

[close window](#)[Web Exclusive](#)[Print](#) 

# Tips I Picked up at the Power Systems Technical University

November 2017 | by [Jaqui Lynch](#)

In October, I attended and spoke at the Power Systems Technical University in New Orleans. While I was there I was able to attend some very valuable sessions and picked up some great tips. I thought this might be an opportunity to share some of those with you. Even though I have been working in the area for a very long time, I find these conferences extremely useful—both for networking and the technical content. At this conference I was focusing on several areas that I wanted to increase my knowledge on: enhanced mode HMC, networking and Spectrum Scale (formerly GPFS). Below are a series of tidbits I picked up. Hopefully you'll find them useful as well.

## Networks

One of the more interesting things I learned (from Alexander Paul) was very basic : how to quickly transfer an IP address without causing an outage. I've been moving my NIM server from fully virtualized to standalone and found this to be very useful.

In my case I was moving from en3 to en1

```
ifconfig en3 192.168.65.150 transfer en1
I then test with ping, etc and if it is working I make it permanent as follows:
chdev -l en1 -a netaddr=192.168.65.150 -a netmask=255.255.255.128 -a state=up
Be careful that you disable en3 before the next reboot or it
will come back up with the same IP address and you will have a duplicate
IP error. You can use chdev to set en3 to state=down or my personal favorite is to
rmdev the adapter and run cfgmgr so it comes back clean:
rmdev -dl et3
rmdev -dl en3
rmdev -dl ent3
cfgmgr
```

I also went to a great session on Wireshark by Roy Spencer. I've been looking at a lot of trace output for VIO servers and the SEA has duplicate packets because the SEA sees the packet come in from the LPAR and the transmit out. It turns out that Wireshark includes a utility to remove these duplicate packets:

```
editcap -F nettle -d 'source' 'dest'
```

This will help me a lot in my analysis especially when trying to produce packet counts.

## Power Masters

The tricks of the Power Masters (Gareth Coates) session is always a great place to find nuggets. In particular, he provided details on how to turn on remote execution, remote operation and remote virtual terminals for the HMC in Enhanced mode. This is not under HMC management—instead you'll find it under users and security, systems and console security. There's a version of this session available at the Power Virtual User Group and a demo of the Enhanced HMC at the UK PowerVM Virtual User Group. I picked up a lot of tips on the HMC in this session including tips around RMC.

## ***VIOS and Other Tips I Picked Up***

export CLI\_DEBUG=33 Then run a VIO command and it will show you the actual AIX commands that are executed behind the scenes.

I decided to test this on my system and here is what I saw:

```
$ export CLI_DEBUG=33
```

```
$ lsnports
```

```
AIX: "/usr/lib/methods/viosmgr -t npiv -f query_fc_ports >/dev/null"
```

name	physloc			fabric	tports	aports	swwpns	awwpns
fcs4	U78C9.001.WZS0234-P1-C6-T1	1	64	57	3088	3065		
fcs5	U78C9.001.WZS0234-P1-C6-T2	1	64	57	3088	3065		

lsnports is a useful way to check if your fibre ports are connected to an NPIC capable switch.

If this is not the case you will see the following message:

```
"System doesn't have NPIV capable ports or no NPIV ports with the link up."
```

```
$ lsmmap -all -vnic
```

```
AIX: "lsdev -c adapter -t IBM,vnic-server -s vdevice -F "name" | wc -l -c"
```

```
AIX: "lsdev -c adapter -t IBM,vnic-server -s vdevice -F "name" "
```

## ***Control channel***

Another useful tip is how to check for the control channel (the new default is 4095 as of 2.2.3)

```
# lsdev -C | grep ent | grep Shared
```

```
ent7          Available          Shared Ethernet Adapter
```

```
# entstat -d ent7 | grep "Control Channel"
```

Control Channel PVID: 4095 Control Channel Adapter: ent4

As you can see, the SEA is on ent7 and the virtual ethernet carrying the traffic is ent4

## ***Path Issues***

Recently I have worked on multiple systems where the paths have been messed up for various reasons. It became necessary to remove the missing or failed paths without destroying the enabled paths. The way to do this is using the lspath command and get it to list all of the paths for that disk.

```
lspath -l hdisk15 -H -F "name:parent:connection:status" | grep Missing (or grep Failed)
```

You can put this into a loop where it can produce the correct output for you.

