
Subject: Re: [PATCH 1/3] change clone_flags type to u64
Posted by [Daniel Hokka Zakrisson](#) on Fri, 11 Apr 2008 08:45:22 GMT
[View Forum Message](#) <> [Reply to Message](#)

Serge E. Hallyn wrote:

> Quoting Daniel Hokka Zakrisson (daniel@hozac.com):

>> Serge E. Hallyn wrote:

>> > Quoting Andi Kleen (andi@firstfloor.org):

>> >> > I guess that was a development rationale.

>> >>

>> >> But what rationale? It just doesn't make much sense to me.

>> >>

>> >> > Most of the namespaces are in

>> >> > use in the container projects like openvz, vserver and probably

>> others

>> >> > and we needed a way to activate the code.

>> >>

>> >> You could just have added it to feature groups over time.

>> >>

>> >> >

>> >> > Not perfect I agree.

>> >> >

>> >> > > With your current strategy are you sure that even 64bit will

>> >> > > be enough in the end? For me it rather looks like you'll

>> >> > > go through those quickly too as more and more of the kernel

>> >> > > is namespaced.

>> >> >

>> >> > well, we're reaching the end. I hope ! devpts is in progress and

>> >> > mq is just waiting for a clone flag.

>> >>

>> >> Are you sure?

>> >

>> > Well for one thing we can take a somewhat different approach to new

>> > clone flags. I.e. we could extend CLONE_NEWIPC to do mq instead of

>> > introducing a new clone flag. The name doesn't have 'sysv' in it,

>> > and globbing all ipc resources together makes some amount of sense.

>> > Similarly has hpa+eric pointed out earlier, suka could use

>> > CLONE_NEWDEV for ptys. If we have net, pid, ipc, devices, that's a

>> > pretty reasonable split imo. Perhaps we tie user to devices and get

>> > rid of CLONE_NEWUSER which I suspect noone is using atm (since only

>> > Dave has run into the CONFIG_USER_SCHED problem). Or not. We could

>> > roll uts into net, and give CLONE_NEWUTS a deprecation period.

>>

>> Please don't. Then we'd need to re-add it in Linux-VServer to support

>> guests where network namespaces aren't used...

>

> So these are networked vservers with a different hostname? Just

> curious, what would be a typical use for these?

Layer 3 isolation will continue to be the default for Linux-VServer.

> Anyway then I guess we won't :) Do you have other suggestions for
> ns clone flags which ought to be combined? Do the rest of what I
> listed make sense to you? (If not, then I guess I'll step out of the
> way and let you and Andi fight it out :)

I think putting mq under CLONE_NEWIPC makes sense, as well as using CLONE_NEWDEV for the ptys. If CLONE_NEWUSER is to be combined with anything, I think it makes more sense to combine it with CLONE_NEWPID than CLONE_NEWDEV.

> thanks,
> -serge
>

--

Daniel Hokka Zakrisson

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