
Subject: 0:tcpsndbuf running full [was: Could NPTL cause problems?]

Posted by [mephisto](#) on Wed, 05 Apr 2006 10:24:17 GMT

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

Hi,

I've run into problems with one system running a development kernel while the other one doesn't (described here: <http://forum.openvz.org/index.php?t=msg&th=374>).

Now one distinct difference between the two systems is that the one that works has everything compiled without NPTL while the one that gets unstable after awhile (just having the openvz modules loaded) is compiled with NPTL. Could this cause the described problems?

Regards,

Mephisto

Subject: Re: Could NPTL cause problems?

Posted by [dev](#) on Wed, 05 Apr 2006 11:50:02 GMT

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

mephisto, can you describe more precisely how unstable it is? Ooopses? Crashes? Lockups? Please, attach some then. It is also helpful to get outputs of Alt-SysRq-p, Alt-SysRq-t, Alt-SysRq-m.

If it is a memory leak or something, need `/proc/slabinfo` and `/proc/meminfo`

I also wonder, have you succeeded building `initrd` in gentoo for `i2o_block` driver or not? I suppose it is easier to use this driver instead of deprecated one.

Subject: Re: Could NPTL cause problems?

Posted by [mephisto](#) on Wed, 05 Apr 2006 11:53:42 GMT

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

Well, I have no physical access to the system (200km away) so it's kinda hard to diagnose. I will lay my hands in a few days and will report how it behaves, also with the SysRQ messages. So far I was unable to boot it with `i2o_block`, `initrd` or not. Will see how this goes, too.

Subject: Re: 0:tcpsndbuf running full [was: Could NPTL cause problems?]

Posted by [mephisto](#) on Fri, 07 Apr 2006 16:57:36 GMT

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

OK, I narrowed it down. First of all I'm still unable to use `i2o_block`, but that's also the case for vanilla kernels, so I gave it up for now.

More importantly I was finally able to diagnose my crashes. I watched the host closely while it was under heavy network load (samba mostly and some rsync). When the errors appeared (network connection very slow, until the host is unreachable). I checked /proc/user_beancounters and found this:

Version: 2.5

uid	resource	held	maxheld	barrier	limit	failcnt
0:	kmemsize	5500576	9732444	2147483647	2147483647	0
	lockedpages	1027	1027	2147483647	2147483647	0
	privvmpages	19654	53732	2147483647	2147483647	0
	shmpages	1326	1345	2147483647	2147483647	0
	dummy	0	0	2147483647	2147483647	0
	numproc	64	153	2147483647	2147483647	0
	physpages	5236	16629	2147483647	2147483647	0
	vmguarpages	0	0	2147483647	2147483647	0
	oomguarpages	5280	16673	2147483647	2147483647	0
	numtcpsock	23	109	2147483647	2147483647	0
	numflock	66	569	2147483647	2147483647	0
	numpty	2	2	2147483647	2147483647	0
	numsiginfo	0	69	2147483647	2147483647	0
	tcpsndbuf	2147163432	2147486264	2147483647	2147483647	2898
	tcprcvbuf	13112	843416	2147483647	2147483647	0
	othersockbuf	44320	2678104	2147483647	2147483647	0
	dgramrcvbuf	0	26960	2147483647	2147483647	0
	numothersock	34	50	2147483647	2147483647	0
	dcachesize	3643446	4167100	2147483647	2147483647	0
	numfile	535	2536	2147483647	2147483647	0
	dummy	0	0	2147483647	2147483647	0
	dummy	0	0	2147483647	2147483647	0
	dummy	0	0	2147483647	2147483647	0
	numiptent	0	0	2147483647	2147483647	0
101:	kmemsize	5550088	5980112	105495756	116045331	0
	lockedpages	0	0	5151	5151	0
	privvmpages	311785	386062	2147483647	2147483647	0
	shmpages	2655	3295	7726	7726	0
	dummy	0	0	0	0	0
	numproc	163	178	8000	8000	0
	physpages	66399	67751	0	2147483647	0
	vmguarpages	0	0	77267	2147483647	0
	oomguarpages	66399	67751	200000	2147483647	0
	numtcpsock	20	31	8000	8000	0
	numflock	14	21	1000	1100	0
	numpty	0	1	512	512	0
	numsiginfo	0	3	1024	1024	0
	tcpsndbuf	3239936	3333776	23792300	35147230	0
	tcprcvbuf	1192	159528	2397252	35165252	0
	othersockbuf	28128	85104	20480000	2147483647	0
	dgramrcvbuf	0	243168	1198626	1198626	0

