
Subject: [PATCH 0/3] Fix problem with static_key decrement
Posted by [Glauber Costa](#) on Thu, 19 Apr 2012 22:49:15 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,

This is my proposed fix for the sock memcg static_key problem raised by Kamezawa. It works for me, but I would Kame, please confirm.

For that to work, I am dependent on two cgroup patches that goes attached. The rationale behind it, is that we can't do static_key updates with the cgroup_mutex held, or we risk deadlocking.

Looking closely, there seem to be no particular reason to hold the cgroup_mutex during destruction. Subsystems that really need it, can hold it themselves.

Tejun, let me know if this is acceptable from your PoV.

Glauber Costa (3):

- don't attach a task to a dead cgroup
- don't take cgroup_mutex in destroy()
- decrement static keys on real destroy time

```
block/blk-cgroup.c      |  2 +
include/net/sock.h      |  9 ++++++
kernel/cgroup.c         | 12 ++++++----
kernel/cpuset.c         |  2 +
mm/memcontrol.c         | 20 ++++++++-----
net/ipv4/tcp_memcontrol.c | 52 ++++++++++++++++++++++++++++++++++++++-----
6 files changed, 83 insertions(+), 14 deletions(-)
```

--

1.7.7.6
