


IBM Systems ^{MEDIA}

What's New With AIX in 2020

AIX expert Jaqui Lynch breaks down the latest updates for the OS.



By Jaqui Lynch

02/03/2020

AIX version 7 has been around for a long time now, so it's time to look into some of the more recent updates that have come out with AIX 7.2 that may be useful in your data centers. Below we look into some of the updates that I have found to be interesting as I review them.

Multipath I/O

Starting with AIX 7.2 tl3, AIX now provides MPIO support for parallel SCSI, iSCSI and Fiber channel devices. MPIO is integrated into AIX with the AIXPCM drivers. SDDPCM is not supported on POWER9 servers. IBM also announced that SDDPCM will not be supported in Spectrum Virtualize after 8.3.0. It's also not supported on the SVC or Storwize hardware from machine code v7.6.1. No new fixes will be provided to SDDPCM at any level after June 30, 2020 (https://www.ibm.com/support/pages/node/1106937?myns=s035&mync=E&cm_sp=s035--NULL--E). This means it's time to migrate from SDDPCM to AIXPCM.

AIXPCM provides path management functions and health checks. The health checks allow AIXPCM to check the paths and determine which paths are usable for sending I/O, enable failed paths once they are returned to service and to check for unused paths that would be used if a failover occurred. It also removes many of the issues with upgrading OSes where the LPAR is booted from the SAN with SDDPCM drivers. With AIXPCM, there are two types of attributes that you need to review - device attributes and path attributes. The defaults are not necessarily what you would normally run with. There is an excellent writeup on all of these in the AIX modernization Redbook from IBM. Key parameters to review are:

```
algorithm = fail_over or round_robin or shortest_queue
```

COPY

The above impact how the I/O will be distributed across the paths. The default is fail_over which means the path with the highest priority will be chosen and other paths are only used if this one fails. round_robin distributes I/O across all the enabled paths depending on priority and whether paths are set to preferred or not. shortest_queue selects paths based on the length of the queue for that path. You can see these settings as follows:

```
lsattr -El hdisk4 | grep algo
```

```
algorithm    round_robin                Algorithm          True+
```

You should also check the settings for `hcheck_mode` and `reserve_policy`:

```
lsattr -El hdisk4 | grep hcheck
```

```
hcheck_mode  nonactive                 Health Check Mode True+
```

`hcheck_mode=nonactive` means the healthcheck command is sent to all paths that do not have any active I/O. This includes paths that are in enabled and failed states. This is the default.

Other modes include `enabled` which sends the command to only enabled and failed paths, or `failed` which only sends the command to failed paths.

```
lsattr -El hdisk4 | grep reserve
```

```
reserve_policy no_reserve                Reserve Policy    True+
```

The default `reserve_policy` is `single_path`. This policy uses SCSI2 reserves and will cause problems in a PowerHA or an LPM environment. `reserve_policy` should be set to `no_reserve` in these environments where disks are shared between hosts. It also should be set if devices have the `shortest_queue` or `round_robin` algorithms configured.

```
lsattr -El hdisk4 | grep queue
```

```
queue_depth  20                    Queue DEPTH      True+
```

`queue_depth` controls the number of concurrent commands that AIX can send to the device concurrently. This should be set according to performance needs and disk vendor recommendations.

You can check path status for an `hdisk` as follows—this shows all the paths for that disk and whether they are optimized (`opt`), selected for the I/O when `lsmpio` was run (`sel`) or non optimized (`non`)

```
lsmpio -l hdisk4
```

```
name  path_id status  path_status  parent connection
hdisk4 0    Enabled Non      fscsi0 500507680d048ef6,6000000000000
hdisk4 1    Enabled Sel,Opt fscsi0 500507680d048ef7,6000000000000
hdisk4 2    Enabled Non      fscsi1 500507680d088ef6,6000000000000
hdisk4 3    Enabled Sel,Opt fscsi1 500507680d088ef7,6000000000000
```

More detailed information on the disk can be found as follows:

```
lsmpio -ql hdisk4
```

Device: hdisk4

Vendor Id: IBM

Product Id: 2145

Revision: 0000

Capacity: 200.00GiB

Machine Type: 2078

Model Number: 124

Host Group: P8NIM

Volume Name: NIM_200GB_nimvg3

Volume Serial: 60050763808100F700000000000000026 (Page 83 NAA)

This provides the serial number, the name in the disk subsystem for the LUN, the size and the machine type and model for the disk subsystem.

