
Subject: Re: semantics for namespace naming

Posted by [Cedric Le Goater](#) on Wed, 20 Dec 2006 09:36:56 GMT

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

Herbert Poetzl wrote:

> On Thu, Dec 14, 2006 at 09:36:31AM -0600, Serge E. Hallyn wrote:

>

> [lot of stuff zapped here]

>

>> Quoting Cedric Le Goater (clg@fr.ibm.com):

>>> how do you enter only a subset of namespaces of a nsproxy/container

>>> and not all of it ?

>> one container corresponds to one nsproxy which is one set of

>> namespaces. Are you asking how you would only switch your pid

>> namespace but keep your network namespace as the original?

>>

>> I'm not sure that's something we want/need to support.

>

> I can, we use this for several purposes, one is to

> extend or modify the VFS namespace by mounting or

> unmounting filesystems _into_ the guest

>

>> (Can you cite a use case?)

>

>> But since I haven't specified how to ask for the nsproxy switch

>> anyway, it's too early to ask how we add a flag to specify a

>> namespace subset :)

>

> we (Linux-VServer) extended the set/enter_namespace

> command to take a flag mask similar to unshare

> (only 64bit wide :) to allow for 'selecting' the

> proper spaces

>

> note: the nsproxy a context is referring to can also

> be changed this way, and a privileged process can

> enter each guest space separately ...

I'll work on a clone64 and unshare64 in January. But did you take a look at the bind_ns syscall to see how it fits with vserver ?

C.

Containers mailing list

Containers@lists.osdl.org

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