
AIX WLM (WORKLOAD MANAGER)



Jaqui Lynch

lynchi@forsythe.com

Presentation can be found at:

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

AGENDA

WLM Introduction

WLM Setup

WLM Commands

WLM and NMON

DEMO

Questions

AIX WLM INTRODUCTION

- Granular control of system CPU, real memory and disk I/O
- Been around since AIX v4.3.3
- Monitors and/or manages CPU, Memory and I/O resources
- Uses classes, shares and tiers
- Assigns threads to a class based on class rules
- Assign can also be done manually
- Assigns minimum and maximum percentages for resources
- Processes can be identified and classified by owner or group id, full path and name for application, process type or by application tags
- Each class is assigned shares and limits
- Classes can be assigned to tiers to group classes
- Can use it in passive or active modes
- Integrated with the AIX kernel – AIX scheduler, virtual memory manager and disk device drivers

RESOURCES MANAGED

- CPU utilization on the threads in a class
Sum of all CPU cycles consumed by every thread in the class
- Physical memory utilization of the processes in a class
Sum of all the memory pages belonging to processes in the class
- Disk I/O bandwidth of the class
Bandwidth in 512 byte blocks per second of all the I/Os started by threads of the class on the disk devices accessed

ALLOWS USERS TO DEFINE HOW RESOURCES ARE ALLOCATED

- Tiers
 - 0 (most favored) through 9 (least favored)
- Classes (superclasses)
- Subclasses
- Targets – relative shares of total available resource
 - Values from 1 to 65535
- Limits – percentages of total available resource
 - Minimum, soft maximum, hard maximum
- Limits take precedence over shares

HOW IT WORKS

- **URAP**

Uniform Resource Access Priority

1. The URAP is a Resource Access Priority – a number where lower means higher priority
2. WLM calculates percent usage for each resource one a second
3. WLM compares the average against class targets and class limits
4. This results in a URAP that is used to favor resources that do not use all their resources or to restrict others
5. URAP is a positive number
 1. Ranges from 60 (P_NICE_DEFAULT) and 255 (P_IDLE)
 2. 255 is used to send processes to the black hole (freeze them)
 3. Calculated per class, looking at superclass and subclass tiers and class resource usage
6. Constraints are given precedence in this order: hard limits, tiers, soft limits, shares
7. WLM reserves priority range for each superclass tier and a sub-range for each subclass tier

HOW IT WORKS

- **CPU**

1. WLM gathers CPU utilization for all threads in each class 10 times per second in AIX v5.3 and higher
2. CPU time used produces a time delayed average for each class
3. WLM compares the average against tier values, class targets and class limits
4. This results in a number that either favors or penalizes each thread when dispatched by the scheduler
5. Constraints are given precedence in this order: hard limits, tiers, soft limits, shares

- **Memory**

1. WLM monitors memory 10 times per second as well as of AIX v5.3
2. VMM enforces percentage based memory limits using LRU (Least recently used)
3. Classes get a URAP priority for accessing new memory pages
4. Memory hard limits can be set (real + paging) as of AIX v5.3 TL5
5. If reach hard limit then LRUD starts stealing pages from the class even if there are plenty of free memory pages

- **Devices**

1. If class needs to be restricted the I/O is delayed by WLM

MODES

- **Passive**

`/usr/sbin/wlmcntrl -d prod95 -p`

- **Active**

`/usr/sbin/wlmcntrl -d prod95`

`# wlmcntrl -q`

WLM is running in passive mode

- **Implementation**

Design classification criteria – don't get complicated

`mkdir /etc/wlm/prod95`

`cp /etc/wlm/template/* /etc/wlm/prod95`

Define classes, limits, shares and rules based on criteria chosen

Bring up in passive mode

Update nmon to use new `-W` flag

Monitor

When ready switch from passive to active mode

Or you can run in passive mode forever

FILES

- At a minimum you need the following files
 - classes
 - Defines each superclass with a description
 - limits
 - Sets the limits for each class
 - rules
 - Defines how each class is determined
 - shares
 - Defines any share limits (usually I leave this empty)
- You can also have the following
 - description
 - Contains a description for each class
 - superclassname
 - Contains subclass definitions for the superclass

CLASSES

- Collection of processes that has a single set of resource limits applied to it
- WLM uses class assignment rules to assign processes then uses shares and limits to control them
- A superclass or subclass is a class that is a set of resources with a single set of resources
- Shares and limits are assigned to a superclass based on total resources in the system
- There are 5 predefined superclasses by default
- You can add 27 more user defined superclasses (total of 32)
- Each superclass can have 12 subclasses (2 predefined and 10 user defined)
- A class name has a maximum of 16 characters made up of only upper and lower case characters
- Superclass names must be unique
- Subclass names must be unique within their superclass

PREDEFINED SUPERCLASSES

1. Default

All non-root processes not assigned elsewhere go here

2. System

All privileged root processes not assigned elsewhere go here

All memory pages belonging to system memory segments, kernel processes and kernel threads are assigned here

Defaults to a memory minimum limit of 1%

Cannot have a hard limit on memory due to potential deadlock situations

3. Shared

All memory pages shared by processes in more than one superclass go here. Includes pages in shared memory regions and pages in files that are used by processes in more than one superclass

Can only have memory shares or limits applied to it

4. Unclassified

Pages that WLM can't figure out how to classify go here

LRUD and wait processes will be here

Pages of a closed file in memory will go here

Cannot have memory shares or limits applied to it and cannot have subclasses

5. Unmanaged

Tracks memory usage for all pinned pages that are not managed by WLM

No processes go here

Cannot have shares or limits applied

PREDEFINED SUBCLASSES

1. Default

All processes not assigned to a specific subclass of the superclass

Can also assign other processes here using rules

2. Shared

All shared memory pages used by processes in more than one subclass of the superclass

Can only have shares and limits for memory applied to it

CLASSIFICATION

- Files are set up in /etc/wlm directory in a subdirectory for the schema
- Default is standard and this is linked as current
- Example schema name might be prod95 for production 9-5
 - /etc/wlm/prod95/
- Subclasses are defined in a configuration file for the superclass
 - /etc/wlm/prod95/superclassname
- Can use cron to automatically switch schemas if the rules are different at certain times or on certain days
- WLM reads rules from top to bottom and assigns a process to the FIRST matching rule
- Looks at superclass first then checks for subclasses
- Processes typically assigned based on userid, group id, or fully qualified path and application name
- Can also use types and tags
- Tags can be up to a 30 character alphanumeric

ASSIGNING THREADS TO CLASSES

This can be based on any of:

- Process owner id
- Process group id
- Executable name (using full path) – these can be Korn Shell patterns
i.e. /usr/oracle/bin/db*
- Application tags
- Type

Attribute value	Process type
32bit	The process is a 32 bit process.
64bit	The process is a 64 bit process.
plock	The process called plock() to pin memory.
fixed	The process is a fixed priority process (SHED_FIFO or SCHED_RR).

