

Free Tools, Hints, and Tricks

Russell Adams
rladams@adamsinfoserv.com

Thursday 9/24/2009

- Houston AIX Users Group
- September 24th, 2009

- Principal Consultant @ Adams Information Services LLC
- <http://www.adamsinfoserv.com/>
- IBM CATE
- Provides Nationwide consulting services for IBM's AIX, pSeries, SAN Storage, and HACMP

- Valuable Free Add-on Software!

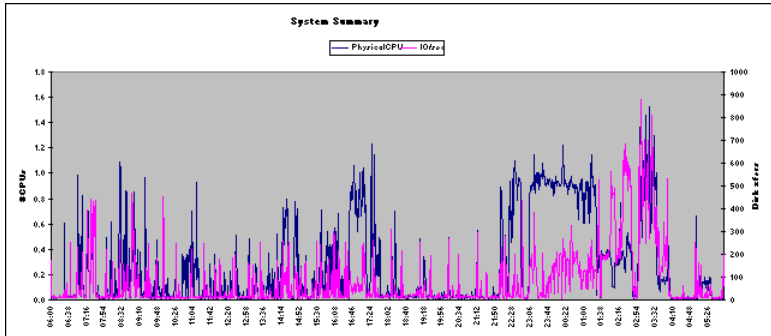
- NMON is an open tool written by Nigel Griffiths
- URL
- Now included in AIX as a fully supported tool!
 - 5.3 TL9
 - 6.1 TL2
 - VIOS 2.1
- Comprehensive Performance Data Collection
- Record performance data before you need it

- Quick to setup
- Create 2GB /var/nmon
- Add cron job
 - `0 0 * * * /usr/bin/nmon -f -m /var/nmon`
 - Collection proceeds for 23:55, named by host and date

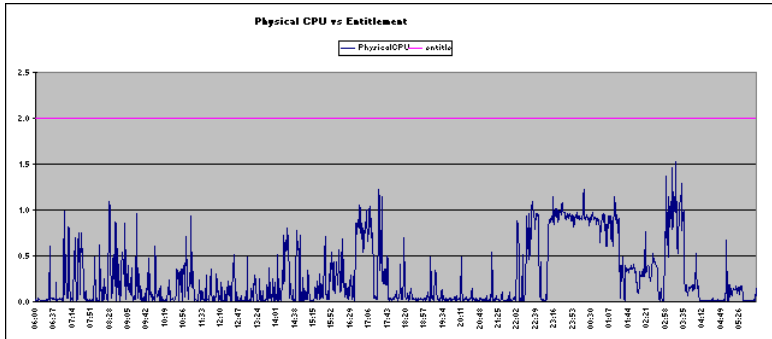
- NMON Analyzer is an open tool written by Stephen Atkins
- URL
- Excel spreadsheet/macro to analyze NMON recording

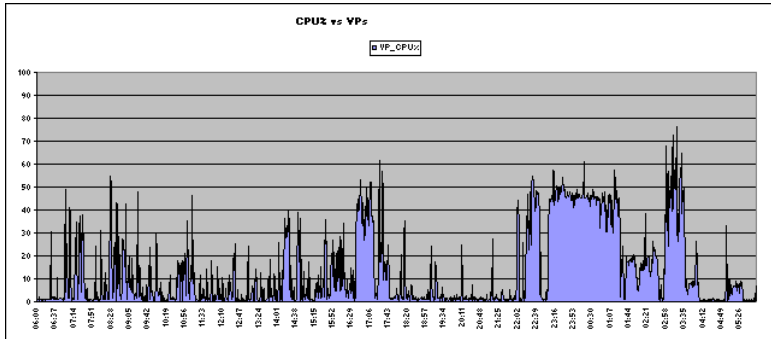
- CPU Utilization, including APV/Micropartitioning
- Memory Utilization
- IO Throughput
- Disk Busy %
- Network Traffic