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:

Adapter Driver: fscsi0 -> AIX PCM

Adapter WWPN: c0507607dbd80028

Link State: Up

Connection Errors

Last 10 Minutes: 0

Last 60 Minutes: 0

Last 24 Hours: 0

A more indepth set of statistics for the hdisk is available using the `lsmPIO -l hdisk4 -S` command. More details are available if you add the `-d` flag `lsmPIO -l hdisk4 -Sd`.

Additionally, `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
```

```
Enabled hdisk4 fscsi0 0
```

```
Enabled hdisk4 fscsi0 1
```

```
Enabled hdisk4 fscsi1 2
```

```
Enabled hdisk4 fscsi1 3
```

You can then narrow it down to only look at path 0 for the disk:

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

```
Enabled hdisk4 fscsi0
```

Or you can ask the system to report on all hdisks connected to path 0, which is useful if you have to take a path offline:

```
lspath -i 0
```

```
Enabled hdisk4 fscsi0
```

```
Enabled hdisk0 fscsi0
```

```
Enabled hdisk1 fscsi0
```

```
Enabled hdisk2 fscsi0
```

```
Enabled hdisk3 fscsi0
```

```
Enabled hdisk5 fscsi0
```

```
Enabled hdisk6 fscsi0
```

```
Enabled hdisk7 fscsi0
```

```
Enabled hdisk8 fscsi0
```

Enabled hdisk9 fscsi0

Enabled hdisk10 fscsi0

If the default settings for options such as `reserve_policy` are not acceptable you can change the default for devices so that new devices are always created with the setting you prefer. An example would be:

```
chdef -t mpioosdisk -c disk -s fcp -a reserve_policy=no_reserve
```

The above changes the default `reserve_policy` for new disks. It is recommended that you run `bosboot` after changing the default. You can check for all attributes that have been changed from defaults using `chdef -H`.

Additionally, paths can have their priorities changed, or be enabled or disabled using the `chpath` command.

It should be noted that when migrating from SDDPCM to AIXPCM all the values that you may have changed on devices will be reset to system defaults. It's important that you reset these after the migration.

iSCSI Support

As of AIX 7.2 TL 3 and above, AIX now has native support for using AIX as an iSCSI initiator by using any network adapter. After updating to AIX 7.2 TL3, there is a new device that is labeled `iscsi0` on your AIX system. This device is the iSCSI Protocol Device, and it is the parent adapter for all your iSCSI targets. It also holds the configuration attributes for your iSCSI software initiator. These values can be updated either by using `smitty` or by running the fast path `smitty chgiscsiw`. Settings can be checked using the `lsdev` and `lsattr` commands.

```
lsdev -C | grep iscsi
```

```
iscsi0      Available      iSCSI Protocol Device
```

```
lsattr -El iscsi0
```

```
disc_filename /etc/iscsi/targets Configuration file
disc_policy   file           Discovery Policy
immediate_data no             Support ImmediateData
initial_r2t   yes            Require InitialR2T
initiator_name iqn.localhost.hostid.7f000001 iSCSI Initiator Name
isns_srvnames auto          iSNS Servers IP Addresses
```

<code>isns_srvports</code>	iSNS Servers Port Numbers
<code>isw_err_recov delayed_fail</code>	iSCSI Network Error Recovery Policy
<code>max_targets 16</code>	Maximum Targets Allowed
<code>max_xfer_size 0x80000</code>	Maximum transfer
<code>num_cmd_elems 200</code>	Maximum number of commands to queue to driver

NIM Improvements

A large number of improvements have been made to NIM since it was first developed. Since AIX 7.2 there have been multiple features added to support new OS functions.

In order to support live update, a new management object class was added. This requires that you install `dsm.core`, but it then provides new object types that include `nas_filer`, `hmc`, `cec`, `vios`, `ivm`, `bcmm` (blade center management module) and `PowerVC`. This allows NIM to define and manage these control points. Live update specifically requires the definition of the HMC and CEC objects as well as an object for the LPAR and a new live update object. The new control points become available for configuration when `dsm.core` is installed.

The use of these new objects allows you to set up associations from the HMC to the CEC to the AIX LPAR so that you can perform operations such as live update which require authentication. There is a new “*dpasswd*” command to set up the authentication component.

In AIX 7.2 NIM now has a new service handler that provides HTTP access to NIM resources called `nimhttp`. The `nimhttp` service is defined in `/etc/services`, and the `nimhttp` daemon listens for requests over port 4901.

If `nimhttp` is configured and active then NIM clients can access some resources using the `http` service on port 4901 (or a port specified on `/etc/services`).

