

Performance Analysis

wangth (2017-2020, CC BY-SA)
? (1996-2016)

交大資工系資訊中心

Computer Center of Department of Computer Science, NCTU

Help! My system is slow!

- https://people.freebsd.org/~kris/scaling/Help_my_system_is_slow.pdf



Help! My system is slow!

Profiling tools, tips and tricks

Kris Kennaway
kris@FreeBSD.org

What you can do to improve performance

- Memory size has a major influence on performance
- Correct the problems of usage
- Load balance appliance
- Organize the system's hard disks and filesystems
- Monitoring your networks
- ...

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 kernel statistics about process, memory, cpu, ..
 - Usage: % `vmstat -c 2 -w 1`
 - **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

```
$ vmstat -c 2 -w 5
procs  memory      page          disks      faults      cpu
r b w  avm  fre  flt  re  pi  po  fr  sr  da0  cd0  in  sy  cs  us  sy  id
0 0 0 2.0T 5.7G  97  0  0  0 106  9  0  0  39 398 457  0  0 100
0 0 0 2.0T 5.7G   0  0  0  0   0 20 21  0  40 156 458  0  0 100
```



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

```
$ vmstat -c 2 -w 5
procs memory          page          disks        faults          cpu
r b w avm   fre      flt  re  pi  po  fr  sr  da0 da1  in  sy  cs   us sy id
3 2 0 50364 1587316  3    0  0  0  3  0  0  0  931 786 181  0  0 100
0 2 0 50368 1587312  5    0  0  0  0  0  0  0  250  91  23  0  0  99
```

- High load, busy http server

```
$ vmstat -c 5 -w 5
procs memory          page          disks        faults          cpu
r b w avm   fre      flt  re  pi  po  fr  sr  da0  in  sy  cs   us sy id
0 0 0 231320 68792   320  4  0  0 264  7  0 2273 3381 952 16  4 80
0 0 0 232984 67100   558  0  0  0 386  0  1 1958 3285 551 11  5 84
1 0 0 228252 69272   192  2  0  0 292  0  5 2787 2626 681 23  4 73
1 0 0 221564 72048   102  0  0  0 229  0  0 1395  556 184  1  2 97
0 0 0 209624 76684    96  0  0  0 306  0  0 1350  935 279  0  2 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
 - Usage: % uptime

```
$ uptime  
2:52AM up 31 days, 6:48, 1 user, load averages: 0.42, 0.28, 0.22
```


System Performance Checkup – Analyzing CPU usage (5)

- top command
 - Display and update information about the top cpu processes
- ps command
 - Show process status
- renice command
 - `renice -n increment -p pid`
 - `renice +1 987 -u daemon root -p 32`

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 -k` (FreeBSD)
 - `swapon -s` (Linux)
 - `swap -l` (Solaris)
- `pstat` command
 - `pstat -s`

```
$ pstat -s
```

Device	1K-blocks	Used	Avail	Capacity
/dev/label/swap-0	1048572	0	1048572	0%
/dev/label/swap-1	1048572	0	1048572	0%
Total	2097144	0	2097144	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

```
$ vmstat -c 3 -w 5
```

procs			memory		page				disks			
r	b	w	avm	fre	flt	re	pi	po	fr	sr	da0	da1
0	3	0	1427M	1196M	224	0	0	0	312	0	0	0
0	3	0	1427M	1196M	3	0	0	0	169	0	12	12
0	3	0	1427M	1196M	3	0	0	0	110	0	15	15

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

```
$ 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

System Performance Checkup – Analyzing network

- The four most common uses of netstat
 - Monitoring the status of network connections
 - netstat -a
 - Inspecting interface configuration information
 - netstat -i

```
$ netstat -i
```

Name	Mtu	Network	Address	Ipkts	Ierrs	Opkts	Oerrs	Coll
bge0	1500	140.113.240.0	derek	2256736153	-	3709378394	-	-
bge0	1500	192.168.7.0	192.168.7.1	1744582	-	49144622	-	-
lo0	16384	your-net	localhost	433424	-	433424	-	-

- Examining the routing table
 - netstat -r -n
- Viewing operational statistics for network protocols

systat

- Display system statistics
 - `systat -ifstat`

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

Interface      Traffic      Peak      Total
  lo0  in      0.000 KB/s    0.000 KB/s    97.325 KB
       out     0.000 KB/s    0.000 KB/s    97.325 KB

  vmx1  in      0.000 KB/s    0.000 KB/s    6.335 GB
       out     0.000 KB/s    0.000 KB/s    2.263 GB

  vmx0  in      0.000 KB/s    0.000 KB/s    2.169 GB
       out     0.000 KB/s    0.000 KB/s   289.635 MB
```

systat

○ systat -vmstat

```

1 users      Load 0.20  0.29  0.28                Dec 20 03:13
Mem usage:   63%Phy 60%Kmem                VN PAGER   SWAP PAGER
Mem:         REAL          VIRTUAL          in  out     in  out
            Tot  Share    Tot  Share    Free  count
Act 209516K 16036K 2097752M 22356K 5821M pages
All 221072K 27492K 2097781M 52088K
Proc:
  r  p  d  s  w  Csw  Trp  Sys  Int  Sof  Flt  6033
                41      445  5K  4K                6K
25.0%Sys  0.0%Intr  0.0%User  0.0%Nice 75.0%Idle
|  |  |  |  |  |  |  |  |  |  |  |
=====
                                dtbuf
Namei      Name-cache  Dir-cache  349995 desvn
  Calls    hits  %    hits  %    349982 numvn
    18      18 100      323816 frevn
                                9644M wire
Disks      da0    cd0 pass0 pass1  52764K act
KB/t      0.00  0.00  0.00  0.00      383M  inact
tps        0      0      0      0          0  laund
MB/s      0.00  0.00  0.00  0.00      5821M free
%busy     0      0      0      0      55808 buf

```



*stat commands

```
$ 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
```

Wait Channels: <https://wiki.freebsd.org/WaitChannels>



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