Security & Performance Tuning

Performance

- “Performance” is a meaningless concept in isolation
- It only makes sense to talk about performance of a particular workload, and according to a particular set of metrics
- The first step is to characterize the workload you care about, and what aspects of its operation are most important to you

“Help! My system is slow!”, meetBSD, 2008, Kris Kennaway <kris@FreeBSD.org>

Factors that affect Performance

- Four major resources
  - CPU Time
  - Memory
  - Hard disk I/O bandwidth
  - Network I/O bandwidth
- Where is the real bottleneck
  - Not CPU, hard disk bandwidth it is !!
  - When memory is not enough, system will do swap, so memory and disk bandwidth are the major suspects

System Performance Checkup – Analyzing CPU usage (1)

- Three information of CPU
  - Overall utilization
    - Help to identify whether the CPU resource is the system bottleneck
  - Load average
  - Per-process consumption
    - Identify specific process's CPU utilization

System Performance Checkup – Analyzing CPU usage (2)

- vmstat command
  - report virtual memory statistics
    - us: user time
      - High us means high computation
    - sy: system time
      - High sy means process are making lots of system call or performing I/O
    - id: cpu idle
  - us and sy time should half-half
  - Monitoring interval should not be too small

```
lucky7:~ -lwhsu- vmstat -c 2 -w 1
 procs      memory      page                    disks     faults         cpu
 r b w     avm    fre   flt  re  pi  po    fr  sr ad0 ad2   in   sy   cs us sy id
 0 0 0    579M   259M  2009  29 103   1  2532  33   0   0  385 2158  791  3  3 95
 0 0 0    579M   259M  6304   0   0   0  6176   0   0   0   85 5745 1190  2  2 96
```

System Performance Checkup – Analyzing CPU usage (3)

- faults (average per second over last 5 seconds)
  - in: device interrupt per interval
  - sy: system calls per interval
  - cs: cpu context switch rate

Nothing to do Server

```
csws1:~ -lwhsu- vmstat -c 2 w 1
 procs      memory      page                    faults      cpu
 r b w     avm    fre   flt  re  pi  po    fr  sr   in   sy   cs us sy id
 0 0 0  385504 109144    29   0   0   0    29   1   31  451  859  0  0 99
 0 0 0  385504 109144     2   0   0   0     0   0   16  272  585  0  1 99
```

High load, busy server

```
lucky7:~ -lwhsu- vmstat -c 5 -w 1
 procs      memory      page                    disks     faults         cpu
 r b w     avm    fre   flt  re  pi  po    fr  sr ad0 ad2   in   sy   cs us sy id
 0 0 0    603M   245M  2010  29 103   1  2533  33   0   0  385 2160  791  3  3 95
 0 0 0    603M   245M  4991   0   0   0  4956   0   0   0   31 4419  797  2  2 96
 0 0 0    603M   245M  4989   0   0   0  4956   0   0   0   27 4425  812  2  2 96
 0 0 0    603M   245M  4989   0   0   0  4956   0   5   0   29 4372  825  2  2 96
 0 0 0    603M   245M  4973   0   1   0  4956   0   1   0   22 4373  792  2  1 97
```

System Performance Checkup – Analyzing CPU usage (4)

- Load average
  - The average number of runnable processes
    - Including processes waiting for disk or network I/O
- uptime command
  - Show how long system has been running and the load average of the system over the last 1, 5, and 15 minutes

```
cshome:~ -lwhsu- uptime
 5:24PM  up 88 days,  5:09, 5 users, load averages: 0.00, 0.00, 0.00
```

System Performance Checkup – Analyzing CPU usage (5)

- top command
  - Display and update information about the top cpu processes
- ps command
  - Show process status

System Performance Checkup – Analyzing memory usage (1)

- When memory is not enough …
  - Memory page has to be "swapped out" to the disk block
  - LRU (Least Recently Used) algorithm
  - Bad situation – "desperation swapping"
    - Kernel forcibly swaps out runnable process
    - Extreme memory shortage
- Two numbers that quantify memory activity
  - Total amount of active virtual memory
    - Tell you the total demand for memory
  - Page rate
    - suggest the proportion of actively used memory

System Performance Checkup – Analyzing memory usage (2)

- To see amount of swap space in use
  - pstat -s or swapinfo (FreeBSD)
  - swapon -s (Linux)
  - swap -l (Solaris)

```
bsd2:~ -lwhsu- pstat -s
Device          1K-blocks     Used    Avail Capacity
/dev/ad4s1b       1048576        0  1048576     0%
/dev/ad8s1b       1048576        0  1048576     0%
Total             2097152        0  2097152     0%
```

System Performance Checkup – Analyzing memory usage (3)

