

# Tuning RHEL

David Wood  
Consultant, Red Hat  
9-13-14

# Tuning

- Why Tune
- Why Not Tune
- Overall RHEL Tuning
- How to test performance
- Disk IO Tuning
- CPU Tuning
- Network Tuning
- Memory Tuning

# Why Tune

- Increase system usage
- Consolidate Workload
- Get your money's worth
- Reduce quantity of servers
- Real Time (rt-kernel)

# Why Not Tune

- Fixing hardware issues
- Correcting for misconfigured networks
- Because I can
- To make it faster

# *Alchemy's First Law of Equivalent Exchange*

*Humankind cannot gain  
anything without first  
giving something in return.*

*To obtain, something of  
equal value must be lost.*

*-Fullmetal Alchemist*

*[www.inspirationboost.com](http://www.inspirationboost.com)*

# Overall RHEL Tuning

- CFQ – Fair share scheduler
- NIC/CPU connection
- Tickless Kernel
- Changing CPU speeds
- Transparent Huge Pages

# How to Test Performance

- CPU
  - kcbench
  - Calculate PI
- Disk IO
  - iostat
  - iotop
- Network
  - Netperf
- Memory
  - free

# Tuning Disk IO

- CFQ
- Deadline
- Noop
- Anticipatory

	ATIME enabled	ATIME disabled
CFQ	10,000 to 116,000 KB/sec	22,000 to 121,000 KB/sec
Deadline	70,000 to 208,000 KB/sec	70,000 to 208,000 KB/sec
Noop	70,000 to 206,000 KB/sec	70,000 to 208,000 KB/sec

