

ksoftird doesn't work like it should (?)

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-03/msg09546.html>

- *From:* "Aritz Bastida" <aritzbastida@xxxxxxxxxx>
 - *Date:* Tue, 28 Mar 2006 12:26:57 +0200
-

Hello everybody.

First of all, thank you all of you that helped me when i was doing my thesis project: olsson, greg k-h, ... I could learn a lot (although i'm still a newbie, of course) in these 8 months, and wrote in this mailing list when i was too lost...

Well, you won't remember me, but I had to design a network analysis system in kernel space. Letting apart that is or is not useful (I didnt choose that, I just had to do it), I made some testing that could be interesting, or that's what i think.

That's why I would like to share it with you. Here I attach a link to a picture, which show the problems they happen when the network load goes high.

<http://www.flickr.com/photos/57861108@N00/119255818/>

This picture shows the throughput of a network analysis system in user space (not mine), with different injection speeds.

x axis: traffic speed (packets per second)

y axis: analyzer throughput (packets per second)

You can see that, when the network traffic is low, the network analyzer can process all the packets (straight line), but when it gets too high, it starts losing packets. If it gets even higher, the ksoftirqd kernel thread starts eating 99.9% CPU time, even if it is supposed that its priority is the lowest possible, so the throughput goes lower. ksoftirqd is supposed that would not starve user programs, but the system almost goes to livelock. The system stabilizes around 1.2 Gpps.

So what we can conclude here is that ksoftirqd doesnt work as it should, or maybe do_softirq() is called from more places apart from the ksoftird kernel thread. In any case, the user process almost starves.

If you'd like, i have more pictures similar to this, and the "solution" i found, which tries to make the throughput to be constant

ksoftird doesn't work like it should (?)

independently of the traffic speed.

Regards

Aritz

–

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxx

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>