

---

Subject: Re: TCP: time wait bucket table overflow  
Posted by [vaverin](#) on Wed, 26 Dec 2007 09:07:23 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

these messages are printed in `tcp_time_wait()` function  
[http://git.openvz.org/?p=linux-2.6.18-openvz;a=blob;f=net/ipv4/tcp\\_minisocks.c;h=e280537ef09aa804f79e733eae1721a07b5aca](http://git.openvz.org/?p=linux-2.6.18-openvz;a=blob;f=net/ipv4/tcp_minisocks.c;h=e280537ef09aa804f79e733eae1721a07b5aca) c;hb=HEAD  
in the following cases:

1) Per-system `tw_count` is greater than per-system `max_tw_buckets` limit:  
`tw_count < sysctl_max_tw_buckets`  
`net.ipv4.tcp_max_tw_buckets = 1800000`  
IMHO very unlikely in your situation

2) Per-VE counter is greater than per-VE `max_tw_buckets` limit:  
`tw_count < sysctl_tcp_max_tw_buckets_ub`  
`net.ipv4.tcp_max_tw_buckets_ub = 165360`

IMHO unlikely too

3) inside VE `tw_buckets` eats too many memory (greater than allowed fraction of `kmemsize`)  
`net.ipv4.tcp_max_tw_kmem_fraction = 384` means 38.4% of `kmemsize`

IMHO very likely

4) `kmemsize` shortage inside VE -- in this case you should have new failcounters for `kmemsize` parameter.  
IMHO unlikely

Resume:

You need to find VE that generates these messages and increase its `kmemsize`. Or You can increase setting for `net.ipv4.tcp_max_tw_kmem_fraction` `sysctl`

ToDo: make this message cleaner  
[http://bugzilla.openvz.org/show\\_bug.cgi?id=767](http://bugzilla.openvz.org/show_bug.cgi?id=767)

---