

Care and Feeding of VIO Servers

- Jaqui Lynch
- Chief Architect
- Jaqui@circle4.com



1

1

Agenda

- **Presentation**
 - Fundamentals before you start
 - PowerVM 3.1 Prerequisites
 - Installation
 - Maintenance and Upgrades
 - Backup and recovery
 - Storage
 - Network
 - Monitoring
- **Documentation**
 - Useful Commands
 - Useful Links
 - Backup Material



2

2

Fundamentals before you start



3

3

Stay Current

<http://www14.software.ibm.com/webapp/set2/flrt/liteTable?prodKey=vios>

Version	Recommended Update	Recommended Upgrade	Release Date	EoS/SPS Date
2.2.6.31	2.2.6.41	3.1.0.21	2018.09.21	2020.09.30
2.2.6.32	2.2.6.41	3.1.0.21	2018.11.16	2020.09.30
2.2.6.41		3.1.0.21	2019.05.08	2020.09.30
2.2.6.51		3.1.0.21	2019.12.13	2020.09.30
3.1.0.10	3.1.0.21		2018.11.09	2021.11.30
3.1.0.21			2019.05.08	2021.11.30
3.1.1.10			2019.11.15	2021.11.30



NOTE all levels prior to 2.2.5 are EOS as of December 2018
 2.2.5 went end of service 9/30/2019
 2.2.6 is end of service 9/30/2020 – TIME TO UPGRADE TO V3

4

4

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
3.1.1 and 3.1.1.10 came out 11/15/2019

Download 3.1.1 base from entitled software: <https://www-05.ibm.com/servers/eserver/ess/ProtectedServlet.wss>

You can upgrade directly from 3.1.0 (Fix Pack) or 3.1.1.0 (service pack) to 3.1.1.10 so:

Download 3.1.1.10 update from Fix Central: <http://www-933.ibm.com/support/fixcentral/>

Release notes for 3.1.1.0: <ftp://ftp.software.ibm.com/systems/power/docs/hw/p9/p9eoo.pdf>

Readme for 3.1.1.10 Service Pack (U/g from 3.1.1.0): <https://www.ibm.com/support/pages/node/1106265>

Readme for 3.1.1.10 Fix Pack (U/g from 3.1.0.0): <https://www.ibm.com/support/pages/node/1106697>

NIM Master needs to be at 7200-04-01 at a minimum for v3.1.1.10

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

5

5

VIO Server 3.1.1.* Hiper

<https://www-01.ibm.com/support/entdocview.wss?uid=isg1IJ22290>

iFix is at: <https://aix.software.ibm.com/aix/ifixes/ij22290/>

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

IJ22290 - I/O failures on LPARs using certain FC adapters

I/O failures or hangs can occur when using the following Fibre Channel adapters on AIX or VIOS:

- PCIe3 32Gb 2-port Fibre Channel Adapter (FC: EN1A/EN1B; CCIN: 578F)
- PCIe3 16Gb 4-port Fibre Channel Adapter (FC: EN1C/EN1D; CCIN: 578E)
- PCIe3 16Gb 2-port Fibre Channel Adapter (FC: EN0A/EN0B; CCIN: 577F)

Only VIOS 3.1.1 and AIX 7.2 TL4 are affected.

Affected VIOS/AIX Levels and Recommended Fixes

Minimum Affected Level	Maximum Affected Level	Fixing Level	Interim Fix
VIOS 3.1.1.0 devices.pciex.df1060e214103404.com 7.2.4.0	VIOS 3.1.1.10 devices.pciex.df1060e214103404.com 7.2.4.0	VIOS 3.1.1.20 IJ22290	iFix
AIX 7200-04-00 devices.pciex.df1060e214103404.com 7.2.4.0	AIX 7200-04-01-1939 devices.pciex.df1060e214103404.com 7.2.4.0	AIX 7200-04-02 IJ22290	iFix

6

6

Use FLRT and check Prereqs

FLRT Home Page:

<http://www14.software.ibm.com/webapp/set2/flrt/home>
<https://www-304.ibm.com/support/customer-care/flrt/>

FLRT Lite

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

VIOS to NIM Master Mapping:

VIO v3.1 is still not on here

<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

AIX Support Lifecycle

<https://www-01.ibm.com/support/docview.wss?uid=isg3T1012517>



7

7

PowerVM 3.1 Prerequisites



8

8

Minimum NIM Master Levels for VIOS Clients

<http://www14.software.ibm.com/support/customercaare/flrt/sas?page=viostable>

If using NIM to backup, install or update a VIOS partition, the NIM master must be greater than or equal to the levels shown below.

VIOS Release	VIOS Level	Minimum NIM master level		
VIOS 3.1.0	VIOS 3.1.0.21	AIX 7200-03-03		
	VIOS 3.1.0.10	AIX 7200-03-02		
VIOS 2.2.6	VIOS 2.2.6.41	AIX 7100-05-04	7200-03-03	
	VIOS 2.2.6.32	AIX 6100-09-12	7100-05-03	7200-03-02
	VIOS 2.2.6.31	AIX 6100-09-12	7100-05-03	7200-03-01

VIOS 3.1.1.10 requires 7.2.4.1

9

9

PowerVM 3.1 Changes

- 3.1.0 went GA 11/9/2018 – 3.1.1 went GA 11/15/2019
- Both levels have service packs out
- **This is the move to AIX 7.2 for the VIO server**
- Base order number changes for V3 to 5765-G34
- Native compatibility mode for POWER8 and POWER9
- Accelerated secure LPM for E950 and E980
- Based on AIX 7.2 TL3 (3.1.0) and AIX 7.2 TL4 (3.1.1)
- **USB Flash drive install for VIOS**
- IVM is removed so you must install a proper VIO server
- Database changes from Solid to Postgres for SSP management data
- Many old packages removed to clean up image
- Storage multipathing enhancements
- iSCSI support
 - Can export iSCSI disks to client LPARs using vSCSI (min FW860.20)
 - Enables MPIO support for iSCSI
 - iSCSI boot is not supported
 - iSCSI not supported for SSP (shared storage pools)
- http://www-01.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/6/897/ENUS218-346/index.html&lang=en&request_locale=en
- **As of 3.1.0.20 Java7 is no longer required**
 - Java6 and Java7 are not automatically removed but Java8 is automatically installed
 - Use updateios –remove to uninstall old versions of Java on the VIO servers

10

10

PowerVM 3.1 Prerequisites

- At least 1 core, 8GB memory (Nigel recommends at least 16GB if SSPs)
- High performance (8, 16 and 10Gb cards or higher) cards require additional core and memory
- At least 30GB for rootvg (I recommend 100GB)
- At least 4GB free in rootvg
- **Add an extra disk to be used for alternate disk upgrades – this extra disk is required if you are upgrading from v2 to v3**
- On one VIO it is helpful to have a 3rd disk to use for File Backed Optical if you use it – gets it out of rootvg
- NIM Master must be at AIX 7200-03-02-1846 for 3.1.0, 7200-03-03 for 3.1.0.21 and 7200-04-01 for 3.1.1 and 3.1.1.10
 - Upgrade your NIM to 7200-04-01 so you are ready for future upgrades to your AIX LPARs as well as your VIO servers
- Must use separate HMC and VIO server - IVM is removed
- **Only supports Power7+ (D models) and above**
- **No blades supported**
 - If you need to keep older servers around, then use 2.2.6.51 VIO servers for those
- viosupgrade command on VIO becomes available at 2.2.6.30, but if you have SSPs you must go to 2.2.6.32 or higher before trying to upgrade. I did all my upgrades from 2.2.6.32
- V3.1 or v3.1.1 base is downloaded from ESS and comes as either 2 x DVDs or a flash drive image
- Server must have access to a NIM server, the HMC, a DVD drive or be able to use a flash drive
- For flash drive install USB drive must be at least 16GB
- Latest link to VIOS Maintenance Strategy
 - <https://www14.software.ibm.com/support/customer/sas/f/vios/svcstrategy.html>

11

11

PowerVM 3.1 SSP Important Changes

Important Changes in 3.1.0.X for SSP users

A database manager change has occurred for Shared Storage Pool (SSP). This change will have no direct impacts on behavior, however it does mean that **non-disruptive** upgrades of an SSP cluster to 3.1.0.10 requires that SSP nodes first are updated to the latest 2.2.6.X version available before upgrading to 3.1.0.0 or 3.1.0.10. As of the time of this writing, that is version 2.2.6.31.

Once all VIOS nodes in the cluster have been updated to the latest 2.2.6.X version, double check that rolling upgrade has completed. This can be done by checking the output of "cluster -status -verbose" while logged in as padmin on one of the VIOS nodes in the cluster. Then, check the output for each node, and check for this field:

```
Node Upgrade Status: 2.2.6.31 ON_LEVEL
```

If all nodes have 2.2.6.31 or newer, and all say that they are "ON_LEVEL," then upgrades to the VIOS to 3.1.0.00 or newer can occur without disruption to the SSP cluster.

Additionally, backup and restore can be used to restore older versions of the SSP cluster to 3.1.0.X versions of the VIOS.

The above is from the readme. There are additional limitations spelled out in the readme file

3.1 release notes: <ftp://ftp.software.ibm.com/systems/power/docs/hw/p9/p9eeo.pdf>

3.1.0.10 readme: <https://www-01.ibm.com/support/docview.wss?uid=ibm10738523>

3.1.1.10 readme: <https://www.ibm.com/support/pages/node/1106697>

Check Nigel Griffiths Blog as he has written extensively on SSPs

<https://www.ibm.com/support/pages/aixpert-blog-nigel-griffiths-mrnmon>

12

12

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)

13

13

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**

14

14

Sizing the VIO

Minimums

- Memory 4GB I never use less than 8GB now due to high performance adapters
- Cores .5 entitlement and 2VPs I usually do 1 full core minimum per VIO
- 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 will probably need more like 6 or 8GB and an entitlement of 1.5 or 2.

Pay attention to adapter placement – adapter slots have different priorities
 Details are in the redbook for each server – look for the technical overview

If using 10Gb network or 8Gb, 16Gb or 32GB HBA adapters you need more memory for buffering and more CPU to handle traffic

i.e. 512MB for each **active** high performance adapter port (NPIV or vSCSI)
 Plus 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 8GB memory and entitlement of 1-2 cores, especially if using SSPs

rootvg needs at least 30GB – give it 100GB disk space
 Add an extra disk if want to use FBO – don't put FBO repository 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>

Server Memory

- Reserved Memory is based on max memory for an LPAR, not on desired
- This is because memory gets reserved for HPTs (hypervisor page tables)

BAD EXAMPLE:

1	Name	Mode	Min GB	Curr GB	Max GB	ExpFact	AMS->	Weight	Prim VIOS	Sec VIOS	Curr VIOS
2	lpar1	ded	2.25	5.00	100.00						
3	lpar2	ded	2.00	15.75	100.00						
4	lpar3	ded	2.00	15.75	100.00						
5	lpar4	ded	10.00	64.00	100.00						
6	vios2	ded	2.00	6.00	100.00						
7	vios1	ded	2.00	6.00	100.00						
8											

Ent_Sys_Pools OnOff CoD CoD Events LPAR_Summary LPAR_Profiles LPAR_CPU **LPAR_MEM** Physical_Slo

BETTER EXAMPLE:

Name	Mode	Min GB	Curr GB	Max GB
lpar1	ded	2.25	5.00	8.00
lpar2	ded	2.00	15.75	20.00
lpar3	ded	2.00	15.75	20.00
lpar4	ded	10.00	64.00	80.00
vio2	ded	2.00	6.00	12.00
vio1	ded	2.00	6.00	12.00

Server Memory

Status	Type Model	Serial	GHz	CPU Type	Tot Cores	Act Cores	Deconf Cores	Curr Avail Cores	Pend Avail Cores	Ded Cores	Pool Size	Virt Procs	#LPAR	Tot GB	Act GB	Deconf GB	Firm GB	Curr Avail GB	
Operating	8286-41A		0	3.72	POWER8	8	8	0	1.40	1.40	0	8	14	15	256.00	256.00	0.00	5.75	126.25

CPU Type	Tot Cores	Act Cores	Deconf Cores	Curr Avail Cores	Pend Avail Cores	Ded Cores	Pool Size	Virt Procs	#LPAR	Tot GB	Act GB	Deconf GB	Firm GB	Curr Avail GB	Pend Avail GB
PowerPC_POWER8	80	26	0	21.64	21.64	0	26	22	28	3,072.00	1,536.00	0.00	23.00	1,447.00	1,447.00
PowerPC_POWER8	80	26	0	0.31	0.31	0	26	52	39	3,072.00	1,536.00	0.00	38.50	967.50	967.50

Look at Firm GB in HMCScanner under System Summary Tab
 Latest is 0.11.42 (5/23/2019)

<https://www.ibm.com/support/pages/node/1117515>

[https://www.ibm.com/support/pages/sites/default/files/inline-files/\\$FILE/hmcScanner-0.11.42.zip](https://www.ibm.com/support/pages/sites/default/files/inline-files/$FILE/hmcScanner-0.11.42.zip)

Installation



Get the files to upgrade to V3.1.1 (1/24/2020)

- Latest ISO download is for VIOS 3.1.1.0
- Flash image ISO is VIOS 3.1.1.10 – use this and save a step
- Note the link below is a new URL for ESS as it just moved:
- <https://www.ibm.com/servers/eserver/ess/index.wss?lnk=msdDO-enss-usen>

Go to ESS

<https://www-05.ibm.com/servers/eserver/ess/ProtectedServlet.wss>
Sign in with your userid and password for ibm

At the left click on my entitled software
Make sure to the right "Brand selection" says AIX

Once it says AIX click on software updates at the left
It will prompt you for customer number and serial number for a system - use a valid one
Then you will select Powervm v3.1 (5765-ve3) and click on continue (NOTE v3 has a different program number)
Click on agree and then select I want to download now

It will take you to software downloads - make sure category says AIX and v7.2 then click on continue
Check the box that says 5765-ve3 - PowerVM Enterprise ED v3 and click on continue
Then check the powervm box (this is really repetitive) and click on continue
On the next page click on I agree and go to the bottom of the next page and click on "click here to use http"

