Subject: Re: [PATCH v7 0/8] Request for inclusion: tcp memory buffers Posted by Glauber Costa on Thu, 13 Oct 2011 20:05:58 GMT

View Forum Message <> Reply to Message

On 10/14/2011 12:00 AM, David Miller wrote:

- > From: Glauber Costa<glommer@parallels.com>
- > Date: Thu, 13 Oct 2011 17:09:34 +0400

>

- >> This series was extensively reviewed over the past month, and after
- >> all major comments were merged, I feel it is ready for inclusion when
- >> the next merge window opens. Minor fixes will be provided if they
- >> prove to be necessary.

>

> I'm not applying this.

Thank you for letting me now about your view of this that early.

- > You're turning inline increments and decrements of the existing memory
- > limits into indirect function calls.

Yes, indeed.

- > That imposes a new non-trivial cost, in fast paths, even when people
- > do not use your feature.

Well, there is a cost, but all past submissions included round trip benchmarks.

In none of them I could see any significant slowdown.

- > Make this evaluate into exactly the same exact code stream we have
- > now when the memory cgroup feature is not in use, which will be the
- > majority of users.

What exactly do you mean by "not in use"? Not compiled in or not actively being exercised? If you mean the later, I appreciate tips on how to achieve it.

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.