TechU

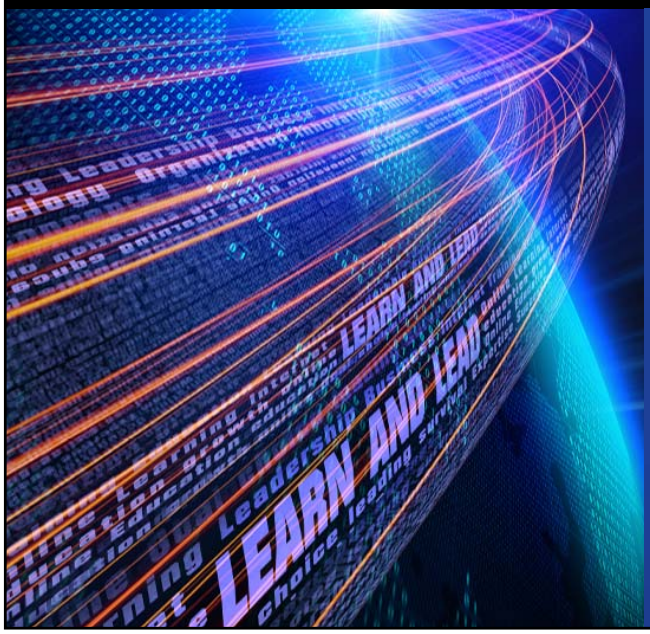

Care and Feeding of VIO Servers
Part 3 - LPM

Jaqui Lynch
Flagship Solutions Group
jlynch@flagshipsg.net

2020 IBM Systems Technical University
October 26-29, 2020| Virtual Conference

IBM

1

Agenda

- Part 3 – LPM
  - Introduction
  - Prerequisites and Planning
  - Storage Needs
  - Remote Migration
  - SRR
- Documentation
  - Useful Commands
  - Useful Links
  - Backup Material

2

2

LPM

3

3

PowerVM Live Partition Mobility (LPM)

- LPM provides the ability to move an AIX, IBM i or Linux partition from one LPM-capable physical server to another *compatible* server
- What is meant by *compatible*?
  - Power Systems server requirements
  - Management console requirements
  - Virtual I/O Server (VIOS) requirements
  - Mobile partition requirements
- How are LPM servers managed?
  - HMC

IBM refers to LPARs as either LPARs or Virtual Servers

Firmware support Matrix for LPM
https://www.ibm.com/support/knowledgecenter/en/POWER9/p9hc3/p9hc3_firmwaresupportmatrix.htm
Migration Modes
https://www.ibm.com/support/knowledgecenter/POWER9/p9hc3/p9hc3_pcmcombosinact.htm

IBM Document that is very useful:
ftp://ftp.software.ibm.com/systems/power/docs/hw/p9/p9hc3.pdf

4

4

Definitions

Active Partition Mobility

Active Partition Migration is the actual movement of a running LPAR from one physical machine to another without disrupting the operation of the OS and applications running in that LPAR.

Inactive Partition Mobility

Inactive Partition Migration transfers a partition that is logically 'powered off' (not running) from one system to another.

Suspended Partition Mobility

Suspended Partition Migration transfers a partition that is suspended from one system to another.

Partition Mobility (Live or Inactive) and Partition Migration (Active or Inactive)

Refer to the same feature.

5

5

Summary of migration phases

Active

- Validate configuration
- Create new LPAR
- Create new virtual resources
- *Migrate state of LPAR in memory*
  - Includes memory, hardware page table, processor state, NVRAM, Time of Day and partition configuration
- Remove old LPAR configuration
- Free up old resources

6

6

Summary of migration phases

Inactive

- Validate configuration
- Create new LPAR
- Create new virtual resources
- Remove old LPAR configuration
- Free up old resources

Migration uses last used partition profile so LPAR must have been activated at least once, even if just in SMS mode

7

7

General LPM Depiction

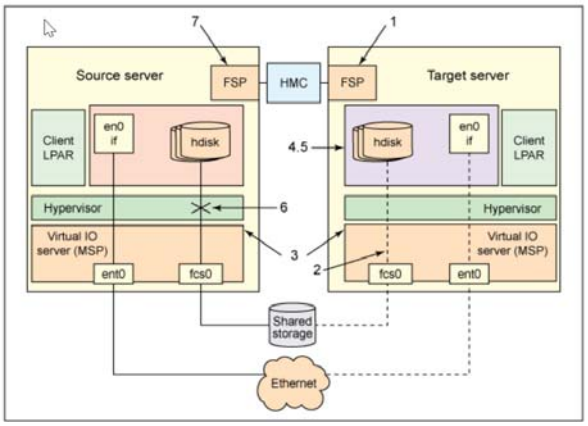

The following describes the general LPM depiction in the diagram

1. Partition profile (presently active) copied from source to target FSP.
2. Storage is configured on the Target.
3. Mover service partitions (MSP) is activated.
4. Partition migration started.
   a. Majority of memory pages moved.
   b. All threads piped down.
5. Activation resumed on target.
   a. Final memory pages moved.
   b. Cleanup storage and network traffic.
6. Storage resources are deconfigured from the source.
7. Partition profile removed from source FSP (Flexible Service Processor).

