

Power Ask The Experts 2013



Tricks of the Power Masters

Advanced Technology Support, Europe.



Gareth Coates
gaz@uk.ibm.com
[@power_gaz](https://twitter.com/power_gaz)

© 2013 IBM Corporation



IBM Power Ask the Experts 2013

09:30 - 10:00	Registration and coffee
10:00 - 11:15	Power Systems Update - Pat O'Rourke: Austin Briefing Centre
11:15 - 12:30	Performance Best Practices with POWER7 - Nigel Griffiths
12:30 - 13:30	Lunch
13:30 - 14:30	Tricks of the Power Masters - Gareth Coates
14:30 - 15:15	Cost Comparison between IBM Power and Intel - David Spurway
15:15 - 15:30	Coffee
15:30 - 16:45	Power Systems Trends and Directions - Pat O'Rourke: Austin Briefing Centre
16:45	Close

© 2013 IBM Corporation

Introduction

- Lots of tips and tricks.
 - Hopefully useful

- Hardware
- Firmware
- HMC
- VIOS
- AIX

- One liners
- Concepts
- Other pointers which may need more research
 - eg: Redbooks

DISCLAIMER

Neither I nor IBM can take any responsibility for any work you carry out without our direct involvement!

Credits

- The tips here came from many sources
- Personal experience
- Other people, including
 - Nigel Griffiths
 - EMEA ATS
 - Andy Thomas
 - Power FTSS, UK
 - Graeme Penman
 - Power FTSS, UK
 - Paul Chapman
 - SSiS Services, UK
 - Pat O'Rourke
 - Briefing Center (sic), Austin, TX
 - Jay Kruemcke
 - Mr AIX, Austin, TX

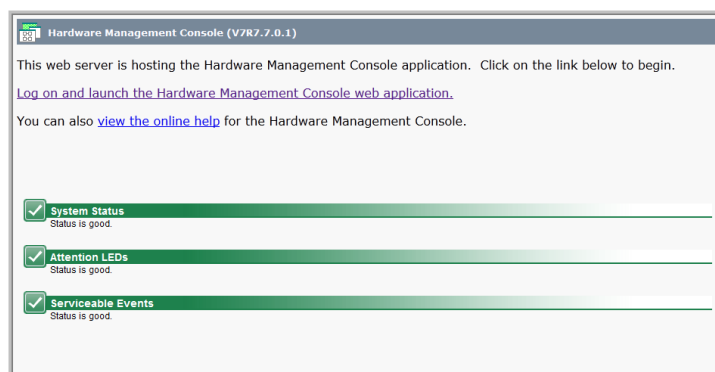
- And many others – too many to mention

A few questions

- Who uses a Web Browser to access/manage Power Servers and HMCs?
 - Internet explorer
 - Firefox
 - Chrome
 - Other

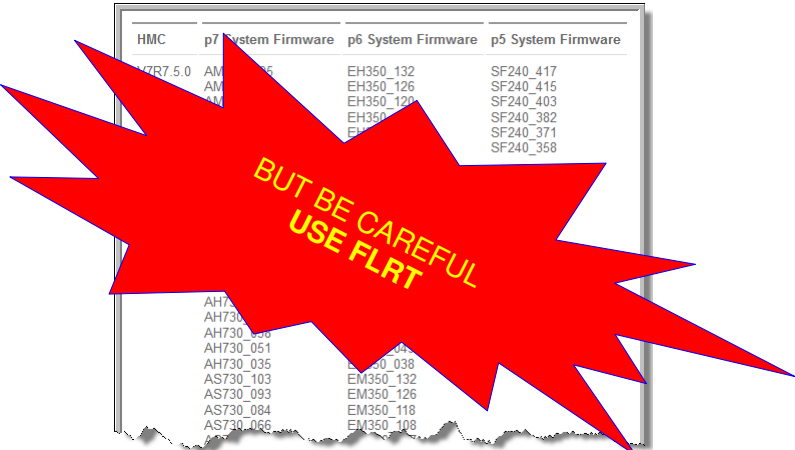
- Who uses:
 - AIX V7.1
 - AIX V6.1
 - AIX V5.3
 - AIX V5.2
 - AIX V5.1
 - AIX V4.3
 - Earlier
 - WPARs
 - vWPARs
 - IBM Systems Director V6.3
 - IBM Systems Director V6.2
 - IBM Systems Director V6.1
 - Earlier

HMC



HMC / Firmware Supported Combinations

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



HMC	p7 System Firmware	p6 System Firmware	p5 System Firmware
POWER7.5.0 AM	EH350_132	SF240_417	
AM	EH350_126	SF240_415	
AM	EH350_120	SF240_403	
	EH350_114	SF240_382	
	EH350_108	SF240_371	
		SF240_358	
AH730_050			
AH730_051			
AH730_035	EM350_038		
AS730_103	EM350_132		
AS730_093	EM350_126		
AS730_084	EM350_118		
AS730_066	EM350_108		

7

monhmc - memory

```
monhmc -r mem -n 0
```

```
Mem: 1028512k total, 1009432k used, 19080k free, 142628k buffers
```

```
monhmc -r mem -n 0
```

```
Mem: 4096084k total, 3854460k used, 241624k free, 429908k buffers
```

- That top HMC only manages one POWER7 750, but even so the amount of free memory is low.
- The bottom HMC has 4GB RAM but even so, there's still not much free.

Richard Wale/UK/IBM

8

monhmc - disk

```
monhmc -r disk -n 0
```

Filesystem on	1K-blocks	Used	Available	Use%	Mounted
/dev/sda2	16121184	8591256	6711016	57%	/
tmpfs	2048432	244	2048188	1%	/dev
tmpfs	2048432	0	2048432	0%	/dev/shm
/dev/sda3	6040320	2998696	2734784	53%	/var
/dev/sda7	8056524	1330396	6316876	18%	/dump
/dev/sda8	250104468	191732	237208136	1%	/extra

9

monhmc - proc & swap

```
monhmc -r proc -n 0
```

```
Cpu0 : 0.0%us, 0.0%sy, 0.0%ni,100.0%id, 0.0%wa, 0.0%hi, 0.0%si, 0.0%st  
Cpu1 : 0.0%us, 0.0%sy, 0.0%ni,100.0%id, 0.0%wa, 0.0%hi, 0.0%si, 0.0%st  
Cpu2 : 0.3%us, 0.0%sy, 0.0%ni, 99.7%id, 0.0%wa, 0.0%hi, 0.0%si, 0.0%st  
Cpu3 : 0.0%us, 0.0%sy, 0.0%ni,100.0%id, 0.0%wa, 0.0%hi, 0.0%si, 0.0%st
```

```
monhmc -r swap -n 0
```

```
Swap: 2040244k total,          0k used, 2040244k  
free, 2354760k cached
```


10

HMC

- Scenario:
 - Managed Server is connected to an HMC
 - HMC is “disconnected”
 - HMC is reconnected
 - HMC does not see the Managed Server
- If you reconnect the managed system to the HMC after you remove the connection, you must use the `mksysconn -o auto` HMC command to clear the connection history on the HMC before reconnecting the managed system.
- If you reconnect the managed system to the HMC without first using the `mksysconn -o auto` command, the HMC does not recognize the managed system.

What servers are connected?

```
hmc11:~ # lssyscfg -r sys -F name
bronze-8203-E4A-SN10E0A21
172.17.254.248
pink-9115-505-SN65080EA
silver-8203-SN10E0A31
172.17.255.246
purple-9117-MMB-SN100525P
plum-8204-E8A-SN105C0B0
oldlace-65BD12E
orange-8203-E4A-SN10E0A51
brown-9115-505-SN6509E5A
grey-9117-MMA-p570-8F
hmc11:~ #
```



My favourite HMC
one-liner

HMC Password policies

- All as `hscroot`, no need for `root`.

