

VIOS **TOP** Tips for Administration



PowerVM Virtual User Group May 19, 2021 jaqui@circle4.com or jlynch@flagshipsg.net

http://www.circle4.com/ptechu/viostoptips-may192021.pdf

Top 10 Tips – (OK there are 12, maybe more)

Stay current – upgrade regularly 1.

- If you don't the upgrades become very complex and take longer 1.
- 2. Install Hiper fixes and fixes to java, ssh, ssl
- 3. Keep it simple
- 4. Document well
- 5. Use dual vio so you get concurrent update
- 6. Use vio commands for maintenance
- emgr for HIPERs is the only exception 1
- 7. Backup regularly
 - 1. New backup function on HMC950 to backup VIO to HMC Give vio plenty of resources
- 8.
- Entitlement, memory, HBA queues, virtual buffers 9. Make sure storage is zoned and mapped for LPM
- Always have a second disk available for updates and quick failback

 alt_disk_copy and alt_disk_mksysb are your friends
 Use NTP to get consistent time

- 12. Don't forget to update I/O firmware (pay attention to patch order)
- Article from 2019 on Using NIM with VIO Servers
 - https://techchannel.com/SMB/09/2019/Using-NIM-with-VIO-Servers
 - Or http://www.circle4.com/jaqui/eserver/usingnimwithVIO-sep2019.pdf

Patch order

- 1. Read the readmes
- 2. Usually do HMC first
- 3. Then server firmware
- 4. NIM server (it should be standalone and needs to be at highest level) plus I/O firmware
- 5. Then VIO servers and any I/O firmware
- 6. LPARs AIX, IBM i, Linux
- 7. Note I document everything
- 8. I write up every step of the upgrade before I do it and then tweak as I go along



Version	Recommended Update	Recommended Upgrade	🕴 Release Date	EoSPS Date	
2.2.6.32	2.2.6.65	3.1.1.39	2018,11.16	2020,09,30	
2.2.6.41	2.2.6.65	3.1.1.30	2019.05.08	2020.09.30	
2.2.6.51	2.2.6.65	3.1.1.30	2019.12.13	2020.09.30	
2.2.6.61	2.2.6.65	3.1.1.30	2020.05.15	2020.09.30	
2.2.6.65		3.1.1.30	2020.07.17	2020.09.30	NOTE all levels prior to 2.2.5 are EOS
3.1.0.10	3.1.0.30	3.1.1.30	2018.11.09	2021.11.30	as of December 2018 2 2 5 went end of service 9/30/2019
3.1.0.21	3.1.0.30	3.1.1.30	2019.05.08	2021.11.30	2.2.6 is end of service 9/30/2020 –
3.1.0.30	3.1.0.40	3.1.1.30	2020.05.15	2021.11.30	All levels prior to 3.1.2.10 go EOSPS by
3.1.0.40		3.1.1.30	2020.07.31	2021.11.30	11/30/2021
3.1.0.50		3.1.1.30	2021.02.12	2021.11.30	
3.1.1.10	3.1.1.25		2019.11.15	2021,11,30	
3.1.1.21	3.1.1.25		2020.05.15	2021.11.30	
3.1.1.25			2020.07.17	2021.11.30	
3.1.1.30			2021.02.12	2021.11.30	
3.1.2.10	3.1.2.21		2020.11.13	2023.11.30	
3.1.2.21			2021.04.16	2023.11.30	5



