
Subject: Re: [PATCH 0/2] resource control file system - aka containers on top of nsproxy!

Posted by [Srivatsa Vaddagiri](#) on Mon, 05 Mar 2007 17:47:50 GMT

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

On Sat, Mar 03, 2007 at 01:22:44PM -0800, Paul Jackson wrote:

> I still can't claim to have my head around this, but what you write
> here, Herbert, writes here touches on what I suspect is a key
> difference between namespaces and resources that would make it
> impractical to accomplish both with a shared mechanism for aggregating
> tasks.

The way nsproxy is structured, its all pointers to actual namespace (or in case of rcfs patch) resource objects. This lets namespaces objects be in a flat hierarchy while resource objects are in tree-like hierarchy. nsproxy itself doesnt decide any hierarchy. Its those objects pointed to by nsproxy which can form different hierarchies. In fact the rcfs patches allows such a combination afaics.

> > on every limit accounting or check? I think that
> > is quite a lot of overhead ...
>
> Do either of these dereferences require locks?

A rcu_read_lock() should be required, which is not that expensive.

--

Regards,
vatsa

Containers mailing list
Containers@lists.osdl.org
<https://lists.osdl.org/mailman/listinfo/containers>
