
Subject: [PATCH v5 00/10] per-cgroup tcp memory pressure
Posted by [Glauber Costa](#) on Mon, 07 Nov 2011 15:26:25 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi all,

This is my new attempt at implementing per-cgroup tcp memory pressure.
I am particularly interested in what the network folks have to comment on
it: my main goal is to achieve the least impact possible in the network code.

Here's a brief description of my approach:

When only the root cgroup is present, the code should behave the same way as
before - with the exception of the inclusion of an extra field in struct sock,
and one in struct proto. All tests are patched out with static branch, and we
still access addresses directly - the same as we did before.

When a cgroup other than root is created, we patch in the branches, and account
resources for that cgroup. The variables in the root cgroup are still updated.
If we were to try to be 100 % coherent with the memcg code, that should depend
on use_hierarchy. However, I feel that this is a good compromise in terms of
leaving the network code untouched, and still having a global vision of its
resources. I also do not compute max_usage for the root cgroup, for a similar
reason.

Please let me know what you think of it.

Glauber Costa (10):

- Basic kernel memory functionality for the Memory Controller
- foundations of per-cgroup memory pressure controlling.
- socket: initial cgroup code.
- per-cgroup tcp buffers control
- per-netns ipv4 sysctl_tcp_mem
- tcp buffer limitation: per-cgroup limit
- Display current tcp memory allocation in kmem cgroup
- Display current tcp memory allocation in kmem cgroup
- Display current tcp memory allocation in kmem cgroup
- Disable task moving when using kernel memory accounting

```
Documentation/cgroups/memory.txt | 38 +++-
include/linux/memcontrol.h       | 40 ++++
include/net/netns/ipv4.h         | 1 +
include/net/sock.h               | 231 +++++
include/net/tcp.h                 | 22 +-
include/net/transp_v6.h          | 1 +
init/Kconfig                     | 14 ++
mm/memcontrol.c                  | 442 +++++
net/core/sock.c                  | 121 +++++--
```

```
net/ipv4/af_inet.c      | 5 +
net/ipv4/proc.c         | 7 +-
net/ipv4/sysctl_net_ipv4.c | 65 +++++-
net/ipv4/tcp.c          | 11 +-
net/ipv4/tcp_input.c    | 12 +-
net/ipv4/tcp_ipv4.c     | 17 ++-
net/ipv4/tcp_output.c   | 2 +-
net/ipv4/tcp_timer.c    | 2 +-
net/ipv6/af_inet6.c     | 5 +
net/ipv6/tcp_ipv6.c     | 13 +-
19 files changed, 962 insertions(+), 87 deletions(-)
```

--

1.7.6.4