From: https://developer.ibm.com/technologies/systems/articles/au-lpm_troubleshooting/

8

8

Ian Robinson – P102956 – Slide 13

PowerVM Live Partition Mobility (LPM)

PowerVM

LPM virtualizes all storage and network resources

Process transfers the processor state, memory, virtual devices and connected users

An identical VM is created on the target server and memory pages are copied

Any memory changes on the source server are tracked to be re-copied

After memory state is copied, processing is stopped on source and started on target

Remaining 'dirty' memory pages are copied and cleanup completed

9

9

Requirements for LPM

- **PLANNING IS CRITICAL**
- http://www14.software.ibm.com/webapp/set2/sas/f/pm/component.html
- Hardware POWER6 Only and above
- HMC v7.3.2 with MH01062
- Firmware E*340_039 min
  - http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/p7hc3/p7hc3firmwaresupportmatrix.htm
- AIX v5.3 5300-07-01
- AIX v6.1 6100-00-01
- VIO Server 1.5.2.1-FP-11.1 or v2.1
- RHEL5 Update 1 and SLES10 Update 1 supported (or later)
- HMC v7.3.4 introduces remote migration
  - Partitions can migrate between systems managed by different HMCs

Mobility between HMC and FSM requires HMC v7r7.1.0 or later

- Virtualized SAN storage (rootvg and all other vgs)
- Virtualized Ethernet (SEA)
- LPAR being moved cannot be using the HEA/IVE (VIO can though)
- Check the prereq site:
  - https://www-912.ibm.com/e_dir/eserverprereq.nsf
- **No dedicated anything at the time of the move**
- **No virtual opticals**
- **IBM i LPAR must be in restricted I/O mode which means you cannot dlpar in or add any real I/O devices**

10

10

Other LPM Prereqs

- Two servers – POWER6, 6+, 7, 7+ (or mix thereof), 8 or 9
- POWER9 does not support POWER6 modes
- PowerVM Enterprise Edition on all servers
- All LPARs on the same Open network with RMC established to HMC
  - **RMC must be working**
- LPARs must be under control of VIOS
- Both systems HMC connected on the RMC network
- Check VIOS levels
  - Many new features require v2.1 or higher
- Storage must be virtualized
  - Storage must be zoned to both source and target
  - No LVM based disks
  - No local storage – must be on the SAN
  - hdisks must be external and have reserve_policy=no_reserve
  - After 2.2.4 of VIOS new vioslpm0 pseudo device for checking for hdisks
  - See section 3.7 of the LPM red book SG24-7460
- Must use Shared Ethernet Adapter
  - See section 3.8 of the LPM red book SG24-7460
- All resources must be shared or virtualized prior to migration
- Check operating system level is supported on the new server
- No partition at receiving server can have the same name as the LPAR being migrated from
- **Must have resources available at the target**
  - **Cores, memory, adapters, etc**

11

11

LPM

- Check LPAR on HMC under Capabilities
  - Look for Active and Inactive Partition Mobility Capable=True
- Ensure VIO server is set up as a Mover Service Partition (MSP) under the general tab on the VIO server at each end
  - By default MSP is set to no on a VIO server
- Mover partition must have a VASI (Virtual Asynchronous Services Interface) device defined and configured (done automatically by HMC)
- The pHypervisor will automatically manage migration of CPU and memory
- Dedicated IO adapters must be de-allocated before migration
- cd0 in VIO may not be attached to mobile LPAR as virtual optical device
- Time of Day clocks for VIO servers should be synchronized
- The operating system and applications must be migration-aware or migration-enabled
- LMB (memory region) size must be the same on both servers – check on HMC
  - Requires a whole server reboot to change

Check capability for Active Memory Sharing, suspend/resume and Trusted Boot if mobile partition is configured for them

Note: As of HMC v8.8.4 you can disable LPM migration for an LPAR:
https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Power%20Systems/page/HMC%20per-partition%20disablement%20of%20Live%20Partition%20Mobility%20(LPM)

12

12

Processor compatibility

- I always set compatibility mode to default
- After an active LPM the LPAR will be in the same compatibility mode it was in on the old system
  - This means that if you were in power7 mode you are still in power7 mode
- After an inactive LPM the LPAR will boot in the preferred compatibility mode
  - This means on a power8 it will boot in power8 mode enabling you to set smt8 later
- NOTE if you want to use smt8 and you were smt4 you will need to make the change in the LPAR
- **VIO (pre v3.1) will show effective mode POWER7 and runs in SMT4 as (pre v3) it is AIX v6 under the covers**

The default processor compatibility mode is a preferred processor compatibility mode that enables the hypervisor to determine the current mode for the logical partition. When the preferred mode is set to default, the hypervisor sets the current mode to the most fully featured mode supported by the operating environment. In most cases, this is the processor type of the server on which the logical partition is activated. For example, assume that the preferred mode is set to default and the logical partition is running on a POWER8 processor-based server. Because the operating environment supports the POWER8 processor capabilities, the hypervisor sets the current processor compatibility mode to POWER8.