There are 3 images to be downloaded:

Download all 3 and burn them to DVD

You can also burn the last one (with flash in its name) to a USB stick

19

19

PowerVM v3 Download from ESS – 1/24/2020

03.01.01	2344: IBM PowerVM V3 / VIOS v03.01.01,ENU,DVD	 View
----------	--	--

README for ISO Download

ISO, Virtual I/O Server v3.1.1.0
DVD 1 of 2 (11/2019)

ISO, Virtual I/O Server v3.1.1.0
DVD 2 of 2 (11/2019)

ISO, Virtual I/O Server v3.1.1.10
Flash (11/2019)

Flash image



20

20

Install Options

- Download v3.1.1.0 from ESS
- Download 3.1.1.10 Flash Image – I used this image
- Download the latest expansion pack from Fix Central

- **Fresh install of VIOS 3.1.1 on a new server**
 1. LPM off all LPARS then fresh install of VIOS 3.1.1 on old server
 2. Install from DVD or USB
 3. **Use NIM to do VIO install to an alternate disk – my preference**
 4. Install VIO from repository on HMC

- **Upgrades – assumes you are at least at 2.2.6.32 of VIO**
 - Use NIM viosupgrade to upgrade current server to an alternate disk
 - If using NIM for bosinst install, then VIOS IP cannot be on the SEA
 - Can still install to alt disk though

- **Use VIO viosupgrade to upgrade current server to an alternate disk (my preferred method)**
 1. VIO viosupgrade requires VIOS to be at 2.2.6.30+, SSP requires 2.2.6.32 – recommend going to 2.2.6.32 minimum (or .51 which is latest)
 2. Use viosupgrade –l –q to monitor VIO upgrade status

- Read the readme/description files for all levels
- If you are using SSPs pay attention to the restrictions and rules around upgrades with SSPs in place
- Note if upgrading versus full install - you cannot use updateios for this upgrade – you must use the viosupgrade command

21

21

Things to think about for Fresh Install

- It is not that different from any other VIO server fresh install
- If replacing a current VIOS
 - Use viosbr to backup metadata and copy the files to a remote system
 - Backup anything outside of rootvg on your VIOS to a remote location
 - Backup the VIOS itself or take a clone (alt_disk_copy)
 - If using SSPs then perform the steps in the README for SSPs
- Install from the v3.1.1 DVD – for NIM use the mksysb from the flash image
- If you can, use the flash image which is at 3.1.1.10 (saves an upgrade step)
 - I was able to burn this iso to both a DVD and a flash drive and to create the mksysb for NIM
- Fresh Install can be an overwrite of current disk or to an alternate disk
 - Use the alternate disk if at all possible
- Restore the metadata and anything that was outside of rootvg
- Perform post install SSP steps

- **Before doing any install or upgrade - check for compatibility between HMC, firmware and VIOS levels as well as client operating systems**

22

22

Backup Virtual Definitions

- Check on your viosbr backups

```
$ ls -al /home/padmin/cfgbackups
total 56
-rw-r--r--  1 padmin  staff    7938 Dec 17 10:00 autoviosbr_vio1.tar.gz
-rw-r--r--  1 padmin  staff    7763 Apr 13 2019 vio1_10158152.tar.gz
-rw-r--r--  1 padmin  staff    7929 Jan 26 11:52 viosname.01.tar.gz
```

```
$ viosbr -view -list
autoviosbr_vio1.tar.gz
vio1_10158152.tar.gz
viosname.01.tar.gz
```

I have my viosbr setup to run daily and keep the last 7 copies, so I only have to copy across the most recent one

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

The above is run once to setup the copy

crontab -l (after oem_setup-env) should now show:

```
0 0 * * * (/usr/ios/cli/ioscli viosbr -backup -file viosname -frequency daily -numfiles 7)
```

23

23

Prior to install or Upgrade

- Backup anything outside of rootvg on your VIOS to a remote location
 - FBO library (if in rootvg)
 - LVs or filebacked disks to clients
- Backup LVs for clients
 - **viosbr does not back these up and viosupgrade does not copy them so back them up**
 - Move LVs for clients to a different VG than rootvg prior to upgrade (they should never be in rootvg)
 - Migrate after backup then use cplv to copy them after the migration
 - <http://www-01.ibm.com/support/docview.wss?uid=isg3T1000167>
- Backup the VIOS itself
 - mount /backups (NFS filesystem from my NIM server)
 - mkdir /backups/vio2
 - umount /var/vio/VMLibrary (makes sure I don't accidentally include the media library)
 - su - padmin -c "ioscli backupios -file /backups/vio2 -nomedialib"
 - su - padmin -c "ioscli backupios -file /usr/local/backups/vio2-previo31-mar2719.mksysb -nomedialib -mksysb"
 - mount /var/vio/VMLibrary

24

24

Pre-install Notes

- If upgrading current VIO servers
 - When working on primary you may want to failover the SEA to the secondary VIO
 - If SEA is ent10
 - As padmin:
 - \$schdev -dev ent10 -attr ha_mode=standby
 - Once complete and all updates done and primary VIO has done its final reboot
 - \$schdev -dev ent10 -attr ha_mode=auto
 - You should see messages in errpt that show the changes from primary to backup and back again
- Aggregation and installs and restores
 - You cannot install a VIO server from the HMC or from NIM if the network is aggregated
 - Network installs are only supported over an access port connection
 - This applies to installing any LPAR that has physical network ports that are aggregated
- Installing onto SAN disks
 - The SAN team may need you to light up the adapters so they can do their zoning and mapping

25

25

Full Install

- From DVD or USB – complete install – with PowerVM 3.1 and higher you will be able to use USB
- Basically boot in SMS mode then tell it to boot from DVD, flash, NIM or HMC and follow instructions
- Using NIM
 - <http://www-01.ibm.com/support/docview.wss?uid=isg3T1011386>
 - Minimum NIM levels
 - <http://www14.software.ibm.com/webapp/set2/sas/f/flrt/viostable.html>
- Using HMC - check vios install box
- Commandline - installios:
 - 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>
- **Network between HMC and VIO LPAR must be alive and not aggregated (request an access port)**
- From a mksysb
 - http://pic.dhe.ibm.com/infocenter/flexsys/information/index.jsp?topic=%2Fcom.ibm.acc.psm.resources.doc%2Fvios%2Fsdmc_vios_vios_backup_restore_file_nim.html
- **After install fix the page spaces – depending on the version you will have 1 x 512MB and 1 x 1024MB or 2 x 1024MB on the same hdisk. Get rid of paging00 and make hd6 at least 4 to 6GB**

26

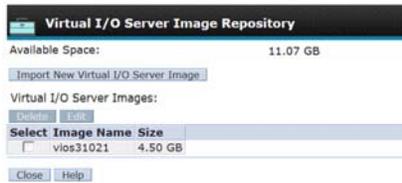
26

VIOS and HMC – Import VIOS31 ISO Images

```
aix1nim:/software/powervm31> du -sg Virt*
4.18 Virtual_IO_Server_Base_Install_3.1.0.0_DVD_1_of_2_112018.iso
1.27 Virtual_IO_Server_Base_Install_3.1.0.0_DVD_2_of_2_112018.iso
4.29 Virtual_IO_Server_Base_Install_3.1.0.10_Flash_112018.iso
4.50 Virtual_IO_Server_Base_Install_3.1.0.21_Flash_052019.iso
```

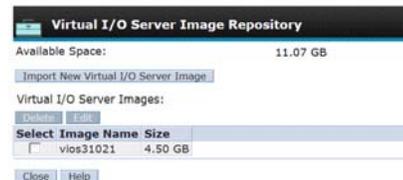
So need 5.45GB minimum if not using the flash image

1. Check repository for space
2. Import the ISO images
3. Message importing



Importing Virtual I/O Server Images from Remote SFTP Server.
This might take several minutes, Please standby.

4. Import complete



31021 is flash image at 3.1.0.21

You can just upload the flash image and use that – it is more current (3.1.1.10 today) and works fine

27

27

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 EVER 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>
- **NIM Master must be at AIX 7200-04-01 for VIOS 3.1.1.10**
- On the NIM Master, use the operation **viosupgrade** in altdisk mode to update the VIOS Server to v3 from v2 to an alternate disk

I had a problem with the flash ISO – you can mount it on AIX using loopmount but you need to mount it as udfs
`loopmount -i /software/powervm31/Virtual_IO_Server_Base_Install_3.1.0.21_Flash_052019.iso -o "-V udfs -o ro" -m /cdrom`
If that fails then open it on windows, extract the mksysb and upload the mksysb to the NIM server

- Using NIM with VIO Servers
<https://ibmsystemsmag.com/Power-Systems/09/2019/Using-NIM-with-VIO-Servers>

28

28

VIOS and NIM

- Add VIOS partition as a NIM client
- Copy the VIOS mksysb image from the CD to your NIM master
 - On VIOS 3.1 base media there are 3 images now across the two DVDs
 - Copy all 3 images individually to a directory and then use cat to combine them

```
cat /export/mksysb/vios3.1/mksysb_image /export/mksysb/vios3.1/mksysb_image2 /export/mksysb/vios3.1/mksysb_image3 >/export/mksysb/nim_vios3.1mksysb
```

OR save yourself time and use the flash image as it is just one mksysb image
- Define the mksysb resource to the NIM master after copying the mksysb into /nim/images
- Define the spot on the NIM master
 - The source for the SPOT will be the combined mksysb or the single flash image mksysb
 - **The SPOT CANNOT be created from an LPP_Source**
 - `nim -o define -t spot -a server=master -a source=mksysb_vios31021 -a location=/nim/spot spotvios31021`
 - `nim -o check spotvios31021`
- Copy the bosinst.data from the DVD and create a viosbosinst resource
- Allocate the mksysb, spot and bosinst resources to the VIO LPAR in NIM and then set it up for a bosinst install from **mksysb**
- You can now use bos_inst to do a mksysb install once the partition profile is defined (fresh install) or NIM's viosupgrade if upgrading
- NOTE syntax for NIM viosupgrade is not the same as the viosupgrade run directly on the VIO server
- https://www.ibm.com/support/knowledgecenter/en/ssw_aix_72/com.ibm.aix.cmds6/viosupgrade.htm

29

29

Cloning disks

If your server has a split backplane then you can make a clone
 After installing VIO1, if you have all the disks in VIO1 you can take a clone to build VIO2
 Make sure the 4 disks are split (2 and 2) across the backplane
 VIO1 is using hdisk0 and hdisk1, hdisk2 and 3 are on the other adapter and will be used for VIO2
 Put all the disks into VIO1 (both adapters)
 Install VIO1 on hdisk0 – from NIM, DVD, USB, HMC

Now clone it to hdisk2

```
alt_disk_copy -B -d hdisk2
```

Check bootlist has not changed after copy finishes

Remove VIO2 hdisks from VIO1, Shutdown VIO1,
 Remove VIO2 resources from VIO1 profile
 Leave VIO1 down

Activate VIO2 (make sure only VIO2 resources are in VIO2 profile)
 Remove any disks, adapters, networks etc that show as defined on VIO2
 Now cleanup VIO2 (see next slide)

It is best to make the clone before you have the network and fibre adapters attached to VIO1 – it makes the post-clone cleanup much easier

30

30

Cleaning up after cloning VIO

If you do not take these steps you will experience RMC issues

Cleanup VIO2:

```
stopsrc -g rsct_rm; stopsrc -g rsct
```

Clear Nodeid

```
chdev -l cluster0 -a node_uuid=00000000-0000-0000-0000-000000000000
```

OR

```
/usr/bin/odmdelete -o CuAt -q 'attribute=node_uuid'
```

Generate new nodeid

```
/usr/sbin/rsct/bin/mknodeid -f
```

```
lsattr -El cluster0
```

```
/usr/sbin/rsct/bin/lsnodeid
```

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

```
lspartition -dlpar
```

```
lssrc -g rsct_rm; lssrc -g rsct
```

You may have to start ctcas – startsrc –s ctcas

Cleanup old VIO1 resources (next slide)

31

31

Cleaning up after cloning VIO

CLEANUP on VIO2

These will vary depending on the server and I/O drawers, etc and whether you had ethernet and fibre adapters in the vio when cloned

```
rmdev all devices showing as defined (fcs, ent, hdisk, etc)
```

```
rmdev -dp hdisk0
```

```
rmdev -dl hdisk0
```

```
rmdev -dp pdisk0
```

```
rmdev -dl pdisk0
```

```
rmdev -dp sissas0
```

```
rmdev -dl sissas0
```

```
rmdev -dp pci0
```

```
rmdev -dp pci1
```

```
rmdev -dp pci2
```

```
rmdev -dp pci3
```

```
rmdev -dp pci4
```

```
rmdev -dl pci0
```

```
rmdev -dl pci1
```

```
rmdev -dl pci2
```

```
rmdev -dl pci3
```

```
rmdev -dl pci4
```

If ethernet adapters were in VIO1 when cloned then you may need to remove all those as well

Once VIO2 is cleaned up reboot it

Then activate VIO1

Clean up VIO1 removing any extra hdisks, pdisks, pci, sissas1, etc that now show as defined. Also remove the adapter definitions for them.