- **Make a new password policy**

```
mkpwdpolicy -i \
"name=gaz4, \
description=, \
min_pwage=1, \
pwage=180, \
min_length=4, \
hist_size=10, \
warn_pwage=7, \
min_digits=0, \
min_uppercase_chars=0, \
min_lowercase_chars=0, \
min_special_chars=0"
```

HMC Password policies

- All as `hscroot`, no need for `root`.

- **Activate it**

```
-chpwdpolicy -o a -n gaz4
```

- **change the password**

```
-chmcur -u gaz -t passwd
```

- **disable the policy**

```
-chpwdpolicy -o d
```

- **remove the policy**

```
-rmpwdpolicy -n gaz4
```

VIOS

Which VIOS provides your vSCSI

- Run this as root in a client LPAR

```
# print "cvai" | kdb | grep vscsi | grep -v read
vscsi0      0x000007 0x000000000000 0x0 plum-viol->vhost0
vscsi1      0x000007 0x000000000000 0x0 plum-viol->vhost1
#
```

- shows the VIOS and vhost for a client vscsi adapter.

Security

- Consider the padmin password to be very sensitive
- padmin can gain root privileges
- can then read data from backing devices for all of its clients



17

© 2013 IBM Corporation

Don't use smit

- padmin cannot use smit
- smit is available in the oem_setup_env environment



18

© 2013 IBM Corporation

Don't use smit

- padmin cannot use smit
- smit is available in the oem_setup_env environment

▪ DO NOT USE IT !!!



Create System Plan Failure (mksysplan)

- <https://www-304.ibm.com/support/docview.wss?uid=isg3T1010961>

If the ldc query command fails on the VIO server when run as padmin then there will be additional information about the failures in the `/home/ios/logs/ioscli_ldc.trace` file. However, before you start looking at that file you should first **be prepared to honestly answer the following questions.**

d. Has configuration been performed on the VIO server outside the padmin command line interface (CLI) after exiting the CLI using `oem_setup_env`?

(1) This tech-note cannot begin to outline the extent of damage well meaning system administrations have inflicted on themselves by performing configuration steps outside of the VIOS CLI command shell. If the answer is yes to the above question then the least of your worries is system plan and ldc query. You can check the `ioscli_ldc.trace` file to see if you can correct some errors, but the track record other users have had dealing with misuse of `oem_setup_env` is not very promising.

When to use oem_setup_env

- So, what is oem_setup_env used for?
- It is an environment to setup OEM device drivers
- ONLY use it to install these device drivers as directed in the release notes and installation guides.



What AIX command did that VIO command actually run?

```
$ export CLI_DEBUG=33
$ lsmmap -vadapter vhost0
AIX: "lsdev -c adapter -t IBM,v-scsi-host -s vdevice -F "name" | wc -l -c"
AIX: "lsdev -c adapter -t IBM,v-scsi-host -s vdevice -F "name""
AIX: "lsdev -C -l vhost0 -F "physloc""
AIX: "lsdev -p vhost0 -F "name" | wc -l -c"
AIX: "lsdev -p vhost0 -F "name""
SVSA Physloc Client Partition ID
-----
vhost0 U8233.E8B.100271P-V2-C11 0x00000001

VTD NO VIRTUAL TARGET DEVICE FOUND

$
```



Hardware



list all the WWPn on a system

```

▪ lshwres -r io --rsubtype slotchildren -m Server-9117-MMB-SN101509A -F
  phys_loc,description,mac_address,wwpn,microcode_version |grep Fibre
U78C0.001.DBJJ568-P2-C1-T1,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9b7cbf0,null
U78C0.001.DBJJ568-P2-C1-T2,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9b7cbf1,null
U78C0.001.DBJJ568-P2-C3-T1,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9c6aa32,null
U78C0.001.DBJJ568-P2-C3-T2,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9c6aa33,null
U78C0.001.DBJJ568-P2-C4-T2,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9b65cc3,null
U78C0.001.DBJJ568-P2-C4-T1,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9b65cc2,null
U78C0.001.DBJJ568-P2-C6-T1,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9b7ca38,null
U78C0.001.DBJJ568-P2-C6-T2,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9b7ca39,null
U78C0.001.DBJJ675-P2-C1-T1,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9b65c40,null
U78C0.001.DBJJ675-P2-C1-T2,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9b65c41,null
U78C0.001.DBJJ675-P2-C3-T1,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9c6aa0e,null
U78C0.001.DBJJ675-P2-C3-T2,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9c6aa0f,null
U78C0.001.DBJJ675-P2-C4-T1,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9b65f78,null
U78C0.001.DBJJ675-P2-C4-T2,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9b65f79,null
U78C0.001.DBJJ675-P2-C6-T1,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9b65e00,null
U78C0.001.DBJJ675-P2-C6-T2,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9b65e01,null
U78C0.001.DBJJ696-P2-C1-T1,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9b7d724,null
U78C0.001.DBJJ696-P2-C1-T2,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9b7d725,null
U78C0.001.DBJJ696-P2-C3-T1,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9c6abf0,null
U78C0.001.DBJJ696-P2-C3-T2,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9c6abf1,null
U78C0.001.DBJJ696-P2-C4-T1,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9c6aefc,null
U78C0.001.DBJJ696-P2-C4-T2,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9c6aef1,null
U78C0.001.DBJJ696-P2-C6-T1,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9b7d652,null
U78C0.001.DBJJ696-P2-C6-T2,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9b7d653,null
U78C0.001.DBJJ704-P2-C1-T1,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9b
U78C0.001.DBJJ704-P2-C1-T2,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9b
U78C0.001.DBJJ704-P2-C3-T1,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9c
U78C0.001.DBJJ704-P2-C3-T2,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9c
U78C0.001.DBJJ704-P2-C4-T1,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9c
U78C0.001.DBJJ704-P2-C4-T2,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9c
U78C0.001.DBJJ704-P2-C6-T1,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9b
U78C0.001.DBJJ704-P2-C6-T2,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,1000000c9b