13

13

Processor compatibility

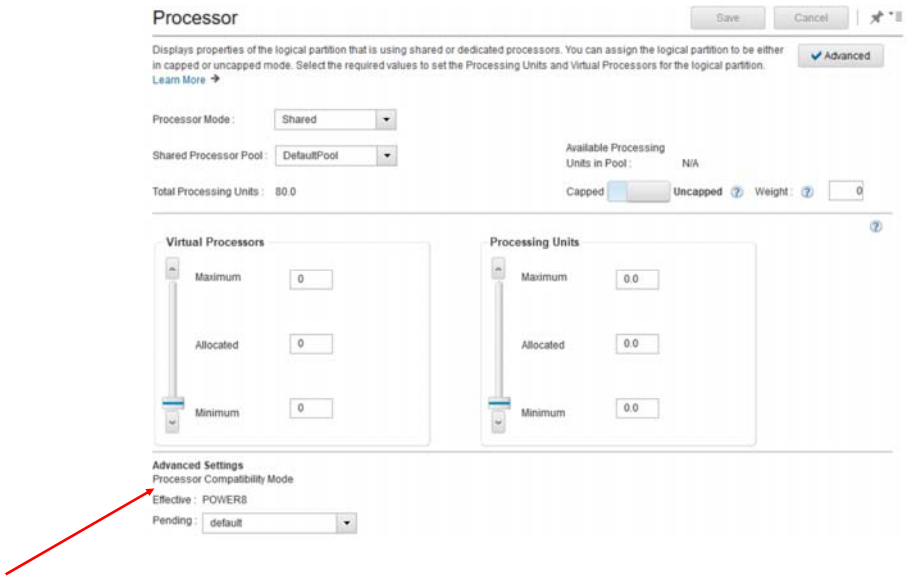

14

14

Processor compatibility

POWER9 compatibility modes

The POWER9 processor has the ability to run in compatibility modes for previous POWER processor generations. This enables older operating systems to run on POWER9 systems. Compatibility modes also allows for live partition migration from systems based on previous generations of POWER processors. The POWER9 processor can run in the following compatibility modes:

- POWER7
- POWER8
- POWER9 Base

https://www.ibm.com/support/knowledgecenter/POWER9/p9hc3/p9hc3_pcmcombosinact.htm

15

15

Live Partition Mobility and Servers

Initially released for Power 6 servers

Now also supports Power7 and Power8 and Power9

Included with PowerVM Enterprise Edition

Can migrate between generations as follows:

- P6 <-> P7, P8 (no P9)
- P7 <-> P6, P8, P9
- P8 <-> P6, P7, P9
- P9 <-> P7, P8 (no P6)

16

16

VIOS requirements checklist

VIOS on both source and destination server must provide:

- Access to the same network (VLAN and subnet) with a Shared Ethernet adapter
- Access to the same physical storage either from:
  - External Fibre Channel storage system using virtual SCSI or NPIV
  - External iSCSI storage using virtual SCSI

Virtual storage:

- VIOS virtual adapters cannot be marked as required and should not be marked for “Any client”
- Destination VIOS must have enough "available" virtual adapters

Other:

- At least one VIOS per server must be configured as an MSP
- For shared memory LPARs, destination must have a paging device available
- Consider configuring a VIOS as a time reference partition

17

17

MSP – Mover Service Partition

Check the box on the VIO servers you want to use to oversee transfers

- Only applies to Active Partition Mobility
- Used to transfer the memory contents of the LPAR from the source to the target server
- The memory contents are moved via the VASI (virtual asynchronous services interfaces) which is a virtual device that is created when the MSP is enabled on a VIOS.
- You can display the VASI on the VIO
- Any VIOS with this checked can be used
- The transfer uses an IP address in the VIO server
- You can use two MSPs at each end (Dual MSP) to increase the bandwidth
- Click on the MSP pairing button during validation to choose the MSP pair to use to get the best performance
- Network for MSPs impacts time to transfer
  - i.e. single MSP on10Gb network transfers 200Gb in about 3-4 minutes depending how active the LPAR is

18

18

Checking the VASI

 # lsdev -C | grep vasi
vasi0 Available Virtual Asynchronous Services Interface (VASI)

 # lsattr -El vasi0
medium_buf_max 256 Maximum Medium Buffers True
medium_buf_size 8192 Medium Buffer Size (in bytes) False
rx_max_pkts 50 Maximum Received Packets Per Interrupt True
small_buf_max 2048 Maximum Small Buffers True
small_buf_size 2048 Small Buffer Size (in bytes) False
tx_buf_min 512 Maximum Transmit Buffers True
tx_buf_size 16384 Transmit Buffer Size (in bytes) False

 # lscfg -vpl vasi0
vasi0 U8286.41A.215D3AV-V1-C32769 Virtual Asynchronous Services Interface (VASI)
Hardware Location Code.....U8286.41A.215D3AV-V1-C32769
PLATFORM SPECIFIC
Name: IBM,VASI
Node: IBM,VASI@30010000
Device Type: IBM,VASI-1
Physical Location: U8286.41A.215D3AV-V1-C32769

19

19

vasistat –all vasi0

As padmin use vasistat command
$vasistat –all vasi0
Provides statistics

```
$ vasistat -all vasi0
-------------------------------------------------------------
VASI STATISTICS (vasi0) :
Elapsed Time: 25 days 13 hours 57 minutes 56 seconds
Transmit to PHYP Statistics:        Receive from PHYP Statistics:
----------------------------        ----------------------------
Packets: 0                          Packets: 0
Bytes: 0                            Bytes: 0
No Buffers: 0                       No Buffers: 1818501198
Transmit Errors: 0                  Receive Errors: 0
Bad Packets: 0                      Bad Packets: 0
Output Calls: 0                     Interupts: 12
                                    Maximum Buffers Per Interrupt: 0
                                    Average Buffers Per Interrupt: 0

Average Time Spent in CRQ Send: 5 microseconds
Maximum Time Spent in CRQ Send: 11 microseconds
Minimum Time Spent in CRQ Send: 1 microseconds

Driver Flags: Up Running 64BitSupport

Maximum Operations: 8
Maximum Receive Pools: 0
Active Operations: 0

DMA Channel: 10010000444D4120
Bus ID: 90000340

Local DMA Window:
        Logical I/O Bus Number: 10010000
        TCE Base: 0000000000000000
        TCE Size: 0000000010000000

Remote DMA Window:
        Logical I/O Bus Number: 00000000
        TCE Base: 0000000000000000
        TCE Size: 0000000000000000

Supported Operation Types: Migration

num of operations: 0
$
```