Can also use 64bit,plock+fixed or plock,fixed,64bit and so on
32 bit and 64 bit are mutually exclusive

CLASS ATTRIBUTES

- *Class name:*

Up to 16 characters long. Can contain only uppercase and lowercase letters, numbers, and underscores.

- *Tier:*

Number between 0 and 9 for class priority ranking.

- *Inheritance:*

Specifies whether or not a child process inherits the class assignment from its parent.

- *Adminuser, admingroup (superclass only):*

Used to delegate the administration of a superclass.

- *Authuser, authgroup:*

Used to delegate the right to manually assign a process to a class.

- *Resource Set:*

Used to limit the set of resources to which a given class has access in terms of CPUs (processor set).

- *Localshm:*

Specifies whether memory segments that are accessed by processes in different classes remain local to the class they were initially assigned to, or if they go to the Shared class.

FROM REDBOOK SG24-5977 Page 17

INHERITANCE FOR CLASSES

Table 1. Inheritance attribute at superclass and subclass level meaning

Superclass level inheritance value	Subclass level inheritance value	Meaning
yes	yes	A child of a process in the subclass will remain in the same subclass upon <i>exec</i> .
yes	no or unspecified	A child of a process in the subclass will remain in the same superclass and will be classified in one of its subclasses according to the assignment rules for the superclass upon <i>exec</i> .
no or unspecified	yes	A child of a process in the subclass will be submitted to the automatic assignment rules for the superclasses upon <i>exec</i> . If the process is classified by the rules in the same superclass, it will remain in the subclass (it will not be submitted to the subclasses assignment rules). If the process is classified by the superclass rules in a different superclass, the subclass assignment rules of the new superclass are applied to determine the subclass of the new superclass to which the process will be assigned.
no or unspecified	no or unspecified	A child of a process in the subclass will be submitted to the standard automatic assignment upon <i>exec</i> .

FROM REDBOOK SG24-5977 Page 18

TIERS

- Tiers are used to prioritize classes relative to other classes
- Can have up to 10 tiers (0 – 9) to prioritize classes
- 0 is most important and 9 is least important
- Default tier is 0
- Tier applies to both the superclass and the subclass
- Process in tier 0 will always have precedence over one in tier 1 or lower

TIERS

To add tiers for a superclass we edit the superclass classes file at:
/etc/wlm/standard/classes and change it to add tiers as follows:

System:

Default:

GrpA:

localshm = yes
adminuser = adminA
authuser = adminA
inheritance = yes

GrpB:

localshm = yes
adminuser = adminB
authuser = adminB
inheritance = yes

SysTools:

localshm = yes

SysBatch:

tier = 1
localshm = yes

Sysbatch goes into tier 1 as it is low priority

All others default to tier 0

Administration of the superclasses is defined by the adminuser and authuser attributes

Inheritance = yes means that any new process started in that class will remain classified in that class

SHARES

- Shares are relative entitlement to available resources
- Superclass shares are compared to other classes in the same tier
- Subclass shares are compared to other subclasses in the same superclass
- Limits can be imposed using shares
- Limits take precedence over shares

SHARES

This file has to be there but can be empty

Or you can use it to assign the initial shares

For the superclass we could edit /etc/wlm/standard/shares

GrpA:

CPU = 3

memory = 3

GrpB:

CPU = 2

memory = 2

The above gives GrpA 3 out of 5 shares for CPU or 60% of the CPU and GrpB 2 out of 5 or 40%

Note if nothing in GrpB then GrpA is really 100%

SysTools, System and Default do not have shares assigned so they have higher priority to the resources and are independent
SysBatch is the only thing in tier1 so shares would be meaningless there. Jobs in tier 1 can only consume resources not being used by tier 0

For the subclass GrpA we edit /etc/wlm/standard/GrpA/shares

Work:

CPU = 5

memory = 5

Monitor:

CPU = 4

memory = 1

Command:

CPU = 1 memory = 1

We do the same for the subclass GrpB

Report is in tier 1 so no need for shares for it – it gets whatever tier 0 is not using

Listen has highest priority by default since it is not assigned shares

Work = 5/10 or 50%, Monitor = 40% and Command = 10%

LIMITS

- Define the minimum and maximum amount of resources to be allocated
- Minimum must be available when requested until target utilization is reached
 - Minimum must be \leq to soft maximum
 - Sum of superclass minimums within a tier cannot exceed 100%
 - Sum of subclass minimums within a superclass cannot exceed 100%
- Maximum is percentage that can be made available when there is contention with other classes
- If no contention then soft maximum can be exceeded
 - Soft maximum must be \leq hard maximum
 - Hard maximum can never be exceeded
- Cannot set a hard maximum on the System class
- Note:
If the maximum limit configured for a class exceeds the system wide limit in /etc/secure/limits, processes in the class will be killed if they reach the system wide limit regardless of the WLM configured limit.

SAMPLE LIMITS

The file must be there but can be empty

For Superclass create /etc/wlm/standard/limits to include:

Default:

CPU = 0%-10%;100%
memory = 0%-10%;100%

SysTools:

CPU = 0%-10%;100%
memory = 0%-5%;100%

System:

CPU = 5%-50%;100%
memory = 5%-50%;100%

In the above we assigned soft limits and allow it to go up to 100%
System is guaranteed a minimum of 5% of cpu and memory

For subclass GrpA create /etc/wlm/standard/GrpA/limits to include:

Listen:

CPU = 10%-30%;100%
memory = 10%-20%;100%

Monitor:

CPU = 0%-30%;100%
memory = 0%-30%;100%

NOTE

Minimum limit for memory

Ensures that pages used by this class will not be stolen by page replacement, resulting in faster execution time.

Minimum limit for CPU

Ensures these processes have the highest-priority access (in the superclass) to the CPU resources.