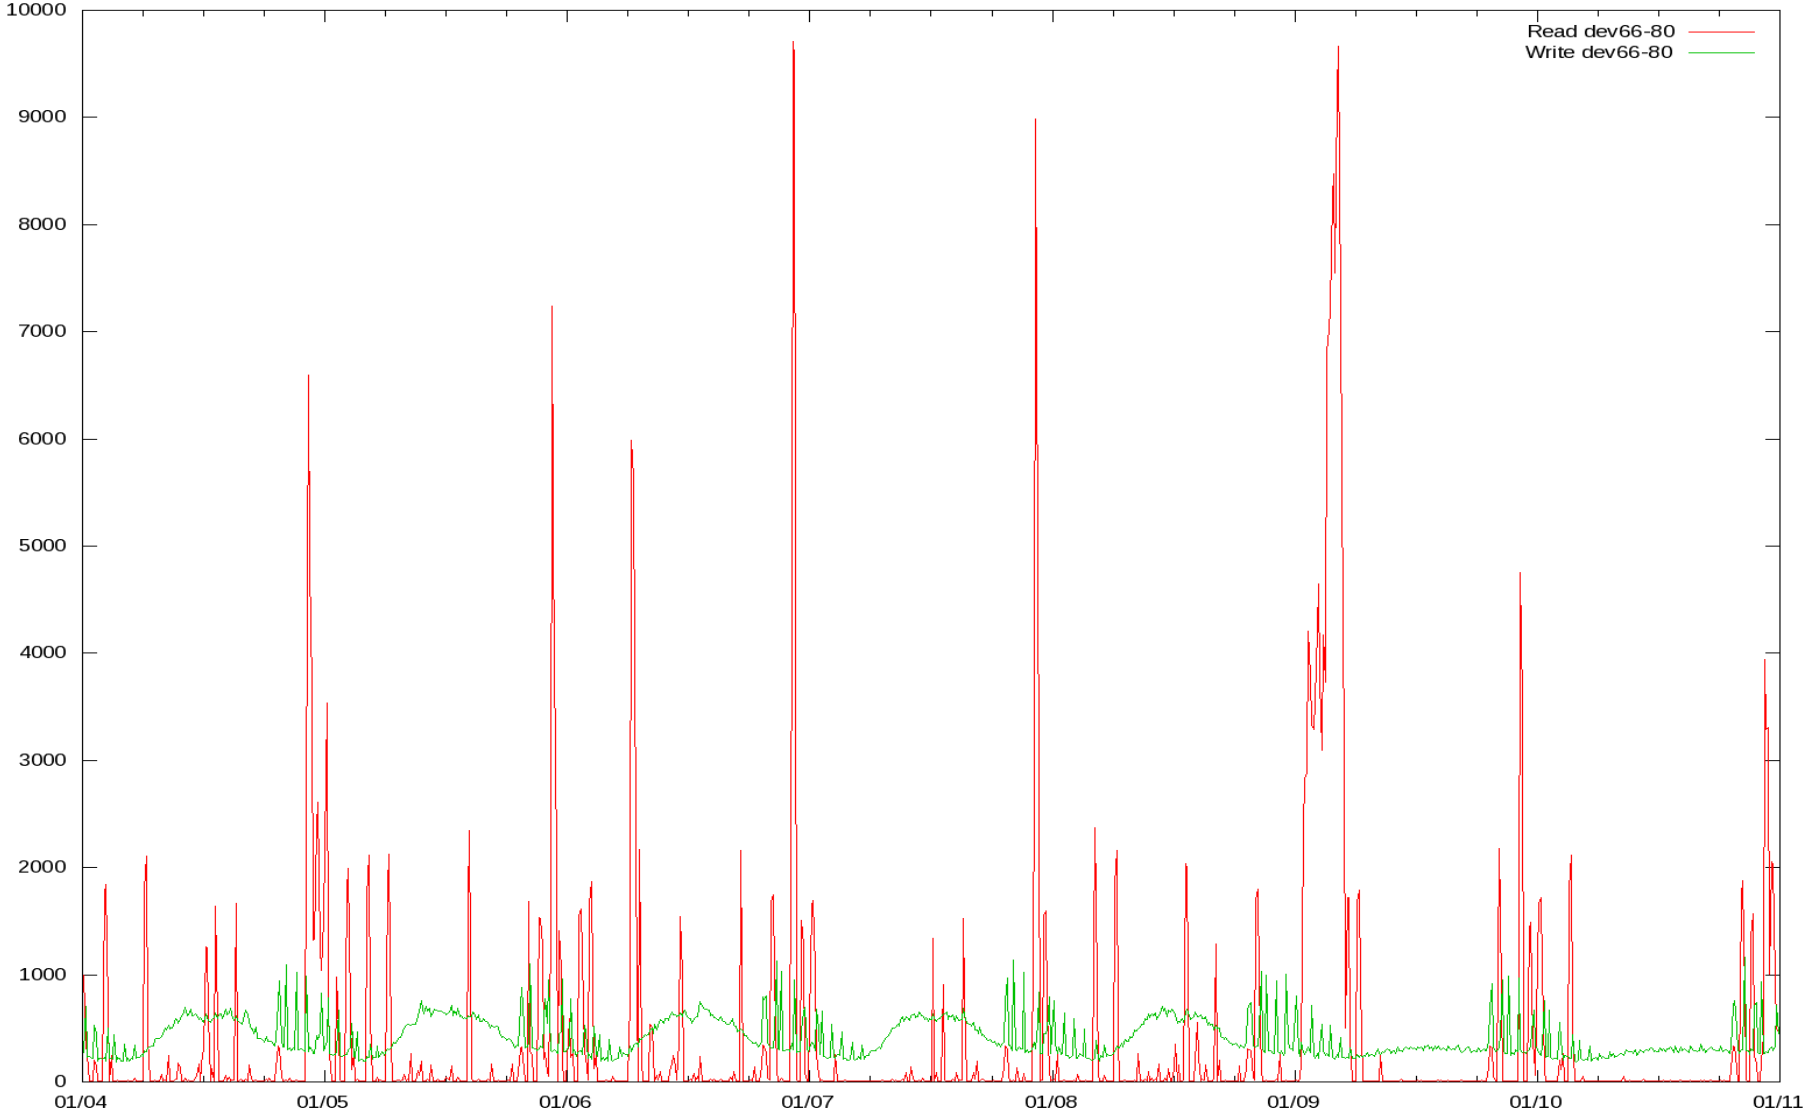


# Tuning Disk IO

- CFQ
- Deadline
- Noop
- Anticipatory
- ATIME, Whats the Deal

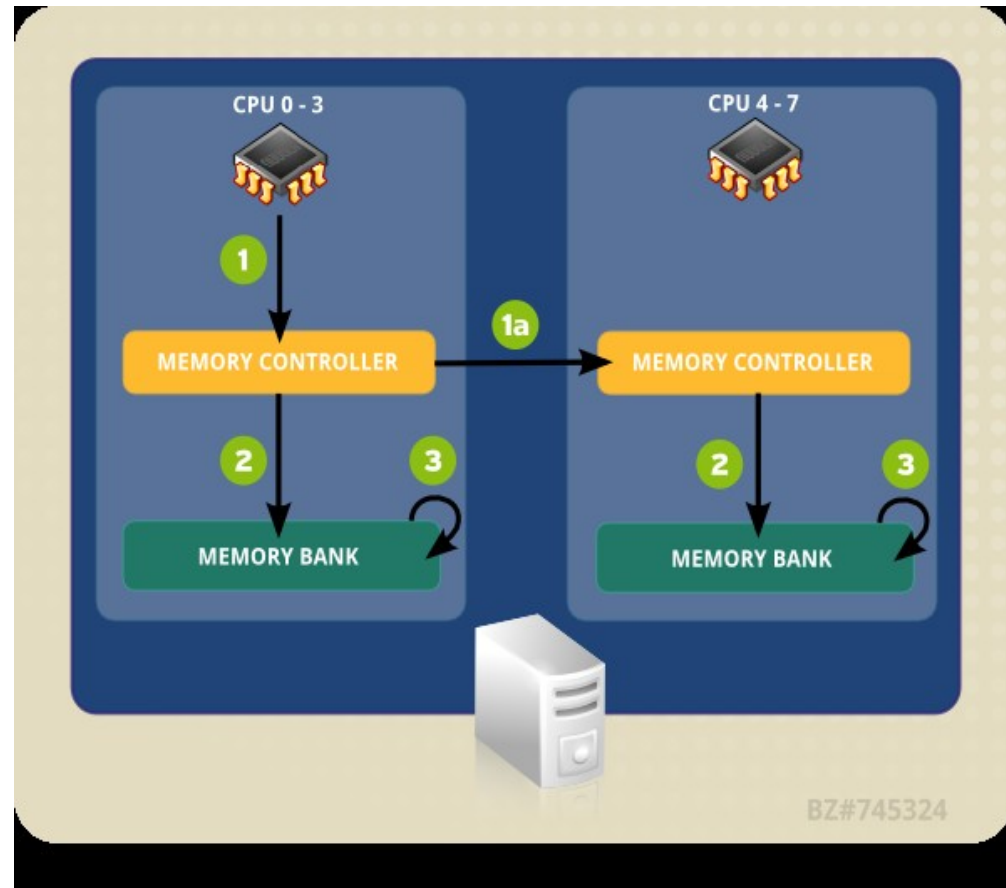
	ATIME enabled	ATIME disabled
CFQ	10,000 to 116,000 KB/sec	22,000 to 121,000 KB/sec
Deadline	70,000 to 208,000 KB/sec	70,000 to 208,000 KB/sec
Noop	70,000 to 206,000 KB/sec	70,000 to 208,000 KB/sec

