
Subject: Re: [PATCH] pidns: Limit kill -1 and cap_set_all
Posted by [ebiederm](#) on Fri, 26 Oct 2007 22:10:04 GMT
[View Forum Message](#) <> [Reply to Message](#)

"Kir Kolyshkin" <kir@swsoft.com> writes:

> Eric, this problem is a known one. Currently Pavel and Sukadev are working on a
> appropriate signal management for namespaces.

I'm fairly certain that the signal issue you they are dealing with
is how to keep children from killing the init of a pid namespace.
At least that is what Suka mentioned when I asked him earlier,
and that is the discussion I have at various times on the containers
list.

This is a totally different issue. This is keeping
kill -1 in a child pid namespace from sending a signal to
every process in the system.

It is a bad bug and this patch is a complete fix for it.
To my knowledge Pavel and Suka just overlooked this case.

> If you'd first sent this patch for discussion/review to containers@, you'd know
> that the patch is very incomplete.

I explicitly asked Suka, Pavel, and Cedric how complete they thought
the pid namespace was before I started sending patches. With
Pavel at a conference it makes sense why he did not reply.

For the most part Suka and Cedric did not know about the issues
that I have been finding and fixing.

So since several of them have been fixable with trivial simple fixes
and because we have passed -rc1 and it is now time for bug fixes
and possibly a few cleanups. I have been making bug fixes.

Kir I am happy from feedback from people who can understand the code
and patches. But I am not going to hold up sending simple obviously
correct fixes for discussion.

Eric

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