

Performance Tuning Lab

Jaqui Lynch
Mainline Information Systems
jaqui.lynch@mainline.com
<http://www.circle4.com/jaqui/papers/perflab.pdf>
Session 6827

Agenda

- Lecture for a bit
- Try out commands

vmstat (AIX)

IGNORE FIRST LINE - average since boot
Run vmstat over an interval (i.e. vmstat 2 30)

Kthr		memory		Page				Faults				CPU				
r	b	avm	fre	re	pi	po	fr	sr	cy	in	sy	cs	us	sy	id	wa
2	3	253292	5	0	1	262	1783	2949	0	941	13281	1877	12	29	0	60
2	3	255503	0	0	0	260	1110	1879	0	911	14725	1049	22	25	0	53
1	2	273017	487	0	0	0	0	0	0	507	12627	422	12	20	65	3

fre should be between minfree and maxfree
fr:sr ratio 1783:2949 means that for every 1783 pages freed 2949 pages had to be examined.

To get a 60 second average try: vmstat 60 2

rc.tune 1/2

```
#
# rc.tune is called by inittab to tune the network and system parameters
#Works on AIX 5.2
#
# To make no changes permanent use the no -opp option
/usr/sbin/no -rp -o rfc1323=1
/usr/sbin/no -rp -o sb_max=1310720
/usr/sbin/no -rp -o subnetsarelocal=1
/usr/sbin/no -rp -o tcp_sendspace=262144
/usr/sbin/no -rp -o tcp_recvspace=262144
/usr/sbin/no -rp -o udp_sendspace=65536
/usr/sbin/no -rp -o udp_recvspace=655360
/usr/sbin/no -rp -o tcp_mssdflt=1448
/usr/sbin/no -rp -o extendednetstats=1
#
/usr/sbin/no -rp -o thewall=131072
/usr/sbin/no -rp -o ipqmaxlen=256
```

rc.tune 2/2

```
#  
# If this works well for you try setting maxperm to 30 instead of 50  
#  
# /usr/samples/kernel/vmtune -p 10 -P 50 -W 32  
# vmtune is being phased out and replaced by ioo and vmo  
#  
/usr/sbin/vmo -p -o minperm%=10  
/usr/sbin/vmo -p -o maxclient%=50  
/usr/sbin/vmo -p -o maxperm%=50  
/usr/sbin/ioo -p -o maxrandwrt=32  
/usr/sbin/nfso -p -o nfs_rfc1323=1  
#  
# If this works well for you try setting maxperm to 30 instead of 50
```

ioo Output

```
minpgahead = 2  
maxpgahead = 8  
pd_npages = 65536  
maxrandwrt = 0  
numclust = 1  
numfsbufs = 186  
sync_release_ilock = 0  
lvm_bufcnt = 9  
j2_minPageReadAhead = 2  
j2_maxPageReadAhead = 8  
j2_nBufferPerPagerDevice = 512  
j2_nPagesPerWriteBehindCluster = 32  
j2_maxRandomWrite = 0  
j2_nRandomCluster = 0  
hd_pvs_opn = 1  
hd_pbuf_cnt = 320
```

vmo Output

```
memory_frames = 65536
maxfree = 128
minfree = 120
minperm% = 20
minperm = 11813
maxperm% = 80
maxperm = 47255
strict_maxperm = 0
maxpin% = 80
maxpin = 52429
maxclient% = 80

lrubucket = 131072
defps = 1
nokilluid = 0
numpsblks = 131072
npskill = 1024
npwarn = 4096
v_pinshm = 0
pta_balance_threshold = 50
pagecoloring = 0
framesets = 0
mempools = 0
lgpg_size = 0
lgpg_regions = 0
num_spec_dataseg = n/a
spec_dataseg_int = n/a
memory_affinity = n/a
```

