

CARE AND FEEDING OF VIO SERVERS

HANDOUT AT:

[HTTP://WWW.CIRCLE4.COM/FORSYTHE/VIOSMAINT-MAY2015.PDF](http://www.circle4.com/forsythe/viosmaint-may2015.pdf)

Jaqui Lynch
lynchj@forsythe.com



AGENDA

- Best Practices Setup
- Installation
- Maintenance and Upgrades
- Backup and recovery
- Monitoring
- HMC and Firmware Maintenance
- Wrap-up/Questions
- Backup Material
 - Useful VIOS and HMC Commands



BEST PRACTICES SETUP

FUNDAMENTALS BEFORE
YOU START



3



JUST PUBLISHED

VIOS 101 – CPU and Processors

<http://www.ibmssystemsmag.com/aix/administrator/virtualization/vios-101-part1/>

VIOS 101 – Memory

<http://www.ibmssystemsmag.com/aix/administrator/virtualization/vios-101-part2/>

Part 3 on I/O and Part 4 on Network underway

4



STAY CURRENT

VIOS Lifecycle

Version	GA	EOM	EOS/EOL
1.5	11/07	2008	09/11
2.1	11/08	2010	09/12
2.2.0 0	9/10	2011	09/13
2.2.1	10/11	10/12	04/15
2.2.2	10/12	10/13	09/16
2.2.3	4Q13		
2.2.4	2Q15		

Latest release:

2.2.3.4 FP27 SP04

Basically AIX 6100-09-044-1441

Requires NIM at 6100-09-04 or 7100-03-04

Can be applied to 2.2.3.0, 2.2.3.1, 2.2.3.2 or 2.2.3.4

Note there are three Interim fixes on Fix Central for this leave

Download updates from Fix Central:

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

Download base from entitled software:

<https://www-05.ibm.com/servers/eserver/ess/ProtectedServlet.wss>

HMC latest version is v8.8.2.0M1 (MH01455) with patches MH01499, MH01507 and MH01515

5



USE FLRT AND CHECK PREREQS

FLRT Home Page:

<http://www14.software.ibm.com/webapp/set2/flrt/home>

VIOS to NIM Master Mapping:

<http://www14.software.ibm.com/webapp/set2/sas/f/flrt/viostable.html>

System Software Maps for VIOS:

<http://www-01.ibm.com/support/docview.wss?uid=ssm1platformvios>

AIX/VIOS Security Tables:

http://www14.software.ibm.com/webapp/set2/sas/f/flrt3/Sec_APARs.html

VIOS Hiper Tables:

http://www14.software.ibm.com/webapp/set2/flrt/doc?page=hiper#vios_hiper

Also check MPIO driver versions as there are specific requirements for each VIO release

6



GENERAL

- Keep it simple
- Ensure LMB is the same on all servers if you want to use LPM
- Use hot pluggable adapters rather than built in ones
 - Easier maintenance
- Use dual VIO to allow for concurrent updates
- All adapters should be desired, not required
- Don't mix multipath drivers / HBAs
- Run HMC Scanner and/or Sysplan before and after all changes
- Plan for at least one update per year (IBM normally puts out 2)
- Separate VIOs for production and non prod on large systems
- Test failover (SEA failover and disk if vio goes down)
- Use VIO commands wherever possible rather than going into oem_setup_env
- mirror vio rootvg
- NOTE – v2 requires at LEAST 30GB in rootvg
- Fix Paging- By default VIO has a 512MB hd6 and a 1.5GB paging00 on the same LUN
- Add logging and set up dump devices properly
- Run VIOS Advisor regularly
- Check errpt regularly
- NEVER run at 100% entitlement – ensure it is high enough and there are plenty of VPs and memory
- Backup regularly – use NIM or scripts

7



SIZING

Use Systems Planning Tool

- Plan and design configuration
- <http://www-947.ibm.com/systems/support/tools/systemplanningtool/>

Try Workload Estimator

- <http://www-947.ibm.com/systems/support/tools/estimator/index.html>

VIOS and Virtualization Performance Advisors

- <https://www-304.ibm.com/support/docview.wss?uid=aixtools159f1226>
- <https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Power%20Systems/page/PowerVM%20Virtualization%20Performance%20Advisor>

Minimums

- Memory 4GB
- Cores .5 entitlement and 2VPs
- BUT remember that the more VFCs and high performance adapters the more memory and CPU you will need
- Also VIO servers perform based on entitlement not VPs
- So you could need more like 6 or 8GB and an entitlement of 1.5 or 2.

Pay attention to adapter placement – adapter slots have different priorities

8



MORE ON SIZING

If using 10Gb or 8Gb adapters need more memory for buffering and more CPU to handle traffic

i.e. 512MB for each active high performance adapter port
140MB per VFC client in the VIO

vSCSI uses more CPU in the VIO than NPIV

High values for VIO adapter slots can also increase memory needs

Not uncommon to see a VIO now needing 6-8GB memory and entitlement of 1-2 cores

rootvg needs at least 30GB

Add an extra disk if want to use FBO – don't put it in rootvg as it will make backups of rootvg enormous

VIOS Sizing Considerations:

<http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/perf.html>



MEMORY PLANNING

<http://www.circle4.com/ptechu/memoryplan.xlsx>

Note div 64 – is 128 for p7+ and p8

Memory Planning Worksheet										This gives a rough estimate	
Power7 770										Assumes LMB size is 256MB	
Max RAM Capacity										Each active IVE port adds 102 MB	
	786432	Ram installed	393216	Ram Active	131072						
		GB	384	MB LMB =	256	LMB below in MB					
Change the LMB size on this line to match MRO on HMC										Used the largest to show worst possible	
Extra high performance ports per VIO										8 NPIV VFCs per VIO	
MB LMB =										12	
LPAR NAME	Desired Memory MB	Maximum Memory MB	Ohead Max Div 64	OH/LMB MB	Roundup OH MB	Actual Ohead (MB)	Memory Needed	Extra high if NPIV Perf ports			
VIOS1	3172	4096	64	0.25	1	256		4096	1680		
VIOS2	3172	4096	64	0.25	1	256		4096	1680		
LPAR1	12032	16384	256	1.00	1	256					
LPAR2	20224	24576	384	1.50	2	512					
LPAR3	14336	16384	256	1.00	1	256					
LPAR7	17152	17152	268	1.05	2	512					
LPAR8	65536	71680	1120	4.38	5	1280					
LPAR9	32768	36864	576	2.25	3	768					
HYPERVISOR								768			
IVE								102			
I/O drawer (I use 512 per 2)								512			
Safety Net								512			
										MB Total	
MB Total	168392	191232	2988	11.671875	16	5990	174382	8192	3360		
GB Total	164					5.85	170	8.00	3.28		
										GB Total	
Hypervisor requires 6GB minimum for overhead with these settings										Add High Perf	
LPARs require 164GB so the total active needed is at least 170GB										Or add NPIV	
Need to add NPIV and high speed adapter memory needs as well - could be 178GB											
Look at AME potential as well											
8GB and 10GB extra high performance adapters											
For each active port add 512MB											
If NPIV then 140MB per VFC adapter per client											
i.e. 20 ports per VIO without NPIV would be 20 * 512 = 10GB plus VIOS base for each VIOS											
If NPIV then we allocate per client so if there are 20 clients on each VIO then each											
VIO needs 20 * 140 = 2.8GB above the base											



HBA SETTINGS



11



HBA TUNING

- Make the same tuning changes you would make on AIX
- Set num_cmd_elems and max_xfer_size on the fiber adapters on VIO


```
chdev -l fcs0 -a max_xfer_size=0x200000 -a num_cmd_elems=1024 -P
chdev -l fcs1 -a max_xfer_size=0x200000 -a num_cmd_elems=1024 -P
```

Check these numbers are supported by your disk vendor

- If NPIV also set on clients
- Client setting cannot be higher than the VIOs
- Pay attention to adapter layout and priorities
- NOTE – as of AIX v7.1 tl2 (or 6.1 tl8) num_cmd_elems is limited to 256 on the VFCs so set num_cmd_elems to the high number on the VIO but to no more than 256 on the NPIV clients
- See: <http://www-01.ibm.com/support/docview.wss?uid=isg1IV63282>

12



ADAPTER TUNING 1/2

fcs0				
bus_intr_lvl	115	Bus interrupt level	False	
bus_io_addr	0xdfc00	Bus I/O address	False	
bus_mem_addr	0xe8040000	Bus memory address	False	
init_link	al	INIT Link flags	True	
intr_priority	3	Interrupt priority	False	
lg_term_dma	0x800000	Long term DMA	True	
max_xfer_size	0x100000	Maximum Transfer Size	True	(16MB DMA)
num_cmd_elems	200	Maximum number of COMMANDS to queue to the adapter	True	
pref_alpa	0x1	Preferred AL_PA	True	
sw_fc_class	2	FC Class for Fabric	True	

Changes I often make (test first)

max_xfer_size	0x200000	Maximum Transfer Size	True	128MB DMA area for data I/O
num_cmd_elems	1024	Maximum number of COMMANDS to queue to the adapter	True	