NMON Analyzer

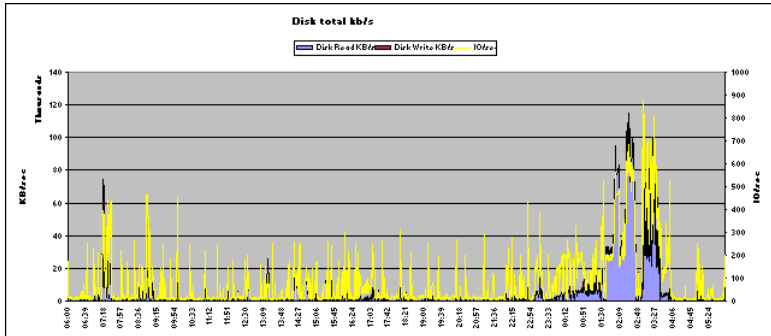


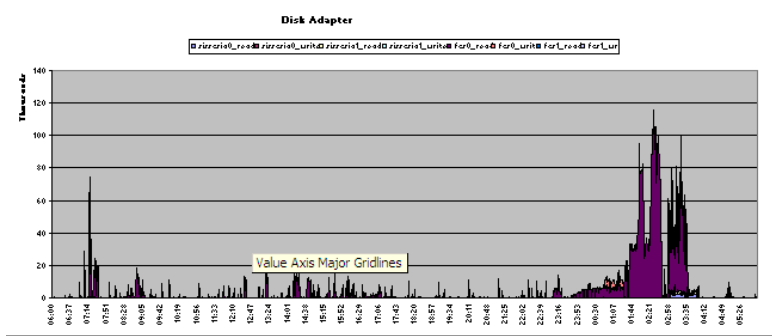
NMON Analyzer

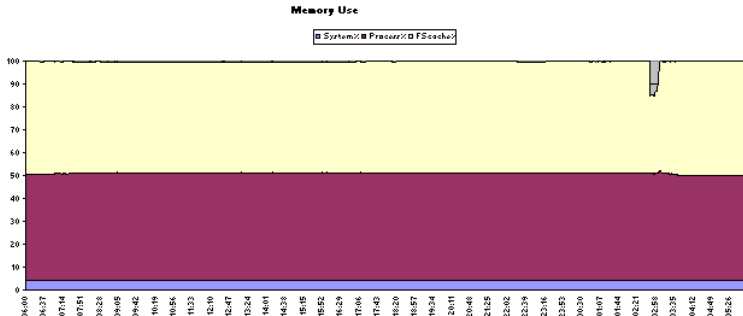




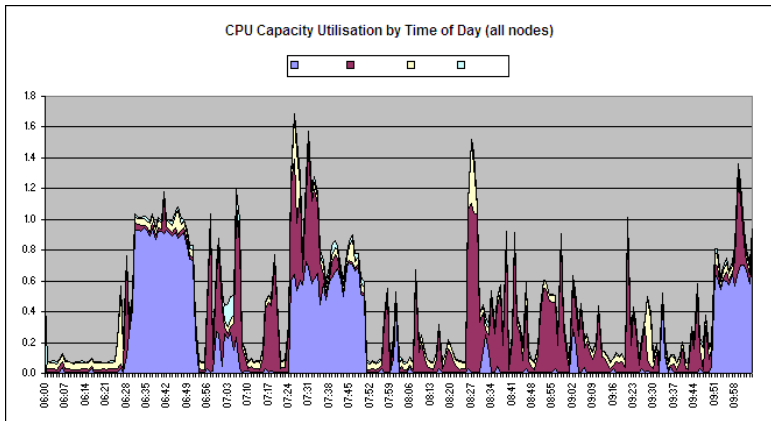
NMON Analyzer

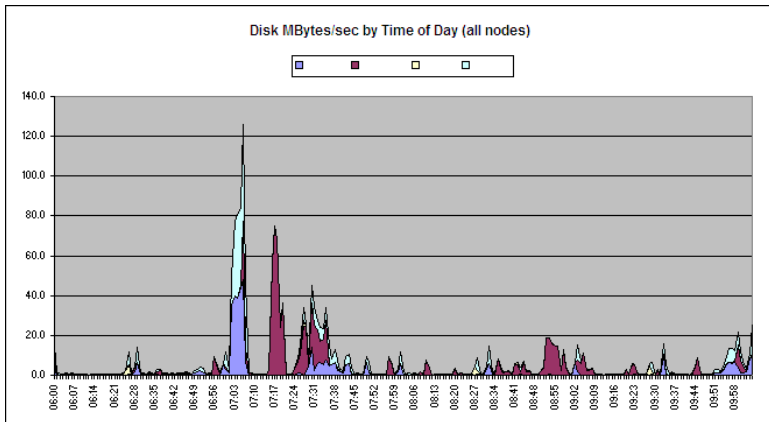


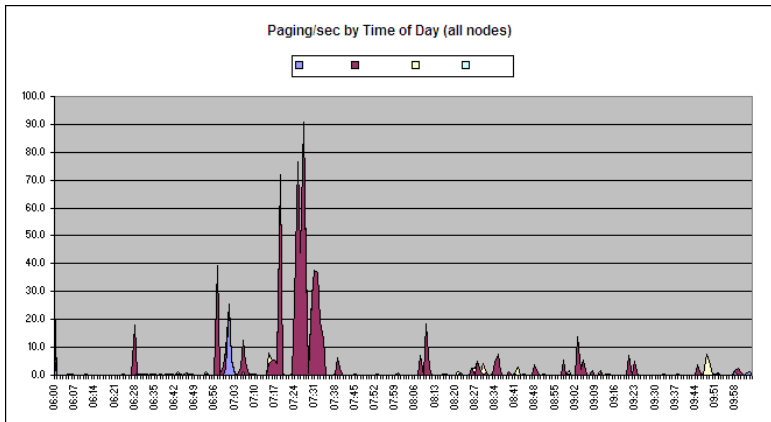




- Another Excel tool written by Stephen Atkins
- URL
- Combine NMON recordings for all LPARs in a CEC
- Narrow focus on CEC Utilization
 - CPU / Memory / IO







- URL
- Fantastic OSS Tools
- Precompiled for AIX
- Easy to install RPMs
- Ships with AIX on CD

- URL
- Migrated out of Toolkit to Expansion Pack
- Recommended as more secure replacement of:
 - telnet
 - ftp
 - r-commands (rcp, rsh, etc...)
- Key-based authentication can minimize password usage
- OpenSSH provided on AIX Linux Toolkit CD

- URL
- Migrated out of Toolkit to Expansion Pack
- Lists open files
- Additional functionality over `fuser`

- URL
- Differential Remote File Copy Tool
- Standalone Daemon or over SSH
- Maintain Mirrors, DR copies, Backups
- Data Migration Tool

- URL
- Allow non-root users to run specific commands
- Free alternative to tools like Powerbroker

- URL
- Multiplex your terminal window
