
Subject: Re: [PATCH 1/3] change clone_flags type to u64

Posted by [serue](#) on Thu, 10 Apr 2008 17:14:53 GMT

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

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.

-serge

Containers mailing list

Containers@lists.linux-foundation.org

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