
Subject: Re: unshare() pid ns
Posted by [Cedric Le Goater](#) on Thu, 30 Aug 2007 16:34:29 GMT
[View Forum Message](#) <> [Reply to Message](#)

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>