- Use to allow multiple viewers of a session
- Leave long running jobs to connect to later

- URL
- Smart tab completion
- Arrow key support
- Can be made default shell for regular users
- **Do not change root's default shell**

- URL
- Less is more
- Smart paging
- Arrow key support

- URL
- Fetch files via URL at the command line
- Skip downloading to PC then FTP
- ie: Updating invscout catalog

- Emacs: URL
- VIM: URL
- Color syntax highlighting
- Modern editors for sophisticated editing

- Wireshark: URL
 - Packet sniffing and Network debugging
- Apache: URL
 - Basic webserver
- Coreutils: URL
 - Color ls output
- Diffutils: URL
 - Unified Diff

- Zip/Unzip: URL
 - Share zip files with PC's
- Expect: URL
 - TCL extension for scripting text applications
- RCS: URL
 - Simple single file version control
 - URL

- Org-mode: [URL](#)
 - Note taking, organizer, TODO lists, etc.
- Microsoft Virtual Desktop Manager: [URL](#)
 - Split into desktops to organize windows

- Backup and Restoring AIX
 - URL
- Stacking mksysb and savevg
 - URL
- Automating HMC Backup
 - `bkconsdata -r ftp -h hostname -u username -passwd password`
 - Can trigger remotely via ssh w/ keys in cronjob
 - URL