As an example

```
lspath -l hdisk15 -H -F "name:parent:connection:status"
```

```
name:parent:connection:status
```

```
hdisk15:fscsi4:500507680d088ef6,10000000000000:Enabled
```

```
hdisk15:fscsi4:500507680d088ef7,10000000000000:Enabled
```

```
hdisk15:fscsi5:500507680d048ef6,10000000000000:Enabled
```

```
hdisk15:fscsi5:500507680d048ef7,10000000000000:Enabled
```

If the first path on fcscli4 was missing you could then remove it using:

```
rmpath -dl hdisk15 -p fcscli4 -w 500507680d088ef6,10000000000000
```

There are several examples of how to script production of the output in an article by David Tansley.

## WWPNs

Gareth also included a script to list the WWNs for the LPAR—I modified it slightly to add the hostname and adapter slot information:

```
# cat seewwpns.sh
hostj=`hostname`
lsdev -Ccoadapter | grep fcs | awk '{print $1,$3}' | while read fcs slot
do
wwpn=`lscfg -vl $fcs | grep -i network | cut -c 37-60`
wwpn1=`lscfg -vl $fcs | grep -i Hardware | cut -c 47-64`
echo $hostj $fcs $slot $wwpn1 $wwpn
done

# ./seewwpns.sh
vio2 fcs4 05-00 WZS0234-P1-C6-T1 10000090FA530975
vio2 fcs5 05-01 WZS0234-P1-C6-T2 10000090FA530976
vio2 fcs6 06-04 WZS0234-P1-C8-T1 10000090FA740155
vio2 fcs7 06-05 WZS0234-P1-C8-T2 10000090FA740156
```

This is a great way to get the WWPNs in order to give them to the storage admins for zoning.

There's also a -client flag on fcstat (as of VIO 2.2.2.2) that will get you information on statistics for all the WWNs and WWPNs that the VIO sees. You can run just fcstat -client to get full statistics or you can use the command to get the WWPNs and WWNs. You can also use the viostat -adapter command to get statistics. To get the WWNs and WWPNs, as padmin on vio2 I ran: fcstat -client | cut -c 1-50 Results were:

hostname	dev	wwpn
vio2	fcs4	0x10000090FA530975
aixlnim	fcs2	0xC0507607DBD8002C
gpfs1	fcs2	0xC0507607DBD80034
vio2	fcs5	0x10000090FA530976
aixlnim	fcs3	0xC0507607DBD8002E
gpfs1	fcs3	0xC0507607DBD80036

From the HMC you can also try:

```
lshwres -r io --rsubtype slotchildren -m Server-8286-41A-SN215D3AV -F phys_loc,description,mac_address,wwpn (you can add a | grep Fibre)
```

I don't use the grep at the end on Fibre as I also want to see what FCOe adapters I have

```
hscroot@HMC:~> lshwres -r io --rsubtype slotchildren -m Server-8286-41A-SN215D3AV -F phys_loc,description,mac_address,wwpn | grep PCI
```

```
U78C9.001.WZS0234-P1-C6-T1,PCIe2 16Gb 2-Port Fibre Channel Adapter,null,10000090fa530975
U78C9.001.WZS0234-P1-C6-T2,PCIe2 16Gb 2-Port Fibre Channel Adapter,null,10000090fa530976
U78C9.001.WZS0234-P1-C8-T2,"PCIe2 4-port (2 10GbE+2 1GbE) FCoE+NIC SRIOV Adapter, FCoE PF",null,10000090fa740156
```

```
U78C9.001.WZS0234-P1-C8-T1,"PCIe2 4-port (2 10GbE+2 1GbE) FCoE+NIC SRIOV Adapter, NIC PF",0090fa740151,null
U78C9.001.WZS0234-P1-C8-T1,"PCIe2 4-port (2 10GbE+2 1GbE) FCoE+NIC SRIOV Adapter, FCoE PF",null,10000090fa740155
U78C9.001.WZS0234-P1-C8-T4,"PCIe2 4-port (2 10GbE+2 1GbE) FCoE+NIC SRIOV Adapter, NIC PF",0090fa740154,null
```

As you can see there are two excellent ways to get the same information, depending on where you want to request it from and what format you need it to be in.