Often I raise this to 2048 – **check with your disk vendor**

lg_term_dma is the DMA area for control I/O

13



ADAPTER TUNING 2/2

Check these are ok with your disk vendor!!!

```
chdev -l fcs0 -a max_xfer_size=0x200000 -a num_cmd_elems=1024 -P
chdev -l fcs1 -a max_xfer_size=0x200000 -a num_cmd_elems=1024 -P
```

At AIX 6.1 TL2 VFCs will always use a 128MB DMA memory area even with default max_xfer_size

Remember make changes too both VIO servers and client LPARs if using NPIV VIO server setting must be at least as large as the client setting

Remember VFCs on the client are limited to num_cmd_elems=256 after AIX 6.1 tl8 or 7.1 tl2

See Dan Braden Techdoc for more on tuning these:

<http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/TD105745>

14



MY VIO SERVER AND NPIV CLIENT ADAPTER SETTINGS

VIO SERVER

```
#lsattr -El fcs0
```

lg_term_dma	0x800000	Long term DMA	True
max_xfer_size	0x200000	Maximum Transfer Size	True
num_cmd_elems	2048	Max number of COMMANDS to queue to the adapter	True

NPIV Client (running at defaults before changes)

```
#lsattr -El fcs0
```

lg_term_dma	0x800000	Long term DMA	True
max_xfer_size	0x200000	Maximum Transfer Size	True
num_cmd_elems	256	Maximum Number of COMMAND Elements	True

NOTE NPIV client must be <= to settings on VIO

15



NETWORK



16



VIRTUAL ETHERNET

Link aggregation

Put vio1 aggregate on a different switch to vio2 aggregate
 Provides redundancy without having to use NIB
 Allows full bandwidth and less network traffic (NIB is pingy)
 Basically SEA failover with full redundancy and bandwidth

Pay attention to entitlement

VE performance scales by entitlement not VPs

If VIOS only handling network then disable network threading on the virtual Ethernet

chdev -dev ent? thread=0
 Non threaded improves LAN performance
 Threaded (default) is best for mixed vSCSI and LAN

<http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/perf.html>

Turn on large send on VE adapters

chdev -dev ent? -attr large_send=yes

Turn on large send on the SEA

chdev -dev entx -attr largesend=1

NOTE do not do this if you are supporting Linux or IBM i LPARs with the VE/SEA

17



STARTER SET OF TUNABLES NETWORK

Typically we set the following:

NETWORK

```
no -p -o rfc1323=1
no -p -o tcp_sendspace=262144
no -p -o tcp_recvspace=262144
no -p -o udp_sendspace=65536
no -p -o udp_recvspace=655360
```

Also check the actual NIC interfaces and make sure they are set to at least these values

You can't set udp_sendspace > 65536 as IP has an upper limit of 65536 bytes per packet

Check sb_max is at least 1040000 – increase as needed

18



My VIO SERVER SEA

```
# ifconfig -a
en6:
flags=1e080863,580<UP,BROADCAST,NOTRAILERS,RUNNING,SIMPLEX,MULTICAST,
GROUPRT,64BIT,CHECKSUM_OFFLOAD(ACTIVE),CHAIN>

    inet 192.168.2.5 netmask 0xfffff00 broadcast 192.168.2.255
    tcp_sendspace 262144 tcp_recvspace 262144 rfc1323 1

lo0:
flags=e08084b,1c0<UP,BROADCAST,LOOPBACK,RUNNING,SIMPLEX,MULTICAST
,GROUPRT,64BIT,LARGESEND,CHAIN>
    inet 127.0.0.1 netmask 0xff000000 broadcast 127.255.255.255
    inet6 ::1%1/0
    tcp_sendspace 131072 tcp_recvspace 131072 rfc1323 1
```

19



NETWORK PERFORMANCE AND THROUGHPUT

Depends on:

- Available CPU power
 - **Scales by entitlement not by VPs**
- MTU size
- Distance between receiver and sender
- Offloading features
- Coalescing and aggregation features
- TCP configuration
- Firmware on adapters and server
- Ensuring all known fixes are on for 10GbE issues

Network Performance Presentation at:

<http://youtu.be/8pth2ujGWK0>

<http://www.circle4.com/movies/networkperf/networkperf.pdf>

20



ALEXANDER PAUL – SESSION APE1271 AT EDGE 2015

Approximate Theoretical Virtual Ethernet Throughput depends on Capping and MTU

	Capped at .4 MTU 1500	Capped at .4 MTU 9000	Uncapped 4 VPs MTU 1500	Uncapped 4 VPs MTU 9000	Capped at .4 MTU 65390	Uncapped 4 VPs MTU 65390
9117-MMB	451 Mb/s	1.6 Gb/s	1.25 Gb/s	6.5 Gb/s		
9117-MMD	991 Mb/s	1.9 Gb/s	1.6 Gb/s	7 Gb/s	3.09 Gb/s	19.7 Gb/s
S824	1.4 Gb/s	4.7 Gb/s	2.55 Gb/s	13 Gb/s		49.2 Gb/s
E870	1.3 Gb/s		2.9 Gb/s			66.3 Gb/s

Approximations taken from slides in presentation

21



VIO 2.2.3 SEA CHANGES

Traditional SEA setup

Ent4 is the virtual adapter defined at the HMC with external access

(SEA goes here)

VIO1 is priority 1 and VIO2 is priority 2

Ent5 is the virtual adapter on Vlan 1 with no external

(IP will go here)

Ent6 is the control channel on vlan 255

```
mkvdev -sea ent0 -vadapter ent4 -default ent4 -defaultid 1 -attr ha_mode=auto ctl_chan=ent6
```

Creates ent7 as the SEA

Do not mess up priorities or ctl_chan or you will cause a spanning tree loop

Update with 2.2.3

See chapter 4 of SG248198- Redbook on 2.2.3 Enhancements

SEA setup has been simplified

Requirement removed for dedicated control channel and VLAN ID for each SEA failover configuration

Multiple SEA pairs can now share VLAN 4095 within the same virtual switch and no ctl_chan is needed

HMC (>= 7.8) reserves 4095 for internal management traffic

Requires VIOS 2.2.3, HMC 7.7.8 and firmware 780 or higher

Not available on 770/780 B models

22



INSTALLATION



23



INSTALL OPTIONS

From DVD – complete install

Using NIM

<http://www-01.ibm.com/support/docview.wss?uid=isg3T1011386>

Using HMC - check vios install box

Commandline - installios:

[http://www-](http://www-01.ibm.com/support/knowledgecenter/POWER7/p7hb11/iphb1_vios_configuring_installhmc.htm?cp=POWER7%2F14-8-0-2-2-1-1)

[01.ibm.com/support/knowledgecenter/POWER7/p7hb11/iphb1_vios_configuring_installhmc.htm?cp=POWER7%2F14-8-0-2-2-1-1](http://www-01.ibm.com/support/knowledgecenter/POWER7/p7hb11/iphb1_vios_configuring_installhmc.htm?cp=POWER7%2F14-8-0-2-2-1-1)

GUI:

<http://ibmsystemsmag.blogs.com/aixchange/2013/05/vios-installation-via-gui.html>

From a mksysb

http://pic.dhe.ibm.com/infocenter/flexsys/information/index.jsp?topic=%2Fcom.ibm.acc.psm.resources.doc%2Fvios%2Fsdm_c_vios-vios_backup_restore_file_nim.html

24



VIOS AND NIM

Use of NIM to back up, install, and update the VIOS is supported.

Note: For install, always create the SPOT resource directly from the VIOS **mksysb** image. Do NOT update the SPOT from an LPP_SOURCE.

Use of NIM to update the VIOS is supported as follows:
Ensure that the NIM Master is at the appropriate level to support the VIOS image.
<http://www14.software.ibm.com/webapp/set2/sas/f/flrt/viostable.html>

On the NIM Master, use the operation **updateios** to update the VIOS Server.
"nim -o updateios -a lpp_source=lpp_source1"

On the NIM Master, use the operation **alt_disk_install** to update an alternate disk copy of the VIOS Server.

**"nim -o alt_disk_install -a source=rootvg -a disk=target_disk
-a fix_bundle=(Value)"**

If NIM is not used to update the VIOS, only the **updateios** or the **alt_root_vg** command from the padmin shell can be used to update the VIOS.

