
Subject: Re: [PATCH 1/1] containers: implement nsproxy containers subsystem
Posted by [Pavel Emelianov](#) on Tue, 05 Jun 2007 13:04:58 GMT
[View Forum Message](#) <> [Reply to Message](#)

Serge E. Hallyn wrote:

> Quoting Pavel Emelianov (xemul@openvz.org):
>> Serge E. Hallyn wrote:
>>> >From 190ea72d213393dd1440643b2b87b5b2128dff87 Mon Sep 17 00:00:00 2001
>>> From: Serge E. Hallyn <serue@us.ibm.com>
>>> Date: Mon, 4 Jun 2007 14:18:52 -0400
>>> Subject: [PATCH 1/1] containers: implement nsproxy containers subsystem
>>>
>>> When a task enters a new namespace via a clone() or unshare(), a new
>>> container is created and the task moves into it. This enables
>> I have a design question.
>>
>> How the child that has a new namespace guesses what id
>> this namespace has in containers?
>
> parse /proc/\$\$/container

Ok.

> So more likely the parent would have to grab the cloned pid of the
> child, parse /proc/\$\$/container, then rename the container.

Child can happen to die before this and we'll have an orphaned
container. I mean, it will be deletable, but its name will be unknown.

Maybe its better to get the containers id from the pid of new task?

> -serge
> -
> To unsubscribe from this list: send the line "unsubscribe linux-kernel" in
> the body of a message to majordomo@vger.kernel.org
> More majordomo info at <http://vger.kernel.org/majordomo-info.html>
> Please read the FAQ at <http://www.tux.org/lkml/>
>