# Tuning Disk IO

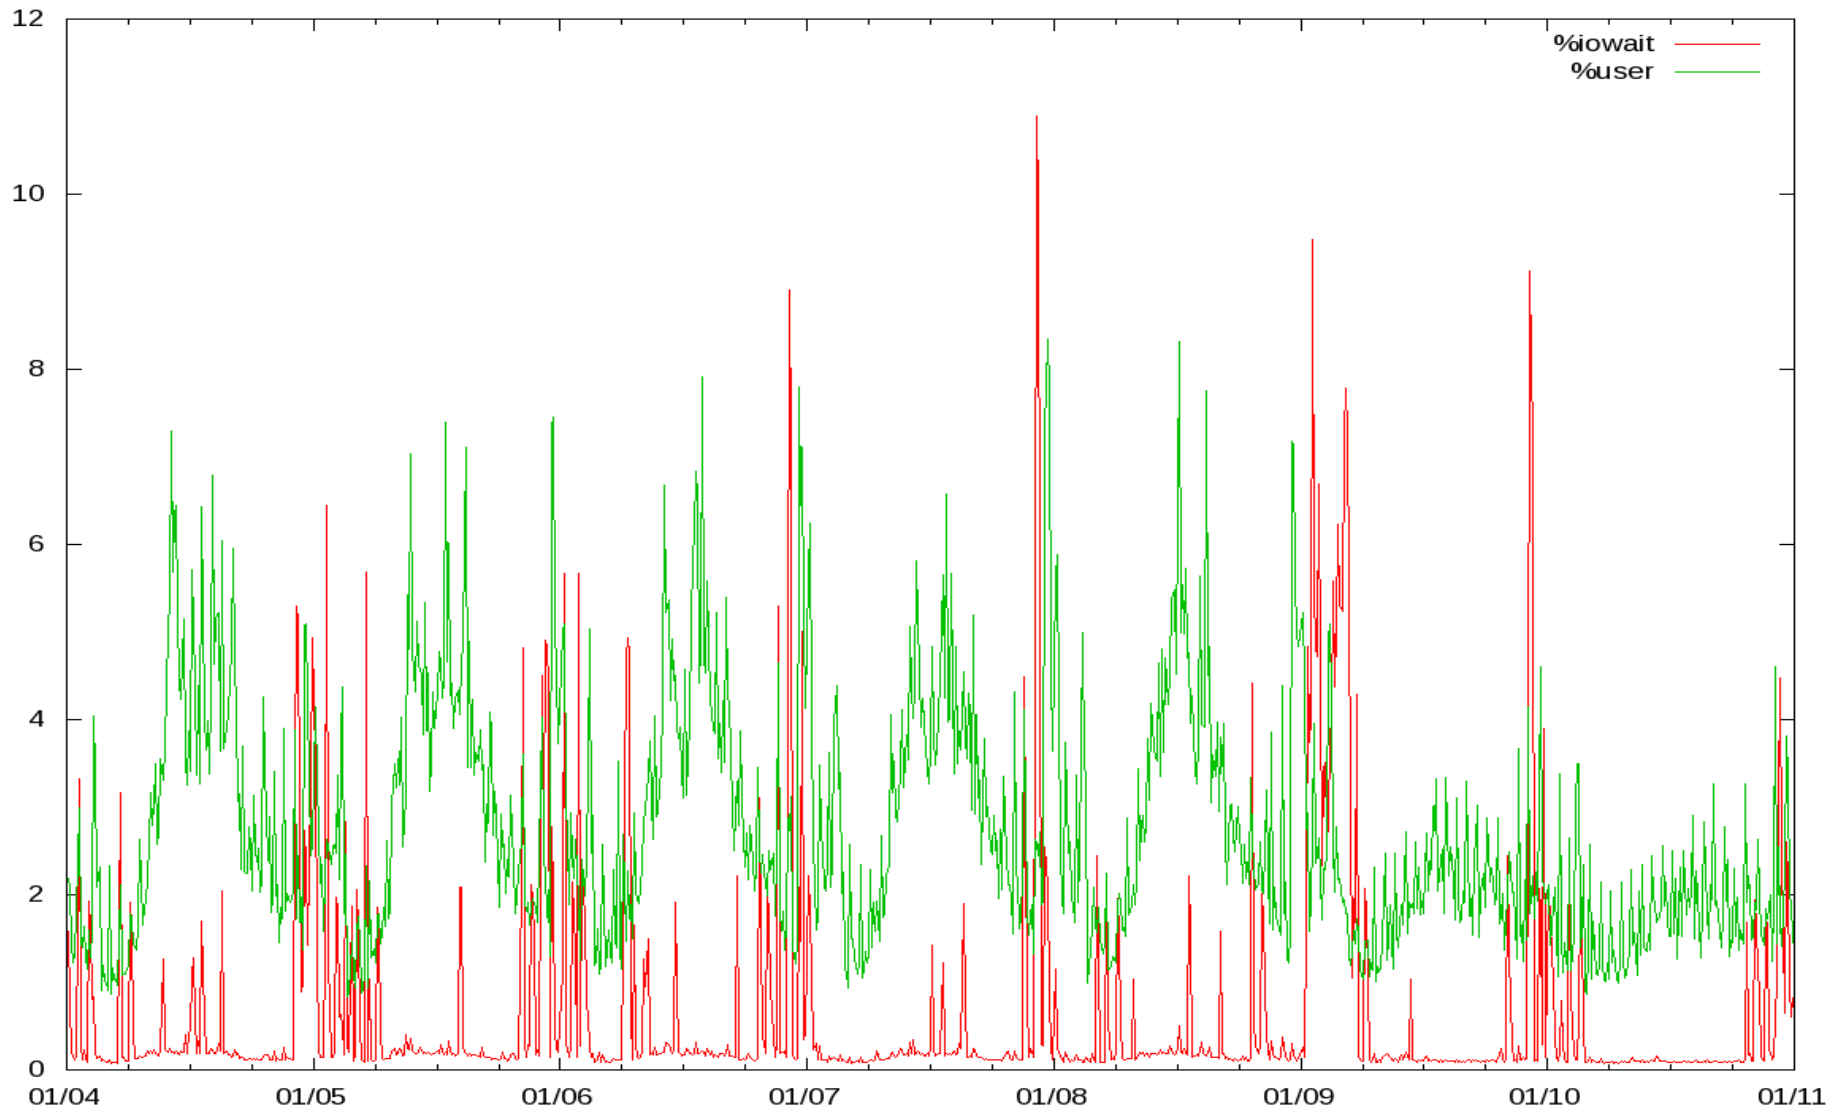


# Tuning the CPU

- Ondemand
- Userspace
- Performance
- NUMA
- CGroups



# Tuning the CPU

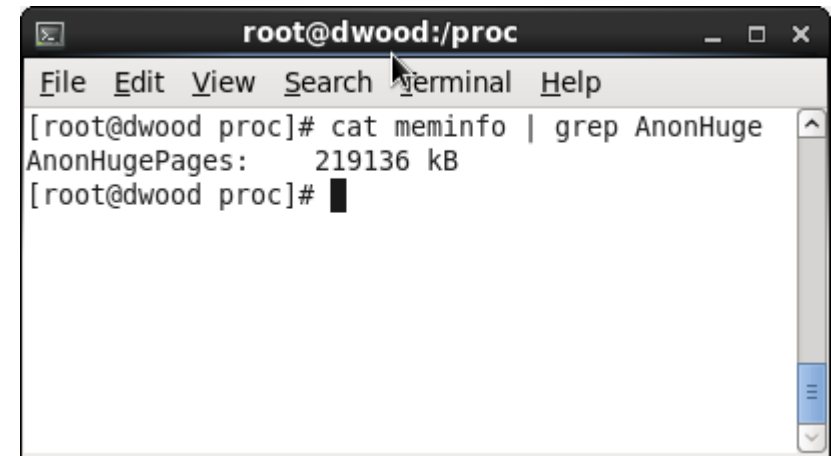


# Tuning the Network

- Auto-negotiate Duplex and Speed
- Default Buffer Size
- Max Buffer Size
- Changing keep alive
- TCP Port range limits
