
Subject: Re: [patch -mm 08/17] nsproxy: add hashtable
Posted by [ebiederm](#) on Thu, 14 Dec 2006 21:08:58 GMT
[View Forum Message](#) <> [Reply to Message](#)

Cedric Le Goater <clg@fr.ibm.com> writes:

>>> Let me try : you would add a 'struct pid pid' field to all namespaces and
>>> assign that 'pid' field with the struct pid of the task creating the
>>> namespace ?

>>

>> Yes a struct pid *pid field, that we did the proper reference counting
>> on.

>

> sure.

>

>> As for which pid to assign, that is a little trickier. The struct pid
>> of the task creating the namespace is the obvious choice and that will
>> always work for clone. For unshare that would only work if we added
>> the restriction you can't unshare if someone already has used that pid
>> for that kind of namespace.

>

> clone will also be an issue if more than one namespace is unshared. Do
> you use the same 'struct pid*' for each namespace ? hmm, it feels also
> wrong.

>

> having an id field in the namespace and using a bind_ns like syscall
> to let the user assign whatever id he wants to, doesn't seem to be
> such a bad idea.

I think I would probably have suggested simply taking the next
available id in that case in practice. Partly it depends on exactly
what we are trying to do with these.

But I do agree getting that last little details right so some corner
case doesn't feel wrong is hard. That is why I try and put off
this kind of things until as much is known about how we are going
to use it as possible. So we can make a good decision and solve
practical problems, and not theoretical ones.

Eric

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