```

Thanks to Nigel Griffiths and Chris Gibson

<https://www.ibm.com/developerworks/mydeveloperworks/blogs/cgaix>



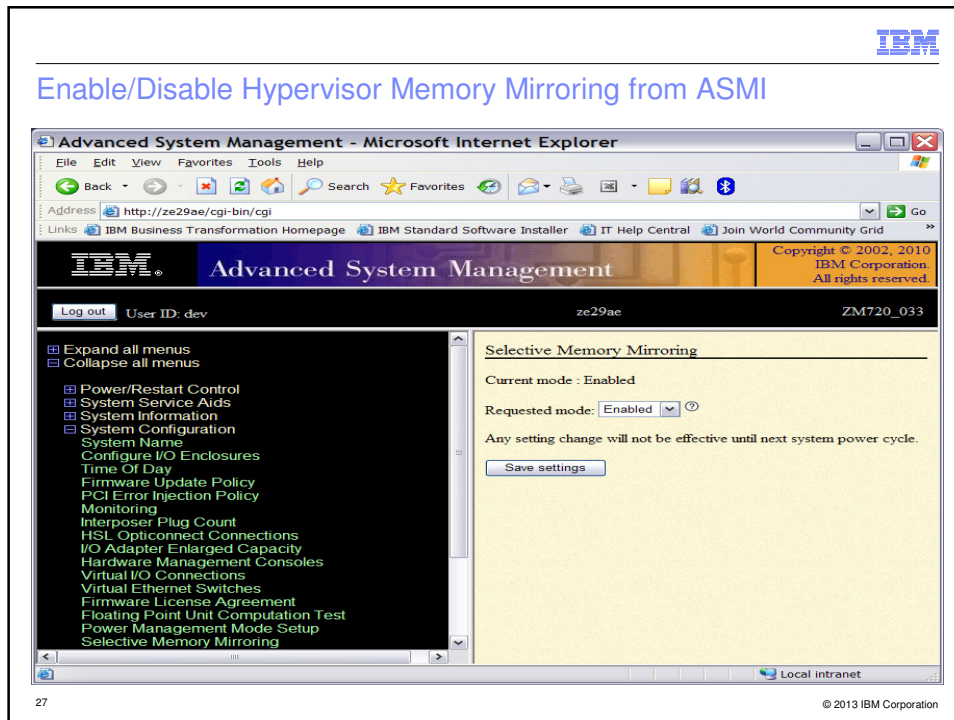
LMB size

- In order to carry out LPM (Live Partition Mobility)
- Many pre-requisites
- LMB (Logical memory Block) size must be the same
- Changing (effective) LMB size requires a power cycle – OF THE ENTIRE SYSTEM
- Choose a size and set all systems to it – now
–It will become effective at the next power cycle.

LMB size

- Check/change it in asmi
 - Check it on the command line
- ```
hmc11:~ # lshwres -r mem -m plum-8204-E8A-
SN105C0B0 --level sys -F mem_region_size
128
hmc11:~ #
```

## Enable/Disable Hypervisor Memory Mirroring from ASMI

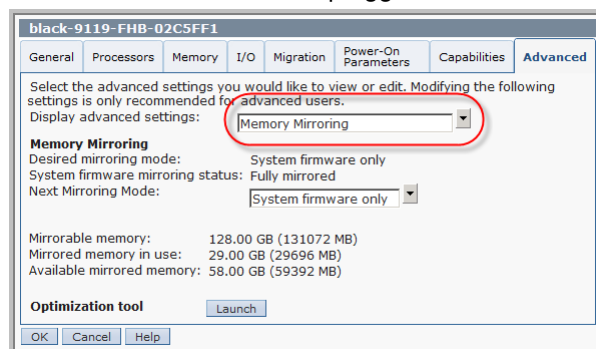


27

© 2013 IBM Corporation

## HMC Mirroring Status

- Advanced Tab on CEC properties
- Desired Mirroring mode
  - Off or System Firmware only
- Option to enable/disable firmware mirroring
- Mirrorable memory
  - Total amount of physical memory that can be mirrored
  - based on the DIMMs that are plugged



28

© 2013 IBM Corporation

## DPO - Dynamic platform optimiser

- EMEA ATS received a Power 760
  - on loan from Austin, Texas in November 2012.
- The product was not announced till 5<sup>th</sup> February 2013
- I was really keen to test DPO
- I gave feedback to the developers and the following slide shows my headline comments.
- The presentation I sent, (pre-announce) was of course, IBM Confidential, but this one is not.
- It was internal IBM communication, so I could be blunt!



29

## DPO - Observations

- Working on DPO has been like reading one of those books that you can't put down!
- Shuffling the VMs (LPARs) by hand, checking the affinity and then watching the optimiser fix it all; has been great.
  - I have certainly consolidated my understanding of POWER7 and POWER7+ LPAR placement and affinity implications
- As the system used for testing is not yet announced, this document is **IBM Confidential**.
  - And that's the only reason that I haven't been tweeting hard about this technology too – **it is really great!**



Now released so no longer confidential!

30

## DPO - motivation

- Partition placement can become sub-optimal
  - Dynamic creation and deletion of partitions
  - DLPAR operations
  - Partition Mobility
  - Hibernation
- Platform will provide a mechanism to optimize partition placement dynamically
- Benefits include
  - Improved performance in a cloud environment
  - Dynamically adjust topology after mobility
  - Simple to use and predicted “score”



Think of it  
as 52 card  
pickup -  
and sort

31

## DPO – HMC commands

```
lssyscfg -r sys -F name
zg23ae
zg24he

lsmemopt -m zg24he -o currscore
curr_sys_score=84

lsmemopt -m zg24he -o calcscore
curr_sys_score=84,predicted_sys_score=93,"requested_lpar_ids=1,2
,17,105,106,107,108,109,110,111",protected_lpar_ids=none

optmem -m zg24he -t affinity -o start

lsmemopt -m zg24he
in_progress=0,status=Finished,type=affinity,opt_id=2,progress=0,
requested_lpar_ids=none,protected_lpar_ids=none,"impacted_lpar
_ids=106,110"

lsmemopt -m zg24he -o currscore
curr_sys_score=95
```


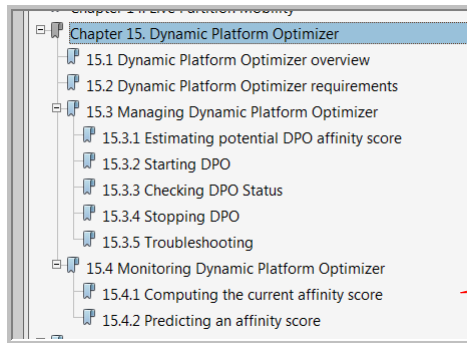


32



## More information

- <http://www.redbooks.ibm.com/redpieces/abstracts/sg247590.html>
- RedBook
- March 2013 update



## POWER7+ (and beyond) Accelerators

- The command to enable crypto to use the accelerator is:
- `acfo -t nx_enabled=0` (disable accelerator)
- `acfo -t nx_enabled=1` (re-enable it)

## POWER7+ (and beyond) Accelerators - HMC commands

- The following are the HMC commands to disable use of the compression/encryption coprocessors for a partition.

The disable commands are:

```
chhwres -m <managed system> -r mem -a "hardware_mem_encryption=0" -o s --id <partition id>
```

```
chhwres -m <managed system> -r mem -a "hardware_mem_expansion=0" -o s --id <partition id>
```

Expansion used to be called compression, so on an older 760 HMC:

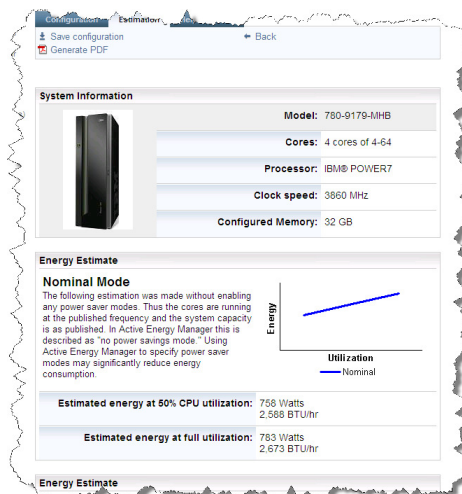
```
chhwres -m <managed system> -r mem -a "hardware_mem_encryption=0" -o s --id <partition id>
chhwres -m <managed system> -r mem -a "hardware_mem_compression=0" -o s --id <partition id>
```

- Replacing the 0 with a 1 will enable the accelerator usage.
- For the change to take affect, the partition needs to be IPL'ed.

