

---

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

Posted by [ebiederm](#) on Wed, 08 Feb 2006 16:48:14 GMT

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

---

Kirill Korotaev <dev@sw.ru> writes:

>>>Eric W. Biederman wrote:

>>>So it seems the clone( flags ) is a reasonable approach to create new namespaces. Question is what is the initial state of each namespace?

>>>In pidspace we know we should be creating an empty pidmap !

>>>In network, someone suggested creating a loopback device

>>>In uts, create "localhost"

>>>Are there examples where we rather inherit ? Filesystem ?

>> Of course filesystem is already implemented, and does inheret a full

>> copy.

>

> why do we want to use clone()? Just because of its name and flags?

> I think it is really strange to fork() to create network context. What has

> process creation has to do with it?

Agreed. Although clones brother unshare takes process creation out of the picture, but otherwise preserves the same interface.

> After all these clone()'s are called, some management actions from host system

> are still required, to add these IPs/routings/etc.

> So? Why mess it up? Why not create a separate clean interface for container

> management?

If we need additional arguments besides create the thing. We have a clear argument that clone is completely the wrong interface.

However. So far I have not seen an instance where using the existing standard configuration mechanisms from inside the namespace is not the proper way to set things up. The only thing I know that needs to happen from outside is to pass the container a network interface. And if it is a physical interface that is all that must happen.

Eric

---