

---

Subject: Re: Summary of resource management discussion  
Posted by [Srivatsa Vaddagiri](#) on Fri, 16 Mar 2007 01:40:24 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Thu, Mar 15, 2007 at 12:12:50PM -0700, Paul Menage wrote:  
> There are some things that benefit from having an abstract  
> container-like object available to store state, e.g. "is this  
> container deleted?", "should userspace get a callback when this  
> container is empty?".

IMO we can still get these bits of information using nsproxy itself (I admit I haven't looked at the callback requirement yet).

But IMO a bigger use of 'struct container' object in your patches is to store hierarchical information and avoid /repeating/ that information in each resource object (struct cpuset, struct cpu\_limit, struct rss\_limit etc) a 'struct container' is attached to (as pointed out here : <http://lkml.org/lkml/2007/3/7/356>). However I don't know how many controllers will ever support such hierarchical res mgmt and thats why I said option 3 [above URL] may not be a bad compromise.

Also if you find a good answer for my earlier question "what more task-grouping behavior do you want to implement using an additional pointer that you can't reusing ->task\_proxy", it would drive home the need for additional pointers/structures.

> > >a. Paul Menage's patches:  
> > >  
> > > (tsk->containers->container[cpu\_ctlr.subsys\_id] - X)->cpu\_limit  
> > >  
> > So what's the '-X' that you're referring to  
> >  
> > Oh ..that's to seek pointer to begining of the cpulimit structure (subsys  
> > pointer in 'struct container' points to a structure embedded in a larger  
> > structure. -X gets you to point to the larger structure).  
> >  
> > OK, so shouldn't that be listed as an overhead for your rcfs version  
> > too?

X shouldn't be needed in rcfs patches, because "->ctlr\_data" in nsproxy can directly point to the larger structure (there is no 'struct container\_subsys\_state' equivalent in rcfs patches).

Container patches:

(tsk->containers->container[cpu\_ctlr.subsys\_id] - X)->cpu\_limit

rcfs:

tsk->nsproxy->ctrl\_data[cpu\_ctrl.subsys\_id]->cpu\_limit

> >Yes me too. But maybe to keep in simple in initial versions, we should  
> >avoid that optimisation and at the same time get statistics on duplicates?.  
>  
> That's an implementation detail - we have more important points to  
> agree on right now ...

yes :)

Eric, did you have any opinion on this thread?

--

Regards,  
vatsa

---

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

---