Subject: kernel: TCP: time wait bucket table overflow
Posted by lakostis on Mon, 02 Oct 2006 19:42:29 GMT

View Forum Message <> Reply to Message

Hi!

After updating to test018 OpenVZ patch my syslog full of those messages.
/proc/user_beancounters shows nothing unusual comparing with previous
kernel. Various docs says that this message indicates low memory situation
or simple DoS attack and consider increase net.ipv4.tcp_max_tw_buckets but
it didn't help. I'm found another sysctl - net.ipv4.tcp_max_tw_buckets_ub,
but it undocumentend in ovz docs and looking in code shows that it's
internal switch for ubc counters. So what I can do now and how ovz kernel
affects this? :) Okay, it's heavy loaded system with 2 Gigs of RAM/1.4 P4
CPU and 1x100Mbit e100 NIC. According nginx internal status there is not
so many connections:

Active connections: 1067
server accepts handled requests
 6494 6494 9525
 Reading: 1 Writing: 6 Waiting: 1060

~# uptime
23:36:22 up 2 days, 10:25, 1 user, load average: 0.80, 2.06, 2.51

~# free
        total   used   free   shared   buffers   cached
Mem: 2072844 1825712   247132   0   128736   951484
-/+ buffers/cache: 745492   1327352
Swap: 995896   3780   992116

PS I can attach kernel buildconfig and dmesg if needed.

--
WBR et al.