

Manage AIX Workloads More Effectively Using WLM

December 2013 | by [Jaqui Lynch](#)

Since AIX V4.3.3, the free, built-in offering called Workload Manager (WLM) has allowed AIX administrators to consolidate workloads into one OS instance. WLM manages heterogeneous workloads, providing granular control of system CPU, real memory and disk I/O. It does this using percentages of resources and a combination of classes, shares and tiers to manage CPU time, memory and I/O bandwidth. WLM is integrated with the AIX kernel including the scheduler, Virtual Memory Manager and disk device drivers—running in either active or passive modes. Effectively a resource manager, WLM tracks:

- The sum of all CPU cycles consumed by every thread in the class
- The physical memory utilization of the processes in each class by looking at the sum of all memory pages belonging to the processes, and
- The disk I/O bandwidth in 512-byte blocks per second for all I/O started by threads in the class

A combination of targets and limits can be used to determine how resources are allocated. Targets are basically shares of available resources and range from 1 to 65,535, whereas limits are percentages of available resources with a minimum, soft maximum and hard maximum setting. Limits take priority over shares so, normally, only limits get defined. Hard limits have precedence, then tiers, soft limits and, finally, shares.

Tiers

WLM uses tiers to define class importance relative to other classes. You can define up to 10 tiers (0-9) to prioritize classes, with 0 being most important and 9 least important. WLM assigns resources to the highest tier process that's ready to run. Processes default to tier 0, but it's common to assign batch or less important workloads to tier 1 to prioritize online response.

Classes

A class is a set of processes with a single set of resource limits. An individual class is either a superclass or a subclass. Resource shares and limits are assigned to a superclass based on the total resources in the system. Subclasses can be defined to further divide a superclass's assigned resources among its assigned jobs.

Five predefined superclasses exist by default, and you can add up to 27 user-defined superclasses. Each superclass can have 12 subclasses: two predefined and 10 user-defined. Each class is assigned a name with a maximum of 16 characters. Superclass names must be unique, and subclass names must be unique within their assigned superclass.

Classification

The files necessary to classify and describe classes and tiers are defined in the `/etc/wlm` directory, beneath which is a directory created for the schema to be used. For example, if the `/etc/wlm/prodsys95` directory is the production system's definition for 9-5, an additional directory could be created for outside those hours. A simple cronjob command could switch between the two definition sets. Classification requires several files, including:

- Classes list each class, its description, the tier to which it belongs and other class attributes.
- Rules are where control is exerted over the class resources.
- Limits and shares define resource limits and shares, respectively.
- The optional description file includes a definition of the classes.

Threads are assigned to a class based on class rules. This can be done automatically using a rules file or manually by a superuser. Each class assigns minimum and maximum amounts for CPU, memory and I/O throughput. To correctly assign a process to a class, WLM goes through process identification, analyzing the process's attributes to see how it matches up with the class definitions. Processes can be identified and classified by owner or group ID, the full application path and name, the process type or a series of application tags. WLM assigns each class a set of resource shares and limits. Additionally classes can be assigned to tiers to further group and prioritize classes.

WLM reads the rules file from top to bottom and assigns a process to the first matching rule, which makes it important to list rules from specific (top) to more general (bottom). Processes are generally assigned to a class based on the user ID, group or fully qualified path and application name. Type and tag fields can also be used. Type can be 32bit, 64bit, "plock" (the process is locked to pin memory) or fixed (a fixed-priority process).

What WLM Does

For CPU, WLM gathers utilization for all threads in each class 10 times per second in AIX V5.3 and beyond. It then produces a time-delayed average for CPU for each class. That average is compared against tier values, class targets and class limits, and a number is produced that results in either favoring or penalizing each thread being dispatched. WLM also monitors memory 10 times per second and enforces memory limits using least recently used (LRU) algorithms. As of AIX V5.3, TL05, it's possible to set hard memory limits. However, if that limit is reached, the LRU daemon will start stealing pages from the class—even if memory pages are free. For I/O, WLM enforces limits by delaying I/O when the limit is reached.

Activation and Monitoring

WLM can be started in one of two modes: passive or active. Always start with passive so the effects can be monitored without making it active. Do this using the command `/usr/sbin/wlmcntrl -d prodsys -p` where `prodsys` is the directory name in which the configuration files live. Once the configuration has been proofed and monitored, WLM can be started in active mode by leaving off the `-p`.

Steps to Implementation

The first, most critical step is to design the classification criteria. Do so by evaluating workloads and determining how tiers and classes will be broken down, along with what limits should be applied to them. The second step is defining the class, limits, shares and rules files, and starting WLM in passive mode. Then refine the definitions before activating them. Once everything is tested, the final step is to restart WLM in active mode so it not only monitors—but also manages—the system.

WLM enables the consolidation of workloads into one AIX instance while ensuring that applications get the percentage of resources required to provide the performance necessary for success. This allows

workloads to be combined into the same OS instance to take better advantage of system resources. It's worth the effort to classify workloads running on your systems now, even if WLM only ever runs in passive mode. This provides an additional means of monitoring what's happening as well as a potential management tool, should the need arise.

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

Other WLM Commands

[< Return to main article](#)

Print Email

acctcom

Updated with a `-w` flag to list WLM information and a `-c` flag to list only specific classes, many other commands, such as `ps`, were also updated.

nmon

nmon has been updated to gather WLM statistics (use the `-w` flag).

nmon analyzer

The analyzer reports on WLM statistics, adding data in the BBBP tab and adding three new tabs that contain WLM information – WLMBIO, WLMCPU and WLMMEM.

ps -ae -o pid,user,class,pcpu,tag,thcount,vsz,wchan,args

This command provides a list of processes that includes WLM class information.

[< 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

Systems

Connect With Us:



Magazine Archives

Search

IBM i

LINUX ON POWER

MAINFRAME

POWER

AIX

ADMINISTRATOR

TRENDS

CASE STUDIES

TIPS & TECHNIQUES

STORAGE

PRODUCT NEWS

Resources

< Return to main article

Print Email

For more information on WLM, check out:

WLM Presentation and Replay

<http://www.circle4.com/papers/wlm-jul2013.pdf>

<http://www.circle4.com/movies/wlm/index.html>

IBM Systems Magazine Article on WLM by Jaqui Lynch (2004)

<http://www.ibmsystemsmag.com/aix/administrator/systemsmanagement/AIX-EXTRA--A-Look-at-Workload-Manager/>

WLM commands

<http://pic.dhe.ibm.com/infocenter/aix/v6r1/topic/com.ibm.aix.cmds/doc/aixcmds6/wlmassign.htm>

<http://pic.dhe.ibm.com/infocenter/aix/v6r1/topic/com.ibm.aix.cmds/doc/aixcmds6/wlmcheck.htm>

<http://pic.dhe.ibm.com/infocenter/aix/v6r1/topic/com.ibm.aix.cmds/doc/aixcmds6/wlmcntrl.htm>

<http://pic.dhe.ibm.com/infocenter/aix/v6r1/topic/com.ibm.aix.cmds/doc/aixcmds6/wlmstat.htm>

<http://pic.dhe.ibm.com/infocenter/aix/v6r1/topic/com.ibm.aix.cmds/doc/aixcmds6/wlmmon.htm>

<http://pic.dhe.ibm.com/infocenter/aix/v6r1/topic/com.ibm.aix.cmds/doc/aixcmds6/wlmpref.htm>

< 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.

