
Subject: Re: Containers: css_put() dilemma
Posted by [Paul Menage](#) on Tue, 17 Jul 2007 17:44:08 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 7/17/07, Balbir Singh <balbir@linux.vnet.ibm.com> wrote:

>
> That sounds correct. I wonder now if the solution should be some form
> of delegation for deletion of unreferenced containers (HINT: work queue
> or kernel threads).

What a great idea. In fact, that's exactly what the release agent
patch already does.

>
> > Adding a synchronize_rcu in container_diput() guarantees that the
> > container structure won't be freed while someone may still be
> > accessing it.
> >
>
> Do we take rcu_read_lock() in css_put() path or use call_rcu() to
> free the container?

Good point, we ought to add rcu_read_lock() (even though it doesn't
actually do anything on architectures other than alpha, right?)

Using call_rcu to do the container kfree rather than synchronize_rcu()
would be a possible future optimization, yes.

Paul

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
