Subject: Re: Re: [PATCH v5 00/10] per-cgroup tcp memory pressure
Posted by davem on Thu, 17 Nov 2011 21:35:01 GMT
View Forum Message <> Reply to Message

From: James Bottomley <jbottomley@parallels.com>
Date: Tue, 15 Nov 2011 18:27:12 +0000

> Ping on this, please. We're blocked on this patch set until we can get
> an ack that the approach is acceptable to network people.

__sk_mem_schedule is now more expensive, because instead of short-circuiting the majority of the function's logic when "allocated <= prot->sysctl_mem[0]" and immediately returning 1, the whole rest of the function is run.

The static branch protecting all of the cgroup code seems to be enabled if any memory based cgroup'ing is enabled. What if people use the memory cgroup facility but not for sockets? I am to understand that, of the very few people who are going to use this stuff in any capacity, this would be a common usage.

TCP specific stuff in mm/memcontrol.c, at best that's not nice at all.

Otherwise looks mostly good.