25



VIOS AND NIM

Add VIOS partition as a NIM client

Copy the VIOS mksysb image from the CD to your NIM master

- On VIOS 2.2 media there are 3 images now – the 3rd is on DVD 2
- Copy all 3 images individually to a directory and then use cat to combine them
cat /export/mksysb/vios2.2/mksysb_image /export/mksysb/vios2.2/mksysb_image2
/export/mksysb/vios2.2/mksysb_image3 >/export/mksysb/nim_vios2.2mksysb

Define mksysb resource to NIM master

Define spot on NIM master

- The source for the SPOT will be the combined mksysb
- The SPOT CANNOT be created from an LPP_Source

Copy the bosinst.data from the DVD and create a viosbosinst resource

You can now use bos_inst to do a mksysb install once the partition profile is defined

<http://www-01.ibm.com/support/docview.wss?uid=isg3T1011386>

26



MAINTENANCE AND UPGRADES



27



UPDATING VIOS

Run Isvopt and make sure no one is using the FBO devices

1. Normally upgrade HMC first then firmware then VIOS and then AIX
2. BUT – check the readme for all of the above first to make sure there is not a different required order
3. Download the updates and cross-check compatibility using FLRT
4. Read the readme again
5. Run errpt to check for problems, check there are no stale partitions, missing disks or paths, etc
 - Isvg rootvg checks for stale PPs and physical volumes.
 - Isvg -p rootvg looks for missing disks.
 - lspath - checks for missing paths.
 - errpt checks for errors.
6. Ensure all paths on clients are redundant so LPARs will stay up when this VIOS is rebooted
7. Run HMC Scanner or sysplan to document prior to changes
8. Backup the VIOS
9. Mount the NFS filesystem or DVD or FBO image to be used for update
10. If using SSPs there are specific additional steps outlined in the README

28



MIGRATION 1/2

Back the VIO up before doing anything and again when done!
If migrating from a pre v2 level ensure VP folding is turn off after the migration

1. In order to migrate to v2.* your HMC must be at v7 or later at least 7.7.4
If VIOS is lower than v2.1 then you must migrate to 2.1.0 using the migration DVD

2. Migrating from prior to v1.3
Basically this is a reinstall

3. Migrating from v1.3 or v1.4
Need the migration DVD for VIOS 1.5 or the updates
Need to update to VIOS 1.5.2.6-FP-11.1 SP-02 prior to upgrade to v2

4. Migrating from v1.5.2.6-FP-11.1 SP-02 or higher
Need the migration DVD for VIOS v2
Boot from the DVD in SMS mode and tell it to do a migration upgrade

Note – once at v2.1 you need to update to 2.2.3.1 prior to applying 2.2.3.4
2.2.3.4 requires a minimum release of 2.2.3.0 in order to be applied
Single step update requires VIO between 2.2.1.1 and 2.2.2.x
NIM allows you to create a single merged lpp_source to get around this but cannot be used with SDDPCM

29



MIGRATION 2/2

5. See Power VM Managing and Monitoring Redbook – Chapter 11
<http://www.redbooks.ibm.com/redbooks/pdfs/sg247590.pdf>

NOTE IBM has a simplified migration offering
http://www.ibmsystemsmag.com/ibmi/trends/ibmannouncements/vios_migration/

Once you are on v2.1 then upgrades are all done using updateios or nim
There are specific concerns around updates if you are running SSPs (Shared storage pools)

30



UPDATING VIOS WITH FIXPACKS OR SPs

From 2.2.3.2 to 2.2.3.3

As padmin run “updateios –commit” to ensure any uncommitted updates are committed

Check to ensure there are no missing filesets prior to updates

```
$ ioslevel
```

```
2.2.3.2
```

```
$cat /usr/ios/cli/ios.level
```

```
$cat /usr/ios/cli/SPLEVEL.TXT
```

The above two will get you the IOS level and the SP

```
$ updateios -commit
```

All updates have been committed.

```
$ oem_setup_env
```

```
# /usr/sbin/emgr -P
```

There is no efix data on this system.

Now run checks

31



PRE INSTALL CHECKS FOR VIOS 2.2.3.2 TO 2.2.3.3 UPDATE

Did VIO2 (secondary VIO first):

```
$ ioslevel
```

```
2.2.3.2
```

```
$ oem_setup_env
```

```
#df -g - make sure no filesystems are full
```

```
#oslevel -s
```

```
6100-09-02-1412
```

```
# instfix -i | grep ML
```

```
All filesets for 6.1.0.0_AIX_ML were found.
```

```
All filesets for 6100-00_AIX_ML were found.
```

```
All filesets for 6100-01_AIX_ML were found.
```

```
All filesets for 6100-02_AIX_ML were found.
```

```
All filesets for 6100-03_AIX_ML were found.
```

```
All filesets for 6100-04_AIX_ML were found.
```

```
All filesets for 6100-05_AIX_ML were found.
```

```
All filesets for 6100-06_AIX_ML were found.
```

```
All filesets for 6100-07_AIX_ML were found.
```

```
All filesets for 6100-08_AIX_ML were found.
```

```
All filesets for 6100-09_AIX_ML were found.
```

```
# lppchk -v
```

```
# lppchk -vm3
```

```
# oslevel -s -l 6100-09-02-1412
```

```
#errpt | more – check there are no errors
```

32



CONTINUE 2.2.3.3 UPDATE BACKUP 1/2

```

Back it up:
# ./save-viostuff.sh
mkdir: 0653-358 Cannot create /home/padmin/saveit.
/home/padmin/saveit: Do not specify an existing file.

# ls -l /home/padmin/saveit
total 824
-rw-r--r-- 1 root staff 118 Jul 22 12:33 b740vio2.disktmp.txt
-rw-r--r-- 1 root staff 24 Jul 22 12:33 b740vio2.ioslevel.txt
-rw-r--r-- 1 root staff 16 Jul 22 12:33 b740vio2.oslevel.txt
-rw-r--r-- 1 root staff 8038 Jul 22 12:33 b740vio2.vioadapter.txt
-rw-r--r-- 1 root staff 4528 Jul 22 12:33 b740vio2.viodisk.txt
-rw-r--r-- 1 root staff 59593 Jul 22 12:33 b740vio2.viodisks.txt
-rw-r--r-- 1 root staff 8800 Jul 22 12:33 b740vio2.violsdevv.txt
-rw-r--r-- 1 root staff 11967 Jul 22 12:33 b740vio2.violsmapall.npiv.txt
-rw-r--r-- 1 root staff 19363 Jul 22 12:33 b740vio2.violsmapall.txt
-rw-r--r-- 1 root staff 4595 Jul 22 12:33 b740vio2.vioslots.txt
-rw-r--r-- 1 root staff 227944 Jul 22 12:33 b740vio2.viovpd.txt
-rw-r--r-- 1 root staff 37 Jul 22 12:33 cfgname.txt
-rw-r--r-- 1 root staff 0 Jul 22 12:33 entstat.txt
-rw-r--r-- 1 root staff 240 Jul 22 12:33 firewall.txt
-rw-r--r-- 1 root staff 652 Jul 22 12:33 hostmap.txt
-rw-r--r-- 1 root staff 5970 Jul 22 12:33 optimize.txt
-rw-r--r-- 1 root staff 713 Jul 22 12:33 routinfo.txt
-rw-r--r-- 1 root staff 240 Jul 22 12:33 user.txt
-rw-r--r-- 1 root staff 15071 Jul 22 12:33 view.txt

```

33



CONTINUE 2.2.3.3 UPDATE BACKUP 2/2

```

$ viosbr -backup -file /home/padmin/saveit/b740vio2-backup
Backup of this node (b740vio2) successful

oem_setup_env
# mount /usr/local/backups
# su - padmin -c "ioscli backupios -file /usr/local/backups/b740vio2-jul2214.mksysb -mksysb"
/usr/local/backups/b740vio2-jul2214.mksysb doesn't exist.
Creating /usr/local/backups/b740vio2-jul2214.mksysb
*** Here it is doing a savevgstructs for rootclients_vg *****
Creating information file for volume group rootclients_vg.
Creating list of files to back up.
Backing up 6 files
6 of 6 files (100%)
0512-038 savevg: Backup Completed Successfully.
Backup in progress. This command can take a considerable amount of time
to complete, please be patient...

Creating information file (/image.data) for rootvg.
Creating list of files to back up.
Backing up 160374 files.....
39229 of 160374 files (24%).....
160374 of 160374 files (100%)
0512-038 savevg: Backup Completed Successfully.

```

34



CONTINUE 2.2.3.3 UPDATE INSTALL 1/3

Download from Fix Central the iso image for 2.2.3.3 – I do this to my NIM server
 It came down as H52175995.iso
 mkdir /cdrom
 loopmount -i H52175995.iso -o "-V cdrfs -o ro" -m /cdrom
 smitty bffcreate – I do this on my NIM server and create a directory to put the files in that the VIO has access to
 In this case /usr/local/soft/vios2233

Normally I copy the files locally to the VIO in case I lose the network during the install

35



CONTINUE 2.2.3.3 UPDATE INSTALL 2/3

Now on the VIO:
 \$ updateios -accept -install -dev /usr/local/soft/vios2233

```
*****
installp PREVIEW: installation will not actually occur.
*****
```

```
+-----+
| Pre-installation Verification... |
+-----+
```

```
Verifying selections...done
Verifying requisites...done
Results...
SUCSESSES
```

```
-----
Filesets listed in this section passed pre-installation verification
and will be installed.
```

```
Mandatory Fileset Updates
```

```
-----
(being installed automatically due to their importance)
bos.rte.install 6.1.9.16 # LPP Install Commands
<< End of Success Section >>
```