Reboot VIO1 to ensure changes are good

32

32

Maintenance and Upgrades



<https://ibmsystemsmag.com/Power-Systems/05/2019/powervm-experience>

33

33

Upgrading VIOS

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

If using NPIV tape drives make sure they are not in use (or activated on IBM i)

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
 2. As an example P9 FW940 requires that all I/O adapters be updated and the HMC be at v9r1m940 prior to installing the firmware
3. Download the updates and cross-check compatibility using FLRT
4. Read the readmes again
5. Run errpt to check for problems, check there are no stale partitions, missing disks or paths, etc
 - lsvg rootvg checks for stale PPs and physical volumes.
 - lsvg -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 – copy files locally if you can
10. If using SSPs there are specific additional steps outlined in the README
11. After upgrading and rebooting the first VIOS check that all your LPARs are back to dual paths (lspath)
12. Only after that should you upgrade the second VIOS

34

34

Things to think about when Upgrading

- Migrating from v2 to VIO 3.1.0 or 3.1.1 is an **upgrade** not an update. You cannot use updateios
- Use viosbr to backup metadata and copy the files to a remote system
- Create a filestosave.txt file that you keep a list of critical files to be saved in
- Backup anything outside of rootvg on your VIOS
 - FBO library
 - LVs for filebacked disks to clients
- Backup the VIOS itself or take a clone (alt_disk_copy – but rename the clone)
- If using SSPs then perform the steps in the README for SSPs
- Perform the upgrade
- Restore the metadata (upgrade should do this)
- Restore anything that was outside of rootvg
- Perform post upgrade SSP steps

- Make sure no NPIV tapes are assigned
- Make sure no virtual optical is loaded and assigned

35

35

Create a filestosave.txt

This file is used to save files that you may need to get information out of after the v3 upgrade is complete

My filestosave.txt is in /home/padmin and consists of:

```

/etc/environment
/etc/group
/etc/hosts
/etc/inetd.conf
/etc/inittab
/etc/motd
/etc/netsvc.conf
/etc/passwd
/etc/profile
/etc/resolv.conf
/etc/syslog.conf
/etc/security/limits
/etc/security/login.cfg
/etc/security/passwd
/etc/tunables/nextboot
/etc/tunables/rc-tunevio.sh
/usr/local/bin/runnmon.sh
/etc/ssh/sshd_config
/home/padmin/.profile
/home/padmin/filestosave.txt
/usr/local/bin/viobackup.sh
/etc/ntp.conf
/etc/rc.tcpip
/home/padmin/config/ntp.conf

```

These will get saved to /home/padmin/backup_files and will be available later on the upgraded system

36

36

Upgrade

Find a spare disk and clean it off

```
$ lspv
NAME          PVID              VG      STATUS
hdisk0        00f95d3a1b679a90  fbovg   active
hdisk2        00f95d3a42550d49  fbovg   active
hdisk3        00f95d3a0de356cd  altinst_rootvg
hdisk1        00f95d3a42550ec9  rootvg   active
```

```
$ lspv -size | head
NAME          PVID              SIZE(megabytes)
hdisk0        00f95d3a1b679a90  51200
hdisk2        00f95d3a42550d49  51200
hdisk3        00f95d3a0de356cd  102400
hdisk1        00f95d3a42550ec9  102400
```

- Then as padmin look for free or unused disks

```
$ lspv -unused
$ lspv -free
```

- Check for mappings

```
$ lsmmap -all | grep hdisk
```

In the above all disks are assigned (none unused or free). lsmmap also shows none are mapped to clients using vSCSI.
We also have altinst_rootvg which is not allowed

37

37

Forgot to clear the disk (VIO viosupgrade)

My base install mkysyb file is the mkysyb file that I grabbed from the flash image iso file

```
viosupgrade -l -i /usr/local/soft/nim_powervm31_basemkysyb -a hdisk3 -g /home/padmin/filestosave.txt
```

Welcome to viosupgrade tool.

Operation triggered for given node(s).

Broadcast message from root@vio2 (pts/0) at 15:59:29 ...

WARNING!!! VIOS Upgrade operation is in progress. Kindly Refrain from making any configuration changes...

Please wait for completion..

The provided disk 'hdisk3' is in use.

Go back and make sure lspv -free and lspv -unused show the disk you want to use as available

38

38

Get a disk

altinst_rootvg cannot exist prior to the upgrade so either export and reimport with a new name or delete it

```
#exportvg altinst_rootvg
#importvg -y rootvgcopy hdisk3
Or just rename it:
#alt_rootvg_op -v alt_disk_jan20 -d hdisk3
```

```
OR delete it:
exportvg altinst_rootvg
OR
alt_rootvg_op -X altinst_rootvg
```

Recommended method is always to use alt_rootvg_op

AFTER delete:

```
# lspv
hdisk0    00f95d3a1b679a90      fbovg    active
hdisk2    00f95d3a42550d49      fbovg    active
hdisk3    00f95d3a0de356cd      None     active
hdisk1    00f95d3a42550ec9      rootvg   active
```

alt_rootvg_op

https://www.ibm.com/support/knowledgecenter/en/ssw_aix_71/a_commands/alt_rootvg_op.html

Managing multiple instances of altinst_rootvg

<https://www.ibm.com/support/pages/managing-multiple-instances-altinstrootvg-and-applying-fixes-them>

39

39

Get a disk

Just because a disk shows as not being in a volume group does not mean it is available

You will probably need to clear the owning volume manager from the disk

```
# chpv -C hdisk3
```

You can also clear the boot image as well

```
# chpv -c hdisk3
```

```
# lspv
hdisk0    00f95d3a1b679a90      fbovg    active
hdisk2    00f95d3a42550d49      fbovg    active
hdisk3    00f95d3a0de356cd      None     active
hdisk1    00f95d3a42550ec9      rootvg   active
```

```
$ lspv -free
```

```
NAME      PVID                SIZE(megabytes)
hdisk3    00f95d3a0de356cd   102400
```

OK NOW we can use hdisk3 for our upgrade

40

40

Upgrading VIOS to V3.1

You need to have your VIO at 2.2.6.30 or higher to use the VIO server viosupgrade command – recommend at least 2.2.6.32
 If you are using SSPs then you have to be at 2.2.6.32
 I recommend going to 2.2.6.32 (or .51 which is latest) regardless and use that as a starting point

As padmin run “updateios –commit” to ensure any uncommitted updates are committed
 Check to ensure there are no missing filesets prior to updates
 Check repository has nothing loaded

```
$ ioslevel
2.2.6.32
$cat /usr/ios/cli/ios.level
2.2.6.32
```

```
$ updateios -commit
All updates have been committed.
```

```
$ oem_setup_env
# /usr/sbin/emgr -P
There is no efix data on this system.
If there are any ifixes remove them
```

Now run checks

41

41

Upgrade – get the code

- You should already have downloaded the flash image for 3.1.0.21 from ESS (entitled software) – as of 11/2019 it is 3.1.1.10
- It comes down as an iso image
- Either
 1. Open it on your windows desktop and extract the mksysb image
 2. Upload the mksysb image as binary to your VIO or NIM server
- OR
 1. Use loopmount on AIX

```
loopmount -i /software/powervm31/ Virtual_IO_Server_Base_Install_3.1.0.21_Flash_052019.iso -o "-V udfs -o ro" -m /cdrom
```

```
mkdir /usr/local/soft/powervm31110
cp -r /cdrom /usr/local/soft/powervm31110
```

Find the mksysb image and put it into /usr/local/soft/powervm31110
- I put it in /usr/local/soft/ which is a local directory on the VIO
- I also copy the .iso and mksysb images to /usr/local/soft
- At the same time I download all the Java, SSH, SSL and I/O adapter updates that I need

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

42

42

What does the VIO viosupgrade command do?

This is my VIO 3.1.0.21 upgrade from 2.2.6.32 – files were all in /usr/local/soft/vios31021

It does the config backup for you then it builds vios 3.1.0.21 on the new disk

It migrates the config

It sets the bootlist

It will then reboot – you have 60 seconds to stop it

```
viosupgrade -l -i /usr/local/soft/vios31021-flash-mksysb_image -a hdisk3 -g /home/padmin/filestosave.txt
```

Below is the syntax

```
viosupgrade -l
```

Flags:

- l Specifies local Node Installation.
- i Specifies image file for the alternate disk installation.
- a Specifies alternate disk to install the provided image.
- c Specify if the node is part of the cluster.
- g Specifies the filename having the list of files to be copied to newly installed rootvg.
- q Queries the status of VIOS restore operation after booting the VIOS with newly installed image.

43

43

Upgrade Attempt 1 - failed

Now on the VIO:
\$updateios -commit

My first attempt at the upgrade:

```
viosupgrade -l -i /usr/local/soft/nim_powervm31_basemksysb -a hdisk3 -g /home/padmin/filestosave.txt
```

Welcome to viosupgrade tool.

Operation triggered for given node(s).

Broadcast message from root@vio2 (pts/0) at 16:04:34 ...

WARNING!!! VIOS Upgrade operation is in progress. Kindly Refrain from making any configuration changes...

Please wait for completion..

Initiating VIOS configuration backup..

VIOS configuration backup successful.

Initiating installation on alternate disk(s)..

Installation on alternate disk(s) failed.

I checked the log at: /var/adm/ras/ioslogs/viosupg_global.log

Restoring mksysb image to alternate disk(s).

restore: 0511-110 There is an unpacking error.

restore: 0511-708 There is an internal unpacking error: decode failure

restore: 0511-108 There was an error during the unpacking of ./opt/IBM/ldap/V6.4/lib64/libibmldapn.a

Turns out I had messed up my mksysb image (initially I concatenated the 3 images on the 3.1.0 DVDs)

This is when I went and got the single image from the Flash ISO

44

44

Upgrade Attempt 2

```
viosupgrade -l -i /usr/local/soft/vios31010-flash-mkysyb_image -a hdisk3 -g /home/padmin/filestosave.txt
```

```
Welcome to viosupgrade tool.
Operation triggered for given node(s).
Broadcast message from root@vio2 (pts/0) at 13:46:39 ...
WARNING!!! VIOS Upgrade operation is in progress. Kindly Refrain from making any configuration changes...
Please wait for completion..
Initiating VIOS configuration backup..
VIOS configuration backup successful.
Initiating installation on alternate disk(s)..
Installation on alternate disk(s) successful.
Copying files to allinst_rootvg.
Waking up allinst_rootvg successful.
Putting volume group allinst_rootvg to sleep ...
forced unmount of /alt_inst/var/adm/ras/livedump
forced unmount of /alt_inst/var/adm/ras/livedump
forced unmount of /alt_inst/var
forced unmount of /alt_inst/var
forced unmount of /alt_inst/usr
forced unmount of /alt_inst/usr
forced unmount of /alt_inst/tmp
forced unmount of /alt_inst/tmp
forced unmount of /alt_inst/opt
forced unmount of /alt_inst/opt
forced unmount of /alt_inst/home
forced unmount of /alt_inst/home
forced unmount of /alt_inst/admin
forced unmount of /alt_inst/admin
forced unmount of /alt_inst
forced unmount of /alt_inst
Fixing LV control blocks...
Fixing file system superblocks...
VIOS will be rebooted after '60' seconds to boot from the newly installed disk.
```

Press contr+c to terminate.

```
VIOS metadata restore (viosbr -restore) will be automatically resumed
after the reboot.
VIOS may be rebooted once during this restore process. Refrain from making
any changes to the VIOS virtual configurations during the restore process.
You can verify the restore status using 'viosupgrade -l -q' command and
resume your operation after the completion of the restore process.
```

I hit ctrl -c to stop the reboot

45

45

Check status (after I hit ctrl-c)

```
$ viosupgrade -l -q
Welcome to viosupgrade tool.
Getting status of node(s):
```

viosupgrade is in progress

Please see the viouprgrade status:

```
=====
Thu Mar 28 13:46:39 2019|STARTED
Thu Mar 28 13:51:14 2019|TRIGGERED
```

Please see the viosbr restore status:

```
=====
```

```
$ viosbr -view -list
vio2_13303902.tar.gz          this is the image created by the upgrade I just did
vio2m.01.tar.gz
```

```
$ ls -alt cfbgbackups
total 120
-rw-r--r--  1 padmin  staff   6922 Mar 27 16:05 vio2_13303902.tar.gz  this is the image created by the upgrade I just did
-rw-r--r--  1 padmin  staff   6960 Mar 27 10:55 vio2m.01.tar.gz
```

46

46

Hold off on reboot

I was not ready to boot the upgrade until my maintenance window when I planned to completely redo the upgrade, so I hit ctrl-C when prompted then I did the following to make sure I did not accidentally boot the upgraded copy:

```
# bootlist -m normal -o
hdisk3 blv=hd5 pathid=0
hdisk3 blv=hd5 pathid=1
hdisk3 blv=hd5 pathid=2
hdisk3 blv=hd5 pathid=3
```

Set it back to the current (unupgraded) disk – hdisk1:

```
# lspv
hdisk0 00f95d3a1b679a90 fbovg active
hdisk2 00f95d3a42550d49 fbovg active
hdisk3 00f95d3a0de356cd altinst_rootvg
hdisk1 00f95d3a42550ec9 rootvg active
```

```
# bootlist -m normal hdisk1
```

```
# bootlist -m normal -o
hdisk1 blv=hd5 pathid=0
hdisk1 blv=hd5 pathid=1
hdisk1 blv=hd5 pathid=2
hdisk1 blv=hd5 pathid=3
```

47

47

After reboot

I reran the process (exported the disk and started the upgrade again) during my maintenance window to make sure I was up to date

A broadcast message is sent out

WARNING!!! VIOS Upgrade operation is in progress.

Kindly Refrain from making any configuration changes...

Then it reboots from the alternate disk

There are at least 2 reboots before the VIO server stays up

After the reboot it will require you to change your password (remember this is an overwrite install even if you upgrade)

Then you have to accept the license:

Indicate by selecting the appropriate response below whether you accept or decline the software maintenance terms and conditions.

