

# More Tricks of the Power Masters

Starting at 10:00 am UK time

Gareth Coates



Smart Meeting → Put questions into the Chat box  
or AT&T Toll Free phone for better audio

- 0800-368-0638 = UK Toll Free
- 0203-059-6451 = UK but you pay for the call
- Then 6403785# Participant Code
- Other countries see chat box for the website
- Please Mute with \*6



## Previous Sessions:

POWER8 from hands-on  
Power up your Linux  
PowerVC  
PowerVP  
SSP4  
Best Practices  
Tricks of Power Masters  
IBMi and External Storage  
Monitoring with ITM  
Whole Machine Monitoring  
Electric Server Agent  
RDX Removable disks  
And more.....

## Future Sessions →

- 16th July - PowerKVM Deep Dive
- Suggestions Welcome



Gareth Coates @power\_gaz  
Jyoti Dodhia @JyotiDodhia



Nigel Griffiths @mr\_nmon  
Mandie Quartly @mandieq

Website: <http://tinyurl.com/PowerSystemsTechnicalWebinars>  
Youtube Channel: <http://tinyurl.com/IBMPowerVUGYyoutubeChannel>

**Speaker Name** Gareth Coates  
**Job Role** Power - Advanced Technology Support  
**Company** IBM UK Limited

---



# Tricks of the Power Masters

(updated 10<sup>th</sup> June 2014)



# Session objectives

---

- 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
    - Power ATS Europe
  - Pat O'Rourke
    - Briefing Center (sic), Austin, TX
  - Jay Kruemcke
    - Mr AIX, Austin, TX
  - Chris Gibson
    - Australian AIX guru
  - And many others – too many to mention

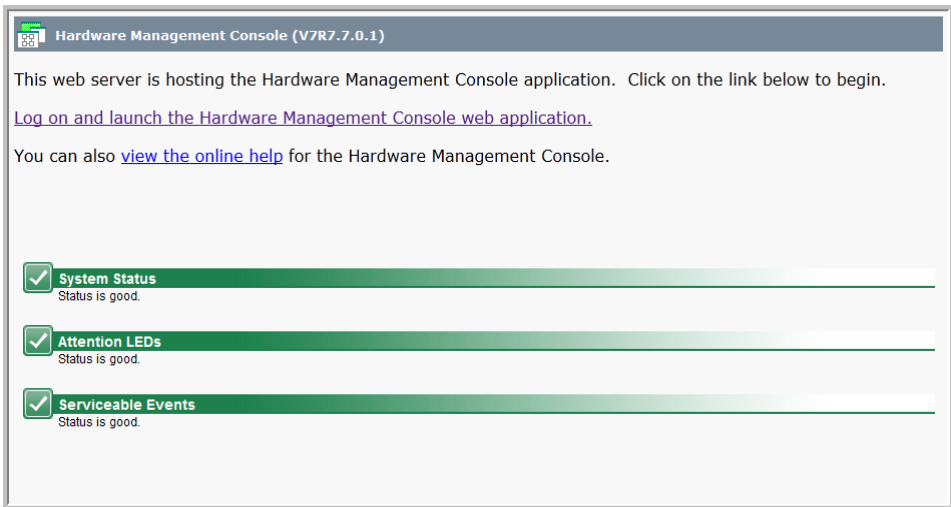
## Some pretty big ones

---

- You really should use these
  - Or at least consider them very seriously

NIM  
PowerVP  
PowerVC  
DPO  
LPM  
SSP

# HMC



Hardware Management Console (V7R7.7.0.1)

This web server is hosting the Hardware Management Console application. Click on the link below to begin.  
[Log on and launch the Hardware Management Console web application.](#)

You can also [view the online help](#) for the Hardware Management Console.

- System Status**  
Status is good.
- Attention LEDs**  
Status is good.
- Serviceable Events**  
Status is good.


# Refresh O/S level information on the HMC

---

- Force a refresh of the operating system level information on the HMC once AIX is upgraded on the partitions
- `lssyscfg -r lpar -m <managed system> --osrefresh`
- When listing partitions, specify this option to refresh the current operating system version information for the partitions first.
- If a partition does not have an active RMC connection to the management console, Unknown will be displayed for that partition's operating system version information.
- Specifying this option may cause this command to take a long time to complete if many partitions are being listed or there are network issues.

# Frame is locked

- Trying to do a firmware update on a 9119-FHB gave this:



Clive Benjamin  
Principal Engineer  
IBM UK Limited

- If, and only if you are sure ....
- see which hmc has the lock
- `lslock -e <frame_name>`
- to clear it :
- `rmlock -e <frame_name>`
- should then be OK

- Question:
- Why did this happen?
- Answer:
- An attempted upgrade had failed part way through



# HMC Upgrades

---

- Download code from FixCentral, burn some media etc etc
- OR
- Download Network install images from
  - <http://www14.software.ibm.com/webapp/set2/sas/f/netinstall/home.html>

