
Subject: Re: The issues for agreeing on a virtualization/namespaces implementation.

Posted by [Alexey Kuznetsov](#) on Wed, 08 Feb 2006 00:43:25 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hello!

> >2) What is the syscall interface to create these namespaces?

> > - Do we add clone flags?

> > (Plan 9 style)

>

> Like that approach .. flexible .. particular when one has well specified
> namespaces.

>

> > - Do we add a syscall (similar to setsid) per namespace?

> > (Traditional unix style)?

>

> Where does that approach end .. what's wrong with doing it at clone() time ?

That most of those namespaces need a special setup rather than a plain copy?

F.e. what are you going to do with NETWORK namespace? The only valid thing to do is to prepare a new context and to configure its content (addresses, routing tables, iptables...) later. So that, in this case it is natural to inherit the context through clone() and to create new context with a separate syscall.

Seems, only PID space needs to be setup at clone time. All the rest of suggested namespaces are more convenient to change with separate syscalls.

I would suggest to combine both approaches. Those namespaces, which can be naturally copied while clone() (f.e. the best example is already existing CLONE_NEWNS) deserve a clone() flag. The rest are preserved through clone() and forked and configured later.

Alexey