/usr/samples/kernel/vmtune

AIX Specific

In AIX this is part of the bos.adt.samples fileset

Run vmtune and look at the following:

```
minperm          default is 20%
maxperm          default is 80%
numperm          This is what percent of real memory is currently being used for
                  caching files - if it is high reduce maxperm to 30 to 50%
```

Reducing maxperm stops file caching affecting programs that are running

Also maxrandwrt - default is 0 - try 32

vmtune output

vmtune: current values:

-p	-P	-r	-R	-f	-F	-N	-W	
minperm	maxperm	minpgahead	maxpgahead	minfree	maxfree	pd_npages	maxrandwrt	
13104	65524	2	8	120	128	524288	32	
-M	-w	-k	-c	-b	-B	-u	-l	-d
maxpin	npswarn	npskill	numclust	numfsbufs	hd_pbuf_cnt	lvm_bufcnt	lrubucket	defps
104839	4096	1024	1	93	208	9	131072	1
-s	-n	-S	-h					
sync_release_ilock	nokilluid	v_pinshm	strict_maxperm					
0	0	0	0					

number of valid memory pages = 131048
maximum pinable=80.0% of real memory
number of file memory pages = 100060

maxperm=50.0% of real memory
minperm=10.0% of real memory
numperm=76.4% of real memory

lsps -a (similar to pstat)

- Ensure all page datasets the same size although hd6 can be bigger - ensure more page space than memory
- Only includes pages allocated (default)
- Use lsps -s to get all pages (includes reserved via early allocation (PSALLOC=early))
- Use multiple page datasets on multiple disks

lsps output

```
lsps -a
Page Space Physical Volume Volume Group Size %Used Active Auto Type
paging05   hdisk9         pagvg01    2072MB  1  yes  yes  lv
paging04   hdisk5         vgpaging01 504MB   1  yes  yes  lv
paging02   hdisk4         vgpaging02 168MB   1  yes  yes  lv
paging01   hdisk3         vgpagine03 168MB   1  yes  yes  lv
paging00   hdisk2         vgpaging04 168MB   1  yes  yes  lv
hd6        hdisk0         rootvg     512MB   1  yes  yes  lv
```

```
lsps -s
Total Paging Space Percent Used
3592MB             1%
```

Bad Layout above
Should be balanced
Make hd6 the biggest by one lp or the same size as the others

iostat (AIX)

Run iostat ver an interval (i.e. iostat 2 30)

tty:	tin	tout	avg cpu:	%user	%sys	%idle	%iowait
(kp)	17.5	2.9		0.0	75.4	2.9	23.4
(te)	0	0		9	8	24	59

Disks:	%tm_act	Kbps	tps	Kb_read	Kb_wrtn
hdisk1	3.5	32	3.5	0	64
hdisk0	59	27.6	59.5	44	508
hdisk2	12	64	16	0	128
hdisk3	3.5	32	3.5	0	64

%user + %sys <= 80%

Similar info from ps au or sar -u

%iowait – time waiting on disk I/O – includes local and nfs

Historical disk I/O collection is off by default on AIX 5.1

netstat

- netstat -m
 - Requires extendednetstats=1 (no)
- netstat -rn
- netstat -v

Other Commands

- sar
 - sar -A -o sar.data int count
 - sar -u -f sar.data
 - -a file access
 - -u cpu
 - -b buffers
 - -d devices
 - -k kernel
 - man sar
- Schedtune
 - /usr/samples/kernel
 - Requires bos.adt.samples
- nmon

NMON

nmon v8b [H for help] Hostname=serv1 Refresh=2.0secs 15:57.02

Use these keys to toggle statistics on/off:

c = CPU l = Long-term CPU - = Faster screen updates
m = Memory k = Kernel Stats + = Slower screen updates
d = Disks a = Adapters (disk only) v = Verbose hints
r = RS6000/pSeries n = Network U = command arguments
e = ESS Disks g = Disk Groups (see -g) . = only busy disks/procs
j = JFS t = Top-proc's (1,2,3,4,5 - different data)

h =more options

Summary