35

## IBM Systems Energy Estimator

- <http://www-912.ibm.com/see/EnergyEstimator>



## Power saving mode

- POWER6 and POWER7 can go static Power saving mode – configure this in ASMI

```
lparstat -E 1 2
```

```
System configuration: type=Shared mode=Uncapped smt=On lcpu=4
mem=4096MB ent=0.40 Power=Static
```

Physical Processor Utilisation:

```
-----Actual----- -----Normalised-----
user sys wait idle freq user sys wait idle
----- ----- ----- ----- ----- ----- ----- ----- -----
0.002 0.005 0.000 0.393 3.6GHz [86%] 0.002 0.004 0.000 0.394
0.003 0.005 0.000 0.392 3.6GHz [86%] 0.003 0.004 0.000 0.393
or Power=Disabled
```

## Power saving mode

```
lparstat -h 1 2
```

```
System configuration: type=Shared mode=Uncapped smt=On lcpu=4 mem=4096MB
psize=4 ent=0.40
```

```
%user %sys %wait %idle physc %entc lbusy vcsw phint %hypv hcalls %nsp
----- ----- ----- ----- ----- ----- ----- ----- ----- -----
0.6 4.0 0.0 95.4 0.02 5.1 0.0 181 1 30.3 297 86
0.3 0.9 0.0 98.8 0.01 1.7 0.0 187 1 94.4 333 85
```

**Running at 86% clock speed**

## Power saving mode

- %nsp
- Indicates the current average processor speed as a percentage of nominal speed
- ASMI -> System Configuration -> Power Management Mode Setup Enable|Disable

## POWER7 Power saving

- How fast are my processors actually running?
- pmcycles
- pmcycles -m

## POWER7 Power saving

```
black1:~# pmcycles
This machine runs at 4004 MHz
black1:~#
black1:~# pmcycles -m
CPU 0 runs at 4004 MHz
CPU 1 runs at 4004 MHz
CPU 2 runs at 4004 MHz
CPU 3 runs at 4004 MHz
CPU 4 runs at 4004 MHz
CPU 5 runs at 4004 MHz
CPU 6 runs at 4004 MHz
CPU 7 runs at 4004 MHz
CPU 8 runs at 4004 MHz
CPU 9 runs at 4004 MHz
CPU 10 runs at 4004 MHz
CPU 11 runs at 4004 MHz
CPU 12 runs at 4004 MHz
CPU 13 runs at 4004 MHz
CPU 14 runs at 4004 MHz
CPU 15 runs at 4004 MHz
black1:~#
```

## How do I know how many Virtual Processors are active?

- It is a common question
  - There is no tool or metric that shows active Virtual Processor count
  - There are ways to guess this, and looking a physical consumption (if folding is activated), physc count should roughly equal active VPs
  - nmon Analyser makes a somewhat accurate representation, but over long intervals (with a default of 5 minutes), it does not provide much resolution
  - For an idea at a given instant, you can use: `echo vpm | kdb`



Be VERY careful with kdb



# Network

## How fast

- To find out the network bandwidth start an FTP session
- then do:

```
put "|dd if=/dev/zero bs=32k count=10000" /dev/null
```

- should get throughput which is almost 95% of the wire-speed
- This works OK for 1Gb/s
- For vEthernet and 10Gb/s FTP may have too much overhead to give valid data

## RMC

- On AIX
- To stop the daemons:  
`/usr/sbin/rsct/bin/rmcctrl -z`
- To start the daemons:  
`/usr/sbin/rsct/bin/rmcctrl -A`
- To enable the daemons for remote client connections (HMC to LPAR and vice versa):  
`/usr/sbin/rsct/bin/rmcctrl -p`

- On HMC

```
hmcshutdown -rt now
```

## RMC

- The whole system can be reinitialized.
- Stops RMC processes and resource managers.  
All ACL files deleted.  
New files are copied from templates.
- Fastest way to fix a broken configuration or to clear out files after cloning.

```
/usr/sbin/rsct/install/bin/recfgct
```

- Also useful if you disconnect an HMC and connect a new one





49



## WLM (Workload Manager)

- A great tip from Peter Nutt (via Andy Thomas)
- Set up WLM on AIX with default in tier 1 and System in tier 0
- If user process runs away you still get a response on the root prompt to sort it out.
- Also, a very handy WLM command is
- ```
# wlmassign myclass $$
```
- Assigns my process and all my children to "myclass" - handy to put this in the first line of a startup script for job you want controlled by WLM

50

© 2013 IBM Corporation

smitty mktcpip

```

Minimum Configuration & Startup

To Delete existing configuration data, please use Further Configuration menus

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

[Entry Fields]
* HOSTNAME [plum-gaz.aixncc.uk.ibm.com]
* Internet ADDRESS (dotted decimal) [9.137.62.161]
  Network MASK (dotted decimal) [255.255.255.0]
* Network INTERFACE en1
  NAMESERVER
    Internet ADDRESS (dotted decimal) [9.137.62.2]
    DOMAIN Name [aixncc.uk.ibm.com]
  Default Gateway
    Address (dotted decimal or symbolic name) [9.137.62.1]
    Cost [0] #
    Do Active Dead Gateway Detection? no +
  Your CABLE Type N/A +
  START Now no +

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

51

© 2013 IBM Corporation

HOSTNAME means FQDN

```

Minimum Configuration & Startup

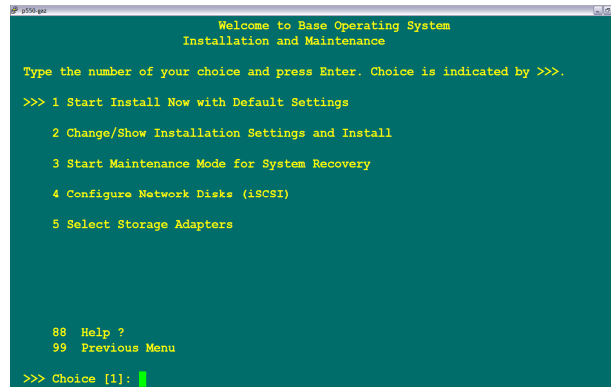
-----+-----
                CONTEXTUAL HELP
Type or s| Press Enter or Cancel to return to the application.
Press Ent|
| [TOP]
* HOSTNAM| Specify the host name you want to assign to the current machine.
* Interne| Enter the host name in the following format:
Network|
Network| host name
* Network|
NAMESER|
| When specifying the host name, use ASCII characters, preferably
Default| alpha-numeric only. Do not use a period in the host name. Avoid
Ad| using hexadecimal or decimal values as the first character (for
Co| example "3Comm," where "3C" might be interpreted as a hexadecimal
Do| character). The unqualified host name should be less than 32
| characters, for compatibility with older hosts. #
Your CA|
START N| #
| #
| [MORE...22]
F1=Help | F1=Help      F2=Refresh   F3=Cancel
F5=Reset | F8=Image     Enter=Do
F9=Shell |
-----+-----
  
```

52

© 2013 IBM Corporation

911 install

- Boot from Media/NIM etc as normal
- Select console
- Select Language



```
p550-gaz
Welcome to Base Operating System
Installation and Maintenance

Type the number of your choice and press Enter. Choice is indicated by >>>.

