Subject: unshare() pid ns
Posted by Sukadev Bhattiprolu on Thu, 30 Aug 2007 05:29:28 GMT

Pavel,

unshare() of pid ns seems to fail with -EINVAL in 2.6.23-rc3-mm1.
I thought we supported it in the earlier patchsets. I guess
I missed that in the review of recent patchsets.

Did we remove/disable it on purpose ? I am not particular that we
support it though. Just want to make sure.

Suka

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

Subject: Re: unshare() pid ns
Posted by Pavel Emelianov on Thu, 30 Aug 2007 08:10:51 GMT

sukadev@us.ibm.com wrote:
> Pavel,
>
> unshare() of pid ns seems to fail with -EINVAL in 2.6.23-rc3-mm1.
> I thought we supported it in the earlier patchsets. I guess
> I missed that in the review of recent patchsets.

I disabled unsharing of pid namespaces because it's almost
impossible. Look - you have to reattach all the pids to the
task with saving its ids as seen in previous namespaces.

> Did we remove/disable it on purpose ? I am not particular that we
> support it though. Just want to make sure.
>
> Suka
>

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

Subject: Re: unshare() pid ns
Posted by serue on Thu, 30 Aug 2007 13:34:22 GMT
View Forum Message <> Reply to Message

Quoting Pavel Emelyanov (xemul@openvz.org):
> sukadev@us.ibm.com wrote:
>> Pavel,
>> unshare() of pid ns seems to fail with -EINVAL in 2.6.23-rc3-mm1.
>> I thought we supported it in the earlier patchsets. I guess
>> I missed that in the review of recent patchsets.
>
> I disabled unsharing of pid namespaces because it's almost
> impossible. Look - you have to reattach all the pids to the
> task with saving its ids as seen in previous namespaces.

We agree, but thought you for some perverse reason preferred unshare to
clone for pidns :)

Thanks for clarifying.

-serge

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

Subject: Re: unshare() pid ns
Posted by Pavel Emelianov on Thu, 30 Aug 2007 14:11:27 GMT
View Forum Message <> Reply to Message

Serge E. Hallyn wrote:
> Quoting Pavel Emelyanov (xemul@openvz.org):
>> sukadev@us.ibm.com wrote:
>>> Pavel,
>>> unshare() of pid ns seems to fail with -EINVAL in 2.6.23-rc3-mm1.
>>> I thought we supported it in the earlier patchsets. I guess
>>> I missed that in the review of recent patchsets.
>> I disabled unsharing of pid namespaces because it's almost
>> impossible. Look - you have to reattach all the pids to the
>> task with saving its ids as seen in previous namespaces.
>
> We agree, but thought you for some perverse reason preferred unshare to
> clone for pidns :)

I did that in my first version of patches, but then realized
that such problem (the need in reattaching pids) makes the
unsharing ugly.

BTW, unsharing of a pid namespace is a valid operation, so I
think I will enable it in the nearest future. I have some
thought on how to make such a reattach ;)

> Thanks for clarifying.
>
> -serge
>

_____

Subject: Re: unshare() pid ns
Posted by Sukadev Bhattiprolu on Thu, 30 Aug 2007 14:19:40 GMT

View Forum Message <> Reply to Message

Pavel Emelianov [xemul@openvz.org] wrote:
| sukadev@us.ibm.com wrote:
| >Pavel,
| >
| >unshare() of pid ns seems to fail with -EINVAL in 2.6.23-rc3-mm1.
| >I thought we supported it in the earlier patchsets. I guess
| >I missed that in the review of recent patchsets.
|
| I disabled unsharing of pid namespaces because it's almost
| impossible. Look - you have to reattach all the pids to the
| task with saving its ids as seen in previous namespaces.

Agree. Eric pointed that out early on.
|
| >Did we remove/disable it on purpose ? I am not particular that we
| >support it though. Just want to make sure.
| >
| >Suka
| >

_____

Subject: Re: unshare() pid ns
Posted by serue on Thu, 30 Aug 2007 14:35:18 GMT

Quoting Pavel Emelyanov (xemul@openvz.org):
> Serge E. Hallyn wrote:
> > Quoting Pavel Emelyanov (xemul@openvz.org):
> >> sukadev@us.ibm.com wrote:
> >>> Pavel,
> >>> unshare() of pid ns seems to fail with -EINVAL in 2.6.23-rc3-mm1.
> >>> I thought we supported it in the earlier patchsets. I guess
> >>> I missed that in the review of recent patchsets.
> >> I disabled unsharing of pid namespaces because it's almost
> >> impossible. Look - you have to reattach all the pids to the
> >> task with saving its ids as seen in previous namespaces.
> >
> > We agree, but thought you for some perverse reason preferred unshare to
> > clone for pidns :)
>
> I did that in my first version of patches, but then realized
> that such problem (the need in reattaching pids) makes the
> unsharing ugly.
>
> BTW, unsharing of a pid namespace is a valid operation, so I
> think I will enable it in the nearest future. I have some
> thought on how to make such a reattach ;)

Alrighty :)

Of course it's not just the kernel ugliness, but also the userspace
ugliness, for instance the rumored (I haven't looked to confirm) caching
of pids by glibc.

But in the end if we can achieve symmetry between all the CLONE_NEW*
flags all the better.

-serge

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

Subject: Re: unshare() pid ns
Posted by Cedric Le Goater on Thu, 30 Aug 2007 16:34:29 GMT

Hello Pavel !

Pavel Emelyanov wrote:

> Serge E. Hallyn wrote:
>> Quoting Pavel Emelyanov (xemul@openvz.org):
>>> sukadev@us.ibm.com wrote:
>>>> Pavel,
>>>> unshare() of pid ns seems to fail with -EINVAL in 2.6.23-rc3-mm1.
>>>> I thought we supported it in the earlier patchsets. I guess
>>>> I missed that in the review of recent patchsets.
>>> I disabled unsharing of pid namespaces because it's almost
>>> impossible. Look - you have to reattach all the pids to the
>>> task with saving its ids as seen in previous namespaces.
>> We agree, but thought you for some perverse reason preferred unshare to
>> clone for pidns :)
>
> I did that in my first version of patches, but then realized
> that such problem (the need in reattaching pids) makes the
> unsharing ugly.
>
> BTW, unsharing of a pid namespace is a valid operation, so I
> think I will enable it in the nearest future. I have some
> thought on how to make such a reattach ;)

that would be also very useful to 'enter' such a namespace.

BTW, did you take a look at the sys_hijack() serge sent ?

Thanks,

C.

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