VIOS levels VIOS 3.1 came out 11/9/2018 plus a minipack 3.1.0.10 and then 3.1.0.21 in May 2019 Latest is 3.1.2.21 which is AIX 7.2.5.2 based as of 4/6/2021 You can install directly from the flash copy which is at 3.1.2.21. You can also use this copy to upgrade directly from VIOS 3.1.2.0 There is a single step process documented for upgrading from 3.1.0.0 to 3.1.2.21 If you are running 2.2.6 the VIO must be upgraded to v3.1.0 before putting on this update Download 3.1.2.21 flash copy from entitled software: https://www.ibm.com/servers/eserver/ess/index.wss Current Flash copy is called: Virtual IO Server Base Install 3.1.2.21 Flash 042021 LCD8250307.iso You can download the 3.1.2.21 update with all prerequisites from Fix Central: http://www-933.ibm.com/support/fixcentral/ Release Notes for 3.1.2.21 Service Pack - https://www.ibm.com/support/pages/node/6443445 NIM Master needs to be at 7200-05-02 at a minimum for v3.1.2.21 Check required HMC and firmware levels Minimum server level is POWER7+ (D model) and above Service strategy: http://www-304.ibm.com/webapp/set2/sas/f/vios/svcstrategy.html Lifecycle: http://www-01.ibm.com/support/docview.wss?uid=isg3T1023504 When installing a new server read the redbook to ensure your VIO level, HMC, etc are supported As an example the E980 requires a minimum of VIO 2.2.6.31, S922 and S924 are 2.2.6.21 S950 is 2.2.6.23 - these are MINIMUMs

7

VIO Server 3.1.2.21 Hiper

VIOS 3.1.2.10 had 9 HIPER fixes to put on – there is only one for 3.1.2.21 It is IJ31191 also known as IJ31936 <u>https://www.ibm.com/support/pages/apar/IJ31191</u> This patch resolved issues around NPIV I/O errors iFix is at: <u>https://aix.software.ibm.com/aix/ifixes/ij31191/</u>

All VIOS Hipers: http://www14.software.ibm.com/webapp/set2/flrt/doc?page=hiper&os=vios_hiper

Just added: IJ31604 <u>https://www.ibm.com/support/pages/apar/IJ31604</u> Zoning changes in Cisco MDS may cause 7.2.5.1 client to lose connectivity







Efixes and ifixes

Many security patches are put on using efixes or ifixes The VIO server also needs these to be applied – **use FLRTVC to determine what fixes are needed** Run flrtvc and download and install the ifixes that are needed <u>https://www-304.ibm.com/webapp/set2/sas/f/flrt/flrtvc.html</u> You should do this on AIX LPARs too

/usr/sbin/emgr –I lists them emgr –P lists the patches and the packages they affect To apply a fix change into the directory it is in and then run it in preview mode: cd /usr/local/soft/vios31fixes/bind_fix17 emgr -p -e IJ25927s2a.200708.epkg.Z Remove the –p and run again for real if the preview was successful: emgr -e IJ25927s2a.200708.epkg.Z

If you run emgr –I and there are no fixes listed then you most likely have security holes that need patching, specifically Java, openssh and openssl.

Efixes and ifixes	
For 3.1.1.25 as of 9/14/2020 one patch is needed once SSH, SSL and Java are updated:ij25927 ftp://aix.software.ibm.com/aix/efixes/security/bind_fix17.tar	
After you untar the downloaded file you can check the efix name to put on by changing into the director #grep 3.1.1.25 *.asc 3.1.1.25 IJ25927s2a.200708.epkg.Z key_w_fix	y and then:
# emgr -l ID_STATE LABELINSTALL TIMEUPDATED BY ABSTRACT	
Image: second	
It will vary by O/S level and SP. This was for 3.1.1.25	
To remove an efix or ifix: # /usr/sbin/emgr -r -L <efix label=""> emgr -r -L IJ16586s3a</efix>	
Running flrtvc.ksh on 9/14/2020 showed no other efixes to go on	13





- If your upgrade goes south
- Know how to get a perfpmr
 - <u>https://www.ibm.com/support/pages/readme-collecting-perfpmr</u>
 - Keep an up to date copy downloadd so you don't have to download it in a panic
- Know how to run snap
 - <u>https://www.ibm.com/support/pages/how-collect-snap-powervm-virtual-io-server-vios</u>
 - <u>https://www.ibm.com/support/pages/node/670105</u>
 - Uploading snap to Testcase
 - https://www.ibm.com/support/pages/node/681391
- Opening a support ticket
 - <u>https://www.ibm.com/mysupport/s/my-cases</u>
- Have your VIO predocumented as you will have to provide lots of information