- TCP segmentation offload (TSO)
- Jumbo frames

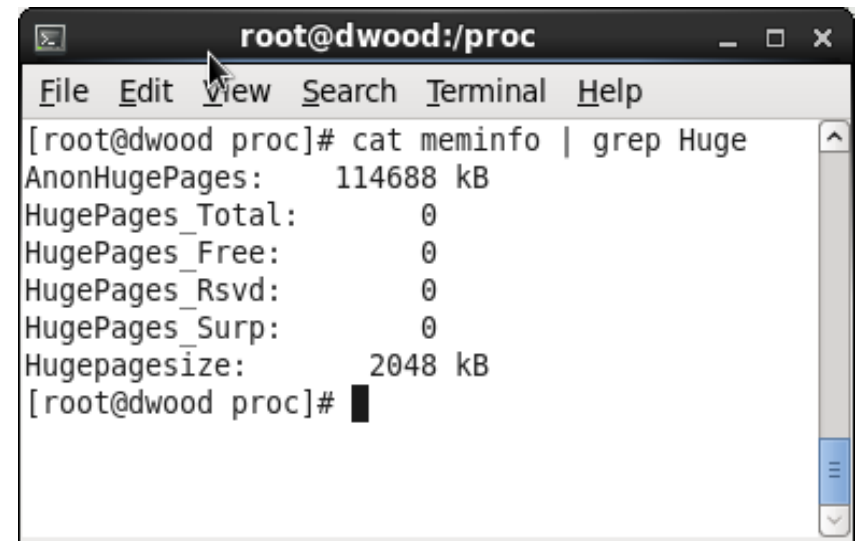
# Tuning Memory

- Transparent Huge Pages
  - `grep AnonHuge /proc/meminfo`



```
root@dwood:/proc
File Edit View Search Terminal Help
[root@dwood proc]# cat meminfo | grep AnonHuge
AnonHugePages:    219136 kB
[root@dwood proc]#
```