SAMPLE LIMITS

System:

CPU = 5%-100%;100%
memory = 5%-100%;100%

SAS:

memory = 0%-60%;60%

Stata:

memory = 0%-60%;60%

In the above the system is guaranteed 5% of memory and CPU
It is soft capped so can go to 100%

For both SAS and Stata they are hard capped at 60% of memory
No CPU limits are defined

/ETC/WLM CONTENTS

```
# cd /etc/wlm
# ls -al
total 64
drwxr-xr-x  6 root  system    256 Jul 15 19:23 .
drwxr-xr-x 36 root  system   28672 Jul 17 10:23 ..
-----  1 root  system     0 Jun 18 11:04 .lock
dr-xr-xr-x  3 root  system    256 Jul 15 18:46 .running
lrwxrwxrwx  1 root  system    17 Jul 15 19:23 current -> /etc/wlm/standard
drwxr-xr-x  2 root  system    256 Jul 15 17:58 prod95
drwxr-xr-x  3 root  system   4096 Jul 15 18:43 standard
drwxr-xr-x  2 root  system    256 Aug 27 2011 template
```

/ETC/WLM/STANDARD CONTENTS

ls -al standard – this is the superset

total 48

```
drwxr-xr-x  3 root  system    4096 Jul 15 18:43 .
drwxr-xr-x  6 root  system    256 Jul 15 19:23 ..
-rw-r--r--  1 root  system      0 Jul 15 19:23 .status
-rw-r--r--  1 root  system    729 Aug 27 2011 README
-rw-r----- 1 root  system    729 Jun 18 11:03 README.txt
drwxr-xr-x  2 root  system    256 Jul 15 18:46 SSHD
-rw-r--r--  1 root  system    213 Jul 15 18:42 classes
-rw-r--r--  1 root  system     79 Jul 15 18:09 limits
-rw-r--r--  1 root  system    963 Jul 15 18:43 rules
-rw-r--r--  1 root  system      0 Jun 18 11:03 shares
```

SSHD is a directory so this means it is defining a subclass

/ETC/WLM/STANDARD/SSHD CONTENTS

```
# ls -al standard/SSHD – this is the subclass for SSHD
total 24
drwxr-xr-x  2 root  system    256 Jul 15 18:46 .
drwxr-xr-x  3 root  system   4096 Jul 15 18:43 ..
-rw-r--r--  1 root  system     0 Jul 15 18:43 .lock
-rw-r--r--  1 root  system     0 Jul 15 19:23 .status
-rw-r--r--  1 root  system    56 Jul 15 18:21 classes
-rw-r--r--  1 root  system     0 Jul 15 18:44 limits
-rw-r--r--  1 root  system   370 Jul 15 18:28 rules
-rw-r--r--  1 root  system     0 Jul 15 18:44 shares
```

MY STRUCTURE FOR THIS TEST

Superclasses:

System

Default

Shared

SSHD

Monit

Backups

These will be in tier1 (all others is Tier 0)

Subclasses:

SSHD.Jaqui

Userid jaqui

Limits:

System

Guarantee 5% memory and 1% CPU, soft cap

Monit

Hard cap at 60% memory

SETTINGS 1/3

/etc/wlm/standard/classes

System:

Default:

Shared:

SSHD:

description = "SSH"

inheritance = "yes"

Monit:

description = "Monitoring"

inheritance = "yes"

Backups:

description = "mksysbs"

tier = 1

inheritance = "yes"

SETTINGS 2/3

/etc/wlm/standard/limits

System:

CPU = 1%-100%;100%

memory = 5%-100%;100%

Monit:

memory = 0%-60%;60%

SETTINGS 3/3

/etc/wlm/standard/rules

```
* class resvd user group application type tag
SSHD - - - /usr/sbin/sshd - -
Backups - - - /usr/bin/mksysb - -
Monit - - - /usr/bin/netstat - -
Monit - - - /usr/bin/iostat - -
Monit - - - /usr/bin/vmstat - -
Monit - - - /usr/bin/topasrec - -
Monit - - - /usr/sbin/sar - -
Monit - - - /usr/bin/topas_nmon - -
Monit - - - /usr/lib/sa/sadc - -
System - root - - - -
Default - - - - -
```

/etc/wlm/standard/shares

Is empty

UPDATING STANDARD CONFIG

Go into /etc/wlm/standard and edit the files

```
# wlmctrl -q  
WLM is running in passive mode  
# wlmctrl -o  
# wlmctrl -q  
WLM is stopped  
# wlmctrl -p -d standard  
# wlmctrl -q  
WLM is running in passive mode
```

```
# wlmstat
```

CLASS	CPU	MEM	DKIO
Unclassified	0.00	24.95	0.00
Unmanaged	0.00	25.73	0.00
Default	0.35	11.29	0.00
Shared	0.00	7.69	0.00
System	0.46	48.53	0.00
SSHD	0.11	0.08	0.00
Monit	0.01	0.93	0.00
TOTAL	0.93	93.47	0.00

UPDATING STANDARD CONFIG – ADD SUBCLASSES

Go into `/etc/wlm/standard/SSHD` and edit the files `classes` & `rules`

Can use tiers in these subclasses – I did not bother

/etc/wlm/SSHD/classes

Jaqui:

description = "Jaqui work"

inheritance = "yes"

/etc/wlm/SSHD/rules

* class resvd user group application type tag

Jaqui - jaqui - - - -

Also touch `.lock`, `limits`, `shares` – they have to be there

```
# wlmcheck -d standard
```

```
WLM is running in passive mode, Rset bindings active.
```

```
Checking classes and rules for 'standard' configuration...
```

```
System
```

```
Default
```

```
Shared
```

```
SSHD
```

```
Monit
```

```
Backups
```

Stop and start again

COMMANDS

- smitty wlm
- wlmassign—Manually assign a process to a class.
 wlmassign super1.sub2 -S pid1
- wlmcheck—Check assignment rules to see how to classify a resource.
- wlmcntrl—Stop and start WLM, switch it from passive to active mode and query what mode it's running in.
- wlmstat—Show WLM per class resource utilization.
- wlmmon and wlmperf—Provide graphical views of resource activities by class.
- 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 – updated to gather wlm statistics (-W flag)
- nmon analyzer – updated to report on WLM statistics
 New data in the BBBP tab
 Three new tabs – WLM BIO, WLM CPU, WLM MEM

SMITTY WLM

```
File Edit View Window Help
Quick Connect Profiles
Workload Manager
Move cursor to desired item and press Enter.
Manage time-based configuration sets
Work on alternate configurations
Work on a set of Subclasses
Show current focus (Configuration, Class Set)
List all classes
Add a class
Change / Show Characteristics of a class
Remove a Class
Class assignment rules
Start/Stop/Update WLM
Assign/Unassign processes to a class/subclass
F1=Help F2=Refresh F3=Cancel Esc+8=Image
Esc+9=Shell Esc+0=Exit Enter=Do
Connected to bpicd SSH2 - aes128-cbc - hmac-md5 - r 141x30
```

PS COMMAND

- ps command

```
ps -ae -o class,pid,user,args >outfilea  
sort outfilea -o $outfilea-sorted.txt
```

The above provides a list of processes sorted by class – you can add other flags to get memory etc information (pid,user,class,pcpu,tag,thcount,vsz,wchan,args)