3. Keep it simple and consistent

15

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
- Check VLANs on trunk adapters match between the 2 VIO servers that are paired
 - Second VIO server won't boot if they don't match
- Don't mix multipath drivers on 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)
- At least two VIO servers, but can also separate VIOs for production and non prod, or network from storage on large systems
- Test failover (SEA failover and disk if VIO goes down)

General

• Use VIO commands wherever possible rather than going into oem_setup_env and using smitty

- Mirror VIO rootvg if on internal disk
- Have a spare disk in your VIO to use for cloning prior to updates
- NOTE v2 requires at LEAST 30GB in rootvg give it 100GB
- Fix Paging- By default VIO has a 512MB hd6 and a 1.5GB paging00 on the same LUN
 - On some systems it is 2 x 1GB page spaces
- Add logging and set up dump devices properly
- Run VIOS Advisor (part) 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



Documentation is Critical

- Use HMCScanner and Sysplan
- Put together spreadsheet of documentation
- All equipment and serial numbers
- UAK expiration dates on servers
- Customer numbers
- Server and I/O firmware levels, VIO and O/S levels
- HMC information including version, BMC and PNOR, networking etc
- IP addresses
- Resource profiles
- Adapter allocations
- Standards used for network, vSCSI, NPIV mappings
- Actual vSCSI assignments
- Actual NPIV vfcmaps
- Vfchosts and their associated WWPNs
- SEA and virtual ethernet VLAN assignments
- Switchports for SAN and network
- Power needs and PDU mapping
- Anything else you can think of

19



Dual VIO

• Use dual VIO to allow for concurrent updates

• If properly setup you can update one vio and reboot it while the other is managing the environment

- Test failover (SEA failover and disk if VIO goes down)
- Requires a split backplane if you are using internal disks
- Provides greater redundancy and flexibility

• Can have one pair managing everything or can have a pair for network and a pair for disk I/O



Use VIO commands For updates use updateios updateios -accept -install -dev /usr/local/soft/vios31221 For upgrades (v2 to v3) use viosupgrade If you need to update java, ssh and ssl Put patches in their own directory and then use updateios against that directory updateios -accept -install -dev /usr/local/soft/javasshssl-vio-mar162021 Exceptions emgr for hiper patches emgr -p -e IJ31604s2a.210414.AIX72TL05SP02.epkg.Z I/O firmware updates which are done using diag



Use al	t_disk_copy prior	r to changes				
 Have two disks so you can take a clone If rootvg is mirrored you will need to unmirror for maintenance No need to mirror if on SAN, always mirror rootvg if on internal disk 						
# lspv gro hdisk0 hdisk1	ep root 00ce48c008314b9f 00ce48c03c8f2115	rootvg active altinst_rootvg				
# bootinfo hdisk0	-b					
exportvg a alt_disk_co	ltinst_rootvg opy −V −B −d hdisk1					
l always do Recovery i	o a bosboot and rewrite the f issues with the upgrade is	e bootlist before any reboot to point the bootlist to the new disk and reboot				
		25				

Table 1. Backup and restoration n	nethods for the VIOS	
Backup method	Media	Restoration method
To tape	Таре	From tape
To DVD	DVD-RAM	From DVD
To remote file system	nim_resources.tar image	From an HMC using the Network Installation Management (NIM) on Linux facility and the installios command
To remote file system	mksysb image	From an AIX 5L™ NIM server and a standard mksysb system installation
Tivoli Storage Manager	mksysb image	Tivoli Storage Manager

