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

Posted by [Srivatsa Vaddagiri](#) on Sat, 10 Mar 2007 03:19:09 GMT

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

On Sat, Mar 10, 2007 at 07:32:20AM +0530, Srivatsa Vaddagiri wrote:

> Ok, let me see if I can convey what I had in mind better:

```
>
>   uts_ns pid_ns ipc_ns
> \  |  /
> -----
>   | nsproxy |
>   -----
>       / | \ \ <-- 'nsproxy' pointer
> T1 T2 T3 ...T1000
> | | | | <-- 'containers' pointer (4/8 KB for 1000 task)
> -----
>   | container_group |
>   -----
> /
> -----
> | container |
> -----
> |
> -----
> | cpu_limit |
> -----
```

[snip]

> We save on 4/8 KB (for 1000 tasks) by avoiding the 'containers' pointer
> in each task_struct (just to get to the resource limit information).

Having the 'containers' pointer in each task-struct is great from a non-container res mgmt perspective. It lets you dynamically decide what is the fundamental unit of res mgmt.

It could be {T1, T5} tasks/threads of a process, or {T1, T3, T8, T10} tasks of a session (for limiting login time per session), or {T1, T2 ..T10, T18, T27} tasks of a user etc.

But from a vserver/container pov, this level flexibility (at a -task- level) of deciding the unit of res mgmt is IMHO not needed. The vserver/container/namespace (tsk->nsproxy->some_ns) to which a task belongs automatically defines that unit of res mgmt.

--

Regards,
vatsa

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