- vmstat command
  - procs
    - r: in run queue
    - b: blocked for resource
    - w: runnable or short sleeper but swapped
  - memory
    - avm: active virtual pages
    - fre: size of the free list
  - page (averaged each five seconds, given in units per second)
    - flt: total number of page faults
    - pi: pages paged in
    - po: pages paged out
      - 50 page-out cause about 1 seconds latency
    - fr: pages freed per second

```
lucky7:~ -lwhsu- vmstat -c 3 -w 5
 procs      memory      page                    disks     faults         cpu
 r b w     avm    fre   flt  re  pi  po    fr  sr ad0 ad2   in   sy   cs us sy id
 0 0 0    628M   233M  2011  28 103   1  2534  33   0   0  385 2161  791  3  3 95
 0 0 0    628M   233M  4101   0   0   0  4075   0   0   0   29 5879  768  2  1 97
 0 0 0    628M   233M  5095   0   0   0  5066   0   0   0   30 6704  787  2  1 97
```

systat

- display system statistics on a crt

```
/0   /1   /2   /3   /4   /5   /6   /7   /8   /9   /10
     Load Average   ||

      Interface           Traffic               Peak                Total
            sk0  in      99.399 KB/s          99.399 KB/s          17.963 GB
                 out      3.647 MB/s           3.647 MB/s         526.633 GB
```

systat -ifstat

systat -vmstat

```
 7 users    Load  0.18  0.07  0.02                  Jan  6 13:48

Mem:KB    REAL            VIRTUAL                       VN PAGER  SWAP PAGER
        Tot   Share      Tot    Share    Free           in  out     in  out
Act   94652   10240   285808    35956   63392 count
All 1007780   14884115816224    56360         pages
                                                       Interrupts
Proc:r  p  d  s  w    Csw  Trp  Sys  Int  Sof  Flt        cow    1280 total
     1     2 76         77  128 2560 1357   26  102  201900 wire   1152 0: clk
                                                 100188 act          1: atkb
99.9%Sys   0.0%Intr  0.0%User  0.0%Nice  0.0%Idl   659228 inact   128 8: rtc
|    |    |    |    |    |    |    |    |          48524 cache        10: em0
==================================================  14868 free         11: fxp
                                                          daefr        14: ata
Namei         Name-cache    Dir-cache                     prcfr        15: ata
    Calls     hits    %     hits    %                     react
        4        4  100                                   pdwake
                                          77 zfod         pdpgs
Disks   ad0   ad1   ad2   ad3            77 ozfod         intrn
KB/t   0.00  0.00  0.00  0.00              %slo-z   114304 buf
tps       0     0     0     0              tfree       33 dirtybuf
MB/s   0.00  0.00  0.00  0.00                       70236 desiredvnodes
% busy    0     0     0     0                      54930 numvnodes
                                                   17559 freevnodes
```

System Performance Checkup – Analyzing disk I/O

- iostat command
  - Report I/O statistics
  - Usage: iostat –w 1 –c 5
    - tin/tout: characters read from /write to terminal
    - KB/t: kilobytes per transfer
    - tps: transfers per second
    - MB/s: megabytes per second

```
FreeBSD:~ -lwhsu- iostat da0 -w 1
      tty             da0              cpu
 tin tout  KB/t tps  MB/s  us ni sy in id
   0  258 59.78 253 14.77   3  0  4  0 94
   0  127 63.13 501 30.89   3  0  4  0 93
   0   43 62.58 346 21.14   5  0  5  0 90
   0   42 62.40 289 17.63   3  0  5  0 92
   0   43 61.19 720 43.02   1  0  2  0 97
```

*stat commands

```
lucky7:/bin -lwhsu- ls -al {,/usr}{/bin,/sbin}/*stat
-r-xr-xr-x  1 root  wheel  -  49976 Jan  2 18:52 /sbin/ipfstat*
-r-xr-xr-x  1 root  wheel  -   7264 Jan  2 18:52 /sbin/kldstat*
-r-xr-sr-x  1 root  kmem   -  11872 Jan  2 18:53 /usr/bin/btsockstat*
-r-xr-sr-x  1 root  kmem   -  20432 Jan  2 18:53 /usr/bin/fstat*
-r-xr-sr-x  1 root  kmem   - 144208 Jan  2 18:53 /usr/bin/netstat*
-r-xr-xr-x  1 root  wheel  -  12352 Jan  2 18:53 /usr/bin/nfsstat*
-r-xr-xr-x  1 root  wheel  -  16912 Jan  2 18:53 /usr/bin/procstat*
-r-xr-xr-x  1 root  wheel  -  15696 Jan  2 18:53 /usr/bin/sockstat*
-r-xr-xr-x  2 root  wheel  -  15560 Jan  2 18:53 /usr/bin/stat*
-r-xr-xr-x  1 root  wheel  -  82424 Jan  2 18:53 /usr/bin/systat*
-r-xr-xr-x  1 root  wheel  -  25552 Jan  2 18:53 /usr/bin/vmstat*
-r-xr-xr-x  1 root  wheel  -  15760 Jan  2 18:53 /usr/sbin/gstat*
lrwxr-xr-x  1 root  wheel  -     21 Jan  2 18:53 /usr/sbin/hoststat@ ->
                                                        /usr/sbin/mailwrapper
-r-xr-x---  1 root  wheel  -  11504 Jan  2 18:53 /usr/sbin/ifmcstat*
-r-xr-xr-x  1 root  wheel  -  19808 Jan  2 18:53 /usr/sbin/iostat*
-r-xr-xr-x  1 root  wheel  -  39376 Jan  2 18:53 /usr/sbin/pmcstat*
-r-xr-xr-x  2 root  wheel  -  13040 Jan  2 18:53 /usr/sbin/pstat*
lrwxr-xr-x  1 root  wheel  -     21 Jan  2 18:53 /usr/sbin/purgestat@ ->
                                                        /usr/sbin/mailwrapper
-r-xr-xr-x  1 root  wheel  -  10048 Jan  2 18:53 /usr/sbin/slstat*
```