Sa	ckup Script to put in crontab
±	
#	‡
r	nachine=`uname -n`
r	nount /usr/local/backups
r	nkdir /usr/local/backups/\$machine
ι	umount /var/vio/VMLibrary
s	su - padmin -c "ioscli backupios -file /usr/local/backups/\$machine -nomedialib"
s	su - padmin -c "ioscli backupios -file /usr/local/backups/vio-mksysbs/\$machine.mksysb -nomedialib -mksysl
r	nount /var/vio/VMLibrary
#	ŧ
e	exit 0
Ν	NOTES
Т	The above can be put in root's crontab to run regularly
0	Don't forget to set up an NFS mount to the VIO from your NIM or NFS server
C	Do not allow ANY NFS mount to mount automatically at boot in case the NIM or NFS server is down at the time of boot
A	Also, regularly grab an HMCScanner report
	https://www.ibm.com/support/pages/hmc-scanner-power-server-config-and-performance-stats
	https://www.ipin.com/support/pages/sites/uerauit/files/ifiline-files/sfitz/fiffitscafiner-0.11.42.2ip







Your vio is your most loved one

- Give it the highest weight (I like 254)
- Make sure you give it plenty of entitlement (more than it needs)
- Tune the virtual buffers
- Tune the HBA settings (num_cmd_elems and max_xfer_size)
- Monitor with nmon and part





Sample /etc/tunables/rc-tunebufs.sh

This tunes buffer settings for the two virtual adapters – assumes ent4, ent5 are virtuals

Isdev –C | grep ent will show the adapters so you can pick the right ones

#! /bin/ksh

#
chdev -l ent4 -a buf_mode=min -P
chdev -l ent5 -a buf_mode=min -P
chdev -l ent4 -a max_buf_tiny=4096 -P
chdev -l ent4 -a max_buf_small=4096 -P
chdev -l ent5 -a max_buf_tiny=4096 -P
chdev -l ent5 -a max_buf_small=4096 -P
chdev -l ent5 -a max_buf_small=4096 -P
chdev -l ent5 -a max_buf_small=4096 -P

33

HBA Tuning Make the same tuning changes you would make on AIX, but VIO must be set at least as high as clients 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 · VIO must be rebooted to at least the client value prior to client change. 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: <u>http://www-01.ibm.com/support/docview.wss?uid=isg1IV63282</u> Increased again to 2048 in July 2016 http://www-01.ibm.com/support/docview.wss?uid=isg1IV76270 This upper limit is set in the client LPAR not the VIO server BUT the client setting MUST NOT be larger than what is set in the VIO server 34

HBA max_xfer_size					
The default is 0x100000 /* Default io_dma of 16MB */					
After that, 0x200000,0x400000,0x80000 gets you 128MB					
After that 0x1000000 checks for bus type, and you may get 256MB, or 128MB					
There are also some adapters that support very large max_xfer sizes which can possibly allocate	512MB				
VFC adapters inherit this from the physical adapter (generally)					
Unless you are driving really large IO's, then max_xfer_size on the HBA is rarely changed beyond which provides a 128MB DMA	0x200000				
Client setting cannot be higher than the VIOs were booted with					
	35				



Zoning and Mapping

• Zoning

- This is when the switch is configured to allow the switch port to talk to the storage and the WWPN for the LPAR or server
- Mapping (masking)
 - This is when the storage is updated to allow the host (LPAR or server) WWPNs access to the specific LUNs provisioned
- LUNs must be provisioned at the storage, then mapped and zoned before they can be used in an LPAR
- For direct attach we zone and map the WWNs for the real adapters, for NPIV we use the WWPNs on the virtual adapters
- WWNs tend to start with 10 or 20
- WWPNs (NPIV) start with CO
- These can be found in an HMCScanner report or by logging onto the LPAR or VIO or from the HMC
- Check the VIO connection to the switch is NPIV enabled:

ame	physloc	fabric	tports	aports	swwpns	awwpns
cs0	U78C9.001.WZS0234-P1-C12-T1		64	63	2048	2046
cs1	U78C9.001.WZS0234-P1-C12-T2		64	63	2048	2046
cs2	U78C9.001.WZS0234-P1-C6-T1		64	56	3088	3062
Ecs3	U78C9.001.WZS0234-P1-C6-T2		64	56	3088	3062