20

20

Setting VIOS to MSP

Click on VIOS properties and check the Mover service partition box in advanced

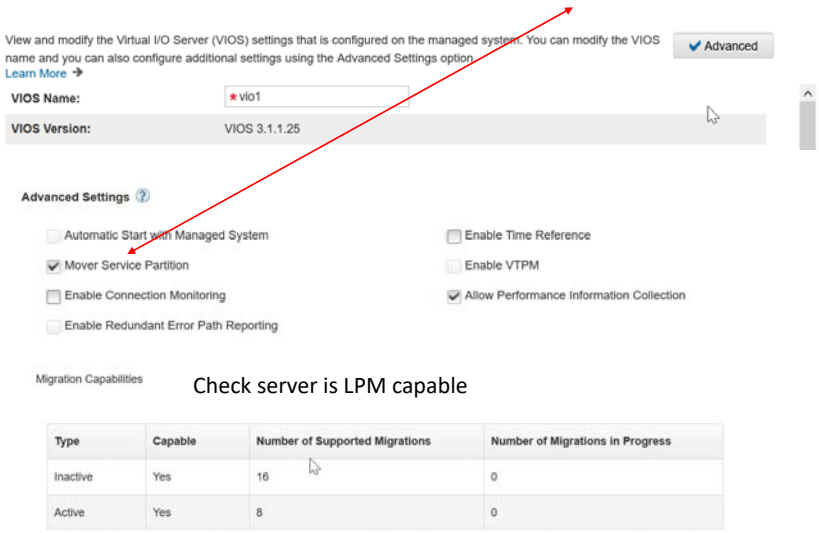

Migration Capabilities

Check server is LPM capable

Type	Capable	Number of Supported Migrations	Number of Migrations in Progress
Inactive	Yes	16	0
Active	Yes	8	0

21

21

MSP

https://www.ibm.com/support/pages/lpm-enhancements-high-resiliency-performance-powervm-version-225

Specific firmware requirements (on P8 FW860 or higher) need to be met

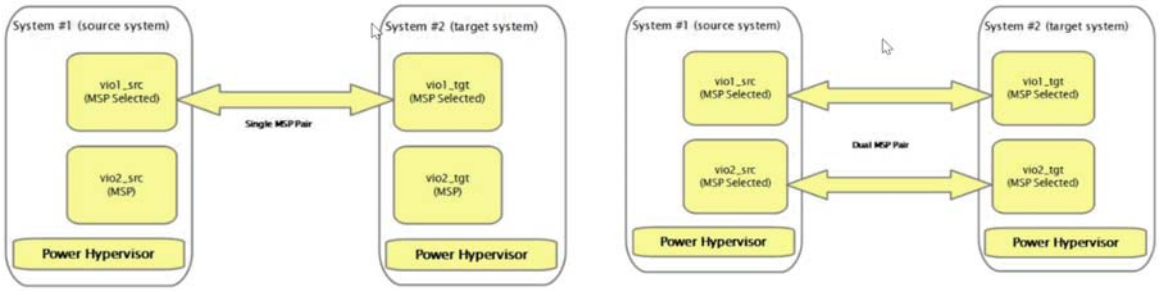

22

22

Select MSP Pairing

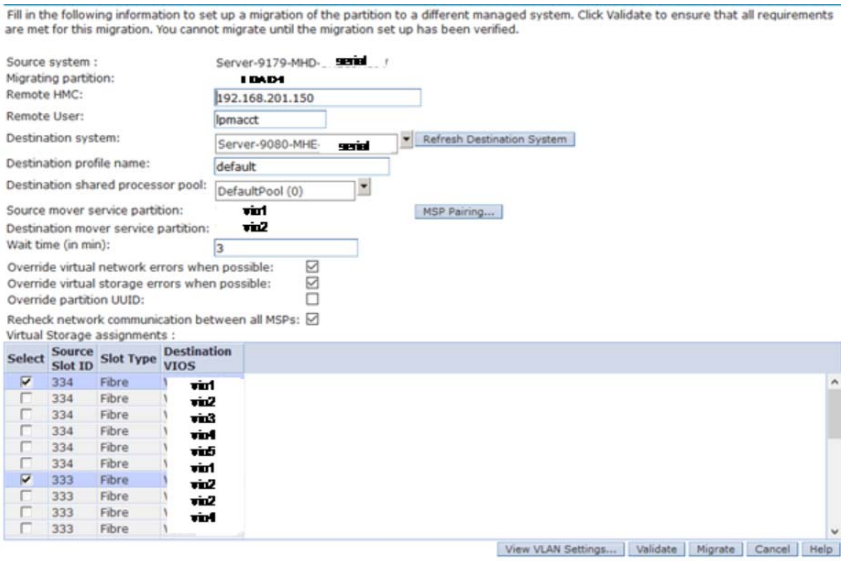

23

23

Mobile partition requirements checklist 1/3

Check mobile partition configuration for:
Operating system support
Functioning RMC (resource management controller) daemons
No workload group
No open consoles (warning only)
Storage on external storage unit and accessible by both VIOS LPARs (Zoned and mapped)
Storage set to no_reserve in LPAR
No adapters set to required in profile
Check that LPAR name will be unique on destination server
Check MAC address will be unique on destination server
Not a service partition
No logical ports configuration on Integrated Virtual Ethernet (IVE)/Host Ethernet Adapter (HEA)
    AIX 6100-05 (or higher) clients can convert LHEA to virtual Ethernet during migration
Verify valid processor compatibility mode
Virtual network configuration compatible with destination VIOS