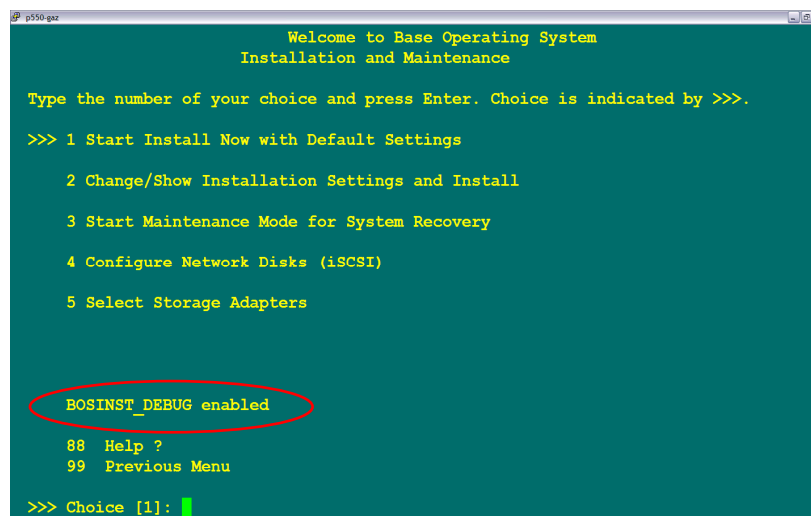
>>> 1 Start Install Now with Default Settings
    2 Change/Show Installation Settings and Install
    3 Start Maintenance Mode for System Recovery
    4 Configure Network Disks (iSCSI)
    5 Select Storage Adapters

88 Help ?
99 Previous Menu

>>> Choice [1]:
```

- Enter 911 at this screen. This will enable the install debugging for this session. The screen will refresh, and continue with option 1.

911 Install



```
p550-gaz
Welcome to Base Operating System
Installation and Maintenance

Type the number of your choice and press Enter. Choice is indicated by >>>.

>>> 1 Start Install Now with Default Settings
    2 Change/Show Installation Settings and Install
    3 Start Maintenance Mode for System Recovery
    4 Configure Network Disks (iSCSI)
    5 Select Storage Adapters

BOSINST_DEBUG enabled

88 Help ?
99 Previous Menu

>>> Choice [1]:
```

911 install

```
p550-gaz
+ ../usr/lib/methods/showled 0xA52
+ Log Initialize_Disk_Environment
+ dspmsg ../usr/lib/nls/msg/C/BosMenus.cat -s 10 37 Initializing disk environme
nt.\n
+ Change_Status 0 Initializing disk environment.
Initializing disk environment.
+ LIBPATH=/SPOT/usr/ccs/lib:/SPOT/usr/lib:/mnt/usr/ccs/lib:/mnt/usr/lib:/SPOT/us
r/lpp/xlC/lib:
+ PATH=/SPOT/usr/bin:/SPOT/usr/sbin:/SPOT/usr/lpp/bosinst:/mnt/usr/bin:/mnt/usr/
sbin:/mnt/usr/lpp/bosinst:
+ export PATH LIBPATH
+ [ -n ]
+ + ../SPOT/usr/lpp/bosinst/bidata -i -g image data -f PRODUCT TAPE
PT=yes
+ [ yes = no ]
+ [ yes = no ]
+ [ yes = no ]
+ [ yes = yes ]
+ ETC=usr/lpp/bos/inst_root/etc
+ DEV=usr/lpp/bos/inst_root/dev
+ VAR=usr/lpp/bos/inst_root/var
+ [ -n ]
+ ln -sf /mnt/usr/bin/pwd /usr/bin/pwd
+ 1> /dev/null 2>& 1
```

- You will get more debug than you could possibly want
- You will be able to get the text from your PuTTY log ☺

dfvg

```
p550-gaz:/usr/local/bin# cat dfvg
#
# a noddy script to do dfs for each volume groups
#
lsvg | while read VG
do
echo "disk usage in $VG filesystems: "
lsvg -l $VG | awk ' !/N\|A\|LV\|$VG/ { print $7 }' | xargs df -gP
echo
done
```

dfvg

```
p550-gaz:/usr/local/bin# time dfvg
```

```
disk usage in rootvg filesystems:
```

Filesystem	GB blocks	Used	Available	Capacity	Mounted on
/dev/hd4	19.75	16.98	2.77	86%	/
/dev/hd2	3.25	1.47	1.78	46%	/usr
/dev/hd9var	1.00	0.24	0.76	24%	/var
/dev/hd3	6.00	4.05	1.95	68%	/tmp
/dev/hd1	24.00	16.07	7.93	67%	/home
/dev/hd10opt	3.00	0.31	2.69	11%	/opt
/dev/hd11admin	1.00	0.00	1.00	1%	/admin
/dev/livedump	0.25	0.00	0.25	1%	/var/adm/ras/livedump
/dev/lv00	1.00	0.03	0.97	4%	/export/installios

```
disk usage in nimvg filesystems:
```

Filesystem	GB blocks	Used	Available	Capacity	Mounted on
/dev/export	367.12	357.17	9.96	98%	/export
nfs:/export	187.50	185.65	1.85	100%	/nfs/export

```
real    0m0.10s
```

```
user    0m0.05s
```

```
sys     0m0.01s
```

```
p550-gaz:/usr/local/bin#
```

© 2013 IBM Corporation

duvg

```
p550-gaz:/usr/local/bin# cat duvg
```

```
#
```

```
# a noddly script to estimate disk usage in filesystems in all volume groups
```

```
#
```

```
lsvg | while read VG
```

```
do
```

```
GIGS=$(lsvg -l $VG | awk ' !/N/A|LV|$VG/ { print $7 }' | xargs df -gP | sumcol 3)
```

```
echo "
```

```
    Estimated disk usage in $VG filesystems is $GIGS GB
```

```
"
```

```
done
```

58

© 2013 IBM Corporation

duvg

```
p550-gaz:/usr/local/bin# time duvg
```

```
Estimated disk usage in rootvg filesystems is 39.15 GB
```

```
Estimated disk usage in nimvg filesystems is 542.82 GB
```

```
real    0m0.14s
user    0m0.04s
sys     0m0.02s
```

sumcol

```
# SumCol - sum of a column
#
#SYNOPSIS
# SumCol N <data>
#
#DESCRIPTION
# Add up column N1 of stdin.  If you pipe the output of "ls -l" to
# this, it will add up the sizes, which are in column 4, our
# default.  You can specify the column for other kinds of files.
#
#BUGS
# There's a better version available in perl; awk is obsolete.
#
#AUTHOR
# John Chambers <jc@trillian.mit.edu>

if [ $# -lt 1 ];then set 4;fi
awk '{s += '$1'} END {print s}'
```

Mounting an ISO image in AIX

- Have you ever used loop devices eg: in Linux?
- Do you wish you could do the same in AIX?

▪ Well, why don't you?

Mounting an ISO image in AIX

- Firstly, we can see that nothing is mounted on /mnt

```
/# mount | grep mnt
```

- Then we mount an ISO image

```
/# loopmount -i /tmp/AIX616_dvd1.iso -m /mnt -o "-V cdrfs -o ro"
```

- And we can now see that it is mounted

```
/# mount | grep mnt
/dev/loop0          /mnt                cdrfs  02 Nov 17:54 ro
/# ls /mnt
6100-06      image.data  OSLEVEL  README.aix  RPMS  .Version
bosinst.data installp    ppc      root        usr
```

- Then we can unmount it

```
/# loopumount -l loop0 -m /mnt
/# ls /mnt
/# mount | grep mnt
/# ls mnt
/#
```

Mounting an ISO image in AIX

- The commands are in the devices.loopback.rte fileset

```

/# type loopmount
loopmount is /usr/sbin/loopmount
/# lslpp -w /usr/sbin/loopmount
  File                                Fileset                                Type
  -----
  /usr/sbin/loopmount                 devices.loopback.rte                   File
/# lslpp -w /usr/sbin/loopumount
  File                                Fileset                                Type
  -----
  /usr/sbin/loopumount                devices.loopback.rte                   File
/#

```

Mounting an ISO image in AIX

- So, with loopmount and virtual optical devices,
- You hardly ever need to write media and can do installations, upgrades etc remotely.

Linux (and AIX)

65

rpm

- How to install a new package, or upgrade it if it already exists?