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

PID	USER	CLASS	%CPU	TAG	THCNT	VSZ	WCHAN	COMMAND
1	root	System	0.0	-	1	776	-	/etc/init
1835150	root	System	0.0	-	1	396	EVENT	/usr/ccs/bin/shlap64
2293978	root	System	0.0	-	17	588	*	/usr/sbin/syncd 60
3211418	root	System	0.0	-	1	1016	EVENT	/usr/sbin/cron
3276810	dirinst1	Default	0.0	-	67	47252	*	db2sysc 0
3473568	root	System	0.0	-	1	332	EVENT	/opt/freeware/cimom/pegasus/bin/cimssys platform
3604670	root	System	0.0	-	1	864	-	/usr/sbin/srcmstr
3735676	root	System	0.0	-	1	1000	-	/usr/sbin/snmpd
3801280	root	System	0.0	-	1	5244	EVENT	/opt/ibm/director/cimom/bin/tier1slp
5898476	root	Monit	0.0	-	1	19240	EVENT	/usr/bin/topasrec -L -s 300 -R 1 -r 6 -o /etc/p
ystart_time=13:23:19,Jan02,2013								
6226014	dirinst1	Default	0.0	-	7	6160	*	db2fmp (C) 0
6291682	root	System	0.0	-	1	548	-	/bin/ksh /pconsole/lwi/bin/lwistart_src.sh
7340262	root	SSHD.Default	0.0	-	1	752	-	/usr/sbin/sshd
7602428	root	System	0.0	-	8	6880	*	/usr/sbin/rsct/bin/rmcd -a IBM.LPCommands -r
10289152	root	SSHD.Default	0.0	-	1	572	-	-ksh
15466626	root	SSHD.Default	0.0	-	1	892	-	ps -ae -o pid,user,class,pcpu,tag,thcount,vsz,wcl
17170512	root	SSHD.Default	0.0	-	1	1188	-	sshd: root@pts/0
19464238	root	SSHD.Default	0.0	-	1	560	EVENT	-ksh
20119560	root	Monit	0.0	-	1	2620	EVENT	/usr/bin/topas_nmon -ft -A -M -L -^ -W -s 150 -c
ed_ystart_time=23:59:01,Jan16,2013								

PS -AE

```
# ps -ae -o class,pid,user,args | grep jaqui
Monit          6619388  jaqui /usr/bin/topas_nmon
SSHD.Jaqui     8650894  jaqui sshd: jaqui@pts/1
SSHD.Default  14024954  root sshd: jaqui [priv]
Default       16449536  jaqui /usr/bin/perl /usr/bin/nmon
System        17629326  root grep jaqui
Default       19464398  jaqui -ksh
```

PS COMMAND – CHECK SUPERCLASS AND SUBCLASS

ps -c SSHD

PID	TTY	TIME	CMD
7340262	-	0:00	sshd
10289152	pts/1	0:00	ksh
12386504	pts/2	0:00	ksh
13959358	pts/2	0:00	perl
15401150	pts/1	0:00	ps
15466550	-	0:00	sshd
15597588	-	0:00	sshd
15859748	-	0:00	sshd
17170512	-	0:00	sshd
17629226	pts/2	0:00	topas_nmon
19464238	pts/0	0:00	ksh

ps -c SSHD.Jaqui

PID	TTY	TIME	CMD
12386504	pts/2	0:00	ksh
13959358	pts/2	0:00	perl
15466550	-	0:00	sshd
17629226	pts/2	0:00	topas_nmon

WLMCNTRL AND WLMCHECK

- wlmcntrl

wlmcntrl -q

WLM is running in passive mode

wlmcntrl -o stops WLM

wlmcntrl -p -d standard start WLM in passive mode using the standard config

wlmcntrl -u -d standard Refresh the settings for WLM

If you just type in wlmcntrl you will accidentally put it in active mode so be careful

- wlmcheck

wlmcheck -d config -a attributes -q

wlmcheck -a "- jaqui"

Above shows me the processes owned by jaqui and what categories they fall into

WLMCHECK

```
# wlmcheck -d standard
```

```
WLM is running in passive mode, Rset bindings active.
```

```
Checking classes and rules for 'standard' configuration...
```

```
System
```

```
Default
```

```
Shared
```

```
SSHD
```

```
Monit
```

```
Backups
```

```
# wlmcheck -a "- jaqui"
```

```
SSHD.Jaqui
```

```
Backups
```

```
Monit
```

```
Default
```

WLMSTAT

wlmstat or wlmstat -a

CLASS	CPU	MEM	DKIO
Unclassified	0.00	24.62	0.00
Unmanaged	0.00	25.74	0.00
Default	0.06	11.59	0.00
Shared	0.00	7.69	0.00
System	0.91	48.66	0.00
SSHD	0.12	0.21	0.00
SSHD.Default	0.12	0.19	0.00
SSHD.Shared	0.00	0.00	0.00
SSHD.Jaqui	0.00	0.02	0.00
Monit	0.01	0.96	0.00
Backups	0.00	0.00	0.00
TOTAL	1.10	93.73	0.00

WLMSTAT FLAGS

You can use the following flags to limit the columns

- b *only show DKIO*
- B hdisk2 *only show DKIO for hdisk2*
- B "" *only show DKIO but for all disks*
- c *only show CPU*
- l class *display class statistics (lower case L)*
- l SSHD *show all 3 fields for the SSHD class*
- m *only show MEM*
- M *display real/virtual memory statistics*
adds RMSIZ, VMSIZ, RMLIM, VMLIM, LGPGSIZ, LGPGLIM
if limit is undefined you may see a –
- s *show only subclass statistics*
- S *show only superclass statistics*
- t 0 *show only tier 0 statistics*
- T *display resource utilization totals since WLM was started*
- v *verbose mode – adds a number of columns*
- u *show per tier and total unused resources*
- w *show the memory high water mark – max pages a class has used since WLM started*

WLMSTAT -M

wlmstat -M

CLASS	RMSIZ	RMLIM	VMSIZ	VMLIM	LGPGSIZ	LGPG LIM
Unclassified	531	-	531	-	0	-
Unmanaged	0	-	0	-	0	-
Default	343	-	343	-	0	-
Shared	233	-	233	-	0	-
System	2078	-	2078	-	0	-
SSHD	3	-	3	-	0	-
SSHD.Default	3	-	3	-	0	-
SSHD.Shared	0	-	0	-	0	-
SSHD.Jaqui	0	-	0	-	0	-
Monit	21	-	21	-	0	-
Backups	0	-	0	-	0	-
TOTAL	3209	-	3209	-	0	-

Display real/virtual memory statistics

WLMSTAT -T WITH TIERS

```
# wlmstat -T
```

CLASS	CPU	MEM	DKIO	PROCESSES	THREADS	LOGINS
Unclassified	0	0	1816	0	0	0
Unmanaged	23	0	0	0	0	0
Default	3105	0	33840	9	117	3
Shared	0	0	0	0	0	0
System	9794	0	32072	90	717	52
SSHD	0	0	0	4	4	3
SSHD.Default	0	0	0	3	3	3
SSHD.Shared	0	0	0	0	0	0
SSHD.Jaqui	0	0	0	1	1	0
Monit	1226	0	1296	2	2	0
Backups	0	0	0	0	0	0
TOTAL	14148	0	69024	105	840	58

Displays the total numbers for resource utilization since WLM was started or the class was created, whichever is the latter. The units are:

CPU The total processor time, in milliseconds, consumed by a class

MEM Unused

DKIO The total number of 512 byte blocks sent/received by a class for all the disk devices accessed.