- Huge Pages
  - `Grep Huge /proc/meminfo`



```
root@dwood:/proc
File Edit View Search Terminal Help
[root@dwood proc]# cat meminfo | grep Huge
AnonHugePages:    114688 kB
HugePages_Total:    0
HugePages_Free:    0
HugePages_Rsvd:    0
HugePages_Surp:    0
Hugepagesize:      2048 kB
[root@dwood proc]#
```

# Tuning Memory – Page Faults

- Sar -B

```
root@dwood:/proc
File Edit View Search Terminal Help
05:10:01 PM      0.92      6.95     157.13      0.02     633.17      0.00      0.00      0.00      0.00
05:20:02 PM      0.00     10.81     192.61      0.00     702.53      0.00      0.00      0.00      0.00
05:30:01 PM      0.16      6.25     104.03      0.00     605.61      0.00      0.00      0.00      0.00
05:40:01 PM      0.00      6.28     161.66      0.00     639.23      0.00      0.00      0.00      0.00
05:50:01 PM     74.01     31.48     739.42      0.80    1344.74      0.00      0.00      0.00      0.00
06:00:01 PM      2.20     32.72     140.02      0.01     393.66      0.00      0.00      0.00      0.00
06:10:01 PM      0.00      3.54     102.56      0.00     607.81      0.00      0.00      0.00      0.00
06:20:01 PM      5.45      5.91     109.19      0.05     486.63      0.00      0.00      0.00      0.00
06:30:01 PM      5.73      5.46     160.05      0.05     505.37      0.00      0.00      0.00      0.00
06:40:01 PM      0.00      3.72     115.09      0.00     195.47      0.00      0.00      0.00      0.00
06:50:01 PM     72.17     27.58     319.11      0.62     389.17      0.00      0.00      0.00      0.00
07:00:01 PM      1.19     71.45     385.60      0.02     780.38      0.00      0.00      0.00      0.00
07:10:01 PM      0.00     37.38     119.08      0.00     645.73      0.00      0.00      0.00      0.00
07:20:01 PM      0.07     37.56      92.28      0.00     629.93      0.00      0.00      0.00      0.00
07:30:02 PM      0.00     41.29      97.63      0.00     635.97      0.00      0.00      0.00      0.00
07:40:01 PM      0.00     37.74     121.30      0.00     615.08      0.00      0.00      0.00      0.00
07:50:01 PM      0.45     52.39     255.25      0.01     555.50      0.00      0.00      0.00      0.00

07:50:01 PM  pgpgin/s pgpgout/s  fault/s  majflt/s  pgfree/s  pgscank/s  pgscand/s  pgsteal/s  %vmeff
08:00:02 PM      0.00     11.78     101.60      0.00     125.33      0.00      0.00      0.00      0.00
08:10:01 PM      0.01      9.15     116.24      0.00     120.62      0.00      0.00      0.00      0.00
08:20:01 PM     45.09     28.23     330.63      0.25     281.86      0.00      0.00      0.00      0.00
Average:      23.49     629.64     499.31      0.04     945.74     83.29      1.10     83.25     98.65
[root@dwood proc]#
```

