
Subject: Re: [patch -mm 10/17] nsproxy: add unshare_ns and bind_ns syscalls
Posted by [Cedric Le Goater](#) on Mon, 11 Dec 2006 17:05:16 GMT
[View Forum Message](#) <> [Reply to Message](#)

Herbert Poetzl wrote:

> On Fri, Dec 08, 2006 at 12:26:49PM -0700, Eric W. Biederman wrote:

>> 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.

>

> no, but a different nevertheless related question:

> does anybody, except for 'us' use the unshare() syscall?

>

> because if not, then why not simply extend that one

> to 64bit and be done, we probably won't need a clone64()

> but if we find we do (at some point) adding that with

> the new flags would be trivial ...

>

> OTOH, we could also just add an unshare64() too

>

> anyway, we _will_ run out of flags in the near future

yes. that's probably the way to go. I'll rework unshare_ns() in a
unshare64(). it will give some air to the 32bits clone() and unshare() and
will let us use the >32bits flags for namespaces.

thanks,

C.

Containers mailing list

Containers@lists.osdl.org

<https://lists.osdl.org/mailman/listinfo/containers>