Prompts you to reply Y which you do and it installs them

36



CONTINUE 2.2.3.3 UPDATE INSTALL 3/3

After bos.rte.install is installed it then prompts you re installing the other 272 fixes
 Check estimated space needed and free space and if all is good then:
 Reply Y and they begin installing – takes about 2 hours depending

```
$ioslevel
Shows as 2.2.3.3
```

```
$oem_setup_env
# oslevel -s
6100-09-03-1415
```

```
lspv | grep rootvg
hdisk0    00f6934cc34a30f3      rootvg    active
hdisk1    00f6934c30e34699      rootvg    active
```

```
bosboot -a -d hdisk0
bosboot -a -d hdisk1
bootlist -m normal hdisk0 hdisk1
```

Now reboot and then run post install tests

37



POST INSTALL CHECKS

```
$ ioslevel
2.2.3.3
$ oem_setup_env

# oslevel -s
Should show: 6100-09-03-1415
6100-09-03-1415
# instfix -i | grep ML
All filesets for 6100-00_AIX_ML were found.
All filesets for 6100-01_AIX_ML were found.
All filesets for 6100-02_AIX_ML were found.
All filesets for 6100-03_AIX_ML were found.
All filesets for 6100-04_AIX_ML were found.
All filesets for 6100-05_AIX_ML were found.
All filesets for 6100-06_AIX_ML were found.
All filesets for 6100-07_AIX_ML were found.
All filesets for 6.1.0.0_AIX_ML were found.
All filesets for 6100-08_AIX_ML were found.
All filesets for 6100-09_AIX_ML were found.
# lppchk -v
# lppchk -vm3
# oslevel -s -l 6100-09-03-1415
#errpt | more – check there are no errors
```

Once all checks are passed and VIO2 is back up then go do the same upgrade to VIO1

38



UPDATING VIOS PROBLEMS

```

oem_setup_env
oslevel -s
6100-00-00-0000
instfix -i | grep ML
  All filesets for 6100-07_AIX_ML were found.
  All filesets for 6.1.0.0_AIX_ML were found.
  Not all filesets for 6100-08_AIX_ML were found.
This means there are missing filesets
# oslevel -sq
Known Service Packs
-----
6100-08-02-1316
6100-08-01-1245

# oslevel -s -l 6100-08-02-1316
Fileset                Actual Level    Service Pack Level
-----
bos.alt_disk_install.boot_images  6.1.8.0        6.1.8.15
bos.loc.utf.ES_ES        6.1.7.15        6.1.8.15

These filesets should be corrected prior to updating
Either use updateios to update them or to remove them

```

39



REMOVE OR UPDATE PROBLEM FILESETS

```

DO NOT USE SMITTY – use updateios

Issues with bos.suma
updateios -remove bos.suma

# oslevel -s -l 6100-08-02-1316
Fileset                Actual Level    Service Pack Level
-----
bos.alt_disk_install.boot_images  6.1.8.0        6.1.8.15
bos.loc.utf.ES_ES        6.1.7.15        6.1.8.15

updateios -remove bos.loc.utf.ES_ES

Upgrade alt disk
Copy images to be updated into a directory (/usr/local/soft/missing)
Run inutoc .

updateios -commit
Updateios -accept -install -dev /usr/local/soft/missing

Also remove efixes prior to updates:
/usr/sbin/emgr -P lists them

To remove:
# /usr/sbin/emgr -r -L <EFIX label>
emgr -r -L IV46869m3a

```

40



BACKUP AND RECOVERY



41



BACKING UP VIOS

Use `viosbr` to backup user defined virtual and logical resources on the VIO

Make sure to save that backup in `rootvg`

- `viosbr -backup -file /tmp/viosabkupbr`
- You can also use `viosbr` to view or restore
- <http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7hcg/viosbr.htm>

You may also want to use `snap` to grab other critical data

Mount NFS filesystem to backup to (in my case `/backups`)

```
mkdir /backups/viosa
```

Then as `padmin` run `backupios` which automatically calls `savevgstruct`:

```
backupios -file /backups/viosa
```

The above creates a `nim_resources.tar` package in that directory and it can be used to clone or restore VIO servers by NIM or `installios`

You can also back it up as a `mksysb` file that is easy to restore

```
backupios -file /backups/viosa.mksysb -mksysb
```

If the media library is large and is on `rootvg` then you can add the `-nomedialib` flag

42



BACKING UP VIOS FROM ROOT

As root run viosave.sh (see next slide)

lsvg – if you have other VGs i.e. datavg then for each

```
su - padmin -c "ioscli savevgstruct datavg"
```

```
su - padmin -c "ioscli viosbr -backup -file /tmp/viosabr.backup"
```

Mount the NFS repository for the backups (/nfsmnt)

```
su - padmin -c "ioscli backupios -file /nfsmnt/vio2-jul2114.mksysb -mksysb"
```

This backs it up to a bootable mksysb file

If using NIM to clone VIO servers don't forget:

```
su - padmin -c "ioscli backupios -file /nfsmnt/nimbkups"
```

This creates a nim_resources.tar file that can be used for restores described at:

http://public.dhe.ibm.com/software/server/vios/docs/backupios_mod.pdf

Create a daily backup once a day and keep up to 7 in /home/padmin/cfgbackups

```
viosbr -backup file viobkup -frequency daily numfiles 7
```

43



DOCUMENT VIO INFORMATION – VIOSAVE.SH

```
#!/bin/sh
#
day="/bin/date +%d"
month="/bin/date +%m"
year="/bin/date +%y"
set -- Dec Jan Feb Mar Apr May Jun Jul Aug Sep Oct Nov Dec
shift $month
lmonth="$1"
machine=$(uname -n)
directory="/bin/date +%m%d%Y_%H%M"
machine_directory=$(printf "%s_%s" $machine $directory)
mkdir /home/padmin/saveit
cd /home/padmin/saveit
logit="/home/padmin/saveit/$machine"
logit1="/home/padmin/saveit/$machine"
su - padmin -c "ioscli ioslevel" >>$logit1.ioslevel.txt
su - padmin -c "ioscli lsdev -type disk" >>$logit1.viodisk.txt
su - padmin -c "ioscli lsdev -type adapter" >>$logit1.vioadapter.txt
su - padmin -c "ioscli lsdev -vpd" >>$logit1.viovpd.txt
su - padmin -c "ioscli lsdev -slots" >>$logit1.vioslots.txt
su - padmin -c "ioscli lsmap -all" >>$logit1.viosmapall.txt
su - padmin -c "ioscli lsmap -all -npiv" >>$logit1.viosmapall.npiv.txt
su - padmin -c "ioscli lsdev -virtual" >>$logit1.viosdevv.txt
su - padmin -c "ioscli cfgrnamesrv -ls" >>cfgrname.txt
su - padmin -c "ioscli entstat -all ent9" >>entstat.txt
su - padmin -c "ioscli hostmap -ls" >>hostmap.txt
su - padmin -c "ioscli lsuser" >>user.txt
su - padmin -c "ioscli netstat -routinfo" >>routinfo.txt
su - padmin -c "ioscli optimizenet -list" >>optimize.txt
su - padmin -c "ioscli viosecure -firewall view" >>firewall.txt
su - padmin -c "ioscli viosecure -view -nonint" >>view.txt
oslevel -s >>$logit1.oslevel.txt
getlvodm -C >> $logit1.disktmp.txt
while read label line
do
echo "$line" >>$logit1.viodisks.txt
echo "Hdisk is $label" >>$logit1.viodisks.txt
echo "$line" >>$logit1.viodisks.txt
su - padmin -c "ioscli lsdev -dev $label -attr" >>$logit1.viodisks.txt
done <"$logit1.disktmp.txt"
#
exit 0
```

44



MONITORING



45



CPU AND MEMORY

Remember VIO scales by entitlement not VPs
Ensure sufficient entitlement
Watch for VCSWs – this is a sign of entitlement shortage
If running close to entitlement on average increase entitlement
If running close to VPs on average increase entitlement and VPs
Consider running dedicated

NEVER EVER let your VIO server page
Clean up the VIO server page spaces

46



NMON MONITORING

nmon -ft -AOPV^dMLW -s 15 -c 120

- Grabs a 30 minute nmon snapshot
- A is async IO
- M is mempages
- t is top processes
- L is large pages
- **O is SEA on the VIO**
- P is paging space
- V is disk volume group
- d is disk service times
- ^ is fibre adapter stats
- W is workload manager statistics if you have WLM enabled

If you want a 24 hour nmon use:

nmon -ft -AOPV^dMLW -s 150 -c 576

May need to enable accounting on the SEA first – this is done on the VIO
chdev -dev ent* -attr accounting=enabled

Can use entstat/seastat or topas/nmon to monitor – this is done on the vios
topas -E
nmon -O

