

Re: Changing p_swtime and td_slptime to ticks

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/arch/2007-09/msg00022.html>

- *From:* Jeff Roberson <jroberson@xxxxxxxxxxxxxxxx>
 - *Date:* Tue, 18 Sep 2007 01:25:52 -0700 (PDT)
-

On Tue, 18 Sep 2007, Kostik Belousov wrote:

On Mon, Sep 17, 2007 at 05:05:57PM -0700, Jeff Roberson wrote:

Enclosed is a patch that fixes swapping with ULE. ULE has never properly set p_swtime and td_slptime which are used by the swapout/swapin code to select the appropriate thread to swap.

In 4BSD these two variables are increment once per-second as schedcpu() iterates over all threads. ULE does not have a once per-second loop iterating over all threads. So I have changed p_swtime to p_swtick and td_slptime to td_slptick. These record the value of 'ticks' when the thread slept or was last swapped in or out.

For backwards compatibility I leave the values in kinfo_proc with the legacy meaning by subtracting from ticks and dividing by hz. I perform a similar transformation in the swapout code to convert to seconds. This change does make it possible to use sub-second granular decisions in the swap code, however I'm not sure if that's really necessary.

So that I did not disturb the 4BSD mechanism I kept the original td_slptime in the td_sched area. It should be possible to use td_slptick directly but especially this close to release I did not want to change 4BSD.

Feedback and testing welcome.

Purely cosmetic request: please make the ticks and hz variables in the libkvm/kvm_proc.c static. Or, even better, move it into the struct __kvm.

Thanks, will do.

Jeff

frebsd-arch@xxxxxxxxxxxxx mailing list

Re: Changing p_swtime and td_slptime to ticks

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

To unsubscribe, send any mail to "freebsd-arch-unsubscribe@xxxxxxxxxxxx"