
Subject: Re: Re: Network virtualization/isolation

Posted by [Herbert Poetzl](#) on Wed, 06 Dec 2006 18:30:36 GMT

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

On Wed, Dec 06, 2006 at 02:54:16PM +0300, Kirill Korotaev wrote:

> >>> If there is a better and less intrusive while still being obvious
> >>> method I am all for it. I do not like the OpenVZ thing of doing the
> >>> lookup once and then stashing the value in current and the special
> >>> casing the exceptions.
> >>
> >> Why?
> >
> >
> > I like it when things are obvious and not implied.
> >
> > The implementations seems to favor fewer lines of code touched over
> > maintainability of the code. Which if you are maintaining out of
> > tree code is fine. At least that was my impression last time
> > I looked at the code.

> FYI, when we started doing networking virtualization many years ago
> we tried both approaches.

> Over time, context notion looked much more natural and easier for us.

> Even Alexey Kuznetsov tells that he prefers `exec_env` as the logic

> becomes very clear and little mess is introduced.

>

> > I know there are a lot of silly things in the existing implementations

> > because they were initially written without the expectation of being

> > able to merge the code into the main kernel. This resulted in some

> > non-general interfaces, and a preference for patches that touch

> > as few lines of code as possible.

> Sure, but OpenVZ code is being constantly cleaned from such code

> and we are open for discussion. No one pretends that code is perfect

> from the beginning.

>

> > Anyway this has bit has been discussed before and we can discuss it

> > seriously in the context of patch review.

> Let me explain when explicit context like `exec_env` IMHO is cleaner:

> - context is a natural notion of linux kernel. e.g. `current`.

> why not pass '`current`' to all the functions as an argument

> starting from `entry.S`?

> `in_atomic()`, `in_interrupt()` etc. all these functions deal with

> `current` context. IMHO when one needs to pass an argument too many

> times like '`current`'

> it is better to use a notion of the context.

> - e.g. NFS should set networking context of the mount point or socket.

how would that work for a 'shared' NFS partition?

(shared between different context)

> But, ok, it is not the real point to argue so much imho and waste our
> time instead of doing things.

well, IMHO better talk (and think) first, then implement
something ... not the other way round, and then start
fixing up the mess ...

best,
Herbert

> Thanks,

> Kirill

>

> _____

> Containers mailing list

> Containers@lists.osdl.org

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