
Subject: 2.6.18-028test007.1 crush
Posted by [dagr](#) on Thu, 21 Dec 2006 01:27:16 GMT
[View Forum Message](#) <> [Reply to Message](#)

kernel 2.6.18-028test007.1 , rhel 4
I posted its messages right before hanging here :

http://bugzilla.openvz.org/show_bug.cgi?id=412

Looks like scheduler bug. Important thing - no VPSs were up. Just kernel and production services on HN ie VPS 0 .

Subject: Re: 2.6.18-028test007.1 crush
Posted by [dagr](#) on Thu, 21 Dec 2006 01:42:54 GMT
[View Forum Message](#) <> [Reply to Message](#)

BTW - used to got another error messages before.
Server was under high number connections load.
After starting server under openvz kernel we start get messages in /var/log/messages like this
TCP: time wait bucket table overflow
printk: 116 messages suppressed.

.....

.....

Looks like some VPS 0 restrictions, although its unlimited

sudo cat /proc/user_beancounters (this is under 30% of average load)

uid resource	held	maxheld	barrier	limit	failcnt
0: kmemsize	65457707	67576465	2147483647	2147483647	0
lockedpages	1019	1019	2147483647	2147483647	0
privvmpages	393767	409333	2147483647	2147483647	0
shmpages	4615	5271	2147483647	2147483647	0
dummy	0	0	2147483647	2147483647	0
numproc	928	965	2147483647	2147483647	0
physpages	80131	84819	2147483647	2147483647	0
vmguarpages	0	0	2147483647	2147483647	0
oomguarpages	80131	84819	2147483647	2147483647	0
numtcpsock	4225	4282	2147483647	2147483647	0
numflock	3	7	2147483647	2147483647	0
numpty	5	6	2147483647	2147483647	0
numsiginfo	0	4	2147483647	2147483647	0
tcpsndbuf	38173792	38625464	2147483647	2147483647	0
tcprcvbuf	70973068	71957612	2147483647	2147483647	0
othersockbuf	1156384	1237324	2147483647	2147483647	0
dgramrcvbuf	0	9800	2147483647	2147483647	0

numothersock	549	586	2147483647	2147483647	0
dcachesize	0	0	2147483647	2147483647	0
numfile	23462	23895	2147483647	2147483647	0
dummy	0	0	2147483647	2147483647	0
dummy	0	0	2147483647	2147483647	0
dummy	0	0	2147483647	2147483647	0
numiptent	113	113	2147483647	2147483647	0

Subject: Re: 2.6.18-028test007.1 crush

Posted by [dagr](#) on Thu, 21 Dec 2006 02:50:56 GMT

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

restarted with kernel 2.6.9 023stab037.3
got just same error in message log

Checked sysctl

```
sudo sysctl -a | grep -i tcp_max
```

```
net.ipv4.netfilter.ip_conntrack_tcp_max_retrans = 3
net.ipv4.tcp_max_syn_backlog = 1024
net.ipv4.tcp_max_tw_buckets_ve = 2048
net.ipv4.tcp_max_tw_kmem_fraction = 128
net.ipv4.tcp_max_tw_buckets = 180000
net.ipv4.tcp_max_orphans = 131072
```

```
set net.ipv4.tcp_max_syn_backlog = 2048
```

No more error messages so far.

First question - why net.ipv4.tcp_max_tw_buckets_ve and net.ipv4.tcp_max_tw_buckets - differ so much ? It looks like this parameter affects HN as well , ie VPS 0 .

How far is it safe to increase this parameter in case i have no running VPSs at all, is it safe to set it 180000 ?

Still wondering if that kernel panic was caused by this thing , it used to work ok for 2 day with that error messages.