WLMSTAT -T

```
# wlmstat -t 1
  CLASS    CPU    MEM    DKIO
Backups    0.00   0.00   0.00
TOTAL      0.00   0.00   0.00
```

```
# wlmstat -t 0
  CLASS    CPU    MEM    DKIO
Unclassified 0.00  24.88  0.00
Unmanaged    0.00  25.74  0.00
Default      0.05  11.40  0.00
Shared       0.00   7.69  0.00
System       0.27  48.42  0.00
SSHD         0.06   0.21  0.00
SSHD.Default 100.00 90.48  -
SSHD.Shared  0.00   0.00  -
SSHD.Jaqui   0.00   9.52  -
Monit        0.00   0.94  0.00
TOTAL        0.38  93.54  0.00
```

```
# wlmstat
  CLASS    CPU    MEM    DKIO
Unclassified 0.00  24.96  0.00
Unmanaged    0.00  25.74  0.00
Default      0.06  11.32  0.00
Shared       0.00   7.69  0.00
System       0.28  48.42  0.00
SSHD         0.09   0.21  0.00
SSHD.Default 100.00 90.48  -
SSHD.Shared  0.00   0.00  -
SSHD.Jaqui   0.00   9.52  -
Monit        0.00   0.94  0.00
Backups      0.00   0.00  0.00
TOTAL        0.43  93.54  0.00
```

WLMSTAT -V

-v Verbose mode. Intended for trouble shooting. Displays some class attributes, resource shares and limits and other WLM parameters, including internal parameter values intended for AIX(R) support personnel.

Column Headers:

tr tier number (0 to 9)
I Value of the inheritance attribute: 0 = no, 1 = yes.
#pr Number of processes in the class. If a class has no (0) process assigned to it, the values shown in the other columns might not be significant.
CPU Processor utilization of the class (%).
MEM Physical memory utilization of the class (%).
DKIO Disk IO bandwidth utilization for the class (%).
Sha Number of shares ('-' is represented as -1)
Min Resource minimum limit (%)
Smx Resource soft maximum limit (%)
Hmx Resource hard maximum limit (%)
Des (desired): percentage goal (target) calculated by WLM using the shares numbers (%)
Npg Number of memory pages owned by the class.

The other columns are for internal use

This format is better used with a resource selector (-c, -m, or -b), otherwise the lines might be too long to fit into a line of a display terminal.

WLMSTAT -VC

```
# wlmstat -vc
```

CLASS	tr	i	#pr	CPU	sha	min	smx	hmx	des	rap	urap	pri
Unclassified	0	0	0	0.00	-1	0.00	100.00	100.00	100.00	100	0	0
Unmanaged	0	0	0	0.00	-1	0.00	100.00	100.00	100.00	100	0	0
Default	0	0	10	0.05	-1	0.00	100.00	100.00	95.00	100	0	0
Default.Default	0	0	10	100.00	-1	0.00	100.00	100.00	100.00	100	0	0
Default.Shared	0	0	0	0.00	-1	0.00	100.00	100.00	100.00	100	0	0
Shared	0	0	0	0.00	-1	0.00	100.00	100.00	100.00	100	0	0
Shared.Default	0	0	0	-	-1	0.00	100.00	100.00	100.00	100	0	0
Shared.Shared	0	0	0	-	-1	0.00	100.00	100.00	100.00	100	0	0
System	0	0	90	0.36	-1	5.00	100.00	100.00	100.00	100	0	0
System.Default	0	0	90	100.00	-1	0.00	100.00	100.00	100.00	100	0	0
System.Shared	0	0	0	0.00	-1	0.00	100.00	100.00	100.00	100	0	0
SSHD	0	1	4	0.10	-1	0.00	100.00	100.00	95.00	100	0	0
SSHD.Default	0	0	3	100.00	-1	0.00	100.00	100.00	100.00	100	0	0
SSHD.Shared	0	0	0	0.00	-1	0.00	100.00	100.00	100.00	100	0	0
SSHD.Jaqui	0	1	1	0.00	-1	0.00	100.00	100.00	100.00	100	0	0
Monit	0	1	3	0.01	-1	0.00	100.00	100.00	95.00	100	0	0
Monit.Default	0	0	3	100.00	-1	0.00	100.00	100.00	100.00	100	0	0
Monit.Shared	0	0	0	0.00	-1	0.00	100.00	100.00	100.00	100	0	0
Backups	1	1	0	0.00	-1	0.00	100.00	100.00	100.00	100	97	97
Backups.Default	0	0	0	-	-1	0.00	100.00	100.00	100.00	100	97	97
Backups.Shared	0	0	0	-	-1	0.00	100.00	100.00	100.00	100	97	97

It shows that Backups is in tier 1

Remember:

CPU Processor utilization of the class (%).

WLMSTAT -VM

```
# wlmstat -vm
```

CLASS	tr	i	#pr	MEM	sha	min	smx	hmx	des	rap	urap	npg
Unclassified	0	0	0	24.69	-1	0.00	100.00	100.00	98.00	59	8191	192110
Unmanaged	0	0	0	25.74	-1	1.00	100.00	100.00	99.00	58	855	269905
Default	0	0	10	11.71	-1	0.00	100.00	100.00	98.00	78	448	91123
Default.Default	0	0	10	100.00	-1	0.00	100.00	100.00	100.00	78	448	91123
Default.Shared	0	0	0	0.00	-1	0.00	100.00	100.00	100.00	100	224	0
Shared	0	0	0	7.69	-1	0.00	100.00	100.00	98.00	85	305	59856
Shared.Default	0	0	0	100.00	-1	0.00	100.00	100.00	100.00	85	305	59856
Shared.Shared	0	0	0	0.00	-1	0.00	100.00	100.00	100.00	100	152	0
System	0	0	90	48.58	-1	1.00	100.00	100.00	99.00	34	1344	378128
System.Default	0	0	90	100.00	-1	0.00	100.00	100.00	100.00	34	1344	378128
System.Shared	0	0	0	0.00	-1	0.00	100.00	100.00	100.00	100	672	0
SSHD	0	1	4	0.21	-1	0.00	100.00	100.00	98.00	99	20	1626
SSHD.Default	0	0	3	90.48	-1	0.00	100.00	100.00	100.00	99	20	1499
SSHD.Shared	0	0	0	0.00	-1	0.00	100.00	100.00	100.00	100	10	0
SSHD.Jaqui	0	1	1	9.52	-1	0.00	100.00	100.00	100.00	99	20	127
Monit	0	1	3	1.01	-1	0.00	60.00	60.00	60.00	96	81	7854
Monit.Default	0	0	3	100.00	-1	0.00	100.00	100.00	100.00	96	81	7854
Monit.Shared	0	0	0	0.00	-1	0.00	100.00	100.00	100.00	100	40	0
Backups	1	1	0	0.00	-1	0.00	100.00	100.00	100.00	100	4096	0
Backups.Default	0	0	0	-	-1	0.00	100.00	100.00	100.00	100	4096	0
Backups.Shared	0	0	0	-	-1	0.00	100.00	100.00	100.00	100	4096	0

It shows that Backups is in tier 1

Remember:

MEM Physical memory utilization of the class (%).