Accept (a) | Decline (d) | View Terms (v) > a

Now run the `viosupgrade -l -q` to check what happened – see next slide:

You should see started, triggered, restore, restore and then completed and it then shows the `viosbr restore status`
It shows the restore that happened and provides information on devices it could not restore

Now run all your post upgrade checks

48

48

viosupgrade -l -q

```
$ viosupgrade -l -q
Welcome to viosupgrade tool.
Getting status of node(s):
```

viosupgrade is in progress

Please see the vioupgrade status:

```
=====
Sat Apr 13 23:54:22 2019|STARTED
Sat Apr 13 23:58:58 2019|TRIGGERED
Sat Apr 13 23:09:26 2019|RESTORE
```

Please see the viosbr restore status:

```
=====
```

viosbr restore timestamp:
Sat Apr 13 23:09:26 CDT 2019

License acceptance is successful

Restoring the backup..
Lots more messages then:

I logged in too soon so it rebooted at least once more – if you wait 10 minutes after the first reboot it will do the reboot

49

49

Status after upgrade

```
ioslevel
3.1.0.10
```

```
# oslevel -s
7200-03-02-1846
```

```
# oslevel -s -l 7200-03-02-1846
```

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

```
All filesets for 7.2.0.0_AIX_ML were found.
All filesets for 7200-00_AIX_ML were found.
All filesets for 7200-01_AIX_ML were found.
All filesets for 7200-02_AIX_ML were found.
All filesets for 7200-03_AIX_ML were found.
```

Now I had to go make all my customizations for filesystems, paging, logging, /etc/environment, etc

Also had to rerun tunables for virtual ethernet, no commands, etc plus add back in scripts such as my nmon and backup scripts

This upgrade is basically a fresh install but it restores your virtual definitions

50

50

Post Upgrade or install (after reboot)

- Fix page spaces if you have not already done so
- Use updateios to upgrade to 3.1.0.21 or 3.1.1.10 if you did not install at that level
- Update Java7 to 7.0.0.655 or the latest – better yet remove it
- Make sure Java8 is installed at 8.0.0.600 or higher
- Install ssl-1.0.2.2001 or higher
- Install ssh 7.5.102.2000 or higher (v8 is now out)
- Be aware that ssh v7 & v8 tighten up security over v6 so you may need the saved copy of /etc/ssh/sshd_config
- If you have Java6 then update it to 6.0.0.655 or the latest
 - Better yet remove it
- You get the Java updates from Fix Central
 - Run `lspp -l | grep ava` to find out what you have installed
- SSH and SSL are obtained from the Web Download Pack which has moved to:
 - <https://www-01.ibm.com/marketing/iwm/iwm/web/pickUrxNew.do?source=aixbp>

51

51

Java, etc Patched Levels as at 1/26/2020

Java7.jre	7.0.0.655	COMMITTED	Java SDK 32-bit Java Runtime
Java7_64.jre	7.0.0.655	COMMITTED	Java SDK 64-bit Java Runtime
Java8.jre	8.0.0.600	COMMITTED	Java SDK 32-bit Java Runtime
Java8_64.jre	8.0.0.600	COMMITTED	Java SDK 64-bit Java Runtime
openssl.base	1.0.2.2001	COMMITTED	Open Secure Socket Layer
openssh.base.client	7.5.102.2000	COMMITTED	Open Secure Shell Commands
openssh.base.server	7.5.102.2000	COMMITTED	Open Secure Shell Server

52

52

Updating Java, SSH and SSL

- SSH and SSL are obtained from the Web Download Pack which has moved to:
 - <https://www-01.ibm.com/marketing/iwm/iwm/web/pickUrxNew.do?source=aixbp>
- Untar the files and put all ssh, ssl and java files (Java6 through 8 if needed) into a directory. I used /usr/local/soft/javasshssl
- \$updateios -commit
- \$updateios -accept -install -dev /usr/local/soft/javasshssl
 - There are about 96 to go on
- #lspp -l | grep Java8
- Make sure Java8.sdk and Java8_64.sdk are on
- \$updateios -commit
- **\$updateios -remove Java6**
 - Removes 7 filesets
- **\$updateios -remove Java6_64**
 - Removes 7 filesets
- As of 3.1.0.21 you can also remove Java7 the same way you remove Java6 above

53

53

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

<https://www-304.ibm.com/webapp/set2/sas/f/flrt/flrtvc.html>

You should do this on AIX LPARs too

```
/usr/sbin/emgr -l lists them
```

To apply a fix change into the directory it is in and then run it in preview mode:

```
cd /usr/local/soft/vios31fixes/ij16586
```

```
emgr -p -e IJ16586s3a.190607.VIOS3.1.0.21.epkg.Z
```

Remove the -p and run again for real if the preview was successful:

```
emgr -e IJ16586s3a.190607.VIOS3.1.0.21.epkg.Z
```

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

54

54

Efixes and ifixes

For 3.1.0.21 as of 10/2/2019 three patches are needed once SSH, SSL and Java are updated:

```
ij16586      ftp://aix.software.ibm.com/aix/ifixes/ij16586
ij17505      ftp://aix.software.ibm.com/aix/ifixes/ij17505
ntp_fix12    ftp://aix.software.ibm.com/aix/efixes/security/ntp\_fix12.tar
```

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

```
# emgr -l
```

```
ID STATE LABEL      INSTALL TIME    UPDATED BY ABSTRACT
=====
```

```
1  S  IJ16586s3a 08/27/19 19:52:26      Ifix for APAR IJ16586
```

It will vary by O/S level and SP. This was for 3.0.1.21

To remove an efix or ifix:

```
# /usr/sbin/emgr -r -L <EFIX label>
emgr -r -L IJ16586s3a
```

55

55

VIOS Patches at 3.1.1.10

At 3.1.1.10 there are four additional patches to go on once openssl, openssh and java are all updated to the latest patch levels

Patches identified by FLRTVC for 3.1.1.10 of VIO as of 1/26/2020

Fileset	Current	Type	Abstract	APARs
bos.net.tcp.tcpdump	7.2.4.0	sec	NOT FIXED - There is a vulnerability in tcpdump that affects AIX.	IJ20786
devices.pci.df1000f7.com	7.2.4.0	hiper	NOT FIXED - Possible undetected data corruption with certain Fibre Channel adapters	IJ21527
devices.pci.df1000f7.com	7.2.4.0	hiper	NOT FIXED - Possible undetected data corruption with certain Fibre Channel adapters	IJ21527
pool.basic.rte	7.2.4.1	hiper	NOT FIXED - VIOS 3.1.1 SSP may not recover from network loss	IJ21564

APARs	Bulletin URL	Download URL	Reboot Required	Last Update	Fixed In
IJ20786	http://aix.software.ibm.com/aix/efixes/security/tcpdump_advisory5.asc	ftp://aix.software.ibm.com/aix/efixes/security/tcpdump_fix5.tar	NO	1/8/2020	7200-04-02-2015
IJ21527	http://www-01.ibm.com/support/docview.wss?uid=isg1IJ21527	http://aix.software.ibm.com/aix/ifixes/ij21580/	YES	12/18/2019	4/2/7200
IJ21527	https://www-01.ibm.com/support/docview.wss?uid=isg1IJ21527	http://aix.software.ibm.com/aix/ifixes/ij21527/	YES	12/18/2019	3.1.0.20
IJ21564	http://www-01.ibm.com/support/docview.wss?uid=isg1IJ21564	http://aix.software.ibm.com/aix/ifixes/ij21564/	YES	12/14/2019	3.1.1.20

56

56

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

57

57

Sample /etc/tunables/rc-tunevio.sh

```
#!/bin/ksh
#
# First we set the network tuneables
#
/usr/sbin/no -p -o rfc1323=1
/usr/sbin/no -p -o tcp_sendspace=262144
/usr/sbin/no -p -o tcp_recvspace=262144
/usr/sbin/no -p -o udp_sendspace=65536
/usr/sbin/no -p -o udp_recvspace=655360
#
vmo -p -o minfree=1024
vmo -p -o maxfree=2048

#Run ifconfig -a and check the en values - = you may have to do this on the sea (assuming sea is en6):
#chdev -l en6 -a tcp_recvspace=262144 -P
#chdev -l en6 -a tcp_sendspace=262144 -P
#chdev -l en6 -a rfc1323=1 -P
#
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
chdev -l fcs2 -a max_xfer_size=0x200000 -a num_cmd_elems=1024 -P
chdev -l fcs3 -a max_xfer_size=0x200000 -a num_cmd_elems=1024 -P
```

58

58

Sample /etc/tunables/rc-tunebufs.sh

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

lsdev -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 ent4 -a max_buf_medium=512 -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_medium=512 -P
```

59

59

POST Install Checks

```
$ ioslevel
3.1.0.21

$ oem_setup_env
#oslevel -sq
Known Service Packs
-----
7200-03-03-1914
7200-03-03-1913
.....

#oslevel -s
7200-03-03-1914

#oslevel -s -l 7200-03-03-1914
Should show nothing
# instfix -i | grep ML
All filesets for 7.2.0.0_AIX_ML were found.
All filesets for 7200-00_AIX_ML were found.
All filesets for 7200-01_AIX_ML were found.
All filesets for 7200-02_AIX_ML were found.
All filesets for 7200-03_AIX_ML were found.
```

60

60

POST Install Checks

```
#instfix -icqk 7200-03_AIX_ML | grep :-:
```

```
#lppchk -v
```

```
#lppchk -vm3
```

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

Once all checks are passed and VIO2 is back up check your client LPARs to make sure they see all their paths again

Then go do the same upgrade to VIO1

Don't forget to clean up inetd.conf and other files and then remirror rootvg once you are committed

Back up both VIO servers when done – the backups seem smaller now

61

61

Log Files

On VIOS after viosupgrade command from VIOS

- viosupgrade command logs: /var/adm/ras/ioslogs/*

Look at the viosupg_global.log

- viosupgrade restore logs: /home/ios/logs/viosupg_restore.log
- viosupgrade restore logs: /home/ios/logs/viosupg_status.log
- viosbr backup logs: /home/ios/logs/backup_trace*
- viosbr restore logs: /home/ios/logs/restore_trace*

62

62

Updating - VIOS Problems at 2.2.6.32

```
oem_setup_env
oslevel -s
6100-00-00-0000
Or 7200-00-00-0000
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
```

Using vios 2.2.6 examples as so far no problems with 3.1 upgrade but this will give you the idea

```
# oslevel -sq
Known Service Packs
-----
Top one should be: 6100-09-11-1810
```

63

63

Updating - VIOS Problems at 2.2.6.32

```
# oslevel -s -l 6100-09-11-1810
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
DirectorCommonAgent     6.3.3.1       6.3.5.0
DirectorPlatformAgent   6.3.3.1       6.3.5.0
adde.v2.common.ddk      6.1.9.0       6.1.9.100
adde.v2.ethernet.ddk    6.1.9.15      6.1.9.300
adde.v2.rdma.ddk        6.1.9.100     6.1.9.300
```

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

64

64

Remove or update problem filesets

DO NOT USE SMITTY – use updateios

Issues with bos.suma
updateios –remove bos.suma

```
# oslevel -s -l 6100-09-11-1810
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

65

65

Backup and recovery



66

66

IBM Supported Backup and Restore Methods for VIO Servers

- https://www.ibm.com/support/knowledgecenter/TI0002C/p8hb1/p8hb1_vios_backup_backup.htm
- Note that IBM does not support (even at v3) backup and restore with **USB sticks**

Table 1. Backup and restoration methods for the VIOS

Backup method	Media	Restoration method
To tape	Tape	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

You can backup and restore AIX from USB but VIO is not supported – see below for AIX information
<https://www.ibm.com/support/pages/using-and-taking-advantage-usb-devices-and-aix>

67

67

Backing up VIOS

- The VIO server version of viosupgrade will take an extra viosbr backup for you
- You may also want to use snap to grab other critical data
 - #snap -gtknc
- Mount NFS filesystem to backup to (in my case /backups)
- mkdir /backups/viosa
- umount /var/vio/VMLibrary if you are using the media library or your mksysb will be huge
- Then as padmin run backupvios which automatically calls savevgstruct:
- **backupios -file** {File name} [-mksysb] [-nopak] [-nosvg] [-nomedialib]
- 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 (from HMC)
- You can also back it up as a mksysb file that is easy to restore
- backupios -file /backups/viosa.mksysb -mksysb
- When the **-mksysb** flag is used, the NIMOL resources are not saved in the image.
- If the media library is large and is on rootvg, then you can add the -nomedialib flag, but still unmount it prior to the backup

68

68

Backing up VIOS from root

As root (login as padmin then oem_setup_env) run viosave.sh (see next slide)

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

Mount the NFS repository for the backups (/backups)

```
#mount /backups
```

```
#su - padmin -c "ioscli backupios -file /backups/vio2-sep0919.mksysb -mksysb -nomedialib"
```

This backs it up to a bootable mksysb file

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

```
#mkdir /backups/nimbkups
```

```
#su - padmin -c "ioscli backupios -file /backups/nimbkups -nomedialib"
```

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

```
#su - padmin -c "ioscli viosbr -backup -file viobkup -frequency daily numfiles 7"
```

If you use alt_disk_copy to clone your rootvg disk you have a very fast failback – consider using this as a preupdate backup

69

69

Backup

If you have set the system up to automatically do a daily viosbr backup then you don't need to do it here

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

```
#su - padmin -c "ioscli viosbr -backup -file viobkup -frequency daily numfiles 7"
```

I usually do 2 backups

A regular mksysb

A backup that is for NIM

Both are done to an NFS mount

If you are using the file backed optical (media repository) there seems to be a bug where specifying nomedialib does not stop it backing up the media library. Since mine is over 100GB that is a problem so here is what I do

(vio2dir is a directory)

```
#mkdir /usr/local/backups/vio2dir
```

Create directory if not there

```
#umount /var/vio/VMLibrary
```

```
#su - padmin -c "ioscli backupios -file /usr/local/backups/vio2-previo31-sep0919.mksysb -mksysb -nomedialib"
```

```
#su - padmin -c "ioscli backupios -file /usr/local/backups/vio2dir -nomedialib"
```

My 2.2.6.32 backups are around 19GB but during the backup they can need as much as 32GB – don't ask me why ☹

Once you are done you can remount /var/vio/VMLibrary

70

70

Check the virtual backup

The following adds a cron entry and will backup your VIO virtual definitions every day and keep the last 7 copies in /home/padmin/cfgbackups
You only need to run it once

```
$viosbr -backup -file vio2m -frequency daily -numfiles 7
```

```
$ viosbr -view -list
vio2m.01.tar.gz
```

```
$ ls -al /home/padmin/cfgbackups
total 72
```

```
drwxr-xr-x  2 padmin  staff      256 Mar 27 10:55 .
drwxr-x---  7 padmin  system    4096 Nov 27 12:51 ..
-rw-r--r--  1 padmin  staff     6960 Mar 27 10:55 vio2m.01.tar.gz
```

