
Subject: Re: Summary of resource management discussion
Posted by [Paul Menage](#) on Thu, 15 Mar 2007 11:24:37 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 3/12/07, Srivatsa Vaddagiri <vatsa@in.ibm.com> wrote:

- > - (subjective!) If there is a existing grouping mechanism already (say
- > tsk->nsproxy[->pid_ns]) over which res control needs to be applied,
- > then the new grouping mechanism can be considered redundant (it can
- > eat up unnecessary space in task_struct)

If there really was a grouping that was always guaranteed to match the way you wanted to group tasks for e.g. resource control, then yes, it would be great to use it. But I don't see an obvious candidate. The pid namespace is not it, IMO. Resource control (and other kinds of task grouping behaviour) shouldn't require virtualization.

- >
- > a. Paul Menage's patches:
- >
- > (tsk->containers->container[cpu_ctlr.subsys_id] - X)->cpu_limit

Additionally, if we allow mature container subsystems to have an id declared in a global enum, then we can make the cpu_ctlr.subsys_id into a constant.

- >
- > b. rcfs
- > tsk->nsproxy->ctlr_data[cpu_ctlr.subsys_id]->cpu_limit

So what's the '-X' that you're referring to

- > 3. How are cpusets related to vserver/containers?
- >
- > Should it be possible to, lets say, create exclusive cpusets and
- > attach containers to different cpusets?

Sounds reasonable.

- >
- > 6. As tasks move around namespaces/resource-classes, their
- > tsk->nsproxy/containers object will change. Do we simple create
- > a new nsproxy/containers object or optimize storage by searching
- > for one which matches the task's new requirements?

I think the latter.

- >
- > - If we don't support hierarchy in res controllers today

> but were to add that support later, then
> user-interface shouldn't change. That's why
> designining -atleast- the user interface to support
> hierarchy may make sense

Right - having support for a hierarchy in the API doesn't mean that individual controllers have to support being in a hierarchy.

Paul

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