Re: 2.6.12 Performance problems

Source: http://linux.derkeiler.com/Mailing−Lists/Kernel/2005−08/5524.html

---

From: Danial Thom (*danial_thom_at_yahoo.com*)
Date: 08/21/05

```
Date:   Sun, 21 Aug 2005 10:07:19 −0700 (PDT)
To: linux−kernel@vger.kernel.org
```

−−− Jesper Juhl <jesper.juhl@gmail.com> wrote:

> *On 8/21/05, Danial Thom <danial_thom@yahoo.com>*
> *wrote:*
> > *I just started fiddling with 2.6.12, and*
> *there*
> > *seems to be a big drop−off in performance*
> *from*
> > *2.4.x in terms of networking on a*
> *uniprocessor*
> > *system. Just bridging packets through the*
> > *machine, 2.6.12 starts dropping packets at*
> > *~100Kpps, whereas 2.4.x doesn't start*
> *dropping*
> > *until over 350Kpps on the same hardware*
> *(2.0Ghz*
> > *Opteron with e1000 driver). This is pitiful*
> > *prformance for this hardware. I've*
> > *increased the rx ring in the e1000 driver to*
> *512*
> > *with little change (interrupt moderation is*
> *set*
> > *to 8000 Ints/second). Has "tuning" for MP*
> > *destroyed UP performance altogether, or is*
> *there*
> > *some tuning parameter that could make a*
> *4−fold*
> > *difference? All debugging is off and there*
> *are*
> > *no messages on the console or in the error*
> *logs.*
> > *The kernel is the standard kernel.org dowload*
> > *config with SMP turned off and the intel*
> *ethernet*
> > *card drivers as modules without any other*
> > *changes, which is exactly the config for my*

> 2.4
> > kernels.
> >
>
> If you have preemtion enabled you could disable
> it. Low latency comes
> at the cost of decreased throughput − can't
> have both. Also try using
> a HZ of 100 if you are currently using 1000,
> that should also improve
> throughput a little at the cost of slightly
> higher latencies.
>
> I doubt that it'll do any huge difference, but
> if it does, then that
> would probably be valuable info.
>
Ok, well you'll have to explain this one:

"Low latency comes at the cost of decreased
throughput − can't have both"

Seems to be a bit backwards. Threading the kernel
adds latency, so its the additional latency in
the kernel that causes the drop in throughput. Do
you mean that kernel performance has been
sacrificed in order to be able to service other
threads more quickly, even when there are no
other threads to be serviced?

Danial

_____
Start your day with Yahoo! − make it your home page
http://www.yahoo.com/r/hs

–
To unsubscribe from this list: send the line "unsubscribe linux−kernel" in
the body of a message to majordomo@vger.kernel.org
More majordomo info at http://vger.kernel.org/majordomo−info.html
Please read the FAQ at http://www.tux.org/lkml/