- Backup VIOS
 - URL
 - URL
 - URL
 - Create NFS mount w/ option root=vioserver
 - After first `nimresources.tar` file, use `-mksysb` option for speed

- Automate download of software updates
- Ideal for NIM server LPP sources

SUMA Example

```
suma -x  
-a Action=Download  
-a RqType=SP  
-a RqName=5300-05-05  
-a FilterML=5300-05  
-a DLTarget=/usr/sys/inst.images/5300-05-05  
-a FilterSysFile=/dev/null  
-a FilterDir=/usr/sys/inst.images  
-a Supersedes=n  
-a ResolvePE=Never  
-a Extend=y  
> /usr/sys/inst.images/5300-05_5300-05-05.txt
```

- Poor man's documentation
- Essentially a compressed tarball
- `snap -ac` captures **all** system data

- filesystem
 - df -k
 - mount
 - lsfs
 - lsvg
 - lspv
 - lslv

- general
 - lsattr -El
 - lspp.hac
 - oslevel.info
 - Copies of common system files from /etc
- kernel
 - bootinfo
 - vmo -L
 - lssrc

- lsvirt
 - lsvirt.out all Virtual Device Mappings
- lvm
 - lspv
 - lsvg

- svCollect
 - VIO server mapping data
- tcpip
 - netstat
 - Common tcpip config files

- Database memory usage conflicts with JFS cache
 - Symptom is high non-computational memory (FS cache)
 - URL
- AIX < 5.2 ML5 or AIX 5.3 ML2
 - Lower minperm% from 20 to 5
 - Lower maxperm% and maxclient% from 80 to 20
- AIX >= 5.2 ML5 or 5.3 ML2
 - Toggle lru_file_repage to 0 (on)
- AIX >= 6.1
 - lru_file_repage now defaults to 0 (on)

- Use INLINE logs for JFS2 Filesystems
- Use diag to light identity LEDs for maintenance
 - diag -> Task Selection -> Identify and Attention Indicators
 - Choose the hardware location from `lscfg -vl <device>`

- Free firmware reports from IBM
- URL
- Requires up to date catalog file
 - URL
 - Copy to `/var/adm/invscout/microcode/catalog.mic`
- Run `invscout`, creates `/var/adm/invscout/hostname.mup`
- Upload `.mup` to URL

Invscout Sample Report

action will be "Update". In some cases it may not be possible to determine the installed level or detect if the installed level is the latest, in this scenario the suggested actions will be "Inspect" or "Research" respectively.

Before installing any microcode, be sure to review its README file.

Microcode by Host

localhost
IP Addr: 127.0.0.1 , Aliases: localhost
Model: 7044-170 Serial: 108C62C
Data taken at: 2009.09.23 05:27:00 Microcode catalog: 2009.09.17

Devices found on localhost :				
Device	Logical Device	Suggested Action	Installed Level	Latest Available
7044-170	group	None	system:SPH05195 service:sh040616	system:SPH05195 service:sh040616
9 1/18.2GB SCSI Drive Model DPSS	hdisk0	Update	53393347	53395241
4.5/9.1/18.2 GB SCSI Drive Model DNES	hdisk1	Update	53414730	53414755
4.5/9.1/18.2 GB SCSI Drive Model DNES	hdisk2	None	53414755	53414755

Invscout Sample Report

<u>Device</u>	<u>Logical Device</u>	<u>Suggested Action</u>	<u>Installed Level</u>	<u>Latest Available</u>
7044-170	group	None	system:SPH05195 service:sh040616	system:SPH05195 service:sh040616
9.1/18.2GB SCSI Drive Model DPSS	hdisk0	Update	53393347	53395241
4.5/9.1/18.2 GB SCSI Drive Model DNES	hdisk1	Update	53414730	53414755
4.5/9.1/18.2 GB SCSI Drive Model DNES	hdisk2	None	53414755	53414755