WLMSTAT -S AND -S

wlmstat -S (superclass only)

CLASS	CPU	MEM	DKIO
Unclassified	0.00	24.69	0.00
Unmanaged	0.00	25.74	0.00
Default	0.09	11.71	0.00
Shared	0.00	7.69	0.00
System	3.85	48.58	0.00
SSHD	0.09	0.21	0.00
Monit	0.05	1.01	0.00
Backups	0.00	0.00	0.00
TOTAL	4.08	93.89	0.00

wlmstat -s (subclass only)

CLASS	CPU	MEM	DKIO
SSHD.Default	90.91	90.48	-
SSHD.Shared	0.00	0.00	-
SSHD.Jaqui	9.09	9.52	-
TOTAL	0.11	0.21	0.00

WLMSTAT I/O

wlmstat -b (show only I/O)

CLASS	DKIO
Unclassified	0.00
Unmanaged	0.00
Default	0.00
Shared	0.00
System	0.00
SSHD	0.00
SSHD.Default	-
SSHD.Shared	-
SSHD.Jaqui	-
Monit	0.00
Backups	0.00
TOTAL	0.00

WLMSTAT I/O

wlmstat -B "" (show only I/O but break it down by all disks)

CLASS	DKIO	hdisk1	hdisk3	hdisk2	hdisk0
Unclassified	0.00	0.00	0.00	0.00	0.00
Unmanaged	0.00	0.00	0.00	0.00	0.00
Default	0.00	0.00	0.00	0.00	0.00
Shared	0.00	0.00	0.00	0.00	0.00
System	0.00	0.00	0.00	0.00	0.00
SSHD	0.00	0.00	0.00	0.00	0.00
SSHD.Default	-	0.00	0.00	0.00	0.00
SSHD.Shared	-	0.00	0.00	0.00	0.00
SSHD.Jaqui	-	0.00	0.00	0.00	0.00
Monit	0.00	0.00	0.00	0.00	0.00
Backups	0.00	0.00	0.00	0.00	0.00
TOTAL	0.00	0.00	0.00	0.00	0.00

WLMSTAT CLASSES

```
# wlmstat -l Jaqui  
Class Jaqui not found
```

```
# wlmstat -l SSHD.Jaqui  
CLASS      CPU      MEM  DKIO  
SSHD.Jaqui  0.00    9.52  -
```

```
# wlmstat -l SSHD  
CLASS      CPU      MEM  DKIO  
SSHD       0.08    0.21  0.00  
SSHD.Default 100.00  90.48  -  
SSHD.Shared  0.00    0.00  -  
SSHD.Jaqui  0.00    9.52  -
```

WLMSTAT -W AND -U

wlmstat -w (high watermarks)

CLASS	CPU	MEM	DKIO
Unclassified	0.00	194293	0.00
Unmanaged	0.00	269989	0.00
Default	0.00	91400	0.00
Shared	0.00	59856	0.00
System	0.00	380972	0.00
SSHD	0.00	1626	0.00
SSHD.Default	0.00	1499	0.00
SSHD.Shared	0.00	0	0.00
SSHD.Jaqui	0.00	127	0.00
Monit	0.00	8837	0.00
Backups	0.00	0	0.00

wlmstat -u (per tier and total unused resources)