71

71

Document VIO Information – save-viostuff.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
• imonth="$1"
• machine=$(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"
• 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" >>$logit1.vioslots.txt
• su - padmin -c "ioscli lsmmap -all" >>$logit1.viosmapall.txt
• su - padmin -c "ioscli lsmmap -all -npiv" >>$logit1.viosmapall.npiv.txt
• su - padmin -c "ioscli lsdev -virtual" >>$logit1.viosdevv.txt
• su - padmin -c "ioscli cfnamesrv -ls" >>cfnames.txt
• su - padmin -c "ioscli entstat -all ent9" >>entstat.txt
• su - padmin -c "ioscli hostmap -ls" >>hostmap.txt
• su - padmin -c "ioscli luser" >>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 >>$logit1.oslevel.txt
• gethvdmm -c >>$logit1.disktmp.txt
• while read label line
• do
•   echo "$line" >>$logit1.viodisks.txt
•   echo "Hdisk is $label" >>$logit1.viodisks.txt
•   echo " " >>$logit1.viodisks.txt
•   su - padmin -c "ioscli lsdev -dev $label -attr" >>$logit1.viodisks.txt
• done <<$logit1.disktmp.txt
• #
• exit 0

```

72

72

Continue settings Backup

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.violslots.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
```

73

73

Backup Script to put in crontab

```
#!/bin/sh
#
# viobackup.sh
#
machine=`uname -n`
mount /backups
mkdir /backups/$machine
umount /var/vio/VMLibrary
su - padmin -c "ioscli backupios -file /backups/$machine -nomedialib"
su - padmin -c "ioscli backupios -file /backups/vio-mksysbs/$machine.mksysb -nomedialib -mksysb"
mount /var/vio/VMLibrary
#
exit 0
```

NOTES

The above can be put in root's crontab to run regularly

Don't forget to set up an NFS mount to the VIO from your NIM or NFS server

Do not allow the NFS mount to mount automatically at boot in case the NIM or NFS server is down at the time of boot

Also, regularly grab an HMCScanner report

74

74

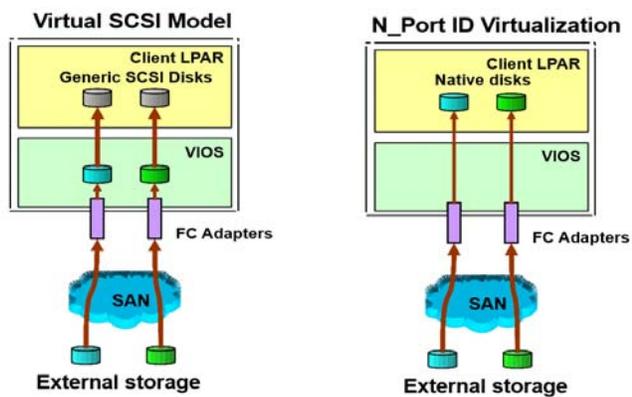
Storage



75

75

The context: PowerVM disk access



This is relevant for LPM as for:

- Zoning
- LUN-masking

76

76

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**
 - 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 C0
- These can be found in an HMCScanner report or by logging onto the LPAR or VIO and using:
- `lscfg -vpl fcs? | grep Network`

77

77

LPM Zoning and mapping - NPIV

- **Do not confuse zoning with mapping**
- Regular and LPM WWPNs must be zoned at the switch and mapped at the storage
- Each virtual fibre adapter for an LPAR has 2 x WWPNs
 - The first is the default one that is used
 - The second is used by LPM – it normally does not login unless LPM has been used
 - Both WWPNs must be zoned and mapped
- If they are not mapped at the storage and you do an LPM you will damage your boot image
 - You can avoid this problem after 2.2.4 by setting 2 parameters on `vioslpm0` on all VIO LPARs
- You should also do your zoning by zoning all WWPNs for the LPAR to both switches. Keep zoning simple and have a zone that is LPARname and all the WWPNs. This will avoid problems during LPM when you allocate fiber ports to each VIO for dual VIO systems.

78

78

vSCSI and NPIV

- vSCSI

Use WWN of the VIO server when zoning then map at VIO server to the client LPAR
 lscfg -vpl fcs0 | grep Network will show something like: 10000090fa530975
 Those WWNs belong to the VIO not the client LPAR
 MPIO drivers are installed in the VIO
 Mirrored in client LPAR
 Disks are seen at and assigned from the VIO
 View using "lsmmap -all"
 Shows as vSCSI in client LPAR
 FCS tunables set in VIO

- NPIV

Use WWPNS that are created when client LPAR is created
 You can find them in the client's profile for the virtual adapters and they look like: c0507607dbd80028
 Those WWPNS (also called VFCs) migrate with the client LPAR
 Disks are not seen at the VIO so MPIO drivers are installed in the client LPAR
 Fibre adapters get mapped from the VIO
 View using "lsmmap -all -npiv"
 Shows as fibre adapters (FCS?) in the client LPAR
 FCS tunables set in VIO and client LPAR – client LPAR settings must be <= to settings in the VIO servers
 On V5000 and some other storage arrays you may also have to zone the real WWNs for the adapters,
 not just the client WWPNS

79

79

NPIV Finding WWPNS

On the HMC

Virtual Fibre Channel Adapter Properties: aix1nim

Virtual Fibre Channel adapter
 Adapter: +70
 Type of adapter: Client
 WWPNS: c0507607dbd80028
 c0507607dbd80029
 This adapter is required for partition activation.
 Server partition: vio1(1)
 Server adapter ID: 70
 OK Cancel Help

Virtual Fibre Channel Adapter Properties: aix1nim

Virtual Fibre Channel adapter
 Adapter: +90
 Type of adapter: Client
 WWPNS: c0507607dbd8002c
 c0507607dbd8002d
 This adapter is required for partition activation.
 Server partition: vio2(2)
 Server adapter ID: 90
 OK Cancel Help

On the Client LPAR

```
aix1nim: /> lsdev -C | grep fcs
fcs0      Available 70-T1  Virtual Fibre Channel Client Adapter
fcs1      Available 71-T1  Virtual Fibre Channel Client Adapter
fcs2      Available 90-T1  Virtual Fibre Channel Client Adapter
fcs3      Available 91-T1  Virtual Fibre Channel Client Adapter
fcs4      Available 80-T1  Virtual Fibre Channel Client Adapter
fcs5      Available 81-T1  Virtual Fibre Channel Client Adapter
aix1nim: /> lscfg -vpl fcs* | grep Network
Network Address.....C0507607DBD80028
Network Address.....C0507607DBD8002A
Network Address.....C0507607DBD8002C
Network Address.....C0507607DBD8002E
Network Address.....C0507607DBD80048
Network Address.....C0507607DBD8004A
aix1nim: />
```

80

80

LPM's use of the two WWPNS

- Each virtual fibre adapter for an LPAR has 2 x WWPNS
 - The first is the default one that is used
 - The second is used by LPM – it normally does not login unless LPM has been used
- Prior to an LPM the default WWPNS is used
- After the LPM the second WWPNS is used
- After the next LPM it goes back to the default WWPNS
- i.e. it flip flops between them
- EXCEPT
 - If you perform an inactive LPM then it stays with whatever the WWPNS were that it used last

81

81

NPIV

<https://www.ibm.com/support/knowledgecenter/TI0002C/p8edm/chnportlogin.html>

```
chnportlogin -o login -m ServerName --id 3
```

Above logs in all ports for LPAR 3

Can also use -d (0-5 default is 1) to get more detail and -v for verbose mode

```
-o logout          logs them out
-n profilename     logs in only those in the specified profile otherwise it uses the current running profile
```

When performing a login operation, all inactive WWPNS will be activated, including the second WWPNS in the pair assigned to each virtual Fibre Channel client adapter. When performing a logout operation, all WWPNS not in use will be deactivated.

<https://www.ibm.com/support/knowledgecenter/TI0002C/p8edm/lspnportlogin.html>
 lspnportlogin -m Server-8286-41A-SN123452X --filter "lpar_names=jaqui" -F lpar_name:wwpn:wwpn_status

```
wwpn_status
```

The WWPNS status. Possible values are:

- 0 - WWPNS is not activated
- 1 - WWPNS is activated
- 2 - WWPNS status is unknown

82

82

NPIV Zoning

<http://www-01.ibm.com/support/docview.wss?uid=isg3T1024487>

The above shows you how to use the HMC Enhanced GUI and login the WWPNS from the profile for the LPAR

On the HMC go to the profile (action, profiles, manage profile then select the profile)

Then virtual adapters

Check all the virtual fibre adapters (called client fibre channel)

Then actions, advanced, login/logout fibre

Click on login to log them all in or logout to logout any not being used

83

83

Simple Zoning for LPM

SIMPLE ZONE

SWITCH1

```
zone:      NPIV_AIXNIM
           c0:50:76:03:ca:b6:00:28
           c0:50:76:03:ca:b6:00:29
           c0:50:76:03:ca:b6:00:2a
           c0:50:76:03:ca:b6:00:2b
           50:05:07:68:02:16:2f:c3
           50:05:07:68:02:16:2f:c4
           10:00:00:90:fa:19:04:40
```

SWITCH2

```
zone:      NPIV_AIXNIM
           c0:50:76:03:ca:b6:00:28
           c0:50:76:03:ca:b6:00:29
           c0:50:76:03:ca:b6:00:2a
           c0:50:76:03:ca:b6:00:2b
           50:05:07:68:02:26:2f:c3
           50:05:07:68:02:26:2f:c4
           10:00:00:90:fa:19:15:a9
           20:02:00:0e:11:13:06:67
```

I zone all the WWPNS for the client on both switches – reduces problems with LPM
If you don't do this, you have to be certain to assign the right adapter to the right VIO when using LPM
The only difference is the storage subsystem zones (50:05) and tape drive units they can see

84

84

Get rid of annoying FCS errors on 10/1Gb cards

- The 4 port 10Gb/1Gb cards can act as network cards or can be used for San. Most people use them as network cards and then see lots of FCS errors.
- You can stop this from happening as follows:

```
#lsdev -C | grep fcs
```

Look for the 10Gb cards - in my case they showed as fcs4 and fcs5

As padmin:

```
$rmdev -dev fcs4 -recursive -ucfg
```

```
$rmdev -dev fcs5 -recursive -ucfg
```

```
$chdev -dev fcs4 -attr autoconfig=defined
```

```
$chdev -dev fcs5 -attr autoconfig=defined
```

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-aug0519a.txt
```

```
#errpt -a >/usr/local/logs/errpta-aug0519a.txt
```

```
#errclear 00
```

```
#cfgmgr
```

```
#errpt
```

85

85

HBA Settings



86

86

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: <http://www-01.ibm.com/support/docview.wss?uid=isg1IV63282>
- **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

87

87

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 0x200000 which provides a 128MB DMA

88

88

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

89

89

Adapter Tuning 2/2

Check these are ok with your disk vendor!!! And also for the adapter.

```
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

DMA area (max_xfer_size) controls the max IO size the adapter can send to the disk subsystem (default is 16MB). To use full bandwidth of adapter this needs to be 128MB.

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

Remember VFCs on the client may be 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/atsmastr.nsf/WebIndex/TD105745>

90

90

My VIO Server and NPIV Client Adapter Settings

VIO SERVER connected to V7000

#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

91

91

MPIO

• MPIO

- IBM is now recommending using the AIXPCM rather than SDDPCM. No new fixes will be provided for SDDPCM after June 30, 2020. See:
 - https://www.ibm.com/support/pages/node/1106937?myns=s035&mync=E&cm_sp=s035- -NULL- -E
- They have a new MPIO best practices document at:
 - <https://www.ibm.com/developerworks/aix/library/au-aix-mpio/index.html>
- There is a good description of AIXPCM here:
 - https://www.ibm.com/support/knowledgecenter/en/ssw_aix_72/com.ibm.aix.osdevice/devmpio.htm
- Migration notes
 - <http://www-01.ibm.com/support/docview.wss?uid=ssg1S1010646>
- Article on MPIO resiliency and problem determination
 - <https://www.ibm.com/developerworks/aix/library/au-aix-multipath-io-mpio/index.html>
- If you need to change defaults (i.e. from single_path to no_reserve)
 - chdef -t mpioosdisk -c disk -s fcp -a reserve_policy=no_reserve
 - Use chdef -H to check for all attributes that have been changed from defaults

92

92

Ensure you install correct drivers and MPIO

Install correct drivers for SAN and other disks
Ensure all disks are set to no_reserve and round_robin (or shortest_queue)

```
lsattr -El hdisk0
```

Check MPIO drivers
lspp -l | grep mpio

```
manage_disk_drivers -l
Switch to new AIXPCM driver
manage_disk_drivers -d IBMSVC -o AIX_AAPCM
manage_disk_drivers -l
```

```
bosboot -a -d hdisk0
bootlist -m normal hdisk0
Shutdown and reactivate
```

Now correct any single paths if still needed:
chdev -l hdisk0 -a algorithm=round_robin -a reserve_policy=no_reserve -P
chdev -l hdisk1 -a algorithm=round_robin -a reserve_policy=no_reserve -P
etc
bosboot -a -d hdisk0
bootlist -m normal hdisk0
shutdown -r now

93

93

Checking Paths with AIXPCM

```
# lspv | grep hdisk0
hdisk0      00f95d3a1b679a90      fbovg      active
```

```
# lsdev -C | grep hdisk0
hdisk0      Available 01-01-01  MPIO IBM 2076 FC Disk
```

```
# lsmpio -ql hdisk0
Device: hdisk0
  Vendor Id: IBM
  Product Id: 2145
  Revision: 0000
  Capacity: 50.00GiB           Shows size
  Machine Type: 2078
  Model Number: 124
  Host Group: Power8_VIO2     Disk subsystem Group name
  Volume Name: P8_VIO2_UTIL1   Disk subsystem volume name
  Volume Serial: 60050763808100F70000000000000004 (Page 83 NAA) Sows serial number for LUN
```

94

94

lsmPIO

Shows if paths are optimized or not

```
# lsmPIO -l hdisk0
```

name	path_id	status	path_status	parent	connection
hdisk0	4	Enabled	Non	fscsi2	500507680d108ef6,0
hdisk0	5	Enabled	Sel,Opt	fscsi2	500507680d108ef7,0
hdisk0	6	Enabled	Non	fscsi3	500507680d0c8ef6,0
hdisk0	7	Enabled	Sel,Opt	fscsi3	500507680d0c8ef7,0