File name	Description	Destination path	Size (bytes)	Checksum
<a href="#">bzImage</a>	Kernel image	/var/ftp/hmc	2730176	02364
<a href="#">initrd.gz</a>	Ram Disk file system	/var/ftp/hmc	34185788	06816
<a href="#">disk1.img</a>	Base Image	/home/hmc	817065984	55470
<a href="#">disk2.img</a>	Base HMC Image	/home/hmc	1456427008	33312
<a href="#">disk3.img</a>	Information Center Image	/home/hmc	873922560	45622
<a href="#">hmcnetworkfiles.sum</a>	hmcnetworkfiles.sum	/home/hmc	78 bytes	n/a

# HMC upgrades over the network

---

- Why bother?
- Network upgrades take far less time than using media
  - Assuming a half decent LAN eg: Gbit
- You do not need to be in the computer room

# HMC upgrades over the network

---


- Always read the instructions, but here is an overview:
- On the HMC, save upgrade:
  - **saveupgdata -r disk**
- Copy the files over to the bootable disk partition on the HMC.
  - **getupgfiles -h <hostname> -u <userid> -d <directory>**
- After the files are copied over, run the following command:  
**chhmc -c altdiskboot -s enable --mode upgrade**
- Reboot the HMC
  - **hmcshutdown -r -t now**

---

# VIOS

# HMC one liner to run a command on every VIOS

Brian Smith



```
• for sys in `lssyscfg -r sys -F name`; do for vio in  
  `lssyscfg -r lpar -m $sys -F name,lpar_env | grep  
  vioserver | cut -f 1 -d, | sort`; do echo $vio;  
  viosvr cmd -m $sys -p $vio -c errlog; done; done
```

• [https://www.ibm.com/developerworks/community/blogs/brian/entry/hmc\\_one\\_liner\\_to\\_run\\_a\\_command\\_on\\_every\\_vio\\_server?lang=en](https://www.ibm.com/developerworks/community/blogs/brian/entry/hmc_one_liner_to_run_a_command_on_every_vio_server?lang=en)

# viosbr

---

- A useful command to backup VIOS configuration
- Recommended for inclusion
  - by [Daniel.Loftus@mckesson.co.uk](mailto:Daniel.Loftus@mckesson.co.uk)
- Performs the operations for backing up the virtual and logical configuration, listing the configuration, and restoring the configuration of the Virtual I/O Server.
- The viosbr command can be run only by the padmin user.


# viosbr -backup -file viosbr.out

---

```
$ viosbr -backup -file viosbr.out
```

```
Backup of this node (purpleviol1.aixncc.uk.ibm.com) successful
```

```
$ viosbr -view -file viosbr.out.tar.gz
```



# viosbr -view -file viosbr.out.tar.gz

---

Controllers:

=====

Name	Phys Loc
----	-----
iscsi0	
sissas0	U5802.001.00H1960-P1-C1-T1
sissas1	U5802.001.00H1960-P1-C4-T1
pager0	U9117.MMB.100525P-V3-C32769-L0-L0
vasi0	U9117.MMB.100525P-V3-C32769
vbsd0	U9117.MMB.100525P-V3-C32769-L0
sata0	U5802.001.00H1960-P1-C1-T1
sata1	U5802.001.00H1960-P1-C4-T1
lhea0	
lhea1	U78C0.001.DBJ0440-P2
fcs0	U5802.001.00H1960-P1-C10-T1
fcs1	U5802.001.00H1960-P1-C10-T2
sfwcomm0	U5802.001.00H1960-P1-C10-T1-W0-L0
sfwcomm1	U5802.001.00H1960-P1-C10-T2-W0-L0
sfwcomm2	SAS
sfwcomm3	SAS
fscsi0	U5802.001.00H1960-P1-C10-T1
ent0	U78C0.001.DBJ0440-P2-C8-T2
fscsi1	U5802.001.00H1960-P1-C10-T2
ent1	U9117.MMB.100525P-V3-C44-T1
ent2	U9117.MMB.100525P-V3-C45-T1
sas0	U5802.001.00H1960-P1-C1-T1
sas1	U5802.001.00H1960-P1-C4-T1
sfw0	
fcnet0	U5802.001.00H1960-P1-C10-T1
fcnet1	U5802.001.00H1960-P1-C10-T2



# viosbr -view -file viosbr.out.tar.gz

---

## Physical Volumes:

```
=====
Name          Phys Loc
----          -
hdisk20       U5802.001.00H1960-P1-C10-T1-W500507680210AC8C-L2000000000000
hdisk21       U5802.001.00H1960-P1-C10-T1-W500507680210AC8C-L3000000000000
hdisk22       U5802.001.00H1960-P1-C10-T1-W500507680210AC8C-L4000000000000
hdisk23       U5802.001.00H1960-P1-C10-T1-W500507680210AC8C-L5000000000000
hdisk24       U5802.001.00H1960-P1-C10-T1-W500507680210AC8C-L6000000000000
hdisk0        U5802.001.00H1960-P3-D1
hdisk1        U5802.001.00H1960-P3-D2
hdisk2        U5802.001.00H1960-P3-D3
hdisk3        U5802.001.00H1960-P3-D4
hdisk4        U5802.001.00H1960-P3-D5
hdisk5        U5802.001.00H1960-P3-D6
hdisk6        U5802.001.00H1960-P3-D7
hdisk7        U5802.001.00H1960-P3-D8
hdisk8        U5802.001.00H1960-P3-D9
hdisk9        U5802.001.00H1960-P3-D10
hdisk10       U5802.001.00H1960-P3-D11
hdisk11       U5802.001.00H1960-P3-D12
hdisk12       U5802.001.00H1960-P3-D13
hdisk13       U5802.001.00H1960-P3-D14
hdisk14       U5802.001.00H1960-P3-D15
hdisk15       U5802.001.00H1960-P3-D16
hdisk16       U5802.001.00H1960-P3-D17
hdisk17       U5802.001.00H1960-P3-D18
hdisk18       U5802.001.00H1960-P1-C10-T1-W500507680210AC8C-L0
hdisk19       U5802.001.00H1960-P1-C10-T1-W500507680210AC8C-L1000000000000
```