TIER	CPU	MEM	DKIO
0	99.45	0.00	100.00
1	99.45	0.00	100.00

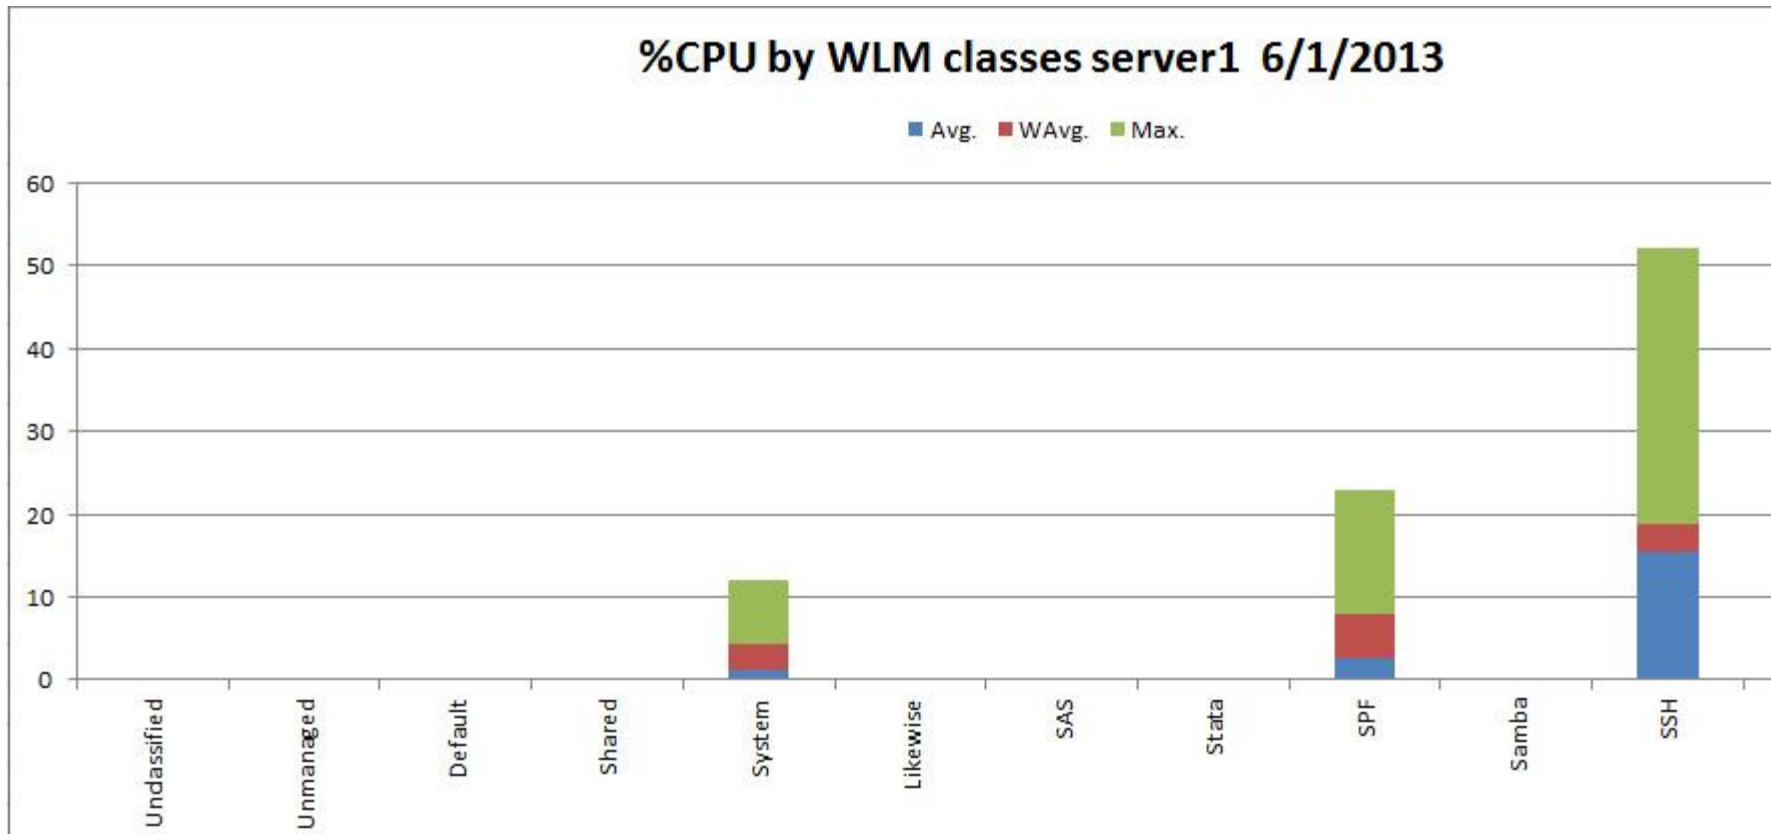
NMON COMMANDS

- To take a 30 minute nmon snap
`/usr/bin/nmon -ft -A -M -L -^ -W -s 150 -c 576`
- Will now get additional fields in the nmon analyzer report