numothersock	24	53	8000	8000	0
dcachesize	402711	456183	23031052	23721984	0
numfile	3530	4337	41184	41184	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
numiptent	0	0	200	200	0

So the host's tcpsndbuf is running full. I should mention that I moved eth1 to the vps with NETDEV="eth1". This is the major configuration difference to my test system that runs fine with a venet0 interface.

The error also occurred without the vps loaded and traffic occurring just on the not virtualized eth0 interface. I've seen the error with 2.6.15-025stab014 and 2.6.16-026stab007. Do you need any more information?

Subject: Re: 0:tcpsndbuf running full [was: Could NPTL cause problems?]

Posted by [dev](#) on Sat, 08 Apr 2006 18:45:32 GMT

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

Quote:So the host's tcpsndbuf is running full.

looks like a memory leak... bug! can you please post one to bugzilla with a ref to forum and descriptions of your setup, network card etc.

Quote:I should mention that I moved eth1 to the vps with NETDEV="eth1". This is the major configuration difference to my test system that runs fine with a venet0 interface.

oh, this is important info, will try to reproduce it locally then. thanks!

Quote:

The error also occurred without the vps loaded and traffic occurring just on the not virtualized eth0 interface. I've seen the error with 2.6.15-025stab014 and 2.6.16-026stab007. Do you need any more information?

sorry, the last statement looks a bit contradictory to the prev quote. Can you make it clear for me? do you see the same bug, w/o eth1 delegated to VPS?

Subject: Re: 0:tcpsndbuf running full [was: Could NPTL cause problems?]

Posted by [mephisto](#) on Sat, 08 Apr 2006 19:12:36 GMT

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

dev wrote on Sat, 08 April 2006 20:45

sorry, the last statement looks a bit contradictory to the prev quote. Can you make it clear for me? do you see the same bug, w/o eth1 delegated to VPS?

You're right, it is. If I think about it, the virtualized eth1 most likely has nothing to do with it. I tried

to reproduce the error locally on another computer but the tcpsndbuf goes back to 0 while there is no traffic, so it might be related to some obscure kernel setting or to the hardware.
I also should mention that I removed the start_net section from my vz initscript (Gentoo that is) because I don't need the venet0 interface.
I'll open a bug tomorrow.

Subject: Re: 0:tcpsndbuf running full [was: Could NPTL cause problems?]
Posted by [dev](#) on Sun, 09 Apr 2006 05:52:47 GMT
[View Forum Message](#) <> [Reply to Message](#)

any way, mephisto, I would appreciate if you could report the described information (AltSysRq-M, /proc/meminfo, /proc/slabinfo, slabtop) after some time machine is working, but before it goes unresponsive/totally slow. Can you do it please? slab information will help quickly to determine where the leak occurs, so I will be able to make a patch for you.

Subject: Re: 0:tcpsndbuf running full [was: Could NPTL cause problems?]
Posted by [mephisto](#) on Sun, 09 Apr 2006 08:20:58 GMT
[View Forum Message](#) <> [Reply to Message](#)

OK, will do. This will most likely take till Tuesday, though.

Subject: Re: 0:tcpsndbuf running full [was: Could NPTL cause problems?]
Posted by [mephisto](#) on Tue, 11 Apr 2006 14:51:53 GMT
[View Forum Message](#) <> [Reply to Message](#)

I opened the bug report here:
http://bugzilla.openvz.org/show_bug.cgi?id=135
