
Subject: Too many orphaned sockets in CT0
Posted by [tobia](#) on Tue, 16 Mar 2010 09:28:46 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello

I'm getting bursts of hundreds of these messages every few hours:

Mar 16 08:58:08 vishnu kernel: [8604106.249004] TCP: too many of orphaned sockets (125 in CT0)

Mar 16 08:58:13 vishnu kernel: [8604112.427568] __ratelimit: 22 messages suppressed

Mar 16 08:58:13 vishnu kernel: [8604112.427568] TCP: too many of orphaned sockets (125 in CT0)

Mar 16 08:58:18 vishnu kernel: [8604118.668693] __ratelimit: 18 messages suppressed

Mar 16 08:58:18 vishnu kernel: [8604118.668693] TCP: too many of orphaned sockets (125 in CT0)

Mar 16 08:58:23 vishnu kernel: [8604124.151381] __ratelimit: 13 messages suppressed

They are interleaved with bursts of these other messages:

Mar 16 09:02:03 vishnu kernel: [8604394.874281] Orphaned socket dropped (124,248 in CT0)

Mar 16 09:02:08 vishnu kernel: [8604400.663420] __ratelimit: 11 messages suppressed

Mar 16 09:02:08 vishnu kernel: [8604400.663446] Orphaned socket dropped (112,224 in CT0)

Mar 16 09:02:16 vishnu kernel: [8604410.213555] __ratelimit: 5 messages suppressed

Mar 16 09:02:16 vishnu kernel: [8604410.213580] Orphaned socket dropped (106,212 in CT0)

Mar 16 09:02:19 vishnu kernel: [8604414.421430] __ratelimit: 2 messages suppressed

This is a Debian Lenny 2.6.26-2-openvz-amd64 with 6 VE plus VE0. Network configuration is a standard routed venet0. The only iptables rule is fail2ban-ssh, on the INPUT filter chain. The only unusual configuration is that the host's default gateway is a point-to-point address:

```
auto eth0
```

```
iface eth0 inet static
```

```
address 78.46.xx.x
```

```
netmask 255.255.255.255
```

```
gateway 78.46.xx.1
```

```
pointopoint 78.46.xx.1
```

Here are common socket usage numbers on the VEs, from netstat output:

```
VE 0 tcp 4 udp 6 unix 77
```

```
VE 302 tcp 158 udp 2 unix 6
```

```
VE 303 tcp 36 udp 2 unix 58
```

```
VE 304 tcp 31 udp 2 unix 6
```

```
VE 305 tcp 2 udp 2 unix 3
```

```
VE 311 tcp 4 udp 2 unix 6
```

```
VE 312 tcp 127 udp 2 unix 78
```

VE0 is completely unlimited, as it should be:

0: kmemsize	8910573	14586823	9223372036854775807
9223372036854775807	0		

lockedpages	0	8 9223372036854775807
9223372036854775807	0	
privvmpages	17995	22896031752 9223372036854775807
9223372036854775807	0	
shmpages	829	2765 9223372036854775807
9223372036854775807	0	
dummy	0	0 9223372036854775807 9223372036854775807
0		
numproc	130	165 9223372036854775807
9223372036854775807	0	
physpages	5466	97726 9223372036854775807
9223372036854775807	0	
vmguarpages	0	0 9223372036854775807
9223372036854775807	0	
oomguarpages	5655	97915 9223372036854775807
9223372036854775807	0	
numtcpsock	8	30 9223372036854775807
9223372036854775807	0	
numflock	5	17 9223372036854775807 9223372036854775807
0		
numpty	1	5 9223372036854775807 9223372036854775807
0		
numsiginfo	0	7 9223372036854775807 9223372036854775807
0		
tcpsndbuf	170304	6152224 9223372036854775807
9223372036854775807	0	
tcprcvbuf	131072	6699216 9223372036854775807
9223372036854775807	0	
othersockbuf	187272	4406680 9223372036854775807
9223372036854775807	0	
dgramrcvbuf	0	5648 9223372036854775807
9223372036854775807	0	
numothersock	148	212 9223372036854775807
9223372036854775807	0	
dcachesize	1786335	4051908 9223372036854775807
9223372036854775807	0	
numfile	2469	4637 9223372036854775807
9223372036854775807	0	
dummy	0	0 9223372036854775807 9223372036854775807
0		
dummy	0	0 9223372036854775807 9223372036854775807
0		
dummy	0	0 9223372036854775807 9223372036854775807
0		
numiptent	14	16 9223372036854775807
9223372036854775807	0	

I tried doubling the values in tcp_mem, to no avail:

```
# cat /proc/sys/net/ipv4/tcp_mem
2297728 2301824 2305920
# echo 4595456 4603648 4611840 > /proc/sys/net/ipv4/tcp_mem
# cat /proc/sys/net/ipv4/tcp_mem
4595456 4603648 4611840
```

What can I do to debug it? This server has been up for 3 months without showing any error. The services are still up and running. Should I expect downtime anytime soon?

Subject: Re: Too many orphaned sockets in CT0
Posted by [TheStig](#) on Tue, 16 Mar 2010 12:04:26 GMT
[View Forum Message](#) <> [Reply to Message](#)

I had a similar problem a while ago (same kernel) when using drbd8 with ocfs2. as it was a test system and ocfs2 (the file locks nearly crashed the VPSs with two systems/apaches accessing the same files) didn't work nearly as well as expected I quit using this setup and the orphaned sockets/ratelimits disappeared.

as you mention the only unusual config is the P-T-P gw, maybe the crux lies in a slow or somehow "damaged" P-T-P?! just guessing...