## Optical Devices:

```
=====
Name          Phys Loc
----          -
```

## Tape Devices:

```
=====
Name          Phys Loc
----          -
```

```
viosbr -view -file viosbr.out.tar.gz
```

---

```
Ethernet Interfaces:
```

```
=====
```

```
Name
```

```
----
```

```
en0
```

```
en1
```

```
en2
```

```
en3
```

```
Storage Pools:
```

```
=====
```

```
SP Name
```

```
-----
```

```
rootvg
```

```
clientsvg
```

```
PV Name
```

```
-----
```

```
hdisk0
```

```
hdisk1
```

```
hdisk2
```

```
hdisk3
```

```
hdisk4
```

```
hdisk5
```

```
hdisk6
```

```
hdisk7
```

```
hdisk8
```

```
hdisk9
```

```
hdisk10
```

```
hdisk11
```

```
hdisk12
```

```
hdisk13
```

```
hdisk14
```

```
hdisk15
```

```
hdisk16
```

```
hdisk17
```

**viosbr -view -file viosbr.out.tar.gz**

---

Optical Repository:

=====

Name	Parent SP
----	-----
VMLibrary	clientsvg

Shared Ethernet Adapters:

=====

Name	Physical Adapter	Default Adapter	Virtual Adapters
----	-----	-----	-----
ent3	ent0	ent1	ent1

# viosbr -view -file viosbr.out.tar.gz

---

## Virtual Server Adapters:

=====

SVSA	Phys Loc	VTD
----	-----	---
vhost13	U9117.MMB.100525P-V3-C24	purple14_1.vtd
vhost14	U9117.MMB.100525P-V3-C25	
vhost15	U9117.MMB.100525P-V3-C26	
vhost16	U9117.MMB.100525P-V3-C32	vtscsi7 vtopt1
vhost0	U9117.MMB.100525P-V3-C11	vtscsi1 vtscsi0 vtscsi8 vtopt3 vtscsi4
vhost1	U9117.MMB.100525P-V3-C12	
vhost2	U9117.MMB.100525P-V3-C13	purple3_1.vtd
vhost3	U9117.MMB.100525P-V3-C14	vtopt0 vtscsi10 vtscsi9 purple4_1.vtd
vhost4	U9117.MMB.100525P-V3-C15	purple5_1.vtd
vhost5	U9117.MMB.100525P-V3-C16	purple6_1.vtd
vhost6	U9117.MMB.100525P-V3-C17	vtopt2 purple7_1.vtd
vhost7	U9117.MMB.100525P-V3-C18	vtopt4 vtscsi6 vtscsi5
vhost8	U9117.MMB.100525P-V3-C19	
vhost9	U9117.MMB.100525P-V3-C20	vtopt6 vtscsi11
vhost10	U9117.MMB.100525P-V3-C21	vtscsi12
vhost11	U9117.MMB.100525P-V3-C22	
vhost12	U9117.MMB.100525P-V3-C23	

```
viosbr -view -file viosbr.out.tar.gz
```

---

```
RSCT:
```

```
=====
```

```
ct_node_id
```

```
-----
```

```
9df352ea78a5d4e9
```

```
Cluster:
```

```
=====
```

```
Cluster          State
```

```
-----          ----
```

```
Virtual Log Repositories:
```

```
=====
```

```
Virtual Log Repository  State
```

```
-----                ----
```

```
vlogrepo0          AVAILABLE
```

```
$
```

