
Subject: Re: [PATCH v5 00/10] per-cgroup tcp memory pressure
Posted by [Glauber Costa](#) on Wed, 09 Nov 2011 18:02:55 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 11/07/2011 01:26 PM, Glauber Costa wrote:

> 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.

Dave, Eric,

Can you let me know what you think of the general approach I've followed in this series? The impact on the common case should be minimal, or at least as expensive as a static branch (0 in most arches, I believe).

I am mostly interested in knowing if this a valid pursue path. I'll be happy to address any specific concerns you have once you're ok with the general approach.

Thanks!