VIOS performance advisor also reports on the SEAs

47



SHARED PROCESSOR POOL MONITORING

Turn on "Allow performance information collection" on the LPAR properties
This is a dynamic change

topas -C

Most important value is app – available pool processors
This represents the current number of free physical cores in the pool

nmon option p for pool monitoring

To the right of PoolCPUs there is an unused column which is the number of free
pool cores

nmon analyser LPAR Tab

lparstat

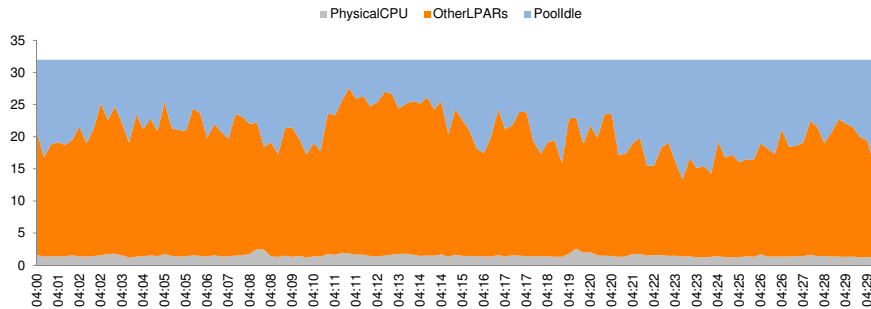
Shows the app column and poolsize

48



NMON ANALYSER LPAR TAB

Shared Pool Utilisation - b750n1 10/4/2013



49



NPIV STATISTICS

Normally need to use nmon to get information at each client LPAR

Could also use -O when recording

BUT as of v2.2.3

VIOS Performance advisor supports NPIV aggregation information

As of v2.2.2

<http://www-01.ibm.com/support/knowledgecenter/POWER7/p7hcg/fcstat.htm?cp=POWER7%2F1-8-3-8-2-60>

fcstat -n wwpn device_name

i.e. fcstat -n C05012345678000 fcs0

Provides statistics at the WWPn for the virtual adapter

You can also try fcstat -client

Also check out NPIVGRAPH for visualizing NPIV mappings:

<http://npivgraph.sourceforge.net/>

50



NETSTAT -V VIO

SEA

Transmit Statistics:

Packets: 83329901816
 Bytes: 87482716994025
 Interrupts: 0
 Transmit Errors: 0
 Packets Dropped: 0

Max Packets on S/W Transmit Queue: 374
 S/W Transmit Queue Overflow: 0
 Current S/W+H/W Transmit Queue Length: 0

Elapsed Time: 0 days 0 hours 0 minutes 0 seconds

Broadcast Packets: 1077222
 Multicast Packets: 3194318
 No Carrier Sense: 0
 DMA Underrun: 0
 Lost CTS Errors: 0
 Max Collision Errors: 0

Receive Statistics:

Packets: 83491933633
 Bytes: 87620268594031
 Interrupts: 18848013287
 Receive Errors: 0
Packets Dropped: 67836309
 Bad Packets: 0

Broadcast Packets: 1075746
 Multicast Packets: 3194313
 CRC Errors: 0
 DMA Overrun: 0
 Alignment Errors: 0

No Resource Errors: 67836309 check those tiny, etc Buffers

Virtual I/O Ethernet Adapter (I-Han) Specific Statistics:

Hypervisor Send Failures: 4043136
 Receiver Failures: 4043136
 Send Errors: 0

Hypervisor Receive Failures: 67836309

"No Resource Errors" can occur when the appropriate amount of memory can not be added quickly to vent buffer space for a workload situation.
 You can also see this on LPARs that use virtual Ethernet without an SEA

51



BUFFERS

Virtual Trunk Statistics

Receive Information

Receive Buffers

Buffer Type	Tiny	Small	Medium	Large	Huge
Min Buffers	512	512	128	24	24
Max Buffers	2048	2048	256	64	64
Allocated	513	2042	128	24	24
Registered	511	506	128	24	24
History					
Max Allocated	532	2048	128	24	24
Lowest Registered	502	354	128	24	24

"Max Allocated" represents the maximum number of buffers ever allocated

"Min Buffers" is number of pre-allocated buffers

"Max Buffers" is an absolute threshold for how many buffers can be allocated

chdev -l <veth> -a max_buf_small=4096 -P

chdev -l <veth> -a min_buf_small=2048 -P

Above increases min and max small buffers for the virtual ethernet adapter configured for the SEA above

Needs a reboot

Max buffers is an absolute threshold for how many buffers can be allocated

Use entstat -d (-all on vio) or netstat -v to get this information

52



HMC MAINTENANCE



53



UPGRADING HMC FROM 7.7.7.0 TO 7.7.8

ssh to HMC with 2 sessions

OUR HMC is 7042-cr6 installed at 7.7.7.0 SP2
Upgrading to HMC v7.7.8 MH01388

Step 1 Save upgrade data and then backup to USB stick or remote FTP using GUI

Step 2 check we have plenty of memory

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

Mem: 4095732k total, 3978304k used, 117428k free, 311480k buffers

So our server has 4GB memory

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

Check if filesystems are full

If they are in use a lot then

```
chhmcfs -o f -d 0
```

The above clears out all temp files

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

Also lshmcfs shows all filesystems

Check for profile sizing:

<http://www-01.ibm.com/support/docview.wss?uid=nas8N1019821>

54



UPGRADING HMC FROM 7.7.7.0 TO 7.7.8

ssh to HMC with 2 sessions

Since it is an upgrade we need to either use the media or do it via the CLI

On the first of the two ssh sessions: Login and cd /hmcdump

```
getupgfiles -h ftp.software.ibm.com -u anonymous --passwd ftp -d /software/server/hmc/network/v7780
```

On second ssh session:

```
ls -la /hmcdump
```

You will see files being loaded into the directory

Once everything is downloaded you will no longer see files in this directory

Exit this connection

On the first ssh session

```
chhmc -c altdiskboot -s enable --mode upgrade
```

The above tells it to set up to upgrade on boot

```
hmcshutdown -r -t now
```

Causes it to do the upgrade and takes about 20 minutes

HMC 778 is now apar MB03715 PTF MH01377

55



UPGRADING HMC FROM 7.7.7.0 TO 7.7.8

Once it is back up we can do the updates:

In the GUI select Updates, Update HMC

Server information is:

```
ftp.software.ibm.com
```

```
anonymous login with your email as password
```

```
/software/server/hmc/fixes
```

Or for service packs

```
/software/server/hmc/updates
```

Mandatory fix apar MB03754 PTF MH01388

REBOOT HMC

Then do MH01404 is latest update (requires MH01388) using same process as above

After the reboot put in a new USB stick (if that is how backup was done)

Save upgrade data and then backup to USB stick or FTP server using the GUI

DVD has been disabled at one of the versions so you now need to backup to an FTP server or the 8GB USB stick that you may have purchased with the server.

56



HMC v8

Required for POWER8

Runs on cr5 or C08 or higher

Will not run on earlier HMCs

Validates entitlement for POWER8

Introduces new Performance and Capacity Monitoring Task

Provides reports on resource utilization

NIST support – updates to JVM

LPM improvements to vSCSI performance

SR-IOV support

Dynamic partition remote restart can be changed when LPAR deactivated, not just at creation time

Absolute values for DLPAR

57



UPGRADING TO HMC v8

Check memory and hardware prereqs

i.e.no POWER5, etc

HMC must already be at v7.780 with fixpack MH01402 or HMC v7.7.9 prior to upgrade

PowerVM 2.2.3.0 is required for the new performance metrics

Check prereqs if using redundant HMCs

Process:

Back it up

Get the upgrade files

Reboot to upgrade to v8

Apply first mandatory PTF (can do via GUI)

Reboot

Repeat till you run out of fixes

Backup again after the last reboot

58



USEFUL HMC CLI COMMANDS

monhmc -r mem -n 0	shows total, used and free mamory of HMC
monhmc -r disk -n 0	shows filesystems and usage info (same as "df -k")
monhmc -r proc -n 0	shows cpu usage of each processor
monhmc -r swap -n 0	shows paging space usage
vtmenu	Get a console for an LPAR
getupgfiles -h ftp.software.ibm.com -u anonymous --passwd ftp -d /software/server/hmc/network/v8810	
chhmc -c altdiskboot -s enable --mode upgrade	Boot from install image to upgrade
hmcshutdown -r -t now	Reboot now
lshmc -V	Show HMC version
chhmcfs -o f -d 0	Clear out old logfiles
Lshmcfs	List HMC filesystems

59



HMC SCANNER

Latest HMC Scanner is available at <http://tinyurl.com/HMCscanner>

Java program that uses SSH to connect to HMC, FSM or IVM to gather information about the system configuration – latest is 0.11.17 as of January 2015

I run it on one of the AIX Systems as follows:

- ./hmcScanner.ksh servername hscroot -p password -stats --sanitize
- Sanitize causes it to produce two spreadsheets – one that has been cleansed of identifying data

Information is organized in tabs in an excel spreadsheet:

- System summary: name, serial number, cores, memory, service processor IP for each server
- LPAR Summary: list of all LPAR by serve with status, environment, version, processor mode
- LPAR CPU: processor configuration of each LPAR
- LPAR MEM: memory configuration of each LPAR
- Physical Slots: list of all slots of each system with LPAR assignment, description, physical location and drc_index
- Virtual Ethernet: network configuration of each virtual switch and each LPAR
- Virtual SCSI: configuration of all virtual SCSI adapters, both client and server
- VSCSI Map: devices mapped by each VIOS to partitions
- Virtual Fibre: virtual fibre channel configuration of client and server with identification of physical adapter assigned
- **SEA**: SEA configuration and statistics for all VIOS
- SW Cores: LPAR and virtual processor pool configuration matrix to compute the number of software licenses. Simulation of alternative scenarios is possible.
- CPU Pool Usage: monthly average history of CPU usage of each system. Based on last 12 months of lslparutil data.
- Sys RAM Usage: monthly average history of physical memory assignment to each LPAR. Based on last 12 months of lslparutil data.
- LPAR CPU Usage: monthly average history of CPU usage of each LPAR. Based on last 12 months of lslparutil data.
- **CPU Pool Daily Usage**: 1 year of CPU usage of every pool and subpools of each system. Based on daily averages.
- **LPAR Daily Usage**: 1 year of CPU usage of every LPAR of each system. Based on daily averages.
- **CPU Pool Hourly Usage**: 2 months of CPU usage of every pool and subpools of each system. Based on hourly averages.
- **LPAR Hourly Usage**: 2 months of CPU usage of every LPAR of each system. Based on hourly averages.

