
Subject: Re: [patch -mm 10/17] nsproxy: add unshare_ns and bind_ns syscalls
Posted by [ebiederm](#) on Fri, 08 Dec 2006 19:26:49 GMT
[View Forum Message](#) <> [Reply to Message](#)

clg@fr.ibm.com writes:

> From: Cedric Le Goater <clg@fr.ibm.com>
>
> The following patch defines 2 new syscalls specific to nsproxy and
> namespaces :
>
> * unshare_ns :
>
> enables a process to unshare one or more namespaces. this
> duplicates the unshare syscall for the moment but we
> expect to diverge when the number of namespaces increases

Are we out of clone flags yet? If not this is premature.

> * bind_ns :
>
> allows a process to bind
> 1 - its nsproxy to some identifier
> 2 - to another nsproxy using an identifier or -pid

NAK

Don't use global identifiers. Use pids. i.e. struct pid * for your
identifiers. Is there is a reason pids are unsuitable?

I'm also worried about the security implications of switching namespaces
on a process. That is something that needs to be looked at very closely.

These two changes certainly don't belong in a single patch, and they
certainly use a bit more explanation. syscalls are not something to add lightly.
Because they must be supported forever.

Eric

Containers mailing list
Containers@lists.osdl.org
<https://lists.osdl.org/mailman/listinfo/containers>
