
Subject: Re: [PATCH v7 0/8] Request for inclusion: tcp memory buffers

Posted by [davem](#) on Thu, 13 Oct 2011 20:12:21 GMT

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

From: Glauber Costa <glommer@parallels.com>

Date: Fri, 14 Oct 2011 00:05:58 +0400

> Also, I kind of dispute the affirmation that !cgroup will encompass
> the majority of users, since cgroups is being enabled by default by
> most vendors. All systemd based systems use it extensively, for
> instance.

I will definitely advise people against this, since the cost of having
this on by default is absolutely non-trivial.

People keep asking every few releases "where the heck has my performance
gone" and it's because of creeping features like this. This socket
cgroup feature is a prime example of where that kind of stuff comes
from.

I really get irritated when people go "oh, it's just one indirect
function call" and "oh, it's just one more pointer in struct sock"

We work really hard to `_remove_` elements from structures and make them
smaller, and to remove expensive operations from the fast paths.

It might take someone weeks if not months to find a way to make a
patch which compensates for the extra overhead your patches are adding.

And I don't think you fully appreciate that.
