

---

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

Posted by [Srivatsa Vaddagiri](#) on Tue, 13 Mar 2007 02:22:34 GMT

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

---

On Tue, Mar 13, 2007 at 12:31:13AM +0100, Herbert Poetzl wrote:

> just means that the current Linux-VServer behaviour  
> is a subset of that, no problem there as long as  
> it really \_is\_ a subset :) we always like to provide  
> more features in the future, no problem with that :)

Considering the example Sam quoted, doesn't it make sense to split resource classes (some of them atleast) independent of each other? That would also argue for providing multiple hierarchy feature in Paul's patches.

Given that and the mail Serge sent on why nsproxy optimization is usefull given numbers, can you reconsider your earlier proposals as below:

- pid\_ns and resource parameters should be in a single struct (point 1c, 7b in [1])
- pointers to resource controlling objects should be inserted in task\_struct directly (instead of nsproxy indirection) (points 2c in [1])

[1] <http://lkml.org/lkml/2007/3/12/138>

--

Regards,  
vatsa

---

Containers mailing list  
[Containers@lists.osdl.org](mailto:Containers@lists.osdl.org)  
<https://lists.osdl.org/mailman/listinfo/containers>

---