Microcode by Type

<u>Device</u>	<u>Logical Device</u>	<u>Readme</u>	<u>Latest Available</u>	<u>Release Date</u>	<u>Suggested Action</u>	<u>Impact</u>	<u>Severity</u>
7044-170	group	Readme	SPH05195 sh040616	07/28/2005 11:54:14 AM	None	Serviceability	SPE
9.1/18.2GB SCSI Drive Model DPSS	hdisk0	Readme	53395241	09/18/2001 07:47:29 PM	Update	Performance	ATT
4.5/9.1/18.2 GB SCSI Drive Model DNES	hdisk1	Readme	53414755	05/23/2000 10:53:46 AM	Update	Function	ATT
4.5/9.1/18.2 GB SCSI Drive Model DNES	hdisk2	Readme	53414755	05/23/2000 10:53:46 AM	None	Function	ATT

- `lsdev` default contains software/logical code

ent0 Available Logical Host Ethernet Port (lp-hea)

ent1 Available 01-08 2-Port 10/100/1000 Base-TX PCI-X Adapter (14108902)

ent2 Available 01-09 2-Port 10/100/1000 Base-TX PCI-X Adapter (14108902)

fcnet0 Defined 02-00-01 Fibre Channel Network Protocol Device

fcnet1 Defined 03-00-01 Fibre Channel Network Protocol Device

fcs0 Available 02-00 4Gb FC PCI Express Adapter (df1000fe)

fcs1 Available 03-00 4Gb FC PCI Express Adapter (df1000fe)

fscsi0 Available 02-00-02 FC SCSI I/O Controller Protocol Device

fscsi1 Available 03-00-02 FC SCSI I/O Controller Protocol Device

hdisk2 Available 02-00-02 FC 2145

hdisk3 Available 02-00-02 FC 2145

- `lsdev -F` “name status location physloc description”

hdisk0 Available 10-60-00-4,0 P1/Z1-A4 16 Bit LVD SCSI Disk Drive

hdisk1 Available 10-60-00-5,0 P1/Z1-A5 16 Bit SCSI Disk Drive

hdisk2 Available 10-60-00-6,0 P1/Z1-A6 16 Bit SCSI Disk Drive

- Parent/Child Relationship
 - Find base digits of child device
 - ie: hdisk0 10-60-00-4,0 -> 10-60
- ie: `lsdev | grep 10-60`

cd0 Available 10-60-00-1,0 16 Bit SCSI Multimedia CD-ROM Drive

hdisk0 Available 10-60-00-4,0 16 Bit LVD SCSI Disk Drive

hdisk1 Available 10-60-00-5,0 16 Bit SCSI Disk Drive

hdisk2 Available 10-60-00-6,0 16 Bit SCSI Disk Drive

scsi0 Available 10-60 Wide/Ultra-2 SCSI I/O Controller

- `lscfg -vl <device>`

```
hdisk0 U78A0.001.xxxxxxx-P2-D5 SAS Disk Drive (146800 MB)  
Manufacturer.....IBM  
Machine Type and Model.....HUS153014VLS300  
FRU Number.....10N7232  
Serial Number.....JFVTXXXX  
Part Number.....42R4234  
Device Specific.(Z0).....000005329F003002  
Hardware Location Code.....U78A0.001.DNWGSZB-P2-D5
```