In BBBP it lists the status of WLM, classes, rules, limits and shares settings


It adds four tabs: WLMBIO, WLMCPU, WLMPCPU and WLMMEM

NMON ANALYZER WLMCPU



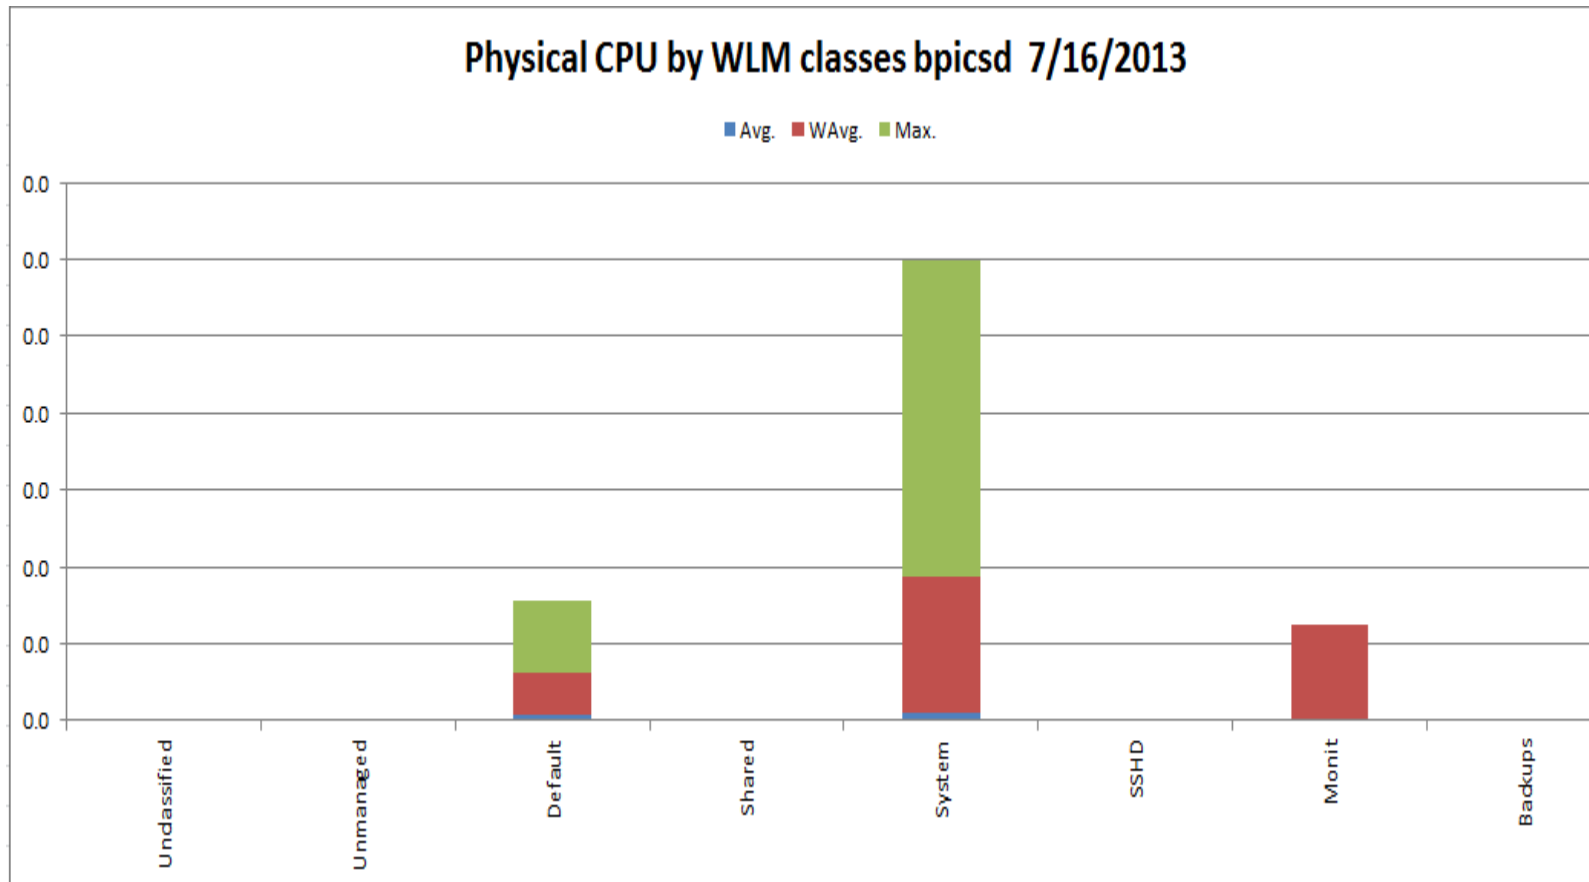
NMON ANALYZER WLMCPU

%CPU by WLM classes




NMON ANALYZER WLMPCPU

Physical CPU by WLM classes

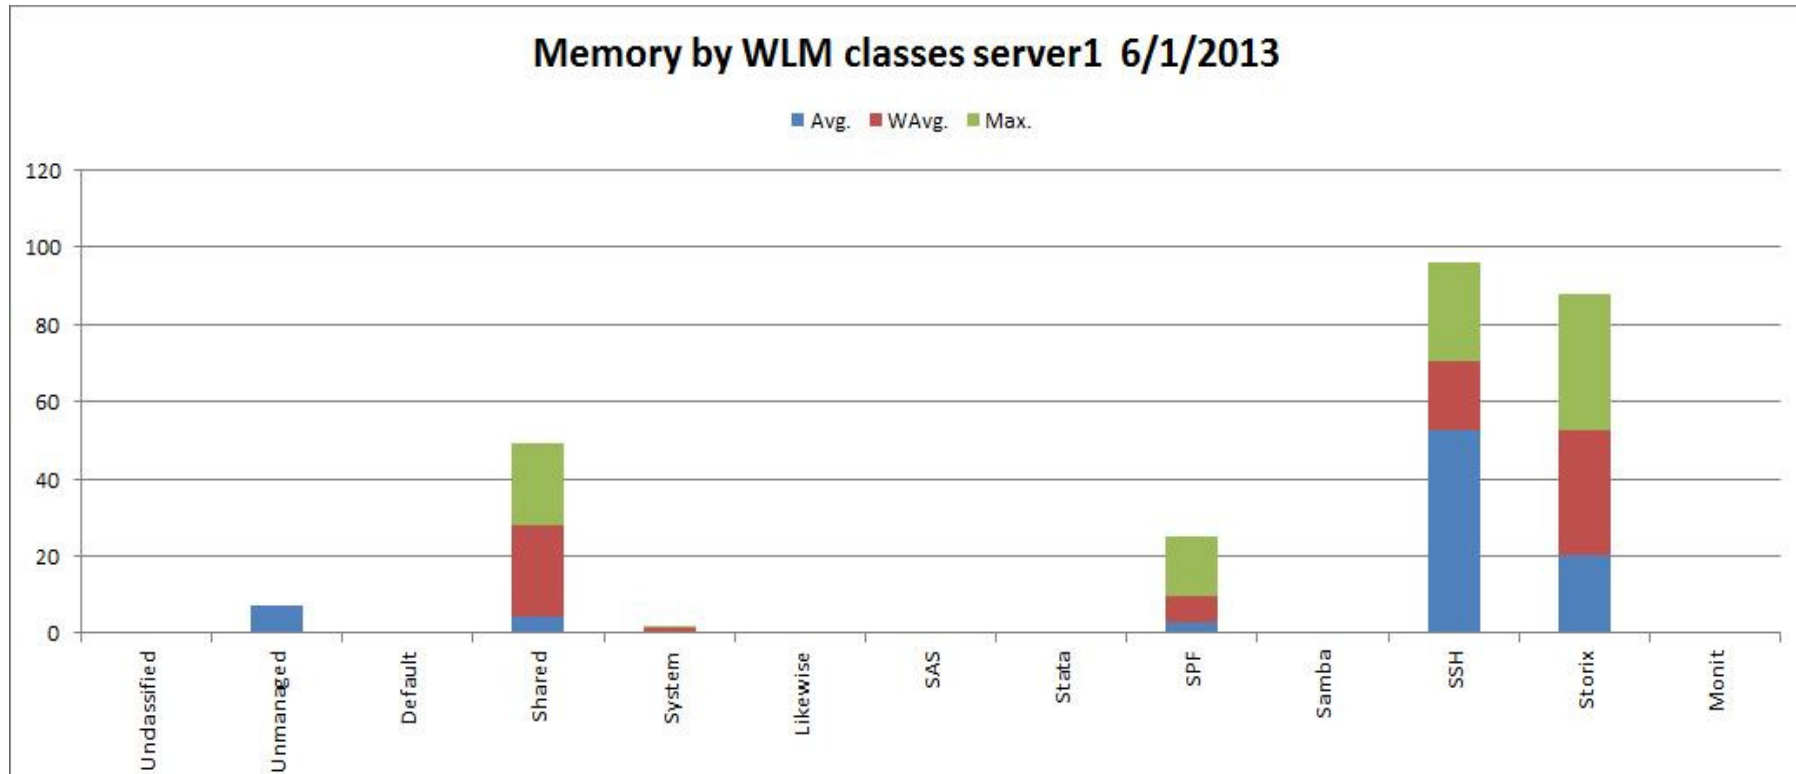


NMON ANALYZER WLMMEM

Memory by WLM classes




NMON ANALYZER WLMOMEM




NMON ANALYZER WLMMEM

Memory by WLM classes server1 6/1/2013



NMON ANALYZER WLMBIO



EXAMPLE OF ACCTCOM (REQUIRES ACCOUNTING TO BE ON)

```
# acctcom -w | head -10
```

COMMAND			START	END	REAL	CPU	MEAN	
NAME	USER	CLASS	TTYNAME	TIME	TIME	(SECS)	(SECS)	SIZE(K)
#acctcom	root	System.Default	pts/9	15:25:55	15:25:55	0.00	0.00	364.00
#bash	root	System.Default	pts/9	15:25:55	15:25:55	0.03	0.00	0.00
#rm	root	System.Default	pts/9	15:25:56	15:25:56	0.02	0.00	0.00
#rm	root	System.Default	pts/9	15:25:56	15:25:56	0.02	0.00	332.00
#rm	root	System.Default	pts/9	15:25:56	15:25:56	0.02	0.00	0.00
#bash	root	System.Default	pts/9	15:25:55	15:25:55	0.11	0.02	225.00
#bash	root	System.Default	pts/9	15:25:55	15:25:55	0.23	0.02	0.00

NOTE ON REPORTING

AIX and WLM gather statistics a little differently
May result in differences in out put from WLM and sar
sar reports raw values whereas WLM uses devayed averages

AIX performance monitoring tools

Use a sampling approach.
Every clock tick (10 ms), AIX charges one CPU tick to the current process.
It determines the mode of execution (user/system)
And increments the user or system time accordingly

WLM

Uses microsecond timing
Records precisely how long a thread has been active.
Records user/system breakouts
And accumulates the times on a per class basis only
It does not save individual thread CPU usage information

SNAP

If using snap to gather info ensure you add the `-w` flag to get WLM related info
kdb also has WLM options that can be used

REFERENCES

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

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

AIX 5L WLM redbook – 2001

Out of date but provides some ideas

<http://www.redbooks.ibm.com/redbooks/pdfs/sg245977.pdf>

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>

Article on WLM, DLPAR and WPAR

<http://www.rootvg.net/content/view/148/>

THANK YOU FOR YOUR TIME



If you have questions please email me at:
lynchj@forsythe.com

Presentation can be found at:
<http://www.circle4.com/papers/wlm-jul2013.pdf>