24

24

Mobile partition requirements checklist 2/3

*** Okay for inactive migrations but not active migrations**
No huge memory pages*****
No BSR arrays*****
No redundant error reporting*
No physical I/O*

If mobile LPAR using AME (Active Memory Expansion)
Make sure it's supported on the destination server
If the mobile partition is suspend-resume capable, make sure the target has a reserved storage pool greater than or equal to 110 percent of the lpar size

NOTE target server must have enough CPU and memory resources for the LPAR to be moved

25

25

Mobile partition requirements checklist 3/3

If shared processor pool with an entitlement below 0.1 processing units and greater than or equal to 0.05 (POWER7+), make sure the destination server also shares that capability
LUNs using NPIV need to be zoned and mapped to both world wide port names on each client Virtual Fibre Channel adapter
All I/O resources must be shared or virtualized prior to migration – dedicated devices need to be removed including vtopt
No vSCSI or NPIV adapters can be set as required in the profile

https://ibmsystemsmag.com/Power-Systems/10/2018/guide-live-partition-mobility

Migrating IBM i LPARs

1. Verify the destination server supports the migration of IBM i mobile partitions and the restricted I/O mode
2. Verify the IBM i mobile partition is in the restricted I/O mode
3. Restricted I/O mode requires an LPAR reboot and disable DLPAR of physical I/O into the IBM i LPAR
4. No tape drives can be active

26

26

Zoning and Mapping (masking)

- ***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 or masking***
  - This is when the storage is updated to allow the host (LPAR or server) WWPNs access to the specific LUNs provisioned
- LUNs must be provisioned at the storage, then mapped and zoned before they can be used in an LPAR
- For direct attach we zone and map the WWNs for the real adapters, for NPIV we use the WWPNs on the virtual adapters
- WWNs tend to start with 10 or 20
- WWPNs (NPIV) start with C0
- These can be found in an HMCScanner report or by logging onto the LPAR or VIO or from the HMC
- Check the VIO connection to the switch is NPIV enabled:

```
$ lsnports
name             physloc                        fabric tports aports swwpns  awwpns
fcs0             U78C9.001.WZS0234-P1-C7-T1          1     64     56   3088    3062
fcs1             U78C9.001.WZS0234-P1-C7-T2          1     64     56   3088    3062
```

27

27

LPM Zoning and mapping - NPIV

- **Do not confuse zoning with mapping (aka masking)**
- 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.

28

28

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 "lsmap –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 "lsmap –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

29

29

Finding WWPNs for NPIV

On HMC, go to virtual adapters in LPAR profile, select a VFC then actions, advanced, login logout fibre

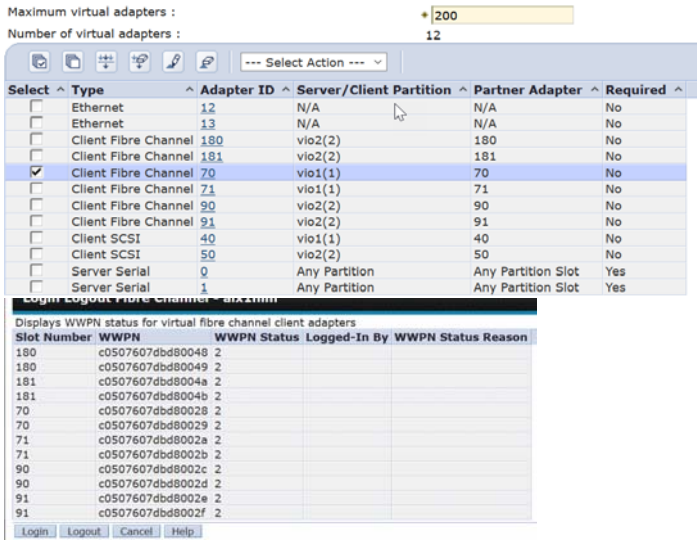

Or go into the LPAR profile
Select a VFC and actions, properties

30

30

Simple Zoning for LPM

SIMPLE ZONE

SWITCH1
zone: NPIV_AIX1
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_AIX1
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

31

31

NPIV Mapping at VIO as padmin

NPIV

$ vfcmap -vadapter vfchost20 -fcp fcs0

$ lsmap -vadapter vfchost20 -npiv
Name Physloc ClntID ClntName ClntOS
------------- ---------------------------------- ------ -------------- -------
vfchost20 U8286.41A.215D3AV-V1-C108 17 aixtest1 AIX

Status:LOGGED_IN
FC name:fcs0 FC loc code:U78C9.001.WZS0234-P1-C7-T1
Ports logged in:3
Flags:a<LOGGED_IN,STRIP_MERGE>
VFC client name:fcs0 VFC client DRC:U8286.41A.215D3AV-V17-C108

