
Subject: Re: [ckrm-tech] [PATCH 7/7] containers (V7): Container interface to nsproxy subsystem

Posted by [serue](#) on Wed, 28 Mar 2007 15:26:19 GMT

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

Quoting Srivatsa Vaddagiri (vatsa@in.ibm.com):

> On Mon, Mar 26, 2007 at 04:57:55PM -0500, Serge E. Hallyn wrote:

> > 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 :)

>

> I think you are right here, in that CONFIG_UTS_NS was not turned on,

> although I was thinking it was on.

>

> However as a result of this experiment, I found this anomaly:

>

> - On a kernel with CONFIG_UTS_NS=n, a test which does

> clone(CLONE_NEWUTS) works fine. clone() succeeds and the child

> starts running with no error.

> - On the same kernel, if ns container hierarchy is mounted, then

> the test fails. clone() returns failure and child is never created.

> As soon as the ns container hierarchy is unmounted, the test works

> again.

>

> I would have expected a consistent behavior here, irrespective of

> whether ns hierarchy is mounted or not. Is this difference in behavior

> acceptable? Returning -EINVAL in copy_utsname() when CONFIG_UTS_NS=n, as

> you say above, would fix this anomaly.

Will, not would, fix the anomaly :)

2.6.21-rc5-mm2 has the correct behavior. Returning 0 was a bug.

thanks,

-serge

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