# 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-vio1->vhost0
vscsi1      0x000007 0x000000000000 0x0 plum-vio1->vhost1
#
```

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



---

# 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,10000000c9b7cbf0,null
U78C0.001.DBJJ568-P2-C1-T2,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,10000000c9b7cbf1,null
U78C0.001.DBJJ568-P2-C3-T1,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,10000000c9c6aa32,null
U78C0.001.DBJJ568-P2-C3-T2,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,10000000c9c6aa33,null
U78C0.001.DBJJ568-P2-C4-T2,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,10000000c9b65cc3,null
U78C0.001.DBJJ568-P2-C4-T1,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,10000000c9b65cc2,null
U78C0.001.DBJJ568-P2-C6-T1,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,10000000c9b7ca38,null
U78C0.001.DBJJ568-P2-C6-T2,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,10000000c9b7ca39,null
U78C0.001.DBJJ675-P2-C1-T1,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,10000000c9b65c40,null
U78C0.001.DBJJ675-P2-C1-T2,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,10000000c9b65c41,null
U78C0.001.DBJJ675-P2-C3-T1,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,10000000c9c6ac0e,null
U78C0.001.DBJJ675-P2-C3-T2,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,10000000c9c6ac0f,null
U78C0.001.DBJJ675-P2-C4-T1,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,10000000c9b65f78,null
U78C0.001.DBJJ675-P2-C4-T2,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,10000000c9b65f79,null
U78C0.001.DBJJ675-P2-C6-T1,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,10000000c9b65e00,null
U78C0.001.DBJJ675-P2-C6-T2,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,10000
U78C0.001.DBJJ696-P2-C1-T1,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,10000
U78C0.001.DBJJ696-P2-C1-T2,8 Gigabit PCI-E Dual Port Fibre Channel Adapter,null,10000
```

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 (see the Redbook)
- 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:~ #
```

# LPM Checklists

---

- Here are some Checklists for setting up LPM
- **Live Partition Mobility Setup Checklist**

<http://www.redbooks.ibm.com/abstracts/tips1184.html?Open>

## **Live Partition Mobility Preparation Checklist**

<http://www.redbooks.ibm.com/abstracts/tips1185.html?Open>

# 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.

## 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`



# # echo vpm | kdb

---

(0)> vpm

## VPM Tunables

```
vpm_interval..... 000003E8  schedp.vpm_fold_threshold.. 00000050
vpm_fold_limit..... 00000008  vpm_unfold_limit..... 00000002
vpm_be_threshold..... 0000000A
```

## VPM Data

```
.....vpm_mgmt :01
.....vpm_flags :00
.....vpm_prev_time :1350612958965291578 ns
.....vpm_interval :1000636881 ns
.....vpm_prev_intr :6571448426312 ns
.....vpm_intr :1408406 ns
```

## VPM Core Buckets

INDX	ID	TYPE	SR	CORE	LAST	ACTI	HEAD	UTIL	STATE
0	0	SYST	1	4	2	2	4	1004877499	
1	4	6	0	0	0	GP			
2	2	6	0	0	0	GP ACTIVE			
3	0	2	0	0	0	GP ACTIVE			

...



```
# echo vpm | kdb
```

---



Be VERY careful with kdb

```
VPM Sleep Data.
```

```
.....vsd_state :01 SLEEP_INITIALIZED
.....vsd_sleep_specifier :00000000
.....vsd_sleep_latency :0000000000000000 sec 00000000 nsec
.....vsd_min_sleep :0000000000000000 sec 1DCD6500 nsec
.....vsd_prod_lock :@F1000A00002E4E00
.....vsd_prod_head :00000000024A9DD0
```

```
VSD Thread State.
```

CPU	VP_STATE	SLEEP_STATE	PROD_TIME: SECS	NSECS	CEDE_LAT
0	ACTIVE	AWAKE	0000000000000000	00000000	00
1	ACTIVE	AWAKE	0000000000000000	00000000	00
2	ACTIVE	AWAKE	0000000000000000	00000000	00
3	ACTIVE	AWAKE	00000000508069D0	26F6B8F6	00
4	DISABLED	SLEEPING	000000005080B7E1	0AD9D9B5	00
5	DISABLED	SLEEPING	000000005080B7E1	0ADAE51C	00
6	DISABLED	SLEEPING	000000005080B7E1	0ADC2A67	00
7	DISABLED	AWAKE	0000000000000000	00000000	00

```
(0)> plum-gaz:/#
```



---

# Resources

Index of ftp://www.oss4aix.org/latest/aix71/

Up to higher level directory

**Name**

- a2ps-4.14-1.aix5.1.ppc.rpm
- a52dec-0.7.4-1.aix5.1.ppc.rpm
- a52dec-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-lbs-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
- argtable2-13-1.aix5.1.ppc.rpm

My Open Source packages for AIX and Linux on POWER - Mozilla F

File Edit View History Bookmarks Tools Help

My Open Source packages for AIX and Linux... +

www.perzl.org

Most Visited Electrabel ibmtedu tiny URL stgdam2013 IBM W

Here you find the following stuff:

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

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



# From Andy Thomas (IBM UK)

---

- 64k page size
- If you are not doing AMS - use a 64k pagesize - can improve performance by 10%
- ```
# ldedit -btextpsize=4K -bdatapsize=64K -bstacksize=64K  
$ORACLE_HOME/bin/oracle
```
- Spikey workload on a system with lots of busy LPARs not getting resources
- Set the `vpm_xvcpus` to `-1` disables VCPU folding.
- (If Physical resources are not immediately available - then CPU's won't unfold by default - even if weighting is the same as other LPARs)
- If using AMS - set `maxclient%` and `maxperm%` to 50% or 60%
  - guarantee some "free memory" for AMS to work with
  - Good if LPAR has a 2GB+ EC and is not an NFS server or filesystem based database etc

# How to debug NIM problems

---

- START NIMESIS IN DEBUG MODE
  - To debug nimesis problems on the NIM Master do the following:

- 1. Stop nimesis if it still running:

```
stopsrc -s nimesis
```

- 2. Run:

```
export NIMBUG=/tmp/nimesis.log
```

Note: Choose another location for the debug output if you wish.

- 3. Run:

```
touch $NIMBUG
```

- 4. Run:

```
/usr/sbin/nimesis -ds
```

- NOTE: Please make sure ou keep an eye on /tmp/nimesis.out. The /tmp filesystem can get filled up.

# How to debug NIM problems

---

- LOOKING AT NIMLOG
- The "nimlog" file lives in `/var/adm/ras`.
- nimlog is an alog file. Do the following to make the contents ASCII readable, run:

1. `cd /var/adm/ras`

2. `alog -f nimlog -o > /tmp/nimlog.out`

3. View `/tmp/nimlog.out` with your favourite editor

# How to debug NIM problems

---

- BOOTP PROBLEMS - Make sure you check the following:

- Check if the bootpd is active on the master:

```
#lssrc -t bootps
```

- Check if the bootp entry for the client is correct.

```
#more /etc/bootptab
```

- If you are doing a pull, make sure you enter exactly what is in the bootp entry.

- Kill any run-away bootp processes and start again.

- Make sure that your networking makes sense:

- gateway is on same network as client
- if there is a gateway make sure the client and the master are not on the same network
- if the client and master are on the same net, make sure the gateway is
  - 0.0.0.0
  - or the ip address of the master.

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

# Bootp debugging

---

- ENABLING BOOTP DEBUGGING

1. `vi /etc/inetd.conf`

`# Comment out the following line:`

```
bootps dgram udp wait root /usr/sbin/bootpd bootpd /etc/bootptab
```

`# Save changes and exit.`

2. `refresh -s inetd`

3. `ps -ef | grep bootpd`

`# If it is running, kill the process with a "-15".`

`# If it is still running, kill it with a "-9".`

4. `bootpd -d -d -d -d -s > /tmp/bootpd.out 2>&1 &`

`# The output will be logged to /tmp/bootpd.out.`



# Bootp debugging

---

## DISABLING BOOTP DEBUGGING

1. `ps -ef | grep bootpd`

# Kill the debug bootpd process with a "-15".

# If it is still running, kill it with a "-9".

2. `vi /etc/inetd.conf`

# Uncomment out the following line:

```
bootps dgram udp wait root /usr/sbin/bootpd bootpd /etc/bootptab
```

# Save changes and exit.

3. `refresh -s inetd`

# bootpd will not show up in the process table until a bootp request is made.

# When AIX is out of memory it kills processes

- You can protect things if you want..

– ( AIX OOM ) <- out of memory

```
# grep ssh /etc/passwd
```

```
sshd:*:202:201::/var/empty:/usr/bin/ksh
```

```
# vmo -o nokilluid=202
```

From the man page:

Thanks to **Keld Norman**

Senior IT Specialist,  
UNIX (AIX), Linux and  
Windows

IBM Denmark

User IDs lower than this value will be exempt from getting killed due to low page-space conditions.

Values:

Default: 0

Range: 0 - 4294967295

Type: Dynamic

Unit: uid

Tuning:

A value of 0 indicates off. Useful when system is out of paging space and system administrator's processes are getting killed. Either set this tunable to 1 in order to protect specific user ID processes from getting killed due to low page space or ensure there is sufficient paging space available.

# Test RSCT connection

---

- # `/usr/sbin/rsct/bin/rmcdomainstatus -s ctrmc`

Management Domain Status: Management Control Points

I A 0xdac61f3c784c692b 0001 9.137.62.146

I UP

O Down

Thanks to **Keld Norman**

Senior IT Specialist,  
UNIX (AIX), Linux and  
Windows

IBM Denmark

# Putty and the shell

---

- The following instructions describe how to you can automatically configure your home area environment on the Unix server so that the hostname is automatically displayed in your PuTTY terminal window.

Insert the following lines into your `.profile`:

```
# Display hostname as PuTTY window title
echo "\033]0;`hostname`\007"
```

```
# Reset PuTTY window title when logging off
trap "$HOME/.logout" EXIT
```

Create a file called `logout` in your home directory. Add the following lines:

```
tsmlpar:/home/langc > cat .logout
BLANK=""
export BLANK
echo "\033]0;$BLANK\007"
tsmlpar:/home/langc >
```

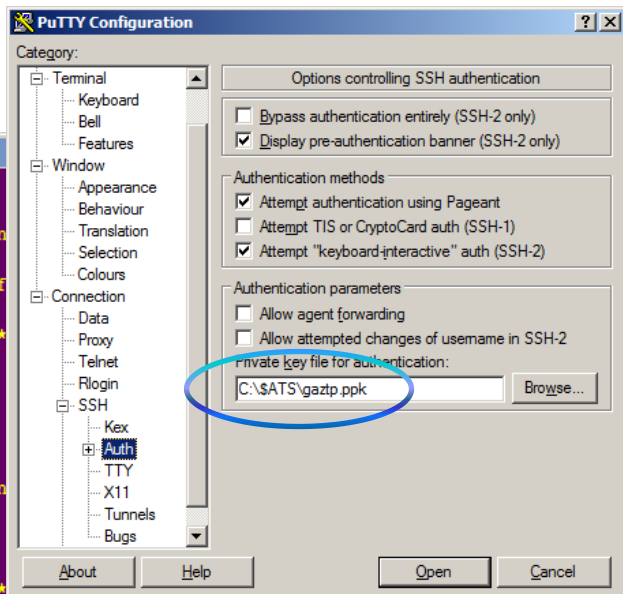
Chris Lang

Celerity UK

Make the `logout` file executable.

# Putty with ssh keys