## ***Useful IBM Tools***

I have mentioned these before, but there are a group of tools that I use all the time. They consist of nmon, nmon analyser, HMC Scanner, FLRT (fix level recommendation tool, FLRTVC (FLRT vulnerability checker) and fcstat -D.

### ***nmon***

nmon has been part of the operating system since AIX 6.1 and is used by admins all over the world to gather performance data. I like to run it using the following flags so that I don't miss anything:

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

The above takes a 30 minute snapshot (120 x 15 second snaps) and includes asynchronous IO (A), the SEA (O) paging (P), volume groups (V), fibre adapter statistics (^), disk service times (d), memory pages (M), large pages (L) and workload manager statistics (W) if you are running workload manager. I run nmon all the time so my normal cron job runs for 24 hours using "-s 150 -c 576". Depending on how granular you need to get you can tweak these values.

### ***nmon analyser***

nmon analyser goes hand in hand with nmon. It's an Excel spreadsheet that processes nmon files and produces graphs of what is happening. It's not a total performance monitoring solution, but it provides some valuable information for day to day performance work. I supplement this with my own data gathering scripts that gather more in depth data.

### ***HMC Scanner***

The HMC Scanner is a great tool for documenting your system in one place. The current version is 0.11.35—you actually download 0.11.24 and then replace the .jar file to get to 0.11.35. You can install it on AIX and run it from there against the HMC or you can run it from your desktop, assuming you have the correct Java installed. I normally run it from AIX using:

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

This produces a bunch of files, but the one you want is the one that ends .xls. This is the Excel spreadsheet with all the systems connected to the HMC documented.

### ***FLRT and FLRTVC***

FLRT (Fix level recommendation tool) has been around for several years now and it's something I use whenever I am planning an upgrade or trying to determine how long I can stave off an upgrade for. The most recent version provides you with the release date for the versions you are running and their end of service date (if announced). It also provides those for the recommended upgrade levels. It will also provide links to the readme files for the recommended updates.

FLRTVC is a vulnerability checker. You download the script and run it on the system to be reviewed. It uses wget or curl to try to download a file called apar.csv from IBM and it then checks known issues against your software levels. The most common things it finds are back levels of SSH, SSL and Java - by default these are at an insecure level and new levels need to be installed. At a minimum, Openssh should be at 7.1.102.1100 and openssl should be at 1.0.2.1100 to avoid security problems. Java should be at the August 2017 level which

for Java7 is 7.0.0.610. I replace these on all my systems including my VIO servers. FLRTVC identifies the efixes and ifixes that need to go on, provides links to the readmes and also to the actual download. You can send the output to a .txt file, download it and open it in Excel (it's in csv format) to make it easier to view. I typically run FLRTVC on my critical systems once a month.

### **fcstat**

If you are looking at fibre adapter statistics then you will want to be familiar with the fcstat command, especially the fcstat -D. If you run fcstat -D against a fibre port it will provide the WWN/WWPN as well as the firmware version, the adapter speed that it is running:

```
FIBRE CHANNEL STATISTICS REPORT: fcs4
Device Type: FC Adapter (adapter/pciex/df1000e21410f10)
Serial Number: 1A34200935

ZA: 10.2.252.1919                               This is the microcode - lsmcode -A shows
it as: fcs4!df1000e21410f103.00010000020025201919
World Wide Node Name: 0x20000090FA530975
World Wide Port Name: 0x10000090FA530975       WWPN for zoning
.....
Port Speed (supported): 16 GBIT                 Can run up to 16Gbit
Port Speed (running): 8 GBIT                   But is connected at 8Gbit
Port FC ID: 0x010300                            Adapter FCID - useful to figure out where you are connected
Port Type: Fabric
Attention Type: Link Up
Topology: Point to Point or Fabric
```

Further down you can review statistics that will provide indications of whether there is queuing at the adapter or protocol drivers.

