
Subject: [PATCH v6 00/10] Request for inclusion: per-cgroup tcp memory pressure controls

Posted by [Glauber Costa](#) on Fri, 25 Nov 2011 17:38:06 GMT

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

Hi Dave,

I hope the following series is in an acceptable state: I modified the tests in `__sk_mem_schedule()` in a way that we should still leave the function pretty soon under no pressure conditions.

Also, I managed to remove almost everything tcp related from `memcontrol.c`: the only thing left is a simple function to calculate the address of the tcp sub-structure in the main memcg structure - I hope this is acceptable, since besides being simple, it is not related to the protocol itself.

I also believed that by now, all other comments so far were addressed. Let me know if there are any blocking concerns to this, and I'll address them as soon as I can.

Thanks

Glauber Costa (10):

- Basic kernel memory functionality for the Memory Controller foundations of per-cgroup memory pressure controlling.

- socket: initial cgroup code.

- Account tcp memory as kernel memory

- per-netns ipv4 `sysctl_tcp_mem`

- tcp buffer limitation: per-cgroup limit

- Display current tcp memory allocation in kmem cgroup

- Display current tcp failcnt in kmem cgroup

- Display maximum tcp memory allocation in kmem cgroup

- Disable task moving when using kernel memory accounting

```
Documentation/cgroups/memory.txt | 38 +++++-
include/linux/memcontrol.h       | 19 +++
include/net/netns/ipv4.h         | 1 +
include/net/sock.h               | 232 +++++
include/net/tcp.h                 | 4 +-
include/net/tcp_memcg.h          | 20 +++
init/Kconfig                     | 14 ++
mm/memcontrol.c                  | 209 +++++
net/core/sock.c                  | 106 +++++
net/ipv4/Makefile                | 1 +
net/ipv4/af_inet.c               | 2 +
net/ipv4/proc.c                  | 7 +-
net/ipv4/sysctl_net_ipv4.c       | 65 +++++
net/ipv4/tcp.c                   | 17 +--
net/ipv4/tcp_input.c             | 12 +-

```

```
net/ipv4/tcp_ipv4.c      | 13 ++-
net/ipv4/tcp_memcg.c     | 263 ++++++
net/ipv4/tcp_output.c    |  2 +-
net/ipv4/tcp_timer.c     |  2 +-
net/ipv6/af_inet6.c      |  2 +
net/ipv6/tcp_ipv6.c      |  7 +-
21 files changed, 955 insertions(+), 81 deletions(-)
create mode 100644 include/net/tcp_memcg.h
create mode 100644 net/ipv4/tcp_memcg.c
```

--

1.7.6.4
