Table des matières

CPU

- Lancer un **sar 1 10** pour avoir une vue globale de la machine :

```
root@server1104092:~# sar 1 10
Linux 2.6.9-67.ELsmp (server1104092)     11/12/08

14:57:54          CPU     %user     %nice   %system   %iowait     %idle
14:57:55          all     11.60      0.00      4.49     25.94     57.98
14:57:56          all     10.88      0.00      3.75     25.00     60.38
14:57:57          all     10.59      0.00      3.86     24.41     61.15
14:57:58          all     10.76      0.00      4.13     24.66     60.45
14:57:59          all     11.36      0.00      4.99     24.22     59.43
14:58:00          all     11.51      0.00      3.25     27.91     57.32
```

- On peut utiliser **top** :

```
top - 15:07:42 up 10 days,  7:32, 20 users,  load average: 1.87, 2.72, 2.58
Tasks: 448 total,   1 running, 447 sleeping,   0 stopped,   0 zombie
Cpu(s):  7.6% us,  0.9% sy,  0.0% ni, 87.3% id,  3.1% wa,  0.1% hi,  0.9% si
Mem:  65747848k total, 65217884k used,   529964k free,   168860k buffers
Swap: 65535992k total,      976k used, 65535016k free, 62379852k cached

  PID USER      PR  NI %CPU    TIME+  %MEM  VIRT  RES  SHR S COMMAND
32368 root      16   0   54  37:38.27  0.0 20456 5028 3604 S dsmc
17209 sybase    16   0    5 252:52.99  7.3 7327m 4.6g 4.6g S RPO_dataserver
17211 sybase    16   0    4 305:02.35  7.5 7327m 4.7g 4.7g S RPO_dataserver
26930 root      16   0    1  20:23.75  0.0  2900 1828  520 S cmaperfd
14326 sybase    15   0    1 394:40.93  9.3 7327m 5.8g 5.8g S IRD_dataserver
18525 sybase    15   0    1 232:11.65  5.4 7331m 3.4g 3.4g S RPS_dataserver
30714 root      16   0    1   0:00.07  0.0  6424 1428  848 R top
14323 sybase    15   0    0 419:17.35  9.3 7327m 5.8g 5.8g S IRD_dataserver
15498 sybase    15   0    0 522:40.79  9.1 7578m 5.7g 5.7g S IRS_dataserver
16354 sybase    15   0    0 291:18.70 11.0 7337m 6.9g 6.9g S RPO_dataserver
17207 sybase    15   0    0 346:40.57  7.5 7327m 4.7g 4.7g S RPO_dataserver
18003 sybase    15   0    0 266:00.77 11.0 7341m 6.9g 6.9g S RPS_dataserver
22149 root      15   0    0   3:32.06  0.0     0    0    0 S kjournald
25904 root      15   0    0  22:04.93  0.0     0    0    0 S kpanfs_thpool
25905 root      15   0    0  22:03.25  0.0     0    0    0 S kpanfs_thpool
```

→ **SHIFT + F** pour changer la colonne de tri, etc. **h** pour l'aide.

Top process

```
ps -eo pcpu,pid,user,args | sort -k 1 -r | head -10
ps -eo pcpu,pid,user,args | sort -r -k1 | less
```

cf. http://www.cyberciti.biz

Mémoire

- Vérifier la mémoire avec un **sar -r 1 10** et **free -m** :

```
root@server1104092:~# sar -r 1 10
Linux 2.6.9-67.ELsmp (server1104092)     11/12/08

15:05:25    kbmemfree kbmemused  %memused kbbuffers  kbcached kbswpfree kbswpused  %swpused  kbswpcad
15:05:26       488156  65259692     99.26    170472  62419380  65535016       976      0.00         0
15:05:27       479004  65268844     99.27    170488  62426708  65535016       976      0.00         0
15:05:28       471644  65276204     99.28    170504  62436008  65535016       976      0.00         0
15:05:29       463452  65284396     99.30    170512  62444228  65535016       976      0.00         0
15:05:30       454292  65293556     99.31    170540  62453448  65535016       976      0.00         0
15:05:31       445596  65302252     99.32    170552  62462344  65535016       976      0.00         0
```

```
root@server1104092:~# free -m
             total       used       free     shared    buffers     cached
Mem:         64206      63597        609          0        165      60824
-/+ buffers/cache:       2607      61599
Swap:        63999          0      63999
```

→ ici on a *2607 Mo* de ram utilisée et *60824 Mo* de cache (utilisable par le kernel en cas de besoin). La valeur *63597 Mo* correspond à la mémoire allouée, *61599* correspond à la mémoire disponible.