Use `chpath` if you need to change priorities for paths

You can also find the parent for an `hdisk` and get statistics on each of the paths using:

```
lsmPIO -are -l hdisk4
```

This will show a long report that includes all the paths. An example of part of one of those reports is below:

Other options:

```
lsmPIO -l hdisk4 -sd
```

Detailed statistics for the `hdisk`

`lspath` was updated with new `-t` and `-i` flags. The `-t` flag ensures the `pathid` is listed at the end, and the `-i` flag allows you to specify the `pathid` and only get information for devices on that specific path. The `mkpath` and `rmpath` commands were also updated to include the option to process specific `pathids`.

```
lspath -t -l hdisk4
```

Include `pathid` in report

```
lspath -l hdisk4 -i 0
```

Only show `path 0`

```
lspath -i 0
```

Report on all disks on `path 0`:

95

95

Network



96

96

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 (in VIO and client)

If VIOS is only handling the 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>

97

97

Large Receive and Send

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

See <http://tinyurl.com/gpe5zgd> for information for Linux and Large send/receive
 This provides information on correctly using Large send and Large receive with Linux

Also <http://tinyurl.com/lm6x5er> for info for large send in general and also IBM i

LRO (large receive offload) is enabled by default on virtual ethernet resources
 The issue with IBM i and LRO was resolved in the base code for IBM i 7.2 so you are safe to enable LRO at that level
 Below IBM i 7.2 there are patches you need to install

98

98

VIO 2.2.3 and above SEA Changes

Traditional SEA setup

ent0-3 are the physical adapters

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 or you can leave this out and let it default to 4095 on mkvdev

OLD

Add a virtual network to the profile to be used for the control channel (used vlan 255 in this case)

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

Creates ent7 as the SEA and uses ent6 for the control channel

NEW

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

Above creates ent7 as SEA and defaults to vlan 4095 for control channel

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

99

99

Correct setup of switch ports for SEA Failover

- Incorrect setup can cause SEA failover delays which may disconnect users
- <https://www.ibm.com/support/pages/sea-failover-delayed-and-supported-methods-test-failover>
- If Spanning Tree is turned on, then Portfast also should be turned on.
 - Portfast is a Cisco term
 - The IBM switch equivalent is Spanning Tree Edge
- To reduce fallback time you may want to turn Spanning Tree off

100

100

SEA Failover Testing

After you set up a VIO pair you should test that failover is working
 The quick way is to ssh to the primary and reboot the secondary and see if you lose your ssh connection
 Then ssh to the secondary and reboot the primary and see if you lose the ssh connection

If the switch ports are set up correctly and the SEA is defined correctly there should be no issues

Also check errpt to make sure that the primary and secondary (backup) correctly become primary and backup during the process

When performing maintenance on the primary network VIO LPAR or rebooting it, I normally force a failover manually by doing the following:

If the SEA is set up correctly then this is not necessary, but it ensures the connectivity of the client partitions that are using the SEA in PRIMARY state do not lose connectivity temporarily.

Prior to maintenance on the primary VIOS

1. Set ha_mode to standby on primary VIOS with chdev command:
`$ chdev -dev entX -attr ha_mode=standby`

After reboot or maintenance is complete:

2. Reset it back to auto and the SEA should fail back to the primary VIOS:
`$ chdev -dev entX -attr ha_mode=auto`

101

101

Starter set of tunables - Network

Typically we set the following – they are not set by default:

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

102

102

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>

103

103

More complex Networking

Resources, All Systems, Server, Virtual Networks

The hypervisor provides an IEEE 802.1Q virtual LAN style virtual Ethernet switch. When you add a virtual network, you can add a virtual switch. The default is ethernet0.

You can add Virtual networks, switches, bridges and VLANs

You can also create your etherchannels here

Networks can be internal or bridged

They can have no tagging or 802.1q tagging

They can be set for single VIO, dual VIO, loadsharing (load groups) or just regular

You can create VLANs and spread them across networks

https://www.ibm.com/support/knowledgecenter/TI0002C/p8efd/p8efd_virt_switch_concept.htm

https://www.ibm.com/support/knowledgecenter/en/9119-MHE/p8efd/p8efd_add_new_virt_net_wizard_task.htm

There are also options such as vnic and SR-IOV – Alexander Paul has great presentations dedicated to these

NOTE – map out all your networks and vlans before you even consider going down this path – it gets complicated very fast

104

104

SEA Loadsharing

Allows you to have some VLANs primary on VIO1 and backed up on VIO2, with other VLANs primary on VIO2 and backed up on VIO1.

This lets you take advantage of the full bandwidth of the adapters

<https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Power%20Systems/page/SEA%20Load%20Sharing>

<https://www.ibm.com/support/pages/how-setup-sea-failover-load-sharing-configuration>

2 options

1. Use `ha_mode=sharing` as per the above – let you have a single SEA with multiple VLANs
`ha_mode=sharing` must be set on the primary SEA before the backup
SEA evenly divides traffic up by adapter (not VLAN) between the two VIO LPARs
2. Define two SEAs (with their own adapters)
SEA1 is primary on VIO1 (priority=1)
SEA2 is primary on VIO2 (priority=1)
Assign SEAs to clients depending on which vio you want to be primary for that client

105

105

Monitoring



106

106

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
 - Check your %idle first as you can go over entitlement if you have too many VPs, yet not be using all the threads on the cores
- 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
- Plan for cores and memory for VIO servers when sizing systems
 - At least 2-3 cores for a pair and 8Gb minimum each

107

107

nmon Monitoring

- **nmon -ft -AOPV^dML -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 you can add this

If you want a 24 hour nmon use:

```
nmon -ft -AOPV^dML -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

108

108

Shared Processor Pool Monitoring

Turn on “Allow performance information collection” on the LPAR properties

This is a dynamic change

Without this being set on every LPAR the cross LPAR statistics won't be correct

This includes APP and other statistics reported by nmon and lparstat, etc

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

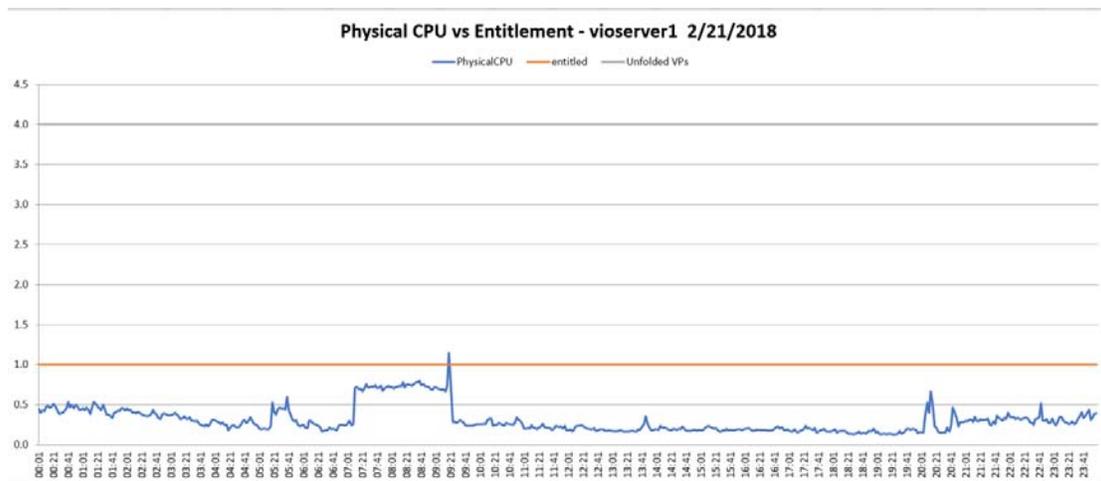
lparstat

Shows the app column and poolsize

109

109

nmon Analyser LPAR Tab



110

110

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
 - <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 as padmin
- Shows all clients for the vio, the WWPns, statistics and error counts
- Also check out NPIVGRAPH for visualizing NPIV mappings:
 - <http://npivgraph.sourceforge.net/>
 - Review options on fcstat – fcstat -d and fcstat -e provide additional statistics on adapter usage
 - https://www.ibm.com/support/knowledgecenter/en/ssw_aix_61/com.ibm.aix.cmds2/fcstat.htm

111

111

SEA Statistics

- Output in part command
- SEA Adapter Statistics
 - Provides details on the meaning of statistics for the SEA
 - https://www.ibm.com/support/knowledgecenter/POWER9/p9hb1/p9hb1_statssea.htm
- SEA Adapter Failover Statistics
- Provides details on meaning of output from entstat -all specific to SEA Failover Statistics
 - https://www.ibm.com/support/knowledgecenter/POWER9/p9hb1/p9hb1_statsseafailover.htm
- VIO Server Network Attributes Information
 - https://www.ibm.com/support/knowledgecenter/POWER9/p9hb1/p9hb1_vios_managing_sea_attr.htm

112

112

netstat -v (vio or client Virtual Ethernet)

```
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

Virtual I/O Ethernet Adapter (I-lan) Specific Statistics:
-----
Hypervisor Send Failures: 4043136
Receiver Failures: 4043136
Send Errors: 0
Hypervisor Receive Failures: 67836309
```

check those tiny, etc Buffers
Check on client LPARs too

"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

113

113

Buffers – check at vio and client

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

entstat -d ent7 (where ent7 is the SEA) gets you the information for ent7 only

114

114

Thank you for your time



If you have questions please email me at:
jaqui@circle4.com or jlynch@flagshipsg.net

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

Copy of presentation at:
<http://www.circle4.com/ptechu/vioscareandfeeding-Jan302020.pdf>

And the Virtual User Group
<https://www.ibm.com/support/pages/node/1120377>

115

115

115

Useful Commands, Links and Documentation



116

116

USEFUL COMMANDS

117

117

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
```

118

118

Useful Commands

vSCSI Commands

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

```
mkvdev -fbo -vadapter vhost0
```

NPIV

Setup NPIV mappings

```
vfcmap -vadapter vfchost0 -fcp fcs0
```

```
lsmap -npiv -all
```

```
lsmap -vadapter vfchost0 -npiv
```

```
lsdev -virtual
```

```
lsnports
```

```
lsdev -slots
```

```
lscfg -vpl vfchost0
```

119

119

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
vsao	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

120

120

Useful Commands

\$ lsmmap -vadapter vhost0

SVSA	Physloc	Client Partition ID
vhost0	U8205.E6B.1093XXX-V1-C21	0x00000003

VTD	b740l1_rv1
Status	Available
LUN	0x8300000000000000
Backing device	lv_b740l1
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

121

121

Useful Commands

\$ lsmmap -vadapter vfchost0 -npiv

Name	Physloc	ClntID	ClntName	ClntOS
vfchost0	U8205.E6B.1093XXX-V1-C31		3	

Status:NOT_LOGGED_IN
 FC name:fcs0 FC loc code:U78AA.001.WZSG8XX-P1-C5-T1
 Ports logged in:0
 Flags:4<NOT_LOGGED>
 VFC client name: VFC client DRC:

\$ lsmmap -vadapter vfchost4 -npiv

Name	Physloc	ClntID	ClntName	ClntOS
vfchost4	U8205.E6B.1093XXX-V1-C36		8 b740nl1	AIX

Status:LOGGED_IN
 FC name:fcs0 FC loc code:U78AA.001.WZSG8XX-P1-C5-T1
 Ports logged in:3
 Flags:<LOGGED_IN,STRIP_MERGE>
 VFC client name:fcs0 VFC client DRC:U8205.E6B.1093XXX-V8-C36

122

122

Useful Commands

\$ lsports

```
name      physloc      fabric tports aports swwpns awwpns
fcs0      U78AA.001.WZSG8XX-P1-C5-T1  1  64  63  2048  2041
```

\$ lsdev -slots

```
# Slot      Description      Device(s)
HEA 1      Logical I/O Slot lhea0 ent0
U8205.E6B.1093XXX-V1-C0  Virtual I/O Slot vsa0
U8205.E6B.1093XXX-V1-C11 Virtual I/O Slot ent5
U8205.E6B.1093XXX-V1-C12 Virtual I/O Slot ent6
U8205.E6B.1093XXX-V1-C13 Virtual I/O Slot ent7
U8205.E6B.1093XXX-V1-C21 Virtual I/O Slot vhost0
U8205.E6B.1093XXX-V1-C22 Virtual I/O Slot vhost1
U8205.E6B.1093XXX-V1-C23 Virtual I/O Slot vhost2
U8205.E6B.1093XXX-V1-C31 Virtual I/O Slot vfchost0
U8205.E6B.1093XXX-V1-C32 Virtual I/O Slot vfchost1
U8205.E6B.1093XXX-V1-C33 Virtual I/O Slot vfchost2
U8205.E6B.1093XXX-V1-C32769 Virtual I/O Slot vasi0
U8205.E6B.1093XXX-V1-C32773 Virtual I/O Slot vasi1
U8205.E6B.1093XXX-V1-C32774 Virtual I/O Slot vasi2
U8205.E6B.1093XXX-V1-C32775 Virtual I/O Slot vasi3
U8205.E6B.1093XXX-V1-C32776 Virtual I/O Slot vasi4
```

123

123

Documentation on VIOS 3.1 upgrades

- What's new in Virtual I/O Server commands
- https://www.ibm.com/support/knowledgecenter/en/9040-MR9/p9hcg/p9hcg_whatsnew.htm
- Virtual I/O Server release notes – include USB Memory/Flash key install
- https://www.ibm.com/support/knowledgecenter/en/9040-MR9/p9eeo/p9eeo_ipeeo_main.htm
 - USB Memory/Flash key install
 - Duff minimum size for a VIOS
- **VIOS viosupgrade** command in VIOS 2.2.6.30
- https://www.ibm.com/support/knowledgecenter/en/9009-42A/p9hcg/p9hcg_viosupgrade.htm
- **NIM viosupgrade** command on the NIM AIX 7.2 TL3 + sp
- https://www.ibm.com/support/knowledgecenter/en/ssw_aix_72/com.ibm.aix.cmds6/viosupgrade.htm
 - This one is buried in the AIX commands reference for AIX Commands of AIX 7.2

