
Subject: Re: containers development plans
Posted by [serue](#) on Mon, 02 Jul 2007 21:53:15 GMT
[View Forum Message](#) <> [Reply to Message](#)

Quoting Erich Focht (efocht@hpce.nec.com):

> Hi,
>
> On Monday 02 July 2007 18:55, Serge E. Hallyn wrote:
> > A list of the people we are currently aware of who are showing interest
> > in these features follows. What I'd like to know is, from this list, do
> > some people know what general or specific areas they plan to or want to
> > work on over the next year?
> >
> > Stakeholders:
> > Eric Biederman
> > google
> > ibm
> > kerlabs
> > openvz
> > osdl (Masahiko Takahashi?)
> > vserver
> >
> > Who is missing from the list?
>
> you could add the XtreamOS project to the list. It's grid project and we
> need to use containers for grid job isolation. We're currently using BLCR
> for checkpointing, and also kerrighed, but aim at switching to containers
> and will work on container checkpointing. XtreamOS is related to
> kerlabs & kerrighed, but has broader interests in containers.

That's cool. I'll go read up on XtreamOS before I ask stupid questions...

> > What more needs to be added to this list?
>
> By the way, there are potential synergies between container checkpointing
> work and the need of the VM to have something like per cpuset or per
> container swap (as explained in the VM BoFS at OLS, there are unnecessary
> OOMs with a common swap).

Ah, cool, I hadn't heard that. At the risk of really exploding the # of container subsystems needed to be composed in order to do checkpoint/restart, this does sound like a separate container subsystem which just attaches a swapfile to it's tasks might then be a useful thing. It can then be composed with a cpuset or a freezer subsystem, or both (at the same time - separately might be weird :).

I guess then for checkpoint we just freeze the subsystem, force a write

of all dirty pages using something like Dave's patch, somehow create a copy-on-write clone as the checkpoint, then unfreeze. (and oh yeah, checkpoint all the other crap too :)

> Regards,
> Erich

thanks,
-serge

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