- **vmstat -s** affiche les infos détaillées sur la mémoire :

```
root@server2311829:/# vmstat -s
     24665472  total memory
     15389220  used memory
     14490508  active memory
       136844  inactive memory
      9276252  free memory
       768868  buffer memory
     10595768  swap cache
     25165816  total swap
            0  used swap
     25165816  free swap
     12203574 non-nice user cpu ticks
        32747 nice user cpu ticks
    141624951 system cpu ticks
   4344770172 idle cpu ticks
       579126 IO-wait cpu ticks
       409728 IRQ cpu ticks
       181938 softirq cpu ticks
      9452723 pages paged in
    690716466 pages paged out
            0 pages swapped in
            0 pages swapped out
   1640151373 interrupts
    765936041 CPU context switches
   1223379772 boot time
    314718391 forks
```

- un **cat /proc/meminfo** donne aussi des infos pertinentes (noyau 2.6 dans notre cas) :

```
root@parcl1104092:~# cat /proc/meminfo
MemTotal:     65747848 kB
MemFree:       4676348 kB
Buffers:        155412 kB
Cached:       58250332 kB
SwapCached:          0 kB
Active:       42665024 kB
Inactive:     17041772 kB
HighTotal:           0 kB
HighFree:            0 kB
LowTotal:     65747848 kB
LowFree:       4676348 kB
SwapTotal:    65535992 kB
SwapFree:     65535016 kB
Dirty:            3148 kB
Writeback:           0 kB
Mapped:       31220740 kB
Slab:           789220 kB
CommitLimit:  98409916 kB
Committed_AS: 33127992 kB
PageTables:     419756 kB
VmallocTotal: 536870911 kB
VmallocUsed:     51352 kB
VmallocChunk: 536819195 kB
HugePages_Total:     0
HugePages_Free:      0
Hugepagesize:     2048 kB
```

→ *Buffers & Cached* : cette mémoire peut-être utilisée par le kernel en cas de besoin, c'est la mémoire disponible (mais allouée)
→ *Active* : mémoire récemment utilisée qui ne sera pas utilisée par le kernel sauf en cas de besoin fort
→ *Inactive* : mémoire utilisée pouvant être utilisée par le kernel si besoin
→ *Dirty* : datas non flushées sur le disque (ou en cours)
→ D'autres infos ici.

- Un petit **vmstat 1 10** donne aussi des infos utiles :

```
root@server1104092:~# vmstat 1 10
procs -----------memory---------- ---swap-- -----io---- --system-- ----cpu----
 r  b   swpd   free   buff  cache   si   so    bi    bo   in    cs us sy id wa
 0  2    976 472404 165332 62007272    0    0   460   550    5     2 12  5 77  6
 2  0    976 467988 165340 62012296    0    0  2208   236 16949 36775 15 12 66  8
 5  0    976 463380 165360 62016900    0    0  2088  1705 9980 36940 16  9 70  5
 3  0    976 457684 165364 62022064    0    0  2264   397 7457 44781 12  9 73  6
```

```
 6  0    976 452564 165376 62027832    0    0  2260   275 6159 45565 20 10 66  4
 1  3    976 447988 165376 62033068    0    0  2376  9536 7778 35145 15 10 68  7
 3  1    976 440628 165384 62038772    0    0  2496   556 8299 71486 21 11 63  5
 0  0    976 433684 165400 62045624    0    0  3108   696 5167 52533 10  9 74  6
 0  1    976 427284 165400 62051948    0    0  2892   216 4757 45852 12  8 74  7
 1  0    976 421468 165404 62056160    0    0  1620   609 6609 69289 11  8 75  6
```

Les champs *so* et *si* permettent de voir l'activité de swap. On peut voir la runqueue *r* et les process *b* qui attendent des I/O (réseau, disque) *b*, cf le man pour les autres infos (cs → context swicth, in → interrupts, etc).