32

32

Fixing RMC Issues

IBM Power cloud:

ifconfig -a - check which en is inet6 - if none then make a note of the next free en
Replace en?? below with that en
odmdelete -o CuAt -q name=cluster0
autoconf6 -i en??
/usr/sbin/rsct/install/bin/uncfgct -n
sleep 5
/usr/sbin/rsct/install/bin/cfgct
sleep 5
/usr/sbin/rsct/bin/rmcctrl -z
/usr/sbin/rsct/bin/rmcctrl -A
/usr/sbin/rsct/bin/rmcctrl -p

Regular AIX:

https://www.ibm.com/support/pages/fixing-no-rmc-connection-error
Check RSCT levels
Check status of daemons
For >= HMCv8
diagrmc --autocorrect –v
Then on LPAR
/usr/sbin/rsct/bin/rmcctrl -z
/usr/sbin/rsct/bin/rmcctrl -A
/usr/sbin/rsct/bin/rmcctrl -p
See document for other commands that may be necessary

33

33

Inactive Migration

- Differences in configuration for inactive migrations:
  - LPAR may be configured with huge memory pages and barrier synchronization registers
  - Any physical I/O configured in the profiles will be removed during the migration
  - The preferred processor compatibility mode must be supported by the destination server
  - Mobile partition OS versions that support virtual devices and POWER6 servers but are not supported for active migrations may be supported for inactive migrations
- Either the last activated partition profile or the last running configuration is used, therefore a partition that has *never* been activated cannot be migrated
  - Workaround for new LPARs: Activate to SMS then shut down
- No log entries on HMC

34

34

Requirements for Remote Migration

- Ability to use LPM between 2 servers on different HMCs

- A local HMC managing the source server
- A remote HMC managing the target server
- Functional RMC daemons
- Version 7.3.4 or later of the HMC software
- Network access to the remote HMC
- SSH key authentication to the remote HMC and all involved LPARs (VIOS and actual LPAR)
- Plus all the other requirements for single HMC migration

- If FSM to HMC or vice versa then need HMC v7.7.1.0 or higher

35

35

Remote Migration SSH keys

- https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Power%20Systems/page/Remote%20LPM%20(Live%20Partition%20Migration)

- Setting up SSH Keys
- Create a userid to use for LPM - I created lpmuser on both systems

- Login to the source HMC to setup the communication channel to the target HMC:
- mkauthkeys –u target_hmc_user_name --ip target_hmc_ip --passwd target_hmcuser_password
- i.e.
- mkauthkeys –u lpmuser --ip 192.168.2.130 --passwd lpmpasswd

- After running mkauthkeys once, remote LPM will work from both HMC GUI and command line.
- Now test the communication channel:
- mkauthkeys –u target_hmc_user_name --ip target_hmc_ip –test
- i.e.
- mkauthkeys –u lpmuser --ip 192.168.2.130 –test

- If you want to setup remote LPM between two HMCs such that either can be a source HMC, you can run the mkauthkeys command on both HMCs, or on just one with the –g option:
- mkauthkeys –u target_hmc_user_name --ip target_hmc_ip --passwd target_hmc_user_password –g

36

36

Simplified Remote Restart (SRR)

Simplified Remote Restart (SRR)

- Introduced on POWER8
- Design for unplanned outages
  - Build a SRR plan ready to execute in the event of an unplanned outage
  - Lets you move and restart an LPAR within minutes after a POWER server has crashed (rare)
  - LPM can't be used when server has crashed but SRR can

- Article by Nigel Griffiths:
  - https://www.ibm.com/support/pages/simplified-remote-restart-hmc-or-powervc
- IBM Tech Doc
  - https://www.ibm.com/support/pages/basics-simplified-remote-restart-srr-hmc

37

37

Setting SRR for an LPAR

Go into LPAR properties
At the bottom of the general section you will see the box for SRR
Check the box

Then on the server properties
Check the box to allow migration with inactive source storage VIOS