```
grep nimhttp /etc/services
```

```
nimhttp    4901/tcp
```

```
nimhttp    4901/udp
```

NIM resources that support HTTP access include:

- `file_res`

- `fix_bundle`
- `installp_bundle`
- `lpp_source`
- `script`

`nimhttp` is enabled using the `nimconfig -h` command which is not covered in the man page.

Cloning Updates

In my system maintenance I make extensive use of AIX clones, specifically with `alt_disk_copy`. I also occasionally use `alt_disk_mksysb` to create an alternate image from a `mksysb` image. There are times when I want to have more than one copy out there but you can only have one `altinst_rootvg` volume group at a time. This means it is necessary to rename the one you want to keep. This is done with `alt_rootvg_op`.

```
alt_rootvg_op -v rootvg1220 -d hdisk1
```

COPY

This renames `altinst_rootvg` on `hdisk1` to be `rootvg1220` and I can now create a second clone on another `hdisk`. If you need to perform operations on the clone after it is created, you can access it using:

```
alt_rootvg_op -W -d hdisk1
```

COPY

Next, put it back to sleep after making the changes using:

```
alt_rootvg_op -S -d hdisk1
```

COPY

These are useful commands to remember if you have to update something like `/etc/hosts` after the clone is taken.

AIX Toolbox for Linux

IBM has now moved to using `yum` to install third-party products such as `samba` or `gcc` that IBM provides using the AIX Linux Toolbox. This makes it far easier to install these products and ensure dependencies are being met. YUM is installed by running the

yum.sh found on the AIX Toolbox for Linux website and more information is provided in the yum readme on how to install and use YUM. If you have previously installed toolbox products outside the toolbox and yum then you may have some cleanup work to do prior to getting YUM to install.

SMT Defaults

With POWER9 and AIX 7.2 tl3 IBM has changed the default for SMT from SMT4 to SMT8. This now allows a maximum of 1536 threads for an LPAR.

AIX and PowerVM Secure Boot and POWER9 Firmware 940

Secure boot was introduced to AIX in 7.2 tl3 (also PowerVM v3). POWER9 hardware and firmware include new PowerVM features to support Secure boot. This allows only appropriately signed firmware components to run on the system processors. Each component of the firmware stack, including HostBoot, the IBM POWER Hypervisor (PHYP), and the Partition Firmware (PFW) is signed by the platform manufacturer and verified as part of the IPL process. AIX Secure Boot extends this to the AIX LPAR.

POWER9 firmware level 940 has code in it that supports Secure Boot at the firmware level. There's a warning in the readme that this requires your HMC to be at HMC v9R1M940 and that any boot capable PCI adapters must be updated to the latest firmware prior to upgrading to firmware 940 as the more recent levels include signed boot driver code. The LPAR will still boot if the adapter is not updated but will produce multiple error entries for BA5400A5 and BA5400A6.

There are also specific warnings that Novalink with certain SRIOV capable adapters is not supported at POWER9 firmware 930 or 940.

PowerVM v3.1

PowerVM v3.1 is the AIX 7.2 tl3 and higher version of the VIO server. The installation takes place using `alt_disk_copy` and requires a second disk to install to. It is a complete rootvg replacement so you will need to do any customizations again after the system has rebooted at the new level. Your virtual definitions will be backed up and restored but any changes you have made to logging, `/etc/hosts` or other rootvg files will need to be redone. You can tell the upgrade to save certain files for you in `/home/padmin` so you can restore the information later if you need to.

This version adds support for iSCSI over vSCSI and it also supports iSCSI MPIO. Additionally, it provides updates for NPIV. The 16GB adapters still have a maximum of 64 active VFCs but the 32GB adapters can now support 255 active VFCs.

Updates to LPM (Live Partition Mobility)

The default compatibility mode on POWER9 has been POWER9-base. After firmware 940 is installed there is a new mode added called POWER9. LPARs running in POWER9 compatibility mode are only supported on POWER9 servers running firmware 940. If you plan to LPM between POWER9 servers at different firmware levels make sure they are set to POWER9-base as a mode.

HMC v9r1m930 and POWER9 now allow for bidirectional concurrent migrations. Prior to these levels, migrations could only go in one direction at a time.

As of HMC v9r1m920 it's possible to use LPM to migrate LPARs between servers that have the same serial number although they must be different machine types and models. Each machine must be on a different HMC as well. This allows you to use LPM to migrate LPARs to a new server when using the same serial number as part of an upgrade; e.g., when upgrading from an E880 to an E980 using serial number protection.

General Updates

You can monitor the frequency your system is running at in several ways. Options include `mpstat`, `lparstat` and `nmon`. With `nmon` in logging mode you add “-y dfreq=on” to the command line. If you are in interactive mode you can use the R flag and then look for the CPU-Current-Speed:

```
CPU-Current-Speed 2047.1 MHz
```

```
mpstat -E 1 1
```

```
System configuration: lcpu=8 ent=0.5 mode=Uncapped
```

vcpu	pbusy	physc	freq	scaled physc
----	-----	-----	----	-----
0	0.0018[0%]	0.0101[2%]	2.0GHz[55%]	0.0056[1%]
4	0.0001[0%]	0.0003[0%]	2.8GHz[75%]	0.0002[0%]
ALL	0.0019[0%]	0.0103[2%]	2.0GHz[55%]	0.0057[1%]

Summary

There are many updates beyond what I have shared in this article that have come out in AIX 7.2. Every technology level adds more to the mix. AIX 7.1 is now coming close to end of life. All levels of AIX 7.1 prior to tl05 went out of support by 12/31/2019.

7.1.5 will still be supported until 4/30/2023. AIX 7.2 has been around since 2015 and IBM is slowly ending support for older technology levels; 7.2.1 went out of support on Nov. 20, 2019 and tl02 goes out of support on Oct. 31, 2020. This is a good time to look at updating to the latest tl (7.2.4.1) so that you are current, supported and can take advantage of all the new features that AIX has to offer. All PowerVM software prior to 2.2.6 also went out of support 9/30/2019 so this is also the time to update to PowerVM v3.1. I highly recommend reading the AIX Enhancements and Modernization Redbook so that you're aware of all of the new features that are now available.

References

AIX Enhancements and Modernization Redbook

<http://www.redbooks.ibm.com/redbooks/pdfs/sg248453.pdf> (<http://www.redbooks.ibm.com/redbooks/pdfs/sg248453.pdf>)

Migrating from SDDPCM to AIXPCM

<https://www.ibm.com/support/pages/how-migrate-sddpcm-aixpcm> (<https://www.ibm.com/support/pages/how-migrate-sddpcm-aixpcm>)

and

<https://www-01.ibm.com/support/docview.wss?uid=ssg1S1010646> (<https://www-01.ibm.com/support/docview.wss?uid=ssg1S1010646>)

Spectrum Virtualize MPIO Support

https://www.ibm.com/support/pages/node/1106937?myns=s035&mync=E&cm_sp=s035-_-NULL-_-E (https://www.ibm.com/support/pages/node/1106937?myns=s035&mync=E&cm_sp=s035-_-NULL-_-E)

AIX Linux Toolbox

<https://www.ibm.com/support/pages/aix-toolbox-linux-applications-downloads-alpha> (<https://www.ibm.com/support/pages/aix-toolbox-linux-applications-downloads-alpha>)

About the author

Jaqui Lynch has over 38 years of experience working with a projects and OSES across vendor platforms, including IBM Z, UNIX systems and more.

Related Content

[Systems management \(/systems-management\) Understanding AIX Versions → \(https://ibmsystemsmag.com/Power-Systems/11/2010/understanding-aix-versions\)](#)

[Systems management \(/systems-management\) Tips for Updating Server Firmware → \(https://ibmsystemsmag.com/Power-Systems/03/2012/tips-updating-server-firmware\)](#)

[IT infrastructure \(/it-infrastructure\) POWER8 Hardware Entitlement – Update Access Keys → \(https://ibmsystemsmag.com/Power-Systems/03/2015/power8-hardware-entitlement-firmware\)](#)



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.

© 2020 Key Enterprises LLC. All rights reserved