### **Create a clone of a boot disk**

```
alt_disk_copy -d hdisk2
bosboot -a -d hdisk2
bootlist -m normal hdisk2
```

The above copies the current rootvg disk to hdisk2 which will become the altinst\_rootvg volume group. You can use this as a backup before making changes to the current system (simple reboot of the new volume group to recover), or you can then work on the new volume group using alt\_disk\_install and alt\_rootvg\_op commands. You can also make minor changes (such as adding a host to the /etc/hosts on the alternat disk) by waking up the volume group:

```
alt_disk_install -W hdisk2
You can now go and vi /alt_inst/etc/hosts
```

When you are done you put the volume group back to sleep:

```
alt_disk_install -S hdisk2
```

Whenever you use alternate disk install or any of its derivatives be sure to point the bootlist to whichever disk you want to boot from next. By default, many of the alternate disk install commands change the bootlist to point to the new volume group and this may not be what you want to do.

### **Summary**

I found the Power Systems Technical University conference to be very valuable in terms of picking up tips to make my life easier as a sysadmin. Hopefully you will find these tips helpful as well. I highly recommend attending these conferences—there are some great sessions, but many tips get picked up during the networking discussions where you can compare notes with others. Many thanks to all those who presented and shared their ideas with myself and others.

IBM Systems Magazine is a trademark of International Business Machines Corporation. The editorial content of IBM Systems Magazine is placed on this website by MSP TechMedia under license from International Business Machines Corporation.

©2019 MSP Communications, Inc. All rights reserved.

# Systems

Connect With Us:



Magazine Archives

Search

IBM i

LINUX ON POWER

MAINFRAME

POWER

**AIX**

ADMINISTRATOR

TRENDS

CASE STUDIES

TIPS & TECHNIQUES

STORAGE

PRODUCT NEWS

## References

< Return to main article

Print Email

**IBM Power Virtual User Group:** <http://www.tinyurl.com/ibmaixvug>

**UK PowerVM Virtual User Group:** <http://tinyurl.com/PowerVUG>

**FLRT:** <http://www14.software.ibm.com/support/customer-care/flrt/>

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

**Using fcstat:** <http://www-01.ibm.com/support/docview.wss?uid=tss1td106122>

**HMC Scanner:** <http://tinyurl.com/HMCscanner>

**David Tansley – Removing failed or missing paths:** <http://ibmsystemsmag.com/aix/tipstechniques/systemsmanagement/removing-missing-paths/>

**Nmon analyser:** [https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Power%20Systems/page/nmon\\_analyser](https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Power%20Systems/page/nmon_analyser)

**IBM Technical Events:** <https://www-03.ibm.com/services/learning/ites.wss/zz-en?pageType=page&c=a0002173>

< Return to main article

ADVERTISEMENT

### POWER SYSTEMS EXTRA

Maximize your IT investment with weekly information from THE source... Power Systems EXTRA eNewsletter.

**SIGN UP TODAY**

**Read Previous Issues**

READ THE CURRENT ISSUE: **DIGITAL** | **ONLINE** | **eNEWSLETTER**

**AIX** | **IBM i** | **LINUX ON POWER** | **MAINFRAME** | **POWER**

Connect With Us:

[Homepage](#) [About Us](#) [Contact Us](#) [Subscriptions](#) [Editorial Calendar](#)

[Advertise With Us](#) [Reprints](#) [Privacy Policy](#) [Terms of Service](#) [Sitemap](#)

IBM Systems Magazine is a trademark of International Business Machines Corporation. The editorial content of IBM Systems Magazine is placed on this website by MSP TechMedia under license from International Business Machines Corporation.

©2019 MSP Communications, Inc. All rights reserved