Disque

- **vmstat -d 1 1**

```
root@server2311829:/# vmstat -d  1 1
disk- ------------reads------------ ------------writes----------- -----IO------
       total merged sectors      ms  total merged sectors      ms    cur    sec
ram0       0      0       0       0      0      0       0       0      0      0
ram1       0      0       0       0      0      0       0       0      0      0
ram2       0      0       0       0      0      0       0       0      0      0
ram3       0      0       0       0      0      0       0       0      0      0
ram4       0      0       0       0      0      0       0       0      0      0
ram5       0      0       0       0      0      0       0       0      0      0
ram6       0      0       0       0      0      0       0       0      0      0
ram7       0      0       0       0      0      0       0       0      0      0
ram8       0      0       0       0      0      0       0       0      0      0
ram9       0      0       0       0      0      0       0       0      0      0
ram10      0      0       0       0      0      0       0       0      0      0
ram11      0      0       0       0      0      0       0       0      0      0
ram12      0      0       0       0      0      0       0       0      0      0
ram13      0      0       0       0      0      0       0       0      0      0
ram14      0      0       0       0      0      0       0       0      0      0
ram15      0      0       0       0      0      0       0       0      0      0
hda        0      0       0       0      0      0       0       0      0      0
sda   194356   5945 4172146  137864 96031752 77260674 1380561782 51965656      0   6458
dm-0   11404      0  509890   13774 28497459      0 227979672 39823426      0    701
dm-1   13623      0  533730   13236 7019350      0 56154800 11378930      0    637
dm-2    2998      0   24154     844 109615714      0 876925712 7213939      0   3883
dm-3    2972      0   23770    1397 280989      0 2247912 3567093      0     42
dm-4    3373      0   78090    1367 2954396      0 23635168 5853703      0    166
dm-5   69064      0 2116530   67730 2682039      0 21456312 8221587      0    204
dm-6    5320      0  114218    7105 2974177      0 23793416 267554994      0    381
dm-7    2909      0   23272     621      0      0       0       0      0      0
dm-8    3222      0   25770    8636 7034863      0 56278904 2119759062      0    217
md0        0      0       0       0      0      0       0       0      0      0
loop0      0      0       0       0      0      0       0       0      0      0
loop1      0      0       0       0      0      0       0       0      0      0
loop2      0      0       0       0      0      0       0       0      0      0
loop3      0      0       0       0      0      0       0       0      0      0
loop4      0      0       0       0      0      0       0       0      0      0
loop5      0      0       0       0      0      0       0       0      0      0
loop6      0      0       0       0      0      0       0       0      0      0
loop7      0      0       0       0      0      0       0       0      0      0
dm-9    2235      0   17874    2412 474271      0 3794168 150902341      0     27
dm-10   3787      0   22986    1879 318459      0  654606  163241      0      9
dm-11   5240      0   34388    4321 197786      0  424060  493783      0      5
dm-12   5274      0   34456    4168 323270      0  675028  963431      0      5
dm-15   1721      0   13762    1684 5671213      0 45369704 8472081      0    458
dm-16   1708      0   13658    8814 469486      0 3755888 64149095      0     21
dm-17   1764      0   14106    3557 4370518      0 34964144 204951066      0    231
dm-14   3378      0   19768    1341  62748      0  138480  209068      0      2
dm-18   3784      0   22980    1701  45660      0  109008   60575      0      3
dm-19    807      0    6450     525 249030      0 1992240 6801428      0      4
dm-20   2868      0   19382    1516  25806      0   69300 1170260      0      1
```

→ Ici on voit l'activité de chaque device. Pour déterminer à quel LV correspond quel device dm-x voir ici.

- **iostat -k -x 1 10**

```
root@server2311828:/# iostat -k -x 1 1
Linux 2.6.9-67.ELsmp (parsl2311828)     12/11/2008

avg-cpu:  %user   %nice    %sys %iowait   %idle
           1.09    0.00   18.40    0.11   80.40
```

```
Device:    rrqm/s wrqm/s   r/s   w/s  rsec/s  wsec/s    rkB/s    wkB/s avgrq-sz avgqu-sz   await  svctm  %util
sda          0.00  62.30  0.05 106.47    0.92  585.84     0.46   292.92     5.51     0.01    0.13   0.05   0.58
dm-0         0.00   0.00  0.00 20.08    0.12  160.66     0.06    80.33     8.00     0.02    0.82   0.02   0.05
dm-1         0.00   0.00  0.00  1.95    0.10   15.56     0.05     7.78     8.04     0.00    0.67   0.05   0.01
dm-2         0.00   0.00  0.00 139.36    0.00  351.81     0.00   175.91     2.52     0.01    0.06   0.04   0.51
dm-3         0.00   0.00  0.00  0.03    0.00    0.24     0.00     0.12     8.00     0.00    0.26   0.05   0.00
dm-4         0.00   0.00  0.00  1.65    0.00   13.20     0.00     6.60     8.00     0.00    2.19   0.04   0.01
dm-5         0.00   0.00  0.01  0.65    0.41    5.22     0.21     2.61     8.44     0.00    1.87   0.06   0.00
dm-6         0.00   0.00  0.00  0.64    0.03    5.12     0.01     2.56     8.03     0.03   51.40   0.08   0.00
dm-7         0.00   0.00  0.00  0.00    0.00    0.00     0.00     0.00     8.00     0.00    0.34   0.21   0.00
dm-8         0.00   0.00  0.00  0.84    0.00    6.70     0.00     3.35     8.00     0.15  184.80   0.03   0.00
dm-9         0.00   0.00  0.00  0.20    0.01    1.60     0.00     0.80     8.01     0.04  205.42   0.04   0.00
dm-10        0.00   0.00  0.00  0.03    0.00    0.07     0.00     0.04     2.18     0.00    0.35   0.03   0.00
dm-11        0.00   0.00  0.00  0.05    0.01    0.10     0.00     0.05     2.20     0.00    0.41   0.03   0.00
dm-12        0.00   0.00  0.00  0.08    0.01    0.16     0.00     0.08     2.12     0.00    0.38   0.03   0.00
dm-15        0.00   0.00  0.00  1.67    0.06   13.35     0.03     6.67     8.02     0.00    0.29   0.08   0.01
dm-16        0.00   0.00  0.00  0.17    0.00    1.33     0.00     0.67     8.00     0.00    0.55   0.04   0.00
dm-17        0.00   0.00  0.00  1.18    0.03    9.44     0.02     4.72     8.00     0.13  108.44   0.07   0.01
dm-14        0.00   0.00  0.00  0.01    0.00    0.03     0.00     0.01     2.39     0.00    1.09   0.06   0.00
dm-18        0.00   0.00  0.00  0.01    0.00    0.02     0.00     0.01     2.73     0.00    0.40   0.10   0.00
dm-19        0.00   0.00  0.00  0.13    0.00    1.01     0.00     0.51     8.00     0.00    0.63   0.05   0.00
dm-20        0.00   0.00  0.00  0.02    0.01    0.05     0.00     0.03     2.34     0.00    0.57   0.04   0.00
```