60



RUNNING HMC SCANNER

I run it from AIX as Windows and Java issues have caused problems

Right now I have HMCScanner11

```
./hmcScanner.ksh hmcname hscroot -p password -stats
```

hmcScanner version 0.11.0

Detecting manager type: HMC

Detecting managed systems: 3 systems present.

Starting managed system configuration collection:

Scanning p720-Server-8202-E4B-SN10934BP: DONE

Scanning p740-Server-8205-E6B-SN10934CP: DONE

Scanning p750-Server-8233-E8B-SN069348P: DONE

Collection successfully finished. Data is in /software/hmcscanner-11/bpichmc/

Performance data collection:

Loading p720-Server-8202-E4B-SN10934BP: . .

Loading p740-Server-8205-E6B-SN10934CP: . .

Loading p750-Server-8233-E8B-SN069348P: . .

..... DONE

61



FIRMWARE MAINTENANCE



ENTITLEMENT

- Starting with POWER8 IBM will be checking entitlement when applying firmware fixes.
- Entitlement requires an HWMA

POWER8 (and later) servers require machine code
“update entitlement at activation”

- POWER8 and later servers contain an “update access key” (UAK)

Machine code update entitlement is checked using the UAK at
each activation / installation


Entitlement check must pass before an update can proceed

- Entitlement is checked based on existing terms and conditions
- Security and safety fixes are exempt from the entitlement check



MAINTAINING YOUR ENVIRONMENT

- A good fix maintenance strategy is an important part of maintaining and managing your server. Regular maintenance of your server, and application of the latest fixes help to maximize server performance, and may reduce the impact of problems if they arise.
- It is recommended that all servers be kept on a supported release and current with latest available fix packages for HMC and server firmware fixes whenever possible.
- The most important scenario to avoid is remaining on a release so long that all subsequent releases that support a single-step upgrade are withdrawn from marketing. Without a single-step upgrade available, there are no supported ways for you to upgrade your server.



GENERAL FIRMWARE STRATEGIES

IBM releases new firmware for the following reasons:

- The addition of new system function.
- To correct or avoid a problem.



There are some natural points at which firmware should be evaluated for potential updates:

- ✓ When a subscription notice advises of a critical or HIPER (highly pervasive) fix, the environment should be reviewed to determine if the fix should be applied.
- ✓ When one of the twice-yearly updates is released.
- ✓ Whenever new hardware is introduced into the environment the firmware pre-reqs and co-reqs should be evaluated.
- ✓ Anytime HMC firmware levels are adjusted.
- ✓ Whenever an outage is scheduled for a system which otherwise has limited opportunity to update or upgrade.
- ✓ When the firmware level your system is on is approaching end-of-service.
- ✓ If other similar hardware systems are being upgraded and firmware consistency can be maximized by a more homogenous firmware level.
- ✓ On a yearly cycle if firmware has not been updated or upgraded within the last year.



ACCESS TO THE WEB

Have web access in computer room to access the fixes and documentation

Having a landline phone available to use for talking with support etc., it is helpful (what happens if your battery dies?)

Have access to documentation for a server somewhere OTHER than on the server (ESPECIALLY restore procedures!)

- Documentation
- Led codes
- Error Records
- Fixes
- TL



USEFUL COMMANDS

67



USEFUL COMMANDS

Command History

```
$ fc -l
725 lsrep
726 backupios -file /usr/local/backups/b750viobkp
727 exit
728 lsmmap -vadapter vhost0
729 fc -l
```

Global command log

```
$ lsgcl | grep "Aug 9 2013"
Aug 9 2013, 08:25:35 root ioslevel
Aug 9 2013, 08:59:22 padmin license
Aug 9 2013, 09:00:29 padmin lsmmap -vadapter vhost0
Aug 9 2013, 09:01:29 padmin lsgcl
```

Redirecting output when running as padmin
lsmmap -all -npiv | tee npivdata.txt

68



USEFUL COMMANDS

vSCSI Commands

```
mkvdev -vdev hdisk2 -vadapter vhost0
mkvdev -fbo -vadapter vhost0
```

NPIV

```
Setup NPIV mappings
vfcmmap -vadapter vfchost0 -fcp fcs0
lsmap -npiv -all
lsmap -vadapter vfchost0 -npiv
lsdev -virtual
lsnports
lsdev -slots
lscfg -vpl vfchost0
```

69



USEFUL COMMANDS

\$ lsdev -virtual

name	status	description
ent5	Available	Virtual I/O Ethernet Adapter (I-lan)
ent6	Available	Virtual I/O Ethernet Adapter (I-lan)
ent7	Available	Virtual I/O Ethernet Adapter (I-lan)
vasi0	Available	Virtual Asynchronous Services Interface (VASI)
vbsd0	Available	Virtual Block Storage Device (VBSD)
vfchost0	Available	Virtual FC Server Adapter
vfchost1	Available	Virtual FC Server Adapter
vhost0	Available	Virtual SCSI Server Adapter
vhost1	Available	Virtual SCSI Server Adapter
vsa0	Available	LPAR Virtual Serial Adapter
b740ios1_rv1	Available	Virtual Target Device - Logical Volume
b740l1_rv1	Available	Virtual Target Device - Logical Volume
vtopt0	Available	Virtual Target Device - File-backed Optical
vtopt1	Available	Virtual Target Device - File-backed Optical
vtscsi0	Available	Virtual Target Device - Disk
vtscsi1	Available	Virtual Target Device - Disk
vtscsi2	Available	Virtual Target Device - Disk
vtscsi3	Available	Virtual Target Device - Disk
ent8	Available	Shared Ethernet Adapter

70



USEFUL COMMANDS

\$ Ismap -vadapter vhost0

SVSA	Physloc	Client Partition ID
vhost0	U8205.E6B.10934CP-V1-C21	0x00000003
VTD	b74011_rv1	
Status	Available	
LUN	0x8300000000000000	
Backing device	lv_b74011	
Physloc		
Mirrored	N/A	
VTD	vtopt0	
Status	Available	
LUN	0x8200000000000000	
Backing device		
Physloc		
Mirrored	N/A	
VTD	vtopt1	
Status	Available	
LUN	0x8100000000000000	
Backing device		
Physloc		
Mirrored	N/A	

71



USEFUL COMMANDS

\$ Ismap -vadapter vfchost0 -npiv

Name	Physloc	ClntID	ClntName	ClntOS
vfchost0	U8205.E6B.10934CP-V1-C31		3	

Status:NOT_LOGGED_IN
 FC name:fcs0 FC loc code:U78AA.001.WZSG8PD-P1-C5-T1
 Ports logged in:0
 Flags:4<NOT_LOGGED>
 VFC client name: VFC client DRC:

\$ Ismap -vadapter vfchost4 -npiv

Name	Physloc	ClntID	ClntName	ClntOS
vfchost4	U8205.E6B.10934CP-V1-C36	8	b740nl1	AIX

Status:LOGGED_IN
 FC name:fcs0 FC loc code:U78AA.001.WZSG8PD-P1-C5-T1
 Ports logged in:3
 Flags:a<LOGGED_IN,STRIP_MERGE>
 VFC client name:fcs0 VFC client DRC:U8205.E6B.10934CP-V8-C36

72



USEFUL COMMANDS

\$ Isnports

name	physloc	fabric	tports	aports	swwpns	awwpns
fcs0	U78AA.001.WZSG8PD-P1-C5-T1	1	64	63	2048	2041

