Download from ftp://ftp.emc.com/pub/symm3000/aix/ODM_DEFINITIONS/
 - Install and reboot
 - If using LPARs, make the following entry in /etc/environment:
PP_LPAR_KEY_FIX=1
2. Have SAN administrator assign EMC LUN(s) to one HBA (or alternatively, assign them to all desired HBAs and disconnect the fibre from all but 1 HBA)
3. Detect new disk(s)
cfgmgr -vl fcsX
4. Assign PVID to new disk(s)
chdev -l hdiskX -a pv=yes
5. Clone rootvg to new hdisk. The alt_disk_install command is nice for this because it copies the data, creates the boot image, and changes the bootlist. Be careful if the existing rootvg is mirrored, as alt_disk_install will then expect the target disk(s) to have enough space to contain two mirrors.
alt_disk_install -C hdiskX
 - reboot
6. Install PowerPath software + applicable patches through SMI T
7. Register PowerPath software
 - **emcprep -install**
 - **'powermt check_registration'** command checks to make sure registration key is valid
8. Initialize the Powerpath hdiskpower devices:
powermt config
9. Enable PowerPath for rootvg:
 - **pprootdev on**
 - reboot
10. Have SAN administrator assign EMC LUN(s) to remaining HBAS (or plug in fibre ports)
11. Run EMC's version of cfgmgr to detect new paths:
/usr/lpp/Symmetrix/bin/emc_cfgmgr
12. Bosboot now needs to be run, but it expects an hdisk device, but instead now has an hdiskpower device. Run the following command to fix this:
pprootdev fix
13. Rebuild boot image on new hdiskpower boot device:
bosboot -ad /dev/ipldevice

14. Manually update bootlist to include ALL hdisks that correspond to the hdiskpower boot device:

- Identify hdisks to hdiskpower relationship using 'powermt display dev=all'
- Update bootlist to include all hdisks corresponding to the hdiskpower boot device:

```
bootlist -m normal hdiskX hdiskY
```

- Reboot

SHARK:

The following are instructions to install & configure Shark LUNs with all except rootvg using SDD.

1. Download & Install 2105 host attachment package:
http://www-1.ibm.com/support/docview.wss?rs=540&context=ST52G7&q=host+scripts&uid=ssq1S4000106&loc=en_US&cs=utf-8&lang=en+en
 - reboot
2. Download & Install SDD package:
http://www-1.ibm.com/support/docview.wss?rs=540&context=ST52G7&q=&uid=ssq1S4000065&loc=en_US&cs=utf-8&lang=en+en
 - no reboot needed
3. Have SAN administrator assign SHARK LUN(s) to appropriate HBAs. In the case of the LUNs for rootvg, make two separate LUNs. Each rootvg LUN should be visible off of only one HBA. For redundancy, make the first LUN for rootvg off of a different HBA than the second LUN. Redundancy should also be planned across switches if possible.
4. Run `'cfdmgr'` to detect LUNs.
5. Run `'lsdev -Cc disk | grep 2105'` to check that LUNs are configured
6. For all devices to be excluded from SDD control (rootvg devices, for example) run `'querysn -l hdiskX'`. This places an entry into `/etc/vpexclude`.
7. Change the `qdepth_enable` field to 'no' if running a database:
 - `/usr/lib/methods/defdpo`
 - `chdev -l dpo -a "qdepth_enable=no"`
 - Verify using `'lsattr -El dpo'`
8. Configure all vpaths. The disks in `/etc/vpexclude` will remain outside of SDD control
`cfallvpath`
9. Assign PVID to new boot disks:
`chdev -l hdiskX -a pv=yes`
10. Clone rootvg to new hdisk. The `alt_disk_install` command is nice for this because it copies the data, creates the boot image, and changes the bootlist. Be careful if the existing rootvg is mirrored, as `alt_disk_install` will then expect the target disk(s) to have enough space to contain two mirrors.
`alt_disk_install -C hdiskX`
reboot
11. Mirror rootvg to target LUN.
`mirrorvg rootvg hdiskY`
`bosboot -ad /dev/hdiskY`
`bootlist -m normal -o hiskX hdiskY`

References:

EMC references:

PowerPath version 3.0 for UNIX Installation and Administration Guide, PN 300-00-511 REV A02, available for download from <http://powerlink.emc.com>

IBM references:

Fibre Channel Marketing Bulletins at http://w3-1.ibm.com/sales/systems/portal/_s.155/254?navID=f220s240&geoID=AM&prodID=pSeries&docID=rscfsk.skit&docType=SalesKit&skCat=DocumentType

IBM TotalStorage Enterprise Storage Server Host Attachment Guide, pub SC26-7446-04, online at <http://publibfp.boulder.ibm.com/epubs/pdf/f2bhs04.pdf>

Subsystem Device Driver User's Guide, pub SC26-7608-00, online at <http://www-1.ibm.com/support/docview.wss?uid=ssg1S7000303&aid=1>

Software Required for Fibre Channel Attachment to AIX:
<http://www-1.ibm.com/support/docview.wss?uid=isg1hintsTips1304>