- `lscfg -vp | grep fcs`

*fcs1 U78A0.001.xxxxxxx-P1-C3-T1 4Gb FC PCI Express Adapter
(df1000fe)*

*fcs0 U78A0.001.xxxxxxx-P1-C2-T1 4Gb FC PCI Express Adapter
(df1000fe)*

- `lscfg -vp | grep vhost`

vhost0 U9117.MMA.10xxxxx-V10-C41 Virtual SCSI Server Adapter

vhost1 U9117.MMA.10xxxxx-V10-C51 Virtual SCSI Server Adapter

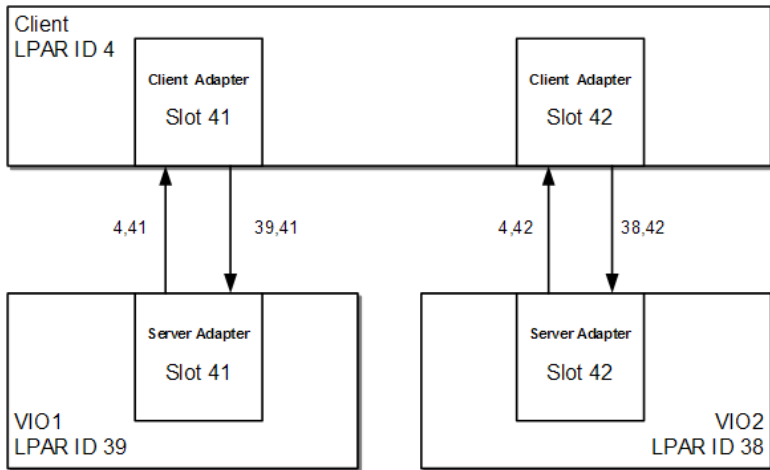
vhost2 U9117.MMA.10xxxxx-V10-C61 Virtual SCSI Server Adapter

vhost3 U9117.MMA.10xxxxx-V10-C71 Virtual SCSI Server Adapter

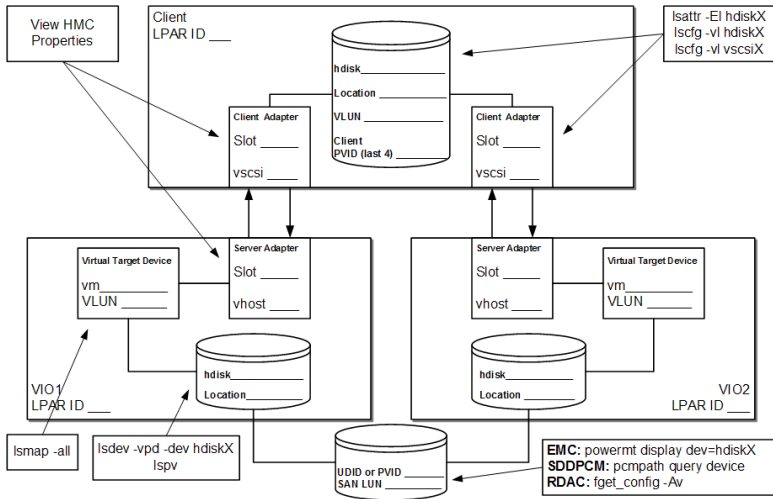
- Create VIO servers at upper limit of LPAR ID's
- 4 CPU box * 10 LPARs / CPU = 40 LPARs max
 - VIO1 LPAR ID 39 (*Note odd!*)
 - VIO2 LPAR ID 38 (*Note even!*)
- Create clients from 2-12
 - 1 is reserved for Full System Partition
- Increase maximum virtual adapters as needed

- Virtual Ethernet in teens (slots 10-19)
- Virtual SCSI numbered by client LPAR ID and VIO Server
 - Client LPAR 4 has VSCSI slot 41 directed to VIO1 slot 41
 - VIO1 has VSCSI slot 41 directed to Client slot 41
- Easier to read and infer client
- LPAR Mobility breaks numbering, uses first available slot on VIOS

Slot Numbering Diagram



VSCSI Trace Worksheet



- Allows transparent migration between direct LUN access and VIO
- Without UDID VIO indexes LUNs by PVID
 - Client writes offset
 - Data “encapsulated”
 - No migration, backup/restore/mirror
- `lsattr -El hdiskX`
 - IBM drivers often list `unique_id` attribute
 - Powerpath hides it, query ODM directly instead
 - `odmget -q attribute=unique_id CuAt`

- `set -o vi` or `set -o emacs`
- Update path
- Update prompt variable (PS1)

- Add to *home/padmin*.profile

```
export ENV=/home/padmin/.kshrc
```

- Create *home/padmin.kshrc*

```
export PS1=""hostname '$ '  
set -o vi
```


HMC Virtual Terminal

- Java Terminal difficult to copy/paste
- ssh to HMC (ie: `ssh hscroot@hmc1.mydomain.com`)
- Old way to access VT via `mkvterm` difficult
- `vtmenu` command presents menu of systems and LPARs
- Opens terminal in ssh session!
- Copy and paste work normally, depending on SSH client
- To exit, `<enter> ~ .`
- May still require GUI to close dead terminals