```
PEACH-GAZ
Using username "root".
Authenticating with public key "rsa-key-20130507"
Last unsuccessful login: Wed  4 Jun 15:30:04 2014 on
m.com
Last login: Tue 10 Jun 06:27:07 2014 on /dev/pts/0 f
.ibm.com
*****
*
* Welcome to AIX Version 7.1!
*
* Please see the README file in /usr/lpp/bos for in
* this release of the AIX Operating System.
*
*****
peach-gaz:/#
```



# Parallel operations using dsh, distributed shell

---

- dsh is available on AIX distribution (dsm filesets)
- dsh requires that commands run without prompting for password
- dshbak sorts output by target partitions

## Setup

```
# DSH_LIST=<file listing target partitions>
# DSH_REMOTE_CMD=/usr/bin/ssh           #If using ssh
# DSH_PATH=$PATH:/usr/local/bin        #setting paths
```

Morten Vagmo  
IBM Norway

# Parallel operations using dsh, distributed shell

---

- `>: dsh date`

```
nimserver.oslo.forum.ibm.com: Mon Oct 14 22:33:06 CEST 2013
520_deneb.oslo.forum.ibm.com: Mon Oct 14 22:33:06 CEST 2013
stg_mortvag.oslo.forum.ibm.com: Mon Oct 14 22:33:06 CEST 2013
520_altair.oslo.forum.ibm.com: Mon Oct 14 22:33:06 CEST 2013
stg_gold.oslo.forum.ibm.com: Mon Oct 14 22:33:06 CEST 2013
510_vega.oslo.forum.ibm.com: Mon Oct 14 22:33:06 CEST 2013
```

- `>: dping -f /home/root/nodes`

```
stg_gold.oslo.forum.ibm.com: ping (alive)
510_vega.oslo.forum.ibm.com: ping (alive)
520_altair.oslo.forum.ibm.com: ping (alive)
nimserver.oslo.forum.ibm.com: ping (alive)
stg_mortvag.oslo.forum.ibm.com: ping (alive)
520_deneb.oslo.forum.ibm.com: ping (alive)
```


# Parallel operations using dcp, distributed copy

- Using ssh authentication and scp remote command
  - # DCP\_NODE\_RCP=/usr/bin/scp
  - (# alias dcp='dcp --node-rcp /usr/bin/scp')
- Example
  - # DSH\_LIST=/home/root/nodes
  - # dcp /etc/hosts /etc

Chapter 5.2 in AIX7 Diff. Guide:  
Distributed System Management


- dsm.core, dsm.dsh filesets
  - dpasswd
  - dconsole
    - dsh
    - dcp

**New from AIX 6.1 TL3**








# Setting up dsh



👉 Install ssh on AIX  
👉 ssh default on vios



1



id\_dsa.pub -> authorized\_keys2

2



3

dsh date



# Dawn May's iCan blog

- [http://ibmsystemsmag.blogs.com/i\\_can/](http://ibmsystemsmag.blogs.com/i_can/)
- Regular good tips

The screenshot shows the IBM Systems Magazine website. At the top, the logo reads "IBM Systems MAGAZINE". Below the logo are navigation tabs for "AIX", "MAINFRAME", and "POWER". A yellow arrow points to the "IBM i" tab. Other tabs include "ADMINISTRATOR", "DEVELOPER", "TRENDS", "TIPS & TECHNIQUES", and "CASE STUDIES". The page is titled "Blog" and features a blue banner with the text "i Can Technical Tips for i". The main article is titled "IBM i Storage Allocation Perspectives" and is dated "March 26, 2014". The author's name "Dawn May" is displayed next to her profile picture. The article text begins with "Quite some time ago, I wrote a blog on Understanding Disk Space Usage. Near the end of that article, I mentioned that Collection Services collects storage usage data in the QAPMJOBM file." and continues with "Late last year, there were additional enhancements to Navigator for i, which included updates to the Performance Data Investigator (PDI). With the Fall 2013 PDI enhancements, you can now visualize these Collection Services storage usage metrics. You will need to install the required PTFs to enable this support, as documented in the 'Fall 2013' section of the Performance on the Web information." and ends with "You will now find some new charts that will show the storage usage on your system over time." Below the article are social media sharing buttons for "Tweet" and "SHARE", and a "Recent Comments" section with a comment from "Dawn on POWER and PowerVM Reference Information" and another from "bryan dietz on CPF1240 and 341 M...".

Jyoti Dodhia

IBM UK

And, of course, Dawn

# The last IPL on an IBM i

---

- Use the following CL command:
- `DSPJOB JOB(SCPF) OPTION(*STSA)`
- If older joblogs exists from previous IPL's, you will see a list like this:

```
                Select Job                                DKIBMASF
                                                    11/08/13 19:33:41

Type option, press Enter.
  1=Select
```

| Option | Job  | User | Number | Type | -----Status----- | Entered<br>System |
|--------|------|------|--------|------|------------------|-------------------|
|        | SCPF | QSYS | 000000 | SYS  | ACTIVE           | 11/07/13          |
|        | SCPF | QSYS | 060989 | SYS  | OUTQ             | 08/29/13          |

- Then you just select the active one with the 1 and you can see last time the partition was IPL'ed, similar to this:

# The last IPL on an IBM i

---

```

      Display Job Status Attributes
      System:      DKIBMASF
Job:   SCPF      User:   QSYS      Number:  000000
      Status of job . . . . . :  ACTIVE
      Current user profile . . . . . :  QSYS
      Job user identity . . . . . :  QSYS
      Set by . . . . . :  *DEFAULT
      Entered system:
      Date . . . . . :  11/07/13
      Time . . . . . :  15:43:33
      Started:
      Date . . . . . :  11/07/13
      Time . . . . . :  15:43:33
      Subsystem . . . . . :
      Subsystem pool ID . . . . . :
      Type of job . . . . . :  SYS
      Special environment . . . . . :  *NONE
      Program return code . . . . . :  0

```

Morten Rasmussen  
IBM DK

• So you have the start time of the system.

# 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

# 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

# HELP!

---

- Please send in IBM i tips and tricks





---

# Scripting etc

# Beep

---

- In a loop with a sleep:

```
while :
do
echo -e \\007      #Linux
# echo \\007      #AIX
lssysconn -r all -F type_model_serial_num
sleep 30
done
```

- you get a beep to tell you to look for an update.

# vi

---

- Run a command

`:!`

- Repeat last command

`:!!`

- Read next file

`:n`

- Re-read current file

`:e %`

- Re-read previous file

`:e #`

- Rewind

`:rew`

`:map lhs rhs`

- eg:

`:map <CTRL-V><F1> :!ps -ef<CTRL-V><CTRL-M>`

# shell

---

- The cd command:

```
/apps/my-apps/tools/hammer/V1/admin/logs/debug# cd V1 V2
/apps/my-apps/tools/hammer/V2/admin/logs/debug
/apps/my-apps/tools/hammer/V2/admin/logs/debug#
```

- Useful for jumping between different versions etc

- Use hard links

```
peach-gaz:/usr/local/bin# ls -l *hmc*
-rwxr-xr-x 6 root system 42 07 Aug 2013 blackhmc
-rwxr-xr-x 6 root system 42 07 Aug 2013 hmc11
-rwxr-xr-x 6 root system 42 07 Aug 2013 hmc12
-rwxr-xr-x 6 root system 42 07 Aug 2013 hmc13
-rwxr-xr-x 6 root system 42 07 Aug 2013 hmc14
```

```
peach-gaz:/usr/local/bin# cat hmc14
HMC=$(basename $0)
```

```
ssh hscroot@${HMC} $@
```

•Saves disk space  
•Edit once

# Shell

---

- Efficient handling of a .tgz
- `gunzip -c /tmp/backuphdr.tgz | pax -s //./ -r`
- If / filesystem is read only, you can remount it read/write
- `mount -o remount,rw`

- **Some variables**

|                   |                                |
|-------------------|--------------------------------|
| <code>\$\$</code> | PID of current shell           |
| <code>\$!</code>  | PID of last background process |
| <code>\$?</code>  | Exit status of last command    |
| <code>\$0</code>  | The name of the shell          |
| <code>%1</code>   | PID of backgrounded job #1     |

# Odds

---

- ksh dry run, very use full to test *postinstall*, *fb\_script* syntax on *NIM* server without running it :
  - # ksh -xvn postinstall.ksh
- An old one but still good to know to show *HMC* connections from system :
  - before AIX 6.1 TL 07
    - # lsrsrc -ls IBM.ManagementServer
  - after AIX 6.1 TL 07
    - # lsrsrc -ls IBM.MCP

@chmod666

# Odds

---

- loop mounts

```
loopmount -i /tmp/AIX616_dvd1.iso -m /mnt -o "-V cdrfs -o ro"  
loopumount -l loop0 -m /mnt
```

- CPU speed