# Tuning Memory – Page Faults

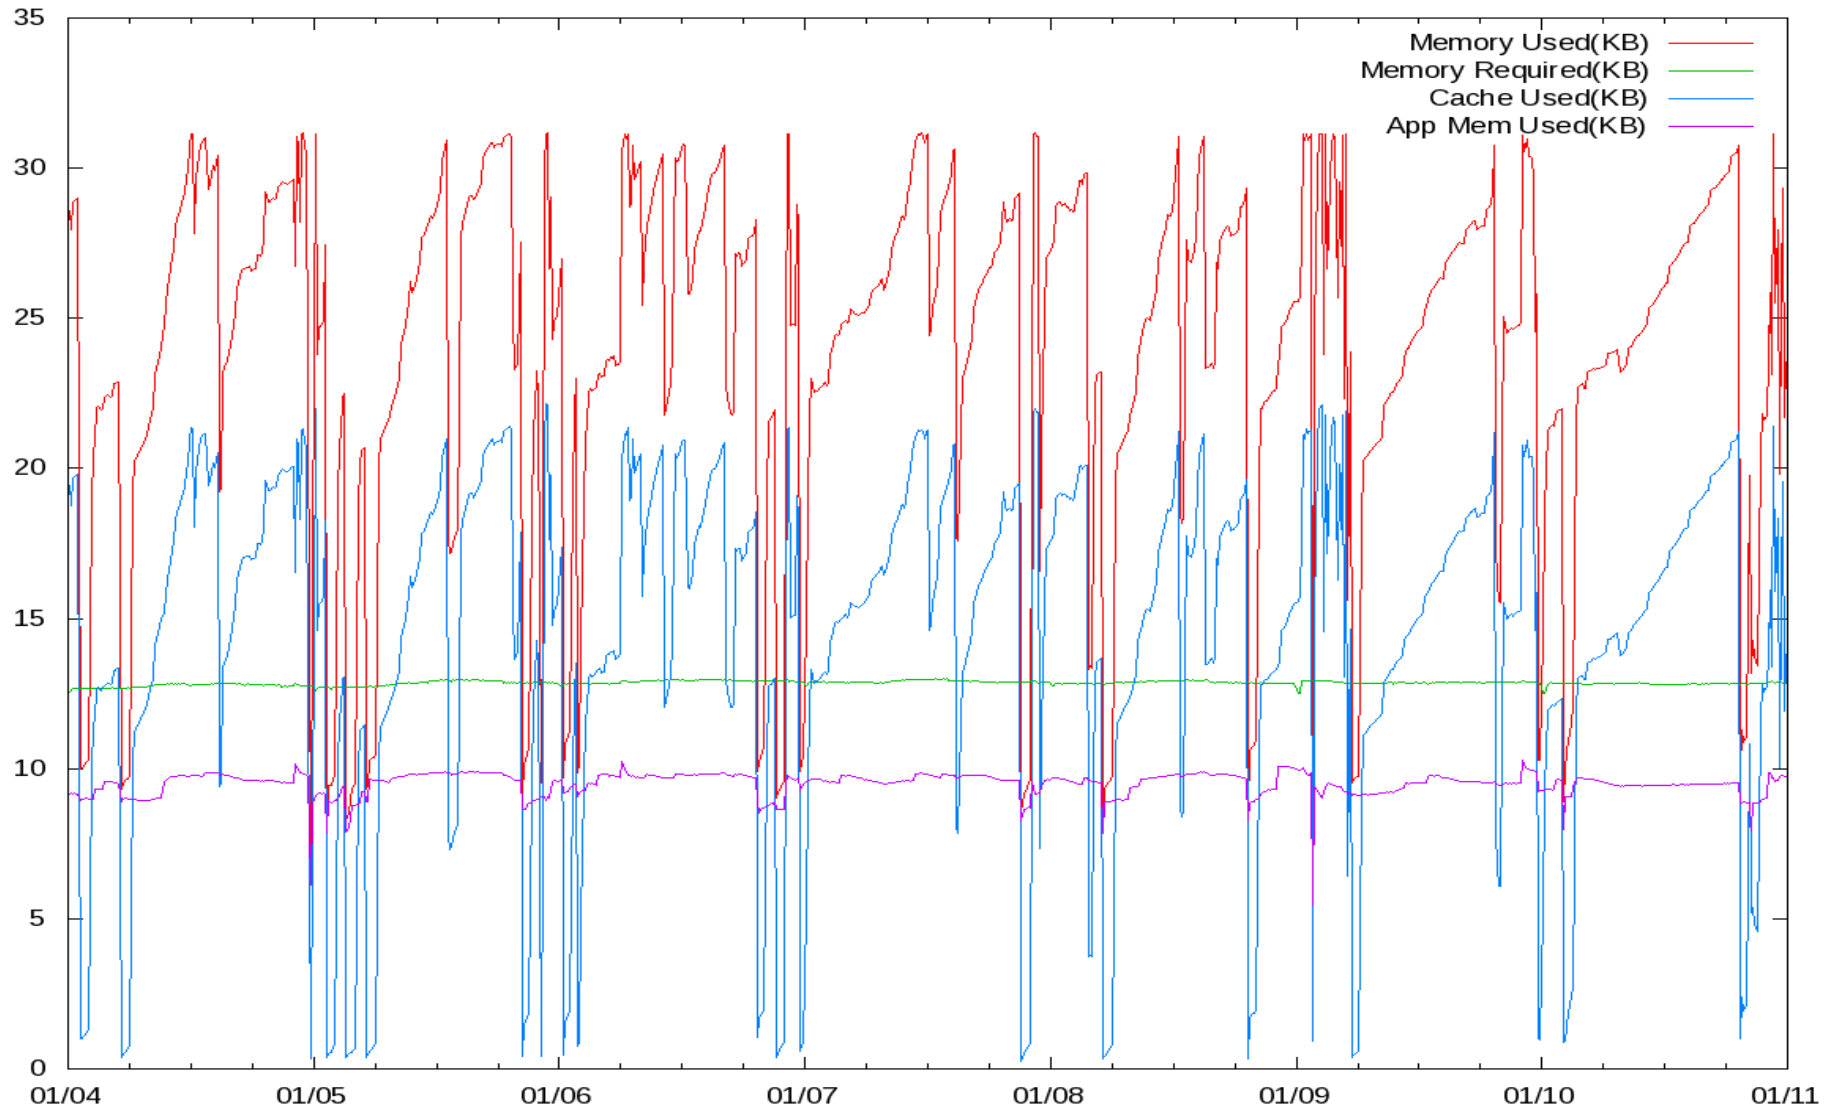
- top

```
root@dwood:/proc
File Edit View Search Terminal Help
top - 20:21:40 up 1 day, 22:27, 6 users, load average: 0.05, 0.04, 0.07
Tasks: 311 total, 2 running, 309 sleeping, 0 stopped, 0 zombie
Cpu(s): 4.5%us, 2.2%sy, 0.0%ni, 91.9%id, 1.4%wa, 0.0%hi, 0.0%si, 0.0%st
Mem: 16216924k total, 2285664k used, 13931260k free, 229968k buffers
Swap: 4194296k total, 0k used, 4194296k free, 1071344k cached

  PID USER      PR  NI  VIRT  RES  SHR  S  %CPU  %MEM    TIME+  nFLT  COMMAND
 23081 dwood    20   0 1139m 310m  46m  S   8.6   2.0   2:56.44  342  firefox
 21206 root      20   0  151m  31m  20m  S   6.0   0.2   3:32.37   18  Xorg
 24512 dwood    20   0 1080m 160m  73m  S   2.6   1.0   0:14.61  148  soffice.bin
 21444 dwood    20   0  340m  17m  12m  S   2.0   0.1   0:02.29   34  gnome-panel
 21176 dwood    20   0  450m 9732 8172  S   1.3   0.1   8:13.45   11  pulseaudio
 23177 dwood    20   0  455m  23m  13m  S   1.3   0.1   1:30.52    0  plugin-containe
 21443 dwood    20   0  433m  13m  10m  S   1.0   0.1   0:01.45    2  metacity
 22431 dwood    20   0  294m  13m 9688  S   1.0   0.1   0:12.24    5  gnome-terminal
