

Hi,

This is my new attempt to fix all the concerns that were raised during the last iteration.

I should highlight:

- 1) proc information is kept intact. (although I kept the wrapper functions)
it will be submitted as a follow up patch so it can get the attention it deserves
- 2) sockets now hold a reference to memcg. sockets can be alive even after the task is gone, so we don't bother with between cgroups movements.
To be able to release resources more easily in this cenario, the parent pointer in struct cg_proto was replaced by a memcg object. We then iterate through its pointer (which is cleaner anyway)

The rest should be mostly the same except for small fixes and style changes.

Glauber Costa (9):

Basic kernel memory functionality for the Memory Controller
foundations of per-cgroup memory pressure controlling.
socket: initial cgroup code.
tcp memory pressure controls
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

```
Documentation/cgroups/memory.txt | 46 ++++++-
include/linux/memcontrol.h       | 23 +++++
include/net/netns/ipv4.h         | 1 +
include/net/sock.h               | 239 +++++
include/net/tcp.h                | 4 +-
include/net/tcp_memcontrol.h     | 19 +++
init/Kconfig                     | 11 ++
mm/memcontrol.c                  | 189 +++++
net/core/sock.c                  | 118 +++++
net/ipv4/Makefile                | 1 +
net/ipv4/af_inet.c               | 2 +
net/ipv4/proc.c                  | 6 +-
net/ipv4/sysctl_net_ipv4.c       | 65 ++++++-
net/ipv4/tcp.c                   | 11 +--
net/ipv4/tcp_input.c             | 12 +-
net/ipv4/tcp_ipv4.c              | 14 +-

```

```
net/ipv4/tcp_memcontrol.c      | 272 ++++++
net/ipv4/tcp_output.c          | 2 +-
net/ipv4/tcp_timer.c           | 2 +-
net/ipv6/af_inet6.c            | 2 +
net/ipv6/tcp_ipv6.c            | 8 +-
21 files changed, 968 insertions(+), 79 deletions(-)
create mode 100644 include/net/tcp_memcontrol.h
create mode 100644 net/ipv4/tcp_memcontrol.c
```

--

1.7.6.4
