

What ideas can FreeBSD borrow from AIX?

Jan Szrednicki
w@wrzask.pl



So, what is this AIX thing?

- IBM's proprietary UNIX-like OS
- Runs solely on POWER/PowerPC architecture
- Targeted for commercial use up to the "enterprise" level
- Actively developed, advanced in many fields, quite specific
- Remarkably consistent command names

Why and what to borrow?

- AIX is closed source and full of patents
- No point in comparing *BSD and AIX
- But we can look at the concepts and adapt them (and make them the better way :)
- Main focus: system administration tools and interfaces
- Make good use of the power of the base system

Background and rationale

- Experience from various tasks and projects at ComArch SA
- Mostly external customers – systems not exactly properly maintained
- AIX is rather resilient to mistreatment :)
- The power of built-in tested tools
- My target: developers and committers



Main topics

- Device management
- Administration interface – SMIT
- Upgrade, cloning and archiving
 - Package management
 - System dump – `mksysb`
 - Root filesystem tricks – `alt_disk_*`
- Error logging



Device management

- All devices accessed with an uniform set of tools (both physical and virtual)
- Full control: gathering data, setting tunables, (re)scanning, diagnosis and removal
- Clean device hierarchy
- FreeBSD already has *newbus* (and `devinfo` tool) for a good start

Device management, *cont.*

- Persistent device names!
- Data stored in the ODM (Object Data Manager) database
- Device states: available, stopped, defined
- Particularly helpful for hotplug devices
- **Tool list:** `lsattr`, `lscfg`, `lsdev`, `cfgmgr`, `rmdev`, `chdev`, `diag`



Usage example

```
# lsdev -c disk
hdisk0 Available 06-08-01-3,0 16 Bit LVD SCSI Disk Drive
hdisk1 Available 06-08-01-4,0 16 Bit LVD SCSI Disk Drive
hdisk2 Available 06-08-01-5,0 16 Bit LVD SCSI Disk Drive
# lscfg -pvs -l hdisk2 | head -n10
  hdisk2          U788C.001.AAA6754-P1-T11-L5-L0
    16 Bit LVD SCSI Disk Drive (73400 MB)

    Manufacturer.....IBM    H0
    Machine Type and Model.....HUS103073FL3800
    FRU Number.....03N5262
    ROS Level and ID.....52505152
    Serial Number.....00579F87
    EC Level.....H17923D
    Part Number.....26K5573
# lsattr -l hdisk2 -E -F 'attribute:value' -a 'unique_id,pvid,size_in_mb'
unique_id:2B0800579F870FHUS103073FL380008IBM    H0scsi
pvid:00cd730d5549ca290000000000000000000
size_in_mb:73400
# rmdev -d -l hdisk2
hdisk2 deleted
# cfgmgr
█
```

The way I'd see it in FreeBSD

- Friendlier syntax :)
- At least two device namespaces:
 - In-kernel, fully dynamic, read only
 - "Normal", persistent, changeable
- Integrated functionality of:
`devinfo, pciconf, atacontrol,`
`camcontrol, usbdevs/usbconfig` etc.

The almighty SMIT

- Menu-driven text (or GUI) all-powerful system administration interface
- Built on top of normal command line utilities (can be displayed)
- A cool tool for novices and when you need to do something quick without digging through command line options
- Relies on the binary ODM database
- Extensions from packages

SMIT example

Software Installation and Maintenance

Move cursor to desired item and press Enter.

Install and Update Software

List Software and Related Information

Software Maintenance and Utilities

Software Service Management

Relocatable Software Installation and Maintenance

Network Installation Management

EZ NIM (Easy NIM Tool)

System Workload Partition Software Maintenance

System Backup Manager

Alternate Disk Installation

EFIX Management

Thin Server Maintenance

F1=Help

F2=Refresh

F3=Cancel

F8=Image

F9=Shell

F10=Exit

Enter=Do

SMIT example

Update Installed Software to Latest Level (Update All)

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

```
[TOP]                                     [Entry Fields]
* INPUT device / directory for software
* SOFTWARE to update
PREVIEW only? (update operation will NOT occur)  no      +
COMMIT software updates?                        yes     +
SAVE replaced files?                            no      +
AUTOMATICALLY install requisite software?       yes     +
EXTEND file systems if space needed?            yes     +
VERIFY install and check file sizes?            no      +
DETAILED output?                                no      +
Process multiple volumes?                       yes     +
ACCEPT new license agreements?                  no      +
Preview new LICENSE agreements?                  no      +
```

[MORE...6]

```
F1=Help          F2=Refresh      F3=Cancel       F4=List
F5=Reset         F6=Command     F7=Edit         F8=Image
F9=Shell         F10=Exit       Enter=Do
```

SMIT in FreeBSD?