```
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
```

---

# Snippets



## A bit of a laugh

---

### Always give 100%

|           |           |
|-----------|-----------|
| Monday    | 12%       |
| Tuesday   | 23%       |
| Wednesday | 40%       |
| Thursday  | 20%       |
| Friday    | <u>5%</u> |
|           | 100%      |

**Unix was not designed to stop you from doing stupid things,  
because that would stop you from doing clever things.**

**Doug Gwyn**

© Copyright IBM Corporation 2014

---

- **largesend**

- padmin:           \$ chdev -dev en10 -attr mtu\_bypass=on
- root:             # chdev -l en0 -a mtu\_bypass=on

# Twitter

---

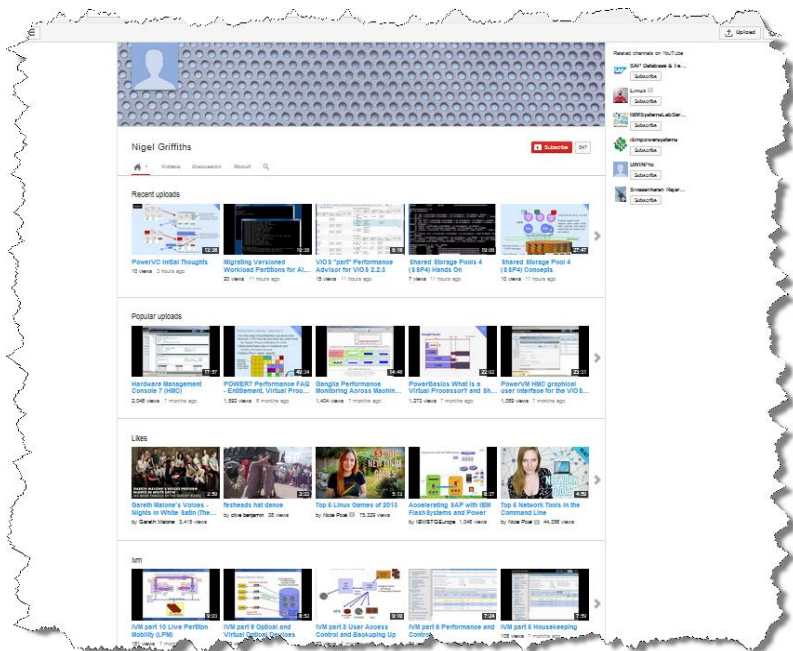
- Follow these twitterers



@power\_gaz  
@mr\_nmon  
@JyotiDodhia  
@cgibbo  
@NicoletteMcF  
@robmcnelly  
@ibmaix  
@chmod666  
@chromeaix  
@IBMPureSystems  
@IBMPowerSystems  
@AIXUserGroup

# Nigel's Channel on YouTube

- <http://www.youtube.com/user/nigelargriffiths>



# Technical University

<http://tinyurl.com/IBMtechu14>

Training - worldwide > Technical events & conferences > Calendar >

2014

## IBM Power Systems Technical University

Featuring IBM PureSystems

20-24 Oct

Registration open: **May 2, 2014**



Overview

Registration and Travel

Agenda

Solution Center

### The right IT infrastructure positions your business to thrive on a smarter planet.

Attend the 2014 IBM Power Systems Technical University to ramp up quickly on smarter computing implementation, to gain an in-depth look at the latest innovations from IBM Systems, to learn about recent product announcements, and the chance to swap insights with colleagues and learn from the experts.



*The conference venue will be confirmed in the coming weeks.*

**Is this event right for you?**

Learn how this event can deepen your skills on your organization's critical technologies.

Resources

[Week at a glance](#) (28.29 KB)

Have questions?

Connect with IBM technical events



**REGISTER NOW!**

# Technical University

<http://tinyurl.com/IBMtechu14>

Training - worldwide > Technical events & conferences > Calendar >

2014

## IBM Power Systems Technical University

Featuring IBM PureSystems

20-24 Oct

Registration open: **May 2, 2014**



Overview

Registration and Travel

Agenda

Solutions

The right IT infrastructure  
positions your business to  
thrive on a smarter planet.



IBM Power Systems Technical University  
20-24 October  
Budapest, Hungary

Attend the 2014 IBM Power Systems  
Technical University to ramp up quickly  
smarter computing implementation, to gain  
in-depth look at the latest innovations from IBM  
Power Systems, to learn about recent product  
announcements, and the chance to swap  
insights with colleagues and learn from the  
experts.

Is this event right for you?

Learn how this event can deepen your  
skills on your organization's critical  
technologies.

Resources

Week at a glance (28.29 KB)

Have questions? +

Connect with IBM technical events



**REGISTER NOW!**

# FEEDBACK

---

IS THIS KIND OF SESSION USEFUL?

If so, please send me your tips ...

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



**Thanks!!!**

**W5: Pillars of Star Formation © NASA**

© Copyright IBM Corporation 2014



Next Time

## 16th July PowerKVM Deep Dive with Dr Michael Perzl



### Previous Sessions:

POWER8 from hands-on  
Power up your Linux  
PowerVC  
PowerVP  
SSP4  
Best Practices  
Tricks of Power Masters  
IBMi and External Storage  
Monitoring with ITM  
Whole Machine Monitoring  
Electric Server Agent  
RDX Removable disks  
And more.....

### Future Sessions

- Suggestions Welcome



Gareth Coates @power\_gaz  
Jyoti Dodhia @JyotiDodhia



Nigel Griffiths @mr\_nmon  
Mandie Quartly @mandieq

Website: <http://tinyurl.com/PowerSystemsTechnicalWebinars>  
Youtube Channel: <http://tinyurl.com/IBMPowerVUGYoutubeChannel>