

---

Subject: Re: [RFC][PATCH 0/4] Container Freezer  
Posted by [serue](#) on Thu, 21 Jun 2007 16:37:59 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

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

> Serge E. Hallyn wrote:

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

> > > This patchset is a prototype using the container infrastructure and  
> > > the swsusp freezer to freeze a group of tasks.

> > >

> > > 2 files are defined by the freezer subsystem in the container

> > > filesystem :

> > >

> > > \* "freezer.freeze"

> > >

> > > writing 1 will freeze all tasks and 0 unfreeze

> > > reading will return the status of the freezer

> > >

> > > \* "freezer.kill"

> > >

> > > writing <n> will send signal number <n> to all tasks

> > >

> > > \* Usage :

> > >

> > > # mkdir /containers/freezer

> > > # mount -t container -ofreezer freezer /containers/freezer

> > > # mkdir /containers/freezer/0

> > > # echo \$some\_pid > /containers/freezer/0/tasks

> >

> > I'd like to point out that by composing this with the nsproxy subsystem

> > (which I've done and tested), you can

> >

> > mount -t container -ons,freezer nsproxy /containers/freezer

> >

> > Then you get a new freezer subsystem automatically for every unshare

> > that you do, avoiding the need to manually do

> >

> > echo \$some\_pid > /containers/freezer/0/tasks

>

> I guess that's true for any subsystem. right ?

Sure, but when this goes to lkml I expect there to be plenty of people who don't know about that quite yet, and might come away thinking this is more intrusive to use than it is.

thanks,  
-serge

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

---