

# Re: Initial 6.1 questions

---

*Source:* <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/performance/2006-06/msg00027.html>

---

- *From:* Danial Thom <[danial\\_thom@xxxxxxxxxx](mailto:danial_thom@xxxxxxxxxx)>
  - *Date:* Tue, 13 Jun 2006 12:48:46 -0700 (PDT)
- 

----- Robert Watson <[rwatson@xxxxxxxxxxxxx](mailto:rwatson@xxxxxxxxxxxxx)> wrote:

On Tue, 13 Jun 2006, Danial Thom wrote:

Two types of measurements are taken: sampled ticks regarding whether the system as a whole is in {user, nice, system, intr, idle}, and then sampling for individual processes. Right now, the system measurements are kept in a simple array of tick counters called cp\_time. John Baldwin and others have changes that make these tick counters per-CPU. The lines at the top of top(1)'s output are derived from those tick counters. Ticks are measured on each CPU, so those are a summary across all CPUs. To add cpustat support, we need to merge John's patch to

Re: Initial 6.1 questions

make `cp_time` per-CPU (ie.,  
different counters for different CPUs) and  
teach the userland tools to  
retrieve them. When you run `top` you'll  
notice that it adjusts the  
measurements each refresh. In effect, what  
it's doing is sampling the  
change in tick counts over the window,  
pulling down the new values and  
calculating the percentages of ticks in each  
"bucket" in the last window.

That doesn't explain why the Top line shows  
99.6% idle, but the `cpu idle`  
threads are showing significant usage.  
I'm getting a constant 6000 Interrupts /  
Second on my em controller, yet `top`  
jumps all over the place; sitting at 99% idle  
for 10 seconds, then jumping  
to 50%, then somewhere in between. It seems  
completely unreliable. The load  
I'm applying is constant.

I can't speak to the details of the  
thread/process use sampling model. `Top`  
uses something called the "weighted `cpu`  
percentage" by default; you can switch  
to "unweighted" using the `-C` argument. The `top`  
documentation fails to  
document the semantics of the percentages, but

Re: Initial 6.1 questions

Re: Initial 6.1 questions

I suspect -C will give you more of what you expect. The weighted CPU measurement takes into account process history, so it takes a while for sudden spike in CPU use to be fully reflected, and you may see seemingly counter-intuitive results, such as the appearance of greater than 100% CPU use. Try out -C and see if you see something that makes more sense?

It seems to work just fine with 1 CPU. Its equally useless with the -C option in SMP mode.

Here's a snip from 'systat -vmstat 1'

```
Proc:r p d s w Csw Trp Sys Int Sof
Flt cow 10009 total
24 18353 1 129 156k 1
17108 wire 6: fdc0

7908 act 14: ata
0.4%Sys 0.4%Intr 0.0%User 0.0%Nice 99.2%Idl
7236 inact 20: em0
|||||||||
cache 6000 21: em1

473456 free 5 24: bge
```

6000 interrupts per second and .4% interrupt usage. Clearly the tools don't work at all in SMP mode. I don't see how you can do development without measurement tools that work.

DT

---

Do You Yahoo!?

Tired of spam? Yahoo! Mail has the best spam protection around  
<http://mail.yahoo.com>

---

freebsd-performance@xxxxxxxxxxx mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-performance>

To unsubscribe, send any mail to "freebsd-performance-unsubscribe@xxxxxxxxxxx"

Re: Initial 6.1 questions