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

Posted by [serue](#) on Mon, 12 Mar 2007 17:25:57 GMT

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

Quoting Srivatsa Vaddagiri (vatsa@in.ibm.com):

> On Mon, Mar 12, 2007 at 10:56:43AM -0500, Serge E. Hallyn wrote:

> > What's wrong with that?

>

> I had been asking around on "what is the fundamental unit of res mgmt
> for vservers" and the answer I got (from Herbert) was "all tasks that are
> in the same pid namespace". From what you are saying above, it seems to
> be that there is no such "fundamental" unit. It can be a random mixture
> of tasks (taken across vservers) whose resource consumption needs to be
> controlled. Is that correct?

If I'm reading it right, yes.

If for vservers the fundamental unit of res mgmt is a vserver, that can surely be done at a higher level than in the kernel.

Actually, these could be tied just by doing

```
mount -t container -o ns,cpuset /containers
```

So now any task in /containers/vserver1 or any subdirectory thereof would have the same cpuset constraints as /containers. OTOH, you could mount them separately

```
mount -t container -o ns /nsproxy
mount -t container -o cpuset /cpuset
```

and now you have the freedom to split tasks in the same vserver (under /nsproxy/vserver1) into different cpusets.

-serge

> > > echo "cid 2" > /dev/cpu/proc/tasks

> >

> > Adding that feature sounds fine,

>

> Ok yes ..that can be a optional feature.

>

> --

> Regards,

> vatsa

Containers mailing list