\$ Isdev -slots

# Slot	Description	Device(s)
HEA 1	Logical I/O Slot	lhea0 ent0
U8205.E6B.10934CP-V1-C0	Virtual I/O Slot	vsa0
U8205.E6B.10934CP-V1-C11	Virtual I/O Slot	ent5
U8205.E6B.10934CP-V1-C12	Virtual I/O Slot	ent6
U8205.E6B.10934CP-V1-C13	Virtual I/O Slot	ent7
U8205.E6B.10934CP-V1-C21	Virtual I/O Slot	vhost0
U8205.E6B.10934CP-V1-C22	Virtual I/O Slot	vhost1
U8205.E6B.10934CP-V1-C23	Virtual I/O Slot	vhost2
U8205.E6B.10934CP-V1-C31	Virtual I/O Slot	vfchost0
U8205.E6B.10934CP-V1-C32	Virtual I/O Slot	vfchost1
U8205.E6B.10934CP-V1-C33	Virtual I/O Slot	vfchost2
U8205.E6B.10934CP-V1-C32769	Virtual I/O Slot	vasi0
U8205.E6B.10934CP-V1-C32773	Virtual I/O Slot	vasi1
U8205.E6B.10934CP-V1-C32774	Virtual I/O Slot	vasi2
U8205.E6B.10934CP-V1-C32775	Virtual I/O Slot	vasi3
U8205.E6B.10934CP-V1-C32776	Virtual I/O Slot	vasi4

73



USEFUL HMC COMMANDS

74



USEFUL HMC COMMANDS

```
hscroot@bpichmc:~>|shmc -b
```

```
"bios=D6E149AUS-1.09
```

```
"
```

```
hscroot@bpichmc:~>|shmc -r
```

```
ssh=enable,sshprotocol=,remotewebui=enable,xntp=disable,xntpserver=127.127.1.0,syslogserver=,syslogtcpserver=,syslogtlserver=,altdiskboot=disable,ldap=disable,kerberos=disable,kerberos_default_realm=,kerberos_realm_kdc=,kerberos_clockskew=,kerberos_ticket_lifetime=,kpasswd_admin=,trace=,kerberos_keyfile_present=,kerberos_allow_weak_crypto=,legacyhmccomm=disable,security=legacy,sol=disabled
```

```
hscroot@bpichmc:~>|shmc -e
```

```
emch=disabled,callhome=enabled,registered_hmcs=
```

On HMC check LMB sizes

```
hscroot@bpichmc:~>|shwres -r mem -m p740-Server-8205-E6B-SN10934CP --level sys -F mem_region_size
256
```

75



USEFUL HMC COMMANDS – HMC UPDATES

ssh to HMC as hscroot or your userid

Use with great care

```
saveupgdata -r disk
```

```
getupgfiles -h public.dhe.ibm.com -u anonymous --passwd anonymous -d /software/server/hmc/network/v8810
```

```
ls -la /hmcdump
```

```
chhmc -c altdiskboot -s enable --mode upgrade
```

```
tail -f /tmp/HmcInstall.log           during upgrade
```

Directories on FTP Server (ftp.software.ibm.com)

Upgrades: /software/server/hmc/network/v8810

Fixes: /software/server/hmc/fixes

Service Packs: /software/server/hmc/updates

76



USEFUL HMC COMMANDS

ssh to HMC as hscroot or your userid

Useful Commands:

lshmc

vtmenu - way better than ascii console

lshwres

monhmc -r mem -n 0 how much memory do I have?

monhmc -r proc -n 0 CPU usage

monhmc -r swap -n 0 Page space

monhmc -r disk -n 0 What is my disk utilization?

chmcfscfs -r disk -n 0 Clear out all temp files

lshmcfs

hmcshutdown -r -t now Reboot HMC

77



USEFUL HMC COMMANDS – 7042-CR6

hscroot@bpichmc:~>monhmc -r mem -n 0

Mem: 4043216k total, 3885308k used, 157908k free, 484132k buffers (has 4GB)

hscroot@bpichmc:~>monhmc -r proc -n 0

Cpu0 : 0.0%us, 0.7%sy, 0.0%ni,98.3%id, 1.0%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 : 0.0%us, 0.0%sy, 0.0%ni,100.0%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

hscroot@bpichmc:~>monhmc -r swap -n 0

Swap: 2040244k total, 137456k used, 1902788k free, 1036824k cached

hscroot@bpichmc:~>monhmc -r disk -n 0

Filesystem	1K-blocks	Used	Available	Use%	Mounted on
/dev/sda2	16121184	7100064	8202208	47%	/
/dev/sda3	6040320	297672	5435808	6%	/var
/dev/mapper/HMCDDataVG-HomeLV	10321208	245052	9551868	3%	/home
/dev/mapper/HMCDDataVG-LogLV	8256952	1292372	6545152	17%	/var/hsc/log
/dev/mapper/HMCDDataVG-DumpLV	123854820	319672	117243692	1%	/dump
/dev/mapper/HMCDDataVG-ExtraLV	20642428	198692	19395160	2%	/extra
/dev/mapper/HMCDDataVG-DataLV	227067260	455376	215077548	1%	/data

hscroot@bpichmc:~>lshmcfs

filesystem=/var,filesystem_size=8063,filesystem_avail=6390,temp_files_start_time=07/14/2014 13:11:00,temp_files_size=783
 filesystem=/dump,filesystem_size=120951,filesystem_avail=114495,temp_files_start_time=07/14/2014 21:09:00,temp_files_size=0
 filesystem=/extra,filesystem_size=20158,filesystem_avail=18940,temp_files_start_time=none,temp_files_size=0
 filesystem=/,filesystem_size=15743,filesystem_avail=8009,temp_files_start_time=07/22/2014 23:18:00,temp_files_size=0

78



USEFUL HMC COMMANDS – 7042-CR7

hscroot@bpic8hmc:-->monhmc -r mem -n 0

Mem: 41263576k total, 3608896k used, 37654680k free, 551600k buffers
Either it has 41GB memory or there is a bug @

hscroot@bpic8hmc:-->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, 0.3%sy, 0.0%ni,99.7%id, 0.0%wa, 0.0%hi, 0.0%si, 0.0%st
Cpu2 : 0.0%us, 0.0%sy, 0.0%ni,100.0%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
Cpu4 : 0.0%us, 0.0%sy, 0.0%ni,100.0%id, 0.0%wa, 0.0%hi, 0.0%si, 0.0%st
Cpu5 : 0.0%us, 0.0%sy, 0.0%ni,100.0%id, 0.0%wa, 0.0%hi, 0.0%si, 0.0%st

hscroot@bpic8hmc:-->monhmc -r swap -n 0

Swap: 2040244k total, 0k used, 2040244k free, 934024k cached

hscroot@bpic8hmc:-->monhmc -r disk -n 0

Filesystem	1K-blocks	Used	Available	Use%	Mounted on
/dev/sda2	16121184	6715032	8587240	44%	/
/dev/sda3	6040320	270112	5463368	5%	/var
/dev/mapper/HMCDDataVG-HomeLV	10321208	244856	9552064	3%	/home
/dev/mapper/HMCDDataVG-LogLV	8256952	479796	7357728	7%	/var/hsc/log
/dev/mapper/HMCDDataVG-DumpLV	61927420	187024	58594668	1%	/dump
/dev/mapper/HMCDDataVG-ExtraLV	20642428	198692	19395160	2%	/extra
/dev/mapper/HMCDDataVG-DataLV	144497320	195428	136961860	1%	/data

hscroot@bpic8hmc:-->lshmcfs

filesystem=/var,filesystem_size=8063,filesystem_avail=7185,temp_files_start_time=07/14/2014 16:33:00,temp_files_size=318
filesystem=/dump,filesystem_size=60475,filesystem_avail=57221,temp_files_start_time=07/14/2014 20:15:00,temp_files_size=0
filesystem=/extra,filesystem_size=20158,filesystem_avail=18940,temp_files_start_time=none,temp_files_size=0
filesystem=/,filesystem_size=15743,filesystem_avail=8385,temp_files_start_time=07/22/2014 22:43:00,temp_files_size=0

79



USEFUL HMC COMMANDS

lshmc

- V - Displays HMC version information.
- v - Displays HMC VPD information.
- r - Displays HMC remote access settings.
- n - Displays HMC network settings.
- b - Displays the BIOS level of the HMC.
- l - Displays the current locale for the HMC.
- L - Displays all supported locales for the HMC.
- h - Displays HMC hardware information.
- i - Displays HMC Integrated Management Module (IMM) settings.
- e - Displays HMC settings for Events Manager for Call Home.
- F [<attribute names>] - delimiter-separated list of the names of the attributes to be listed for the specified HMC setting. If no attribute names are specified, then all attributes will be listed.
- header - prints a header of attribute names when -F is also specified
- help - prints this help

80



USEFUL HMC COMMANDS

ssh to HMC as hscroot or your userid

```
hscroot@bpichmc:~>lshmc -V
*version= Version: 8
*Release: 8.1.0
*Service Pack: 0
HMC Build level 20140602.3
MH01421: Required fix for HMC V8R8.1.0 (06-03-2014)
MH01436: Fix for OpenSSL, GnuTLS (06-11-2014)
MH01441: Fix for HMC V8R8.1.0 (06-23-2014)
", "base_version=V8R8.1.0
"
```

```
hscroot@bpichmc:~>lshmc -v
*vpd=*FC ????????
*VC 20.0
*N2 Wed Jul 23 04:45:57 UTC 2014
*FC ????????
*DS Hardware Management Console
*TM 7042-CR6
*SE 102EEEC
*MN IBM
*PN 0B20PT
*SZ 4140253184
*OS Embedded Operating Systems
*NA 10.250.134.20
*FC ????????
*DS Platform Firmware
*RM V8R8.1.0.0
"
```