```
rpm -Uvh <package.rpm>
```

- If you have a directory containing updated packages, and you wish to upgrade only the packages currently on your system, run:

```
rpm -Fvh *.rpm
```

- To delete an installed package, run:

```
rpm -e <package>
```

- To list all installed packages in your system, run:

```
rpm -qa
```

66

rpm

- You have installed a package, but don't know what it does, run the following command to find out more about the package:

```
rpm -qi <package>
```

- To list files installed by a package, run:

```
rpm -ql <package>
```

- How to get a-yet-to-be-installed package to give you information about itself and the files it would install:

```
rpm -qilp <package.rpm>
```

- To find out what package a particular file/binary is part of, run:

```
rpm -qf <full path to file and filename>
```

67

rpm

- To find out which files are missing from the RPM database, run:

```
rpm -Va
```

- To verify a package, run:

```
rpm -Vv <package>
```

Thanks to Aniket Patel
Technical Computing
(HPC) Services Europe,
CEE & MEA
Ex: Sequent Support

- To compress and rebuild the RPM database, run:

```
rpm --rebuilddb
```

68

Architecting

CHARM architecting

- System:

- Power 770 with one 5802 I/O drawer
 - CEC
 - Split backplane
 - Full of disks
 - 5802
 - also full of disks
 - Lots of adapters



CHARM architecting

- Should we use dual VIO servers?




CHARM architecting

- Derrr – YES!

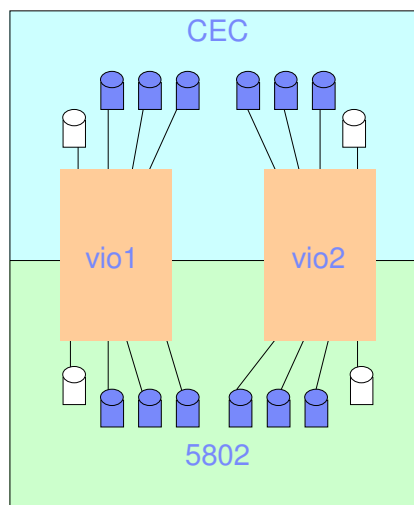
CHARM architecting

- Derrr – YES!
- But where should they be?
- How about:
 - Each VIO server has mirrored rootvg
 - One disk in the CEC and one in the 5802
- Multiple client LPARs
 - get hdisk0 from vio1
 - get hdisk1 from vio2
 - Mirror rootvg




CHARM architecting

- This is good
- We are protected against many types of failure
 - VIOS rootvg disk
 - VIOS rootvg adapter
 - VIOS clientsvg disk
 - VIOS clientsvg adapter
 - 5802 drawer, GX adapter etc



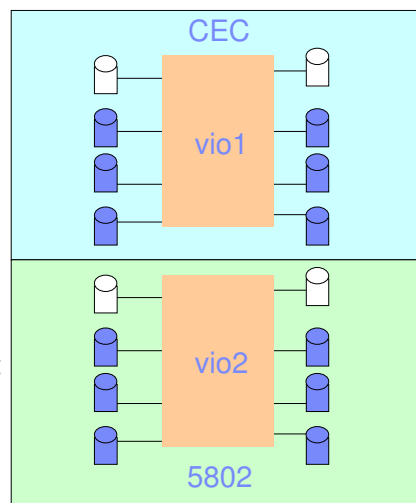
CHARM architecting

- But
- If we wanted to do CHARM on the 5802
- We would need to do a lot of AIX admin
 - Break the mirrors
 - rmdev the adapters and disks
 - etc
 - In BOTH VIOS
 - And ALL the CLIENTS
 - OUCH that will take ages



CHARM architecting

- How about this configuration?
- We still have a very good degree of protection
- If we wanted to do CHARM on the 5802
- We could shutdown VIO2
 - All the client LPARs get stale mirrors
 - Carry out the CHARM action
 - Reboot VIO2
 - Remirror the client LPARs
 - varyonvg
- That would take far less admin work at the time of the CHARM
 - And is less prone to error
- We did this in the European ATS
 - CHARM – GX Adapter replacement
 - **It worked perfectly**



CHARM architecting

- It all comes down to planning



netstat

- On the command line, type in Netstat <CR>
 1. Work with TCP/IP interface status
 2. Display TCP/IP route information
 3. Work with TCP/IP connection status
 4. Work with IPv6 interface status
 5. Display IPv6 route information
 6. Work with IPv6 connection status
- If TCP/IP is up and running, you will get a menu of options.
- If netstat command fails,
 - TCP/IP isn't running
 - need to do a strtcp <CR>, get a message saying tcp/ip starting.

options 1, 2 and 3
are the ones used
most

79

netstat

- The "trick" is everything you ever wanted to know about tcp/ip on an os/400 is here.
- Another "trick", some applications and firewalls use tcp/ip port numbers, how do you now which ports are active?
- so , netstat <enter> , option 3, shows local port as a name, then F24 for more keys, the F14 display port numbers, if the port number you want isn't listed then its not started.

Alun P Davies/UK/IBM


80

Scripting

81

Scripting

Hands up if you write shell scripts



82

JPG2jpg_all

```
plum-gaz:/tmp/demo# ls -l
total 0
-rw-r--r--  1 root  system  0 21 Oct 04:12 a.JPG
-rw-r--r--  1 root  system  0 21 Oct 04:12 A.JPG
-rw-r--r--  1 root  system  0 21 Oct 04:12 b.JPG
-rw-r--r--  1 root  system  0 21 Oct 04:12 B.JPG
-rw-r--r--  1 root  system  0 21 Oct 04:12 c.JPG
-rw-r--r--  1 root  system  0 21 Oct 04:12 C.JPG
plum-gaz:/tmp/demo# JPG2jpg_all
plum-gaz:/tmp/demo# ls -l
total 0
-rw-r--r--  1 root  system  0 21 Oct 04:12 a.jpg
-rw-r--r--  1 root  system  0 21 Oct 04:12 A.jpg
-rw-r--r--  1 root  system  0 21 Oct 04:12 b.jpg
-rw-r--r--  1 root  system  0 21 Oct 04:12 B.jpg
-rw-r--r--  1 root  system  0 21 Oct 04:12 c.jpg
-rw-r--r--  1 root  system  0 21 Oct 04:12 C.jpg
plum-gaz:/tmp/demo#
```

83

JPG2jpg_all

```
for i in $(ls | grep .JPG$ )
do
mv $i ${i%.JPG}.jpg
Done
```

84

to remove a suffix

```
tom:~ # NAME=name.ext
tom:~ # echo $NAME
name.ext
tom:~ # echo ${NAME%.ext}
name
tom:~ #
```

If the pattern matches a trailing portion of the expanded value of parameter, then the result of the expansion is the expanded value of parameter with

- the shortest matching pattern (the ```%`` case)
 - the longest matching pattern (the ```%%`` case)
- deleted.

85

to strip a prefix

```
tom:~ # NAME=name.ext
tom:~ # echo $NAME
name.ext
tom:~ # echo ${NAME#name.}
ext
tom:~ #
```

If the pattern matches the beginning of the value of parameter, then the result of the expansion is the expanded value of parameter with

- the shortest matching pattern (the ```#`` case) or
 - the longest matching pattern (the ```##`` case)
- deleted.

86

Do you use arrays?

Korn shell arrays in AIX

```
set -A MACHINES $(ssh hscroot@${HMC} ' lssyscfg -r sys -F name ' )
```

It really is that easy

Korn shell arrays in Linux (Suse)

```
MACHINES=$(ssh hscroot@${HMC} ' lssyscfg -r sys -F name ' )
```

That is even easier

```
tom:~ # bash --version
GNU bash, version 3.1.17(1)-release (i586-suse-linux)
Copyright (C) 2005 Free Software Foundation, Inc.
tom:~ #
```

Korn shell arrays in AIX

Using the array