124

124

Useful Links

- Jaqui Lynch Articles
 - <http://www.circle4.com/jaqui/eserver.html>
 - <https://ibmsystemsmag.com/Authors/jaqui-lynch>
 - <http://archive.ibmsystemsmag.com/authors/jaqui-lynch/?page=1>
- Nigel Griffiths AIXpert Blog
 - <https://www.ibm.com/support/pages/aixpert-blog-nigel-griffiths-mrnmon>
- Nigel Griffiths Twitter – mr_nmon
 - https://twitter.com/mr_nmon
- Gareth Coates – Tricks of the POWER Masters
 - <https://www.ibm.com/support/pages/node/1116939>
- Gareth Coates Twitter – power_gaz
 - https://twitter.com/power_gaz
- Jaqui's Movie Replays
 - <http://www.circle4.com/movies>
- IBM US Virtual User Group
 - <https://www.ibm.com/support/pages/node/1120377>
- Power Systems UK User Group
 - <https://www.ibm.com/support/pages/node/1110195>

125

125

Useful Links

- ESS Website to download base software
 - <https://www.ibm.com/servers/eserver/ess/index.wss?lnk=msdDO-enss-usen>
 - HMC Scanner
 - <https://www.ibm.com/support/pages/node/1117515>
 - [https://www.ibm.com/support/pages/sites/default/files/inline-files/\\$FILE/hmcScanner-0.11.42.zip](https://www.ibm.com/support/pages/sites/default/files/inline-files/$FILE/hmcScanner-0.11.42.zip)
- NOTE – developerworks is going away 3/31/2020 – some of these URLs will either go away or will be moved so please download anything you may need ASAP**
- Performance Tools Wiki
 - AIX Performance Tools and Commands
 - <https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Power%20Systems/page/AIX%20Performance%20Commands>
 - Performance Monitoring Tips and Techniques
 - <https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Power%20Systems/page/Performance%20Monitoring%20Tips%20and%20Techniques>
 - Other Performance Tools
 - <https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Power+Systems/page/Other+Performance+Tools>
 - Includes new advisors for Java, VIOS, Virtualization
 - VIOS Advisor
 - <https://www.ibm.com/developerworks/community/wikis/home?lang=en#/wiki/Power%20Systems/page/VIOS%20Advisor>
 - https://www.ibm.com/support/knowledgecenter/TI0002C/p8hcg/p8hcg_part.htm

126

126

References

- Technical Introduction and Overview Redbooks
 - Got to <http://www.redbooks.com> and search for the above redbook for your server
 - As an example the E980 Redbook is:
 - <http://www.redbooks.ibm.com/redpapers/pdfs/redp5510.pdf>
- Processor Utilization in AIX by Saravanan Devendran
 - <https://www.ibm.com/developerworks/mydeveloperworks/wikis/home?lang=en#wiki/Power%20Systems/page/Understanding%20CPU%20Utilization%20on%20AIX>
- Rosa Davidson Back to Basics Part 1 and 2 –Jan 24 and 31, 2013
 - <https://www.ibm.com/developerworks/mydeveloperworks/wikis/home?lang=en#wiki/Power%20Systems/page/AIX%20Virtual%20User%20Group%20-%20USA>
- 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>
- SG24-8171 – Power Systems Performance Optimization including POWER8
 - <http://www.redbooks.ibm.com/redbooks/pdfs/sg248171.pdf>
- SG24-8453 - AIX Modernization and Enhancements
 - <http://www.redbooks.ibm.com/redbooks/pdfs/sg248453.pdf>

1/30/2020

127

127

Useful Articles

Conduct an end of year AIX Health Check (Dec 2019)

<https://ibmsystemsmag.com/Power-Systems/12/2019/Conduct-AIX-Systems-Health-Check>

Using NIM with VIO Servers

<https://ibmsystemsmag.com/Power-Systems/09/2019/Using-NIM-with-VIO-Servers>

PowerVM v3 Installation and Upgrade Experience

<https://ibmsystemsmag.com/Power-Systems/05/2019/powervm-experience>

Systems Management Tips

<https://ibmsystemsmag.com/Power-Systems/08/2019/2019-AIX-Systems-Management-Tips>

Using FLRT and FLRTVC

<http://archive.ibmsystemsmag.com/aix/administrator/systemsmanagement/flrt-flrtvc-aix-fixes/>

128

128

Additional Articles

Complete Guide to System Maintenance

<http://archive.ibmssystemsmag.com/aix/tipstechniques/systemsmanagement/the-complete-guide-to-system-maintenance/>

OR <http://tinyurl.com/hbbcefr>

Secure your VIO Server

<http://archive.ibmssystemsmag.com/aix/administrator/security/secure-your-vio-server/>

Upgrading your VIO server – July 2018

- <http://ibmsystemsmag.com/aix/administrator/systemsmanagement/upgrading-your-vio-server/>

Maintaining the HMC

- <http://ibmsystemsmag.com/aix/administrator/systemsmanagement/hmc-maintenance/>

LPM

- <http://ibmsystemsmag.com/aix/administrator/systemsmanagement/A-Step-By-Step-Guide-to-Live-Partition-Mobility/>

HMC Enhanced GUI Links

- <https://www.ibm.com/support/pages/enhanced-gui-links-documentation>

129

129

VIOS Specific References

- 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/developerworks/community/wikis/home?lang=en#/wiki/Power%20Systems/page/SEA%20Load%20Sharing>
- 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>

130

130

VIOS Specific References - Network

- SEA Failover Statistics
 - https://www.ibm.com/support/knowledgecenter/POWER9/p9hb1/p9hb1_statsseafailover.htm
- SEA Statistics
 - https://www.ibm.com/support/knowledgecenter/POWER9/p9hb1/p9hb1_statssea.htm
- Virtual Network Management with enhanced HMC GUI
 - <https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Power%20Systems/page/Virtual%20Network%20Management%20with%20HMC%20Enhanced%20UI/version/59aa40ea-867b-4028-bc6e-786efcff5fa5>
- Using SR-IOV for Optimal Performance
 - <https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Power%20Systems/page/Hybrid%20Network%20Virtualization%20-%20Using%20SR-IOV%20for%20Optimal%20Performance%20and%20Mobility>
- Configure VIO Server using VLAN Tagging
 - https://www.ibm.com/support/knowledgecenter/POWER8/p8hb1/p8hb1_vios_scenarios_network_two.htm
- VLAN Tagging – Load sharing with 10Gb adapters (PPT)
 - https://www.ibm.com/support/knowledgecenter/POWER8/p8hb1/p8hb1_vios_scenarios_network_two.htm

131

131

Backup Slides



132

132

HMC Levels

▲ Version	◆ Recommended Update	◆ Recommended Upgrade	◆ Release Date	◆ EoSPS Date
V8 R870 SP3			2019.04.01	2019.08.31
V9 R1 M910	V9 R1 M921		2018.03.20	2021.04.30
V9 R1 M911	V9 R1 M921		2018.05.25	2021.04.30
V9 R1 M920	V9 R1 M921		2018.08.17	2021.04.30
V9 R1 M921			2018.11.16	2021.04.30
V9 R1 M930			2019.05.17	2021.04.30
V9 R1 M931			2019.09.11	2021.04.30
V9 R1 M940			2019.09.22	2021.04.30

133

133

HMC levels

All HMC levels prior to v9 went out of service by 8/31/2019

V9 goes out of service 4/30/2021

<http://www14.software.ibm.com/webapp/set2/flrt/liteTable?prodKey=hmc>

HMC latest version is v9R1M940 - (9/22/2019) – prereq is v9R1.910.0 min.

Can upgrade to v9r1.910.0 from v8.8.6.0 sp1 or later

V9R1M910 (MH01733 – x86 or MH01735 – PPC):

<https://delivery04.dhe.ibm.com/sar/CMA/HMA/07hbb/6/MH01735.readme.html>

v9R1M940 – (MH01837 – PPC, MH01836 – x86)

<https://delivery04.dhe.ibm.com/sar/CMA/HMA/08mn9/0/MH01836.readme.html>

No patches currently for M940. If you are at m921 you must go to m930 before updating to m940.

Note - v9.1.m940 is the last HMC release that will support x86 HMCs

V9.1 requires the HMC to be a CR7 or higher if Intel, or the new POWER HMC

V9.1 does not support any server prior to POWER7

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>

NOTE – once HMC is at v9r1m920 or higher you can upload VIOS and other images from flash drive to the HMC

V9 only supports the enhanced mode GUI

NOTE there is new BMC and PNOR code as of 12/3/2019

https://delivery04.dhe.ibm.com/sar/CMA/SFA/08nhu/1/7063-CR1_OpenPowerReadme.op825.40.xhtml

134

134

Memory Planning

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

Note div - use 64 for all pre p7+ and IBM I, - 128 for p7+ and p8

POWER Systems Memory Overhead Approximation Calculator
USE AS IS - NO GUARANTEES - UPDATED 8/24/2017
 Complete the information below so that calculations will be accurate

Memory Installed in box in MB	393216
Memory active in box in MB	194560
LMB size for server	256
Extra High performance adapter ports per VIO	8 This is active 10Gb net, 8gb fibre etc ports (not adapters)
These include 10Gb network and 8Gb fibre	
VFCs (NPIV) per VIO server	12 Each NPIV client
I/O drawers attached	2
POWER6 only - IVE/HEA ports active	0 Change to number of ports in use
safety net for memory in MB	512
Active memory mirroring?	2 Set to 2 if using mirroring
Divisor	128 Set to 128 if p7+ or P8

Spreadsheet assumes 2 x VIO servers configured equally

This spreadsheet is an approximation - the author takes no responsibility for the output
Use at your own risk
Output should be compared to the output from:
 IBM SPT <http://www-947.ibm.com/systems/support/tools/systemplanningtool/>
 IBM WLE <http://www-912.ibm.com/wle/EstimatorServlet>
 Questions can be sent to Jaqui@circle4.com

Cover Sheet

135

135

Memory Planning Worksheet
USE AS IS - NO GUARANTEES - UPDATED 8/24/2017

Max RAM Capacity	786432	Ram Installed	393216	Ram Active	194560	This gives a rough estimate Assumes LMB size is 256MB Each active IVE port adds 102 MB		
		GB	384	GB	190			
Change the LMB size on this line to match MRO on HMC				MB LMB =	256	Used the largest to show worst possible		
				Extra high performance ports per VIO	8	NPIV VFCs per VIO		
					12			
LPAR NAME	Desired Memory MB	Maximum Memory MB	Overhead Max Div 64 or 128	OH/LMB MB	Roundup OH MB	Actual Overhead (MB) OH * LMB	Memory Needed	Extra High If NPIV Perf ports
VIOS1	3172	4096	32	0.13	1	256	4096	1680
VIOS2	3172	4096	32	0.13	1	256	4096	1680
LPAR1	12032	16384	128	0.50	1	256		
LPAR2	20224	24576	192	0.75	1	256		
LPAR3	14336	16384	128	0.50	1	256		
LPAR4	16384	24576	192	0.75	1	256		
LPAR5	3072	4096	32	0.13	1	256		
LPAR6	2048	4096	32	0.13	1	256		
LPAR7	17152	17152	134	0.52	1	256		
LPAR8	65536	71680	560	2.19	3	768		
LPAR9	32768	36864	288	1.13	2	512		
HYPERVISOR							1536	
IVE							0	
I/O drawer (1 use 512 per 2)							512	
Safety Net							512	
MB Total	189896	224000	1750	6.8359375	14	6144	196040.00	8192
GB Total	185					6.00	191.45	8.00
							CB Total	3.28
							Total when Add High Perf	199.45
							Or add NPIV	194.73
							Or BOTH	202.73
							Combined New Memory needed including Overhead total	202.73
							Shortfall (needed - active)	12.73

Actual Data

NOTES
 Hypervisor requires 7GB minimum for overhead with these settings for maximum memory
 LPARs require 185GB so the total active needed is at least 192GB just to cover maximum memory setting overhead
 Need to add NPIV and high speed adapter memory needs 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

This spreadsheet is an approximation - the author takes no responsibility for the output
Use at your own risk
Output should be compared to the output from:
 IBM SPT <http://www-947.ibm.com/systems/support/tools/systemplanningtool/>
 IBM WLE <http://www-912.ibm.com/wle/EstimatorServlet>

136

136

Changes to Fix Central

- IBM has moved from anonymous FTP to Secure FTP
- <http://www-01.ibm.com/support/docview.wss?uid=isg3T1024541>
- On AIX this means you will be provided with a userid and password to login when you request the fixes
- `ftp -s -i delivery04-mul.dhe.ibm.com`
- When prompted for userid and password use the ones provided
- `passive` (to set passive mode)
- `binary` (to download as binary)
- `mget *` (to download fixes)
- `Quit`

You can also use `sftp` – i.e. once they give you a userid and password:

`sftp user@delivery04-mul.dhe.ibm.com`

Put in password when prompted then type in “`mget *`” then quit when done

137

137

MPIO Algorithms

I/O algorithms

One attribute that can be changed is the `algorithm` attribute. It determines the way I/O is distributed across the paths that are available for a device. There are three different algorithms for I/O distribution:

<code>fail_over</code>	In this mode, I/O is sent over one enabled path, and another path is used only if this path fails. The initial path that is selected is the one with the highest priority (the lowest path priority value). If this path fails, the next highest priority path is selected for I/O operations. This is the default algorithm for iSCSI and FC devices.
<code>round_robin</code>	In this mode, I/O is distributed across multiple enabled paths. For any device that has preferred and non-preferred paths or active and passive paths, only a subset of paths is used. Paths that have a higher path priority receive a larger share of I/O operations in this mode.
<code>shortest_queue</code>	In this mode, an I/O path is distributed across multiple enabled paths. For any device that has preferred and non-preferred paths or active and passive paths, only a subset of paths is used. In this mode, path priority is ignored, and paths are selected based on the number of pending I/O operations only. The path with the lowest number of pending I/O operations is selected for I/O.