→ Pour afficher les infos sur l'utilisation des disques

Réseau

- **mii-tool** et **ethtool** pour vérifier le statut des cartes :

```
root@server1104092:~# mii-tool
eth0: 100 Mbit, full duplex, link ok
eth1: no link
eth2: 100 Mbit, full duplex, link ok
eth3: negotiated, link ok
eth4: 100 Mbit, full duplex, link ok
eth5: negotiated, link ok
eth6: 100 Mbit, full duplex, link ok
eth7: negotiated, link ok
```

```
root@server1104092:~# ethtool eth7
Settings for eth7:
        Supported ports: [ TP ]
        Supported link modes:   10baseT/Half 10baseT/Full
                                100baseT/Half 100baseT/Full
                                1000baseT/Full
        Supports auto-negotiation: Yes
        Advertised link modes:  1000baseT/Full
        Advertised auto-negotiation: Yes
        Speed: 1000Mb/s
        Duplex: Full
        Port: Twisted Pair
        PHYAD: 0
        Transceiver: internal
        Auto-negotiation: on
        Supports Wake-on: umbg
        Wake-on: g
        Current message level: 0x00000007 (7)
        Link detected: yes
```

- **netstat -taunpo|grep -c ESTA** pour vérifier le nombre de connexions établies

Sémaphores

Pour voir la conf des sémaphores :

```
root@server9010504:/# cat /proc/sys/kernel/sem
1024    32000   100     606

SEMMSL - semaphores per ID
SEMMNS - (SEMMNI*SEMMSL) max semaphores in system
SEMOPM - max operations per semop call
```

```
SEMMNI - max semaphore identifiers
```

Ca se modifie à chaud avec la commande :

```
echo "1024    32000   100     606" > /proc/sys/kernel/sem
```

Il ne faut pas oublier de modifier la ligne kernel.sem du fichier */etc/sysctl.conf* pour la prise en compte aux prochains reboots :

```
root@server9010504:/# grep kernel.sem /etc/sysctl.conf
kernel.sem = 1024        32000   100    606
```

On retrouve les infos aussi avec la commande ipcs :

```
root@server9010504:/# ipcs -s -l

------ Semaphore Limits --------
max number of arrays = 606
max semaphores per array = 1024
max semaphores system wide = 32000
max ops per semop call = 100
semaphore max value = 32767
```

Description

Semaphore	Description	Minimum
SEMMSL	maximum number of semaphores per array	128
SEMMNS	maximum semaphores system-wide	
SEMOPM	maximum operations per semop call	
SEMMNI	maximum arrays	

Calcul des valeurs

Source [1)]

- Calculate the minimum total semaphore requirements using the following formula:

```
sum (process parameters of all database instances on the system) + system and other application requirements
```

Set semmns (total semaphores systemwide) to this total.
Set semmsl (semaphores per set) to 256.
Set semmni (total semaphores sets) to semmns / semmsl rounded up to the nearest multiple of 1024.

The following formula can be used as a guide, although in practice, SEMMNS and SEMMNU can be much less than SEMMNI * SEMMSL because not every program in the system needs semaphores.

SEMMNS = SEMMNU = (SEMMNI * SEMMSL)

[1)]
https://gerardnico.com/wiki/linux/semaphore