- Setup temporary NFS share, not preserved on reboot
 - `exportfs -i /directory`
- Recommend using CIFS mounts instead of Samba
 - Instead of serving Windows clients, be a Windows client
 - Push to Windows servers
 - Offload security issues

- smitty eznim
- Setup NIM in under an hour
- Requires AIX CD/DVD
- Creates default SPOT and LPP sources
- Great way to get started with NIM
- URL

KSH Keybindings: Vi Mode

- `set -o vi`
- URL
- Prior command: `Esc k`
- Search history: `Esc /string`
- Filename completion: `Esc \`
- Inline editing movement
 - Forward Char / Backward: `h / l`
 - Forward Word / Backward: `w / b`
 - BOL / EOL: `0 / $`
 - Edit @ point / after point: `i / a`
 - End editing: `Esc`

- `set -o emacs`
- URL
- Prior command: Control-P
- Search history: Control-R string
- Filename completion: Esc Esc
- Inline editing movement
 - Forward Char / Backward: C-f / C-b
 - Forward Word / Backward: Esc-f / Esc-b
 - BOL / EOL: C-a / C-e

- Vi mode users frustrated?
- Smit uses Emacs key bindings!
- Backspace: C-h
- Delete to EOL from point: C-k
- Page down: C-v

- `grep -p`
 - Grep for paragraphs matching regexp
 - ie: `grep -p /var /etc/filesystems`

/var:

dev = /dev/hd9var

vol = "/var"

mount = automatic

check = false

free = false

vfs = jfs2

log = /dev/hd8

type = bootfs

- Iterate Commands across a list of items using `xargs` or `while read`
 - `xargs` applies list to end of supplied command by default
 - `xargs` common flags
 - `-t` echos each command as it runs
 - `-n1` means one line of input per command
 - Space sensitive!

- Iterate Commands across a list of items using `xargs` or `while read`
 - `while read` more flexible
 - Multiple commands
 - Not limited to applying list item at end of command
 - Can use quoting to eliminate whitespace sensitivity

- Always echo the commands first!
- Listing devices example:
 - `lsdev -F name | grep ^fscsi | xargs -tn1 lsattr -El`
- Bulk set attributes example:
 - `lsdev -F name|grep ^fscsi|while read X; do chdev -P -l "$X" -a dyntrk=yes ; done`
- Mount all FS in a VG
 - `lsvgfs myvg | sort | xargs -tn1 mount`

- Remote login monitoring
- Group work / demonstrations
- `portmir` command
- Root only, Only one session per OS
- Find pty via `who` or `ps -ef`
- `portmir -t pts/7`
- End via `portmir -o`

- Drill down method:
 - `du -sm * | sort -n`
 - `cd` to largest suspect and repeat
- 15 Largest files in current directory tree:
 - `find . -type f -size +10000 | while read X ; do du -sm "$X" ; done | sort -n | tail -n 15`

- URL
- Have errdaemon email you copies of errpt entries in real time
- Configure sendmail so that mail works
- Smart host? (DS line in /etc/sendmail/sendmail.cf)
- Create root /.forward file with comma delimited list of emails to notify
- `chmod 400 /.forward`

- Add ODM entry which matches all events and triggers email

errnotify:

en_pid = 0

en_name = "ZZZ_mail_err"

en_persistenceflg = 1

en_method = "errpt -a -l \$1 | mail -s 'Errpt \$1 \$4 \$3 \$9' root"

- Old way to add custom startup script was to add an rc.local to inittab
- Now create a script in /etc/rc.d/rc2.d/S99Local instead

- `echo "Line to append" >> /file/to/append.txt`
 - Easy to copy and paste across systems, or ssh
- `cat > /file/to/overwrite.txt`
 - Paste!
 - Control-D ends
- Compare large lists by checksum (ie: `lslpp`)
 - `lslpp -l | cksum`