From: SG24-8453 - AIX Modernization and Enhancements - <http://www.redbooks.ibm.com/redbooks/pdfs/sg248453.pdf>

138

138

MPIO Health Check Mode

Health check mode and interval

Another key device attribute is the `hcheck_mode` attribute. This attribute determines which paths are probed when the health check capability is used. Health checking is only performed on devices that have a state of open. A device that is not in use does not have its paths health checked. Health checking is also not performed on any disabled or missing paths. There are three health check modes:

Enabled	In this mode, the <code>healthcheck</code> command is sent to all paths that are enabled for the device, which includes paths that failed.
Failed	In this mode, the <code>healthcheck</code> command is sent to all paths that are in a failed state for the device.
Nonactive	In this mode, the <code>healthcheck</code> command is sent to all paths that do not have any active I/O, which includes paths that are in enabled and failed states. This is the default health check mode that is configured on AIX.

Along with `hcheck_mode`, you can also configure how often the health check is performed by configuring the `hcheck_interval` value. This attribute can be set to any value 0 - 3600, and it represents the time in seconds between polling. If a value of 0 is specified, it indicates that health checking should be disabled on the device. The default value for `hcheck_interval` is set to perform health checking every 60 seconds.

From: SG24-8453 - AIX Modernization and Enhancements - <http://www.redbooks.ibm.com/redbooks/pdfs/sg248453.pdf>

139

139

MPIO Reservation Policies

Reservation policies

If your disks require concurrent access from multiple initiators, another attribute you might need to modify is the device attribute `reserve_policy`. This device attribute is required for all MPIO devices regardless of the PCM in use. This value describes the type of reservation policy that is set on a device. For MPIO devices, the following reservation policies exist:

no_reserve	This policy does not apply any reservation on the target device allowing initiators (paths) on the same system, and on other systems, access to the target device. This is the recommended policy for devices where disks are shared between hosts and devices that have the <code>shortest_queue</code> or <code>round_robin</code> algorithms configured.
single_path	This is the default policy when using AIXPCM. This policy places an SCSI2 reserve on a target device so that the device can be accessed only on the path it was reserved on. This policy prevents other paths on the same system from accessing the storage without first sending a bus device reset to release the reserve on the device.
PR_exclusive	This policy applies an SCSI3 persistent-reserve with exclusive-host methodology on the device when the device is opened to exclusively lock it to a single host. A <code>PR_key_value</code> attribute must also be set on the device when using this mode to uniquely identify the host.
PR_shared	This policy applies an SCSI3 persistent-reserve with shared-host methodology when the device is opened. Initiators from other host systems must register before they can access the device. A <code>PR_key_value</code> attribute must also be set on the device when using this mode to uniquely identify the host.

From: SG24-8453 - AIX Modernization and Enhancements - <http://www.redbooks.ibm.com/redbooks/pdfs/sg248453.pdf>

140

140

SEA with link Aggregate and dual VIO Servers

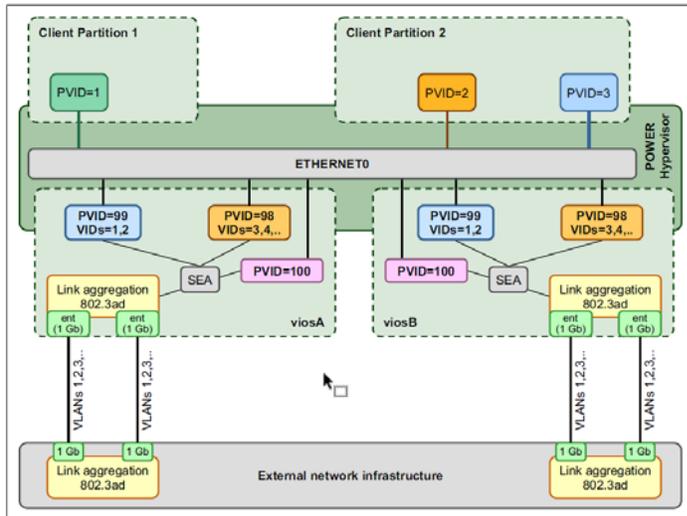


Figure 4-2 Dual Virtual I/O Server configuration with SEA and load balancing

The SEA is connected to two virtual Ethernet adapters; each adapter has a different set of VLAN IDs. Virtual Ethernet adapters on different Virtual I/O Servers must have the same set of VLANs and different trunk priorities.

In addition, the SEA is connected to another virtual Ethernet adapter with a PVID=100 that is used as a SEA control channel.

The control channel is used for SEA heartbeating and exchanging information between the two SEA adapters on the set of VLAN IDs that each SEA bridges.

From PowerVM Best Practices Red Book SG24-8062

NOTE – as of VIO 2.2.3 we let the control channel default to VLAN 4095 – we no longer define it.

141

SEA Loadsharing

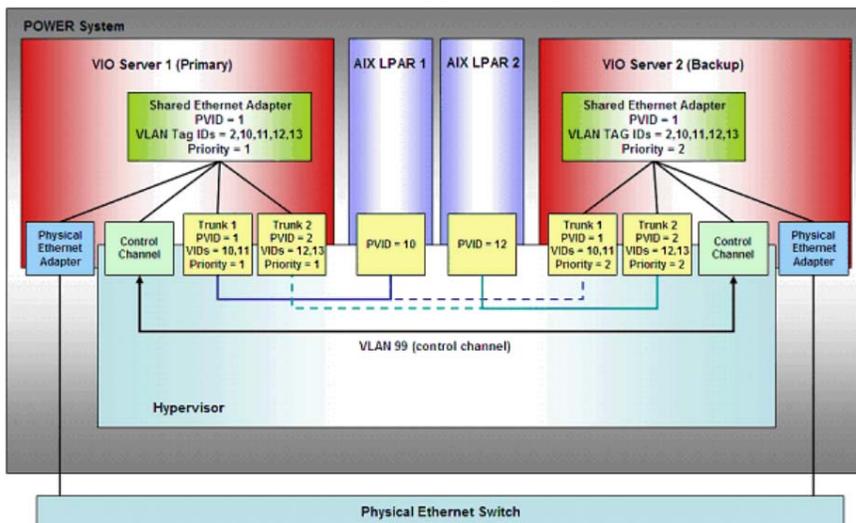


Figure 4: SEA with Load Sharing

142

fcstat -client on VIO1

```
fcstat -client
hostname dev wwpn inreqs outreqs ctrlreqs inbytes outbytes DMA_errs Elem_errs Comm_errs
vio1 fcs0 0x10000090FA530BE2 85802522 87959934 2898600 1166883705621 566506844660 0 0 0
aixlinim fcs0 0xC0507607DBD80028 3234 50126 99840 105863136 246943744 0 0 0
-----
vio1 fcs1 0x10000090FA530BE3 66640483 69920900 2899294 941231200735 458546227132 0 0 0
aixlinim fcs1 0xC0507607DBD8002A 418 50060 83211 22311712 239587328 0 0 0
```

143

143

part command

https://www.ibm.com/support/knowledgecenter/en/POWER8/p8hcg/p8hcg_part.htm

Part is used to gather VIO specific statistics and to create a .xml file that shows the information gathered

On vios as padmin run:

part **{-i interval | -i filename}** **[-t level]** **[-help | -?]**

Minimum of 10 minutes, maximum of 30 minutes

part -i 30

part: Reports are successfully generated in vio1_181007_15_22_16.tar

Runs a 30 minute nmon and gathers other data during that 30 minutes

Creates the .nmon file and a .xml file then tars them up into the current directory

You can run it against a current .nmon file as follows:

part -f filename.nmon

Output fields are described at:

https://www.ibm.com/support/knowledgecenter/en/POWER8/p8hb1/p8hb1_vios_perf_adv_reports.htm

144

144

part command

Advisory Report

System - Configuration

Name	Value
Processor Family	Architecture PowerPC Implementation POWER7_COMPAT_mode 64 bit
Server Model	IBM 8286-41A
Server Frequency	3724.0 MHz
Server - Online CPUs	3.0 cores
Server - Maximum Supported CPUs	3.0 cores
VIOS Level	2.2.6.21
VIOS Advisor Release	0.1

VIOS - I/O Activity

Name	Value
Disk I/O Activity	Insufficient Data from recording
Network I/O Activity	[Average Send: 0 @ 0.0 MBps , Average Receive: 0 @ 0.0MBps] [Peak Send: 0 @ 0.0 MBps , Peak Receive: 0 @ 0.0MBps]

145

145

part command

VIOS - Disk Adapters

Risk/Impact 1=lowest 5=highest

	Name	Measured Value	Suggested Value	First Observed	Last Observed	Risk	Impact
	FC Adapter Count	2		10/07/2018 03:22 PM			
	FC I/O Operations per second	569 @ 38 KB		10/07/2018 03:22 PM	10/07/2018 03:32 PM		
	FC Adapter Utilization	optimal					
	NPIV Client Utilization - fcs1	High: 0.00 % Average: 0.00 %		10/07/2018 03:22 PM	10/07/2018 03:32 PM		
	NPIV Client Utilization - fcs0	High: 0.00 % Average: 0.00 %		10/07/2018 03:22 PM	10/07/2018 03:32 PM		
	FC I/O Operations Blocked	optimal		10/07/2018 03:22 PM	10/07/2018 03:32 PM		
	FC Port Speeds	running at full speed		10/07/2018 03:22 PM	10/07/2018 03:32 PM		

VIOS - Disk Drives

Risk/Impact 1=lowest 5=highest

	Name	Measured Value	Suggested Value	First Observed	Last Observed	Risk	Impact
	Physical Drive Count	19		10/07/2018 03:22 PM			
	I/O Operations Blocked	pass		10/07/2018 03:22 PM	10/07/2018 03:32 PM		
	Long I/O Latency	pass		10/07/2018 03:22 PM	10/07/2018 03:32 PM		

146

146

part command

VIOS - Disk Adapters Risk/Impact 1=lowest 5=highest

Name	Measured Value	Suggested Value	First Observed	Last Observed	Risk	Impact
FC Adapter Count	2		10/07/2018 03:22 PM			
FC I/O Operations per second	569 @ 38 KB		10/07/2018 03:22 PM	10/07/2018 03:32 PM		
fcs1	Average : 282 @ 19 KB		10/07/2018 03:22 PM	10/07/2018 03:32 PM		
fcs0	Average : 287 @ 18 KB		10/07/2018 03:22 PM	10/07/2018 03:32 PM		
FC Adapter Utilization	optimal					
FC Adapter Utilization (fcs1)	high:12.9% (799.0 @ 73.9K)		10/07/2018 03:22 PM	10/07/2018 03:32 PM		
FC Adapter Utilization (fcs0)	high:13.3% (822.0 @ 67.2K)		10/07/2018 03:22 PM	10/07/2018 03:32 PM		
NPIV Client Utilization - fcs1	High: 0.00 % Average: 0.00 %		10/07/2018 03:22 PM	10/07/2018 03:32 PM		
aix1nim	Average 0 iops @ 0 KB Peak: 0 iops @ 0 KB		10/07/2018 03:22 PM	10/07/2018 03:32 PM		
	Average 0 iops @ 0 KB Peak: 0 iops @ 0 KB		10/07/2018 03:22 PM	10/07/2018 03:32 PM		
NPIV Client Utilization - fcs0	High: 0.00 % Average: 0.00 %		10/07/2018 03:22 PM	10/07/2018 03:32 PM		
aix1nim	Average 0 iops @ 0 KB Peak: 0 iops @ 0 KB		10/07/2018 03:22 PM	10/07/2018 03:32 PM		
	Average 0 iops @ 0 KB Peak: 0 iops @ 0 KB		10/07/2018 03:22 PM	10/07/2018 03:32 PM		
FC I/O Operations Blocked	optimal		10/07/2018 03:22 PM	10/07/2018 03:32 PM		
FC Port Speeds	running at full speed		10/07/2018 03:22 PM	10/07/2018 03:32 PM		

147

147

part command

VIOS - Processor Risk/Impact 1=lowest 5=highest

Name	Measured Value	Suggested Value	First Observed	Last Observed	Risk	Impact
CPU Capacity	1.5 ent		10/07/2018 03:22 PM			
CPU consumption	Average:1.4% (cores:0.1) High:2.7% (cores:0.1)		10/07/2018 03:22 PM	10/07/2018 03:32 PM		
Processing Mode	Shared CPU, (UnCapped)		10/07/2018 03:22 PM			
Variable Capacity Weight	255		10/07/2018 03:22 PM			
Virtual Processors	3		10/07/2018 03:22 PM			
SMT Mode	SMT4		10/07/2018 03:22 PM			

System - Shared Processing Pool Risk/Impact 1=lowest 5=highest

Name	Measured Value	Suggested Value	First Observed	Last Observed	Risk	Impact
Shared Pool Monitoring	enabled		10/07/2018 03:22 PM			
Shared Processor Pool Capacity	8.0 ent.		10/07/2018 03:22 PM			
Free CPU Capacity	average_free:7.7 ent. lowest_free:7.3 ent.					

148

148

part command

VIOS - Memory

Risk/Impact 1=lowest 5=highest

	Name	Measured Value	Suggested Value	First Observed	Last Observed	Risk	Impact
✓	Real Memory ?	8.000 GB		10/07/2018 03:22 PM			
i	Available Memory ?	1.867 GB		10/07/2018 03:22 PM			
✓	Paging Rate ?	0.0 MBps Paging Rate		10/07/2018 03:22 PM	10/07/2018 03:32 PM		
✓	Paging Space Size ?	4.000 GB		10/07/2018 03:22 PM	10/07/2018 03:32 PM		
i	Free Paging Space ?	3.880 GB free		10/07/2018 03:22 PM			
⚠	Pinned Memory ?	5.488 GB pinned	less than 4.000 GB pinned	10/07/2018 03:22 PM	10/07/2018 03:32 PM	1	4

VIOS - Shared Ethernet Adapters

Risk/Impact 1=lowest 5=highest

	Name	Measured Value	Suggested Value	First Observed	Last Observed	Risk	Impact
i	SEA Adapter Count ?	1		10/07/2018 03:22 PM			
⚠	SEA (ent7)	Mapping: Physical :(ent2,ent3),Virtual :(ent4,ent4)		10/07/2018 03:22 PM			

149

149