```
ELEMENT=0
while [ ${ELEMENT} -lt ${#MACHINES[*]} ]
do
    echo ${ELEMENT} \t ${MACHINES[${ELEMENT}]}
    ELEMENT=$(( ELEMENT + 1 ))
Done

0      bronze-8203-E4A-SN10E0A21
1      172.17.254.248
2      pink-9115-505-SN65080EA
3      silver-8203-SN10E0A31
4      172.17.255.246
5      purple-9117-MMB-SN100525P
6      plum-8204-E8A-SN105C0B0
7      oldlace-65BD12E
8      orange-8203-E4A-SN10E0A51
9      brown-9115-505-SN6509E5A
10     grey-9117-MMA-p570-8F
```



So what?

POWER Server And LPAR Menu Version: 0.56beta

0	9.137.62.145	hmc10.aixncc.uk.ibm.com
1	9.137.62.146	hmc11.aixncc.uk.ibm.com
2	9.137.62.210	alex00.aixncc.uk.ibm.com
3	9.137.62.148	hmc8.aixncc.uk.ibm.com

Please select an HMC by number: 1

0	Select a Managed Server
1	Get information about the HMC

Please make a selection: 0

93

© 2013 IBM Corporation



So what ?

0	bronze-8203-E4A-SN10E0A21
1	172.17.254.248
2	pink-9115-505-SN65080EA
3	silver-8203-SN10E0A31
4	172.17.255.246
5	purple-9117-MMB-SN100525P
6	plum-8204-E8A-SN105C0B0
7	oldlace-65BD12E
8	orange-8203-E4A-SN10E0A51
9	brown-9115-505-SN6509E5A
10	grey-9117-MMA-p570-8F

Please select a machine by number: 5

94

© 2013 IBM Corporation

So what ?

```
0      Select an LPAR on this Managed Server
1      Get information about this Managed Server
```

Please make a selection: 0

0	purple11-AIX7sp1	21	aixlinux	Not_Activated
1	purple10_RH55	20	aixlinux	Not_Activated
2	purple9_fresh	19	aixlinux	Not_Activated
3	purple8_fresh	18	aixlinux	Not_Activated
4	purple7-AIX7_TL1	17	aixlinux	Running
5	purple6	16	aixlinux	Not_Activated
6	purple5-AIX5	15	aixlinux	Not_Activated
7	purple4-nag-ISD63	14	aixlinux	Running
8	purple3-hpc	13	aixlinux	Not_Activated
9	purple2-nag-ISD63	12	aixlinux	Running
10	purple1-SRAD-test	11	aixlinux	Running
11	mmafull	6	aixlinux	Not_Activated
12	purplevio3	5	vioserver	Not_Activated
13	purplevio2	4	vioserver	Running
14	purplevio1	3	vioserver	Running
15	purple12_IBMi	24	os400	Not_Activated

Please select an LPAR by number: 4

You get the idea

So what ?

```
purple7-AIX7_TL1,17,aixlinux,Running
```

- 1) List resources
- 2) Activate Normal
- 3) Activate SMS
- 4) Open vterm
- 5) Close vterm
- 6) Shutdown immediate
- 7) Shutdown OS

Please make a selection:

You get the idea

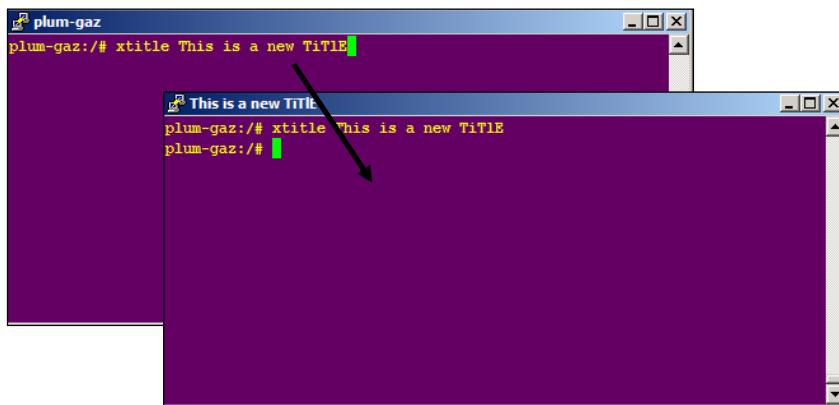
So, here is the *only* configuration file:

```
plum-gaz:/usr/local/lib# cat psalm
9.137.62.145    hmc10          # an HMC 7042-CR5
9.137.62.146    hmc11          # an HMC 7042-CR5
9.137.62.210    alex00         # Alex Abderrazag
9.137.62.148    hmc8           # an old HMC
9.137.62.12     hmc12          # an HMC 7042-CR7