   508 root      20   0    0    0    0  S   0.7   0.0   0:50.62    0  kcryptd
 24675 root      20   0 15164 1416  956  R   0.7   0.0   0:00.33    0  top
   19 root      20   0    0    0    0  S   0.3   0.0   0:20.87    0  events/0
  1126 root      20   0    0    0    0  S   0.3   0.0   0:13.25    0  iwlmwifi
  1932 dbus      20   0 32776 2376  992  S   0.3   0.0   0:36.70    0  dbus-daemon
19953 root      20   0    0    0    0  S   0.3   0.0   0:00.41    0  kblockd/2
19979 root      20   0    0    0    0  S   0.3   0.0   0:00.80    0  events/3
 21425 dwood    20   0  409m  10m 8020  S   0.3   0.1   0:01.19    6  gnome-settings-
 21474 dwood    20   0  328m  12m 9556  S   0.3   0.1   0:01.14    1  wnck-applet
```



# Tuning Memory



# Tuning for Virtual Machines

- Disk
  - CFQ Deadline Noop
- CPU
  - Dynamic Performance Power-save
- Network
  - Buffer-sizing TSO Duplex/Speed
- Memory
  - Transparent Huge Pages (THP)
  - Huge Pages

# Tuning for Virtual Machines

- Disk
  - Deadline/Noop scheduler
- CPU
  - Performance
- Network
  - Depends
- Memory
  - Transparent Huge Pages (THP) depends
  - Huge Pages depends

# Conclusion

- Know why your tuning
- Know the Success Goal
- Test
- Test
- Test

# Questions?