
Subject: Re: Request to nominate someone to represent the containers development community

Posted by [serge](#) on Mon, 11 Jun 2007 19:42:59 GMT

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

Quoting Cedric Le Goater (clg@fr.ibm.com):

> Serge E. Hallyn wrote:

> > Quoting Theodore Tso (tytso@mit.edu):

> > ...

(removing ksummit-2007-pc from cc: list, and adding containers and ckrm-tech lists)

Background for containers@ and ckrm-tech@, we are working to lay out a roadmap for the next year's development in "containers" (both as pertains to namespaces, and as pertains to resource management) for the kernel summit. I figure right now we're discussing things which should go on the roadmap. Then we can talk about dependencies, then we can talk roadmap.

Does that make sense?

(in as much as a roadmap on anything lkml-related can ever make sense :)

> > I took a look back through the current proposed agenda, and it looks
> > like 'containers' is listed under virtualization but explicitly lists
> > resource management. So I'm going to assume that 'containers' refers
> > to both Paul Menage's task containers work and the namespace container
> > work.

> >

> > 1. continuation of namespaces

> > devpts, console and ttydrivers

> > user

> > time

> > namespace management tools

>

> like namespace entering ?

Yes, that's something some people want, and others don't, so that is definately something we'd want to put at least on this list, and probably on the roadmap.

> > more ?

>

> things to complete:

>

> pid namespace

> net namespace

> ro bind mounts (related but should probably be in an other

Ok, so that should go at the very top of '1.'

> > 2. any additional work needed for virtual servers?

> > i.e. in-kernel keyring usage for cross-namespace permissions, etc

> > 3. task containers functionality

> > base features

> > specific containers

> > poll to see who has plans

> > ?

>

> we need a big bullet on resource management

Well isn't that what all of item '3' is ?

> > 4. checkpoint/restart

> > memory c/r

> > (there are a few designs and prototypes)

> > (though this may be ironed out by then)

>

> We have this idea of a per container swap which we would use to snapshot

> processes memory 'a la' swsusp'

Right, we're at an odd time to flesh out this exact item, as we haven't yet decided which overall approach to take. If we go with Dave's approach of using per-container swapfile, combined with yours of using the freezer code to organize the checkpoint/restart, then we can flesh it out according to that. But I don't think we're ready for that yet.

> > overall checkpoint strategy

> > in-kernel

> > userspace-driven

> > overall restart strategy

> >

> > Should we move this conversation to the containers and ckrm lists?

>

> probably yes.

Done.

Have we forgotten entire pieces of functionality which should be added to this outline?

thanks,
-serge

Containers mailing list

