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

Posted by [Paul Menage](#) on Thu, 08 Mar 2007 01:35:58 GMT

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

On 3/7/07, Eric W. Biederman <ebiederm@xmission.com> wrote:

> Pretty much. For most of the other cases I think we are safe referring
> to them as resource controls or resource limits. I know that roughly covers
> what cpusets and beancounters and ckrm currently do.

Plus resource monitoring (which may often be a subset of resource control/limits).

>

> The real trick is that I believe these groupings are designed to be something
> you can setup on login and then not be able to switch out of.

That's going to be the case for most resource controllers - is that the case for namespaces? (e.g. can any task unshare say its mount namespace?)

Paul

Containers mailing list

Containers@lists.osdl.org

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