
Subject: Application is being killed by kernel with signal 11

Posted by [iziz](#) on Tue, 31 Oct 2006 16:44:51 GMT

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

Hello,

I use PostgreSQL server inside OpenVZ VE. From time to time it falls saying that one of the backends exited abnormally (after 11 signal) and possibly corrupted shared memory. This leads to whole server restart which is highly undesirable in my setup.

I checked your wiki and adjusted UBC accordingly to 4GB of RAM I have on physical machine (I set up beancounters to have some free resources for future but also big enough to have PostgreSQL stable working as I thought). Here is my /proc/user_beancounters file:

uid resource	held	maxheld	barrier	limit	failcnt
98076:					
kmemsize	21901424	36869261	70451200	71000000	0
lockedpages	0	0	64	64	0
privvmpages	279568	299439	1560000	1572864	0
shmpages	243748	243764	327680	327680	0
dummy	0	0	0	0	0
numproc	33	45	650	650	0
physpages	266354	275590	0	2147483647	0
vmguarpages	0	0	524288	2147483647	0
oomguarpages	267917	277186	524288	2147483647	0
numtcpsock	20	34	800	800	0
numflock	2	5	400	440	0
numpty	2	2	16	16	0
numsiginfo	0	3	256	256	0
tcpsndbuf	6684	204976	3194880	5242880	0
tcprcvbuf	0	19900	3194880	5242880	0
othersockbuf	122540	498476	1320960	4000000	0
dgramrcvbuf	0	185856	1320960	1320960	0
numothersock	81	90	1000	1000	0
dcachesize	256004	327754	4915200	54067200	0
numfile	1660	2781	12800	12800	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
numiptent	19	19	128	128	0

So the questions are:

1. Why does kernel kill PostgreSQL with SIGSEGV?
2. Why don't I see increments in fail counters?
3. How is that possible to understand what resource limit (if I'm right in my assumption that kernel kills PostgreSQL because of resource limits excess) PostgreSQL tried to exceed given all maxheld values?

4. What's wrong with these UBC?

Forgot to tell:

```
[root@pg1 ~]# uname -r  
2.6.8-022stab078.21.ipsec.1-enterprise
```

Hope that someone can help me.