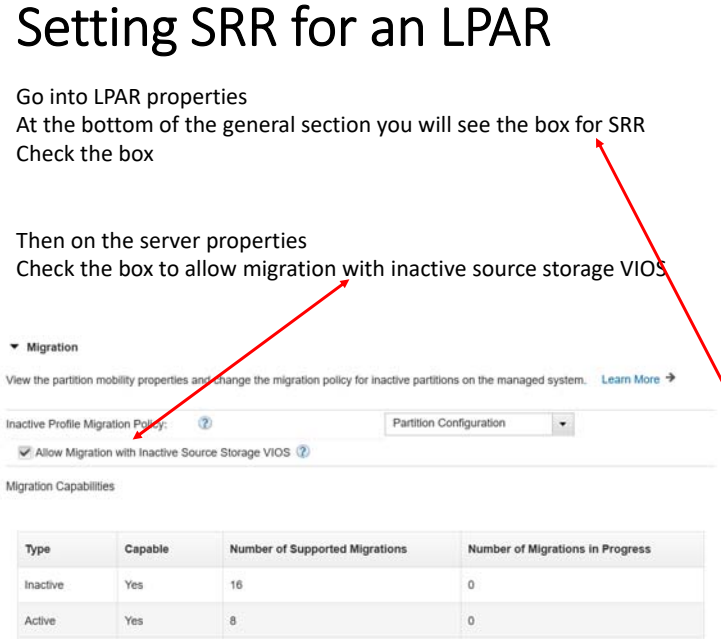

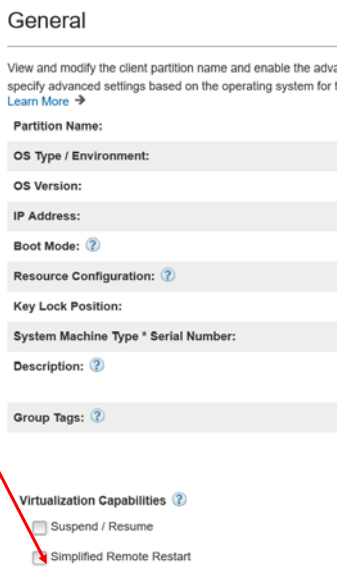

38

38

vioslpm0 – requires PowerVM 2.2.4 or higher

As padmin:
chdev -dev vioslpm0 -attr src_lun_val=on
chdev -dev vioslpm0 -attr dest_lun_val=on

The above changes are what will force it to check all the way through to the disk for any future lpms. This command is entered as padmin on all vio lpars and will never need to be entered again. It is dynamic.

You can check the settings:
oem_setup_env
lsattr –El vioslpm0

Be aware that if you have lots of disks then your validates may take a lot longer as it checks for access to all disks but it is worth it

During the validation it will check for every lun that is seen by the current WWPNs and will give a confusing error
Click on detailed information and you will see an error for every WWPN that cannot see the disk

https://www.ibm.com/support/knowledgecenter/en/POWER8/p8hc3/p8hc3_vioslpmpseudo.htm
https://www.ibm.com/support/knowledgecenter/en/POWER8/p8hc3/p8hc3_npivorlunval.htm

39

39

vioslpm0 new parameters

The **src_lun_val** attribute is used to enable and disable LUN level validation of N_Port ID Virtualization (NPIV) devices. This attribute had two possible values, *on* and *off*. When the attribute is set to *off*, LUN level validation is not performed, and when the attribute is set to *on*, LUN level validation is performed.

The **dest_lun_val** attribute is used to disable LUN level validation of NPIV devices for different operations and is relevant only when **src_lun_val** has the value *on* in the source VIOS. This attribute affects only the destination VIOS that is hosting the NPIV storage for remote restart and partition mobility operations. There are four allowed values for this attribute, *on*, *off*, *restart_off*, and *lpm_off*. By default the attribute is set to *restart_off*. This value disables LUN level validation for remote restart but allows it for partition mobility operations. Setting the attribute to *lpm_off* allows LUN level validation for remote restart operations but disables it for partition mobility operations. A value *on* allows LUN level validation for both partition mobility and remote restart and a value *off* disables LUN level validation for all operations.

40

40

Validation with src_lun_val and dest_lun_val=on

During the validation it will check for every lun that is seen by the current WWPNs and will give a confusing error
Click on detailed information and you will see an error for every WWPN that cannot see the disk

IOS_DETAILED_ERROR
Executed find_devices on VIOS 'vio1' (hostname: vio1)
Client Target WWPNs: 500500000000003e 500500000000003d
Target 0x500500000000003e found
Target 0x500500000000003d found
Matched 2 targets, source has 2 targets, destination has 2 targets
This physical port can not access storage for the client wwpn 'c050760a12340011'
Matched 0 targets, source has 2 targets, destination has 0 targets
Mismatching/Unique WWPNs on:
Source adapter : 0x500507680d04ef3e 0x500507680d08ef3d
Destination adapter :
fscsi2 is not zoned to the same target ports as the source for this client.

List of Logical Units not found on destination (i.e. NOT masked on storage target port = 0x500507680d04ef3e with client's alternate wwpn = 0xc050760a12340011; but masked with client's source wwpn = 0xc050760a12340010) :
Logical Unit 1 : descriptor type = 3, value = 60051234808123489C00000000000026.

41

41

Problems with LPM

Open a PMR with IBM using
https://www.ibm.com/mysupport/s/my-cases

Information is provided below on data you need to gather for Testcase for LPM
http://www-01.ibm.com/support/docview.wss?uid=isg3T1011601
Plus:
https://www.ibm.com/support/pages/complete-guide-must-gather-lpm-data-collection-powervc-vio-aix-linux-and-ibmi

Most problems are related to missing pre-requisites or HMC, VIOS, Firmware levels. Check all of these out prior to opening a PMR

Weird message about processor compatibility – you can get this if you have not activated the LPAR ever. Hint – look at the graphic – if you see 0.0 for PU it will not validate

Also have seen problems with vio lpar where 2 VFCs were mapped to the same FCS. You can do it but LPM won't work.

42

42

LPM and the HMC Command Line

1. To validate the LPM in the Remote HMC, type:
migrlpar -o v -m [source cec] -t [target cec] -p [lpar to migrate] --ip [target hmc] -u [remote user]

2. Get list of Servers
lssyscfg -r sys -F name

3. Get LPAR state
lssyscfg -r lpar -F name,state

4. Check migration state
lslparmigr -r lpar

5. Migrate LPAR
migrlpar -o v -m <srcCecName> -t <dstCecName> -p <lparName> -i "virtual_fc_mappings=
<Client_slot_num>/<target_vios_name>/<target_vios_id>/<target_slot_num>/<vios_fc_port_name>“

6. Can add --npivval portdisk to the end
Makes it validate npiv ports and disks

7. Can specify MSP pair
dest_msp_name=vios1, source_msp_name=vios

43

43

Thank you!

