Chapter 25
Performance Analysis

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**

```
tytsai@u3:/var/log> 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
```

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

```
tytsai@u3:/var/log> 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

```
tytsai@ccbsd3:~> vmstat –c 5 –w 5
 procs      memory          page                    disk   faults      cpu
 r b w     avm    fre      flt   re  pi  po  fr     sr  ad0   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**

```
{tytsai@mgate2}~> uptime
8:22AM  up 6 days, 22:13, 2 users, load averages: 0.06, 0.02, 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**

See Chapter4 pp. 18 ~ 23

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**

```
tytsai@ccduty:~> pstat -s
Device          1K-blocks   Used    Avail   Capacity  Type
/dev/rad0s1b    511608      56      511552  0%        Interleaved
/dev/rad6s1b    505244      68      505176  0%        Interleaved
Total           1016852     124     1016728 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

```
tytsai@ccduty:/var/run> vmstat –c 3 –w 5
procs    memory      page                          disks
r b w    avm    fre      flt      re  pi  po  fr     sr  ad0 ad4
1 0 0   57316  25988   181    0   0   0   165   3  0    0
0 0 0   57316  25988   4      0   0   0   0     0  0    0
0 0 0   57316  25988   3      0   0   0   0     0  0    0
```

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

```
tytsai@u214:~> iostat  -w 5
     tty          da0              da1              cpu
 tin tout   KB/t   tps  MB/s    KB/t  tps  MB/s  us ni sy in id
   0    3   0.00    0  0.00    0.00    0  0.00   0  0  0  0 100
   0   71  63.88   83  5.17    0.00    0  0.00   0  0 10  3 87
   0   63  64.00   83  5.19    0.00    0  0.00   1  0 11  3 85
   0   63  64.00   84  5.25    0.00    0  0.00   0  0 11  3 86
   0   63  63.88   82  5.12    0.00    0  0.00   0  0 10  3 87
```