81



**VIOS ADVISOR
NOW THE PART
COMMAND**



82



PART COMMAND

VIOS Advisor is now the part command

<http://www-01.ibm.com/support/knowledgecenter/POWER7/p7hcg/part.htm?cp=POWER7%2F1-8-3-8-2-142>

part -? Provides usage help

part -i interval

i.e. part -l 30 produces a 30 minute file and shows the name of the tar file created

Download and extract the tar file

Then open the vios_advisor.xml file in a browser to see the results

Or you can point it at a .nmon file:

part -f filename.nmon

Sample reports can be found at:

http://www-01.ibm.com/support/knowledgecenter/POWER7/p7hb1/iphb1_vios_perf_adv_reports.htm?cp=POWER7%2F1-8-3-5-8-8-0

83



VIOS ADVISOR

The ratings and recommendations in the table below were chosen with the following information:

Hostname: vio1. .com

PartitionID: 2

Monitoring Start Time: 03/09 11:45:19

Monitoring Stop Time: 03/09 13:45:19 **Duration:** 120 min

IBM Systems Workload Estimator link: <http://ibm.com/systems/support/tools/estimator> (VIOS Sizings)

SYSTEM - CONFIGURATION	
Name	Value
Processor Family	POWER6
Server Model	IBM,9117-MMA
Server Frequency	4.208 GHz
Server - Online CPUs	10 cores
Server - Maximum Supported CPUs	16 cores
VIOS Level	2.2.0.13-FP24 SP-03
VIOS Advisor Release	121211B

84





VIOS ADVISOR

VIOS - CPU							
	Name	Measured Value	Recommended Value	First Observed	Last Observed	Risk 1=lowest 5=highest	Impact 1=lowes 5=highes
✔	CPU Capacity	1.0 ent	-	03/09 11:45:19	-	n/a	n/a
🔒	CPU Consumption	avg:5.4% (cores:0.1) high:40.2% (cores:0.5)	-	-	-	n/a	n/a
🔒	Processing Mode	Shared CPU, (UnCapped)	-	03/09 11:45:19	-	n/a	n/a
✔	Variable Capacity Weight	200	-	03/09 11:45:19	-	n/a	n/a
✔	Virtual Processors	2 vCPUs	-	03/09 11:45:19	-	n/a	n/a
✔	SMT Mode	SMT2	-	03/09 11:45:19	-	n/a	n/a

VIOS ADVISOR

SYSTEM - SHARED PROCESSING POOL							
	Name	Measured Value	Recommended Value	First Observed	Last Observed	Risk 1=lowest 5=highest	Impact 1=lowes 5=highes
✔	Shared Pool Monitoring	enabled	-	03/09 11:45:19	-	n/a	n/a
🔒	Shared Processing Pool Capacity	10.0 ent.	-	03/09 11:45:19	-	n/a	n/a
✔	Free CPU Capacity	avg_free:9.4 ent. lowest_free:7.7 ent.	-	-	-	n/a	n/a

VIOS ADVISOR

VIOS - I/O ACTIVITY	
Name	Value
 Disk I/O Activity	avg: 229 iops @ 32KB peak: 1916 iops @ 137KB
 Network I/O Activity	[avgSend: 0 iops 0.0MBps , avgRcv: 0 iops 0.0MBps] [peakSend: 0 iops 0.0MBps , peakRcv: 0 iops 0.0MBps]

87






VIOS ADVISOR

VIOS - DISK ADAPTERS							
	Name	Measured Value	Recommended Value	First Observed	Last Observed	Risk 1=lowest 5=highest	Impact 1=lowest 5=highest
	FC Adapter Count	3	-	03/09 11:45:19	-	n/a	n/a
	FC Avg IOps	avg: 77 iops @ 32KB	-	03/09 11:45:19	03/09 13:45:19	n/a	n/a
	FC Adapter Utilization	optimal	-	-	-	n/a	n/a
	FC Port Speeds	running at speed	-	-	-	n/a	n/a

88









VIOS ADVISOR

VIOS - DISK DRIVES							
	Name	Measured Value	Recommended Value	First Observed	Last Observed	Risk 1=lowest 5=highest	Impact 1=lowest 5=highest
	Physical Drive Count	93	-	03/09 11:45:19	-	n/a	n/a
	I/Os Blocked	optimal	-	-	-	n/a	n/a
	Long I/O Latency (hdisk3)	avg:9.7ms (9.7 + 0.0) high:11.5ms (11.5 + 0.0)	Range: 8-12ms	03/09 12:35:58	03/09 13:44:02	n/a	n/a

89



VIOS ADVISOR

VIOS - MEMORY							
	Name	Measured Value	Recommended Value	First Observed	Last Observed	Risk 1=lowest 5=highest	Impact 1=lowest 5=highest
	Real Memory	4.000 GB	7.000 GB	03/09 11:45:19	-	1	2
	Available Memory	0.346 GB	1.5 GB Avail.	03/09 11:45:39	03/09 13:45:05	n/a	n/a
	Paging Rate	0.2 MB/s pg rate	-	-	-	n/a	n/a
	Paging Space Size	8.000 GB	-	03/09 11:45:19	-	n/a	n/a
	Free Paging Space	7.923 GB free	-	-	-	n/a	n/a
	Pinned Memory	1.262 GB pinned	-	-	-	n/a	n/a

90



THANK YOU FOR YOUR TIME



If you have questions please email me at:
lynchj@forsythe.com

Also check out:
<http://www.circle4.com/forsythetalks.html>
<http://www.circle4.com/movies/>

Handout at:
[Http://www.circle4.com/forsythe/viosmaint-may2015.pdf](http://www.circle4.com/forsythe/viosmaint-may2015.pdf)

REFERENCES



USEFUL LINKS

Charlie Cler Articles

- <http://www.ibmssystemsmag.com/authors/Charlie-Cler/>

Jaqui Lynch Articles

- <http://www.ibmssystemsmag.com/authors/Jaqui-Lynch/>
- <https://enterprisesystemsmedia.com/author/jaqui-lynch>
- <http://ourdigitalmags.com/publication/?i=251675&p=66>

Jay Kruemke Twitter – chromeaix

- <https://twitter.com/chromeaix>

Nigel Griffiths Twitter – mr_nmon

- https://twitter.com/mr_nmon

Gareth Coates Twitter – power_gaz

- https://twitter.com/power_gaz

Jaqui's Upcoming Talks and Movies

- Upcoming Talks
 - <http://www.circle4.com/forsyhetalks.html>
- Movie replays
 - <http://www.circle4.com/movies>

IBM US Virtual User Group

- <http://www.tinyurl.com/ibmaixvug>

Power Systems UK User Group

- <http://tinyurl.com/PowerSystemsTechnicalWebinars>

93



USEFUL LINKS

AIX Wiki

- <https://www.ibm.com/developerworks/wikis/display/WikiPtype/AIX>

HMC Scanner

- <http://www.ibm.com/developerworks/wikis/display/WikiPtype/HMC+Scanner>

Workload Estimator

- <http://ibm.com/systems/support/tools/estimator>

Performance Tools Wiki

- <http://www.ibm.com/developerworks/wikis/display/WikiPtype/Performance+Monitoring+Tools>

Performance Monitoring

- <https://www.ibm.com/developerworks/wikis/display/WikiPtype/Performance+Monitoring+Documentation>

Other Performance Tools

- <https://www.ibm.com/developerworks/wikis/display/WikiPtype/Other+Performance+Tools>

• Includes new advisors for Java, VIOS, Virtualization

VIOS Advisor

- <https://www.ibm.com/developerworks/wikis/display/WikiPtype/Other+Performance+Tools#OtherPerformanceTools-VIOSPA>

Capturing Debug output for padmin

- <http://www-01.ibm.com/support/docview.wss?uid=isg3T1012362>

94



VIOS SPECIFIC REFERENCES

SDD and SDDPCM Specific procedures for VIOS

- <http://www-01.ibm.com/support/docview.wss?uid=ssg1S7002686&aid=1>

VIOS 2.2.3.4 FP27 SP04 README

- <http://www-01.ibm.com/support/docview.wss?uid=isg400001971>

SG24-7940 - PowerVM Virtualization - Introduction and Configuration

- <http://www.redbooks.ibm.com/redbooks/pdfs/sg247940.pdf>

SG24-7590 – PowerVM Virtualization – Managing and Monitoring

- <http://www.redbooks.ibm.com/redbooks/pdfs/sg247590.pdf>

SG24-8080 – Power Systems Performance Guide – Implementing and Optimizing

- <http://www.redbooks.ibm.com/redbooks/pdfs/sg248080.pdf>

SG24-8062 – PowerVM Best Practices

- <http://www.redbooks.ibm.com/redbooks/pdfs/sg248062.pdf>

POWERVM Enhancements – what is new in 2013

- <http://www.redbooks.ibm.com/redbooks/pdfs/sg248198.pdf>

Capturing Debug output for padmin

- <http://www-01.ibm.com/support/docview.wss?uid=isg3T1012362>