Jaqui Lynch

jlynch@flagshipsg.net

Please complete the Session Evaluation!

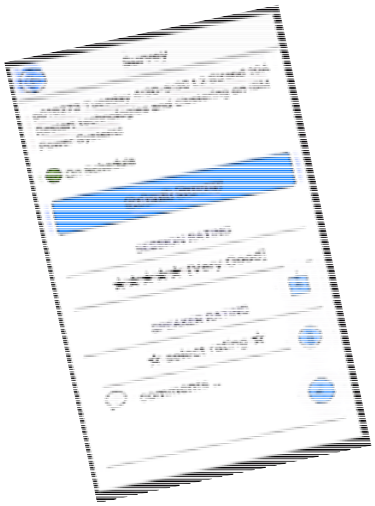

44

44

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/vioscare-part3-oct042020.pdf

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

45

45

45

Useful Commands, Links and Documentation

46

46

USEFUL COMMANDS

47

47

Useful Commands

Command History
$ fc -l
725 lsrep
726 backupios -file /usr/local/backups/b750viobkp
727 exit
728 lsmap -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 lsmap -vadapter vhost0
Aug 9 2013, 09:01:29 padmin lsgcl

Redirecting output when running as padmin
lsmap –all –npiv | tee npivdata.txt

48

48

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

49

49

Useful Commands

$ lsdev -virtual

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

50

50

Useful Commands

$ lsmap -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
```

51

51

Useful Commands

$ lsmap -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:
```

$ lsmap -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:a<LOGGED_IN,STRIP_MERGE>
VFC client name:fcs0            VFC client DRC:U8205.E6B.1093XXX-V8-C36
```

52

52

Useful Commands

$ lsnports

```
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
```

53

53

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
  - 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
  - Hint – upgrade to at least 2.2.6.32 prior to trying to upgrade to v3

- **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

54

54

Useful Links

- Jaqui Lynch Articles
  - http://www.circle4.com/jaqui/eserver.html
  - https://ibmsystemsmag.com/Authors/jaqui-lynch
- 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
- Nigel Griffiths YouTube
  - https://www.youtube.com/nigelargriffiths
- 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

55

55

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
- AIX 7.2 Performance Guide
  - https://www.ibm.com/support/knowledgecenter/ssw_aix_72/performance/performance_pdf.pdf
  - https://www.ibm.com/support/knowledgecenter/en/ssw_aix_72/navigation/performance.html
- VIOS Advisor
  - https://www.ibm.com/support/knowledgecenter/TI0002C/p8hcg/p8hcg_part.htm
  - https://www.ibm.com/support/knowledgecenter/TI0003N/p8hb1/p8hb1_vios_perf_adv.htm
  - https://www.ibm.com/support/knowledgecenter/TI0003M/p8hb1/p8hb1_vios_perf_adv_reports.htm
- 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

56

56

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
- 2019 AIX System Management Tips
  - https://ibmsystemsmag.com/Power-Systems/08/2019/2019-AIX-Systems-Management-Tips
- Secure your VIO Server
  - http://archive.ibmsystemsmag.com/aix/administrator/security/secure-your-vio-server/
- Upgrading your VIO server – July 2018
  - https://ibmsystemsmag.com/Power-Systems/12/2018/powervm-3-1-update
  - https://ibmsystemsmag.com/Power-Systems/05/2019/powervm-experience
- Maintaining the HMC
  - http://ibmsystemsmag.com/aix/administrator/systemsmanagement/hmc-maintenance/
- LPM
  - https://ibmsystemsmag.com/Power-Systems/10/2018/guide-live-partition-mobility
- HMC Enhanced GUI Links
  - https://www.ibm.com/support/pages/enhanced-gui-links-documentation

57

57

VIOS Specific References

- VIO Server Support
  - https://www14.software.ibm.com/support/customercare/sas/f/vios/home.html
- SDD and SDDPCM Specific procedures for VIOS
  - http://www-01.ibm.com/support/docview.wss?uid=ssg1S7002686&aid=1
- SG24-7940 - PowerVM Virtualization - Introduction and Configuration
  - http://www.redbooks.ibm.com/redbooks/pdfs/sg247940.pdf
- SG24-7590 – PowerVM Virtualization – Managing and Monitoring
  - http://www.redbooks.ibm.com/redbooks/pdfs/sg247590.pdf
- SG24-8080 – Power Systems Performance Guide – Implementing and Optimizing
  - http://www.redbooks.ibm.com/redbooks/pdfs/sg248080.pdf
- SG24-8062 – PowerVM Best Practices
  - http://www.redbooks.ibm.com/redbooks/pdfs/sg248062.pdf
- SEA Load Sharing
  - https://www.ibm.com/support/pages/how-setup-sea-failover-load-sharing-configuration
  - https://www.ibm.com/support/pages/shared-ethernet-adapter-sea-fail-over-load-balancing
- POWERVM Enhancements – what is new in 2013
  - http://www.redbooks.ibm.com/redbooks/pdfs/sg248198.pdf
- Capturing Debug output for padmin
  - http://www-01.ibm.com/support/docview.wss?uid=isg3T1012362

58

58

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
- Enhanced GUI Links
  - https://www.ibm.com/support/pages/enhanced-gui-links-documentation
  - Includes many Developerworks documents related to the HMC enhanced GUI
  - Includes how to dynamically add and remove virtual ethernets and VLANs
- 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

59

59