sukadev@us.ibm.com wrote:
> Pavel Emelianov [xemul@openvz.org] wrote:
> | This is "submition for inclusion" of hierarchical, not kconfig
> | configurable, zero overheaded ;) pid namespaces.
> |
> | The overall idea is the following:
> |
> | The namespace are organized as a tree - once a task is cloned
> | with CLONE_NEWPIDS (yes, I've also switched to it :) the new
>
> Can you really clone() a pid namespace all by itself ?
> copy_namespaces() has the following:
>
>
>   if (!(flags & (CLONE_NEWNS | CLONE_NEWUTS | CLONE_NEWIPC |
CLONE_NEWUSER)))
>   return 0;
>
> doesn't it mean you cannot create a pid namespace using clone() unless
> one of the above flags are also specified ?
>
> unshare_nsproxy_namespaces() has the following correct check:
>
>   if (!(unshare_flags & (CLONE_NEWNS | CLONE_NEWUTS | CLONE_NEWIPC |
>   CLONE_NEWUSER | CLONE_NEWPIDS)))
>   return 0;

I have already pointed this out. I attached the previous version of
this patch where I cloned the namespaces via unshare :(However the
copy_pid_ns code looks correct. I will resend it altogether.

My bad. I have to stop working at times I want to sleep...

> BTW, why not use CLONE_NEWPID and drop the 'S' ? We don't have 'S' with
> other namespaces.

CLONE_NEWPID? Hm... I see no difference. OK I'll switch to it.

> Suka
>

_____

Containers mailing list
Containers@lists.linux-foundation.org

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