

[close window](#)

Web Exclusive

Print 

Avoiding Issues Using the 2012 AIX Performance Update

September 2012 | by [Jaqui Lynch](#)

Analyzing AIX performance and avoiding potential problems in the first place are essential for any systems administrator. Several steps should be part of your plans—the first of these is enrolling for Support Notifications. These can be tailored to fit the needs of your organization, and they ensure you notified as soon as IBM publishes fixes, allowing you to be proactive in avoiding problems. Sign up for storage and system notifications [here](#).

IBM has published fixes for a few known performance issues in AIX and POWER7. This is one reason it's important to stay as current as possible with maintenance. No one wants to take an outage over a problem that's already been found and has a fix available. These issues involve: affinity and dispatch, asynchronous I/O (AIO) server deadlocks, Virtual Partition Manager (VPM) serialization crashes, and slow Oracle Automatic Storage Management (ASM) startup. Specific fixes for these are available to various technology levels (TLs) in both AIX V6.1 and V7.1.

Additionally, performance issues are known around some of the older firmware levels, especially those in the 720 release. I monitor the firmware history files for my servers to check for HIPER and performance fixes. I also try to update firmware, Hardware Management Console (HMC), virtual I/O (VIO) servers and AIX at least every six months to avoid problems that have been already resolved. It's not a guarantee but it helps avoid most issues.

For example, in AL720_064 there's a hypervisor dispatch bug—this is fixed in the concurrent update, AL720_101. [The Power Firmware Code Matrix](#) lists the latest firmware version for all servers and provides a link to the supported HMC and server firmware combinations.

Remember, when updating firmware, it's critical to read the firmware description file, especially the sections titled "Minimum HMC Code Level" and "Important Information." If the HMC isn't at the minimum code level listed, then it should be updated prior to any upgrades to firmware.

You can also use the Fix Level Recommendation Tool (FLRT) to make recommendations for system levels. FLRT provides details on updates, including release and end-of-support dates. You can also click through to order the updates with it.

To avoid performance problems, I use a set of starter tunables whenever I set up a system. For AIX V6 and V7, I take the new defaults for the memory tunables (i.e. maxperm%, etc.) These are now restricted tunables that should only be changed if you have a problem management record (PMR) open and IBM tells you to do so. If this happens then you should record the PMR number, what the change was and why it was made.

I also recommend periodically validating that the recommendation is still valid. I work on many performance problems where a change was made because of a problem at a specific TL and subsequent OS updates mean that change is no longer valid and might actually be causing problems.

So what changes do I make in the starter set? The network parameters always need changing because the defaults are rarely correct. Do so using the `no` command to set the defaults as well as the `chdev` command to set the individual network adapters (use `ifconfig -a` to compare). The red parameters in

Table 1 are the new defaults and the system should be allowed to default to these.

Finally, let's look at performance issues. First, it's helpful to understand the process involved in opening a PMR for a performance problem. IBM's perfpmr website not only provides this tool, but also a link that discusses how to report performance problems. Whenever a performance problem occurs, IBM will request you run the perfpmr script, which collects performance data necessary to analyze the problem. I recommend getting familiar with running the tool before an emergency situation arises. However, you should always download the latest version when experiencing an actual problem as it's updated regularly.

Understand that performance tuning is an iterative process—it doesn't happen overnight. Even if all of the aforementioned recommendations have been followed, a quick fix might not be available when an issue arises—some problems take minutes to fix, others take months. So it's important to have a plan.

The first thing I recommend is to keep a baseline. Pick a tool (I like nmon and perfpmr.) and save a baseline of the system on a typical day when it's healthy. Whenever you make system changes, create a new baseline before and after so you have performance points for comparison. When you suspect a problem, rerun the baseline to see if anything is obvious to you. Without a baseline, you have nothing to compare your problem to, which makes resolution challenging. And, of course, good change control and documentation are critical.

The next step is to have a plan. Typically, it will look like this:

1. Describe the problem.
2. Measure where you're at (baseline).
3. Recreate the problem while getting diagnostic data (perfpmr, your own scripts, etc.).
4. Analyze the data.
5. Document potential changes and their expected impact, then group and prioritize them.
6. Make the changes.
7. Measure the results and analyze if they had the expected impact; if not, then why not?
8. Is the problem still the same? If not, return to step 1.
9. If it's the same, return to step 3.

This all sounds like common sense but, in an emergency, many of these steps get missed or rushed through. It's important to document and analyze the changes, especially if you end up in a PMR situation. One small change that only you know about can cause significant problems.

It's also beneficial to group changes into sets and to make them one set at a time. Alternatively, if you make 50 changes at a time, you'll never know which one fixed the problem. While many companies have limited maintenance windows, pushing through changes to everything in the infrastructure in one window is a recipe for disaster. Wherever possible, try to get more (but shorter) maintenance windows and use techniques like alternate disk install that will allow fast recovery if you need to go back after a change.

By putting into place the best practices outlined here you should be able to avoid most problems. And by following these steps when you do encounter an issue, the problem-resolution process should be less painful, especially if you gather good baselines for comparison. This will help IBM (and you) immensely, particularly if you build perfpmr into your baseline. Hopefully, you won't encounter performance issues, but if you do, these practices should help make the resolution process easier.

Starter Set of Tunables

	DEFAULTS			NEW
PARAMETER	AIX V5.3	AIX V6	AIX V7	SET ALL TO

NETWORK (no)				
rfc1323	0	0	0	1
tcp_sendspace	16384	16384	16384	262144 (1Gb)
tcp_recvspace	16384	16384	16384	262144 (1Gb)
udp_sendspace	9216	9216	9216	65536
udp_recvspace	42080	42080	42080	655360
MEMORY (vmo)				
minperm%	20	3	3	3
maxperm%	80	90	90	90
JFS, NFS, VxFS, JFS2				
maxclient%	80	90	90	90
JFS2, NFS				
lru_file_repage	1	0	0	0
lru_poll_interval	?	10	10	10
Minfree	960	960	960	calculation
Maxfree	1088	1088	1088	calculation
page_steal_method	0	0/1 (TL)	1	1
JFS2 (ioo)				
j2_dynamicBufferPreallocation	16	16	16	as needed

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.