
Subject: time wait bucket table overflow

Posted by [Cliff Wells](#) on Tue, 15 Jan 2008 18:10:40 GMT

[View Forum Message](#) <> [Reply to Message](#)

I was doing some benchmarking with ab and noticed these in my syslog:

```
Jan 14 16:20:42 vps1 TCP: time wait bucket table overflow
Jan 14 16:20:42 vps1 TCP: time wait bucket table overflow
Jan 14 16:20:42 vps1 TCP: time wait bucket table overflow
Jan 14 16:20:42 vps1 TCP: time wait bucket table overflow
Jan 14 16:20:43 vps1 printk: 1060 messages suppressed.
```

Google turned up some info on this:

http://bugzilla.openvz.org/show_bug.cgi?id=460

But I've already got that parameter turned up:

```
# cat /proc/sys/net/ipv4/tcp_max_tw_buckets_ub
32768
```

Is there somewhere else I should look?

Note that the benchmarking I was doing was actually in the host OS, not a VE.

Regards,
Cliff

Subject: Re: time wait bucket table overflow

Posted by [vaverin](#) on Wed, 16 Jan 2008 06:35:37 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi Cliff,

I think network stress testing can create much more time wait buckets. It is information message and not an error therefore IMHO it is not a problem for you.

Do you want to allow it? Try to increase
/proc/sys/net/ipv4/tcp_max_tw_buckets (global limit)
and
/proc/sys/net/ipv4/tcp_max_tw_buckets_ub (per-VE limit, but host system itself accounts like usual VE)
values.

thank you,

Subject: Re: Re: time wait bucket table overflow
Posted by [Cliff Wells](#) on Wed, 16 Jan 2008 09:11:07 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Wed, 2008-01-16 at 09:35 +0300, vaverin wrote:

>
> Hi Cliff,
>
> I think network stress testing can create much more time wait buckets.
> It is information message and not an error therefore IMHO it is not a
> problem for you.

I was concerned because I was seeing odd results during testing (the ab test was running 500-1000 concurrent connections). Most of the connections finished in <2ms but a few would take almost as long as the entire test (15s). This may be an ab bug however.

>
> Do you want to allow it? Try to increase
> /proc/sys/net/ipv4/tcp_max_tw_buckets (global limit)
> and
> /proc/sys/net/ipv4/tcp_max_tw_buckets_ub (per-VE limit, but host
> system itself accounts like usual VE)
> values.

>
Yes, I tried this. Apparently I did not go high enough.

Thanks for the help.

Regards,
Cliff