<section-header><list-item><list-item><list-item><list-item><list-item><list-item><list-item><list-item><list-item><list-item><list-item><list-item><list-item><list-item><list-item><list-item><list-item>



As padmin:

\$rmdev -dev fcs4 -recursive -ucfg \$rmdev -dev fcs5 -recursive -ucfg \$chdev -dev fscsi4 -attr autoconfig=defined \$chdev -dev fscsi5 -attr autoconfig=defined After you should see:

# Isdev -C	grep Iscsi							
fscsi0	Available	00-00-01	FC	SCSI	I/0	Controller	Protocol	Device
fscsil	Available	00-01-01	FC	SCSI	I/0	Controller	Protocol	Device
fscsi2	Available	01-00-01	FC	SCSI	I/0	Controller	Protocol	Device
fscsi3	Available	01-01-01	FC	SCSI	I/0	Controller	Protocol	Device
fscsi4	Defined	02-04-01	FC	SCSI	I/0	Controller	Protocol	Device
fscsi5	Defined	02-05-01	FC	SCSI	I/0	Controller	Protocol	Device

As root:

Note I have a log filesystem called /usr/local/logs – change this to wherever you want to save these files #errpt >/usr/local/logs/errpt-may152021.txt #errpt -a >/usr/local/logs/errpta-may152021.txt #errclear 00 #cfgmgr #errpt







Setup NTP

#vi /etc/ntp.conf
Comment out broadcast and add:
server 0.pool.ntp.org
server 1.pool.ntp.org

#vi /home/padmin/config/ntp.conf Add to end: server 0.pool.ntp.org server 1.pool.ntp.org

#ntpdate 0.pool.ntp.org

Update rc.tcpip to start ntp at boot Now start NTP #startsrc -a "-c /home/padmin/config/ntp.conf" -s xntpd

You can substitute your own NTP servers for the ones above if you have them



<section-header><section-header><text><text><text><text><text><text><text><text>

ι	Update I/O Firmware	
Yo ifc ifc rm	ou may have to unconfigure the SEA to do this (SEA here is on ent14) config en14 down config en14 detach ndev -l ent14	
If I ifc ifc rm	IP is on the real then: config en4 down config en4 detach ndev -l ent4	
cfg Ch sm ifc	gmgr neck microcode went on and sea came back nitty tcpip and check ip is now on config -a	
		48
18		





Useful Links • Jaqui Lynch Articles <u>http://www.circle4.com/jaqui/eserver.html</u> Jagui Lynch Youtube https://www.youtube.com/channel/UCYH6OdgB6rV1rPxYt6FWHpw • Jaqui's Movie Replays http://www.circle4.com/movies • Nigel Griffiths AIXpert Blog https://www.ibm.com/support/pages/aixpert-blog-nigel-griffiths-mrnmon • Nigel Griffiths YouTube https://www.youtube.com/nigelargriffiths • Gareth Coates – Tricks of the POWER Masters https://www.ibm.com/support/pages/node/1116939 • IBM US Virtual User Group <u>https://www.ibm.com/support/pages/node/1120377</u> • Power Systems UK User Group https://www.ibm.com/support/pages/node/1110195 51

51

<section-header><section-header><list-item><list-item><list-item><list-item><list-item><list-item><list-item><list-item><list-item><list-item><list-item><list-item><list-item>

VIOS Specific References • VIO Server Support <u>https://www14.software.ibm.com/support/customercare/sas/f/vios/home.html</u> SDD and SDDPCM Specific procedures for VIOS http://www-01.ibm.com/support/docview.wss?uid=ssg1S7002686&aid=1 • 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 • SEA Load Sharing • https://www.ibm.com/support/pages/how-setup-sea-failover-load-sharing-configuration • https://www.ibm.com/support/pages/shared-ethernet-adapter-sea-fail-over-load-balancing • 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 53