plum-gaz:/usr/local/lib#
```

xtitle

- # Change title in an xterm window
- `exec echo "\033]1;${@\007\033]2;${@\007\c"`



Resources

99

The screenshot shows the website www.perzl.org with a directory listing for `ftp://www.oss4aix.org/latest/aix71/`. The listing includes various software packages such as `a2ps-4.14-1.aix5.1.ppc.rpm`, `as2sec-0.7.4-1.aix5.1.ppc.rpm`, `as2sec-devel-0.7.4-1.aix5.1.ppc.rpm`, `aalb-1.4.0-0.1.rc5.aix5.1.ppc.rpm`, `aalb-devel-1.4.0-0.1.rc5.aix5.1.ppc.rpm`, `aalb-libs-1.4.0-0.1.rc5.aix5.1.ppc.rpm`, `adns-1.4-1.aix5.1.ppc.rpm`, `adns-devel-1.4-1.aix5.1.ppc.rpm`, `adns-progs-1.4-1.aix5.1.ppc.rpm`, `aget-0.4.1-1.aix5.1.ppc.rpm`, `agg-2.5-1.aix5.1.ppc.rpm`, `agg-devel-2.5-1.aix5.1.ppc.rpm`, `agrep-0.8.0-2.aix5.1.ppc.rpm`, `analog-6.0-1.aix5.1.ppc.rpm`, `ansifilter-1.7-1.aix5.1.ppc.rpm`, `apachetop-0.12.6-1.aix5.1.ppc.rpm`, `apr-1.4.8-1.aix5.2.ppc.rpm`, `apr-devel-1.4.8-1.aix5.2.ppc.rpm`, `apr-util-1.5.2-1.aix5.1.ppc.rpm`, `apr-util-db4-1.5.2-1.aix5.1.ppc.rpm`, `apr-util-devel-1.5.2-1.aix5.1.ppc.rpm`, `apr-util-freetds-1.5.2-1.aix5.1.ppc.rpm`, `apr-util-gdbm-1.5.2-1.aix5.1.ppc.rpm`, `apr-util-ldap-1.5.2-1.aix5.1.ppc.rpm`, `apr-util-odbc-1.5.2-1.aix5.1.ppc.rpm`, `apr-util-sqlite-1.5.2-1.aix5.1.ppc.rpm`, `archimedes-2.0.0-1.aix5.1.ppc.rpm`, and `argtable2-13-1.aix5.1.ppc.rpm`.

On the right side of the screenshot, a browser window is open to `www.perzl.org`. The address bar shows `www.perzl.org` circled in red. Below the browser window, there is a section titled "Here you find the following stuff:" with two bullet points:

- [Open Source packages for AIX](#)
- [Ganglia packages for AIX and Linux on POWER](#)

Below the browser window, there is a table listing files with their sizes, dates, and times:

58 KB	17/09/2008	00:00:00
146 KB	17/09/2008	00:00:00
108 KB	19/05/2010	00:00:00
14 KB	19/05/2010	00:00:00
23 KB	19/05/2010	00:00:00
14 KB	26/03/2013	12:48:00
778 KB	10/03/2011	00:00:00
449 KB	10/03/2011	00:00:00
11 KB	03/09/2010	00:00:00
668 KB	05/05/2011	00:00:00
207 KB	30/04/2013	07:17:00
33 KB	16/02/2011	00:00:00
265 KB	27/06/2013	11:57:00
490 KB	27/06/2013	11:57:00
227 KB	18/06/2013	21:14:00
12 KB	18/06/2013	21:14:00
260 KB	18/06/2013	21:14:00
32 KB	18/06/2013	21:14:00
12 KB	18/06/2013	21:14:00
15 KB	18/06/2013	21:14:00
59 KB	18/06/2013	21:14:00
27 KB	18/06/2013	21:14:00
92 KB	08/01/2013	00:00:00
2980 KB	09/08/2011	00:00:00

100

PowerVM: Technical Webinar Series

- <http://tinyurl.com/newUK-PowerVM-VUG>
- **A series of technical webinars on Power Systems Virtualisation:**
 - Informal
 - Usually about an hour each
 - Many include demos
- **Aimed at:**
 - Technical audience - operators, systems administrators and technical specialists
 - Those using / planning to use IBM's Power based systems
 - Customers / Business Partners / IBMers
 - UK & Ireland audience primarily, although word spreading to other geographies

101

PowerVM: Technical Webinar Series

- <http://tinyurl.com/newUK-PowerVM-VUG>
- **Sessions:**
 - #1: Exploiting Virtualisation on IBM Power Systems with PowerVM
 - #2: VIOS - how to get going
 - #3: Controlling processor resources in virtualised partitions
 - #4: Deeper dive into Active Memory Sharing
 - #5: Virtualisation Best Practices
 - #6: Deeper dive into Shared Storage Pools and Thin Provisioning
 - #7: Virtualisation and the world of 10Gbit Ethernet
 - #8: Deeper dive into Active Memory Expansion
 - #9: VIOS maintenance
 - #10: Active Memory Deduplication
 - #11: IBM i Suspend/Resume
 - #12: Dual VIOS Upgrade walk-through
 - #13: Shared Storage Pools ... from Experience
 - #14: IBM i Live Partition Mobility
 - #15: Capture and Deploy partitions
 - #16: Virtual Partition Manager for IBM i
 - #17&18: Updating Power Systems, I/O and HMC
 - #19&20: Power7 Affinity and Performance
 - #21: Power Advisors (VIOS, LPAR and Java)
 - #22: PowerSC – What/Why/How
 - #23: Dynamic Platform Optimizer (DPO)
 - And more....

*Join the 450+ who have registered already...
Google 'PowerVM technical webinars'
and register*



102

PowerVM: Technical Webinar Series

- <http://tinyurl.com/newUK-PowerVM-VUG>
- **recent 2013 Sessions**
- **Session 25: Electronic Service Agent**
 - May 22nd 2013, 10:00 - 11:00 (UK time). View the time in your time zone, and optionally add to your calendar --> [Event in my timezone](#)
 - Speaker: Steve Harnett, Client Technical Specialist, Power Systems, IBM Systems & Technology Group, UK & Ireland.
 - What is it, and why use it? The session will include demos.
 - Relevant for AIX/IBM i/Linux environments.
- **Session 26: Active and Dynamic Systems Optimizer (ASO/DSO)**
 - June 12th 2013, 10:00 - 11:00 (UK time). View the time in your time zone, and optionally add to your calendar --> [Event in my timezone](#)
 - Speaker: Nigel Griffiths - Advanced Technology Support (ATS), IBM Power Systems, Europe.
 - Nigel looks at the new Active System Optimizer (ASO) feature where we can use this “expert system” to autonomously & dynamically tune AIX 7 on POWER7. Then looks at the new ASO extension called Dynamic Systems Optimizer (DSO). Both are 'set and forget' in operation. This is like having a Level 3 AIX Support performance guru tuning your systems all day!
 - Relevant for AIX.

103

PowerVM: Technical Webinar Series

- <http://tinyurl.com/newUK-PowerVM-VUG>
- **Jyoti Dodhia**
- **Email:** jyoti_dodhia@uk.ibm.com
- **Replays of previous sessions on the wiki**



104

Snippets

105

snippets

- largesend
 - admin: `$ chdev -dev en10 -attr mtu_bypass=on`
 - root: `# chdev -l en0 -a mtu_bypass=on`
- “PowerVM Virtualization Managing and Monitoring” RedBook (December 2012 update)

106

mount without logging (speed up a transfer)

- `mount -o log=NULL /xxxxx`

107

Twitter

- Follow these twitterers



[@power_gaz](#)
[@mr_nmon](#)
[@JyotiDodhia](#)
[@cgibbo](#)
[@NicoletteMcF](#)
[@robmcnelly](#)
[@ibmaix](#)
[@chmod666](#)
[@chromeaix](#)
[@IBMPureSystems](#)
[@IBMPowerSystems](#)
[@AIXUserGroup](#)

108


A couple of good videos

- One old



<http://t.co/cljsVrvjUp>

- One newer



<http://www.youtube.com/watch?v=1W1M114ppHQ>

109

Power Ask The Experts 2013



FLRT

- <http://www14.software.ibm.com/webapp/set2/flrt/>
- Google: FLRT IBM

ix Level Recommendation tool
For IBM Power Systems administrators

The following consolidated information is for guidance purposes only. This information was obtained from generally available product support documentation. These combinations of product levels are supported by IBM.

Date: 2011.04.05
Model: IBM Power 720 Express (8202-E4B)
Click [here](#) for the latest device firmware for this model.
Clock: 3.0 ghz.

Your selected levels		
Product	Version/Release	Status
AIX	7100-00-00	ⓘ
System Firmware	AL720_082	✓
Virtual I/O Server	2.2.0.10	⚠

Detailed results

ⓘ AIX
Note: You selected a level lower than the recommended level. The current recommended SP level for this AIX TL is 7100-00-02.

⚠ Virtual I/O Server [Set the upgrade](#)
Upgrade recommended: 2.2.0.11
VIOS 2.2.0.11-FP-24 SP-01 is the first Service Pack for VIOS 2.2.0.10-FP-24. It is an optional service pack. Do not apply this service pack to any other level of VIOS except Fix Pack 24 (VIOS 2.2.0.10-FP-24).

110

2013 IBM Corporation

Writing AIX DVDs on Linux

- A command to write DVDs on Linux:

```
cdrecord -dev=/dev/dvd -dao driveropts=burnfree file.iso
```
- How to use mkisofs so that AIX gets the filenames right:

```
mkisofs -o ./ISONAME.iso -L -pad -allow-multidot -rock \  
-untranslated-filenames -joliet -rational-rock -force-rr DIRECTORY
```

FEEDBACK

IS THIS KIND OF SESSION USEFUL?


If so, please send me your tips ...

I will always give **you** the credit 😊



Thank you for attending this session!

IBM Power Ask the Experts 2013

09:30 - 10:00	Registration and coffee
10:00 - 11:15	Power Systems Update - Pat O'Rourke: Austin Briefing Centre
11:15 - 12:30	Performance Best Practices with POWER7 - Nigel Griffiths
12:30 - 13:30	Lunch
13:30 - 14:30	Tricks of the Power Masters - Gareth Coates
 14:30 - 15:15	Cost Comparison between IBM Power and Intel - David Spurway
15:15 - 15:30	Coffee
15:30 - 16:45	Power Systems Trends and Directions - Pat O'Rourke: Austin Briefing Centre
16:45	Close