top

top –m cpu (default)

```
last pid: 61540;  load averages:  0.30,  0.31,  0.32                up 17+09:57:18  13:57:14
242 processes: 1 running, 241 sleeping
CPU states:     % user,     % nice,     % system,     % interrupt,     % idle
Mem: 2195M Active, 7466M Inact, 1574M Wired, 21M Cache, 214M Buf, 619M Free
Swap: 2048M Total, 140K Used, 2048M Free

  PID USERNAME      THR PRI NICE   SIZE    RES STATE  C   TIME   WCPU COMMAND
26091 squid          17  44    0   414M   384M ucond  1  35:51  0.00% squid
11945 bind           11  44    0 71696K 59544K select 1  32:06  0.00% named
11375 root            1  58    0 20960K  3144K select 1   9:35  0.00% sshd
68517 nobody          1  44    0 24472K 14716K select 3   8:00  0.00% rsync
```

top –m io

```
last pid:  9347;  load averages:  0.21,  0.29,  0.32                 up 17+09:58:20  13:58:16
243 processes: 1 running, 242 sleeping
CPU states:  0.5% user,  0.0% nice,  1.2% system,  0.0% interrupt, 98.3% idle
Mem: 2200M Active, 7484M Inact, 1604M Wired, 25M Cache, 214M Buf, 562M Free
Swap: 2048M Total, 140K Used, 2048M Free

  PID USERNAME       VCSW  IVCSW   READ  WRITE  FAULT  TOTAL PERCENT COMMAND
18107 cvsup             0      0      0      0      0      0   0.00% cvsupd
26091 squid            34      0      0      0      0      0   0.00% squid
11945 bind              9      3      0      0      0      0   0.00% named
11375 root              4      0      0      0      0      0   0.00% sshd
```

gstat

```
L(q)  ops/s    r/s   kBps   ms/r    w/s   kBps   ms/w   %busy Name
    0      0      0      0    0.0      0      0    0.0    0.0| acd0
    5    218    218  15756    9.3      0      0    0.0   94.0| da0
    0    111      2    214    5.0    107    933    4.3   23.4| ad4
    0    113      0      0    0.0    111    933    4.3   24.1| ad5
    0    111      2    214    5.0    107    933    4.3   23.5| ad4s1
    0    113      0      0    0.0    111    933    4.3   24.1| ad5s1
    0      0      0      0    0.0      0      0    0.0    0.0| ad6
    0      5      0      0    0.0      5     40    0.6    0.3| ad4s1a
    0      0      0      0    0.0      0      0    0.0    0.0| ad4s1b
    0      0      0      0    0.0      0      0    0.0    0.0| ad4s1c
    0    106      2    214    5.0    102    893    4.7   23.4| ad4s1d
    0      0      0      0    0.0      0      0    0.0    0.0| ad7
    0      5      0      0    0.0      5     40    0.3    0.1| ad5s1a
    0      0      0      0    0.0      0      0    0.0    0.0| ad5s1b
    0      0      0      0    0.0      0      0    0.0    0.0| ad5s1c
    0    108      0      0    0.0    106    893    4.7   24.1| ad5s1d
    0      4      0      0    0.0      4     40    0.8    0.3| mirror/gm0s1a
```

Sysctls

- security.bsd.see_other_uids
  - Unprivileged processes may see subjects/objects with different real uid
- kern.randompid
  - Random PID modulus
- net.inet.ip.random_id
  - Assign random ip_id values
- net.inet.tcp.blackhole
  - Do not send RST on segments to closed ports
- net.inet.udp.blackhole
  - Do not send port unreachables for refused connects

Periodic Jobs

- Check system states
  - daily_status_zfs_enable="YES"
  - daily_status_gmirror_enable="YES"
  - daily_status_ntpd_enable="YES“
  - weekly_noid_enable="YES“
  - ...etc.

Manual Pages

- tuning(7)
  - performance tuning under FreeBSD
- security(7)
  - introduction to security under FreeBSD
- sprog(7)
  - secure programming practices
- sdoc(7)
  - guide to adding security considerations sections to manual pages