Subject: Re: [ckrm-tech] [PATCH 7/7] containers (V7): Container interface to nsproxy subsystem
Posted by serue on Mon, 26 Mar 2007 21:57:55 GMT
View Forum Message <> Reply to Message

Quoting Srivatsa Vaddagiri (vatsa@in.ibm.com):
> On Sat, Mar 24, 2007 at 10:35:37AM +0530, Srivatsa Vaddagiri wrote:
> > > +static int ns_create(struct container_subsys *ss, struct container *cont)
> > > +{
> > > + struct nscont *ns;
> > > +
> > > + if (!capable(CAP_SYS_ADMIN))
> > > + return -EPERM;
> >
> > Does this check break existing namespace semantics in a subtle way?
> > It now requires that unshare() of namespaces by any task requires
> > CAP_SYS_ADMIN capabilities.
>
> I should clarify that I am referring to unshare thr' clone here (and not
> thr' sys_unshare)

That is still not true, see kernel/utsname:copy_utsname().

Now you might have run a userspace testcase in a kernel with
CONFIG_UTS_NS=n, which at the moment erroneously returns 0 rather than
-EINVAL when you clone(CLONE_NEWUTS). But you didn't get a new uts
namespace, you were just lied to :)

-serge

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