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

Posted by [Srivatsa Vaddagiri](#) on Fri, 09 Mar 2007 16:41:18 GMT

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

On Fri, Mar 09, 2007 at 10:04:30PM +0530, Srivatsa Vaddagiri wrote:

> 2. Regarding space savings, if 100 tasks are in a container (I dont know
> what is a typical number) -and- lets say that all tasks are to share
> the same resource allocation (which seems to be natural), then having
> a 'struct container_group *' pointer in each task_struct seems to be not
> very efficient (simply because we dont need that task-level granularity of
> managing resource allocation).

Note that this 'struct container_group *' pointer is in addition to the
'struct nsproxy *' pointer already in task_struct. If the set of tasks
over which resource control is applied is typically the same set of tasks
which share the same 'struct nsproxy *' pointer, then IMHO 'struct
container_group *' in each task_struct is not very optimal.

--

Regards,
vatsa

Containers mailing list

Containers@lists.osdl.org

<https://lists.osdl.org/mailman/listinfo/containers>