- Binary database – bad idea, must rely solely on command line tools
- Good for making novices familiar with the system
- May enforce good habits :)



Package management

- Just one feature worth noting: package state
 - APPLIED – changes are effective, but not final
 - COMMITTED – changes are final
- Easy way to roll back changes when something breaks down
- Can be built into `portupgrade` or `portmaster`

mkysb

- AIX has a good habit of placing the whole system on `rootvg`, while recommending other VGs for data
- One command to archive the whole system, possibly creating bootable media
- Can be used to efficiently clone customized system, package-agnostic

mksysb sample usage

```
# mksysb -i /tmp/mksysb.img

Creating information file (/image.data) for rootvg..

Creating list of files to back up.

Backing up 90871 files.....backup: 0511-089 Cannot open ./var/
tmp/slp_srvreg.lock: Cannot open or remove a file containing a running program.
.....
22258 of 90871 files (24%).....
58685 of 90871 files (64%).....

0512-003 mksysb may not have been able to archive some files.
The messages displayed on the Standard Error contained additional
information.
# mkcd -LS -I /tmp/isos/ -m /tmp/mksysb.img
Initializing mkcd log: /var/adm/ras/mkcd.log...
Verifying command parameters...
Creating temporary file system: /mkcd/cd_fs...
Populating the CD or DVD file system...
Building chrp boot image...
```

mkysb sample usage, cont.

```
Copying backup to the CD or DVD file system...
.....
Creating Rock Ridge format image: /tmp/isos//cd_image_479382.vol1
Running mkisofs ...
.....
.....
.....
mkrr_fs was successful.

Making the CD or DVD image bootable...

Copying the remainder of the backup to the CD or DVD file system...
Creating Rock Ridge format image: /tmp/isos//cd_image_479382.vol2
Running mkisofs ...
.....
mkrr_fs was successful.

Removing temporary file system: /mkcd/cd_fs...
# █
```

alt_disk_copy & co.

- Copies the current system (or `mksysb` image) to a new alternative root volume group
- Can be used to perform updates to a new `rootvg` while preserving the old one
- Uses LVM to manage seamless operation



alt_disk_copy, *cont.*

- One command to create the new VG, one command to revert to the old one
 - fully automated
 - no place for mistakes
 - upgrade without downtime, only reboot with a new volume group



alt_disk_copy, *cont.*

- **Commands:**
 - `alt_disk_copy` – copies the current sys
 - `alt_disk_mksysb` – from `mksysb` image
 - `alt_rootvg_op` – operations on alternative root volume groups
- On FreeBSD we can take advantage of ZFS features to provide similar scripts
- Yet another mechanism to roll back changes when something is wrong

Error logging

- AIX has a central error log facility, independent of syslog
 - `errpt` is used to browse and filter entries
 - `errclear` to remove entries
- Based on templates defined for each type of error out there
- Not just logging – the template mechanism actively collects data from the system and saves them in the log

Error logging example

```
l488pp032_pub[/] > errpt
IDENTIFIER  TIMESTAMP  T C RESOURCE_NAME  DESCRIPTION
A924A5FC    0702145710 P S SYSPROC        SOFTWARE PROGRAM ABNORMALLY TERMINATE
F7FA22C9    0702145710 I O SYSJ2         UNABLE TO ALLOCATE SPACE IN FILE SYST
A6DF45AA    0622044610 I O RMCdaemon     The daemon is started.
2BFA76F6    0622044610 T S SYSPROC        SYSTEM SHUTDOWN BY USER
9D8CFDEE    0622044610 T O errdemon     ERROR LOGGING TURNED ON
l488pp032_pub[/] > errpt -a -j A924A5FC
-----
LABEL:      CORE_DUMP
IDENTIFIER: A924A5FC

Date/Time:  Fri Jul  2 14:57:47 CDT 2010
Sequence Number: 86
Machine Id: 00F604884C00
Node Id:    l488pp032_pub
Class:      S
Type:       PERM
WPAR:      Global
Resource Name: SYSPROC

Description
SOFTWARE PROGRAM ABNORMALLY TERMINATED

Probable Causes
SOFTWARE PROGRAM
```

Error logging example

```
User Causes
USER GENERATED SIGNAL

        Recommended Actions
        CORRECT THEN RETRY

Failure Causes
SOFTWARE PROGRAM

        Recommended Actions
        RERUN THE APPLICATION PROGRAM
        IF PROBLEM PERSISTS THEN DO THE FOLLOWING
        CONTACT APPROPRIATE SERVICE REPRESENTATIVE

Detail Data
SIGNAL NUMBER
        11
USER'S PROCESS ID:
        237790
FILE SYSTEM SERIAL NUMBER
        1
INODE NUMBER
        0          2
CORE FILE NAME
//core
PROGRAM NAME
xmtopas
```

syslog has facilities, why bother?

- Resolved entries can be removed
- syslog leaves all the analysis to the user or external parsing programs
- Templates allow immediate gathering of data which may not be available later
- A quick way to assess the health of the system



Questions?



Links

- <http://www-03.ibm.com/systems/power/software/aix/>
- http://aixwiki.org/wiki/Main_Page
- http://42.pl/u/2mPe_AIX_cmds
<http://publib.boulder.ibm.com/infocenter/aix/v6r1/topic/com.ibm.aix.cmds/alphabeticallistofcommands.htm>
- <https://www.ibm.com/developerworks/wikis/display/WikiPtype/AIX>
- <http://www.linux2aix.com/aix-pseries/tips-tricks/ibm-aix-cheat-sheets.html>

