
Subject: Re: [PATCH 1/3] change clone_flags type to u64
Posted by [Andi Kleen](#) on Thu, 10 Apr 2008 08:25:44 GMT
[View Forum Message](#) <> [Reply to Message](#)

sukadev@us.ibm.com writes:

> From: Sukadev Bhattiprolu <sukadev@us.ibm.com>
> Subject: [lxc-dev] [patch -lxc 1/3] change clone_flags type to u64
>
> This is a preliminary patch changing the clone_flags type to 64bits
> for all the routines called by do_fork().

I must admit I was always a little sceptical of giving every tiny namespaceable kernel feature its own CLONE flag (and it's own CONFIG option). What was the rationale for that again?

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.

Also I think the user interface is very unfriendly. How is a non kernel hacker supposed to make sense of these myriads of flags? You'll be creating another CreateProcess123_extra_args_extended() in the end I fear.

Wouldn't it be better to just partition all this into fewer more understandable larger feature groups? I think that would be much nicer from pretty much all perspectives (kernel maintenance, user interface sanity, not needing clone128/256 in the end etc.)

Some consolidation on the CONFIGs would be good too. I just cannot imagine it really makes sense to configure everything so fine grained and this is just asking for random compile breakage on randconfig.

-Andi

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